

Managing vRealize Automation

21 July 2021

vRealize Automation 7.6

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

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

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

Copyright © 2015-2021 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

| | | |
|----------|---|----------|
| 1 | Maintaining and Customizing vRealize Automation Components and Options | 5 |
| | Broadcast a Message to All Users | 5 |
| | Create a Message Board URL Allowlist | 7 |
| | Starting Up and Shutting Down vRealize Automation | 8 |
| | Start Up vRealize Automation | 8 |
| | Restart vRealize Automation | 9 |
| | Shut Down vRealize Automation | 10 |
| | Updating vRealize Automation Certificates | 11 |
| | Extracting Certificates and Private Keys | 13 |
| | Replace Certificates in the vRealize Automation Appliance | 14 |
| | Replace the Infrastructure as a Service Certificate | 16 |
| | Replace the IaaS Manager Service Certificate | 18 |
| | Update Embedded vRealize Orchestrator to Trust vRealize Automation Certificates | 20 |
| | Update External vRealize Orchestrator to Trust vRealize Automation Certificates | 22 |
| | Updating the vRealize Automation Appliance Management Site Certificate | 23 |
| | Replace a Management Agent Certificate | 28 |
| | Change the Polling Method for Certificates | 31 |
| | Managing the vRealize Automation Postgres Appliance Database | 31 |
| | Configure the Appliance Database | 33 |
| | Three Node Appliance Database Automatic Failover Scenarios | 34 |
| | Scenario: Perform Manual Appliance Database Failover | 37 |
| | Scenario: Perform a Maintenance Database Failover | 38 |
| | Manually Recover Appliance Database from Catastrophic Failure | 40 |
| | Backup and Recovery for vRealize Automation Installations | 42 |
| | The Customer Experience Improvement Program | 42 |
| | Join or Leave the Customer Experience Improvement Program for vRealize Automation | 42 |
| | Configure Data Collection Time | 43 |
| | Adjusting System Settings | 43 |
| | Modify the All Services Icon in the Service Catalog | 43 |
| | Customize Data Rollover Settings | 45 |
| | Adjusting Settings in the Manager Service Configuration File | 47 |
| | Monitoring vRealize Automation | 52 |
| | Monitoring Workflows and Viewing Logs | 52 |
| | Monitoring Event Logs and Services | 53 |
| | Using vRealize Automation Audit Logging | 55 |
| | Viewing Host Information for Clusters in Distributed Deployments | 56 |
| | Monitoring vRealize Automation Health | 59 |

| | |
|--|----|
| Configure System Tests for vRealize Automation | 59 |
| Configure Tenant Tests For vRealize Automation | 61 |
| Configure Tests For vRealize Orchestrator | 63 |
| Custom Test Suite | 65 |
| View the vRealize Automation Health Service Test Suite Results | 67 |
| Troubleshooting the Health Service | 67 |
| Monitoring vRealize Automation Environment Resources Using SNMP | 68 |
| Monitoring and Managing Resources | 69 |
| Choosing a Resource Monitoring Scenario | 69 |
| Resource Usage Terminology | 70 |
| Connecting to a Cloud Machine | 70 |
| Reducing Reservation Usage by Attrition | 73 |
| Decommissioning a Storage Path | 73 |
| Data Collection | 74 |
| Understanding vSwap Allocation Checking for vCenter Server Endpoints | 77 |
| Removing Datacenter Locations | 78 |
| Monitoring Containers | 79 |
| Bulk Import, Update, or Migrate Virtual Machines | 79 |
| Import a Virtual Machine to a vRealize Automation Environment | 80 |
| Update a Virtual Machine in a vRealize Automation Environment | 84 |
| Migrate a Virtual Machine to a Different vRealize Automation Environment | 87 |

Maintaining and Customizing vRealize Automation Components and Options

1

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

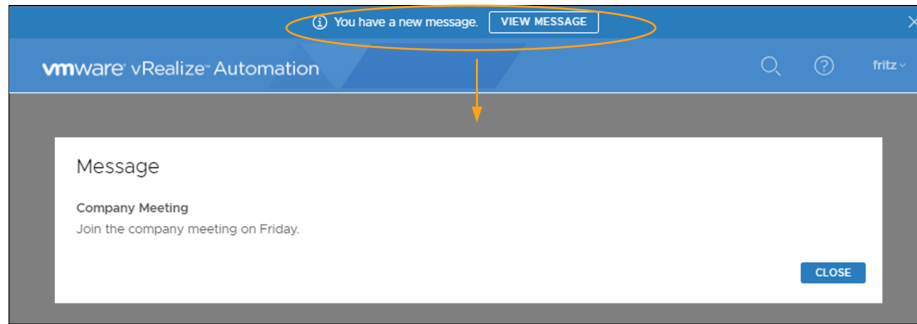
This chapter includes the following topics:

- [Broadcast a Message to All Users](#)
- [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](#)
- [The Customer Experience Improvement Program](#)
- [Adjusting System Settings](#)
- [Monitoring vRealize Automation](#)
- [Monitoring vRealize Automation Health](#)
- [Monitoring vRealize Automation Environment Resources Using SNMP](#)
- [Monitoring and Managing Resources](#)
- [Monitoring Containers](#)
- [Bulk Import, Update, or Migrate Virtual Machines](#)

Broadcast a Message to All Users

As the tenant administrator, you can broadcast a message to all users. The message notification appears at the top of browser page. Your users click the notification to see the message.

As a user, you can access the message from the banner, or from your user drop-down menu on the header.



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

The message board has the following limitations.

Table 1-1. Message Board Limitations

| Option | Limitations |
|----------------------------|--|
| URL message limitations | <ul style="list-style-type: none"> ■ The target URL must be included in the message board allowlist. See Create a Message Board URL Allowlist. ■ 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 websites 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 website. If your target website is one that you control, you can set the X-Frame-Options header to x-Frame-Options: ALLOW-FROM https://<vRealizeAutomationApplianceURL>. ■ Some websites have a redirect to a top-level page that might refresh entire vRealize Automation page. This type of website 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 allows simple markup, but does not support HTML code. For example, you cannot use <href> to link to a website. You must use the URL message option. |

Prerequisites

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

Procedure

- 1 Click the **Administration** tab.
- 2 Select **Notifications > Message Board**
- 3 In the **Type** drop-down menu, select the message type.

| Option | Description |
|-----------------------|---|
| None | Removes the message notification. |
| Custom Message | Enter a plain text message. |
| URL | <p>Enter the page URL.</p> <p>The URL must be included in the message board allowlist. See Create a Message Board URL Allowlist.</p> <p>To log the user into a website, most commonly your internal website, based on their vRealize Automation user ID, select Include user ID. The URL that is passed to the website similar to <code>http://company.com/internal/message?userID=richard_dawson@company.com</code>. This method allows your website to use the <code>window.location.search</code> JavaScript property to provide the current user's ID to your website.</p> |

- 4 Click **OK**.

Results

The message is broadcast as a banner to all your tenant users.

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, select **None** as the Type and click **OK**.

Create a Message Board URL Allowlist

As the security administrator, you configure an allowed list of URLs that can be used in the message board. This allowlist ensures added security.

Prerequisites

Log in to vRealize Automation as a **security administrator**.

Procedure

- 1 Select **Administration > Message Board Whitelist**.
- 2 Click **New**.
- 3 Add a URL and click **OK**.

The URL entries can include the following content:

- IP address or FQDN of a site. For example, `https://docs.vmware.com`.
- Includes `https`.

- Can include allowed ports. If a port is not specified, the allowed ports are 80 and 443.

4 Repeat for each additional entry.

Results

A tenant administrator cannot add a URL to the message board unless it is included in this list.

What to do next

Verify that you can add and broadcast a URL included in your allowlist using the message board. See [Broadcast a Message to All Users](#).

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 after it was powered off for any expected or unexpected reason, you must start components in a specified order.

If you are managing deployment components in vCenter Server, you can start their guest operating systems from there.

Prerequisites

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

Procedure

- 1 If you are using a legacy, standalone PostgreSQL database, start that server.
- 2 In any order, start standalone vRealize Automation MS SQL servers.
- 3 In a deployment that uses load balancers with health checks, deactivate all health checks except pings.
- 4 Start the primary vRealize Automation appliance.
- 5 In the primary vRealize Automation appliance management interface, look under the **Cluster** tab to check whether the system is in synchronous or asynchronous mode. A single-appliance deployment is always asynchronous.
 - If the deployment is synchronous, start the remaining vRealize Automation appliances.
 - If the deployment is asynchronous, go to the primary vRealize Automation appliance management interface, and wait until the licensing service is running and REGISTERED.

Afterward, start any remaining vRealize Automation appliances.

- 6 After all appliances have started, use their management interfaces to verify that services are running and REGISTERED.

It might take 15 or more minutes for appliances to start.

- 7 Start all IaaS Web nodes, and wait 5 minutes.
- 8 Start the primary Manager Service node, and wait 2 to 5 minutes.
- 9 In a distributed deployment with multiple Manager Service nodes, start secondary Manager Service nodes, and wait 2 to 5 minutes.

On secondary machines, do not start or run the Windows service unless you are configured for automatic Manager Service failover.

- 10 In any order, start the DEM Orchestrator, DEM Workers, and all vRealize Automation proxy agents.

You do not need to wait for one startup to finish before starting another.

- 11 If you had to deactivate load balancer health checks, reactivate them.
- 12 Verify that started services are running and REGISTERED.
 - a In a browser, log in to the primary vRealize Automation appliance management interface.
https://vrealize-automation-appliance-FQDN:5480
 - b Click the **Services** tab.
 - c Monitor service startup progress by clicking **Refresh**.

Results

When all services are REGISTERED, the deployment is ready.

Restart vRealize Automation

Restarting vRealize Automation components might help resolve problems. You must restart components in a specified order.

If you are managing deployment components in vCenter Server, you can restart their guest operating systems from there.

If you can't perform a restart, try the instructions in [Shut Down vRealize Automation](#) and [Start Up vRealize Automation](#) instead.

Prerequisites

- Verify that all load balancers that your deployment uses are running.

Procedure

- 1 Verify that the vRealize Automation appliance database is set to asynchronous mode. If necessary, use the management interface to change it to asynchronous mode.

You may return to synchronous mode after completing the whole procedure. See [Managing the vRealize Automation Postgres Appliance Database](#) for more information.

- 2 Restart the primary vRealize Automation appliance, and wait for startup to finish.
- 3 Use the primary vRealize Automation appliance management interface to verify that the licensing service is running and REGISTERED.
- 4 Restart the remaining vRealize Automation appliances at the same time.
- 5 Wait for the appliances to restart, and use their management interfaces to verify that services are running and REGISTERED.

It might take 15 or more minutes for appliances to restart.

- 6 Restart the primary Web node, and wait for startup to finish.
- 7 If you are running a distributed deployment with multiple Web nodes, restart secondary Web nodes, and wait for startups to finish.
- 8 Restart Manager Service nodes, and wait for startups to finish.

If you are running automatic Manager Service failover, and you want to keep the active and passive nodes the same, restart in the following order:

- a Stop the passive Manager Service nodes without restarting them.
- b Completely restart the active Manager Service node.
- c Start the passive Manager Service nodes.

- 9 In any order, restart the DEM Orchestrator, DEM Workers, and all vRealize Automation proxy agents. Wait for all startups to finish.

You do not need to wait for one restart to finish before restarting another.

- 10 Verify that restarted services are running and REGISTERED.
 - a In a browser, log in to the primary vRealize Automation appliance management interface.
https://vrealize-automation-appliance-FQDN:5480
 - b Click the **Services** tab.
 - c Monitor service startup progress by clicking **Refresh**.

Results

When all services are REGISTERED, the deployment is ready.

Shut Down vRealize Automation

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

If you are managing deployment components in vCenter Server, you can shut down their guest operating systems from there.

Procedure

- 1 In any order, shut down the DEM Orchestrator, DEM Workers, and all vRealize Automation proxy agents. Wait for shutdown to finish.
- 2 Shut down Manager Service nodes, and wait for shutdown to finish.
- 3 In distributed deployments with multiple Web nodes, shut down secondary Web nodes, and wait for shutdown to finish.
- 4 Shut down the primary Web node, and wait for shutdown to finish.
- 5 In distributed deployments with multiple vRealize Automation appliances in synchronous mode, use the vRealize Automation appliance management interface to change to asynchronous mode.
- 6 In distributed deployments with multiple vRealize Automation appliances, shut down secondary appliances, and wait for shutdown to finish.
- 7 Shut down the primary vRealize Automation appliance, and wait for shutdown to finish.

The primary vRealize Automation appliance is the one that contains the primary, or writeable, appliance database. Make note of which appliance is primary so that you can start back up in the correct order.
- 8 In any order, shut down any standalone vRealize Automation MS SQL servers, and wait for shutdown to finish.
- 9 If you are using a legacy, standalone PostgreSQL database, shut down that server.

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:

- vRealize Automation appliance
- IaaS website component
- IaaS manager service component

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

Note vRealize Automation uses several third party products, such as Rabbit MQ, to support a variety of functionality. Some of these products use their own self signed certificates that persist even if you replace primary vRealize Automation certificates with certificates supplied by a CA. Because of this situation, users cannot effectively control certificate use on specific ports, such as 5671 which is used by RabbitMQ for internal communication.

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.

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 might need to update to SHA2 certificates due to operating system or browser requirements.

The vRealize Automation appliance management interface provides options for updating or replacing certificates.

In a clustered deployment, you must initiate changes from the primary node interface.

- **Generate certificate** — Have vRealize Automation generate a self-signed certificate.
- **Import certificate** — Use your own certificate.
- **Provide certificate thumbprint** — Provide a certificate thumb print to use a certificate already in the certificate store on IaaS Windows servers.

This option does not transmit the certificate from the vRealize Automation appliance to IaaS Windows servers. The option allows users to deploy existing certificates already on IaaS Windows servers without uploading the certificates in the vRealize Automation appliance management interface.

- **Keep Existing** — Continue to use the current certificate.

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

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

Virtual Machine Templates

After you change vRealize Automation appliance or IaaS Windows server certificates, you must update vRealize Automation guest and software agents on virtual machine templates so that the templates work again in vRealize Automation. If you don't update the agents, deployment requests involving software components fail with an error similar to the following example.

```
The following component requests failed: Linux. Request failed: Machine VM-001:
InstallSoftwareWorkflow. Install software work item timeout.
```

vRealize Orchestrator

After you change vRealize Automation certificates, you must update vRealize Orchestrator to trust the new certificates.

The vRealize Orchestrator component associated with your vRealize Automation deployment has its own certificates, but it must also trust the vRealize Automation certificates. By default, the vRealize Orchestrator component is embedded in vRealize Automation, although a few users elect to use an external vRealize Orchestrator. In either case, see the vRealize Orchestrator documentation for information about updating vRealize Orchestrator certificates.

If you run a multiple-node vRealize Orchestrator deployment behind a load balancer, all vRealize Orchestrator nodes must use the same certificate.

For More Information

For more about certificate troubleshooting, supportability, and trust requirements, see [VMware Knowledge Base article 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 | openssl pkcs12 -in <i>path_to_.pfx</i> <i>certificate_file</i> -nocerts -out key.pem | RSA Private Key |
| PEM File | openssl pkcs12 -in <i>path_to_.pfx</i> <i>certificate_file</i> -clcerts -nokeys -out cert.pem | 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 Log in to the vRealize Automation appliance management interface as root.
`https://vrealize-automation-appliance-FQDN:5480`
- 2 Select **vRA > Certificates**.
- 3 Select the vRealize Automation component for which you are updating the certificate.
- 4 Select the appropriate action 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.

If you want to generate a CSR request for a new certificate that you can submit to a certificate authority, select **Generate Signing Request**. A CSR helps your CA create a certificate with the correct values for you to import.

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 | <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. |
| Generate Signing Request | <ol style="list-style-type: none"> a Select Generate Signing Request. b Review the entries in the Organization, Organization Unit, Country Code, and Common Name text boxes. These entries are populated from the existing certificate. You can edit these entries if needed. c Click Generate CSR to generate a certificate signing request, and then click the Download the generated CSR here link to open a dialog that enables you to save the CSR to a location where you can send it to a certificate authority. d When you receive the prepared certificate, click Import and follow instructions for importing a certificate into vRealize Automation. |
| Import | <ol 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> <ol 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**.

A vRealize Automation appliance certificate update requires vRealize Automation services to gracefully restart. The restart might take anywhere from 15 minutes to an hour depending on the number of vRealize Automation appliances in your environment.

After the restart, 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 deactivate SSH when finished.
- e Select the **Administrator SSH login** check box.
Deselect the check box to deactivate 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.

Results

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.

Procedure

- 1 Log in to the vRealize Automation appliance management interface as root.

`https://vrealize-automation-appliance-FQDN:5480`

- 2 Select **vRA > Certificates**.
- 3 Click **IaaS Web** on the **Component Type** menu.
- 4 Go to the **IaaS Web Certificate** pane.
- 5 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 are 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 | 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**.

An IaaS Windows server certificate update requires vRealize Automation services to gracefully restart. The restart might take anywhere from 15 minutes to an hour depending on the number of vRealize Automation appliances in your environment.

After the restart, 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 Open a Web browser to the vRealize Automation appliance management interface URL.
- 2 Log in with user name **root** and the password you specified when deploying the vRealize Automation appliance.
- 3 Select **vRA > Certificates**.
- 4 Click **Manager Service** from the **Component 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 | <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. |
| Import | <ol 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> <ol 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.

Update Embedded vRealize Orchestrator to Trust vRealize Automation Certificates

If you update or change vRealize Automation appliance or IaaS certificates, you must update vRealize Orchestrator to trust the new or updated certificates.

This procedure applies to all vRealize Automation deployments that use an embedded vRealize Orchestrator instance. If you use an external vRealize Orchestrator instance, see [Update External vRealize Orchestrator to Trust vRealize Automation Certificates](#).

Note This procedure resets tenant and group authentication back to the default settings. If you have customized your authentication configuration, note your changes so that you can re-configure authentication after completing the procedure.

See the vRealize Orchestrator documentation for information about updating and replacing vRealize Orchestrator certificates.

In a clustered configuration, you must complete this procedure on the primary vRealize Automation appliance node and then perform a `join-cluster` against the primary from each replica vRealize Automation appliance node.

Note In a cluster, stop the `vco-configurator` service on all replica nodes until the procedure is completed to avoid unwanted automatic control center synchronization.

If you replace or update vRealize Automation certificates without completing this procedure, the vRealize Orchestrator Control Center may be inaccessible, and errors may appear in the `vco-server` and `vco-configurator` log files.

Problems with updating certificates can also occur if vRealize Orchestrator is configured to authenticate against a different tenant and group than vRealize Automation. For information, see VMware Knowledge Base article [Exception Untrusted certificate chain after replacing vRA certificate \(2147612\)](#).

The trust command syntaxes shown herein are representative rather than definitive. While they are appropriate for most typical deployments, there may be situations in which you need to experiment with variations on the commands.

- If you specify `--certificate` you must provide the path to a valid certificate file in PEM format.
- If you specify `--uri`, you must provide the uri from which the command can fetch a trusted certificate.
- If you specify the `--registry-certificate` option, you indicate that the requested certificate should be treated as the certificate for the component registry and the trusted certificate is added to the truststore under a specific alias used by the component registry certificate.

You can also manage certificates by using SSL Trust Manager workflows in vRealize Orchestrator. For information, see the *Manage Orchestrator Certificates* topic in the [vRealize Orchestrator documentation](#).

Procedure

- 1 Stop the vRealize Orchestrator server and Control Center services.

```
service vco-server stop
service vco-configurator stop
```

- 2 Reset the vRealize Orchestrator authentication provider by running the following command.

```
/var/lib/vco/tools/configuration-cli/bin/vro-configure.sh reset-authentication
ls -l /etc/vco/app-server/
mv /etc/vco/app-server/vco-registration-id /etc/vco/app-server/vco-registration-id.old
vcac-vami vco-service-reconfigure
```

- 3 Check the trusted certificate for the vRealize Orchestrator trust store using the command line interface utility located at `/var/lib/vco/tools/configuration-cli/bin` with the following command.

```
/var/lib/vco/tools/configuration-cli/bin/vro-configure.sh list-trust
```

- Check for the certificate with the following alias: `vco.cafe.component-registry.ssl.certificate`. This should be the vRealize Automation certificate that the vRealize Orchestrator instance uses as an authentication provider.
- This certificate must match the newly configured vRealize Automation certificate. If it does not match, it can be changed as follows:
 - 1 Copy your vRealize Automation signed appliance certificate PEM file to the `/tmp` folder on the appliance.
 - 2 Run the following command adding the appropriate certificate path.

```
./vro-configure.sh trust --certificate path-to-the-certificate-file-in-PEM-format--registry-certificate
```

See the following example command.

```
/var/lib/vco/tools/configuration-cli/bin/vro-configure.sh trust --certificate /var/tmp/test.pem --registry-certificate
```

- 4 You may need to run the following commands to trust the certificate.

```
/var/lib/vco/tools/configuration-cli/bin/vro-configure.sh trust --uri https://vra.domain.com

/var/lib/vco/tools/configuration-cli/bin/vro-configure.sh trust --registry-certificate --uri https://vra.domain.com
```

- 5 Ensure that the vRealize Automation certificate is now injected into the vRealize Orchestrator trust store using the following command.

```
/var/lib/vco/tools/configuration-cli/bin/vro-configure.sh list-trust
```

- 6 Start the vRealize Orchestrator server and control center services.

```
service vco-server start
service vco-configurator start
```

What to do next

You can validate that trust has been updated on a clustered system.

- 1 Log in to the virtual appliance management interface as root.
- 2 Select the Services page.
- 3 Ensure that there are no duplicate vco services listed.
If you see any duplication of the vco services listed, click **Unregister** to remove the services that do not have a state of Registered.
- 4 Ensure that vco-configurator is started on all virtual appliance nodes.
- 5 Log in to the vRealize Orchestrator control center and navigate to the Validate Configuration page to validate the configuration.
- 6 Navigate to the Authentication Provider page, and verify that the auth settings are correct.

You can also test the login credentials on this page.

Update External vRealize Orchestrator to Trust vRealize Automation Certificates

If you update or change vRealize Automation appliance or IaaS certificates, you must update vRealize Orchestrator to trust the new or updated certificates.

This procedure applies to vRealize Automation deployments that use an external vRealize Orchestrator instance.

Note This procedure resets tenant and group authentication back to the default settings. If you have customized your authentication configuration, note your changes so that you can re-configure authentication after completing the procedure.

See the vRealize Orchestrator documentation for information about updating and replacing vRealize Orchestrator certificates.

If you replace or update vRealize Automation certificates without completing this procedure, the vRealize Orchestrator Control Center may be inaccessible, and errors may appear in the vco-server and vco-configurator log files.

Problems with updating certificates can also occur if vRealize Orchestrator is configured to authenticate against a different tenant and group than vRealize Automation. See [Knowledge Base article 2147612](#).

Procedure

- 1 Stop the vRealize Orchestrator server and Control Center services.

```
service vco-configurator stop
```
- 2 Reset the vRealize Orchestrator authentication provider.

```
/var/lib/vco/tools/configuration-cli/bin/vro-configure.sh reset-authentication
```
- 3 Start the vRealize Orchestrator Control Center service.

```
service vco-configurator start
```
- 4 Log in to the Control Center using virtual appliance management interface root credentials.
- 5 Unregister and re-register the authentication provider.

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

If the SSL certificate of the management site service expires, or you started with a self-signed certificate and site policies require a different one, you can replace the certificate.

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

If the SSL certificate of the management site service expires, or you started with a self-signed certificate and site policies require a different one, you can replace the certificate.

You are allowed to reuse the certificate used by the vRealize Automation service on port 443, or use a different one. If you are requesting a new CA-issued certificate to update an existing certificate, a best practice is to reuse the Common Name from the existing certificate.

Note The vRealize Automation appliance uses lighttpd to run its own management site. You secure the management site service on port 5480.

Prerequisites

- The certificate must be in PEM format.
- The certificate must include both of the following, in order, together in one file:
 - a RSA private key
 - b Certificate chain

- The private key cannot be encrypted.
- The default location and file name is `/opt/vmware/etc/lighttpd/server.pem`.

See [Extracting Certificates and Private Keys](#) for more 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 Run the following command to restart the haproxy service.

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

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 host 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 by 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 select the **Cluster** tab.
- 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.
- 3** Import the newly generated certificate into the local computer .personal store, or do not import anything if you want the system to auto-generate a new self-signed certificate.

- 4 Register the Management Agent certificate with the vRealize Automation appliance management site.

- a Open a command prompt as an administrator and navigate to the Cafe directory on the machine on which the Management Agent is installed at `<vra-installation-dir>\Management Agent\Tools\Cafe`, typically `C:\Program Files (x86)\VMware\vCAC\Management Agent\Tools\Cafe`.
- b Enter the `Vcac-Config.exe RegisterNode` command with options to register the Management Agent identifier and certificate in one step. Include the Management Agent identifier you recorded earlier as the value for the `-nd` option.

Table 1-3. Required Options and Arguments for Vcac-Config.exe RegisterNode

[illegible]

The following example shows the command format:

```
Vcac-Config.exe RegisterNode -v -vamih "vra-vr-hostname.domain.name:5480"  
-cu "root" -cp "password" -hn "machine-hostname.domain.name"  
-nd "00000000-0000-0000-0000-000000000000"  
-tp "000000000000000000000000000000000000000000000000"
```

5 Restart the Management Agent.

Example: Command to Register a Management Agent Certificate

```
Vcac-Config.exe RegisterNode -v -vami "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 there are commas in the OU section of the IaaS certificate, you might encounter STOMP WebSocket errors in the Manager Service log files. In addition, virtual machine provisioning might fail. You can remove the commas, or change the polling method from WebSocket to HTTP.

To change the polling method, take the following steps.

Procedure

- 1 Open the following file in a text editor.

```
C:\Program Files (x86)\VMware\VCAC\Server\Manager Service.exe.config.
```

- 2 Add the following lines inside the <appSettings> section.

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

- 3 Save and close Manager Service.exe.config.
- 4 Restart the Manager Service.

Results

For more information about the Manager Service, see *Installing vRealize Automation*.

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 Automation 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 or with multiple nodes to facilitate high availability through failover. The vRealize Automation installer includes a database node on each vRealize Automation appliance installation. So if you install three instances of a vRealize Automation appliance, you have three database nodes. Automatic failover is implemented on applicable deployments. The appliance database requires no maintenance unless a machine configuration changes or, if you use a clustered configuration, you promote a different node for the primary.

Note The database clustered configuration is set up automatically when you join a virtual appliance to the cluster using the Join cluster operation. 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.

For high availability, vRealize Automation uses the PostgreSQL primary-replica model to support data replication. This means that all of the database nodes work in a cluster with one leading node, known as the primary, and several replicating nodes, known as replicas. The primary node handles all database requests and the replica nodes stream and replay transactions from the primary locally.

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

Several high availability appliance database options exist. Selecting 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 if the primary 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. To support high availability database failover, asynchronous mode requires two nodes, whereas synchronous mode requires three nodes. Synchronous mode also invokes automatic failover.

| Replication Mode | Advantages | Disadvantages |
|------------------|--|---|
| Synchronous | <ul style="list-style-type: none"> ■ Minimizes chance of data loss. ■ Invokes automatic failover. | <ul style="list-style-type: none"> ■ Might affect system performance. ■ Requires three nodes. |
| Asynchronous | <ul style="list-style-type: none"> ■ Requires only two nodes. ■ Affects system performance less than synchronous mode. | <ul style="list-style-type: none"> ■ 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 **Cluster** tab on the Virtual Appliance Management Interface enables you to switch synchronization modes and to add database nodes as needed.

When operating in synchronous mode, vRealize Automation invokes automatic failover.

If you begin 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.

Appliance Database Failover

In a high availability configuration, the primary constantly streams transactions to the replica servers. If the primary fails, the active and working replica is ready to proceed with read-only requests. When the new primary is promoted, either manually or automatically, all of the upcoming requests are moved to it.

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 primary 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 primary. 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 | Primary node - no status Replica node - sync Other nodes - potential |
| Asynchronous replication | Primary node - no status Other nodes - potential |

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

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

When operating in synchronous mode, vRealize Automation invokes automatic failover. In the event of primary node failure the next available replica node will automatically become the new primary. The failover operation requires 10 to 30 seconds on a typical vRealize Automation deployment.

Prerequisites

- Install and configure vRealize Automation according to appropriate instructions in *Installing vRealize Automation*.
- Log in to vRealize Automation Appliance Management as **root** using the password you entered when you deployed the vRealize Automation appliance.
- 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 primary 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 primary, click **Promote** in the appropriate column.
- 4 Click **Save Settings** to save your configuration if you have made any changes.

Three Node Appliance Database Automatic Failover Scenarios

There are several appliance database high availability failover scenarios, and vRealize Automation behavior varies depending on appliance database configuration and the number of nodes that fail.

Single Node Failure Scenarios

If one of the three nodes fails, vRealize Automation will initiate an auto failover. No additional auto failover operations can occur until all three nodes are restored.

The following table describes behavior and actions related to a primary node failure in a high availability deployment.

Table 1-5. The Primary Node Fails

| | |
|-------------------|--|
| Expected Behavior | <ul style="list-style-type: none"> ■ The configured sync replica node becomes the primary and automatically picks up appliance database functionality. ■ The potential sync replica becomes the sync standby node. ■ The vRealize Automation deployment functions in read only mode until the automatic failover completes. |
| Further Action | <ul style="list-style-type: none"> ■ When the former primary is recovered, it will be reset as replica automatically by the failover agent repair logic. No manual action is required. ■ If the former primary cannot be recovered, manually set the appliance database to asynchronous mode. |

The following table describes behavior and actions related to a sync replica node failure in a high availability deployment.

Table 1-6. The Sync Replica Fails

| | |
|-------------------|---|
| Expected Behavior | <ul style="list-style-type: none"> ■ The vRealize Automation deployment experiences no downtime. There will be a delay of a couple of seconds for database requests until the potential replica becomes the new sync replica. The appliance database performs this action automatically. |
| Further Action | <ul style="list-style-type: none"> ■ When the former synch replica comes online, it will become a potential replica automatically. No manual action is required. ■ If the former sync replica cannot be repaired, manually set the appliance database to asynchronous mode. |

The following table describes behavior and actions related to a primary node failure in a high availability deployment.

Table 1-7. The Potential Replica Fails

| | |
|-------------------|---|
| Expected Behavior | No deployment downtime. |
| Further Action | <ul style="list-style-type: none"> ■ When the former potential replica comes online, it becomes a potential replica automatically. No manual action is required. ■ If the former potential replica cannot be repaired, set the appliance database to asynchronous mode. |

Two Node Failure Scenarios

If two out of the three nodes fail simultaneously, vRealize Automation switches to read only mode until a manual repair is performed.

The following table describes behavior and actions related to a primary node and potential replica node failure in a high availability deployment.

Table 1-8. The Primary Node and Potential Replica Fail

| | |
|-------------------|--|
| Expected Behavior | <ul style="list-style-type: none"> ■ The sync replica is not promoted to primary automatically. vRealize Automation functions in read only mode as it is able to process read-only transactions until a manual promotion is performed. |
| Further Action | <ul style="list-style-type: none"> ■ Manual promotion is required. Set the appliance database to asynchronous mode. ■ When the primary and potential replica are recovered, manually set them to synchronize against the new primary. At that point, you can switch vRealize Automation back to synchronous mode. ■ When two out of three nodes are down simultaneously, vRealize Automation will switch to read-only mode until you effect a manual repair. If only one database node is available, switch your deployment to asynchronous mode. |

The following table describes behavior and actions related to Sync and Potential node failure in a high availability deployment.

Table 1-9. The Sync and Potential Replicas Fail

| | |
|-------------------|--|
| Expected Behavior | <ul style="list-style-type: none"> ■ vRealize Automation functions in read only mode as it is able to process read-only transactions until a manual repair is performed. |
| Further Action | <ul style="list-style-type: none"> ■ Manual promotion is required. Set the appliance database to asynchronous mode. ■ When the sync and potential replicas are recovered, they should be manually reset to synchronize against the primary. At this point, you can switch vRealize Automation back to synchronous mode. ■ When two out of three nodes are down simultaneously, vRealize Automation will switch to read-only mode until you effect a manual repair. If only one database node is available, switch your deployment to asynchronous mode. |

Links Failures Among Nodes

If a link failure occurs among nodes on a distributed deployment, the automatic failover agent attempts to repair the configuration.

The following table describes behavior and actions related to a link failure between two sites in a high availability deployment with the specified configuration when all nodes remain up and online.

Site A: Primary and potential replica

Site B: Sync replica

Table 1-10. Link Failure Between Two Sites when all Nodes Remain Up and Online

| | |
|-------------------|---|
| Expected Behavior | No downtime for the vRealize Automation deployment. The potential replica automatically becomes the sync replica. |
| Further Action | No manual action is required. |

The following table describes behavior and actions related to a link failure between two sites in a high availability deployment with the specified configuration when all nodes remain up and online.

Site A: Primary

Site B: Sync and potential replica

Table 1-11. Link Failure Between Two Sites when all Nodes Remain Up and Online - Alternate Configuration

| | |
|-------------------|---|
| Expected Behavior | Sync replica becomes the primary and automatically picks up appliance database functionality. Automatic failover agent promotes the potential replica to become the new sync replica. vRealize Automation deployment operates in read only mode until this promotion completes. |
| Further Action | No manual action is required. When the link is recovered, the automatic failover agent resets the former primary as replica. |

Scenario: 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 primary vRealize Automation appliance node fails or stops running.

Note Once a node goes into a unhealthy state, do not attempt to use its virtual appliance management interface for any operations including failover.

Prerequisites

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

Procedure

- 1 Remove the primary 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 Select **Cluster**.
- 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 primary node.

6 Correct issues with the former primary node and add it back to the cluster:

- a Isolate the former primary 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 primary 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 primary 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 primary node management interface as root.

- g Select **Cluster**.

- h Locate the former primary node, and click **Reset**.

- i After a successful reset, restart the former primary node.

- j With the former primary 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 primary node to the external load balancer.

Note If a primary node that was demoted to replica is still listed as primary, 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 primary node is up and running normally. There are two database failover maintenance steps: maintenance of the primary and maintenance of a replica node. When a primary node has been replaced so that it becomes a replica, you should perform maintenance on it so that it is suitable to become the primary again should the need arise.

Note Do not stop or restart the HAProxy service on the applicable host machine while performing a maintenance failover.

Prerequisites

- vRealize Automation is installed and configured according to appropriate instructions in the *Installing vRealize Automation*.
- Log in to vRealize Automation Appliance Management as **root** using the password you entered when you deployed the vRealize Automation appliance.
- Install and configure an appropriate embedded Postgres appliance database cluster.
- If your database uses synchronous replication mode, ensure that there are three active nodes in the cluster.

Procedure

- 1 Remove the primary node IP address from the external load balancer.
- 2 Isolate the primary node.
Disconnect the node from its current network. This should be the network that is routing to the remaining vRealize Automation appliance nodes.
- 3 Select another NIC for management, or manage it directly from the Virtual Appliance Management Interface.
- 4 Select **Cluster** on the Virtual Appliance Management Interface.
- 5 Select the replica node with the lowest priority for promotion to the primary, and click **Promote**.
Replica nodes appear in ascending priority order.
The old primary is demoted to replica status, and the new primary is promoted.
- 6 Perform the appropriate replica maintenance.
- 7 When the maintenance is complete, ensure that the virtual appliance is running with network connectivity and that its HAProxy service is running.
 - a Log in to the vRealize Automation management console as **root**.
 - b Ensure that the replica node can be pinged, resolved by name, and has a recent status in the Virtual Appliance Management Interface **Cluster** tab.

- 8 Click **Reset** for the replica node.

This operation resets the database so that it is configured to replicate to the current primary and re-synchronizes the replica node with the latest haproxy configuration from the primary node.

- 9 Following successful reset, return the replica virtual appliance node IP address to the external virtual appliance load balancer IP address pool.
- 10 Ensure that the replica node appears healthy on the database table and that it can be pinged and resolved by name.

What to do next

Correct issues with the former primary node and add it back to the cluster.

Manually Recover Appliance Database from Catastrophic Failure

If the appliance database fails, and no database nodes are up and running or all replica nodes are out of sync when the primary fails, use the following procedure to attempt to recover the database.

This procedure applies to situations in which no database nodes are operational across a cluster that is running in asynchronous mode. In this scenario, you typically see errors similar to the following on the Virtual Appliance Management Interface page when trying to load or refresh the page:

Error initializing the database service: Could not open JDBC Connection for transaction; nested exception is org.postgresql.util.PSQLException: The connection attempt failed.

Procedure

- 1 Try to recover the database using the Virtual Appliance Management Interface from one of the database nodes.
 - a If possible, open the Virtual Appliance Management Interface **Cluster** page of the node with the most recent state. Typically, this node is the one that was the primary node before the database failed.
 - b If the Virtual Appliance Management Interface for the primary node fails to open, try to open the Interface for other replica nodes.
 - c If you can find a database node with a working Virtual Appliance Management Interface, try to recover it by performing a manual failover.

See [Scenario: Perform Manual vRealize Automation Appliance Database Failover](#).

- 2 If the procedure in step 1 fails, start a shell session and try to determine the node with the most recent state. Start a shell session to all the available cluster nodes and try to start their databases by running the following shell command: `service vpostgres start`

- 3 Use the following procedure for each node that has a running local database to determine the node with the most recent state.

- a Run the following command to determine the node with the most recent state. If the command returns f, then it is the node with most recent state and you can proceed to step 4.

```
su - postgres
psql vcac
vcac=# select pg_is_in_recovery();
pg_is_in_recovery
```

- If this command returns an f, then this node has the most recent state.
- If the node returns a t, run the following command on the node:

```
SELECT pg_last_xlog_receive_location() as receive_loc, pg_last_xlog_replay_location() as
replay_loc, extract(epoch from pg_last_xact_replay_timestamp()) as replay_timestamp;
```

This command should return a result similar to the following.

```
vcac=# SELECT pg_last_xlog_receive_location() as receive_loc, pg_last_xlog_replay_location()
as replay_loc, extract(epoch from pg_last_xact_replay_timestamp()) as replay_timestamp;
 receive_loc | replay_loc | replay_timestamp
-----+-----+-----
 0/20000000 | 0/203228A0 | 1491577215.68858
(1 row)
```

- 4 Compare the results for each node to determine which one has the most recent state.

Select the node with greatest value under the `receive_loc` column. If equal, select the greatest from the `replay_loc` column and then, if again equal, select the node with greatest value of `replay_timestamp`.

- 5 Run the following command on the node with the most recent state: `vcac-vami psql-promote-master -force`
- 6 Open the `/etc/haproxy/conf.d/10-psql.cfg` file in a text editor and update the following line.

```
server masterserver sc-rdops-vm06-dhcp-170-156.eng.vmware.com:5432 check on-marked-up shutdown-
backup-sessions
```

To read as follows with the current node FQDN:

```
server masterserver current-node-fqdn:5432 check on-marked-up shutdown-backup-sessions
```

- 7 Save the file.
- 8 Run the `service haproxy restart` command.

- 9 Open the Virtual Appliance Management Interface **Cluster** page for the most recent node.

This node should appear as the primary node with the other nodes as invalid replicas. In addition, the **Reset** button for the replicas is enabled.

- 10 Click **Reset** and for each replica in succession until the cluster state is repaired.

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 back up and restore vRealize Automation, see the following topics in the [vRealize Suite documentation](#):

- vRealize Automation Preparations for Backing Up
- vRealize Automation System Recovery

The Customer Experience Improvement Program

This product participates in VMware's 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 you on how best to deploy and use our products. You can choose to join or leave the CEIP for vRealize Automation 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>.

Join or Leave the Customer Experience Improvement Program for vRealize Automation

You can join or leave the Customer Experience Improvement Program (CEIP) for vRealize Automation at any time.

vRealize Automation gives you the opportunity to join the Customer Experience Improvement Program (CEIP) when you initially install and configure the product. After installation, you can join or leave the CEIP by following these steps.

Procedure

- 1 Log in as root to the vRealize Automation appliance management interface.
`https://vrealize-automation-appliance-FQDN:5480`
- 2 Click the **Telemetry** tab.
- 3 Check or uncheck the **Join the VMware Customer Experience Improvement Program** option.
When checked, the option activates the Program and sends data to `https://vmware.com`.

4 Click **Save Settings**.

Configure Data Collection Time

You can set the day and time when the Customer Experience Improvement Program (CEIP) sends data to VMware.

Procedure

1 Log in to a console session on the vRealize Automation appliance as root.

2 Open the following file in a text editor.

```
/etc/telemetry/telemetry-collector-vami.properties
```

3 Edit the properties for day of week (dow) and hour of day (hod).

| Property | Description |
|--|--|
| <code>frequency.dow=<day-of-week></code> | Day when data collection occurs. |
| <code>frequency.hod=<hour-of-day></code> | Local time of day when data collection occurs. Possible values are 0–23. |

4 Save and close `telemetry-collector-vami.properties`.

5 Apply the settings by entering the following command.

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

Changes are applied to all nodes in your deployment.

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.
- 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 | <code>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"}'</code> |
| Windows | <code>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\"}"</code> |

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 | <code>export AUTH="Bearer <Auth Token>"</code> |
| Windows | <code>set AUTH=Bearer <Auth Token></code> |

- 4 Add the base64 encoded string for the image.

| Operating System | Command |
|------------------|---|
| Linux/Mac | <code>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>"}</code> |
| Windows | <code>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>\"}"</code> |

Results

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 | <code>curl https://\$VCAC/catalog-service/api/icons/cafe_default_icon_genericAllServices --insecure -H "Authorization: \$AUTH" --request DELETE</code> |
| Windows | <code>curl https://%VCAC%/catalog-service/api/icons/cafe_default_icon_genericAllServices --insecure -H "Authorization: %AUTH%" --request DELETE</code> |

Customize Data Rollover Settings

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

Use the data rollover feature to enable rollover, set the maximum number of days for vRealize Automation to retain data in the IaaS SQL Server database before archiving or deleting it, and other data rollover controls.

By default, the data rollover feature is deactivated.

Configure data rollover settings on the vRealize Automation **Global Settings** page. When activated, 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 AM 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. Note that this process generally runs quickly within a few minutes to an hour. However,

when this feature is first turned on the process may have a lot of data to archive/delete to catch up and thus could take much longer to complete. This process is designed to run until done. It performs its work in small and quick batch sized transactional chunks of work as to not cause concurrency issues. Note that this process can be gracefully stopped as described below.

Note You can stop the DataRollover process by changing the `DataRollover Status` setting of Running to Disabled or Enabled. This causes the currently running process to quit gracefully. No work is lost. All data archived or deleted up to the point of stopping the process is saved.

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

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 |
|--|--|
| <code>DataRollover BatchSize</code> | This is defaulted to 2000 and probably does not need to be changed. However, if there seem to be some performance impacts, then a smaller BatchSize may help. A larger BatchSize may get the job done faster, but will put more pressure on concurrent processing. Valid range is 100 to 20000. |
| <code>DataRollover IsArchiveEnabled</code> | Specifies whether to move rollover data to archive tables after the maximum number of days is reached. By default this value is set to True. If you set this value to False, all data older than that specified in the <code>DataRollover MaximumAgeInDays</code> setting is permanently deleted. |
| <code>DataRollover MaximumAgeInDays</code> | Specifies the maximum number of days that the system retains data in the database before moving it to archive or permanently deleting it. By default this value is set to 90 days. |
| <code>DataRollover Status</code> | Specifies whether to enable data rollover. By default this value is set to Disabled. To enable data rollover, set the value to Enabled. |

| Setting | Description |
|--|--|
| DataRollover VirtualMachineHistory BatchSize | Specifies batch size in the VirtualMachineHistory table in the range of 1 - 5 records. The default is 1. |
| DataRollover UpdateStatistics | The UpdateStatistics is off by default, but is highly recommended to be turned on (set to 1) as updated statistics is good for query performance. This causes the [dbo].[usp_DataRollover] stored procedure to perform update statistics command on the tables after the archival process has run. |

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

The **Value** area for the applicable setting becomes editable.

- 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="3600000"/>`
- `<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 |
|---|---|
| MaxOutstandingResourceIntensiveWorkItems | Concurrent provisioning limit (default is 8) |
| CloneExecutionTimeout | Virtual provisioning execution timeout interval |
| SetupOSExecutionTimeout | Virtual provisioning execution timeout interval |
| CloneTimeout | Virtual provisioning clone delivery timeout interval |
| SetupOSTimeout | Virtual provisioning setup OS delivery timeout interval |
| CloudInitializeProvisioning | Cloud provisioning initialization timeout interval |
| MaxOutstandingDataCollectionWorkItems | Concurrent data collection limit |
| InventoryTimeout | Inventory data collection execution timeout interval |
| PerformanceTimeout | Performance data collection execution timeout interval |
| StateTimeout | 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 *ProcessLeaseWorkflowTimerCallbackIntervalMilliseconds* 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 | <p>Specify a severity level to ignore events below that severity.</p> <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 <p>For example, <code><add key="RepositoryLogSeverity" value="Warning" /></code>.</p> |
| RepositoryLogCategory | <p>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.</p> |

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

Results

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 1-12. 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. | IaaS 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, NSX, reclamation, and reconfigure actions. |
| View the status of scheduled and available Distributed Execution Manager and other workflows. | IaaS 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. | IaaS 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. | IaaS 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. | IaaS 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.

The default retention period for the event logs is 90 days. You can change the period from the `/etc/vcac/vcac.properties` file.

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 1-13. 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 1-14. 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 1-15. 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 1-16. IaaS Services Group

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

Table 1-17. XaaS

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

Using vRealize Automation Audit Logging

vRealize Automation offers audit logging to support collection and retention of important system events.

Currently, vRealize Automation supports audit logging as an extension of event logging. This functionality provides basic auditing information, and retention settings are configurable only using the appropriate vRealize Automation REST API event broker service calls. Audit logging is currently available to tenant administrators and system administrators who can log on to tenants. It provides search and filter capabilities for events.

By default, vRealize Automation supports audit logging for workflow subscription, endpoint, and fabric group create, update, and delete events. vRealize Automation also supports audit logging customization for a variety of IaaS events as well.

vRealize Automation audit logging is deactivated by default. You can switch it on or off by toggling the **Enabled** check box in the Audit Log Integration section on the **vRA > Logs** page of the virtual appliance management interface.

Audit log information appears on the standard Event Logs page. As a tenant admin, select **Administration > Event Logs** to view this page. Audit events are identified in the event log table with the designation Audit in the Event Type field. Each entry shows an Event Description for each event as well as the Tenant, Time, User, and related Service Name.

Enabling audit logging for any other IaaS events requires a custom configuration file and running the appropriate commands on your IaaS host machine. Contact VMware Professional Services for assistance.

You can configure vRealize Automation to export events to an external syslog server, specifically VMware Log Insight.

Configure vRealize Automation for Log Insight Audit Logging

You can configure vRealize Automation to export audit events to VMware Log Insight to facilitate viewing audit events.

Audit logging is deactivated by default and you must enable it to generate and view audit logging events.

If used, SSL is configured on the vRealize Automation appliance where the Log Insight agent resides, and it concerns the connection to the Log Insight Syslog server. To use SSL, you must configure the appropriate certificates and connectivity between vRealize Automation and the Log Insight server installed on your deployment.

Prerequisites

vRealize Automation uses the Log Insight Agent that is installed by default on a vRealize Automation deployment to read log entries for viewing in Log Insight.

Procedure

- 1 Log in to the Virtual Appliance Management Interface as a system administrator.
- 2 Select **vRA > Logs**.
- 3 Verify that the **Enabled** check box for audit logging is selected under the Audit Log Integration heading.
- 4 Enter the **Host** machine name for the Log Insight server under the Log Insight Agent Configuration heading.
 - a Enter the **Host** machine name for the Log Insight agent.
 - b Enter the **Port** to be used for communication with the Log Insight agent.
 - c Select the appropriate communication protocol.
 - d Use the **SSL Enabled** check box to indicate whether SSL will be used for communication between the Log Insight agent and server.

If you choose not to use SSL, you can ignore the remainder of the settings on the page. If SSL is used, you must configure these settings.

- 5 Make the appropriate selections in the SSL Trusted Root Certificates section if you are using SSL.

By default, the vRealize Automation appliance uses a self-signed certificate. If you want to use a Trusted Root certificate, you must import it.

- a Select the appropriate check box to indicate whether you want to use a new certificate or an existing certificate.

See the notes on the Virtual Appliance Management Interface Configure vRealize Automation Logging page for more information.

- 6 Click **Save Settings**.
- 7 Make the appropriate selections in the SSL Server Certificates section.
- 8 Use the Agent Behavior Configuration section to configure how the agent works with log files.

Results

vRealize Automation audit log events are visible from the Log Insight interface.

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

| | Host / Node Name | Version | Last Connected | Type | State* | Valid* |
|---|------------------------------|-------------|----------------|--------|--------|-------------------------|
| ▶ | cava-n-80-175.eng.vmware.com | 7.5.0.378 | 7 minutes ago | MASTER | Up | <button>Delete</button> |
| ▶ | cava-n-85-043.eng.vmware.com | 7.5.0.14528 | 14 seconds ago | IAAS | | <button>Delete</button> |

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 using the Create Support Bundle button on the **vRA > Logs** page. 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. Click the **Delete** button to remove a node from the table.

Collect Logs for Clusters and Distributed Deployments

To support troubleshooting and record keeping activities, you can create a zip file that includes all log files for servers in your deployment.

The Distributed Deployment Information table on the Cluster tab of the Virtual Appliance Management Interface lists the nodes for which log files are collected. You can also delete nodes from this table.

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

Procedure

- 1 Log in to the Virtual Appliance Management Interface as a system administrator.
- 2 Click **vRA > Logs**.
- 3 Click **Create Support Bundle**.

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

Remove a Node from the Distributed Deployment Information Table

Delete a node when you want to remove it from your deployment cluster or when you are replacing a Management Agent certificate.

The Distributed Deployment Information table on the Cluster tab of the Virtual Appliance Management Interface lists the nodes for the applicable cluster. You can click the **Delete** button for any node on the table to remove that node from the cluster, or you can use the following procedure.

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 the **Cluster** tab.

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

- 3 Locate the node ID for the node to be deleted by opening a command prompt and running the following command:

```
/usr/sbin/vcac-config cluster-config-node --action list
```

- 4 Locate the node ID, for example `cafe.node.46686239.17144`, in the JSON output.
- 5 Open a command prompt and type a command of the following form, using the node ID that you obtained in the previous step.

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

For example, enter the following command for the example node ID `cafe.node.46686239.17144`:

```
/usr/sbin/vcac-config cluster-config-node --action delete --id cafe.node.46686239.17144
```

- 6 Click **Refresh**.

The node no longer appears in the display.

Monitoring vRealize Automation Health

The vRealize Automation Health Service assesses the functional health of a vRealize Automation environment.

IaaS administrators configure the Health Service to run test suites that determine if the components are registered and the necessary resources are available. This table shows the test suites provided by the Health Service and some example tests in each suite.

| Health Service Test Suites | Example Tests |
|--------------------------------------|--|
| System tests for vRealize Automation | <ul style="list-style-type: none"> ■ SSO/Identity VA Connection Test ■ vRealize Automation License Check - Is License Expired? ■ vRealize Automation Virtual Appliance Root Password Check - Is Password Expiring Soon? |
| Tenant tests for vRealize Automation | <ul style="list-style-type: none"> ■ Check vSphere Reservation Storage Paths ■ Check Reservation Policy to Reservation Assignments ■ Check the Portal Service Status |
| Tests for vRealize Orchestrator | <ul style="list-style-type: none"> ■ Check number of active vRO nodes ■ Check the utilization of the java memory heap in the vRO nodes ■ Check the status of the vro-server service in the vRO nodes |

After you run a test suite on a virtual machine, the Health Service reports the number of tests that passed or failed. For each failed test, the Health Service provides these links:

| Link | Content |
|-------------|--|
| Cause | Explanation of why the test failed. |
| Remediation | Information that you can use to fix the problem. |

You can configure the Health Service to run tests on a schedule or only on demand.

You can also use Python to create custom tests. See the *vRealize Automation Health Service Extensibility Guide*.

Tenant administrators with a Health Consumer role can view test results for their tenancy but cannot configure or run a test.

Configure System Tests for vRealize Automation

An **IaaS administrator** configures the Health Service to run system tests on a selected vRealize Automation virtual appliance. These tests determine if components, such as the vRealize Automation license, are registered and necessary resources, such as memory, are available on the virtual appliance. When you configure the system tests, the Health page displays the tests as a test card.

To configure the Health Service to run system tests for vRealize Automation, complete this procedure.

Prerequisites

Log in to vRealize Automation as an **laaS administrator**.

Procedure

- 1 Select **Administration > Health**.
- 2 Click **New Configuration**.
- 3 On the Configuration Details page, provide the requested information.

| Option | Description |
|-------------|---|
| Name | Your title for this configuration. This title appears on the test card. |
| Description | A description of the test suite. |
| Product | Select vRealize Automation. |
| Schedule | Select how often the test suite runs. |

- 4 Click **Next**.
- 5 On the Select Test Suites page, select **System Tests for vRealize Automation**.
- 6 Click **Next**.
- 7 On the Configure Parameters page, provide the requested information.

Table 1-18. vRealize Automation Virtual Appliance

| Option | Description |
|---------------------------|--|
| Public Web Server Address | <ul style="list-style-type: none"> ■ For a minimal deployment, the base URL for the vRealize Automation appliance host. For example, <code>https://va-host.domain/</code>. ■ For a high-availability deployment, the base URL for the vRealize Automation load balancer. For example, <code>https://load-balancer-host.domain/</code>. |
| SSH Console Address | Fully qualified domain name of the vRealize Automation appliance. For example, <code>va-host.domain</code> . |
| SSH Console User | root |
| SSH Console Password | The root password. |

Table 1-19. vRealize Automation System Tenant

| Option | Description |
|-----------------------------|-----------------------------|
| System Tenant Administrator | administrator |
| System Tenant Password | The administrator password. |

Table 1-20. vRealize Automation Disk Space Monitoring

| Option | Description |
|----------------------------|---|
| Warning Threshold Percent | Acceptable percent of virtual appliance disk space that is used before the warning test fails. |
| Critical Threshold Percent | Acceptable percent of virtual appliance disk space that is used before the critical test fails. |

- 8 Click **Next**.
- 9 On the Summary page, review the information.
- 10 Click **Finish**.

Tests run according to the selected schedule.

What to do next

[View the vRealize Automation Health Service Test Suite Results](#)

Configure Tenant Tests For vRealize Automation

An **laaS administrator** configures the Health Service to run tenant tests on a selected vRealize Automation virtual appliance. These tests determine if tenant-related components, such as software-service, are registered and necessary resources, such as vSphere virtual machines, are available on the virtual appliance. When you configure the tenant tests, the Health page displays the tests as a test card.

To configure the Health Service to run tenant tests for vRealize Automation, complete this procedure.

Prerequisites

Log in to vRealize Automation as an **laaS administrator**.

Procedure

- 1 Select **Administration > Health**.
- 2 Click **New Configuration**.

- 3 On the Configuration Details page, provide the requested information.

| Option | Description |
|-------------|---|
| Name | Your title for this configuration. This title appears on the test card. |
| Description | A description of the tests. |
| Product | Select vRealize Automation. |
| Schedule | Select how often these tests run. |

- 4 Click **Next**.

- 5 On the Select Test Suites page, select **Tenant Tests for vRealize Automation**.

- 6 Click **Next**.

- 7 On the Configure Parameters page, provide the requested information.

Table 1-21. vRealize Automation Virtual Appliance

| Option | Description |
|---------------------------------|--|
| vRealize Automation Web Address | <ul style="list-style-type: none"> ■ For a minimal deployment, the base URL for the vRealize Automation appliance host. For example, <code>https://va-host.domain/</code>. ■ For a high-availability deployment, the base URL for the vRealize Automation load balancer. For example, <code>https://load-balancer-host.domain/</code>. |
| SSH Console Address | Fully qualified domain name of the SSH host. For example, <code>ssh-host.domain</code> . |
| SSH Console User | root |
| SSH Console Password | Password for root. |
| Max Service Response time (ms) | Maximum amount of time in milliseconds the system waits for a response. |

Table 1-22. vRealize Automation Tenant

| Option | Description |
|-------------------------------|--|
| Tenant Under Test | qe |
| Fabric Administrator Username | <p>Fabric administrator user name.</p> <p>Note This fabric administrator must also have a tenant administrator and an IaaS administrator role in order for all of the tests to run.</p> |
| Fabric Administrator Password | Password for fabric administrator. |

Table 1-23. vRealize Automation System Tenant

| Option | Description |
|-----------------------------|-----------------------------|
| System Tenant Administrator | administrator |
| System Tenant Password | Password for administrator. |

Table 1-24. vRealize Automation Disk Space Monitoring

| Option | Description |
|----------------------------|---|
| Critical Threshold Percent | Acceptable percent of virtual appliance disk space that is used before the critical test fails. |

- 8 Click **Next**.
- 9 On the Summary page, review the information.
- 10 Click **Finish**.

Tests run according to the selected schedule.

What to do next

[View the vRealize Automation Health Service Test Suite Results](#)

Configure Tests For vRealize Orchestrator

An **IaaS administrator** configures the health service to run tests for vRealize Orchestrator on the vRealize Orchestrator host. These tests confirm that components, such as the vro-server service, are registered and necessary resources, such as sufficient Java memory heap, are available are available on the host machine. When you configure the vRealize Orchestrator tests, the Health page displays the tests as a test card.

Prerequisites

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

Procedure

- 1 Select **Administration > Health**.
- 2 Click **New Configuration**.
- 3 On the Configuration Details page, provide the requested information.

| Option | Description |
|-------------|---|
| Name | Your title for this configuration. This title appears on the test card. |
| Description | A description of the tests. |

| Option | Description |
|----------|---------------------------------|
| Product | Select vRealize Orchestrator. |
| Schedule | Select how often the tests run. |

- 4 Click **Next**.
- 5 On the Select Test Suites page, select **Tests for vRealize Orchestrator**.
- 6 Click **Next**.
- 7 On the Configure Parameters page, provide the requested information.

Table 1-25. vRealize Orchestrator Host/Load Balancer

| Option | Description |
|----------------------------|---|
| Client Address | <ul style="list-style-type: none"> ■ For a minimal deployment, the fully qualified domain name of the vRealize Orchestrator host. For example, <i>vro-host.domain</i>. ■ For a high-availability deployment, the base URL for the vRealize Orchestrator load balancer, <i>https://load-balancer-host.domain/</i>. |
| Client Username | administrator |
| Client Password | The administrator password. |
| SSH Console Username | root |
| SSH Console Password | The root password. |
| Heap Utilization Threshold | Acceptable percent of heap space that is used before the warning test fails. |

Table 1-26. vRealize Orchestrator Instances Behind Load Balancer

| Option | Description |
|----------------------|---|
| SSH Console Address | IP address or URL of the vRealize Orchestrator instance behind the load balancer. |
| SSH Console Username | User name with access to this instance. |
| SSH Console Password | The user name password. |

- Click **ADD** to add another vRealize Orchestrator instance to the list.
- Click **REMOVE** to remove a selected vRealize Orchestrator instance from the list of instances behind the load balancer.

- 8 Click **Next**.
- 9 On the Summary page, review the information.

10 Click **Finish**.

Tests run according to the selected schedule.

What to do next

[View the vRealize Automation Health Service Test Suite Results](#)

Custom Test Suite

You can use Python to create a custom test suite for vRealize Automation Health Service.

Creating a custom test suite allows you to extend the tests supplied for the health service by adding a test suite to determine the health of additional vRealize Automation components. For information about creating a custom test suite, see the *vRealize Automation Health Service Extensibility Guide*.

Add a Custom Test Suite

An **laaS administrator** must add a custom test suite to vRealize Automation health service before you run the test suite.

To add a custom test suite for a vRealize Automation asset, complete this procedure.

Prerequisites

- Create a Python wheel for the custom test suite files. For information, see the *vRealize Automation Health Service Extensibility Guide*.
- Log in to vRealize Automation as an **laaS administrator**.

Procedure

- 1** Click **Administration > Health**.
- 2** In the upper right, click the gear icon and select **Extensibility**.
- 3** Click **New Asset**.
- 4** In the Add Asset dialog box, provide the requested information.

| Option | Description |
|-------------------|---|
| Asset Title | The name and version number of the test suite you are running, for example, Infoblox 1.0. |
| Asset Description | A description of the tests contained in the Python wheel. |
| Asset Version | Test suite version number. |
| Asset File | Click Choose File and select your custom test suite file. |

5 Click **Add**.

A new row is added to the asset table with the status **UPLOADED**. When the status changes to **INSTALLED**, your test suite is ready to use. If the install process fails, you see a popup that provides a reason.

Note If the page does not update, click the refresh icon.

What to do next

[Run a Custom Test Suite](#).

Run a Custom Test Suite

An **laaS administrator** configures the health service to run a custom test suite in the vRealize Automation environment. When you configure the custom test suite, the Health page displays the test suite as a test card.

To configure the health service to run a custom test suite for vRealize Automation, complete this procedure.

Prerequisites

- [Add a Custom Test Suite](#).
- Log in to vRealize Automation as an **laaS administrator**.

Procedure

- 1 Select **Administration > Health**.
- 2 Click **New Configuration**.
- 3 On the Configuration Details page, provide the requested information.

| Option | Description |
|-------------|---|
| Name | Your title for this configuration. This title appears on the test card. |
| Description | A description of the test suite. |
| Product | Select the product you want to test from the Product drop-down menu. |
| Schedule | Select how often you want to run this test suite. |

- 4 Click **Next**.
- 5 On the Select Test Suites page, select the custom test suite, and click **Next**.
- 6 On the Configure Parameters page, enter the requested information, and click **Next**.
- 7 On the Summary page, review the information, and click **Finish**.

The custom test suite runs according to the selected schedule.

What to do next

[View the vRealize Automation Health Service Test Suite Results](#)

View the vRealize Automation Health Service Test Suite Results

You can view the health service test results after you run the tests.

The Health page displays each configured test suite as a test card. When a test suite runs, the result appears in the middle of the test card.

The test cards that you see on the Health page are filtered according to your privilege.

- IaaS administrators can see all test cards.
- Tenant administrators with the Health Consumer role can see only the test card for their tenancy.

Prerequisites

- The configured test suite has run on schedule.
- Log in to the vRealize Automation console as an **IaaS administrator** or as a **tenant administrator**.

Procedure

- 1 Select **Administration > Health**.
- 2 If a test is not scheduled to run, click **Run** on the test card.
- 3 Click the center of a test card after the tests are finished.

A page appears that displays the status of each test. To see why a test failed, click **Cause**. To open a topic that explains how to fix the problem, click the **Remediation** link if one is available.

Troubleshooting the Health Service

The Health Service troubleshooting topics provide solutions to problems you might experience when you use the Health Service.

Service Status Test Fails

You can fix a failed service test by changing the test schedule setting.

Problem

If a service status test fails and you click **Cause**, you see this message: Cannot establish SSH connection ; Exception message:[Auth fail].

Cause

When the test suite is scheduled to run every 15 minutes, the system login locks the root user account.

Solution

- ◆ Change the test schedule to **None**, wait 15 minutes, and run the test suite again.

After Upgrade the Health Page in the Appliance Console Is Empty

After you upgrade vRealize Automation, the Health page in the appliance console is empty.

Problem

The health service does not start after upgrade.

Solution

- ◆ On each vRealize Automation virtual appliance, open a command prompt as **root** and run these commands.

- To configure the health service to start automatically, run this command.

```
chkconfig vrhb-service on
```

- To start the health service on this virtual appliance, run this command.

```
service vrhb-service start
```

Monitoring vRealize Automation Environment Resources Using SNMP

As a system administrator familiar with SNMP, you want to use the vRealize Automation vRealize Automation REST API for vSNMP to facilitate how you monitor your vRealize Automation nodes. Using vSNMP, you can use SNMP to act as an encrypted early warning system when vRealize Automation is about to run out of CPU, RAM, disk space so that you avoid slowdowns.

You can manually monitor the SNMP OIDs, or you can actively monitor resources by setting SNMP traps.

For example, if vSNMP sends you an event, such as "High CPU usage detected," you might start gathering information about the processes consuming CPU and determine which one is using excessive resources. You might then correlate the CPU, memory, and other usage to troubleshoot additional problems.

Using the vRealize Automation vSNMP, you can expose the entire Linux tree for monitoring and retrieving data using the REST API, or by using the vSNMPD daemon that is running on your vRealize Automation instances.

vRealize Automation SNMP does not have a general use interface. You must use the REST API or the daemon commands.

For more information, see "Using SNMP to Monitor vRealize Automation" in the vRealize Automation Programming Guide. To locate the Programming Guide, see [vRealize Automation API Documentation](#), and select the version link.

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 1-27. 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 |

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 1-28. 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 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.

Results

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 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.

Results

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 deactivate 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 deactivates the storage path on all reservations that use it. See [Deactivate 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](#).

Deactivate a Storage Path



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

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

Prerequisites

Log in to vRealize Automation 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 deactivate 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 1-29. Data Collection Types

| Data Collection Type | Description |
|---|---|
| Infrastructure Source Endpoint Data Collection | <p>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.</p> <p>Endpoint data collection runs every 4 hours.</p> |
| Inventory Data Collection | <p>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.</p> <p>Inventory data collection runs every 24 hours.</p> <p>The default timeout interval for inventory data collection is 2 hours.</p> |
| State Data Collection | <p>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.</p> <p>State data collection runs every 15 minutes.</p> <p>The default timeout interval for state data collection is 1 hour.</p> |
| Performance Data Collection (vSphere compute resources only) | <p>Updates the record of the average CPU, storage, memory, and network usage for each virtual machine discovered through inventory data collection.</p> <p>Performance data collection runs every 24 hours.</p> <p>The default timeout interval for performance data collection is 2 hours.</p> |
| Network and security inventory data collection (vSphere compute resources only) | <p>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.</p> |
| WMI data collection (Windows compute resources only) | <p>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.</p> |

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 vRealize Automation as an **laaS administrator**.

Procedure

- 1 Select **Infrastructure > Endpoints > Endpoints**.
- 2 Click in the row of the endpoint that you want to data collect.
- 3 Select an available data collection action.

Configure Compute Resource Data Collection

You can activate or deactivate 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 vRealize Automation 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 activate data collection.
 - Select **Off** to deactivate data collection.
- 4 Configure **Inventory** data collection.
 - Select **On** to activate data collection.
 - Select **Off** to deactivate 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 activate data collection.
 - Select **Off** to deactivate 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 activate data collection.
- Select **Off** to deactivate 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 **Snapshot Inventory** data collection.

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

- Select **On** to activate data collection.
- Select **Off** to deactivate 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.

8 Click **OK**.

Update Cost Data for All Compute Resources

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

Prerequisites

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

Procedure

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

Results

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.

Monitoring Containers

You can monitor the status of a container that you create in Containers for vRealize Automation.

After you create your containers based on a template, you can monitor their state. By clicking **Details** on a container, you can monitor the network bandwidth, CPU and memory usage, logs, and properties of that container.

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.

Bulk Imports creates a CSV file that contains defining virtual machine data such as reservation, storage path, blueprint, owner, and any custom properties. You use the CSV file to import virtual machines to your vRealize Automation environment. 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 vRealize Automation environment to another.

Note Only vCloud Director and vSphere are supported for bulk import. Setting the filter to another endpoint type does not generate data in the CSV file.

You can run 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.

Note Bulk machine importing does not bypass normal provisioning steps. Any existing external workflows that are triggered by the Event Broker during provisioning are run for imported machines. You can temporarily deactivate workflows for imported machines by performing one of the following:

- Deactivate all Event Broker subscriptions. If you are deactivating subscriptions, you must schedule a service outage for your vRealize Automation cluster because extensibility will not be applied to any normal machine provisioned during this time.
 - Add a condition to event subscriptions to not trigger when a machine is imported. To add this condition, navigate to Event Subscriptions, select the subscription to deactivate, and add a custom property `VirtualMachine.Imported.ConvergedBlueprint` does not equal `<Id of the import blueprint>`. This condition does not effect normally provisioned machines and instead is only applied to imported machines.
-

Prerequisites

- Log in to vRealize Automation 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 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 **Deployments** tab.

The bulk import option does not support deployments that are provisioned from a blueprint that contains an NSX network and security component or a software component.

Prerequisites

- Log in to vRealize Automation 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 about using a network profile to control IP address ranges, see *Configuring vRealize Automation*.
- If you use bulk import to import a virtual machine with a static IP address that is allocated to another virtual machine, the import fails.

Procedure

- 1 Temporarily deactivate all Event Broker Subscriptions.

Note When disabling subscriptions you must schedule a service outage for your vRealize Automation cluster. During this process, extensibility is not applied to any normally provisioned machine. Failure to deactivate subscriptions can result in data loss and permanent deletion of machines from the backing infrastructure.

- 2 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**.

3 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 Enter only the blueprint ID, not 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. |
| | For imported virtual machines, do not associate a blueprint that includes component profiles. Existing settings in imported virtual machines, such as memory or storage size, might be outside of profile limits. When that happens, validation for any future blueprint-based reconfiguration of the virtual machines fails. |

| Heading | Comment |
|----------------------|---|
| 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. |

If you import a virtual machine with one or more custom properties, you identify each custom property using three comma separated values appended to the line with the values for that machine. Use this format for each custom property.

,Custom.Property.Name, Value, FLAGS

FLAGS are three characters that describe how the property is treated by vRealize Automation. In their order of use, the flags are:

- 1 H or N = Hidden or Not Hidden
- 2 E or O = Encrypted or Not Encrypted
- 3 R or P = Runtime or Not Runtime

For example, you can append a custom property to configure a static IP address for a machine. Using the following format, this custom property allocates an available static IP address from a network profile.

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

You change the variables with the appropriate information for your virtual machine.

- Change # to the number of the network interface being configured with this static IP address. For example, `VirtualMachine.Network0.Address`.
- Change `w.x.y.z` to be the static IP address for the virtual machine. For example, `11.27.42.57`.

The HOP flag string—Hidden, Not encrypted, Not Runtime—sets the visibility of the property. Because this particular property is used only by bulk import, it is removed from the virtual machine after a successful import.

In order for this custom property to work, 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.

- b Save the CSV file.
- 4 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 vRealize Automation 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 vRealize Automation 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 about using a network profile to control IP address ranges, see *Configuring vRealize Automation*.

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 |

| Heading | Comment | Example |
|---------------------------|---|-----------------|
| 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.