

Managing vRealize Automation

vRealize Automation 7.1



vmware®

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

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

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

docfeedback@vmware.com

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

Copyright © 2015 – 2017 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

Managing vRealize Automation 5

Updated Information 6

1 Managing vRealize Automation 7

Broadcast a Message on the Message Board Portlet 7

Starting Up and Shutting Down vRealize Automation 9

Start Up vRealize Automation 9

Restart vRealize Automation 10

Shut Down vRealize Automation 11

Updating vRealize Automation Certificates 11

Extracting Certificates and Private Keys 13

Replace Certificates in the vRealize Automation appliance 13

Replace the Infrastructure as a Service Certificate 15

Replace the IaaS Manager Service Certificate 17

Updating the vRealize Automation Appliance Management Site Certificate 19

Replace a Management Agent Certificate 23

Change the Polling Method for Certificates 26

Managing the vRealize Automation Postgres Appliance Database 26

Configure the Appliance Database 27

Perform Manual vRealize Automation Appliance Database Failover 29

Scenario: Perform a Maintenance Database Failover 30

Backup and Recovery for vRealize Automation Installations 31

Configuring the Customer Experience Improvement Program for vRealize Automation 31

Join or Leave the VMware Customer Experience Improvement Program for
vRealize Automation 31

Configure Data Collection Time 32

Adjusting System Settings 33

Modify the All Services Icon in the Service Catalog 33

Customize Data Rollover Settings 35

Adjusting Settings in the Manager Service Configuration File 36

Monitoring vRealize Automation 42

Monitoring Workflows and Viewing Logs 42

Monitoring Event Logs and Services 42

Viewing Host Information for Clusters in Distributed Deployments 44

Monitoring and Managing Resources 46

Choosing a Resource Monitoring Scenario 46

Resource Usage Terminology 50

- Connecting to a Cloud Machine 50
- Reducing Reservation Usage by Attrition 53
- Decommissioning a Storage Path 53
- Data Collection 54
- Understanding vSwap Allocation Checking for vCenter Server Endpoints 58
- Removing Datacenter Locations 59
- Bulk Import, Update, or Migrate Virtual Machines 59
 - Import a Virtual Machine to a vRealize Automation Environment 60
 - Update a Virtual Machine in a vRealize Automation Environment 64
 - Migrate a Virtual Machine to a Different vRealize Automation Environment 66
- Managing Machines 70
 - Managing Virtual Machines 70
 - Running Actions for Provisioned Resources 85

Managing vRealize Automation

Managing vRealize Automation provides information about maintaining VMware vRealize™ Automation, including how to start and stop a deployment, as well as manage certificates and the appliance database. In addition, it contains information on backing up and restoring vRealize Automation.

Intended Audience

This information is intended for anyone who wants to manage a vRealize Automation deployment. The information is written for experienced Windows or Linux system administrators who are familiar with virtual machine technology and datacenter operations.

VMware Technical Publications Glossary

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

Updated Information

This *Managing vRealize Automation* is updated with each release of the product or when necessary.

This table provides the update history of the *Managing vRealize Automation*.

Revision	Description
EN-002077-04	Removed erroneous topic.
EN-002077-03	Updated Perform Manual vRealize Automation Appliance Database Failover .
EN-002077-02	<ul style="list-style-type: none">Updated Import a Virtual Machine to a vRealize Automation Environment to fix CSV file heading name.Updated Migrate a Virtual Machine to a Different vRealize Automation Environment to fix CSV file heading name.
EN-002077-01	Added a topic that was later removed.
EN-002077-00	Initial 7.1 release.

Managing vRealize Automation

You can manage provisioned machines and other aspects of your vRealize Automation deployment.

This section includes the following topics:

- [Broadcast a Message on the Message Board Portlet](#)
- [Starting Up and Shutting Down vRealize Automation](#)
- [Updating vRealize Automation Certificates](#)
- [Managing the vRealize Automation Postgres Appliance Database](#)
- [Backup and Recovery for vRealize Automation Installations](#)
- [Configuring the Customer Experience Improvement Program for vRealize Automation](#)
- [Adjusting System Settings](#)
- [Monitoring vRealize Automation](#)
- [Monitoring and Managing Resources](#)
- [Bulk Import, Update, or Migrate Virtual Machines](#)
- [Managing Machines](#)

Broadcast a Message on the Message Board Portlet

As the tenant administrator, you use the message board portlet to broadcast a message to all the users who have the portlet on their Home tab.

Any new users that you add to vRealize Automation has the portlet on their Home tab by default. Existing users must add the portlet to receive your messages.

You use the message board portlet to broadcast a text message or a Web page. Depending on the Web page, your users can navigate through the Web site in the message board.

The message board has the following limitations.


Table 1-1. Message Board Portlet Limitations

Option	Limitations
URL message limitations	<ul style="list-style-type: none"> ■ You can only publish content that is hosted on an https site. ■ You cannot use self-signed certificates. The option to accept the certificate does not appear in the message board. ■ The message board URL is embedded in an iframe. Some Web sites do not work in iframe and an error appears. One cause of the failure is the X-Frame-Options DENY or SAMEORIGIN in the header on the target Web site. If your target Web site is one that you control, you can set the X-Frame-Options header to X-Frame-Options: ALLOW-FROM https://<vRealizeAutomationApplianceURL>. ■ Some Web sites have a redirect to a top-level page that might refresh entire vRealize Automation page. This type of Web site does not work in the message board. The refresh is suppressed and a Loading... message appears on the message board. ■ If you display an internal HTML page, the page cannot have the vRealize Automation host as the URL.
Custom message limitations	<ul style="list-style-type: none"> ■ To maintain security, the Custom Message does not support HTML code. For example, you cannot use <href> to link to a Web site. You must use the URL message option.

Prerequisites

Log in to the vRealize Automation console as a **tenant administrator**.

Procedure

- 1 Select the **Home** tab.
- 2 Click the **Edit** icon () in the upper right corner.
- 3 Select **Add Portlets**.
- 4 Locate Message Board and click **Add**.
- 5 Click **Close**.

The portlet is added to the top of the Home tab. If you are a user and a message is broadcast, you see the message until the tenant administrator changes it or removes it. If you are the tenant administrator, you configure the message.

- 6 To configure the message as a tenant administrator, click **Add New Message**.
- 7 Configure one of the following options.

Option	Description
URL	Enter the page URL.
Custom Message	Enter the plain text message.

- 8 Click **Publish**.

The message is broadcast to any tenant users who added the message board portlet to their Home tab.

To change or remove the message, you must be logged in as the tenant administrator. To change the message, repeat the same steps. To remove the message, remove the URL or text and publish the blank message.

Starting Up and Shutting Down vRealize Automation

A system administrator performs a controlled shutdown or startup of vRealize Automation to preserve system and data integrity.

You can also use a controlled shutdown and startup to resolve performance or product behavior issues that can result from an incorrect initial startup. Use the restart procedure when only some components of your deployment fail.

Start Up vRealize Automation

When you start vRealize Automation from the beginning, such as after a power outage, a controlled shutdown or after recovery, you must start its components in a specified order.

Prerequisites

Verify that the load balancers that your deployment uses are running.

Procedure

- 1 Start the MS SQL database machine. If you are using a legacy PostgreSQL standalone database, start that machine as well.
- 2 (Optional) If you are running a deployment that uses load balancers with health checks, disable the health check before you start the vRealize Automation appliance. Only ping health check should be enabled.
- 3 Start all instances of vRealize Automation appliance at the same time and wait for approximately 15 minutes for the appliances to startup. Verify that the vRealize Automation appliance services are up and running.

If you have more than one node and you start only one node, you may have to wait for extra 35 minutes. However, the extra wait time would be canceled out as soon as you start the second node.
- 4 Start the primary Web node and wait for the startup to finish.
- 5 (Optional) If you are running a distributed deployment, start all secondary Web nodes and wait 5 minutes.
- 6 Start the primary Manager Service node and wait for 2 to 5 minutes, depending on your site configuration.

- 7 Start the Distributed Execution Manager Orchestrator and Workers and all vRealize Automation proxy agents.

You can start these components in any order and you do not need to wait for one startup to finish before you start another.

- 8 If you disabled health checks for your load balancers, reenable them.
- 9 Verify that the startup succeeded.
 - a Navigate to the vRealize Automation appliance management console by using its fully qualified domain name, `https://vra-va-hostname.domain.name:5480/`.
 - b Click the **Services** tab.
 - c Click the **Refresh** tab to monitor the progress of service startup.

When all services are listed as registered, the system is ready to use.

Restart vRealize Automation

When you restart more than one vRealize Automation component, you must restart the components in a specified order.

You might need to restart some components in your deployment to resolve anomalous product behavior. If you are using vCenter Server to manage your virtual machines, use the guest restart command to restart vRealize Automation.

If you cannot restart a component or service, follow the instructions in [Shut Down vRealize Automation](#) and [Start Up vRealize Automation](#).

Prerequisites

Verify that load balancers that your deployment uses are running.

Procedure

- 1 Restart the all instances of the vRealize Automation appliance at the same time.
- 2 Restart the primary Web node and wait for the startup to finish.
- 3 If you are running a distributed deployment, start all secondary Web nodes and wait for the startup to finish.
- 4 Restart all Manager Service nodes and wait for the startup to finish.
- 5 Restart the Distributed Execution Manager Orchestrator and Workers and all vRealize Automation agents, and wait for the startup to finish for all components.

You can restart these components in any order.

- 6 Verify that the service you restarted is registered.
 - a Navigate to the vRealize Automation appliance management console by using its fully qualified domain name, `https://vra-va-hostname.domain.name:5480/`.
 - b Click the **Services** tab.
 - c Click the **Refresh** tab to monitor the progress of service startup.

When all services are listed as registered, the system is ready to use.

Shut Down vRealize Automation

To preserve data integrity, you must shut down vRealize Automation in a specified order.

If you are using vCenter Server to manage your virtual machines, use the guest shutdown command to shut down vRealize Automation.

Procedure

- 1 Shut down the Distributed Execution Manager Orchestrator and Workers and all vRealize Automation agents in any order and wait for all components to finish shutting down.
- 2 Shut down virtual machines that are running the Manager Service and wait for the shutdown to finish.
- 3 (Optional) For distributed deployments, shut down all secondary Web nodes and wait for the shutdown to finish.
- 4 Shut down the primary Web node, and wait for the shutdown to finish.
- 5 (Optional) For distributed deployments, shut down all secondary vRealize Automation appliance instances and wait for the shutdown to finish.
- 6 Shut down the primary vRealize Automation appliance and wait for the shutdown to finish.

If applicable, the primary vRealize Automation appliance is the one that contains the master, or writeable, Appliance Database. Make a note of the name of the primary vRealize Automation appliance. You use this information when you restart vRealize Automation.

- 7 Shut down the MSSQL virtual machines in any order and wait for the shutdown to finish.
- 8 If you are using a legacy standalone PostgreSQL database, also shut down that machine.

You shut down your vRealize Automation deployment.

Updating vRealize Automation Certificates

A system administrator can update or replace certificates for vRealize Automation components.

vRealize Automation contains three main components that use SSL certificates in order to facilitate secure communication with each other. These components are as follows:

- vRealize Automation appliance
- IaaS website component

- IaaS manager service component

In addition, your deployment can have certificates for the vRealize Automation appliance management site. Also, each IaaS machine runs a Management Agent that uses a certificate.

Typically, self-signed certificates are generated and applied to these components during product installation. You might need to replace a certificate to switch from self-signed certificates to certificates provided by a certificate authority or when a certificate expires. When you replace a certificate for a vRealize Automation component, trust relationships for other vRealize Automation components are updated automatically.

For instance, in a distributed system with multiple instances of a vRealize Automation appliance, if you update a certificate for one vRealize Automation appliance all other related certificates are updated automatically.

Note vRealize Automation supports SHA2 certificates. The self-signed certificates generated by the system use SHA-256 With RSA Encryption. You may need to update to SHA2 certificates due to operating system or browser requirements.

The vRealize Automation appliance management console provides three options for updating or replacing certificates for existing deployments:

- **Generate certificate** - Use this option to have the system generate a self-signed certificate.
- **Import certificate** - Use this option if you have a certificate that you want to use.
- **Provide certificate thumbprint** - Use this option if you want to provide a certificate thumbprint to use a certificate that is already deployed in the certificate store on the IaaS servers. Using this option will not transmit the certificate from the virtual appliance to the IaaS servers. It enables users to deploy existing certificates on IaaS servers without uploading them in the vRealize Automation management console.

Also, you can select the **Keep Existing** option to keep your existing certificate.

Certificates for the vRealize Automation appliance management site do not have registration requirements.

With one exception, changes to later components in this list do not affect earlier ones. The exception is that an updated certificate for IaaS components must be registered with vRealize Automation appliance.

Note If your certificate uses a passphrase for encryption and you fail to enter it when replacing your certificate on the virtual appliance, the certificate replacement fails and the message `Unable to load private key` appears.

For important information about troubleshooting, supportability, and trust requirements for certificates, see the VMware knowledge base article at <http://kb.vmware.com/kb/2106583>.

Extracting Certificates and Private Keys

Certificates that you use with the virtual appliances must be in the PEM file format.

The examples in the following table use Gnu `openssl` commands to extract the certificate information you need to configure the virtual appliances.

Table 1-2. Sample Certificate Values and Commands (openssl)

Certificate Authority Provides	Command	Virtual Appliance Entries
RSA Private Key	<code>openssl pkcs12 -in <i>path_to_pfx_certificate_file</i> -nocerts -out key.pem</code>	RSA Private Key
PEM File	<code>openssl pkcs12 -in <i>path_to_pfx_certificate_file</i> -clcerts -nokeys -out cert.pem</code>	Certificate Chain
(Optional) Pass Phrase	n/a	Pass Phrase

Replace Certificates in the vRealize Automation appliance

The system administrator can update or replace a self-signed certificate with a trusted one from a certificate authority. You can use Subject Alternative Name (SAN) certificates, wildcard certificates, or any other method of multi-use certification appropriate for your environment as long as you satisfy the trust requirements.

When you update or replace the vRealize Automation appliance certificate, trust with other related components is re-initiated automatically. See [Updating vRealize Automation Certificates](#) for more information about updating certificates.

Procedure

- 1 Navigate to the vRealize Automation appliance management console by using its fully qualified domain name, `https://vra-va-hostname.domain.name:5480/`.
- 2 Log in with user name **root** and the password you specified when deploying the vRealize Automation appliance.
- 3 Select **vRA Settings > Host Settings**.

4 Select the certificate type from the **Certificate Action** menu.

If you are using a PEM-encoded certificate, for example for a distributed environment, select **Import**.

Certificates that you import must be trusted and must also be applicable to all instances of vRealize Automation appliance and any load balancer through the use of Subject Alternative Name (SAN) certificates.

Note If you use certificate chains, specify the certificates in the following order:

- a Client/server certificate signed by the intermediate CA certificate
- b One or more intermediate certificates
- c A root CA certificate

Option	Action
Keep Existing	Leave the current SSL configuration. Select this option to cancel your changes.
Generate Certificate	<ul style="list-style-type: none"> a The value displayed in the Common Name text box is the Host Name as it appears on the upper part of the page. If any additional instances of the vRealize Automation appliance available, their FQDNs are included in the SAN attribute of the certificate. b Enter your organization name, such as your company name, in the Organization text box. c Enter your organizational unit, such as your department name or location, in the Organizational Unit text box. d Enter a two-letter ISO 3166 country code, such as US, in the Country text box.
Import	<ul style="list-style-type: none"> a Copy the certificate values from BEGIN PRIVATE KEY to END PRIVATE KEY, including the header and footer, and paste them in the RSA Private Key text box. b Copy the certificate values from BEGIN CERTIFICATE to END CERTIFICATE, including the header and footer, and paste them in the Certificate Chain text box. For multiple certificate values, include a BEGIN CERTIFICATE header and END CERTIFICATE footer for each certificate. <p>Note In the case of chained certificates, additional attributes may be available.</p> <ul style="list-style-type: none"> c (Optional) If your certificate uses a pass phrase to encrypt the certificate key, copy the pass phrase and paste it in the Passphrase text box.

5 Click **Save Settings**.

After a few minutes, the certificate details for all applicable instances of the vRealize Automation appliance appear on the page.

- 6 If required by your network or load balancer, copy the imported or newly created certificate to the virtual appliance load balancer.

You might need to enable root SSH access in order to export the certificate.

- a If not already logged in, log in to the vRealize Automation appliance Management Console as root.
- b Click the **Admin** tab.
- c Click the **Admin** sub menu.
- d Select the **SSH service enabled** check box.
Deselect the check box to disable SSH when finished.
- e Select the **Administrator SSH login** check box.
Deselect the check box to disable SSH when finished.
- f Click **Save Settings**.

- 7 Confirm that you can log in to vRealize Automation console.

- a Open a browser and navigate to `https://vcac-hostname.domain.name/vcac/`.
If you are using a load balancer, the host name must be the fully qualified domain name of the load balancer.
- b If prompted, continue past the certificate warnings.
- c Log in with `administrator@vsphere.local` and the password you specified when configuring Directories Management.
The console opens to the **Tenants** page on the **Administration** tab. A single tenant named `vsphere.local` appears in the list.

- 8 If you are using a load balancer, configure and enable any applicable health checks.

The certificate is updated.

Replace the Infrastructure as a Service Certificate

The system administrator can replace an expired certificate or a self-signed certificate with one from a certificate authority to ensure security in a distributed deployment environment.

You can use a Subject Alternative Name (SAN) certificate on multiple machines. Certificates used for the IaaS components (Website and Manager Service) must be issued with SAN values including FQDNs of all Windows hosts on which the corresponding component is installed and with the Load Balancer FQDN for the same component.

There are three options for replacing a certificate:

- Generate certificate - Use this option to have the system generate a self-signed certificate.
- Import certificate - Use this option if you have a certificate that you want to use.

- Provide certificate thumbprint - If you accept a certificate that is signed by a CA but that certificate is not trusted by your system, you must determine whether to accept the certificate thumbprint. The thumbprint is used to quickly determine if a presented certificate is the same as another certificate, such as the certificate that was accepted previously.

Also, you can use Keep Existing to keep your existing certificate.

Procedure

- 1 Navigate to the vRealize Automation appliance management console by using its fully qualified domain name, `https://vra-va-hostname.domain.name:5480/`.
- 2 Log in with user name **root** and the password you specified when deploying the vRealize Automation appliance.
- 3 Select **vRA Settings > Certificates**.
- 4 Click **IaaS Web** on the **Component Type** menu.
- 5 Go to the **IaaS Web Certificate** pane.
- 6 Select the certificate replacement option from the **Certificate Action** menu.

If you are using a PEM-encoded certificate, for example for a distributed environment, select **Import**.

Certificates that you import must be trusted and must also be applicable to all instances of vRealize Automation appliance and any load balancer through the use of Subject Alternative Name (SAN) certificates.

Note If you use certificate chains, specify the certificates in the following order:

- a Client/server certificate signed by the intermediate CA certificate
- b One or more intermediate certificates
- c A root CA certificate

Option	Description
Keep Existing	Leave the current SSL configuration. Choose this option to cancel your changes.
Generate Certificate	<ol style="list-style-type: none"> a The value displayed in the Common Name text box is the Host Name as it appears on the upper part of the page. If any additional instances of the vRealize Automation appliance available, their FQDNs are included in the SAN attribute of the certificate. b Enter your organization name, such as your company name, in the Organization text box. c Enter your organizational unit, such as your department name or location, in the Organizational Unit text box. d Enter a two-letter ISO 3166 country code, such as US, in the Country text box.

Option	Description
Import	<p>a Copy the certificate values from BEGIN PRIVATE KEY to END PRIVATE KEY, including the header and footer, and paste them in the RSA Private Key text box.</p> <p>b Copy the certificate values from BEGIN CERTIFICATE to END CERTIFICATE, including the header and footer, and paste them in the Certificate Chain text box. For multiple certificate values, include a BEGIN CERTIFICATE header and END CERTIFICATE footer for each certificate.</p> <hr/> <p>Note In the case of chained certificates, additional attributes may be available.</p> <hr/> <p>c (Optional) If your certificate uses a pass phrase to encrypt the certificate key, copy the pass phrase and paste it in the Passphrase text box.</p>
Provide Certificate Thumbprint	<p>Use this option if you want to provide a certificate thumbprint to use a certificate that is already deployed in the certificate store on the IaaS servers. Using this option will not transmit the certificate from the virtual appliance to the IaaS servers. It enables users to deploy existing certificates on IaaS servers without uploading them in the management interface.</p>

7 Click Save Settings.

After a few minutes, the certificate details appear on the page.

Replace the IaaS Manager Service Certificate

A system administrator can replace an expired certificate or a self-signed certificate with one from a certificate authority to ensure security in a distributed deployment environment.

You can use a Subject Alternative Name (SAN) certificate on multiple machines. Certificates used for the IaaS components (Website and Manager Service) must be issued with SAN values including FQDNs of all Windows hosts on which the corresponding component is installed and with the Load Balancer FQDN for the same component.

The IaaS Manager Service and the IaaS Web Service share a single certificate.

Procedure

- 1 Navigate to the vRealize Automation appliance management console by using its fully qualified domain name, `https://vra-va-hostname.domain.name:5480/`.
- 2 Log in with user name **root** and the password you specified when deploying the vRealize Automation appliance.
- 3 Select **vRA Settings > Certificates**.
- 4 Click **Manager Service** from the **Certificate Type** menu.

5 Select the certificate type from the **Certificate Action** menu.

If you are using a PEM-encoded certificate, for example for a distributed environment, select **Import**.

Certificates that you import must be trusted and must also be applicable to all instances of vRealize Automation appliance and any load balancer through the use of Subject Alternative Name (SAN) certificates.

Note If you use certificate chains, specify the certificates in the following order:

- a Client/server certificate signed by the intermediate CA certificate
- b One or more intermediate certificates
- c A root CA certificate

Option	Description
Keep Existing	Leave the current SSL configuration. Choose this option to cancel your changes.
Generate Certificate	<ul style="list-style-type: none"> a The value displayed in the Common Name text box is the Host Name as it appears on the upper part of the page. If any additional instances of the vRealize Automation appliance available, their FQDNs are included in the SAN attribute of the certificate. b Enter your organization name, such as your company name, in the Organization text box. c Enter your organizational unit, such as your department name or location, in the Organizational Unit text box. d Enter a two-letter ISO 3166 country code, such as US, in the Country text box.
Import	<ul style="list-style-type: none"> a Copy the certificate values from BEGIN PRIVATE KEY to END PRIVATE KEY, including the header and footer, and paste them in the RSA Private Key text box. b Copy the certificate values from BEGIN CERTIFICATE to END CERTIFICATE, including the header and footer, and paste them in the Certificate Chain text box. For multiple certificate values, include a BEGIN CERTIFICATE header and END CERTIFICATE footer for each certificate. <p>Note In the case of chained certificates, additional attributes may be available.</p> <ul style="list-style-type: none"> c (Optional) If your certificate uses a pass phrase to encrypt the certificate key, copy the pass phrase and paste it in the Passphrase text box.
Provide Certificate Thumbprint	Use this option if you want to provide a certificate thumbprint to use a certificate that is already deployed in the certificate store on the IaaS servers. Using this option will not transmit the certificate from the virtual appliance to the IaaS servers. It enables users to deploy existing certificates on IaaS servers without uploading them in the management interface.

6 Click **Save Settings**.

After a few minutes, the certificate details appear on the page.

7 If required by your network or load balancer, copy the imported or newly created certificate to the load balancer.

- 8 Open a browser and navigate to `https://managerServiceAddress/vmpsProvision/` from a server that this running a DEM worker or agent.

If you are using a load balancer, the host name must be the fully qualified domain name of the load balancer.

- 9 If prompted, continue past the certificate warnings.
- 10 Validate that the new certificate is provided and is trusted.
- 11 If you are using a load balancer, configure and enable any applicable health checks.

Updating the vRealize Automation Appliance Management Site Certificate

The system administrator can replace the SSL certificate of the management site service when it expires or to replace a self-signed certificate with one issued by a certificate authority. You secure the management site service on port 5480.

The vRealize Automation appliance uses `lighttpd` to run its own management site. When you replace a management site certificate, you must also configure all Management Agents to recognize the new certificate.

If you are running a distributed deployment, you can update Management Agents automatically or manually. If you are running a minimal deployment, you must update the management agent manually.

See [Manually Update Management Agent Certificate Recognition](#) for more information.

Procedure

- 1 [Find the Management Agent Identifier](#)

You use the Management Agent identifier when you create and register a new management site server certificate.

- 2 [Replace the vRealize Automation Appliance Management Site Certificate](#)

The vRealize Automation appliance uses `lighttpd` to run its own management site. You can replace the SSL certificate of the management site service if your certificate expires or if you are using a self-signed certificate and your company security policy requires you to use its SSL certificates. You secure the management site service on port 5480.

- 3 [Update Management Agent Certificate Recognition](#)

After replacing a vRealize Automation appliance management site certificate, you must update all Management Agents to recognize the new certificate and to reestablish trusted communications between the virtual appliance management site and Management Agents on IaaS hosts.

Find the Management Agent Identifier

You use the Management Agent identifier when you create and register a new management site server certificate.

Procedure

- 1 Open the Management Agent configuration file located at `<vra-installation-dir>\Management Agent\VMware.IaaS.Management.Agent.exe.config`.
- 2 Record the value from the id attribute of the agentConfiguration element.
`<agentConfiguration id="0E22046B-9D71-4A2B-BB5D-70817F901B27">`

Replace the vRealize Automation Appliance Management Site Certificate

The vRealize Automation appliance uses lighttpd to run its own management site. You can replace the SSL certificate of the management site service if your certificate expires or if you are using a self-signed certificate and your company security policy requires you to use its SSL certificates. You secure the management site service on port 5480.

You can choose to install a new certificate or reuse the certificate used by the vCloud Automation Center service on port 443.

When you request a new certificate to update another CA-issued certificate, it is a best practice to reuse the Common Name from the existing certificate.

Prerequisites

- New certificates must be in PEM format and the private key cannot be encrypted. By default, the vRealize Automation appliance management site SSL certificate and private key are stored in a PEM file located at `/opt/vmware/etc/lighttpd/server.pem`.

See [Extracting Certificates and Private Keys](#) if you require information about exporting a certificate and private key from a Java keystore to a PEM file.

Procedure

- 1 Log in by using the appliance console or SSH.
- 2 Back up your current certificate file.

```
cp /opt/vmware/etc/lighttpd/server.pem /opt/vmware/etc/lighttpd/server.pem-bak
```

- 3 Copy the new certificate to your appliance by replacing the content of the file `/opt/vmware/etc/lighttpd/server.pem` with the new certificate information.
- 4 Run the following command to restart the lighttpd server.

```
service vami-lighttpd restart
```
- 5 Log in to the management console and validate that the certificate is replaced. You might need to restart your browser.

The new vRealize Automation appliance management site certificate is installed.

What to do next

Update all management agents to recognize the new certificate.

For distributed deployments, you can update management agents manually or automatically. For minimal installations, you must update agents manually.

- For information about automatic update, see [Automatically Update Management Agents in a Distributed Environment to Recognize a vRealize Automation Appliance Management Site Certificate](#).
- For information about manual update, see [Manually Update Management Agent Certificate Recognition](#).

Update Management Agent Certificate Recognition

After replacing a vRealize Automation appliance management site certificate, you must update all Management Agents to recognize the new certificate and to reestablish trusted communications between the virtual appliance management site and Management Agents on IaaS hosts.

Each IaaS hosts runs a Management Agent and each Management Agent must be updated. Minimal deployments must be updated manually, while distributed deployments can be updated manually or using an automated process.

- [Manually Update Management Agent Certificate Recognition](#)

After replacing a vRealize Automation appliance management site certificate, you must update Management Agents manually to recognize the new certificate to reestablish trusted communications between the virtual appliance management site and Management Agents on IaaS hosts.

- [Automatically Update Management Agents in a Distributed Environment to Recognize a vRealize Automation Appliance Management Site Certificate](#)

After the Management Site certificate is updated in a high-availability deployment, the Management Agent configuration must also be updated to recognize the new certificate and reestablish trusted communication.

Manually Update Management Agent Certificate Recognition

After replacing a vRealize Automation appliance management site certificate, you must update Management Agents manually to recognize the new certificate to reestablish trusted communications between the virtual appliance management site and Management Agents on IaaS hosts.

Perform these steps for each Management Agent in your deployment after you replace a certificate for the vRealize Automation appliance management site.

For distributed deployments, you can update Management Agents manually or automatically. For information about automatic update, see [Automatically Update Management Agents in a Distributed Environment to Recognize a vRealize Automation Appliance Management Site Certificate](#).

Prerequisites

Obtain the SHA1 thumbprints of the new vRealize Automation appliance management site certificate.

Procedure

- 1 Stop the VMware vCloud Automation Center Management Agent service.
- 2 Navigate to the Management Agent configuration file located at `[vcac_installation_folder]\Management Agent\VMware.IaaS.Management.Agent.exe.Config`, typically `C:\Program Files (x86)\VMware\vCAC\Management Agent\VMware.IaaS.Management.Agent.exe.Config`.
- 3 Open the file for editing and locate the endpoint configuration setting for the old management site certificate, which you can identify by the endpoint address.

For example:

```
<agentConfiguration id="C816CFBC-4830-4FD2-8951-C17429CEA291" pollingInterval="00:03:00">
  <managementEndpoints>
    <endpoint address="https://vra-va.local:5480"
thumbprint="D1542471C30A9CE694A512C5F0F19E45E6FA32E6" />
  </managementEndpoints>
</agentConfiguration>
```

- 4 Change the thumbprint to the SHA1 thumbprint of the new certificate.

For example:

```
<agentConfiguration id="C816CFBC-4830-4FD2-8951-C17429CEA291" pollingInterval="00:03:00">
  <managementEndpoints>
    <endpoint address="https://vra-va.local:5480"
thumbprint="8598B073359BAE7597F04D988AD2F083259F1201" />
  </managementEndpoints>
</agentConfiguration>
```

- 5 Start the VMware vCloud Automation Center Management Agent service.
- 6 Login to the virtual appliance management site and go to **vRA Settings > Cluster**.
- 7 Check the Distributed Deployment Information table to verify that the IaaS server has contacted the virtual appliance recently, which confirms that the update is successful.

Automatically Update Management Agents in a Distributed Environment to Recognize a vRealize Automation Appliance Management Site Certificate

After the Management Site certificate is updated in a high-availability deployment, the Management Agent configuration must also be updated to recognize the new certificate and reestablish trusted communication.

You can update vRealize Automation appliance management site certificate information for distributed systems manually or automatically. For information about manually updating Management Agents, see [Manually Update Management Agent Certificate Recognition](#).

Use this procedure to update the certificate information automatically.

Procedure

- 1 When Management Agents are running, replace the certificate on a single vRealize Automation appliance management site in your deployment.
- 2 Wait fifteen minutes for the Management Agent to synchronize with the new vRealize Automation appliance management site certificate.
- 3 Replace certificates on other vRealize Automation appliance management sites in your deployment. Management Agents are automatically updated with the new certificate information.

Replace a Management Agent Certificate

The system administrator can replace the Management Agent certificate when it expires or replace a self-signed certificate with one issued by a certificate authority.

Each IaaS host runs its own Management Agent. Repeat this procedure on each IaaS node whose Management Agent you want to update.

Prerequisites

- Copy the Management Agent identifier in the Node ID column before you remove the record. You use this identifier when you create the new Management Agent certificate and when you register it.
- When you request a new certificate, ensure that the Common Name (CN) attribute in the certificate subject field for the new certificate is typed in the following format:

```
VMware Management Agent 00000000-0000-0000-0000-000000000000
```

Use the string VMware Management Agent, followed by a single space and the GUID for the Management Agent in the numerical format shown.

Procedure

- 1 Stop the Management Agent service from your Windows Services snap-in.
 - a From your Windows machine, click **Start**.
 - b In the Windows Start Search box, enter `services.msc` and press Enter.
 - c Right-click **VMware vCloud Automation Center Management Agent** service and click **Stop** to stop the service.
- 2 Remove the current certificate from the machine. For information about managing certificates on Windows Server 2008 R2, see the Microsoft Knowledge Base article at <http://technet.microsoft.com/en-us/library/cc772354.aspx> or the Microsoft wiki article at <http://social.technet.microsoft.com/wiki/contents/articles/2167.how-to-use-the-certificates-console.aspx>.
 - a Open the Microsoft Management Console by entering the command `mmc.exe`.
 - b Press Ctrl + M to add a new snap-in to the console or select the option from the File drop-down menu.

- c Select **Certificates** and click **Add**.
- d Select **Computer account** and click **Next**.
- e Select **Local computer: (the computer this console is running on)**.
- f Click **OK**.
- g Expand **Certificates (Local Computer)** on the left side of the console.
- h Expand **Personal** and select the Certificates folder.
- i Select the current Management Agent certificate and click **Delete**.
- j Click **Yes** to confirm the delete action.

Example: Command to Register a Management Agent Certificate

```
Vcac-Config.exe RegisterNode -v -vamih "vra-va.eng.mycompany:5480" -cu "root" -cp
"secret" -hn "iaas.eng.mycompany" -nd "C816CFBX-4830-4FD2-8951-C17429CEA291" -tp
"70928851D5B72B206E4B1CF9F6ED953EE1103DED"
```

Change the Polling Method for Certificates

If you use commas in the OU section of the IaaS certificate, you may encounter STOMP WebSocket errors in the Manager Service log files and virtual machine provisioning may fail. You can remove the commas or change the polling method from WebSocket to HTTP to resolve these issues.

See *Installing vRealize Automation 7.1* for more information about the Manager Service.

Procedure

- 1 Open the Manager Service configuration file in a text editor.

The Manager Service configuration file is located at C:\:\Program Files (x86)\VMware\VCAC\Server\Manager Service.exe.config.

- 2 Add the following lines to the <appSettings> section of the Manager Service configuration file.

```
<add key="Extensibility.Client.RetrievalMethod" value="Polling"/>
<add key="Extensibility.Client.PollingInterval" value="2000"/>
<add key="Extensibility.Client.PollingMaxEvents" value="128"/>
```

- 3 Restart the Manager Service.

Managing the vRealize Automation Postgres Appliance Database

vRealize Automation requires the appliance database for system operation. You can manage the appliance database through the vRealize Appliance Virtual Appliance Management Interface.

Note This information applies only to deployments that use an embedded appliance database. It does not apply to deployments that use an external Postgres database.

You can configure the database as a single node system or with multiple nodes to facilitate high availability through failover. The appliance database is set up initially when you install vRealize Automation, and it requires no maintenance unless a machine configuration changes or, in the case of a clustered configuration, you need to promote a different node to be the master.

Note The database clustered configuration is set up automatically when you join a virtual appliance to the cluster using the Join cluster operation. Importantly, however, the database cluster is not directly dependent upon the virtual appliance cluster. For instance, a virtual machine joined to a cluster can operate normally even if the embedded appliance database is not started or has failed.

A clustered configuration contains one master node and one or more replica nodes. The master node is the vRealize Automation appliance node with the master database that supports system functionality. Replica nodes contain copies of the database that can be pulled into service if the master node fails.

Several high availability appliance database options exist. Choosing the replication mode is the most important database configuration option. The replication mode determines how your vRealize Automation deployment maintains data integrity and, for high availability configurations, how it fails over should the master or primary node fail. There are two available replication modes: synchronous and asynchronous.

Both replication modes support database failover, though each has advantages and disadvantages. Note that to support high availability database failover, asynchronous mode requires at least two nodes, while synchronous mode requires at least three nodes.

Replication Mode	Advantages	Disadvantages
Synchronous	Minimizes chance of data loss.	<ul style="list-style-type: none"> ■ Might affect system performance . ■ Requires at least three nodes.
Asynchronous	<ul style="list-style-type: none"> ■ Requires only two nodes. ■ Affects system performance less than synchronous mode. 	Not as robust as synchronous mode in preventing data loss.

vRealize Automation supports both modes, but operates in asynchronous mode by default and provides high availability only if there are at least two appliance database nodes. The **Database** tab on the Virtual Appliance Management Interface enables you to switch synchronization modes and to add database nodes as needed.

If you start with one node in a non-high-availability configuration, you can add nodes later as required to enhance high availability. If you have the appropriate hardware and require maximum protection against data loss, consider configuring your deployment to operate in synchronous mode.

Configure the Appliance Database

You can use the Virtual Appliance Management Interface Database page to monitor or update the configuration of the appliance database. You can also use it to change the master node designation and the synchronization mode used by the database.

The appliance database is installed and configured during vRealize Automation system installation and configuration, but you can monitor and change the configuration from the **Database** tab on the Virtual Appliance Management Interface.

The **Connection Status** text box indicates whether the database is connected to the vRealize Automation system and is functioning correctly.

If your appliance database uses multiple nodes to support failover, the table at the bottom of the page displays the nodes, and their status and indicates which node is the master. The **Replication mode** text box shows the currently configured operation mode for the system, either synchronous or asynchronous. Use this page to update appliance database configuration.

The Sync State* column in the database nodes table shows the synchronization method for the cluster. This column works with the Status column to show the state of cluster nodes. Potential status differs depending on whether the cluster uses asynchronous or synchronous replication.

Table 1-4. Sync State for Appliance Database Replication Modes

Mode	Sync State Message
Synchronous replication	Master node - no status
	Replica node - sync
	Other nodes - potential
Asynchronous replication	Master node - no status
	Other nodes - potential

The Valid column indicates whether replicas are synchronized with the master node. The master node is always valid.

The Priority column shows the position of replica nodes in relation to the master node. The master node has no priority value. When promoting a replica to become the master, select the node with the lowest priority value.

Prerequisites

- Install and configure vRealize Automation according to appropriate instructions in the *Installing vRealize Automation 7.1*.
- Log in to the vRealize Automation management console as **root**.
- Configure an appropriate embedded Postgres appliance database cluster as part of your vRealize Automation deployment.

Procedure

- 1 On the Virtual Appliance Management Interface, select **vRA Settings > Database**.
- 2 If your database uses multiple nodes, review the table at the bottom of the page and ensure that the system is operating appropriately.
 - Ensure that all nodes are listed.
 - Ensure that the appropriate node is the designated master node.

Note Do not click **Sync Mode** to change the synchronization mode of the database unless you are certain that your data is secure. Changing the sync mode without preparation may cause data loss.

- 3 To promote one of the nodes to be the master, click **Promote** in the appropriate column.
- 4 Click **Save Settings** to save your configuration if you have made any changes.

Perform Manual vRealize Automation Appliance Database Failover

When there is a problem with the vRealize Automation appliance Postgres database, you manually fail over to a replica vRealize Automation appliance node in the cluster.

Follow these steps when the Postgres database on the master vRealize Automation appliance node fails or stops running.

Prerequisites

- Configure a cluster of vRealize Automation appliance nodes. Each node hosts a copy of the embedded Postgres appliance database.

Procedure

- 1 Remove the master node IP address from the external load balancer.
- 2 Log in to the vRealize Automation appliance management interface as root.
`https://vrealize-automation-appliance-FQDN:5480`
- 3 Click **vRA Settings > Database**.
- 4 From the list of database nodes, locate the replica node with the lowest priority.
Replica nodes appear in ascending priority order.
- 5 Click **Promote** and wait for the operation to finish.
When finished, the replica node is listed as the new master node.
- 6 Correct issues with the former master node and add it back to the cluster:
 - a Isolate the former master node.
Disconnect the node from its current network, the one that is routing to the remaining vRealize Automation appliance nodes. Select another NIC for management, or manage it directly from the virtual machine management console.
 - b Recover the former master node.
Power the node on or otherwise correct the issue. For example, you might reset the virtual machine if it is unresponsive.
 - c From a console session as root, stop the vpostgres service.
`service vpostgres stop`
 - d Add the former master node back to its original network, the one that is routing to the other vRealize Automation appliance nodes.
 - e From a console session as root, restart the haproxy service.
`service haproxy restart`
 - f Log in to the new vRealize Automation appliance master node management interface as root.

- g Click **vRA Settings > Database**.
- h Locate the former master node, and click **Reset**.
- i After a successful reset, restart the former master node.
- j With the former master powered on, verify that the following services are running.

```
haproxy
horizon-workspace
rabbitmq-server
vami-lighttp
vcac-server
vco-server
```

- k Re-add the former master node to the external load balancer.

Note If a master node that was demoted to replica is still listed as master, you might need to manually re-join it to the cluster to correct the problem.

Scenario: Perform a Maintenance Database Failover

As a vRealize Automation system administrator, you must perform an appliance database maintenance failover operation.

This scenario assumes that the current master node is up and running normally. There are two database failover maintenance steps: maintenance of the master and maintenance of a replica node. When a master node has been replaced so that it becomes a replica, you should perform maintenance on it so that it is suitable to become the master again should the need arise.

Prerequisites

- vRealize Automation is installed and configured according to appropriate instructions in the *Installing vRealize Automation 7.1*.
- Log in to the vRealize Automation management console as **root**.
- You have an appropriate embedded Postgres appliance database cluster installed and configured.
- If your database uses synchronous replication mode, ensure that there are at least three active nodes in the cluster.

Procedure

- 1 Ensure that the current master node is up and running in preparation for maintenance.
- 2 Select the most suitable replica node for promotion to the master, and click **Promote**.
The old master is demoted to replica status, and the new master is promoted.
- 3 Exclude the replica virtual appliance address from the external load balancer pool.
- 4 Perform the appropriate replica maintenance.

- 5 When the maintenance is complete, ensure that the virtual appliance is running with network connectivity and that its haproxy service is running.
- 6 Click **Reset** for the replica node.

This operation resets the database so that it is configured to replicate to the current master and resynchronizes the replica node with the latest haproxy configuration from the master node.
- 7 Following successful reset, return the replica virtual appliance node IP address to the external virtual appliance load balancer IP address pool.

Backup and Recovery for vRealize Automation Installations

To minimize system downtime and data loss in the event of failures, administrators back up the entire vRealize Automation installation on a regular basis. If your system fails, you can recover by restoring the last known working backup and reinstalling some components.

To backup and restore vRealize Automation, see the following topics in the vRealize Suite documentation:

- [vRealize Automation Preparations for Backing Up](#)
- [vRealize Automation System Recovery](#)

Configuring the Customer Experience Improvement Program for vRealize Automation

This product participates in the VMware Customer Experience Improvement Program (CEIP). The CEIP provides VMware with information that enables VMware to improve its products and services, to fix problems, and to advise on how best to deploy and use our products.

As part of the CEIP, VMware collects technical information about your organizations's use of VMware products and services on a regular basis in association with your organizations's VMware license keys. This information does not personally identify any individual. For additional information regarding the CEIP, please see the Trust & Assurance Center at <http://www.vmware.com/trustvmware/ceip.html>

You can modify collection parameters and join or leave the program at any time.

Join or Leave the VMware Customer Experience Improvement Program for vRealize Automation

vRealize Automation participates in VMware's Customer Experience Improvement Program (CEIP). Details regarding the data collected through CEIP and the purposes for which it is used by VMware are set forth at the Trust & Assurance Center at <http://www.vmware.com/trustvmware/ceip.html>.

You choose whether to join or not join the program at product installation or when you upgrade. You can also join or leave the program after you install or upgrade vRealize Automation.

Procedure

- 1 Navigate to the vRealize Automation appliance management console by using its fully qualified domain name, <https://vcac-va-hostname.domain.name:5480/>.
- 2 Log in with the user name root and the password you specified when the appliance was deployed.
- 3 Click the **Telemetry** tab to choose whether to join the VMware Customer Experience Improvement Program (CEIP).

Details regarding the data collected through CEIP and the purposes for which it is used by VMware are set forth at the Trust & Assurance Center at <http://www.vmware.com/trustvmware/ceip.html>.

- Select **Join the VMware Customer Experience Improvement Program** to participate in the program.
 - Deselect **Join the VMware Customer Experience Improvement Program** to not participate in the program.
- 4 Click **Save Settings**.

When you join the program, the vRealize Automation appliance attempts to establish a connection to <https://vmware.com> and to automatically discover any proxy server that you might have configured for your vRealize Automation appliance.

Configure Data Collection Time

This product participates in the VMware Customer Experience Improvement Program (CEIP). When you join the VMware Customer Experience Improvement Program, your vRealize Automation appliance is automatically configured to send data at a specific time each week. You can change this setting at any time.

Details regarding the data collected through CEIP and the purposes for which it is used by VMware are set forth at the Trust & Assurance Center at <http://www.vmware.com/trustvmware/ceip.html>.

For more information, see [Join or Leave the VMware Customer Experience Improvement Program for vRealize Automation](#).

Procedure

- 1 Log in to the vRealize Automation appliance by using SSH.

The user name is root and the password is the password you specified when you deployed the vRealize Automation appliance.

- 2 Open the `/etc/telemetry/telemetry-collector-vami.properties` file in an editor.

- 3 Change the time at which data collections are sent by modifying the values assigned to `frequency.dow` and `frequency.hod`.

Property	Description
<code>frequency.dow=<day_of_week></code>	Specifies the day of the week on which the data collection is sent.
<code>frequency.hod=<hour_of_day></code>	Specifies the hour of the day on which the data collection is sent. Possible values are 0-23 and time is based on local time.

- 4 Save the file.
- 5 Enter the following command at a command prompt to apply the new settings to all nodes in your deployment.

```
#vcac-config telemetry-config-update --update-info
```

Adjusting System Settings

As a system administrator, you adjust logging and customize IaaS email templates. You can also manage settings that appear as defaults for each tenant, such as email servers to handle notifications. Tenant administrators can choose to override these defaults if their tenant requires different settings.

Modify the All Services Icon in the Service Catalog

You can modify the default icon in the service catalog to display a custom image. When you modify the icon, it changes for all tenants. You cannot configure tenant-specific icons for the catalog.

Commands are provided for Linux or Mac and Windows so that you can run the cURL commands on any of those operating systems.

Prerequisites

- Convert the image to a base64 encoded string. You can use a conversion tool such as www.dailycoding.com/UTILS/CONVERTER/IMAGETOBASE64.ASPX.
- cURL must be installed on the machine where you run the commands.
- You must have the credentials for a vRealize Automation user with the system administrator role.

Procedure

- 1 Set the VCAC variable in the terminal session for the cURL commands.

Operating System	Command
Linux/Mac	<code>export VCAC=<VA URL></code>
Windows	<code>set VCAC=<VA URL></code>

2 Retrieve the authentication token for the system administrator user.

Operating System	Command
Linux/Mac	<pre>curl https://\$VCAC/identity/api/tokens --insecure -H "Accept: application/json" -H 'Content-Type: application/json' --data '{"username":"<Catalog Administrator User>","password":"<password>","tenant":"vsphere.local"}'</pre>
Windows	<pre>curl https://%VCAC%/identity/api/tokens --insecure -H "Accept:application/json" -H "Content-Type:application/json" --data "{\"username\": \"<Catalog Administrator User>\", \"password\": \"<password>\", \"tenant\": \"vsphere.local\"}"</pre>

An authentication token is generated.

3 Set the authentication token variable by replacing <Auth Token> with the token string you generated in the previous step.

Operating System	Command
Linux/Mac	<pre>export AUTH="Bearer <Auth Token>"</pre>
Windows	<pre>set AUTH=Bearer <Auth Token></pre>

4 Add the base64 encoded string for the image.

Operating System	Command
Linux/Mac	<pre>curl https://\$VCAC/catalog-service/api/icons --insecure -H "Accept: application/json" -H 'Content-Type: application/json' -H "Authorization: \$AUTH" --data '{"id":"cafe_default_icon_genericAllServices","fileName":"<filename>","contentType":"image/png","image":"<IMAGE DATA as base64 string>"}'</pre>
Windows	<pre>curl https://%VCAC%/catalog-service/api/icons --insecure -H "Accept: application/json" -H "Content-Type: application/json" -H "Authorization: %AUTH%" --data "{\"id\": \"cafe_default_icon_genericAllServices\", \"fileName\": \"<filename>\", \"contentType\": \"image/png\", \"image\": \"<IMAGE DATA as base64 string>\"}"</pre>

The new services icon appears in the service catalog after approximately five minutes.

If you want to revert to the default icon, you can run the following command after you follow steps 1-3..

Operating System	Command
Linux/Mac	<pre>curl https://\$VCAC/catalog-service/api/icons/cafe_default_icon_genericAllServices --insecure -H "Authorization: \$AUTH" --request DELETE</pre>
Windows	<pre>curl https://%VCAC%/catalog-service/api/icons/cafe_default_icon_genericAllServices --insecure -H "Authorization: %AUTH%" --request DELETE</pre>

Customize Data Rollover Settings

You can enable and configure vRealize Automation data rollover settings to control how your system retains, archives, or deletes legacy data.

Use the data rollover feature to configure the maximum number of days for vRealize Automation to retain data in the IaaS SQL Server database before archiving or deleting it. By default, this feature is disabled.

Configure data rollover settings on the vRealize Automation Global Settings page. When enabled, this feature queries and removes data from the following SQL Server database tables:

- UserLog
- Audit
- CategoryLog
- VirtualMachineHistory
- VirtualMachineHistoryProp
- AuditLogItems
- AuditLogItemsProperties
- TrackingLogItems
- WorkflowHistoryInstances
- WorkflowHistoryResults

If you set `DataRolloverIsArchiveEnabled` to `True`, archive versions of the tables are created in the `dbo` schema. For example, the archive version of `UserLog` would be `UserLogArchive`, and the archive version of `VirtualMachineHistory` would be `VirtualMachineHistoryArchive`.

When enabled, the data rollover feature runs once a day at a predetermined time of 3 a.m. according to the vRealize Automation appliance time zone configuration. Using the `DataRolloverMaximumAgeInDays` setting, you can set the maximum number of days that you want to retain the data.

If `DataRolloverIsArchiveEnabled` is set to `True`, data older than that specified in the `DataRolloverMaximumAgeInDays` is moved to the archive tables. If `DataRolloverIsArchiveEnabled` is set to `False`, data is permanently deleted and no data archiving occurs. Deleted data is not recoverable.


Note Consider existing system data and the potential impact on system performance before enabling data rollover. For example, if you enable this feature one year after vRealize Automation began running in your environment, verify that you have set the value of `DataRolloverMaximumAgeInDays` to 300 or greater to ensure that enabling data rollover feature does not impact system performance.

Procedure

- 1 Log in to the vRealize Automation console as a **system administrator**.
- 2 Select **Infrastructure > Administration > Global Settings**.

- 3 On the Global Settings page, locate the Data Rollover section of the table and review and configure settings.

Setting	Description
DataRollover IsArchiveEnabled	<p>Specifies whether to move rollover data to archive tables after the maximum number of days is reached.</p> <p>By default this value is set to True.</p> <p>If you set this value to False, all data older than that specified in the DataRollover MaximumAgeInDays setting is permanently deleted.</p>
DataRollover MaximumAgeInDays	<p>Specifies the maximum number of days that the system retains data in the database before moving it to archive or permanently deleting it.</p> <p>By default this value is set to 90 days.</p>
DataRollover Status	<p>Specifies whether to enable data rollover.</p> <p>To enable data rollover, set the value to Enabled. By default this value is set to Disabled.</p> <p>If you disable this workflow while it is running, the current workflow is not impacted, but the next workflow is disabled.</p>

- 4 Click the **Edit** icon () in the first table column to edit a setting.

The Value field for the applicable setting becomes editable and you can place your cursor within it to change the value.

- 5 Click the **Save** icon () in the first table column to save your changes.

Adjusting Settings in the Manager Service Configuration File

You can use the manager service configuration file (`managerService.exe.config`) to adjust common settings for machine deployments.

The `managerService.exe.config` file is typically located in the `%System-Drive%\Program Files x86\VMware\vCAC\Server` directory. You should always make a copy of the file before editing it.

You can use the following `managerService.exe.config` file settings to control various aspects of machine deployments. Default values are shown.

- `<add key="ProcessLeaseWorkflowTimerCallbackIntervalMilliseconds" value="600000"/>`
- `<add key="BulkRequestWorkflowTimerCallbackMilliseconds" value="10000"/>`
- `<add key="MachineRequestTimerCallbackMilliseconds" value="10000"/>`
- `<add key="MachineWorkflowCreationTimerCallbackMilliseconds" value="10000"/>`
- `<add key="RepositoryConnectionMaxRetryCount" value="100"/>`
- `<add key="MachineCatalogRegistrationRetryTimerCallbackMilliseconds" value="120000"/>`

- `<add key="MachineCatalogUnregistrationRetryTimerCallbackMilliseconds" value="120000"/>`
- `<add key="MachineCatalogUpdateMaxRetryCount" value="15"/>`

Setting Resource-Intensive Concurrency Limits

To conserve resources, vRealize Automation limits the number of concurrently running instances of machine provisioning and data collection. You can change the limits.

Configuring Concurrent Machine Provisioning

Multiple concurrent requests for machine provisioning can impact the performance of vRealize Automation. You can make some changes to limits placed on proxy agents and workflow activities to alter performance.

Depending on the needs of machine owners at your site, the vRealize Automation server may receive multiple concurrent requests for machine provisioning. This can happen under the following circumstances:

- A single user submits a request for multiple machines
- Many users request machines at the same time
- One or more group managers approve multiple pending machine requests in close succession

The time required for vRealize Automation to provision a machine generally increases with larger numbers of concurrent requests. The increase in provisioning time depends on three important factors:

- The effect on performance of concurrent resource-intensive vRealize Automation workflow activities, including the SetupOS activity (for machines created within the virtualization platform, as in WIM-based provisioning) and the Clone activity (for machines cloned within the virtualization platform).
- The configured vRealize Automation limit on the number of resource-intensive (typically lengthy) provisioning activities that can be executed concurrently. By default this is eight. Concurrent activities beyond the configured limit are queued.
- Any limit within the virtualization platform or cloud service account on the number of vRealize Automation work items (resource-intensive or not) that can be executed concurrently. For example, the default limit in vCenter Server is four, with work items beyond this limit being queued.

By default, vRealize Automation limits concurrent virtual provisioning activities for hypervisors that use proxy agents to eight per endpoint. This ensures that the virtualization platform managed by a particular agent never receives enough resource-intensive work items to prevent execution of other items. Plan to carefully test the effects of changing the limit before making any changes. Determining the best limit for your site may require that you investigate work item execution within the virtualization platform as well as workflow activity execution within vRealize Automation.

If you do increase the configured vRealize Automation per-agent limit, you may have to make additional configuration adjustments in vRealize Automation, as follows:

- The default execution timeout intervals for the SetupOS and Clone workflow activities are two hours for each. If the time required to execute one of these activities exceeds this limit, the activity is cancelled and provisioning fails. To prevent this failure, increase one or both of these execution timeout intervals.
- The default delivery timeout intervals for the SetupOS and Clone workflow activities are 20 hours for each. Once one of these activities is initiated, if the machine resulting from the activity has not been provisioned within 20 hours, the activity is cancelled and provisioning fails. Therefore, if you have increased the limit to the point at which this sometimes occurs, you will want to increase one or both of these delivery timeout intervals.

Configuring Concurrent Data Collections

By default, vRealize Automation limits concurrent data collection activities. If you change this limit, you can avoid unnecessary timeouts by changing the default execution timeout intervals for the different types of data collection.

vRealize Automation regularly collects data from known virtualization compute resources through its proxy agents and from cloud service accounts and physical machines through the endpoints that represent them. Depending on the number of virtualization compute resources, agents, and endpoints in your site, concurrent data collection operations may occur frequently.

Data collection running time depends on the number of objects on endpoints including virtual machines, datastores, templates, and compute resources. Depending on many conditions, a single data collection can require a significant amount of time. As with machine provisioning, concurrency increases the time required to complete data collection.

By default, concurrent data collection activities are limited to two per agent, with those over the limit being queued. This ensures that each data collection completes relatively quickly and that concurrent data collection activities are unlikely to affect IaaS performance.

Depending on the resources and circumstances at your site, however, it may be possible to raise the configured limit while maintaining fast enough performance to take advantage of concurrency in proxy data collection. Although raising the limit can increase the time required for a single data collection, this might be outweighed by the ability to collect more information from more compute resources and machines at one time.

If you do increase the configured per-agent limit, you might have to adjust the default execution timeout intervals for the different types of data collection that use a proxy agent—inventory, performance, state, and WMI. If the time required to execute one of these activities exceeds the configured timeout intervals, the activity is canceled and restarted. To prevent cancellation of the activity, increase one or more of these execution timeout intervals.

Adjust Concurrency Limits and Timeout Intervals

You can change the per-agent limits on concurrent provisioning, data collection activities, and the default timeout intervals.

When typing a time value for these variables, use the format hh:mm:ss (hh=hours, mm=minutes, and ss=seconds).

Prerequisites

Log in as an administrator to the server hosting the IaaS Manager Service. For distributed installations, this is the server on which the Manager Service was installed.

Procedure

- 1 Open the `ManagerService.exe.config` file in an editor. The file is located in the vRealize Automation server install directory, typically `%SystemDrive%\Program Files x86\VMware\vCAC\Server`.
- 2 Locate the section called `workflowTimeoutConfigurationSection`.
- 3 Update the following variables, as required.

Parameter	Description
<i>MaxOutstandingResourceIntensiveWorkItems</i>	Concurrent provisioning limit (default is 8)
<i>CloneExecutionTimeout</i>	Virtual provisioning execution timeout interval
<i>SetupOSExecutionTimeout</i>	Virtual provisioning execution timeout interval
<i>CloneTimeout</i>	Virtual provisioning clone delivery timeout interval
<i>SetupOSTimeout</i>	Virtual provisioning setup OS delivery timeout interval
<i>CloudInitializeProvisioning</i>	Cloud provisioning initialization timeout interval
<i>MaxOutstandingDataCollectionWorkItems</i>	Concurrent data collection limit
<i>InventoryTimeout</i>	Inventory data collection execution timeout interval
<i>PerformanceTimeout</i>	Performance data collection execution timeout interval
<i>StateTimeout</i>	State data collection execution timeout interval

- 4 Save and close the file.
- 5 Select **Start > Administrative Tools > Services**.
- 6 Stop and then restart the vRealize Automation service.
- 7 (Optional) If vRealize Automation is running in High Availability mode, any changes made to the `ManagerService.exe.config` file after installation must be made on both the primary and failover servers.

Adjust Execution Frequency of Machine Callbacks

You can change the frequency of several callback procedures, including the frequency that the vRealize Automation callback procedure is run for changed machine leases.

vRealize Automation uses a configured time interval to run different callback procedures on the Model Manager service, such as *ProcessLeaseWorkflowTimerCallbackIntervalMiliSeconds* which searches for machines whose leases have changed. You can change these time intervals to check more or less frequently.

When entering a time value for these variables, enter a value in milliseconds. For example, 10000 milliseconds = 10 seconds and 3600000 milliseconds = 60 minutes = 1 hour.

Prerequisites

Log in as an administrator to the server hosting the IaaS Manager Service. For distributed installations, this is the server on which the Manager Service was installed.

Procedure

- 1 Open the `ManagerService.exe.config` file in an editor. The file is located in the vRealize Automation server install directory, typically `%SystemDrive%\Program Files x86\VMware\vCAC\Server`.
- 2 Update the following variables, as desired.

Parameter	Description
<i>RepositoryWorkflowTimerCallbackMiliSeconds</i>	Checks the repository service, or Model Manager Web Service, for activity. Default value is 10000.
<i>ProcessLeaseWorkflowTimerCallbackIntervalMiliSeconds</i>	Checks for expired machine leases. Default value is 3600000.
<i>BulkRequestWorkflowTimerCallbackMiliSeconds</i>	Checks for bulk requests. Default value is 10000.
<i>MachineRequestTimerCallbackMiliSeconds</i>	Checks for machine requests. Default value is 10000.
<i>MachineWorkflowCreationTimerCallbackMiliSeconds</i>	Checks for new machines. Default value is 10000.

- 3 Save and close the file.
- 4 Select **Start > Administrative Tools > Services**.
- 5 Stop and then restart the vCloud Automation Center service.
- 6 (Optional) If vRealize Automation is running in High Availability mode, any changes made to the `ManagerService.exe.config` file after installation must be made on both the primary and failover servers.

Adjust IaaS Log Settings

You can adjust vRealize Automation to log only the information you want to see in the Manager Service log.

If vRealize Automation is running in high availability mode, and you make changes to the `ManagerService.exe.config` file after installation, you must make the changes on the primary and the failover vRealize Automation servers.

Procedure

- 1 Log in to the vRealize Automation server by using credentials with administrative access.
- 2 Edit the `ManagerService.exe.config` file in `%SystemDrive%\Program Files (x86)\VMware\vCAC\Server`, or in the vRealize Automation server install directory, if it is in a different location.
- 3 Edit the `RepositoryLogSeverity` and `RepositoryLogCategory` keys to configure what types of events get written to your log files.

Option	Description
RepositoryLogSeverity	Specify a severity level to ignore events below that severity. <ul style="list-style-type: none"> ■ <i>Error</i> logs only recoverable errors and higher ■ <i>Warning</i> logs noncritical warnings and higher ■ <i>Information</i> logs all informative messages and higher ■ <i>Verbose</i> logs a debugging trace and can impair performance For example, <code><add key="RepositoryLogSeverity" value="Warning" /></code> .
RepositoryLogCategory	Specify a category to log all events for that category regardless of severity. For example, <code><add key="RepositoryLogCategory" value="MissingMachines,UnregisteredMachines,AcceptMachineRequest,RejectMachineRequest" /></code> logs all events for missing or unregistered machines, and every accepted or rejected machine request.

- 4 Save and close the file.
- 5 Select **Start > Administrative Tools > Services** and restart the vCloud Automation Center service.

You can see how your changes effect logging by viewing the Manager Service log file located in `%SystemDrive%\Program Files (x86)\VMware\vCAC\Server\Logs` on the machine where the Manager Service is installed, or in the vRealize Automation server install directory, if you installed it in a different location.

Monitoring vRealize Automation

Depending on your role, you can monitor workflows or services, view event or audit logs, or collect logs for all the hosts in a distributed deployment.

Monitoring Workflows and Viewing Logs

Depending on your role, you can monitor workflows and view activity logs.

Table 1-5. Monitoring and Log Display Options

Objective	Role	Menu Sequence and Description
Display information about actions that have occurred, such as the action type, date and time of the action, and so on.	laaS administrator	Display default log information or control display content using column and filter options. Select Infrastructure > Monitoring > Audit Log . The audit log provides details about the status of managed virtual machines and activities performed on these machines during reconfiguration. The log includes information about machine provisioning, vCloud Networking and Security, reclamation, and reconfigure actions.
View the status of scheduled and available Distributed Execution Manager and other workflows.	laaS administrator	Display workflow status and optionally open a specific workflow to display its details. Select Infrastructure > Monitoring > DEM Status .
View and optionally export log data.	laaS administrator	Display default log information or control display content using column and filter options. Select Infrastructure > Monitoring > Log .
View the status and history of executed Distributed Execution Manager and other workflows.	laaS administrator	Display workflow history and optionally open a specific workflow to display its execution details. Select Infrastructure > Monitoring > Workflow History .
Display a list of events, including event type, time, user ID, and so on, and optionally display an event details page.	System administrator	View a list of events and their associated attributes, such as run time, event description, tenant name, target type and ID, and other characteristics. Select Administration > Events > Event Logs .
Monitor the status of your requests and view request details.	Tenant administrator or business group manager	Display the status of requests that you are responsible for or own. Click Requests .
View information about recent events.	laaS administrator or Tenant administrator	Display recent events for the currently logged in user. Select Infrastructure > Recent Events

Monitoring Event Logs and Services

You can monitor vRealize Automation event logs and services to determine their current and historic states.

For information about clearing logs by customizing data rollover settings, see *Configuring vRealize Automation*.

vRealize Automation Services

A system administrator can view the status of vRealize Automation services from the Event Log on the system administrator console.

Subsets of services are required to run individual product components. For example, identity services and UI core services must be running before you can configure a tenant.

The following tables tell you which services are associated with areas of vRealize Automation functionality.

Table 1-6. Identity Service Group

Service	Description
management-service	Identity Service Group
sts-service	Single Sign-on Appliance
authorization	Authorization Service
authentication	Authentication
eventlog-service	Event log service
licensing-service	Licensing service

Table 1-7. UI Core services

Service	Description
shel-ui-app	Shell Service
branding-service	Branding Service
plugin-service	Extensibility (Plug-in) Service
portal-service	Portal Service

All the following services are required to run the IaaS component.

Table 1-8. Service Catalog Group (Governance Services)

Service	Description
notification-service	Notification service
workitem-service	Work Item service
approval-service	Approval Service
catalog-service	Service Catalog

Table 1-9. IaaS Services Group

Service	Description
iaas-proxy-provider	IaaS Proxy
iaas-server	IaaS Windows machine

Table 1-10. XaaS

Service	Description
vco	vRealize Orchestrator
advanced-designer-service	XaaS blueprints and resource actions

Viewing Host Information for Clusters in Distributed Deployments


You can collect logs for all nodes that are clustered in a distributed deployment from the vRealize Automation appliance management console.

You can also view information for each host in your deployment. The **Cluster** tab on the vRealize Automation management console includes a Distributed Deployment Information table that displays the following information:

- A list of all nodes in your deployment
- The host name for the node. The host name is given as a fully qualified domain name.
- The time since the host last replied to the management console. Nodes for IaaS components report availability every three minutes and nodes for virtual appliances report every nine minutes.
- The vRealize Automation component type. Identifies whether the node is a virtual appliance or an IaaS server.

Figure 1-1. Distributed Deployment Information table

Collect Logs

 Save logs from all nodes connected to this cluster.

Collect Logs

There are no collected logs.

Node ID	Host	Last Connected	Type
cafe.node.546174677.31946	vcac-be.eng.vmware.com	4 minutes ago	VA
4CBC2D96-03C8-42D1-9927-2161C8CDB572	vcac-vm387.eng.vmware.com	39 seconds ago	IAAS

You can use this table to monitor activity in your deployment. For example, if the Last Connected column indicates a host has not connected recently, that can be an indication of a problem with the host server.

Log Collection

You can create a zip file that contains log files for all hosts in your deployment. For more information, see [Collect Logs for Clusters and Distributed Deployments](#).

Removing Nodes from the Table

When you remove a host from your deployment, remove the corresponding node from the Distributed Deployment Information table to optimize log collection times. .

Collect Logs for Clusters and Distributed Deployments

You can create a zip file that includes all log files for servers in your deployment.

The Distributed Deployment Information table lists the nodes from which log files are collected.

For related information about vRealize Automation appliance deployment configuration, see *Installing vRealize Automation 7.1*.

Procedure

- 1 Log in to the vRealize Automation appliance with user name **root** and the password you specified when deploying the appliance.
- 2 Click **vRA Settings**.
- 3 Click the **Cluster** tab.

The Distributed Deployment Information table displays a list of nodes for the distributed deployment.

- 4 Click **Collect Logs**.

Log files for each node are collected and copied to a zip file.

Remove a Node from the Distributed Deployment Information Table

You delete the entry for a node from the Distributed Deployment Information table when the node is removed from your deployment cluster or when you are replacing a Management Agent certificate.

Procedure

- 1 Log in to the vRealize Automation appliance by using the user name **root** and the password you specified when you deployed the appliance.
- 2 Click **vRA Settings**.
- 3 Click the **Cluster** tab.

The Distributed Deployment Information table displays a list of nodes for the distributed deployment.

- 4 Locate the node ID for the node to be deleted and copy the ID to use in the next step.
- 5 Open a command prompt and type a command of the following form, using the node ID you previously copied.

```
/usr/sbin/vcac-config cluster-config-node
--action delete --id node-UID
```

6 Click **Refresh**.

The node no longer appears in the display.

Monitoring and Managing Resources

Different vRealize Automation roles monitor resource usage and manage infrastructure in different ways.

Choosing a Resource Monitoring Scenario

Fabric administrators, tenant administrators, and business group managers have different concerns when it comes to resource monitoring. Because of this, vRealize Automation allows you to monitor different facets of resource usage.

For example, a fabric administrator is concerned with monitoring the resource consumption of reservations and compute resources, whereas a tenant administrator is concerned with the resource usage of the provisioning groups within a tenant. Depending on your role and the specific resource usage you want to monitor, vRealize Automation allows you different ways to track resource consumption.

Table 1-11. Choose a Resource Monitoring Scenario

Resource Monitoring Scenario	Privileges Required	Location
Monitor the amount of physical storage and memory on your compute resources that is currently being consumed and determine what amount remains free. You can also monitor the number of reserved and allocated machines provisioned on each compute resource.	Fabric Administrator (monitor resource usage on compute resources in your fabric group)	Infrastructure > Compute Resources > Compute Resources
Monitor machines that are currently provisioned and under vRealize Automation management.	Fabric Administrator	Infrastructure > Machines > Managed Machines
Monitor the amount of storage, memory, and machine quota of your reservation that is currently allocated and determine the capacity that remains available to the reservation.	Fabric Administrator (monitor resource usage for reservations on your compute resources and physical machines)	Infrastructure > Reservations > Reservations
Monitor the amount of storage, memory, and the machine quota that your business groups are currently consuming and determine the capacity that remains on reserve for them.	<ul style="list-style-type: none"> ■ Tenant Administrator (monitor resource usage for all groups in your tenant) ■ Business Group Manager (monitor resource usage for groups that you manage) 	Administration > Users & Groups > Business Groups

You can also add resource monitoring portlets to your vRealize Automation homepage to monitor different resource usage statistics.

Managing Resource Reports

You can add real-time resource reports to your Home page to monitor virtual, physical, and cloud resource usage, change their layout, and export their data to other applications.


Add Reports to the Home Page

You can add one or more IaaS reports to your Home page. These real-time reports list your most recent open tasks, catalog requests, provisioned items, and provisioned machines broken down by user, blueprint, compute resource, and business group. Two reports also display updated summaries of reclamation savings.

Prerequisites

Log in to the vRealize Automation console.

Procedure

- 1 Navigate to the **Home** page.
- 2 Click the Edit  icon in the upper-right corner of the page and click **Add Portlets** in the drop-down menu.
- 3 Click **Add** for each report to add to your Home page.
A disabled **Add** button indicates an already added report.
- 4 Click **Close**.

What to do next

[Configure the Report Layout.](#)


Configure the Report Layout

You can configure your Home page to display reports in one, two, three, or four columns. You can move a report from one column to another.

Prerequisites

Log in to the vRealize Automation console.

Procedure

- 1 Navigate to the **Home** page.
- 2 Click the Edit icon () in the upper-right corner of the page and click **Change Layout** in the drop-down menu.

3 Select a report layout.

Option	Description
1 Column	Lay out reports in one column.
2 Columns	Lay out reports in two columns of equal or unequal widths.
3 Columns	Lay out reports in three columns of equal or unequal widths.
4 Columns	Lay out reports in four equal columns.

4 Click **Submit**.

5 Point to the title bar of a report.

The cursor changes to a four-headed cursor.

6 Drag the report to its new location.

The width of the report changes to fit the new location.

Export Report Data

You can save IaaS reports located on your Home page to CSV files where you can customize the data.

Prerequisites

- Log in to the vRealize Automation console.
- [Add Reports to the Home Page](#).

Procedure

1 Navigate to the **Home** page.

2 Click **Export as CSV** in the report to save.

Some browsers save the file immediately. With Firefox, a dialog box appears with selections for opening or saving the report with Microsoft Excel or another application.

3 (Optional) Select whether to open or save the report data, and which application to use.

Resource Reports

Resource reports display data about machines and resources used and reclaimed according to owner, compute resource, and group.

Name	Description
My Inbox	Displays a list of the most recent open tasks in your inbox. Click a row to view the detail page of a task. Click More to open the complete list of inbox tasks.
My Open Requests	Displays a list of your most recent catalog requests. Click a row to view the detail page of a request. Click More to open the complete list of requests.
My Recent Requests	Displays a list of your most recent catalog requests regardless of status. Click a row to view the detail page of a request. Click More to open the complete list of requests.
My Items	Displays a list of your most recently provisioned items. Click a row to view the detail page of an item. Click More to open the complete list of items.

Name	Description
My Group Requests	Displays a list of the most recent catalog requests for users in groups that you manage. Click a row to view the detail page of a request. Click More to open the complete list of requests.
My Groups Items	Displays a list of the most recently provisioned items for users in groups that you manage. Click a row to view the detail page of an item. Click More to open the complete list of items.
New & Noteworthy	Highlights catalog items that were recently made available in the catalog.
Calendar of Events	Displays a calendar view of important events for catalog items that you own, such as lease expiration and machine destruction.
Business Groups Resource Allocation	Displays the resource allocations for business groups in a tenant. If you are a tenant administrator, the portlet displays the resource allocations for all the tenant business groups. If you are business group manager, the portlet displays the resource allocation for your business groups.
IaaS Capacity Usage by Blueprint	Displays the number of machines provisioned from each blueprint and the total resources that those machines used.
IaaS Capacity Usage by Group	Displays the number of machines that users own in each business group and the total resources that those machines use.
IaaS Capacity Usage by Owner	Displays the number of machines that each user owns and the total resources that those machines use.
IaaS Capacity Usage by Compute Resource	Displays the number of machines provisioned on each compute resource and the total resources that those machines use.
My Trips	Displays a sample consumer report.

Add the Business Groups Resource Allocation Portlet to the Home Tab

The Business Group Resource Allocation Portlet is a dashboard portlet that you add to your **Home** tab to monitor resources for business groups.


If you are a tenant administrator, the portlet displays the resource allocations for all the tenant business groups. If you are business group manager, the portlet displays the resource allocation for your business groups.

If you are not a tenant administrator or business group manager, the portlet is not available to install on your **Home** tab.

Prerequisites

Log in to the vRealize Automation console as a **tenant administrator** or **business group manager**.

Procedure

- 1 Select **Home**.
- 2 Click the **Edit** icon () in the upper right corner.
- 3 Select **Add Portlets**.
- 4 Locate Business Groups Resource Allocation and click **Add**.
- 5 Click **Close**.

The portlet is added to the top of the Home tab.

- 6 Click and drag to portlet title bar to move to a different location.

Resource Usage Terminology

vRealize Automation uses explicit terminology to distinguish between resources that are available, resources that have been set aside for specific usages, and resources that are actively being consumed by provisioned machines.

The Resource Usage Terminology table explains the terminology vRealize Automation uses to display resource usage.

Table 1-12. Resource Usage Terminology

Term	Description
Physical	Indicates the actual memory or storage capacity of a compute resource.
Reserved	Indicates the machine quota, memory, and storage capacity set aside for a reservation. For example, if a compute resource has a physical capacity of 600 GB and there are three reservations on it for 100 GB each, then the reserved storage of the compute resource is 300 GB and the storage reserved is 50 percent.
Managed	Indicates that the machine is provisioned and currently under vRealize Automation management.
Allocated	Indicates the machine quota, memory, or storage resources actively being consumed by provisioned machines. For example, consider a reservation with a machine quota of 10. If there are 15 provisioned machines on it, but only 6 of them are currently powered on, the machine quota is 60 percent allocated.
Used	The Used column value always equals the Allocated column value.
Free	Indicates the unused physical capacity on a storage path.

Connecting to a Cloud Machine

The first time you connect to a cloud machine you must log in as Administrator.

You can then add the credentials under which you log in to the vRealize Automation console as a user on the machine, and log in under your vRealize Automation credentials from that point on.

Important If you are using Amazon Web Services, RDP, or SSH must be enabled on the Amazon machine instance and the machines must be in a security group in which the correct ports are open.

Collect User Credentials for an Amazon Machine

To log in to an Amazon machine as an administrator, you must discover the machine's administrator password.

The administrator password is available on the Machine Information Details page. If the Amazon machine image from which the machine was provisioned is not configured to generate the administrator password on every boot, you will need to find the password using an alternate technique. For information about otherwise obtaining the administrator password, search on *Connect to Your Amazon EC2 Instance* topics in Amazon documentation.

If needed, you can create the necessary vRealize Automation user credentials. The user credentials are then valid for subsequent logins to that machine.

Prerequisites

- The Amazon machine has already been provisioned.
- Log in to the vRealize Automation console as a machine owner, business group manager, or support user.
- RDP or SSH is active on the Amazon machine image that will be used for provisioning
- The machines are in a security group in which the correct ports are open.

Procedure

1 Navigate to the **Items** page and filter on the groups you manage or a specific group.

2 Select the Amazon machine in the list of machines.

You can click **View Details** on the **Actions** drop-down menu to display details such as machine type.

3 Select **Edit** in the **Actions** drop-down menu.

4 Click **Show Administrator Password** to obtain the administrator password of the machine.

Alternatively, you can obtain the password using an external Amazon procedure.

5 Click **Connect Using RDP** from the **Actions** drop-down menu.

6 Click **User another account** when prompted for the login credentials.

7 Type **LOCAL\Administrator** when prompted for the user name.

8 Type the administrator password when prompted.

9 Click **OK**.

You are now logged in to the machine as an administrator.

10 Add your vRealize Automation credentials as appropriate. For example, on a Windows server machine, open the server manager and select **Configuration > Local Users and Groups** and add your credentials, using a **DOMAIN\username** format, to the **Remote Desktop Users** group.

Your vRealize Automation user name and password are now valid credentials for subsequent login to this machine.

11 Log out of the Amazon machine.

12 Click **Connect Using RDP** from the **Actions** drop-down menu.

13 When prompted to log in, type your vRealize Automation user name and password credentials to log in to the machine.

Machine owners can now log in to the machine using their vRealize Automation credentials.

Collect User Credentials for a vCloud Machine

To log in to an vCloud Air or vCloud Director machine as an administrator, you must discover the machine's administrator password.

The administrator password is available on the Machine Information Details page. If the machine image from which the machine was provisioned is not configured to generate the administrator password on every boot, you can find the password using an alternate technique. For information about otherwise obtaining the administrator password, see vCloud Air or vCloud Director documentation.

If needed, you can create the necessary vRealize Automation user credentials. The user credentials are then valid for subsequent logins to that machine.

Prerequisites

- The vCloud Air or vCloud Director machine has already been provisioned.
- Log in to the vRealize Automation console as a machine owner, business group manager, or support user.
- RDP or SSH is active on the vCloud Air or vCloud Director machine image that will be used for provisioning
- The machines are in a security group in which the correct ports are open.

Procedure

- 1 Navigate to the **Items** page and filter on the groups you manage or a specific group.
- 2 Select the vCloud Air or vCloud Director machine in the list of machines.
You can click **View Details** on the **Actions** drop-down menu to display details such as machine type.
- 3 Select **Edit** in the **Actions** drop-down menu.
- 4 Click **Show Administrator Password** to obtain the administrator password of the machine.
Alternatively, you can obtain the password using an external vCloud Air or vCloud Director procedure.
- 5 Click **Connect Using RDP** from the **Actions** drop-down menu.
- 6 Click **User another account** when prompted for the login credentials.
- 7 Type **LOCAL\Administrator** when prompted for the user name.
- 8 Type the administrator password when prompted.
- 9 Click **OK**.

You are now logged in to the machine as an administrator.

- 10 Add your vRealize Automation credentials as appropriate. For example, on a Windows server machine, open the server manager and select **Configuration > Local Users and Groups** and add your credentials, using a **DOMAIN\username** format, to the **Remote Desktop Users** group.

Your vRealize Automation user name and password are now valid credentials for subsequent login to this machine.

- 11 Log out of the vCloud Air or vCloud Director machine.
- 12 Click **Connect Using RDP** from the **Actions** drop-down menu.
- 13 When prompted to log in, type your vRealize Automation user name and password credentials to log in to the machine.

Machine owners can now log in to the machine using their vRealize Automation credentials.

Reducing Reservation Usage by Attrition

Fabric administrators can reduce the number of machines on a particular reservation over the long term while keeping the reservation and the existing machines provisioned on it active.

You can reduce the reserved machine quota, memory, and storage of a virtual reservation below the amount currently allocated. This allows management of existing machines to continue without change while preventing provisioning of new machines until allocation falls below the new reserved amount.

Note Because virtual machines that are powered off are not included in allocated memory and machine quota totals, reducing the memory or machine allocation of a reservation might prevent machines that are currently powered off from being powered back on.

For example, consider a business group with a reservation that contains 20 provisioned machines that are set to expire over the next 90 days. If you want to reduce this reservation by attrition to no more than 15 machines, you can edit the reservation to reduce the quota from 20 machines to 15. No further machines can be provisioned on the reservation until the number of machines on the reservation is naturally reduced by the upcoming expirations.

Decommissioning a Storage Path

If you are decommissioning a storage path and moving machines to a new one, a fabric administrator must disable the storage path in vRealize Automation.

The following is a high-level overview of the sequence of steps required to decommission a storage path:

- 1 A fabric administrator disables the storage path on all reservations that use it. See [Disable a Storage Path](#).
- 2 Move the machines to a new storage path outside of vRealize Automation.
- 3 Wait for vRealize Automation to automatically run inventory data collection or initiate inventory data collection manually. See [Configure Compute Resource Data Collection](#).

Disable a Storage Path



Fabric administrators can disable storage paths on reservations when storage paths are decommissioned.

Note For each reservation where you disable a storage path, verify that there is sufficient space remaining on other enabled storage paths.

Prerequisites

Log in to the vRealize Automation console as a **fabric administrator**.

Procedure

- 1 Select **Infrastructure > Reservations > Reservations**.
- 2 Point to the reservation on which the storage path you are decommissioning is used and click **Edit**.
- 3 Click the **Resources** tab.
- 4 Locate the storage path you are decommissioning.
- 5 Click the **Edit** icon ().
- 6 Select the check box in the Disabled column to disable this storage path.
- 7 Click the **Save** icon ().
- 8 Click **OK**.
- 9 Repeat this procedure for all reservations that use the storage path you are decommissioning.

Data Collection

vRealize Automation collects data from infrastructure source endpoints and their compute resources.

Data collection occurs at regular intervals. Each type of data collection has a default interval that you can override or modify. Each type of data collection also has a default timeout interval that you can override or modify.

IaaS administrators can manually initiate data collection for infrastructure source endpoints and fabric administrators can manually initiate data collection for compute resources.

Table 1-13. Data Collection Types

Data Collection Type	Description
Infrastructure Source Endpoint Data Collection	Updates information about virtualization hosts, templates, and ISO images for virtualization environments. Updates virtual datacenters and templates for vCloud Director. Updates Amazon regions and machines provisioned on Amazon regions. Endpoint data collection runs every 4 hours.
Inventory Data Collection	Updates the record of the virtual machines whose resource use is tied to a specific compute resource, including detailed information about the networks, storage, and virtual machines. This record also includes information about unmanaged virtual machines, which are machines provisioned outside of vRealize Automation. Inventory data collection runs every 24 hours. The default timeout interval for inventory data collection is 2 hours.
State Data Collection	Updates the record of the power state of each machine discovered through inventory data collection. State data collection also records missing machines that vRealize Automation manages but cannot be detected on the virtualization compute resource or cloud endpoint. State data collection runs every 15 minutes. The default timeout interval for state data collection is 1 hour.
Performance Data Collection (vSphere compute resources only)	Updates the record of the average CPU, storage, memory, and network usage for each virtual machine discovered through inventory data collection. Performance data collection runs every 24 hours. The default timeout interval for performance data collection is 2 hours.
Network and security inventory data collection (vSphere compute resources only)	Updates the record of network and security data related to vCloud Networking and Security and NSX, particularly information about security groups and load balancing, for each machine following inventory data collection.
WMI data collection (Windows compute resources only)	Updates the record of the management data for each Windows machine. A WMI agent must be installed, typically on the Manager Service host, and enabled to collect data from Windows machines.

Start Endpoint Data Collection Manually

Endpoint data collection runs automatically every 4 hours, but IaaS administrators can manually start endpoint data collection at any time for endpoints that do not require proxy agents.

The **Data Collection** page provides information on the status and age of data collections and allows you to manually start a new endpoint data collection.

Prerequisites

Log in to the vRealize Automation console as an **IaaS administrator**.

Procedure

- 1 Select **Infrastructure > Endpoints > Endpoints**.
- 2 Point to the endpoint for which you want to run data collection and click **Data Collection**.
- 3 Click **Start**.
- 4 (Optional) Click **Refresh** to receive an updated message about the status of the data collection you initiated.
- 5 Click **Cancel** to return to the **Endpoints** page.

Configure Compute Resource Data Collection

You can enable or disable data collection, configure the frequency of data collection, or manually request data collection.

The **Data Collection** page provides information on the status and age of data collections. It also allows you to configure data collection for your compute resources.

Prerequisites

Log in to the vRealize Automation console as a **fabric administrator**.

Procedure

- 1 Select **Infrastructure > Compute Resources > Compute Resources**.
- 2 Point to the compute resource for which to configure data collection and click **Data Collection**.
- 3 Configure **Compute Resource** data collection specifications.
 - Select **On** to enable data collection.
 - Select **Off** to disable data collection.
- 4 Configure **Inventory** data collection.
 - Select **On** to enable data collection.
 - Select **Off** to disable data collection.
 - Enter a number in the **Frequency** text box to configure the time interval (in hours) between inventory data collections.
 - Click **Request Now** to manually start data collection.
- 5 Configure **State** data collection.
 - Select **On** to enable data collection.
 - Select **Off** to disable data collection.
 - Enter a number in the **Frequency** text box to configure the time interval (in minutes) between state data collections.
 - Click **Request Now** to manually start data collection.

6 Configure **Performance** data collection.

This is available only for vSphere integrations.

- Select **On** to enable data collection.
- Select **Off** to disable data collection.
- Enter a number in the **Frequency** text box to configure the time interval (in hours) between performance data collections.
- Click **Request Now** to manually start data collection.

7 Configure **vCNS Inventory** data collection.

This option is available for vSphere integrations configured to use vCloud Networking and Security.

- Select **On** to enable data collection.
- Select **Off** to disable data collection.
- Enter a number in the **Frequency** text box to configure the time interval (in hours) between vCNS Inventory data collections.
- Click **Request Now** to manually start data collection.

8 Configure **Snapshot Inventory** data collection.

This option is available for compute resources managed by vRealize Business Standard Edition.

- Select **On** to enable data collection.
- Select **Off** to disable data collection.
- Enter a number in the **Frequency** text box to configure the time interval (in hours) between snapshot data collections.
- Click **Request Now** to manually start data collection.

9 Configure **Cost** data collection.

This option is available for compute resources managed by vRealize Business Standard Edition.

- Select **On** to enable data collection.
- Select **Off** to disable data collection.
- Enter a number in the **Frequency** text box to configure the time interval (in hours) between cost data collections.
- Click **Request Now** to manually start data collection.

10 Click **OK**.

Update Cost Data for All Compute Resources

Fabric administrators can manually update cost information for all compute resources managed by vRealize Business Standard Edition.

Prerequisites

Log in to the vRealize Automation console as a **fabric administrator**.

Procedure

- 1 Select **Infrastructure > Compute Resources > Compute Resources**.
- 2 Click **Update Cost**.
- 3 Click **Request Now**.

When the cost update is complete, the status changes to successful.

Understanding vSwap Allocation Checking for vCenter Server Endpoints

You can use vSwap to determine swap space availability for the maximum size swap file on a target machine. The vSwap check occurs when you create or reconfigure a virtual machine from vRealize Automation. vSwap allocation checking is only available for vCenter Server endpoints.

vRealize Automation storage allocation checks if there is sufficient space available on the datastore to accommodate virtual machine disks during a create or reconfigure request. However, when the machine is powered on, if enough space is not available to create swap files on the vCenter Server endpoint, the machine fails to power on. When the power on operation fails, any customizations that depend on the machine also fail. The machine may also be disposed of. Depending on the size of the request, feedback that the machine is not powering on or not provisioning is not immediately obvious.

You can use the vSwap allocation check to help overcome these limitations by checking swap space availability for the maximum size swap file as part of the vRealize Automation create and reconfigure process for vCenter Server endpoints. To enable the vSwap allocation check, set the custom property `VirtualMachine.Storage.ReserveMemory` to True in the machine component or overall blueprint.

Consider the following behaviors for vSwap allocation checks:

- The swap file is located on the datastore that contains the virtual machine. Alternate vCenter Server configurations for locating swap files on a dedicated or different datastore are not supported.
- Swap size is considered when creating or reconfiguring a virtual machine . The maximum swap size is the size of the virtual machine's memory.
- Reserved values for vRealize Automation storage reservations in a host must not exceed the physical capacity of the compute resource.
- When creating a reservation, the sum of the reserved values must not exceed the available storage space.

- Resource pool or host level or virtual machine level memory reservations on vSphere are not collected from the vSphere endpoint and not considered during the calculations on vRealize Automation.
- vSwap does not validate the swap space that is available during power on operations for existing machines.
- You must re-run data collection to capture any changes made to the vSphere endpoint relative to vSwap.

Removing Datacenter Locations

To remove a datacenter location from a user menu, a system administrator must remove the location information from the locations file and a fabric administrator must remove location information from the compute resource.

For example, if you add London to the locations file, associate ten compute resources with that location, and then remove London from the file, the compute resources are still associated with the location London and London is still included in the location drop-down list on the Confirm Machine Request page. To remove the location from the drop-down list, a fabric administrator must edit the compute resource and reset the Location to blank for all compute resources that are associated with the location.

The following is a high-level overview of the sequence of steps required to remove a datacenter location:

- 1 A system administrator removes the datacenter location information from the locations file.
- 2 A fabric administrator removes all the compute resource associations to the location by editing the locations of each associated compute resource.

Bulk Import, Update, or Migrate Virtual Machines

You can use the Bulk Imports feature to import, update, or migrate virtual machines to vRealize Automation. Bulk Imports streamlines the management of multiple machines in multiple environments.

The Bulk Imports feature imports virtual machines intact with defining data such as reservation, storage path, blueprint, owner, and any custom properties. Bulk Imports supports the following administrative tasks:

- Import one or more unmanaged virtual machines so that they can be managed in a vRealize Automation environment.
- Make a global change to a virtual machine property, such as a storage path.
- Migrate a virtual machine from one environment to another.

You can execute the Bulk Imports feature commands using either the vRealize Automation console or the CloudUtil command-line interface. For more information about using the CloudUtil command-line interface, see the *Life Cycle Extensibility* documentation.

Prerequisites

- Log in to the vRealize Automation console as a **fabric administrator** and as a **business group manager**.
- If you are importing virtual machines that use static IP addresses, prepare a properly configured address pool.

Import a Virtual Machine to a vRealize Automation Environment

You can import an unmanaged virtual machine to a VMware vRealize™ Automation environment.

An unmanaged virtual machine exists in a hypervisor but is not managed in a vRealize Automation environment and cannot be viewed in the console. After you import an unmanaged virtual machine, the virtual machine is managed using the vRealize Automation management interface. Depending on your privileges, you can see the virtual machine on the **Managed Machines** tab or the **Items** tab.

Prerequisites

- Log in to the vRealize Automation console as a **fabric administrator** and as a **business group manager**.
- If you are importing virtual machines that use static IP addresses, prepare a properly configured address pool. For more information, see [Assigning a Static IP Address Range by Using Network Profiles](#).
- Create a blueprint for the virtual machine that you plan to import. This blueprint must be published, have a valid owner, and be entitled to that owner. The blueprint must have no more than one component.

Procedure

- 1 Generate a virtual machine CSV data file.
 - a Select **Infrastructure > Administration > Bulk Imports**.
 - b Click **Generate CSV File**.
 - c Select **Unmanaged** from the **Machines** drop-down menu.
 - d Select the **Business group** default value from the drop-down menu.
 - e Enter the **Owner** default value.
 - f Select the **Blueprint** default value from the drop-down menu.

The blueprint must be published and added to an entitlement for the import to be successful.

- g Select the **Component machine** default value from the drop-down menu.

If you select a value for **Business group** and **Blueprint**, you might see the following results in the CSV data file:

- Host Reservation (Name or ID) = INVALID_RESERVATION
- Host To Storage (Name or ID) = INVALID_HOST_RESERVATION_TO_STORAGE

These messages appear if you do not have a reservation in the selected business group for the host virtual machine that also hosts the unmanaged virtual machine. If you have a reservation in that business group for the unmanaged virtual machine host, the Host Reservation and Host to Storage values fill in properly.

- h Select one of the available resource types from the **Resource** drop-down menu.

Menu Item	Description
Endpoint	Information required to access a virtualization host.
Compute Resource	Information required to access a group of virtual machines performing a similar function.

- i Select the name of the virtual machine resource from the **Name** drop-down menu.
- j Click **OK**.

2 Edit your virtual machine CSV data file.

- a Open the CSV file, and edit the data categories to match existing categories in the target vRealize Automation environment.

To import virtual machines contained in a CSV data file, each virtual machine must be associated with the following items:

- Reservation
- Storage location
- Blueprint
- Virtual machine component
- Owner that exists in the target deployment

All the values for each virtual machine must be present in the target vRealize Automation environment for the import to succeed. You can change the values for reservation, storage location, blueprint, and owner, or add a static IP address value to individual virtual machines by editing the CSV file.

Heading	Comment
# Import--Yes or No	Change to No to prevent a particular virtual machine from being imported.
Virtual Machine Name	Do not change.
Virtual Machine ID	Do not change.
Host Reservation (Name or ID)	Enter the name or ID of a reservation in the target vRealize Automation environment.
Host To Storage (Name or ID)	Enter the name or ID of a storage location in the target vRealize Automation environment.
Deployment Name	Enter a new name for the deployment, for example, the virtual machine name, you are creating in the target vRealize Automation environment. Note Each virtual machine must be imported to its own deployment. You cannot import a single virtual machine to an existing deployment. You cannot import multiple virtual machines to a single deployment.
Blueprint ID	Enter the ID of the blueprint in the target vRealize Automation environment that you use to import the virtual machine. Note Make sure that you enter only the blueprint ID. Do not enter the blueprint name. You must select a blueprint that contains only a single virtual machine component. The blueprint must be published and added to an entitlement.
Component Machine ID	Enter the name of a virtual machine component that is contained in the blueprint you selected. You cannot import a virtual machine into a blueprint that has more than one component.
Owner Name	Enter a user in the target vRealize Automation environment who is entitled to the blueprint.

- b If you are importing a virtual machine with a static IP address, append a command in the following form to the CSV file.

`,VirtualMachine.Network#.Address, w.x.y.z, HOP`

Configure the command with the appropriate information for your virtual machine.

- Change the # to the number of the network interface being configured with this static IP address. For example, `VirtualMachineNetwork0.Address`.
- Change `w.x.y.z` to be the static IP address for the virtual machine. For example, `11.27.42.57`.
- The `HOP` string, `Hidden`, `Not encrypted`, `Not runtime`, sets the visibility of the property. This default property is removed from the virtual machine after a successful import.

For a successful import, the IP address must be available in a properly configured address pool. If the address cannot be found or is already in use, the import succeeds without the static IP address definition, and an error is logged.

- c Save the CSV file.
- 3 Use the vRealize Automation management interface to import your virtual machine to a vRealize Automation environment.
 - a Select **Infrastructure > Administration > Bulk Imports**.
 - b Click **New**.
 - c Enter a unique name for this task in the **Name** text box, for example, `unmanaged import 10`.
 - d Enter the CSV filename in the **CSV file** text box by browsing to the CSV filename.
 - e Select import options.

Option	Description
Start time	Schedule a future start date. The chosen start time is the local server time and not the local time of the user workstation.
Now	Begin the import process immediately.
Delay (seconds)	If you are importing many virtual machines, select the number of seconds to delay each virtual machine registration. Selecting this menu item slows the import process. Leave blank to select no delay.
Batch size	If you are importing many virtual machines, select the total number of virtual machines to register at a given time. Selecting this menu item slows the import process. Leave blank to select no limit.
Ignore managed machines	Leave unselected.
Skip user validation	Selecting this menu item sets the virtual machine owner to the value listed in the Owner column of the CSV data file without verifying that the user exists. Selecting this menu item can decrease the import time.
Test import	Test the import process without importing the virtual machines so you can test your CSV file for errors.

- f Click **OK**.

The progress of the operation appears on the Bulk Imports page.

Update a Virtual Machine in a vRealize Automation Environment

You can make a change to a virtual machine property, such as a storage path, to update one or more managed virtual machines in a vRealize Automation environment.

A managed virtual machine is a machine that is managed in a vRealize Automation environment and can be viewed in the console.

Prerequisites

- Log in to the vRealize Automation console as a **fabric administrator** and as a **business group manager**.

Procedure

- 1 Generate a virtual machine CSV data file.
 - a Select **Infrastructure > Administration > Bulk Imports**.
 - b Click **Generate CSV File**.
 - c Select **Managed** from the **Machines** drop-down menu.
 - d Select one of the available resource types from the **Resource** drop-down menu.

Option	Description
Endpoint	Information required to access a virtualization host.
Compute Resource	Information required to access a group of virtual machines performing a similar function.

- e Select the name of the virtual machine resource from the **Name** drop-down menu.
- f (Optional) Select **Include custom properties** if you want to migrate the virtual machine custom properties.
- g Click **OK**.

2 Edit your virtual machine CSV data file.

- a Open the CSV file with a text editor and edit the data categories that you want to change globally.

To update virtual machines contained in a CSV data file, each machine must be associated with the following items:

- Reservation
- Storage location
- Blueprint
- Machine component
- Owner that exists in the target deployment

All of the values for each machine must be present in the target vRealize Automation environment for the update to succeed. You can change the values for reservation, storage location, blueprint, and owner, or add a static IP address value to individual machines by editing the CSV file.

- b If you are changing a virtual machine static IP address, append a command in the following form to the CSV file.

```
,VirtualMachine.Network#.Address, w.x.y.z, HOP
```

Configure the command with the appropriate information for your virtual machine.

- Change the # to the number of the network interface being configured with this static IP address. For example, `VirtualMachineNetwork0.Address`.
- Change `w.x.y.z` to be the static IP address for the virtual machine. For example, `11.27.42.57`.
- The `HOP` string, Hidden, Not encrypted, Not runtime, sets the visibility of the property. This default property is removed from the virtual machine after a successful import.

For a successful update, the IP address must be available in a properly configured address pool. If the address cannot be found or is already in use, the update succeeds without the static IP address definition, and an error is logged.

- c Save the CSV file and close your text editor.

3 Use the vRealize Automation management interface to update one or more virtual machines in a vRealize Automation environment.

- a Select **Infrastructure > Administration > Bulk Imports**.
- b Click **New**.
- c Enter a unique name for this task in the **Name** text box, for example, managed global update 10.
- d Enter the CSV file name in the **CSV file** text box by browsing to the CSV file name.

- e Select import options.

Option	Description
Start time	Schedule a future start date. The specified start time is the local server time and not the local time of the user workstation.
Now	Begin the import process immediately.
Delay (seconds)	If you are updating a large number of virtual machines, select the number of seconds to delay each virtual machine update. Selecting this option slows the update process. Leave blank to specify no delay.
Batch size	If you are updating a large number of virtual machines, select the total number of machines to update at a given time. Selecting this option slows the update process. Leave blank to specify no limit.
Ignore managed machines	Leave unselected.
Skip user validation	Selecting this option sets the machine owner to the value listed in the Owner column of the CSV data file without verifying that the user exists. Selecting this option can decrease the update time.
Test import	Leave unselected.

- f Click **OK**.

The progress of the operation appears on the Bulk Imports page.

Migrate a Virtual Machine to a Different vRealize Automation Environment

You can migrate one or more managed virtual machines in a VMware vRealize™ Automation environment to a different vRealize Automation environment.

A managed virtual machine is a virtual machine that is managed in a vRealize Automation environment and can be viewed in the console.

Prerequisites

- Log in to the vRealize Automation console as a **fabric administrator** and as a **business group manager**.
- If you are importing virtual machines that use static IP addresses, prepare a properly configured address pool. For more information, see [Assigning a Static IP Address Range by Using Network Profiles](#).

Procedure

- 1 Generate a virtual machine CSV data file.
 - a Select **Infrastructure > Administration > Bulk Imports**.
 - b Click **Generate CSV File**.
 - c Select **Managed** from the **Machines** drop-down menu.

- d Select one of the available resource types from the **Resource** drop-down menu.

Option	Description
Endpoint	Information required to access a virtualization host.
Compute Resource	Information required to access a group of virtual machines performing a similar function.

- e Select the name of the virtual machine resource from the **Name** drop-down menu.
- f (Optional) Select **Include custom properties**.
You include custom properties when you import a virtual machine into a new deployment with the same properties.
- g Click **OK**.

2 Edit your virtual machine CSV data file.

Whether you must edit the CSV data file depends on the similarity of the source and target environments. If the configuration values in the source environment do not match the values in the target environment, you must edit the CSV data file so that the values match before you begin migration.

- a Open the CSV file, and edit the data categories to match existing categories in the target vRealize Automation environment.

To migrate virtual machines contained in a CSV data file, each virtual machine must be associated with a reservation, storage location, blueprint, machine component, and owner that exists in the target vRealize Automation environment. All the values for each virtual machine must be present in the target vRealize Automation environment for migration to succeed. You can change the values for reservation, storage location, blueprint, and owner, or add a static IP address value to individual virtual machines by editing the CSV file.

Heading	Comment	Example
# Import--Yes or No	Change to No to prevent a particular virtual machine from being imported.	Yes
Virtual Machine Name	Do not change.	MyMachine
Virtual Machine ID	Do not change.	a6e05812-0b06-4d4e-a84a-fed242340426a
Host Reservation (Name or ID)	Enter the name or ID of a reservation in the target vRealize Automation environment.	DevReservation
Host To Storage (Name or ID)	Enter the name or ID of a storage location in the target vRealize Automation environment.	ce-san-1:custom-nfs-2
Deployment Name	Enter a new name for the deployment you are creating in the target vRealize Automation environment. Each virtual machine must be migrated to its own deployment. You cannot import a single virtual machine to an existing deployment. You cannot import multiple virtual machines to a single environment.	ImportedDeployment0001
Converged Blueprint ID	Enter the ID of the blueprint in the target vRealize Automation environment that you use to import the virtual machine. Make sure that you enter only the blueprint ID. Do not enter the blueprint name. You must select a blueprint that contains only a single virtual machine component. The blueprint must be published and added to an entitlement.	ImportBlueprint
Component Blueprint ID	Enter the name of a virtual machine component that is contained in the blueprint you selected. You cannot import a virtual machine into a blueprint that has more than one component.	ImportedMachine
Owner Name	Enter a user in the target vRealize Automation environment.	user@tenant

Example of a complete, properly formatted CSV line: Yes, MyMachine, a6e05812-0b06-4d4e-a84a-fed242340426, DevReservation, ce-san-1:custom-nfs-2, Imported Deployment 0001, ImportBlueprint, ImportedMachine, user@tenant

- b If you are migrating a virtual machine with a static IP address, append a command in the following form to the CSV file.

```
,VirtualMachine.Network#.Address, w.x.y.z, HOP
```

Configure the command with the appropriate information for your virtual machine.

- Change the # to the number of the network interface being configured with this static IP address. For example, VirtualMachineNetwork0.Address.
- Change w.x.y.z to be the static IP address for the virtual machine. For example, 11.27.42.57.
- The *HOP* string, Hidden, Not encrypted, Not runtime, sets the visibility of the property. This default property is removed from the virtual machine after a successful import.

For a successful migration, the IP address must be available in a properly configured address pool. If the address cannot be found or is already in use, the migration succeeds without the static IP address definition, and an error is logged.

- c Save the CSV file.
- 3 Use the vRealize Automation management interface to migrate your virtual machine to a vRealize Automation environment.
 - a Select **Infrastructure > Administration > Bulk Imports**.
 - b Click **New**.
 - c Enter a unique name for this task in the **Name** text box, for example, managed migration 10.
 - d Enter the CSV filename in the **CSV file** text box by browsing to the CSV filename.

- e Select import options.

Option	Description
Start time	Schedule a future start date. The chosen start time is the local server time and not the local time of the user workstation.
Now	Begin the migration process immediately.
Delay (seconds)	If you are migrating many virtual machines, select the number of seconds to delay each virtual machine registration. Selecting this option slows the migration process. Leave blank to select no delay.
Batch size	If you are migrating many virtual machines, select the total number of virtual machines to register at a given time. Selecting this option slows the migration process. Leave blank to select no limit.
Ignore managed machines	Leave unselected.
Skip user validation	Selecting this option sets the virtual machine' owner to the value listed in the Owner column of the CSV data file without verifying that the user exists. Selecting this option can decrease the migration time.
Test import	Test the migration process without migrating the virtual machines so you can test your CSV file for errors.

- f Click **OK**.

The progress of the operation appears on the Bulk Imports page.

Managing Machines

You can manage provisioned machines and deployments by using available action options.

Managing Virtual Machines

Managing virtual machines requires different roles to perform some tasks. For example, only a fabric administrator can change the reservation of a virtual machine, but a machine owner can create a snapshot of a virtual machine.

Reconfigure a Machine

vSphere, vCloud Air, and vCloud Director platforms support reconfiguration of existing machines to modify specifications for CPU, memory, storage, or networks.

Reconfiguration requests are subject to approval based on entitlements, policies, and the actions enabled for the machine component in the blueprint.

If you are entitled to the Cancel Reconfigure (Machine) and Execute Reconfigure (Machine) actions, you can cancel a reconfiguration or retry a failed reconfiguration.

Prerequisites

- Log in to the vRealize Automation console as a **machine owner, support user, tenant administrator, or business group manager**.
- The machine you want to reconfigure has the status On or Off with no active reconfigure status.

Procedure

1 [Start the Operation](#)

You start the reconfigure operation by selecting a provisioned machine with an On or Off status.

2 [Reconfigure CPUs and Memory](#)

You can optionally change the number of CPUs or the amount of memory used by the provisioned machine, within the limits set by the provisioning blueprint.

3 [Reconfigure Storage](#)

You can add, delete, or change the size of a storage volume on a provisioned virtual machine.

4 [Add Custom Properties](#)

Optionally, you can add custom properties to a volume.

5 [Reconfigure Networks](#)

You can add, remove, or edit a network adapter when you reconfigure a virtual machine.

6 [Schedule the Start](#)

You can start the reconfiguration immediately or schedule it to start at a particular day and time. You can also specify the power option for the machine before reconfiguring it.

Start the Operation

You start the reconfigure operation by selecting a provisioned machine with an On or Off status.

Prerequisites

- Log in to the vRealize Automation console as a **machine owner**, **support user**, **tenant administrator**, or **business group manager**.
- The machine you want to reconfigure must have the status On or Off with no active reconfigure status.

Procedure

- 1 Select **Items > Machines**.
- 2 Select the row of the machine to reconfigure.
- 3 Select **Reconfigure** from the **Actions** drop-down menu.

What to do next

[Reconfigure CPUs and Memory](#).

Reconfigure CPUs and Memory

You can optionally change the number of CPUs or the amount of memory used by the provisioned machine, within the limits set by the provisioning blueprint.

Prerequisites

[Start the Operation](#).

Procedure

- 1 (Optional) Type the number of CPUs in the **# CPUs** text box.
The allowable range appears next to the text box.
- 2 (Optional) Type the amount of memory in the **Memory (MB)** text box.
The allowable range appears next to the text box.

What to do next

[Reconfigure Storage.](#)

Reconfigure Storage


You can add, delete, or change the size of a storage volume on a provisioned virtual machine.

You cannot reconfigure storage for the IDE disk type.


Prerequisites

[Reconfigure CPUs and Memory.](#)

Procedure

- 1 Click the **Storage** tab.
The allowable range for storage appears below the Storage volumes table.
- 2 (Optional) Add a volume.
 - a Click **New Volume**.
 - b Type the capacity in the **Capacity (GB)** text box.
 - c (Optional) Select a storage reservation policy from the **Storage reservation policy** drop-down menu.
 - d Click the **Save** icon ().


- 3 (Optional) Delete a volume.

- a Locate the volume.
- b Click the **Delete** icon (.

An unselectable icon indicates an undeletable volume such as one from a linked clone.

- 4 (Optional) Increase the size of a volume.

You cannot reduce the size of existing volumes. Volume size is limited by the total amount of storage specified in the blueprint, less the amount allocated to other volumes.

- a Locate the volume.
- b Click the **Edit** icon (.

- c Type the new size in the **Capacity (GB)** text box.
- d Click the **Save** icon (✔).

What to do next

[Add Custom Properties.](#)

Add Custom Properties

Optionally, you can add custom properties to a volume.

You cannot use custom properties to enter values for volume disk number, capacity, label, or storage reservation policy. You must enter these values in their required locations by adding or editing a volume in the Storage volumes table.

Prerequisites

[Reconfigure Storage.](#)

Procedure

- 1 In the **Custom Properties** column of the **Storage volumes** table, click **Edit** for the volume receiving the custom property.
- 2 Click **New Property**.
- 3 Enter the name of the custom property in the **Name** text box.
- 4 Enter the value for the custom property in the **Value** text box.
- 5 Select the **Encrypted** check box to encrypt the value.
- 6 Select the **Prompt user** check box to prompt users for the value when they request the machine.

What to do next

[Reconfigure Networks.](#)

Reconfigure Networks

You can add, remove, or edit a network adapter when you reconfigure a virtual machine.

Prerequisites

[Add Custom Properties.](#)

Procedure


- 1 Click the **Network** tab.
- 2 (Optional) Add a network adapter.
 - a Click **New Network Adapter**.
 - b Select a network from the **Network Path** drop-down menu.

All networks selected on the machine's reservation are available.

- c Type a static IP address for the network in the **Address** text box.
The IP address must be unallocated in the network profile assigned in the reservation.



- d Click the **Save** icon ()

3 (Optional) Remove a network adapter.

- a Locate the network adapter.
- b Click the **Delete** icon ()

You cannot remove network adapter 0.

4 (Optional) Edit a network adapter.

- a Locate the network adapter.
- b Click the **Edit** icon ()
- c Select a network from the **Network Path** drop-down menu.
- d Click the **Save** icon ()

What to do next

[Schedule the Start.](#)

Schedule the Start

You can start the reconfiguration immediately or schedule it to start at a particular day and time. You can also specify the power option for the machine before reconfiguring it.

Prerequisites

[Reconfigure Networks.](#)

Procedure

- 1 Click the **Execution** tab.
- 2 (Optional) Select an option from the **Execute request** drop-down menu.

Option	Description
Immediate	Start reconfiguration as soon as possible after approval.
Scheduled	Start reconfiguration at the specified date and time. Type or select the date and time in the text boxes that appear.

The scheduled time is the local time where the vRealize Automation Web server is located. If **Execute request** is not available, reconfiguration starts immediately.

3 (Optional) Select a power action from the **Power action** drop-down menu.

Option	Description
Reboot if required	(Default) If required, reboot the machine before reconfiguring it.
Reboot	Reboot the machine before reconfiguring it, regardless of whether reboot is required.
Do not reboot	Do not reboot the machine before reconfiguring it, even if reboot is required.

The following conditions require that the machine be rebooted before reconfiguration:

- CPU change where hot add is not supported or is disabled
- Memory change where hot add is not supported or is disabled
- Storage change where hot add is not supported or is disabled
- Hardware reconfiguration

If the machine is in the shutdown state, it is not rebooted.

Note You can disable the vSphere hot add option by using the `VirtualMachine.Reconfigure.DisableHotCpu` custom property.

4 Click **OK**.

What to do next

You can monitor the progress of the reconfiguration by observing the workflow states displayed in the user interface. See [Workflow States of Reconfigure Operations](#).

Workflow States of Reconfigure Operations

When reconfiguration starts and as it progresses through the workflow, you can monitor the progress from the Edit page.

Table 1-14. Workflow States of Reconfigure Operations

State	Description
Reconfigure pending	The State Operation was created.
Scheduled	A scheduled workflow has been created for the Distributed Execution Manager (DEM).
Reconfiguring	The interface-specific workflow is being executed.
Reconfigure failed, waiting to retry	The reconfigure failed, waiting for the owner to request a retry. If the machine owner is entitled to the actions execute reconfigure or cancel reconfigure, the owner can retry or cancel a reconfiguration.
ReconfigureFailed	The reconfigure failed, waiting for the RVG workflow to perform the next action.
ReconfigureSuccessful	The reconfigure was successful, waiting for the RVG workflow to perform the next action.

Table 1-14. Workflow States of Reconfigure Operations (Continued)

State	Description
Canceled	The user has canceled the reconfiguration. Machine owners who are entitled to the cancel reconfigure action can cancel a reconfiguration.
Complete	The completion workflow sets this state after completing the cleanup, so that the RVG workflow can proceed to clean up the state operations and approvals. A status of complete indicates that the request from vRealize Automation is finished, but it does not indicate that the machine reconfiguration completed successfully.

Configure a Metrics Provider

You can configure vRealize Automation to use vRealize Operations Manager health and resource metrics for vSphere virtual machines.

For more information about vRealize Operations Manager health badges and metrics, see the vRealize Operations Manager documentation.

Prerequisites

- Log in to the vRealize Automation console as a **tenant administrator**, **business group manager**, or **machine owner**.
- Create a vRealize Operations Manager user account with view and resource metrics query privileges for all vSphere servers that you integrate with vRealize Automation.
- Create vRealize Operations Manager adapter instances for all vSphere servers you add as endpoints in vRealize Automation. For information about creating adapter instances, see the vRealize Operations Manager documentation.

Procedure

- 1 Select **Administration > Reclamation > Metrics Provider**.
- 2 Select a metrics provider.

Option	Description
(Default) vRealize Automation metrics provider	If you do not have a vRealize Operations Manager instance, vRealize Automation provides basic machine metrics.
vRealize Operations Manager endpoint	Provide connection information for the vRealize Operations Manager instance you want to use as your metrics provider for vSphere virtual machines.

- 3 Click **Test Connection**.
- 4 Click **Save**.

Tenant administrators, machine owners, and business group managers of the group in which the machine resides can view health badges and health alerts on the item details pages for vSphere virtual machines. They can also view vRealize Operations Manager metrics and health badges when they filter by the platform type vSphere on the reclamations page.

What to do next

[Send Reclamation Requests.](#)

Send Reclamation Requests

You can send reclamation requests to virtual machine owners. A reclamation request specifies a new lease length in days, the amount of time given for a machine owner's response, and which machines to target for reclamation.

Prerequisites

- Log in to the vRealize Automation console as a **tenant administrator**.
- (Optional) To see health badges or view metrics provided by vRealize Operations Manager, see [Configure a Metrics Provider](#).

Procedure


- 1 Select **Administration > Reclamation > Tenant Machines**.

2 Find virtual machines that match your search criteria.

You must select platform type vSphere to view metrics provided by vRealize Operations Manager.

- a Click the **Advanced Search** down arrow to open the search box.
- b Enter or select one or more search values.

Option	Action
Virtual Machine name contains	Enter one or more characters in the text box to find virtual machine names that match.
Owner name contains	Enter a name in the text box to find owner names that match.
Business group names contains	Enter a name in the text box to find business group names that match.
Platform Type (required for vRealize Operations Manager)	Select a platform type from the drop-down menu. Select vSphere to view metrics provided by vRealize Operations Manager.
Power State	Select a power state value from the drop-down menu to find virtual machines with a matching power state.
Expiration date between	Click the calendar icons and select start and end dates to find expiration dates inside the range.
CPU usage	Select a value from the drop-down menu to find virtual machines with High CPU use, above 80%, Low CPU use below 5%, or None, no value. If you are querying vRealize Operations Manager metrics, you cannot use this filter to query, and you cannot sort results by CPU usage.
Mem usage	Select a value from the drop-down menu to find virtual machines with High Memory use, above 80%, Low Memory use, below 10%, or None, no value. If you are querying vRealize Operations Manager metrics, you cannot use this filter to query, and you cannot sort results by memory usage.
Disk usage	Select a value from the drop-down menu to find virtual machines with Low Hard Disk use, less than 2 KBs per second or None, no value. If you are querying vRealize Operations Manager metrics, you cannot use this filter to query, and you cannot sort results by disk usage.
Network usage	Select a value from the drop-down menu to find virtual machines with Low Network use, less than 1 KB per second, or None, no value. If you are querying vRealize Operations Manager metrics, you cannot use this filter to query, and you cannot sort results by network usage.
Complex metric	Select a value from the drop-down menu to find virtual machines based on complex metrics. For example, select idle to find machines that have CPU, network, memory, and disk usage values all under 20%. You cannot use this filter if you are querying vRealize Operations Manager metrics.

- c Click the search icon ().

3 From the current results page, select one or more virtual machines to reclaim.

For manageability, only machines selected on the current results page are reclaimed.

4 Click **Reclaim Virtual Machine**.

Virtual machines selected on the current results page are included in the request.

Note The Reclamation page can list machines that are not available for reclamation, such as machines for which the lease has expired. If you specify a machine that is not available for reclamation, you receive the following error:

```
Selection Error: Virtual machine name is not in valid state for reclamation.
```

5 Enter the duration of the new lease in the **New lease length (days)** text box.

The minimum is 1 day, the maximum is 365 days, and the default is 7 days.

6 Enter how many days the machine owner has to respond to the reclamation request in the **Wait before forcing lease (days)** text box.

At the end of that time, the machine gets a new lease with the new lease length. The minimum waiting period is 1 day, the maximum is 365 days, and the default is 3 days.

7 Enter a reason for the request in the **Reason for request** text box.

8 Click **Submit**.

9 Click **OK**.

When you send a reclamation request, it appears in the Inbox of the owner of the virtual machine. If the owner does not respond to the request in the required number of days, the virtual machine gets a new lease of the specified length, unless its current lease is shorter. If the owner clicks **Item in Use** on the reclamation request, the virtual machine's lease remains unchanged. If the owner clicks **Release for Reclamation**, the virtual machine lease expires immediately.

What to do next

[Track Reclamation Requests.](#)

Track Reclamation Requests

A tenant administrator can track the current state of reclamation requests and other details.

Prerequisites


Log in to the vRealize Automation console as a **tenant administrator**.

Procedure

1 Select **Administration > Reclamation > Reclamation Requests**.

- 2 Find the virtual machines that match your search criteria.
 - a Click the **Advanced Search** down arrow to open the search box.
 - b Type or select one or more search values.

Option	Action
Virtual Machine name contains:	Type one or more characters in the text box to find virtual machine names that match.
Owner name contains:	Type one or more characters in the text box to find owner names that match.
Request Reason contains:	Type one or more characters in the text box to find a request reason that matches.
Request State:	Select a request state value from the drop-down menu to find virtual machines with a matching request state.

- c Click the **Search** icon () or press Enter to start the search.
 - d Click the **Advanced Search** up arrow to close the search box.
- 3 (Optional) Click **Refresh Data** to update the display of reclamation requests.

Change the Reservation of a Managed Machine

You can change the reservation of a managed machine. This ability is useful when a machine moves to a new storage path that is not available in its current reservation.

You can change the machine’s current compute resource. You can also move it to any reservation on that compute resource, including one belonging to a different business group. You must be a business group manager of the original and the target business groups to use this function.

Prerequisites

Log in to the vRealize Automation console as a **fabric administrator**.

Procedure

- 1 Select **Infrastructure > Managed Machines**.
- 2 Locate the machine with the reservation to change.
- 3 Click **Change Reservation** in the drop-down menu.
- 4 Select values from the drop-down menus as required.
- 5 Click **OK**.

Create a Snapshot of Your Machine

Depending on how your administrators have configured your environment, you might be able to create a snapshot of your virtual machine. A snapshot is an image of a virtual machine at a specific time. It is a space-efficient copy of the original VM image. Snapshots are an easy way to recover a system from damage, data loss, or security threats. After you create a snapshot of your virtual machine, you can apply it and reset your system back to the point where the snapshot was taken.

When you create a memory snapshot, the snapshot captures the state of the virtual machine power settings and, optionally, the virtual machine's memory. When you capture the virtual machine's memory state, the snapshot operation takes longer to complete. You might also see a momentary lapse in response over the network.

Prerequisites

- An existing virtual machine that is powered on, off, or suspended.
- If your virtual machine is configured for one or more independent disks, power off the machine before creating a snapshot. You cannot create a snapshot when it is powered on. For disk configuration information, see *Custom Properties V Table*.
- Your tenant administrator or business group manager entitled you to the snapshot action.

Procedure

- 1 Select **Items > Machines**.
- 2 Locate the machine to snapshot.
- 3 In the Actions column, click the down arrow and click **View Details**.
- 4 Click **Create Snapshot** in the Actions menu.
- 5 Enter a name and, optionally, a description.
- 6 If you want to capture the memory and power settings of the machine, select **Include memory**.
- 7 Click **Submit**.

Configuring Remote Consoles for vSphere with Untrusted SSL Certificates

If your vRealize Automation deployment uses untrusted certificates, before you can use remote consoles with VMRC, you must configure your client browser to trust the certificate. The steps to do this vary by browser.

If vRealize Automation is configured with a trusted SSL certificate for your environment, then VMRC does not require additional configuration on client browsers. When a vRealize Automation appliance certificate is replaced and is a trusted certificate, there is no need to update certificate information for the Web browser client.

If you want to replace the certificate, see the topic on replacing a vRealize Automation appliance certificate in the *System Administration* guide for vRealize Automation.

Remote connections using VMRC for machines provisioned on vSphere are secured by vRealize Appliance certificates through a proxy console. VMRC requires WebSockets support in the browser and browsers must trust the vRealize Appliance certificate. The certificate can be obtained by going to the root-level virtual appliance at an address of the form `https://vra-va.eng.mycompany.com/`.

For information about support requirements for browsers and vSphere, see the VMware vRealize Support Matrix.

Configure Firefox to Trust a Certificate for vRealize Automation

Untrusted vRealize Automation appliance certificates must be manually imported to client browsers to support VMware Remote Console on clients provisioned on vSphere.

For information about supported versions of Firefox, see the *VMware vRealize Support Matrix* on the VMware Web site.

Note If vRealize Automation is configured with a trusted SSL certificate for your environment, then VMware Remote Console does not require additional configuration on client browsers.

Procedure

- 1 In a Firefox browser, log in to the vRealize Automation appliance.
A message appears saying that the certificate is not trusted.
- 2 Choose the option to display the current connection information. Click **View Certificate** to display the current SSL certificate and click **Detail** in the Certificate Viewer.
- 3 Click **More Information** and click the **Security** tab on the Page Info page.
- 4 Select a certificate from the Certificate Hierarchy pane.

Option	Action
Certificate Authority issued certificates	Select the top-level vRealize Automation certificate.
Self-signed certificates	Select the vRealize Automation certificate.

- 5 Click **Export**.
- 6 Configure the certificate information in the Save Certificate To File dialog box.
 - a Enter a certificate name in the **Save As** text box. The certificate name must end in `.crt`, `.cert`, or `.cer`.
 - b Select a location in which to save the file.
 - c Select **X.509 Certificate (PEM)** as the format.
- 7 Click **Save**.

- Click the **Authorities** tab in the Certificate Management dialog box.

Option	Action
Windows	Select Preference > Advanced > Certificates from the Firefox menu.
iOS	Select Preference > Advanced > Certificates from the Firefox menu and click View Certificates .

- Click the **Authorities** tab and click **Import**.
- Select the certificate file you saved earlier and click **Open** in the dialog box.
- Edit the trust settings.

Option	Action
Self-signed certificates	Select This certificate can identify websites .
Certificates issued by a Certificate Authority	Select Trust this CA to identify websites .

- Click **OK** and restart the browser.

You can connect to the remote console without certificate errors.

Configure Internet Explorer to Trust a Certificate for vRealize Automation Appliance

Untrusted vRealize Automation appliance certificates must be manually imported to client browsers to support VMware Remote Console on clients provisioned on vSphere.

Note If vRealize Automation is configured with a trusted SSL certificate for your environment, then VMware Remote Console does not require additional configuration on client browsers.

The steps in this procedure apply for self-signed certificates and certificates issued by a Certificate Authority.

For information about supported versions of Internet Explorer, see the *VMware vRealize Support Matrix* on the VMware Web site.

Procedure

- In an Internet Explorer browser, log in to the vRealize Automation appliance.
- Click **View Certificate** on the certificate error message that appears in the browser address bar.
- Click the **General** tab of the Certificate Information window..
- Verify that the information about the certificate is correct and click **Install Certificate**.
- Select **Place all certificates in the following store** in the Certificate Store dialog box.
- Click **Browse** to locate the certificate store.
- Select **Trusted Root Certification Authority** and click **OK**.
- Click **Next** on the Certificate Store dialog box.

- 9 Click **Yes** in the Security Warning dialog box to install the certificate.
- 10 Restart the browser.

You can connect to the remote console without certificate errors.

Configure Chrome to Trust a Certificate for vRealize Automation Appliance

Untrusted vRealize Automation appliance certificates must be manually imported to client browsers to support VMware Remote Console on clients provisioned on vSphere.

For information about supported versions of Chrome, see the *VMware vRealize Support Matrix* on the VMware Web site.

Note If vRealize Automation is configured with a trusted SSL certificate for your environment, then VMware Remote Console does not require additional configuration on client browsers.

On Windows, Chrome and Internet Explorer use the same certificate store. This means that certificates that are trusted by Internet Explorer are also trusted by Chrome. To establish trusted certificates for Chrome, import them through Internet Explorer. For information about this procedure, see [Configure Internet Explorer to Trust a Certificate for vRealize Automation Appliance](#).

When you complete the procedure, restart Chrome.

To permanently trust a certificate on the Macintosh operating system, download the certificate file and install the certificate as trusted in your certificate management tool.

Procedure

- 1 In a Chrome browser, log in to the vRealize Automation appliance.
- 2 Click the icon in the address bar.
- 3 Click the certificate information link.
- 4 Save the certificate by dragging the certificate icon to the desktop.
- 5 Start the Keychain Access application.
- 6 Select **File > Import Items**.
- 7 On the Keychain Access screen, select the certificate file you saved earlier.
Set the value of **Destination Key** to **System**.
- 8 Click **Open** to import the certificate.
- 9 Restart the browser.

Connect Remotely to a Machine

You can connect remotely to a machine from the vRealize Automation console.

Prerequisites

- Log in to the vRealize Automation console as a **machine owner**, **tenant administrator**, or **business group manager**.
- Verify that VMware Tools is installed.

VMware Tools must be installed on your vRealize Automation client to support fully functioning access when connecting with VMware Remote Console. If VMware Tools is not installed, problems occur, such as the mouse pointer and mouse keys not working after connecting to the target machine. For information about supported VMware Tools versions, see *vRealize Automation Support Matrix*.

- Verify that the provisioned machine is powered on.

Procedure

- 1 Select **Items > Deployment**.
- 2 Click **Actions** in the machine name row or select the machine and click **Actions** on its machine page.
- 3 Select the remote connection method.
 - Select **Connect Using RDP** to connect by using RDP.
 - Select **Connect to remote console** to connect by using VMware Remote Console.
 Respond to any prompts.
- 4 Click **Connect** and log in to the machine as directed.
- 5 When finished, log out and close the browser window.

Running Actions for Provisioned Resources

The actions that are available for a provisioned resource depend on the type of resource, how the action was configured and made available for provisioned items, and the operational state of the item.

The configured actions that are available for a provisioned machine or deployment appear in the **Actions** menu for the selected resource on the **Items** tab.

If the item was provisioned by IaaS using an IaaS machine blueprint, the list of available actions is determined by what was selected on the **Actions** tab for the machine type component when the blueprint was created, and then by what is applicable based on machine type or state.

If the item was provisioned using an XaaS blueprint, the resource actions must be created, published, and entitled in the same service that is used to provision the item. The list of available actions is determined by the item type and the current state of the item.

The available actions for an item that was provisioned as an IaaS machine might also include XaaS resource actions if the actions are mapped to the item.

Action Menu Options for Provisioned Resources

Actions are changes that you can make to provisioned resources. The vRealize Automation actions are used to manage the life cycle of the resources.

The options on the **Action** menu for a provisioned item include the actions specified on the blueprint and might include custom menu operations created by your service architects. The available actions depend on how your business group manager or tenant administrator configured the entitlement that contains the resource on which the actions run.

Table 1-15. Action Menu Options

Action	Resource Type	Description
Associate Floating IP	Machine (OpenStack)	Associate a floating IP address with an Openstack machine.
Cancel Reconfigure	Machine	Cancel a running reconfiguration action.
Change lease	Deployment and Machine	Change the number of days remaining in the lease for either a specific machine or for all resources included in a deployment. If you do not provide a value, the lease does not expire.
Change Owner	Deployment	Change the owner of the deployment and all the included resources. Only Business group managers and support users can change the ownership of a deployment.
Connect using VMRC	Machine	Connect to the virtual machine using a VMRC 8.x application. To use this action, the VMRC application must be installed on the local system of the service catalog user who is running the action. For installation and user instructions, see VMware Remote Console Documentation . To download, see Download VMware Remote Console . The VMRC 8.x replaces the previous VMware Remote Console.
Connect to remote console	Machine	Connect to the selected machine using VMware Remote Console. The virtual machine console appears in the browser. The VMRC 8.x replaces the VMware Remote Console.

Table 1-15. Action Menu Options (Continued)

Action	Resource Type	Description
Connect using Console Ticket	Machine (OpenStack and KVM)	Connect to the OpenStack or KVM virtual machine using a console ticket for a VMware Remote Console connection.
Connect using ICA	Machine (Citrix)	Connect to the Citrix machine using the Independent Computing Architecture.
Connect using RDP	Machine	Connect to the machine by using Microsoft Remote Desktop Protocol.
Connect using SSH	Machine	<p>Connect to the selected machine by using SSH.</p> <p>The Connect Using SSH option requires that your browser has a plug-in that supports SSH, for example the FireSSH SSH terminal client for Mozilla Firefox and Google Chrome. When the plug-in is present, selecting Connect Using SSH displays an SSH console and prompts for your administrator credentials.</p> <p>To use this action, the <code>Machine.SSH</code> custom property must be included and set to true in the blueprint's machine component in either a property group or individual custom property.</p>
Connect using Virtual Desktop	Machine	Connect to the selected machine using Microsoft virtual desktop.
Create Snapshot	Virtual Machine	Create a snapshot of the virtual machine. If you are allowed only two snapshots and you already have them, this option is not available until you delete a snapshot.
Delete Snapshot	Virtual Machine	Delete a snapshot of the virtual machine.

Table 1-15. Action Menu Options (Continued)

Action	Resource Type	Description
Destroy	Cloud Machine, Deployment, Virtual Machine, and VMware NSX Edge	<p>Immediately destroy a provisioned resource. You must run this action to destroy XaaS resources, even if they are part of a deployment you are destroying. Other resources are destroyed when their lease or their archival period ends.</p> <p>With the exception of XaaS, destroying components of a deployment is not a recommended best practice. It is recommended to use the scale in action to reduce the number of machines in your deployment, or destroy the entire deployment.</p> <p>The Destroy action is not available for the following deployment situations:</p> <ul style="list-style-type: none"> ■ physical machine deployments ■ deployments with an NSX existing network or NSX existing security resource ■ deployments with an NSX on-demand load balancer resource <p>Because an NSX load balancer is property of an on-demand NSX edge, when an NSX edge is destroyed, the load balancer resource is also destroyed and resources are released. When a machine tier that is load balanced is destroyed, it is removed from the load balancer pool on the respective NSX edge.</p> <hr/> <p>Note Note: Storage and memory that are assigned to a provisioned machine by a reservation are released when the machine to which they are assigned is deleted in vRealize Automation by the Destroy action. The storage and memory are not released if the machine is deleted on the vCenter Server.</p> <hr/> <p>When destroying a deployment that contains an Amazon machine component, all EBS volumes that were added to the machine during its life cycle are detached, rather</p>

Table 1-15. Action Menu Options (Continued)

Action	Resource Type	Description
		than destroyed. vRealize Automation does not provide an option for destroying the EBS volumes.
Destroy Existing Network	Existing Network	Destroy the network.
Destroy VMWare NSX Load Balancer	VMWare NSX Load Balancer	Destroy the NSX load balancer.
Destroy VMWare NSX Network	VMWare NSX Network	Destroy the NSX network.
Destroy VMWare NSX Security Group	VMWare NSX Security Group	Destroy the NSX security group.
Destroy VMWare NSX Security Tag	VMWare NSX Security Tag	Destroy the NSX security tag.
Disassociate Floating IP	Machine (Openstack)	Remove the floating IP from the Openstack machine.
Execute Reconfigure	Machine	Override a scheduled reconfiguration, or rerun or reschedule a failed reconfiguration.
Expire	Deployment and Machine	Terminate the deployment or machine lease for all resources included in the deployment.
Export Certificate	Machine	Export the certificate from a Cloud machine.
Get Expiration Reminder	Machine	Downloads a calendar event file for the current lease expiration date.
Install Tools	Machine	Install VMware Tools on a vSphere virtual machine.
Power Cycle	Machine	Power off the machine, then power it back on.
Power Off	Machine	Power off the machine without shutting down the guest operating system.
Power On	Machine	Power on the machine. If the machine was suspended, normal operation resumes from the point at which the machine was suspended.
Reboot	Machine	Reboot the guest operating system on a vSphere virtual machine. VMware Tools must be installed on the machine to use this action.

Table 1-15. Action Menu Options (Continued)

Action	Resource Type	Description
Reconfigure	Machine	<p>A business group manager, support user, or machine owner can reconfigure the following machine settings on a vSphere machine:</p> <ul style="list-style-type: none"> ■ Change description ■ Change CPU, memory, network, and disk settings ■ Add, edit, and delete properties ■ Reconfigure shutdown ■ Change machine owner (Business group managers and support users only) <p>You cannot change a storage reservation policy if doing so would change the storage profile on a disk.</p>
Register VDI	Virtual Machine (XenServer)	Register the virtual disk image on XenServer items.
Reprovision	Machine	<p>Destroys the machine, then initiates the provisioning workflow to create a new machine with the same name.</p> <p>When you request that a machine be reprovisioned, a known issue might cause vRealize Automation to display the reprovisioning status as Complete in the catalog, when the actual state is In Progress. After you submit a request to reprovision a machine, you can use any of the following sequences to check the status of the reprovisioned machine:</p> <ul style="list-style-type: none"> ■ Infrastructure > Managed Machines ■ Items > Item Details ■ Administration > Events > Event Logs
Revert Snapshot	Virtual Machine	Revert to a previous snapshot of the machine. You must have an existing snapshot to use this action.

Table 1-15. Action Menu Options (Continued)

Action	Resource Type	Description
Scale In	Deployment	<p>Destroy unneeded instances of machines in your deployment to adjust to reduced capacity requirements. Machine components and any software components installed on them are destroyed. Dependent software components and networking and security components are updated for the new deployment configuration. XaaS components are not scalable and are not updated during scale operations.</p> <p>You can try to repair partially successful scale operations by attempting to scale the deployment again. However, you cannot scale a deployment to its current size, and fixing a partially successful scale this way does not deallocate the dangling resources. You can view the request execution details screen and find out which tasks failed on which nodes to help you decide whether to fix the partially successful scale with another scale operation. Failed and partially successful scale operations do not impact the functionality of your original deployment, and you can continue to use your catalog items while you troubleshoot any failures.</p>

Table 1-15. Action Menu Options (Continued)

Action	Resource Type	Description
Scale Out	Deployment	<p>Provision additional instances of machines in your deployment to adjust to growing capacity requirements. Machine components and any software components installed on them are provisioned. Dependent software components and networking and security components are updated for the new deployment configuration. XaaS components are not scalable and are not updated during scale operations.</p> <p>You can try to repair partially successful scale operations by attempting to scale the deployment again. However, you cannot scale a deployment to its current size, and fixing a partially successful scale this way does not deallocate the dangling resources. You can view the request execution details screen and find out which tasks failed on which nodes to help you decide whether to fix the partially successful scale with another scale operation. Failed and partially successful scale operations do not impact the functionality of your original deployment, and you can continue to use your catalog items while you troubleshoot any failures.</p>
Shutdown	Machine	<p>Shut down the guest operating system and power off the machine. VMware Tools must be installed on the machine to use this action.</p>
Suspend	Machine	<p>Pause the machine so that it cannot be used and does not consume any system resources other than the storage it is currently using.</p>
Unregister	Machine	<p>Remove the machine from the inventory without destroying it. Unregistered machines are not usable.</p>
Unregister VDI	Virtual Machine (XenServer)	<p>Unregister the virtual disk image on XenServer items.</p>

Troubleshooting Missing Actions in the Resource Actions Menu

As a machine or resource owner, you do not see all entitled actions for a provisioned item.

Problem

In an environment where you know that an action was entitled for your user or business group, you expect to see all actions when you select an item in your **Items** list.

Cause

The availability of actions depends on the type of provisioned resource, operational state of the resource, and how it was configured and made available. The following list provides some reasons why you do not see all configured actions.

- The action is not applicable based on the current state of the provisioned resource. For example, Power Off is available only when the machine is powered on.
- The action is not applicable to the selected item type. If the item does not support the action, it does not appear in the list. For example, the Create Snapshot action is not available for a physical machine, and the Connect by Using RDP action is not available if the selected item is a Linux machine.
- The action is applicable for the provisioned resource type, but the action is disabled in the Infrastructure blueprint. If the action is disabled, it never appears as an available action for any of the items that were provisioned using the blueprint.
- The action is not included in the entitlement used to provision the item on which you need to run the action. Only entitled actions, either as part of an IaaS blueprint or as an XaaS resource action, can appear in the Actions menu.
- The action is created as an XaaS resource action but was not included in the entitlement used to provision the item on which you need to run the action. Only entitled actions appear in the Actions menu.
- The action might be limited based on the configured target criteria for XaaS resource actions or resource mappings to provisioned IaaS machines.

Solution

- Verify that the action is applicable to the provisioned item or the state of the provisioned item.
- Verify that the action is configured and included in the entitlement used to provision the item.