Installation and Configuration

vRealize Code Stream 1.2

This document supports the version of each product listed and supports all subsequent versions until the document is replaced by a new edition. To check for more recent editions of this document, see http://www.vmware.com/support/pubs.

EN-001754-00

vmware®

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

http://www.vmware.com/support/

The VMware Web site also provides the latest product updates.

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

docfeedback@vmware.com

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

Contents

Installation and Configuration

The *Installation and Configuration* guide provides information about how to install and configure VMware vRealize Code Stream to automate the release of applications.

Intended Audience

This information is intended for anyone who wants to install vRealize Automation, and configure the environment to automate the release applications in development environments. The information is written for experienced developers and operation teams who are familiar with release automation.

VMware Technical Publications Glossary

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

Introducing vRealize Code Stream 1

vRealize Code Stream automates the software release process by modeling all of the necessary tasks in pipeline templates.

A release pipeline is a sequence of stages. Each stage is composed of multiple tasks and environments that the software has to complete before it is released to production. The stages can include development, functional testing, user acceptance test (UAT), load testing (LT), systems integration testing (SIT), staging, and production. Release managers, typically build and release engineers, model pipeline templates.

Each stage in a pipeline includes a set of activities such as provisioning a machine, retrieving an artifact, deploying software, running a test, creating a manual task, or running a custom workflow or script. The software changes are promoted to the next stage in the pipeline when they satisfy a set of rules called gating rules. The gating rules include testing rules and compliance rules. Gating rules that are associated with a pipeline are specific to an organization or an application. Users can define gating rules when a pipeline template is created. The environments do not need to be aware of these gating rules.

 Introduction to vRealize Code Stream (http://bcove.me/kw7b6zvs)

Figure 1-1. Main Components of vRealize Code Stream

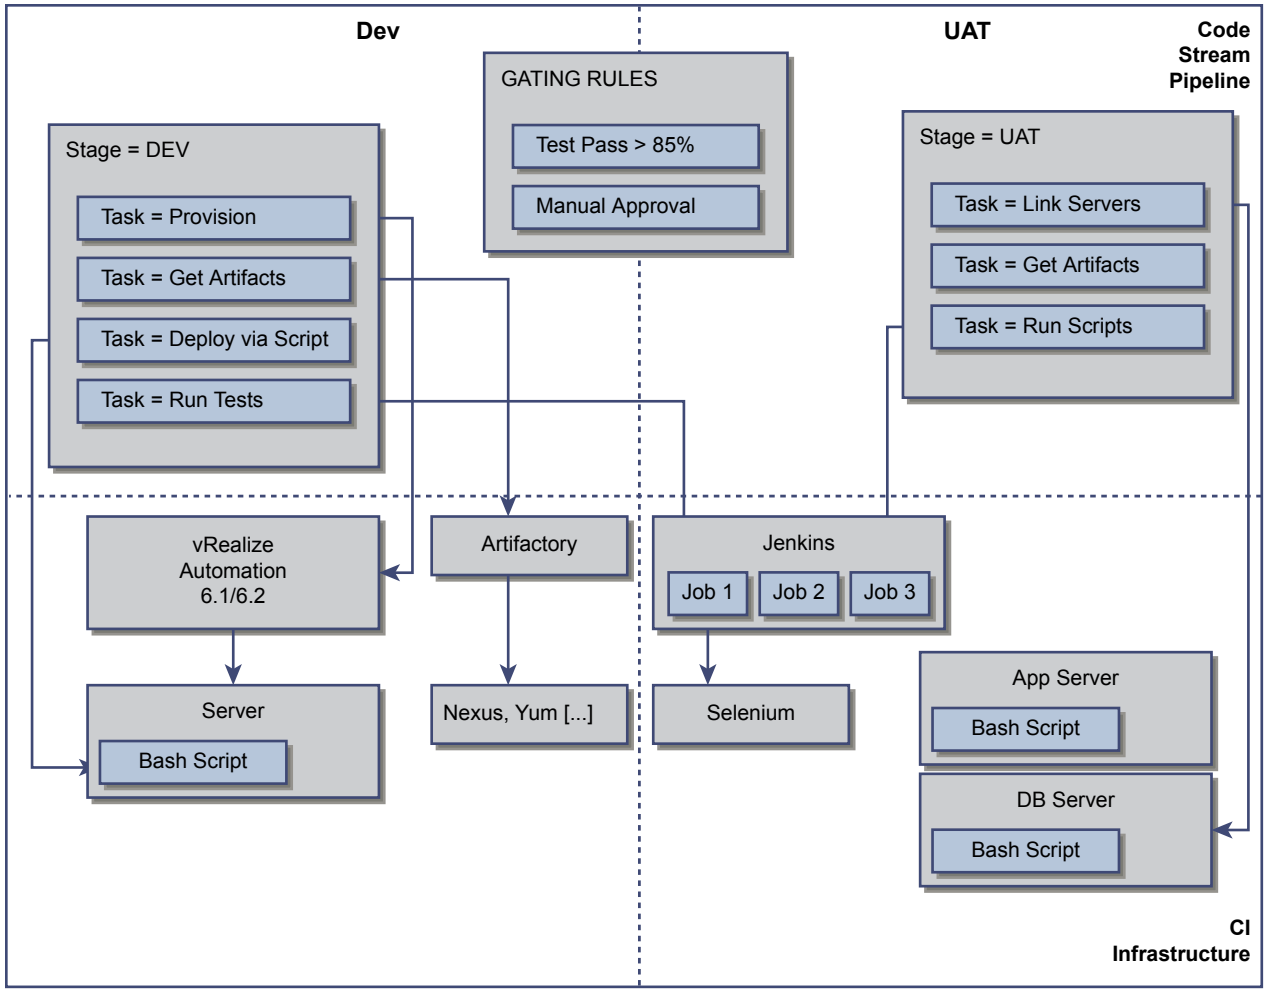

This chapter includes the following topics:

- "Core Architectural Principles," on page 8
- "Roles and Responsibilities of Personas," on page 9
- "Integrating vRealize Code Stream with External Systems," on page 10
- "Key Release Automation Concepts," on page 11

Core Architectural Principles

vRealize Code Stream works with deployment engines such as vRealize Automation (formerly VMware vCloud Automation Center), vCenter Server, Chef, Puppet, or scripts. vRealize Automation can also integrate with continuous integration frameworks and testing frameworks.

vRealize Code Stream includes an approval engine that can integrate with IT service management products and various commercial or custom approval systems. Based on the type of integration, vRealize Code Stream uses the JFrog Artifactory or VMware vCloud Orchestrator (vRO) (formerly vCO) plug-ins for extensibility. Both of the approval and extensibility components are embedded in the vRealize Automation virtual appliance.

For the supported vRealize Code Stream integrations, see "Integrating vRealize Code Stream with External Systems," on page 10.

Deployment Engines

vRealize Code Stream integrates with a number of provisioning and deployment solutions including vRealize Automation. It can also trigger scripts or vRealize Orchestrator workflows. Support for other provisioning solutions is delivered by plug-ins that VMware, partners, or users publish.

Testing Frameworks

vRealize Code Stream integrates with Jenkins to trigger Jenkins jobs, including test routines. The user interface for test automation and for extensibility uses the vRealize Orchestrator plug-in framework.

A Jenkins job can run test cases that are configured for an application. The Test Acceptance Threshold workflow in the gating rule verifies the results of the Jenkins job and returns the response to the vRealize Code Stream server. Based on the results of the test and the gating rules that have been defined, the build either proceeds to the next stage of the release pipeline or it fails.

Approval Systems

vRealize Code Stream uses vRealize Orchestrator plug-ins for integration with approval systems. Manual approval tasks can be created within the vRealize Automation inbox, but vRealize Code Stream can also integrate with BMC Remedy ITSM, HP Service Manager, ServiceNow, and other ticketing systems. The approval systems integration requires downloading and installing the appropriate vRealize Orchestrator plug-in from the VMware Solution Exchange.

Roles and Responsibilities of Personas

A tenant administrator can assign the release manager, release engineer, and the release dashboard user roles, which are an integral part of release automation.

These roles have various responsibilities when they interact with the product. See "Configuring Additional Tenants," on page 27.

Each stage in the release pipeline defines a set of activities such as initialize, deploy, test, approval, custom tasks, and troubleshoot. The following table lists the roles and responsibilities of the personas.

Table 1-1. Roles and Responsibilities in vRealize Code Stream

Role	Responsibility
Release Manager	■ Define and create a pipeline template ■ Define different stages in a pipeline template ■ Define access permissions for release engineers ■ Define gating rules ■ Register repositories ■ Register CI and test frameworks ■ Approve promotion to different stages
Release Engineer	■ Retrieve artifacts ■ Deploy artifacts to an environment ■ Trigger a pipeline ■ Monitor pipeline execution ■ View test results ■ Trigger tests ■ Review test results
Release Dashboard User	■ View the artifact, machine, and pipeline data on the dashboard ■ Troubleshoot a failed pipeline

Integrating vRealize Code Stream with External Systems

vRealize Code Stream includes an extensibility framework that supports modular integrations with external systems, without changing the core platform.

Based on the type of external system, different mechanisms are recommended.

Figure 1-2. Supported Integration with External Systems

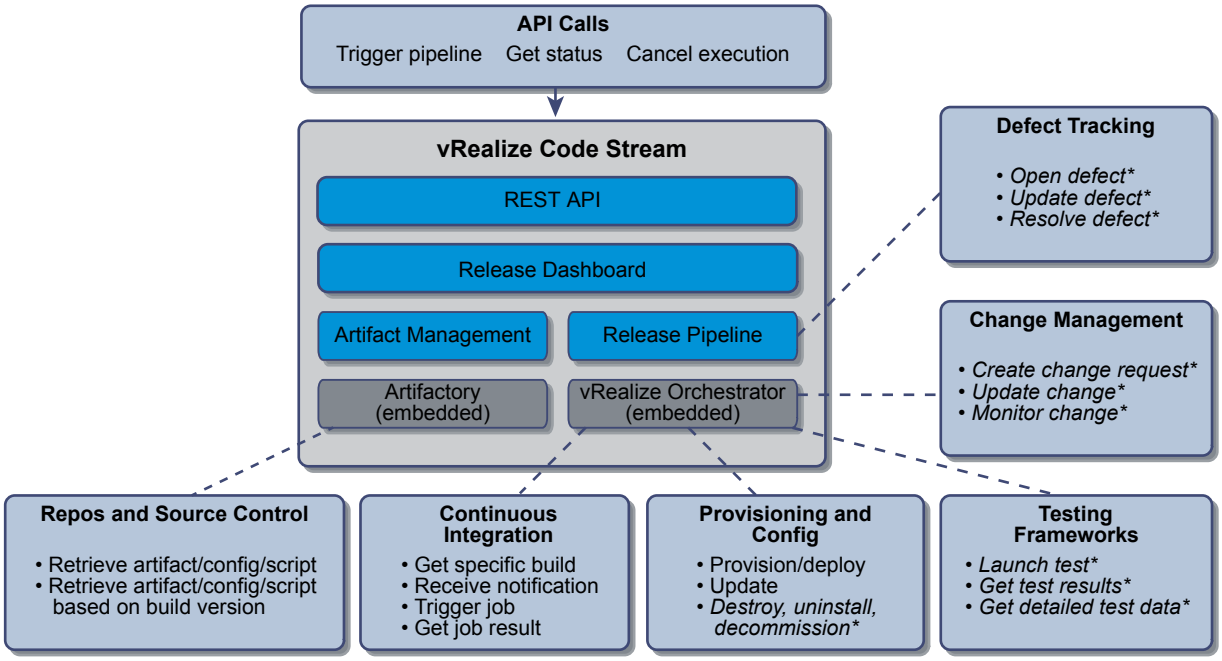

*Items in italics require custom workflows or scripts

Release Pipeline Integrations

Release pipeline templates support various tasks that can trigger actions in a wide category of systems such as continuous integration, testing frameworks, or defect tracking systems.

You can download vRealize Orchestrator plug-ins from the VMware Solution Exchange. The Artifactory plug-ins are available on the JFrog Web site.

Some integrations are supported natively. Some integrations require downloading an Artifactory or vRealize Orchestrator plug-in. Others require creating custom workflows by using a vRealize Orchestrator plug-in protocol such as HTTP-REST, SOAP, or SSH.

Table 1-2. Supported Integrations

System Category	Integration Mechanism	Native Support	Requires Separate Plug-in	Requires Custom Workflows
Repository	JFrog Artifactory	■ Sonatype Nexus ■ Yum ■ HTTP-browsable repository system	N/A	
Continuous Integration	vRealize Orchestrator	Jenkins	N/A	■ Atlassian Bamboo using the REST plug-in ■ JetBrains TeamCity using the REST plug-in ■ Microsoft Team Foundation Server using the REST plug-in
Provisioning and configuration management	vRealize Orchestrator	■ vRealize Automation (single-machine blueprints) ■ Scripts (BASH or Windows PowerShell)	■ vCenter Server ■ VMware vCloud Director	■ Chef using the REST plug-in ■ Puppet using the REST plug-in
Testing frameworks	vRealize Orchestrator	Any testing frameworks exposed as Jenkins jobs	N/A	Any testing framework that offers a REST or SOAP API
Change management systems	vRealize Orchestrator	N/A	■ BMC Remedy ITSM ■ HP Service Manager ■ ServiceNow	Any commercial or custom change management system that offers a REST or SOAP API
Defect tracking systems	vRealize Orchestrator	N/A	N/A	■ Atlassian JIRA using the REST plug-in ■ Bugzilla using the REST plug-in

Key Release Automation Concepts

Use the following definitions to help you understand the release pipeline modeling and the artifact management workflow.

artifact

A script or the output of a build process. The script can be deployed or upgraded in a given stage.

Artifact types can be configuration files, application bits, or third-party software.

artifact management

A service that manages the artifacts over a range of local and remote repositories.

For example, managing a WAR file stored in the Maven repository.

category	A task type. Some supported categories are Provision, Custom, Artifact, Deploy, and Test. A task belongs to a provider and a category.
gating rule	A set of rules that must be completed before the software changes are promoted and the next set of tasks starts in the subsequent stage. The gating rules include testing rules and compliance rules. Gating rules that are associated with a pipeline are specific to an organization and applications.
instance	A vRealize Orchestrator plug-in scenario that captures specific configurations of a provider. The instance is created by using a vRealize Orchestrator client.
pipeline	A collection of all the stages or environments in which a software change has to pass through independently before it is released into production. For example, development, test, user acceptance test, load test, staging, and production.
provider	Plug-in vendors that support the categories. For example, the Provision category is supported by vRealize Automation and vCenter Server providers.
stage	Every stage in the pipeline defines a set of activities. For example, deploy, test, approval through gating rules, and custom tasks.
task	An activity in a given stage. For example, provision the machines, resolve the artifact, deploy the artifact, run the test, and so on. Opening the port in a firewall is a manual task.

vRealize Code Stream Installation 2

vRealize Code Stream shares a platform and common services with vRealize Automation 6.2.

To install vRealize Code Stream, you must set up, configure, and deploy a vRealize Automation appliance. You also need to configure a tenant to assign user roles in vRealize Code Stream.

During installation, you can apply the vRealize Code Stream license to enable the features in the vRealize Automation appliance. To use the vRealize Code Stream function, a system administrator can deploy or upgrade a vRealize Automation 6.2 appliance and apply a vRealize Code Stream license key.

This chapter includes the following topics:

Using vRealize Code Stream Installation Checklist

The installation checklist provides a high-level overview of the sequence of tasks you must perform to complete the vRealize Code Stream installation.

Installation Checklist

Use the checklist to track your work as you complete the installation tasks in the order they are listed.

Table 2-1. Installation Tasks

Task	Details
Complete the installation worksheet	"Installation Worksheets," on page 14
Set up your Identity Appliance	"Deploy and Configure the Identity Appliance," on page 16
Configure Single-Sign On for the Identity Appliance	"Configure the Identity Appliance," on page 18
Deploy the vRealize Automation appliance	"Deploy the vRealize Appliance," on page 20
Set up the vRealize Automation appliance	"Configure the vRealize Appliance," on page 22 Complete tasks up to step 12.

Table 2-1. Installation Tasks (Continued)

Task	Details
Add the vRealize Code Stream license key	"Apply a vRealize Code Stream License to an Appliance," on page 25
Set up the Artifactory Server password	"Set Up the Artifactory Server Password," on page 25
Configure Tenant	"Configuring Additional Tenants," on page 27
Add vRealize Code Stream roles	"Assign Roles to Identity Store Users or Groups," on page 35 Assign the Release Manager, Release Engineer, and Release Dashboard user roles for modeling and publishing a pipeline.
Register the Default Artifactory Server	"Register an Artifactory Server for Artifact Management," on page 41

Installation Worksheets

You can use these worksheets to record important information for reference during the installation process.

One copy of each worksheet is given here. Create additional copies as you need them. Settings are case sensitive.

Table 2-2. Identity Appliance Information

Variable	Value	Example
Host Name (FQDN)		vcac-sso.mycompany.com
SSO service over HTTPS Incoming Port	7444 (do not change)	7444
IP		192.168.1.104
Username	administrator@vsphere.local (default)	administrator@vsphere.local
Password		vmware

Table 2-3. Leading cluster vRealize Appliance Information

Variable	Value	Example
Host Name (FQDN)		vcac-va.mycompany.com
SSO service over HTTPS Outgoing Port (default)	7444 (do not change)	7444
IP		192.168.1.105
Username	administrator@vsphere.local (default)	administrator@vsphere.local
Password		vmware

Table 2-4. Additional vRealize Appliance Information

Variable	Value	Example
Host Name (FQDN)		vcac-va2.mycompany.com
SSO service over HTTPS Outgoing Port (default)	7444 (do not change)	7444
IP		192.168.1.110
Username	administrator@vsphere.local (default)	administrator@vsphere.local
Password		vmware

Table 2-5. IaaS Database Passphrase

Variable	Value	Example
Passphrase (reused in IaaS Installer, Upgrade, and Migration)		myPassphrase

Table 2-6. IaaS Website

Variable	Value	Example
Host Name (FQDN)		iaas-web.mycompany.com
SSO service over HTTPS Outgoing Port (default)		
IP		192.168.1.106
Username		
Password		

Table 2-7. IaaS Model Manager Data

Variable	Value	Example
Host Name (FQDN)		iaas-model-man.mycompany.com
SSO service over HTTPS Outgoing Port (default)		
IP		192.168.1.107
Username		
Password		

Table 2-8. IaaS Model Service

Variable	Value	Example
Host Name (FQDN)		iaas-model-service.mycompany.com
SSO service over HTTPS Outgoing Port (default)		
IP		192.168.1.108
Username		
Password		

Table 2-9. Distributed Execution Managers

Unique Name	Orchestrator/Worker
ex. myuniqueorchestratorname	Orchestrator: Worker:
	Orchestrator: Worker:
	Orchestrator: Worker:
	Orchestrator: Worker:

Deploy and Configure the Identity Appliance

Download and configure the Identity Appliance to provide Single Sign-On (SSO) capability for the vRealize Automation environment.

You can use the Identity Appliance SSO provided with vRealize Automation or some versions of the SSO provided with vSphere. For information about supported versions, see *vRealize Automation Support Matrix* on the VMware Web site.

NOTE PSC version 6.0, the vSphere SSO component introduced in vSphere 6.0, allows you to specify a tenant name other than vsphere.local. vRealize Automation requires vsphere.local as the name of the default tenant because you cannot enter the name of the tenant on the SSO tab of the management console when you configure vRealize Automation. If you have used another name, rename the tenant to vsphere.local.

1 Deploy the Identity Appliance on page 16

The Identity Appliance is a preconfigured virtual appliance that provides single sign-on capabilities. You download the Identity Appliance and deploy it into vCenter Server or ESX/ESXi inventory.

2 Enable Time Synchronization on the Identity Appliance on page 17

You must synchronize the clocks on the Identity Appliance server, the vRealize Automation server, and Windows servers to ensure a successful installation.

3 Configure the Identity Appliance on page 18

The Identity Appliance provides Single-Sign On (SSO) capability for vRealize Automation users. SSO is an authentication broker and security token exchange that interacts with the enterprise identity store (Active Directory or OpenLDAP) to authenticate users. A system administrator configures SSO settings to provide access to the vRealize Automation.

Deploy the Identity Appliance

The Identity Appliance is a preconfigured virtual appliance that provides single sign-on capabilities. You download the Identity Appliance and deploy it into vCenter Server or ESX/ESXi inventory.

Exact steps for this procedure vary depending on whether you use the native or Web vSphere client. Also, specific steps can vary depending on the your data center configuration. If you are using VSphere Single-Sign (SSO), you can skip to "Configure the Identity Appliance," on page 18.

Prerequisites

- Download the Identity Appliance from the VMware Web site.
- Log in to the vSphere client as a user with **system administrator** privileges.

Procedure

1 In the vSphere client, select **File > Deploy OVF Template**.

2 Browse to the Identity Appliance file with the `.ova` or `.ovf` extension and click **Open**.

3 Click **Next**.

4 Click **Next** on the OVF Template Details page.

5 Accept the license agreement and click **Next**.

6 Type a unique virtual appliance name according to the IT naming convention of your organization in the **Name** text box, select the datacenter and location to which you want to deploy the virtual appliance, and click **Next**.

7 Follow the prompts until the Disk Format page appears.

8 Verify on the Disk Format page that enough space exists to deploy the virtual appliance and click **Next**.

9 Follow the prompts to the Properties page.

The options that appear depend on your vSphere configuration.

10 Click **Next**.

11 Restart the host machine.

- If **Power on after deployment** is available on the Ready to Complete page.
  a Select **Power on after deployment** and click **Finish**.
  b Click **Close** after the file finishes deploying into vCenter.
  c Wait for the machine to restart. This could take up to five minutes.
- If **Power on after deployment** is not available on the Ready to Complete page.
  a Click **Close**.
  b Restart the machine. This could take up to five minutes.

After a few moments, a success message appears.

12 Verify that the fully qualified domain name can be resolved against the IP address of the Identity Appliance by opening a command prompt and pinging the FQDN.

Enable Time Synchronization on the Identity Appliance

You must synchronize the clocks on the Identity Appliance server, the vRealize Automation server, and Windows servers to ensure a successful installation.

If you see certificate warnings during this procedure, continue past them.

Prerequisites

"Deploy the Identity Appliance," on page 16.

Procedure

1 Navigate to the Identity Appliance management console by using its fully qualified domain name, https://*identity-hostname.domain.name*:5480/.

2 Log in by using the user name root and the password you specified when you deployed the Identity Appliance.

3 Select **Admin > Time Settings**.

4 Select an option from the **Time Sync Mode** menu.

Option	Action
Use Time Server	Select **Use Time Server** from the **Time Sync Mode** menu to use Network Time Protocol . For each time server that you are using, type the IP address or the host name in the **Time Server** text box.
Use Host Time	Select **Use Host Time** from the **Time Sync Mode** menu to use VMware Tools time synchronization. You must configure the connections to Network Time Protocol servers before you can use VMware Tools time synchronization.

5 Click **Save Settings**.

6 Click **Refresh**.

7 Verify that the value in **Current Time** is correct.

You can change the time zone as required from the Time Zone Setting page on the **System** tab.

Configure the Identity Appliance

The Identity Appliance provides Single-Sign On (SSO) capability for vRealize Automation users. SSO is an authentication broker and security token exchange that interacts with the enterprise identity store (Active Directory or OpenLDAP) to authenticate users. A system administrator configures SSO settings to provide access to the vRealize Automation.

NOTE If you plan to use the vRealize Automation migration tool, you must specify a Native Active Directory when you configure the appliance.

Native Active Directories have the following characteristics:

- Use Kerberos to authenticate
- Do not require a search base, making it easier to find the correct Active Directory store
- Can be used only with the default tenant

Prerequisites

"Enable Time Synchronization on the Identity Appliance," on page 17.

Procedure

1 Navigate to the Identity Appliance management console by using its fully qualified domain name, https://*identity-hostname.domain.name*:5480/.

2 Continue past the certificate warning.

3 Log in with the user name root and the password you specified when the appliance was deployed.

4 Click the **SSO** tab.

The red text is a prompt, not an error message.

5 Type the password to assign to the system administrator in the **Admin Password** and **Repeat password** text boxes.

The **System Domain** text field has the value vsphere.local, which is the local default domain for the Identity Appliance. The default tenant is created with this name and the system administrator is administrator@vsphere.local. Record the user name and password in a secure place for later use.

6 Click **Apply**.

It can take several minutes for the success message to appear. Do not interrupt the process.

7 When the success message appears, click the **Host Settings** tab.

8 Verify that the **SSO Hostname** does not include the SSO port, :7444.

9

10 Select the certificate type from the **Choose Action** menu.

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

Certificates that you import must be trusted and must also be applicable to all instances of vRealize Appliance and any load balancer by using Subject Alternative Name (SAN) certificates.

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

- The client/server certificate signed by the intermediate CA certificate
- One or more intermediate certificates
- A root CA certificate

Option	Action
Import PEM Encoded Certificate	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. c (Optional) If your certificate uses a pass phrase to encrypt the certificate key, copy the pass phrase and paste it in the **Pass Phrase** text box.
Generate Self-Signed Certificate	a Type a common name for the self-signed certificate in the **Common Name** text box. You can use the fully qualified domain name of the virtual appliance (*hostname.domain.name*) or a wild card, such as *.*mycompany*.com. b Type your organization name, such as your company name, in the **Organization** text box. c Type your organizational unit, such as your department name or location, in the **Organizational Unit** text box. d Type a two-letter ISO 3166 country code, such as **US**, in the **Country** text box.
Keep Existing	Leave the current SSL configuration. Select this option to cancel your changes.

11 Click **Apply Settings**.

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

12 Join the Identity Appliance to your Native Active Directory domain.

For migration, you must configure Native Active Directory. If you are not migrating, Native Active Directory is optional.

a Click the **Active Directory** tab.

b Type the domain name of the Active Directory in **Domain Name**.

c Enter the credentials for the domain administrator in the **Domain User** and **Password** text boxes.

d Click **Join AD Domain**.

13 Click the **Admin** tab.

14 Verify that the SSH settings are correct.

When **SSH service enabled** is selected, SSH is enabled for all but the root user. Select or uncheck **Administrator SSH login enabled** to enable or disable SSH login for the root user.

The SSO host is initialized. If Identity Appliance does not function correctly after configuration, redeploy and reconfigure the appliance. Do not make changes to the existing appliance.

Deploy and Configure the vRealize Appliance

The vRealize Appliance is a preconfigured virtual appliance that deploys the vRealize Appliance server and Web console (the user portal). It is delivered as an open virtualization format (OVF) template. The system administrator downloads the appliance and deploys it into the vCenter Server or ESX/ESXi inventory.

1 Deploy the vRealize Appliance on page 20

To deploy the vRealize Appliance, a system administrator must log in to the vSphere client and select deployment settings.

2 Enable Time Synchronization on the vRealize Appliance on page 21

Clocks on the Identity Appliance server, vRealize Automation server, and Windows servers must be synchronized to ensure a successful installation.

3 Configure the vRealize Appliance on page 22

To prepare the vRealize Appliance for use, a system administrator configures the host settings, generates an SSL certificate, and provides SSO connection information.

Deploy the vRealize Appliance

To deploy the vRealize Appliance, a system administrator must log in to the vSphere client and select deployment settings.

Prerequisites

- Download the vRealize Appliance from the VMware Web site.
- Log in to the vSphere client as a user with **system administrator** privileges.

Procedure

1 Select **File > Deploy OVF Template** from the vSphere client.

2 Browse to the vRealize Appliance file you downloaded and click **Open**.

3 Click **Next**.

4 Click **Next** on the OVF Template Details page.

5 Accept the license agreement and click **Next**.

6 Type a unique virtual appliance name according to the IT naming convention of your organization in the **Name** text box, select the datacenter and location to which you want to deploy the virtual appliance, and click **Next**.

7 Follow the prompts until the Disk Format page appears.

8 Verify on the Disk Format page that enough space exists to deploy the virtual appliance and click **Next**.

9 Follow the prompts to the Properties page.

The options that appear depend on your vSphere configuration.

10 Configure the values on the Properties page.

a Type the root password to use when you log in to the virtual appliance console in the **Enter password** and **Confirm password** text boxes.

b Type the fully qualified domain name of the virtual machine in the **Hostname** text box, even if you are using DHCP.

c Configure the networking properties.

d Select or uncheck the **SSH service** checkbox to choose whether SSH service is enabled for the appliance.

This value is used to set the initial status of the SSH service in the appliance. You can change this setting from the appliance management console when you configure the appliance.

11 Click **Next**.

12 Restart the host machine.

- If **Power on after deployment** is available on the Ready to Complete page.
  a Select **Power on after deployment** and click **Finish**.
  b Click **Close** after the file finishes deploying into vCenter.
  c Wait for the machine to restart. This could take up to five minutes.
- If **Power on after deployment** is not available on the Ready to Complete page.
  a Click **Close**.
  b Restart the machine. This could take up to five minutes.

After a few moments, a success message appears.

13 Open a command prompt and ping the FQDN to verify that the fully qualified domain name can be resolved against the IP address of vRealize Appliance.

Enable Time Synchronization on the vRealize Appliance

Clocks on the Identity Appliance server, vRealize Automation server, and Windows servers must be synchronized to ensure a successful installation.

If you see certificate warnings during this process, continue past them to finish the installation.

Prerequisites

"Deploy the vRealize Appliance," on page 20.

Procedure

1 Navigate to the vRealize Appliance management console by using its fully qualified domain name, https://*vra-va-hostname.domain.name*:5480/.

2 Log in with the user name root and the password you specified when the appliance was deployed.

3 Select **Admin > Time Settings**.

4 Select an option from the **Time Sync Mode** menu.

Option	Action
Use Time Server	Select **Use Time Server** from the **Time Sync Mode** menu to use Network Time Protocol . For each time server that you are using, type the IP address or the host name in the **Time Server** text box.
Use Host Time	Select **Use Host Time** from the **Time Sync Mode** menu to use VMware Tools time synchronization. You must configure the connections to Network Time Protocol servers before you can use VMware Tools time synchronization.

5 Click **Save Settings**.

6 Click **Refresh**.

7 Verify that the value in **Current Time** is correct.

You can change the time zone as required from the Time Zone Setting page on the **System** tab.

8 (Optional) Click **Time Zone** from the **System** tab and select a system time zone from the menu choices.

The default is Etc/UTC.

9 Click **Save Settings**.

Configure the vRealize Appliance

To prepare the vRealize Appliance for use, a system administrator configures the host settings, generates an SSL certificate, and provides SSO connection information.

Prerequisites

"Enable Time Synchronization on the vRealize Appliance," on page 21.

Procedure

1 Navigate to the vRealize Appliance management console by using its fully qualified domain name, https://*vra-va-hostname.domain.name*:5480/.

2 Continue past the certificate warning.

3 Log in with user name root and the password you specified when you deployed vRealize Appliance.

4 Select **vRA Settings > Host Settings** and select **Resolve Automatically** to view the name of the currently specified host.

5 (Optional) If you want to change the host name, select **Update Host** and enter the fully qualified domain name, *vra-hostname.domain.name*, of the vRealize Appliance in the **Host Name** text box. If you are using a load balancer, enter the fully qualified domain name for the load balancer server.

6

7 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 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:

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

Option	Action
Import	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. 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.
Generate Certificate	a Type a common name for the self-signed certificate in the **Common Name** text box. You can use the fully qualified domain name of the virtual appliance (*hostname.domain.name*) or a wild card, such as *.*mycompany*.com. If you use a load balancer, you need to specify the FQDN of the load balancer or a wildcard that matches the name of the load balancer. If the name is the same as the host name for the virtual appliance, you can leave the text box empty. Do not accept a default value if one is shown, unless it matches the host name of the virtual appliance. b Type your organization name, such as your company name, in the **Organization** text box. c Type your organizational unit, such as your department name or location, in the **Organizational Unit** text box. d Type a two-letter ISO 3166 country code, such as **US**, in the **Country** text box.
Keep Existing	Leave the current SSL configuration. Select this option to cancel your changes.

8 Click **Save Settings** to save host information and SSL configuration.

9 Configure the SSO settings that the vRealize Appliance uses to interact with the Identity Appliance. These settings must match the settings you entered when configuring the Identity Appliance.

a Click **SSO**.

b Type the fully qualified domain name of the Identity Appliance, *identity-va-hostname.domain.name* in the **SSO Host** text box.

For example, **`vcac-sso.mycompany.com`**. Do not use an https:// prefix.

c Edit the default port number, 7444, in the **SSO Port** text box if you are using a nondefault port.

For example, for vCenter Platform Services Controller SSO, specify 443.

d Do not modify the default tenant name, **`vsphere.local`**, in the **SSO Default Tenant** text box.

e Type the default administrator name **`administrator@vsphere.local`** in the **SSO Admin User** text box.

f Type the SSO administrator password in the **SSO Admin Password** text box. The password must match the password you specified in the SSO settings for the Identity Appliance.

g Click **Save Settings**.

After a few minutes, a success message appears and SSO Status is updated to Connected.

h (Optional) Select **Apply Branding** to apply vRealize Automation branding to your installation.

Use this option if you are installing from vCenter and want to use vRealize Automation instead of vCenter branding.

i (Optional) If the spinner does not stop within a few minutes, exit the appliance, close the browser, and log in again.

10 Click **Messaging**. The configuration settings and status of messaging for your appliance is displayed. Do not change these settings.

11 Click the **Telemetry** tab.

You can choose to participate in the Customer Experience Improvement Program. You can unsubscribe from the program at any time.

- Select **Enable** to activate the Program.
- Deselect **Enable** to unsubscribe from the Program.

When you enable the Program, vRealize Automation attempts to establish a connection to https://vmware.com and to automatically discover any proxy server you might have configured for your vRealize Automation deployment.

12 Click **Services** and verify that services are registered.

Depending on your site configuration, this can take about 10 minutes.

NOTE You can log in to the appliance and run `tail –f /var/log/vcac/catalina.out` to monitor startup of the services.

---

13 Configure the license to enable the Infrastructure tab on the vRealize Automation console.

a Click **vRA Settings > Licensing**.

b Click **Licensing**.

c Type a valid vRealize Automation license key that you downloaded when you downloaded the installation files, and click **Submit Key**.

NOTE If you experience a connection error, you might have a problem with the load balancer. Check network connectivity to the load balancer.

---

14 Confirm that you can log in to the vRealize Automation console.

a Open a browser and navigate to https://*vcac-hostname.domain.name*/vcac.

b Accept the vRealize Automation certificate.

c Accept the SSO certificate.

d Log in with administrator@vsphere.local and the password you specified when you configured SSO.

The console opens to the Tenants page on the **Administration** tab. A single tenant named vsphere.local appears in the list.

You have finished the deployment and configuration of your vRealize Appliance. If the appliance does not function correctly after configuration, redeploy and reconfigure the appliance. Do not make changes to the existing appliance.

Apply a vRealize Code Stream License to an Appliance

When you apply the vRealize Code Stream standalone license to a vRealize Automation Appliance, you enable the vRealize Code Stream functions. If you enter the vRealize Code Stream and vRealize Automation licenses, you can enable the vRealize Automation and vRealize Code Stream features.

You can use the vRealize Code Stream standalone license to enable the Artifact Management, Release Management, Release Dashboard, Approval Services, and Advanced Service Designer features. You can apply the vRealize Code Stream license in conjunction with the vRealize Automation Standard, Advanced, or Enterprise licenses.

Overlapping features in vRealize Automation and vRealize Code Stream are combined. For example, if you apply a vRealize Code Stream license to an existing vRealize Automation appliance, the common Advanced Service Designer feature works as is in the appliance.

Prerequisites

Verify that the vRealize Automation appliance is set up. See "Configure the vRealize Appliance," on page 22.

Procedure

1 Open the vRealize Automation Appliance management console with the fully qualified domain name, https:// vra-va-*hostname.domain.name*:5480/.

2 Log in as the root user.

3 Select **vRA Settings > Licensing**.

4 Enter a valid vRealize Code Stream license key and click **Submit Key**.

The default Artifactory server is enabled when the valid license key is accepted.

The vRealize Code Stream license includes the Artifactory Pro version.

5 Confirm that you can log in to the vRealize Automation console.

a Open a Web browser.

b Navigate to https://vra-*hostname.domain.name*/vcac.

What to do next

Configure a repository in the Artifactory server. See the JFrog Web site https://www.jfrog.com/confluence/display/RTF/Configuring+Repositories.

Set Up the Artifactory Server Password

The default Artifactory server is enabled when you apply the vRealize Code Stream license to the appliance. For security purposes, change the default login credentials.

The vRealize Code Stream license includes the Artifactory Pro version.

Procedure

1 Open a Web browser.

2 Type the **`https://vra-hostname/artifactory/`** URL.

3 Log in using the default username **`vmadmin`** and password **`vmware`**.

4 Click the **Admin** tab and click **Users** in the left pane.

5 Select the **vmadmin** user name.

6 Change the default login credentials.

7 Click **Save.**

8 Log out.

Install Artifactory on a Separate Artifactory Server

When you apply the vRealize Code Stream license, the embedded Artifactory Pro version in the vRealize Automation appliance is enabled. In some cases, you might want to install Artifactory on a separate Artifactory server that provides additional storage for artifacts.

Prerequisites

- Create a vRealize Automation appliance. See "Deploy and Configure the Identity Appliance," on page 16.
- Apply the vRealize Code Stream license. See "Apply a vRealize Code Stream License to an Appliance" on page 25.

Procedure

1 Open a Web browser.

2 Type the **`https://vra-hostname/artifactory/`** URL to verify that Artifactory is running.

3 Open a terminal and SSH to the vRealize Automation appliance.

4 Stop the services that are running on the appliance.

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

5 Disable the services so that they do not restart after a reboot.

```
chkconfig vcac-server off
chkconfig vco-server off
chkconfig vco-configurator off
```

6 Verify that the Artifactory server is running.

Configuring Components 3

You must configure components such as vRealize Automation tenants, assign roles to the identity store, and configure the Artifactory server before you can use vRealize Code Stream.

This chapter includes the following topics:

Configuring Additional Tenants

You create the default tenant when you install vRealize Automation, but you can create additional tenants to represent business units in an enterprise or companies that subscribe to cloud services from a service provider.

Tenancy Overview

A tenant is an organizational unit in a vRealize Automation deployment. A tenant can represent a business unit in an enterprise or a company that subscribes to cloud services from a service provider.

Each tenant has its own dedicated configuration. Some system-level configuration is shared across tenants.

Table 3-1. Tenant Configuration

Configuration Area	Description
Login URL	Each tenant has a unique URL to the vRealize Automation console. ■ The default tenant URL is in the following format: https://*hostname*/vcac ■ The URL for additional tenants is in the following format: https://*hostname*/vcac/org/*tenantURL*
Identity stores	Each tenant requires access to one or more directory services, such as OpenLDAP or Microsoft Active Directory servers, that are configured to authenticate users. You can use the same directory service for more than one tenant, but you must configure it separately for each tenant.
Branding	A tenant administrator can configure the branding of the vRealize Automation console including the logo, background color, and information in the header and footer. System administrators control the default branding for all tenants.

Table 3-1. Tenant Configuration (Continued)

Configuration Area	Description
Notification providers	System administrators can configure global email servers that process email notifications. Tenant administrators can override the system default servers, or add their own servers if no global servers are specified.
Business policies	Administrators in each tenant can configure business policies such as approval workflows and entitlements. Business policies are always specific to a tenant.
Service catalog offerings	Service architects can create and publish catalog items to the service catalog and assign them to service categories. Services and catalog items are always specific to a tenant.
Infrastructure resources	The underlying infrastructure fabric resources, for example, vCenter servers, Amazon AWS accounts, or Cisco UCS pools, are shared among all tenants. For each infrastructure source that vRealize Automation manages, a portion of its compute resources can be reserved for users in a specific tenant to use.

About the Default Tenant

When the system administrator configures single sign-on during the installation of vRealize Automation, a default tenant is created with the built-in system administrator account to log in to the vRealize Automation console. The system administrator can then configure the default tenant and create additional tenants.

The default tenant supports all of the functions described in Tenant Configuration. In the default tenant, the system administrator can also manage system-wide configuration, including global system defaults for branding and notifications, and monitor system logs.

The default tenant is the only tenant that supports native Active Directory authentication. All other tenants must use Active Directory over OpenLDAP.

User and Group Management

All user authentication is handled through single sign-on. Each tenant has one or more identity stores, such as Active Directory servers, that provide authentication.

The system administrator performs the initial configuration of single sign-on and basic tenant setup, including designating at least one identity store and a tenant administrator for each tenant. Thereafter, a tenant administrator can configure additional identity stores and assign roles to users or groups from the identity stores.

Tenant administrators can also create custom groups within their own tenant and add users and groups defined in the identity store to custom groups. Custom groups, like identity store groups and users, can be assigned roles or designated as the approvers in an approval policy.

Tenant administrators can also create business groups within their tenant. A business group is a set of users, often corresponding to a line of business, department or other organizational unit, that can be associated with a set of catalog services and infrastructure resources. Users, identity store groups, and custom groups can be added to business groups.

Comparison of Single-Tenant and Multitenant Deployments

vRealize Automation supports deployments with either a single tenant or multiple tenants. The configuration can vary depending on how many tenants are in your deployment.

System-wide configuration is always performed in the default tenant and can apply to one or more tenants. For example, system-wide configuration might specify defaults for branding and notification providers.

Infrastructure configuration, including the infrastructure sources that are available for provisioning, can be configured in any tenant and is shared among all tenants. The infrastructure resources, such as cloud or virtual compute resources or physical machines, can be divided into fabric groups managed by fabric administrators. The resources in each fabric group can be allocated to business groups in each tenant by using reservations.

Single-Tenant Deployment

In a single-tenant deployment, all configuration can occur in the default tenant. Tenant administrators can manage users and groups, configure tenant-specific branding, notifications, business policies, and catalog offerings.

All users log in to the vRealize Automation console at the same URL, but the features available to them are determined by their roles.

Figure 3-1. Single-Tenant Example

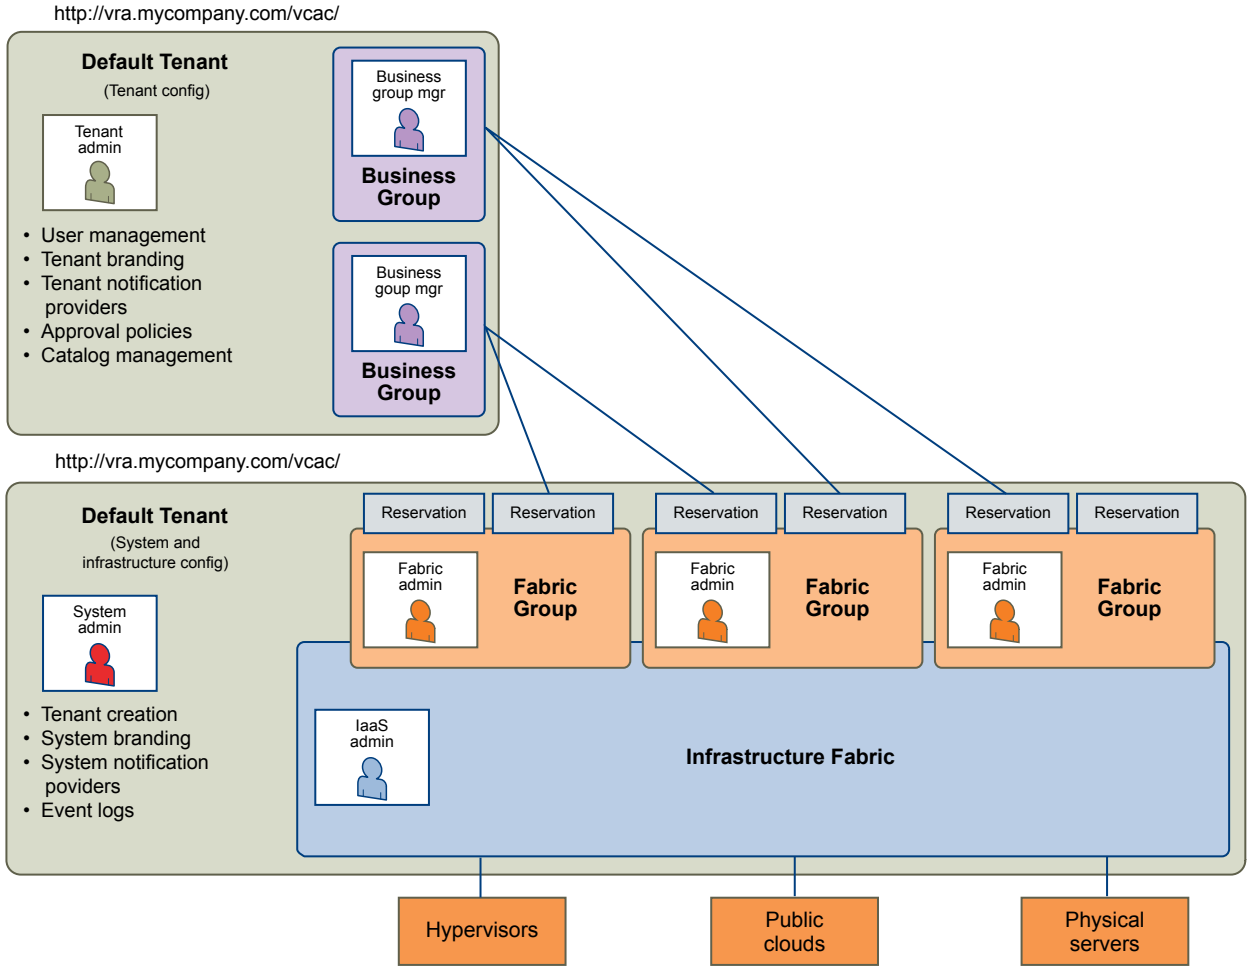

NOTE In a single-tenant scenario, it is common for the system administrator and tenant administrator roles to be assigned to the same person, but two distinct accounts exist. The system administrator account is always administrator@vsphere.local. The tenant administrator must be a user in one of the tenant identity stores, such as *username*@mycompany.com.

Multitenant Deployment

In a multitenant environment, the system administrator creates tenants for each organization that uses the same vRealize Automation instance. Tenant users log in to the vRealize Automation console at a URL specific to their tenant. Tenant-level configuration is segregated from other tenants and from the default tenant. Users with system-wide roles can view and manage configuration across multiple tenants.

There are two main scenarios for configuring a multi-tenant deployment.

Table 3-2. Multitenant Deployment Examples

Example	Description
Manage infrastructure configuration only in the default tenant	In this example, all infrastructure is centrally managed by IaaS administrators and fabric administrators in the default tenant. The shared infrastructure resources are assigned to the users in each tenant by using reservations.
Manage infrastructure configuration in each tenant	In this scenario, each tenant manages its own infrastructure and has its own IaaS administrators and fabric administrators. Each tenant can provide its own infrastructure sources or can share a common infrastructure. Fabric administrators manage reservations only for the users in their own tenant.

The following diagram shows a multitenant deployment with centrally managed infrastructure. The IaaS administrator in the default tenant configures all infrastructure sources that are available for all tenants. The IaaS administrator can organize the infrastructure into fabric groups according to type and intended purpose. For example, a fabric group might contain all virtual resources, or all Tier One resources. The fabric administrator for each group can allocate resources from their fabric groups. Although the fabric administrators exist only in the default tenant, they can assign resources to business groups in any tenant.

NOTE Some infrastructure tasks, such as importing virtual machines, can only be performed by a user with both the fabric administrator and business group manager roles. These tasks might not be available in a multitenant deployment with centrally managed infrastructure.

Figure 3-2. Multitenant Example with Infrastructure Configuration Only in Default Tenant

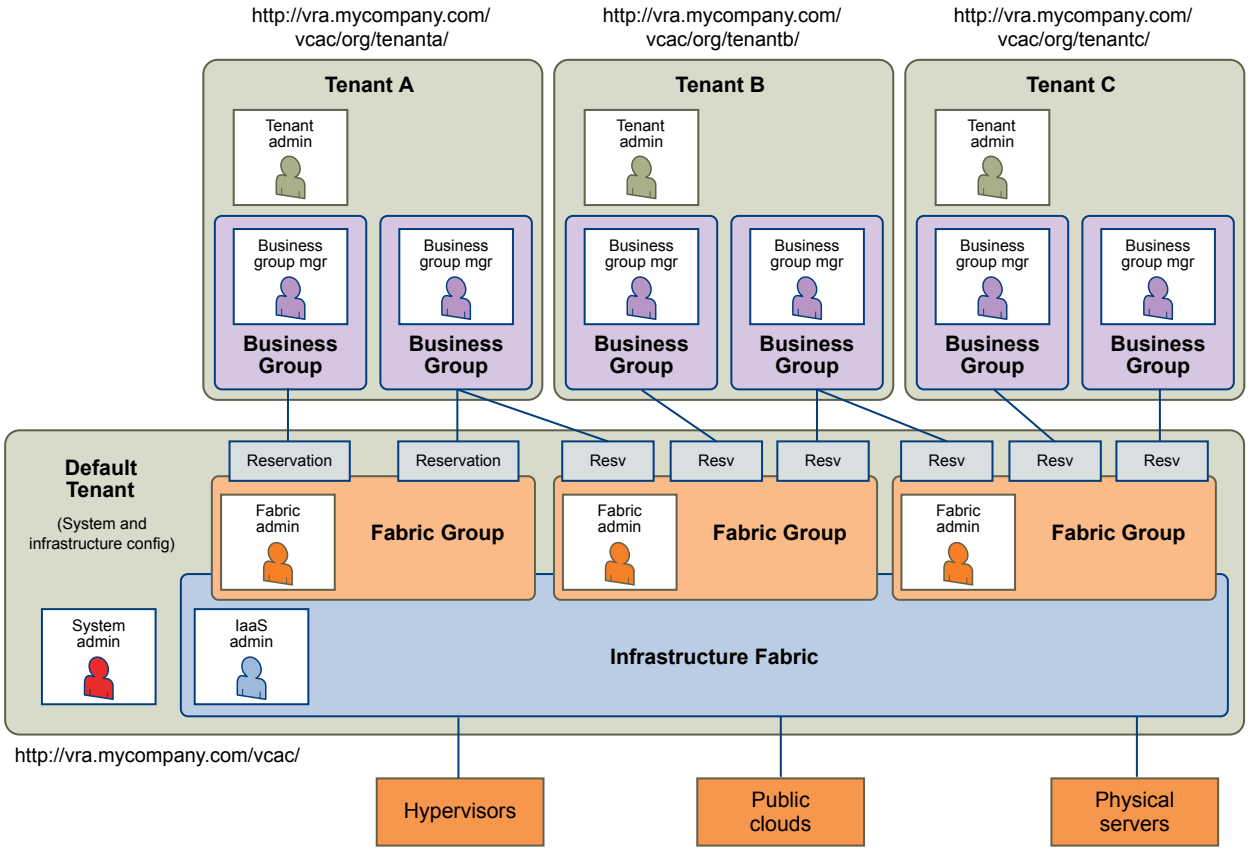

The following diagram shows a multitenant deployment where each tenant manages their own infrastructure. The system administrator is the only user who logs in to the default tenant to manage system-wide configuration and create tenants.

Each tenant has an IaaS administrator, who can create fabric groups and appoint fabric administrators with their respective tenants. Although fabric administrators can create reservations for business groups in any tenant, in this example they typically create and manage reservations in their own tenants. If the same identity store is configured in multiple tenants, the same users can be designated as IaaS administrators or fabric administrators in each tenant.

Figure 3-3. Multitenant Example with Infrastructure Configuration in Each Tenant

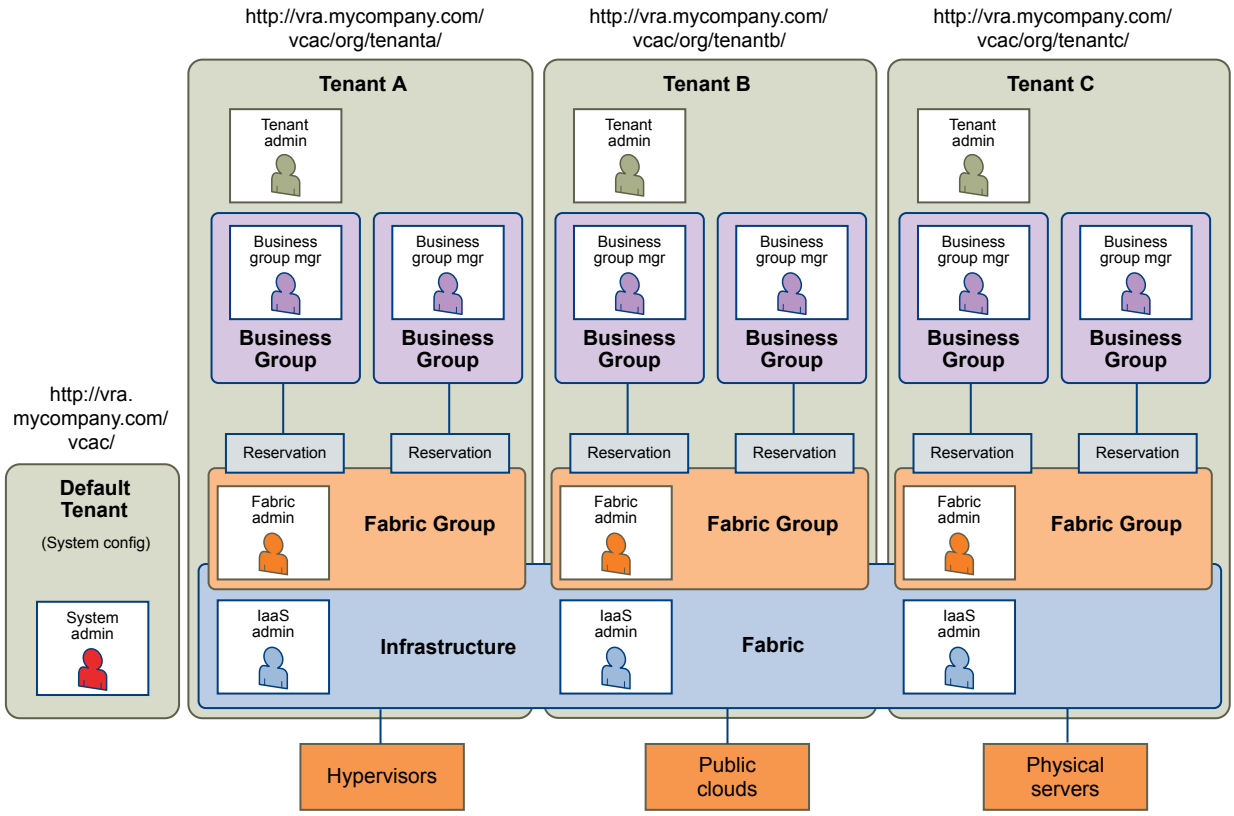

Create and Configure a Tenant

System administrators create tenants and specify basic configuration such as name, login URL, identity stores, and administrators.

Prerequisites

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

Procedure

1 Specify Tenant Information on page 33

The first step to configuring a tenant is to add the new tenant to vRealize Automation and create the tenant-specific access URL.

2 Configure Identity Stores on page 33

Each tenant must be associated with at least one identity store. Identity stores can be OpenLDAP or Active Directory. Use of Native Active Directory is also supported for the default tenant.

3 Appoint Administrators on page 34

You can appoint one or more tenant administrators and IaaS administrators from the identity stores you configured for a tenant.

Specify Tenant Information

The first step to configuring a tenant is to add the new tenant to vRealize Automation and create the tenant-specific access URL.

Prerequisites

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

Procedure

1 Select **Administration > Tenants**.

2 Click the **Add** icon ( ).

3 Enter a name in the **Name** text box.

4 (Optional) Enter a description in the **Description** text box.

5 Type a unique identifier for the tenant in the **URL Name** text box.

This URL token is used to create tenant-specific URLs to access vRealize Automation.

6 (Optional) Type an email address in the **Contact Email** text box.

7 Click **Submit and Next**.

Your new tenant is saved and you are automatically directed to the **Identity Stores** tab for the next step in the process.

Configure Identity Stores

Each tenant must be associated with at least one identity store. Identity stores can be OpenLDAP or Active Directory. Use of Native Active Directory is also supported for the default tenant.

Prerequisites

"Specify Tenant Information," on page 33.

Procedure

1 Click the **Add** icon ( ).

2 Enter a name in the **Name** text box.

3 Select the type of identity store from the **Type** drop-down menu.

4 Type the URL for the identity store in the **URL** text box.

For example, **`ldap://ldap.mycompany.com:389`** .

5 Type the domain for the identity store in the **Domain** text box.

6 (Optional) Type the domain alias in the **Domain Alias** text box.

The alias allows users to log in by using *userid@domain-alias* rather than *userid@identity-store-domain* as a user name.

7 Type the Distinguished Name for the login user in the **Login User DN** text box.

Use the display format of the user name, which can include spaces and is not required to be identical to the user ID.

For example, **`cn=Demo Admin,ou=demo,dc=dev,dc=mycompany,dc=com`**.

8 Type the password for the identity store login user in the **Password** text box.

9 Type the group search base Distinguished Name in the **Group Search Base DN** text box.

For example, `ou=demo,dc=dev,dc=mycompany,dc=com`.

10 (Optional) Type the user search base Distinguished Name in the **User Search Base DN** text box.

For example, `ou=demo,dc=dev,dc=mycompany,dc=com`.

11 Click **Test Connection**.

Check that the connection is working.

12 Click **Add**.

13 (Optional) Repeat Step 1 to Step 12 to configure additional identity stores.

14 Click **Next**.

Your new identity store is saved and associated with the tenant. You are directed to the **Administrators** tab for the next step in the process.

Appoint Administrators

You can appoint one or more tenant administrators and IaaS administrators from the identity stores you configured for a tenant.

Tenant administrators are responsible for configuring tenant-specific branding, as well as managing identity stores, users, groups, entitlements, and shared blueprints within the context of their tenant. IaaS Administrators are responsible for configuring infrastructure source endpoints in IaaS, appointing fabric administrators, and monitoring IaaS logs.

Prerequisites

- "Configure Identity Stores," on page 33.
- Before you appoint IaaS administrators, you must install IaaS. For more information about installation, see *Installation and Configuration*.

Procedure

1 Type the name of a user or group in the **Tenant Administrators** search box and press Enter.

Repeat this step to appoint additional tenant administrators.

2 Type the name of a user or group in the **Infrastructure Administrators** search box and press Enter.

Repeat this step to appoint additional IaaS administrators.

3 Click **Add**.

Managing Users

Tenant administrators create and manage custom groups and grant and manage user access rights to the vRealize Automation console.

Add Identity Store

vRealize Automation uses identity stores to authenticate users. Each tenant is associated with at least one identity store when it is created, but you can add new ones if necessary.

When you delete an identity store, this removes the roles assigned to users from this store, the roles assigned to users from custom groups, and the information about which services are available to this user. Entries for entitlements and business groups are not affected.

Prerequisites

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

Procedure

1 Select **Administration > Identity Stores**.

2 Click the **Add** icon ( ).

3 Enter a name in the **Name** text box.

4 Select the type of the identity store from the **Type** drop-down menu.

5 Enter the following Identify Store configuration options.

Option	Action
URL	Enter the URL for the identity store. For example, `ldap://10.141.64.166:875`.
Domain	Enter the domain for the identity store.
(Optional) Domain Alias	Enter the domain alias.
Login User DN	Enter the login user Distinguished Name. For example, `cn=demoadmin,ou=demo,dc=dev,dc=mycompany,dc=com`.
Password	Enter the password for the identity store login user.
Group Search Base DN	Enter the group search base Distinguished Name. For example, `ou=demo,dc=dev,dc=mycompany,dc=com`.
User Search Base DN	Enter the user search base Distinguished Name.

6 Click **Test Connection**.

7 Click **Add**.

What to do next

"Assign Roles to Identity Store Users or Groups," on page 35.

Assign Roles to Identity Store Users or Groups

Tenant administrators grant users access rights by assigning roles to users or groups.

Prerequisites

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

Procedure

1 Select **Administration > Users & Groups > Identity Store Users & Groups**.

2 Enter a user or group name in the **Search** box and press Enter.

Do not use an at sign (@), backslash (\), or slash (/) in a name. You can optimize your search by typing the entire user or group name in the form user@domain.

3 Click the name of the user or group to which you want to assign roles.

4 Select one or more roles from the Add Roles to this User list.

The Authorities Granted by Selected Roles list indicates the specific authorities you are granting.

5 (Optional) Click **Next** to view more information about the user or group.

6 Click **Update**.

Users who are currently logged in to the vRealize Automation console must log out and log back in to the vRealize Automation console before they can navigate to the pages to which they have been granted access.

What to do next

Optionally, you can create your own custom groups from users and groups in your identity stores. See "Create a Custom Group," on page 36.

Create a Custom Group

Tenant administrators can create custom groups by combining other custom groups, identity store groups, and individual identity store users.

You can assign roles to your custom group, but it is not necessary in all cases. For example, you can create a custom group called Machine Specification Approvers, to use for all machine pre-approvals. You can also create custom groups to map to your business groups so that you can manage all groups in one place. In those cases, you do not need to assign roles.

Prerequisites

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

Procedure

1 Select **Administration > Users & Groups > Custom Groups**.

2 Click the **Add** icon ( ).

3 Enter a group name in the **New Group Name** text box.

Custom group names cannot contain the combination of a semicolon (;) followed by an equal sign (=).

4 (Optional) Enter a description in the **New Group Description** text box.

5 Select one or more roles from the Add Roles to this Group list.

The Authorities Granted by Selected Roles list indicates the specific authorities you are granting.

6 Click **Next**.

7 Add users and groups to create your custom group.

a Enter a user or group name in the **Search** box and press Enter.

Do not use an at sign (@), backslash (\), or slash (/) in a name. You can optimize your search by typing the entire user or group name in the form user@domain.

b Select the user or group to add to your custom group.

8 Click **Add**.

Users who are currently logged in to the vRealize Automation console must log out and log back in to the vRealize Automation console before they can navigate to the pages to which they have been granted access.

Configure a Tenant to a User in the Artifactory Server

For organizations that have Artifactory with multiple repositories, a tenant administrator can map a vRealize Automation tenant to the corresponding user in the Artifactory server from the vRealize Code Stream user interface. The mapping lets the organization control the access to the Artifactory repository among its users.

The Artifactory administrator can set up permissions which include the group, user, and the Artifactory repository.

IMPORTANT The Artifactory users must belong to the vRealize group in Artifactory.

Prerequisites

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

  If you are mapping users for a different tenant such as Dev-Artifactory tenant, log out and log in as that tenant.
- Verify that the Artifactory repositories, users, and groups are configured. See the *Artifactory User Guide* on the jFrog Web site.

Procedure

1 Select **Administration > Artifact Management**.

2 Enter a name for the Artifactory server.

  You can add notes that apply to the server in the Description section.

3 Enter the URL of the Artifactory server.

  The server URL format is https://vra-*hostname*/artifactory.

4 Enter the Artifactory user name to map that user to the tenant.

  You can add notes that apply to the user and tenant in the Description section.

5 Enter a name for the Artifactory server.

  You can add notes that apply to the server in the Description section.

6 Enter the Artifactory password set for the user.

7 Select the artifact and click **Test Connection** to verify that the appliance can connect to the Artifactory server.

8 Click **Update** to save your changes.

What to do next

You can create a remote repository and map it to the remote repository in the embedded vRealize Code Stream Artifactory server. To create a remote repository on the jFrog Web site, see the *Configuring Artifactory* guide.

Configure an External Disk Partition for the Artifactory Server

If you plan to store large binaries you can configure external disk partition to the default Artifactory server and synchronize the existing artifacts.

The local storage in the appliance is 25 GB.

Prerequisites

- Verify that you have a vRealize Automation host.
- Verify that the NFS share path is available with adequate storage.

Procedure

1 Open a terminal and SSH in the vRealize Automation host using root@*VRA-host-name*.

2 Type the vRealize Automation appliance password.

3 Stop the artifactory server.

```
/opt/jfrog/artifactory/bin/artifactroy.sh stop
```

4 Mount the NFS share to the vRealize Automation host.

```
mount NFS-server-host-name:/host/NFS-share-host-path /mount/VRA-path
```

5 Navigate to the `/opt/jfrog/artifactory/misc/db/postgresql.properties` file.

6 Uncomment the binary.provider.type entry in the file.

```
binary.provider.type=filesystem
```

7 Add the new NFS share path in the file.

```
binary.provider.filesystem.dir=NFS-share-path
```

8 Restart the artifactory server.

```
/opt/jfrog/artifactory/bin/artifactroy.sh start
```

What to do next

Synchronize the existing artifacts to the new external disk partition. See Changing the Default Storage topic on the JFrog Web site.

Create a Service Blueprint

A blueprint is a complete specification for a service. With service blueprints, you can publish predefined and custom vRealize Orchestrator workflows as catalog items for requesting.

Blueprints for requesting run workflows with no provisioning and provide no options for managing a provisioned item. You cannot perform post-provisioning operations on this type of provisioned resource. For example, you can create a service blueprint for adding a Puppet Master.

Prerequisites

Log in to the vRealize Automation console as a **service architect**.

Procedure

1 Select **Advanced Services > Service Blueprints**.

2 Click **Add** ( ).

3 Navigate through the vRealize Orchestrator workflow library and select a workflow.

You can see the name and description of the selected workflow, and the input and output parameters as they are defined in vRealize Orchestrator.

For example, you can select the Puppet workflow.

4 Click **Next**.

5 Enter a name and, optionally, a description.

The **Name** and **Description** text boxes are prepopulated with the name and description of the workflow as they are defined in vRealize Orchestrator.

6 (Optional) If you do not want to prompt consumers to enter a description and reason for requesting this resource action, select the **Hide catalog request information page** check box.

7 (Optional) Add the version number of the workflow if you have existing iterations of the workflow.

8 Click **Next**.

9 Accept the default service blueprint.

By default, the service blueprint form is mapped to the vRealize Orchestrator workflow presentation.

10 Click **Next**.

11 Accept the default No provisioning output parameter.

The service blueprint does not add new items on the **Items** tab.

12 Click **Add**.

The new service blueprint appears on the Service blueprints page.

Publish a Service Blueprint as a Catalog Item

After you create a service blueprint, it is in a draft state and you can publish it as a catalog item.

Prerequisites

Log in to the vRealize Automation console as a **service architect**.

Procedure

1 Select **Advanced Services > Service Blueprints**.

2 Select the row of the service blueprint to publish, and click **Publish**.

The status of the service blueprint changes to Published.

3 (Optional) Select **Administration > Catalog Management > Catalog Items** to view the published catalog item.

4 Registering Components

You must register various components such as plug-ins, endpoints, and the Artifactory server before you can model and run release pipelines in vRealize Code Stream.

This chapter includes the following topics:

Register an Artifactory Server for Artifact Management

To use artifact management in a release pipeline, you must connect to an Artifactory server.

With artifact management, you can specify an artifact by name and search type from the server, but not by location or unique identifier. Artifact management monitors the physical location and identity of artifacts and supplies the required artifact during the pipeline execution.

Prerequisites

- Log in to the vRealize Automation console as a **tenant administrator**.
- Verify that the Artifactory server is not using the default login credentials. See "Set Up the Artifactory Server Password," on page 25.
- Verify that the tenant role is assigned for the **Release Automation** and **Release Dashboard** tabs to appear in the appliance user interface. See "Configuring Additional Tenants," on page 27.
- Verify that the Artifactory server is configured. See the *Artifactory User Guide* on the jFrog Web site.

Procedure

1 Select **Administration > Artifact Management**.

2 Enter a name for the Artifactory server.

You can add notes that apply to the server in the Description section.

3 Enter the URL of the Artifactory server.

The server URL format is https://vra-*hostname*/artifactory.

4 Enter the Artifactory user name to map that user to the tenant.

You can add notes that apply to the user and tenant in the Description section.

5 Enter the Artifactory password set for the user.

6 Select the artifact and click **Test Connection** to verify that the appliance can connect to the Artifactory server.

7 Click **Update** to save your changes.

The release pipeline can access the Artifactory server.

Registering Plug-In Instances and Endpoints for a Release Pipeline

You must define and configure the plug-in instance or endpoint before you create a task in a release pipeline.

A plug-in instance allows parameters to be defined in the custom task. You can access the custom workflow in the vRealize Orchestrator client when you execute the pipeline.

An endpoint allows the parameters defined in the provision or test tasks in the release pipeline to access the vRealize Automation server or Jenkins server when you run the pipeline.

Register a Jenkins Server Endpoint

You can run tests or other jobs by using the Jenkins plug-in that allows you to use custom automation and scripts.

You can use any Jenkins job in the Jenkins server in the release pipeline with this endpoint enabled. You can use this endpoint to invoke a Jenkins build job during the modeling of a release pipeline and execute the job as part of the release pipeline.

Register a Jenkins Server Endpoint (http://link.brightcove.com/services/player/bcpid2296383276001?bctid=ref:video_vrcs_register_jenkins)

Prerequisites

- Verify that the Jenkins server is available and configured with or without SSL.
- Verify that the Jenkins server version is 1.561 or later.
- Verify that the Jenkins jobs are created in the Jenkins server with the input string parameter, VRCSTestExecutionId.
- Log in to the vRealize Automation console as a **system administrator** or **release manager**.

Procedure

1 Select **Administration > Orchestration Configuration > Endpoints**.

2 Click **Add**.

3 Select **Jenkins (Code Stream)** from the **Plug-In** drop-down menu and click **Next**.

4 Enter a Jenkins server endpoint name and an applicable description.

For example, in the description section you can add the release pipeline name that uses this Jenkins server endpoint.

5 Click **Next**.

6 Enter the Jenkins server configuration details.

Option	Description
Jenkins Instance Name	Enter the Jenkins instance name. For example, qe-jenkins.test.com.
User Credentials	User name and password for the Jenkins server.
URL	Enter the host URL as *protocol://host:port*. Sample host URL, 192.10.121.13:8080.

Option	Description
Polling Interval	Time that the task must wait to check the progress.
Request Retry Count	Number of times to retry the scheduled build request for the Jenkins server.
Retry Wait Time	Seconds to wait before retrying the build request for the Jenkins server.
Offline Creation	Require a validation and certificate acceptance when the endpoint is created. You can accept the default setting or select **Yes** from the drop-down menu to enable this configuration.

7 Click **Add**.

8 Click the **Code Stream** tab in vRealize Automation to continue with the task configuration.

What to do next

Create a test task to use this endpoint in the release pipeline. See the *Using vRealize Code Stream* guide.

Register a vRealize Automation Server Endpoint

When you register a vRealize Automation, vRealize Code Stream invokes a local or remote vRealize Automation 6.1. or 6.2 instance to provision infrastructure in a specific environment

vRealize Code Stream can also invoke multiple vRealize Automation instances.

From a provision task, the release manager can select infrastructure services to be provisioned on resources reserved for that tenant. Each virtual machine uses the resources specified in the machine blueprint when the release pipeline runs.

Register a vRealize Automation Server Endpoint (http://link.brightcove.com/services/player/bcpid2296383276001?bctid=ref:video_vrcs_register_vra)

Prerequisites

- Log in to the vRealize Automation console as a **system administrator** or **release manager**.
- Verify that you have configured the following components in vRealize Automation.
- Choosing an Endpoint Scenario see *IaaS Configuration for Virtual Platforms*.
- Create a Fabric Group see *IaaS Configuration for Virtual Platforms*.
- Create a Business Group see *IaaS Configuration for Virtual Platforms*.
- Create a Reservation see *IaaS Configuration for Virtual Platforms*.
- Create a Reservation Policy see *IaaS Configuration for Physical Machines*.
- Create a Network Profile see *IaaS Integration for Multi-Machine Services*.
- Create a Blueprint see *IaaS Configuration for Virtual Platforms*.
- Publish a Blueprint see *IaaS Configuration for Virtual Platforms*.

Procedure

1 Select **Administration > Orchestration Configuration > Endpoints**.

2 Click **Add**.

3 Select **vRA Provisioning (Code Stream)** from the Plug-in drop-down menu and click **Next**.

4 Enter a vRealize Automation endpoint name and an applicable description.

5 Click **Next**.

6 Enter the vRealize Automation server configuration details.

Option	Description
vRA Instance Name	Enter the vRealize Automation instance name.
Polling Interval	Time that the task must wait to check the progress.
Maximum Wait for IP assignment	Time that the task must wait for the IP addresses to be assigned to a machine before it can report the failure.

7 Click **Next**.

8 Define the user credentials for the vRealize Automation instance authentication.

- Select **Shared session** from the drop-down menu to use provisioned blueprints from a remote vRealize Automation instance.

  You must provide the secure vRealize Automation host and port URL, the tenant name, and the user credentials of the user sharing the session for instance authentication.

  NOTE Select share session to use provisioned blueprints from a remote vRealize Automation instance using a different identity appliance than vRealize Code Stream.

- Select **Per User session** from the drop-down menu to use provisioned blueprints in this vRealize Automation.

  The logged in user credentials are applied during instance authentication.

  NOTE This session does not work if the vRealize Automation instance you are registering is using a different identity appliance than vRealize Code Stream.

9 Click **Add**.

10 Select the **Code Stream** tab invRealize Automation to continue with the task configuration.

What to do next

Create a provision task to use this endpoint in the release pipeline. See the *Using vRealize Code Stream* guide.

Register a vRealize Orchestrator Workflow for a Custom Task

With the vRealize Orchestrator workflow plug-in for the custom task, you can use any of the custom workflows in vRealize Automation. You can select the custom workflow, configure it, and use it to model the release pipeline, and to execute the pipeline.

Prerequisites

- Verify that the vRealize Automation server is installed with vRealize Code Stream.
- Verify that the workflow for a custom workflow is created.

Procedure

1 Log in to the vRealize Orchestrator client to create a workflow.

2 Select **Library** > **Tagging** > **Tag workflow**.

The Tag workflow is required for custom user workflows.

By default, the Manual Task workflow is provided.

3 Right-click the **Tag workflow** and select **Start Workflow**.

4 Click the **Tagged Workflow** text box to select the custom workflow.

For example, the tagged workflow can be Deploy Spring Travel.

5 Enter the tag as **vRCS_CUSTOM**.

6 Enter the value as **vRCS_CUSTOM**.

7 Click **Yes** for the Global tag and click **Submit**.

8 Click the **Code Stream** tab in vRealize Automation to continue with the task configuration.

Register a vRealize Orchestrator Workflow for a Gating Rule

With the vRealize Orchestrator workflow plug-in for gating rules, you can use any workflow as a gating rule workflow from release automation. You can select the custom workflow, configure it with the vRCS_GATING_RULE tag, and use it to model the release pipeline and execute the release pipeline.

Prerequisites

- Verify that the vRealize Automation server is installed with vRealize Code Stream.
- Verify that the workflow for vRCS_GATING_RULE Workflow is created.

Procedure

1 Log in to the vRealize Orchestrator client to create a workflow.

2 Select **Library** > **Tagging** > **Tag workflow**.

3 Right-click the **Tag workflow** and select **Start Workflow**.

4 Click the **Tagged Workflow** text box to select the **vRCS_GATING_RULE** workflow.

5 Enter the tag as **vRCS_GATING_RULE**.

6 Enter the value as **vRCS_GATING_RULE**.

7 Click **Yes** for the Global tag and click **Submit**.

8 Click the **Code Stream** tab in vRealize Automation to continue with the task configuration.

Register a Team Foundation Server Endpoint

You can connect to the Team Foundation Server plug-in to manage version control, track defects and work items, and manage your build projects.

Prerequisites

Verify that you have installed and configured Visual Studio Team Foundation Server 2013.

Procedure

1 Select **Administration > Orchestration Configuration > Endpoints**.

2 Click **Add**.

3 Select **Team Foundation Server (Code Stream)** from the Plug-in drop-down menu and click **Next**.

4 Enter an Team Foundation Server endpoint name and an applicable description.

5 Click **Next**.

6 Enter the Team Foundation Server configuration details.

Option	Description
Team Foundation Server Instance Name	Enter the Team Foundation Server instance name. For example, qe-tfs-test
User Credentials	User name and password for the Team Foundation Server.
Domain Name	Domain name for the Team Foundation Server. If the Team Foundation Server is not in a domain, you can use the Windows server host name.
URL	Enter the host URL as *protocol://host:port*/tfs. For example, 192.10.121.12:8080/tfs
Polling Interval	Time that the task must wait to check the progress.
Offline Creation	Require a validation and certificate acceptance when the endpoint is created. You can accept the default setting or select **Yes** from the drop-down menu to enable this configuration.

7 Click **Add**.

8 Select the **Code Stream** tab in vRealize Automation to continue with the task configuration.

What to do next

Create a provision task to use this endpoint in the release pipeline. See the *Using vRealize Code Stream* guide.

Index

T

U

V