

IaaS Configuration for Cloud Platforms

vRealize Automation 6.2

vmware[®]

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

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

The VMware Web site also provides the latest product updates.

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

docfeedback@vmware.com

Copyright © 2008–2016 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

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

Contents

IaaS Configuration for Cloud Platforms	7
Updated Information	7
Using the Goal Navigator	8
1 Configuring IaaS	9
Configuring IaaS for Cloud Checklist	9
Custom Properties in Machine Provisioning, Blueprints, and Build Profiles	10
Order of Precedence for Custom Properties	11
Bringing Resources Under vRealize Automation Management	12
Store User Credentials for Cloud Endpoints	12
Choosing an Endpoint Scenario	13
Create an Amazon AWS Endpoint	13
Create an OpenStack or PowerVC Endpoint	14
Preparing an Endpoint CSV File for Import	16
Import a List of Endpoints	16
Create a Fabric Group	17
Configure Machine Prefixes	18
Create a Business Group	18
Managing Key Pairs	19
Create a Key Pair	20
Upload the Private Key for a Key Pair	20
Export the Private Key from a Key Pair	21
Cloud Reservations	21
Cloud Reservation Selection	21
Choosing a Reservation Scenario	22
Create an Amazon AWS Reservation (non-VPC)	22
Create an Amazon AWS Reservation (VPC)	25
Create a Red Hat OpenStack Reservation	28
2 Optional Configurations	31
Build Profiles and Property Sets	31
Create a Property Set	32
Create a Build Profile	33
Configuring Reservation Policies	34
Configure a Reservation Policy	35
Static IP Address Assignment	36
Create a Network Profile for Static IP Address Assignment	36
Assign a Network Profile to a Reservation	38
Cost Information for Cloud Machines	39
Using Optional Amazon Features	39
Using Security Groups for Amazon Web Services	39
Understanding Amazon Web Service Regions	40

- Using Virtual Private Cloud 40
- Using Elastic Load Balancers 41
- Using Elastic IP Addresses 41
- Using Elastic Block Storage 42
- Using Optional Red Hat OpenStack Features 42
 - Using Security Groups 42
 - Using Floating IP Addresses 42
- 3 Preparing for Provisioning 43**
 - Choosing a Cloud Provisioning Scenario 43
 - Preparing for Amazon Provisioning 44
 - Understanding Amazon Machine Images 44
 - Understanding Amazon Instance Types 45
 - Add an Amazon Instance Type 45
 - Preparing for OpenStack Provisioning 46
 - Preparing for Virtual Machine Image Provisioning 46
 - Preparing for Linux Kickstart Provisioning 47
 - Preparing for WIM Provisioning 49
- 4 Creating a Cloud Blueprint 57**
 - Choosing a Blueprint Scenario 57
 - Create an Amazon AWS Blueprint 58
 - Specify Amazon AWS Blueprint Information 58
 - Specify Amazon AWS Blueprint Build Information 59
 - Configure Amazon Machine Resources 60
 - Add Amazon AWS Blueprint Custom Properties 61
 - Configure Amazon AWS Blueprint Actions 61
 - Create a Basic Red Hat OpenStack Blueprint 62
 - Specify Basic Red Hat OpenStack Blueprint Information 62
 - Specify Basic Red Hat OpenStack Blueprint Build Information 63
 - Add Basic Red Hat OpenStack Blueprint Custom Properties 64
 - Configure Basic Red Hat OpenStack Blueprint Actions 65
 - Create a Red Hat OpenStack Blueprint for Linux Kickstart Provisioning 65
 - Specify Linux Kickstart Blueprint Information 66
 - Specify Linux Kickstart Blueprint Build Information 67
 - Add Required Properties to a Linux Kickstart Blueprint 68
 - Configure Linux Kickstart Blueprint Actions 71
 - Create an OpenStack Blueprint for WIM Provisioning 71
 - Specify WIM Blueprint Information 72
 - Specify WIM Blueprint Build Information 73
 - Add Required Properties to a WIM Blueprint 74
 - Configure WIM Blueprint Actions 78
 - Publish a Blueprint 78
- 5 Configuring Advanced Blueprint Settings 79**
 - Reservation Policies 79
 - Add a Reservation Policy to a Blueprint 80

Configuring Network Settings	80
Add a Network Profile for Static IP Assignment to a Blueprint	80
Custom Properties for Networking	81
Enabling Visual Basic Scripts in Provisioning	84
Call a Visual Basic Script from a Blueprint	85
Add Active Directory Cleanup to a Blueprint	86
Enabling Remote Desktop Connections	87
Configure Connections Using RDP	88
Enable Connections Using RDP	89
Enable Connections Using SSH	89
6 Managing Cloud Infrastructure	91
Managing Resource Usage	91
Resource Usage Terminology	91
Choosing a Resource Monitoring Scenario	92
Data Collection	93
Start Endpoint Data Collection Manually	93
Configure Compute Resource Data Collection	94
Manage Amazon EBS Volumes	95
Connecting to a Cloud Machine	96
Collect User Credentials for an Amazon Machine	96
7 Monitoring Workflows and Viewing Logs	99
8 Machine Life Cycle and Workflow States for Cloud Platforms	101
Index	103

IaaS Configuration for Cloud Platforms

IaaS Configuration for Cloud Platforms provides information about integrating cloud platforms such as Amazon Web Services and Red Hat Enterprise Linux OpenStack Platform with VMware vRealize™ Automation.

This documentation guides you through the following processes:

- Bringing resources under vCloud Automation Center management
- Configuring IaaS features and policies
- Preparing for provisioning
- Creating machine blueprints

All of the IaaS configuration steps that you must complete before machine provisioning are included in this document. For information about how to manage provisioned machines, see *Tenant Administration*.

NOTE Not all features and capabilities of vRealize Automation are available in all editions. For a comparison of feature sets in each edition, see <https://www.vmware.com/products/vrealize-automation/>.

Intended Audience

This information is intended for IaaS administrators, fabric administrators, and business group managers of vRealize Automation. This content is written for experienced Windows or Linux system administrators who are familiar with virtualization technology and the basic concepts described in *Foundations and Concepts*.

VMware Technical Publications Glossary

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

Updated Information

This *IaaS Configuration for Cloud Platforms* is updated with each release of the product or when necessary.

This table provides the update history of the *IaaS Configuration for Cloud Platforms*.

Revision	Description
001644-04	<ul style="list-style-type: none"> ■ Updates to “Create an OpenStack or PowerVC Endpoint,” on page 14. ■ Update to clarify static IP for SCVMM is only supported for cloning Linux machines, and only if the guest agent is installed on your templates. See “Add a Network Profile for Static IP Assignment to a Blueprint,” on page 80.
001644-03	Updated the following topics: <ul style="list-style-type: none"> ■ “Add Required Properties to a WIM Blueprint,” on page 74 ■ “Custom Properties for WIM Blueprints,” on page 75 ■ “Preparing for Amazon Provisioning,” on page 44
001644-02	Updated “Add Amazon AWS Blueprint Custom Properties,” on page 61.
001644-01	Updated the following topics regarding Amazon instance types: <ul style="list-style-type: none"> ■ “Understanding Amazon Instance Types,” on page 45 ■ “Configure Amazon Machine Resources,” on page 60
001644-00	Initial 6.2 release.

Using the Goal Navigator

The goal navigator guides you through high-level goals that you might want to accomplish in vRealize Automation.

The goals you can achieve depend on your role. To complete each goal, you must complete a sequence of steps that are presented on separate pages in the vRealize Automation console.

The goal navigator can answer the following questions:

- Where do I start?
- What are all the steps I need to complete to achieve a goal?
- What are the prerequisites for completing a particular task?
- Why do I need to do this step and how does this step help me achieve my goal?

The goal navigator is hidden by default. You can expand the goal navigator by clicking the icon on the left side of the screen.

After you select a goal, you navigate between the pages needed to accomplish the goal by clicking each step. The goal navigator does not validate that you completed a step, or force you to complete steps in a particular order. The steps are listed in the recommended sequence. You can return to each goal as many times as needed.

For each step, the goal navigator provides a description of the task you need to perform on the corresponding page. The goal navigator does not provide detailed information such as how to complete the forms on a page. You can hide the page information or move it to a more convenient position on the page. If you hide the page information, you can display it again by clicking the information icon on the goal navigator panel.

Configuring IaaS

Preparation is required by IaaS administrators, tenant administrators, and fabric administrators to bring resources under vRealize Automation management, allocate resources to users, prepare for provisioning machines, and create machine blueprints.

This chapter includes the following topics:

- [“Configuring IaaS for Cloud Checklist,”](#) on page 9
- [“Custom Properties in Machine Provisioning, Blueprints, and Build Profiles,”](#) on page 10
- [“Bringing Resources Under vRealize Automation Management,”](#) on page 12
- [“Configure Machine Prefixes,”](#) on page 18
- [“Create a Business Group,”](#) on page 18
- [“Managing Key Pairs,”](#) on page 19
- [“Cloud Reservations,”](#) on page 21

Configuring IaaS for Cloud Checklist

IaaS administrators, fabric administrators, tenant administrators, and business group managers perform required and optional configurations to fully implement and customize vRealize Automation.

The Configuring IaaS Checklist provides a high-level overview of the sequence of steps required to have a fully functioning IaaS instance.

Table 1-1. Configuring IaaS checklist

Task	Required Role
<input type="checkbox"/> Store credentials and create endpoints to bring resources under vRealize Automation management. See “Bringing Resources Under vRealize Automation Management,” on page 12.	IaaS administrator
<input type="checkbox"/> Configure the machine prefixes used to create names for machines provisioned through vRealize Automation. See “Configure Machine Prefixes,” on page 18.	Fabric administrator
<input type="checkbox"/> Create at least one business group of users who need to request machines. See “Create a Business Group,” on page 18.	Tenant administrator
<input type="checkbox"/> Create at least one reservation to allocate resources to a business group. See “Choosing a Reservation Scenario,” on page 22.	Fabric administrator
<input type="checkbox"/> Configure optional policies and settings. See Chapter 2, “Optional Configurations,” on page 31.	Fabric administrator

Table 1-1. Configuring IaaS checklist (Continued)

Task	Required Role
<input type="checkbox"/> Depending on your intended method of provisioning, preparation outside of vRealize Automation might be required before you can start creating machine blueprints. For example, you might need to create the machine images required for your blueprints. See "Choosing a Cloud Provisioning Scenario," on page 43.	Outside of vRealize Automation
<input type="checkbox"/> Create and publish machine blueprints. See "Choosing a Blueprint Scenario," on page 57.	<ul style="list-style-type: none"> ■ Tenant administrator ■ Business group manager

Before users can request machines, a tenant administrator must configure the service catalog. See *Tenant Administration*.

Custom Properties in Machine Provisioning, Blueprints, and Build Profiles

Custom properties are name-value pairs used to specify attributes of a machine or to override default specifications.

Different custom properties are used for different provisioning methods, types of machines, and machine options. Custom properties can be used as described in the following examples:

- Specify a particular type of guest OS
- Enable WIM-based provisioning, in which a Windows Imaging File Format (WIM) image of a reference machine is used to provision new machines
- Customize the behavior of Remote Desktop Protocol when connecting to a machine
- Register a virtual machine with a XenDesktop Desktop Delivery Controller (DDC) server
- Customize a virtual machine's system specifications, such as adding multiple hard disk drives
- Customize the guest OS for a machine, for instance, by including specified users in selected local groups
- Enable cleanup of a the Active Directory account of a machine after it is destroyed

Specifying the characteristics of the machines to be provisioned is generally done by adding properties to blueprints and build profiles. You can make custom properties available to multiple blueprints and all business groups by placing them in build profiles.

Any property specified in a blueprint overrides the same property specified in the incorporated build profile. This enables a blueprint to use most of the properties in a profile while differing from the profile in some limited way. For more information, see *Order of Precedence for Custom Properties*.

For example, a blueprint that incorporates a standard developer workstation profile might override the US English settings in the profile with UK English settings. On the other hand, if no appropriate profile is available all the needed properties can be specified in the blueprint itself. This arrangement ensures that the number and complexity of blueprints remain manageable.

At new machine request time, vRealize Automation has not yet allocated a reservation and the compute resource and endpoint are unknown as well. Therefore, only custom properties from a build profile, blueprint and business group are reconciled and presented when the machine is requested.

Order of Precedence for Custom Properties

When the same property exists in more than one source, a specific order is followed when applying properties to the machine.

You can add custom properties that apply to provisioned machines to the following elements:

- A reservation, to apply the custom properties to all machines provisioned from that reservation
- A business group, to apply the custom properties to all machines provisioned by business group members
- A global or local blueprint, to apply the custom properties to all machines provisioned from the blueprint
- Build profiles, which can be incorporated into any global or local blueprint, to apply the custom properties to all machines provisioned from the blueprint
- A machine request, if you are a business group manager, to apply the custom properties to the machine being provisioned
- The applicable approval policy, if any exist and if advanced approval support is enabled, to require approvers to provide the values to be applied to the machine being approved

The full order of precedence for custom properties is that any property value specified in a source later in the list overrides values for the same property specified in sources earlier in the list. The order is shown in the following list:

- 1 Build profile
- 2 Blueprint
- 3 Business group
- 4 Compute resource
- 5 Reservations
- 6 Endpoint
- 7 Runtime

For vApps, the order is similar, as shown in the following list:

- 1 Build profile, specified on a vApp component blueprint
- 2 vApp component blueprint
- 3 Build profile, specified on a vApp blueprint
- 4 vApp blueprint
- 5 Business group
- 6 Compute resources
- 7 Reservations
- 8 Endpoint
- 9 Runtime specified on a vApp
- 10 Runtime specified on a component machine

Any runtime property takes higher precedence and overrides a property from any source. A custom property is marked as runtime if the following conditions exist:

- The property is marked as Prompt User, which specifies that the user must supply a value for it when requesting a machine. This requires that the machine requestor customize individual characteristics of each machine, or gives them the option of doing so when a default value is provided for the required property.
- A business group manager is requesting a machine and the property appears in the custom properties list on the Properties tab of the Confirm Machine Request page.

Custom properties in reservations and business groups may be applied to many machines so they should be used carefully. Their use is typically limited to purposes related to their sources, such as resource management, line of business accounting, and so on. Specifying the characteristics of the machine to be provisioned is generally done by adding properties to blueprints and build profiles.

Each blueprint of any type can optionally incorporate one or more build profiles and thereby inherit the custom properties in those profiles. Build profiles are especially useful for applying common sets of properties for specific purposes to a wide range of blueprints. For example, your site might want to add a second disk to, customize Microsoft Remote Desktop Protocol behavior for, and enable Active Directory cleanup for a wide variety of machines. If a build profile with the necessary properties is created, it can be incorporated into all of your blueprints, local or global.

When creating and managing build profiles, a fabric administrator can load a number of predefined property sets to add several related properties all at once, instead of one by one.

Bringing Resources Under vRealize Automation Management

For vRealize Automation to manage your infrastructure, IaaS administrators must create endpoints, store administrator-level user credentials for those endpoints, and add compute resources to a fabric group.

Depending on your environment, the procedure for creating endpoints and storing credentials differs slightly.

For cloud platforms, compute resources represent regions rather than specific virtualization hosts. vRealize Automation collects information about the regions available on each cloud endpoint and an IaaS administrator can add the regions to a fabric group.

Store User Credentials for Cloud Endpoints

An IaaS administrator stores administrator-level credentials so that vRealize Automation can communicate with your cloud endpoints. Because the same credentials can be used for multiple endpoints, credentials are managed separately from endpoints and associated when endpoints are created or edited.

Prerequisites


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

Procedure

- 1 Select **Infrastructure > Endpoints > Credentials**.
- 2 Click **New Credentials**.
- 3 Enter a name and, optionally, a description.

- 4 Type the credentials in the **User name** and **Password** text boxes.

Cloud Platform	Description
Amazon AWS	Type the access key ID into the User name text box and the Secret access key into the Password text boxes. For example: <ul style="list-style-type: none"> ■ User name: AKIAIOSFODNN7EXAMPLE ■ Password: wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY For information about obtaining your access key ID and secret access key, see the Amazon documentation.
Red Hat OpenStack	Type the user name and password for Red Hat OpenStack user account. The credentials you provide must be a user in the Red Hat OpenStack tenant associated with the endpoint. If you have multiple Red Hat OpenStack endpoints associated with different Red Hat OpenStack tenants, you can store the credentials for a single user who is an administrator in all the tenants, or separate administrator users for each Red Hat OpenStack tenant.

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

What to do next

Now that your credentials are stored, you are ready to create an endpoint.

Choosing an Endpoint Scenario

IaaS administrators are responsible for creating the endpoints that allow vRealize Automation to communicate with your infrastructure. Depending on your environment, the procedure to create the endpoints differs slightly.

Table 1-2. Choosing an Endpoint Scenario

Scenario	Procedure
Connect to an Amazon AWS cloud service account.	“Create an Amazon AWS Endpoint,” on page 13
Connect to a Red Hat OpenStack tenant.	“Create an OpenStack or PowerVC Endpoint,” on page 14
Connect to a PowerVC tenant.	“Create an OpenStack or PowerVC Endpoint,” on page 14
Import a list of endpoints.	<ul style="list-style-type: none"> ■ “Preparing an Endpoint CSV File for Import,” on page 16 ■ “Import a List of Endpoints,” on page 16

Create an Amazon AWS Endpoint

An IaaS administrator creates an endpoint to connect to an Amazon Web Services instance.

Prerequisites

- Log in to the vRealize Automation console as an **IaaS administrator**.
- [“Store User Credentials for Cloud Endpoints,”](#) on page 12.

Procedure

- 1 Select **Infrastructure > Endpoints > Endpoints**.
- 2 Select **New Endpoint > Cloud > Amazon EC2**.
- 3 Enter a name and, optionally, a description.

Typically this name indicates the Amazon Web Services account that corresponds to this endpoint.

- 4 Select the **Credentials** for the endpoint.
Only one endpoint can be associated with an Amazon access key ID.
- 5 (Optional) Click the **Use proxy server** check mark box to configure additional security and force connections to Amazon Web Services to pass through a proxy server.
 - a Type the host name of your proxy server in the **Hostname** text box.
 - b Type the port number to use for connecting to the proxy server in the **Port** text box.
 - c (Optional) Click the **Browse** icon next to the **Credentials** text box.
Select or create credentials that represent the user name and password for the proxy server, if required by the proxy configuration.
- 6 (Optional) Add any custom properties.
- 7 Click **OK**.

After the endpoint is created, vRealize Automation begins collecting data from the Amazon Web Services regions.

What to do next

Add the compute resources from your endpoint to a fabric group.

See [“Create a Fabric Group,”](#) on page 17.

Create an OpenStack or PowerVC Endpoint

An IaaS administrator creates an endpoint to allow vRealize Automation to communicate with your OpenStack or PowerVC instance.

Prerequisites

- Log in to the vRealize Automation console as an **IaaS administrator**.
- [“Store User Credentials for Cloud Endpoints,”](#) on page 12.
- Verify that your vRealize Automation DEMs are installed on a machine that meets the Openstack or PowerVC requirements. See *Installation and Configuration*.

Procedure

- 1 Select **Infrastructure > Endpoints > Endpoints**.
- 2 Select **New Endpoint > Cloud > OpenStack**.
- 3 Enter a name and, optionally, a description.
- 4 Type the URL for the endpoint in the **Address** text box.

Option	Description
PowerVC	The URL must be of the format https://FQDN/powervc/openstack/service . For example: https://openstack.mycompany.com/powervc/openstack/admin .
Openstack	The URL must be of the format FQDN:5000 or IP_address:5000 . Do not include the /v2.0 suffix in the endpoint address. For example: https://openstack.mycompany.com:5000 .

- 5 Select the **Credentials** for the endpoint.
The credentials you provide must have the administrator role in the Red Hat OpenStack tenant associated with the endpoint.

- 6 Type a Red Hat OpenStack tenant name in the **OpenStack project** text box.
If you set up multiple endpoints with different Red Hat OpenStack tenants, create reservation policies for each tenant. This ensures that machines are provisioned to the appropriate tenant resources.
- 7 (Optional) Add any custom properties.
- 8 Click **OK**.

What to do next

Add the compute resources from your endpoint to a fabric group.

See [“Create a Fabric Group,”](#) on page 17.

Custom Properties for Openstack Endpoints

vRealize Automation includes custom properties you might want to use when you configure your Openstack endpoints in vRealize Automation.

Table 1-3. Custom Properties for Openstack Endpoints

Custom Property	Description
<code>VirtualMachine.Admin.ConnectAddress.Regex</code>	Used by a vRealize Automation administrator to define a regular expression to match an IP address for terminal connections, such as an RDP connection. If matched, the IP address is saved under the <code>VirtualMachine.Admin.ConnectAddress</code> custom property. Otherwise, the first available IP address is designated. For example, setting the property value to <code>10.10.0.</code> allows selection of an IP address from a <code>10.10.0.*</code> subnet that is assigned to the virtual machine. If the subnet has not been assigned, the property is ignored. This property is available for use with Openstack.
<code>VirtualMachine.NetworkN.AddressM</code>	Defines additional <i>M</i> IP address allocated for an Openstack instance for network <i>N</i> , excluding the IP address set specified by the <code>VirtualMachine.NetworkN.Address</code> property. Additional addresses are displayed on the Network tab in the Additional Addresses column. This property is used by Openstack machine state data collection. While this property is only data-collected by the OpenStack endpoint, it is not specific to OpenStack and can be used for lifecycle extensibility by other endpoint types.
<code>VMware.Endpoint.Openstack.IdentityProvider.Version</code>	For 6.2.4 and 6.2.5, specifies the version of Openstack Identity provider (Keystone) to use when authenticating an Openstack endpoint. Configure a value of 3 to authenticate with Keystone version 3 OpenStack Identity Provider. If you use any other value, or do not use this custom property, authentication defaults to Keystone version 2.
<code>VMware.Endpoint.Openstack.Release</code>	Specifies the OpenStack release, for example Havana or Icehouse, when creating an OpenStack endpoint. Required for 6.2.1, 6.2.2, and 6.2.3 OpenStack provisioning. Deprecated as of 6.2.4.

Preparing an Endpoint CSV File for Import

Instead of adding endpoints one at a time by using the vRealize Automation console, you can import a list of endpoints of various types by uploading a CSV file.

The CSV file must contain a header row with the expected fields. Fields are case sensitive and must be in a specific order. You can upload multiple endpoints of varying types with the same CSV file. For vCloud Director, system administrator accounts are imported, rather than organization administrator endpoints.

Table 1-4. CSV File Fields and Their Order for Importing Endpoints

Field	Description
InterfaceType	(Required) You can upload multiple types of endpoints in a single file. <ul style="list-style-type: none"> ■ AmazonEC2 ■ Openstack ■ vCloud ■ vCO ■ CiscoUCS ■ DellIldrac ■ HPilo ■ NetAppOnTap ■ SCVMM ■ KVM ■ vSphere
Address	(Required for all interface types except Amazon AWS) URL for the endpoint. For information about the required format for your platform type, see the appropriate procedure to create an endpoint for your platform.
Credentials	(Required) Name you gave the user credentials when you stored them in vRealize Automation.
Name	(Required) Provide a name for the endpoint. For RedHat Openstack, the address is used as the default name.
Description	(Optional) Provide a description for the endpoint.
DataCenter	(Optional) For physical machines, you can provide the datacenter where the machine is located.
Row	(Optional) For physical machines, you can provide the row where the machine is located.
Rack	(Optional) For physical machines, you can provide the rack where the machine is located.
OpenstackProject	(Required for RedHat Openstack only) Provide the tenant name for the endpoint.

Import a List of Endpoints

Importing a CSV file of endpoints can be more efficient than adding endpoints one at a time by using the vRealize Automation console.

Prerequisites

- Log in to the vRealize Automation console as an **IaaS administrator**.
- Store the credentials for your endpoints.
- Prepare an Endpoint CSV file for import.

Procedure

- 1 Select **Infrastructure > Endpoints > Endpoints**.
- 2 Click **Import Endpoints**.
- 3 Click **Browse**.
- 4 Locate the CSV file that contains your endpoints.
- 5 Click **Open**.

A CSV file opens that contains a list of endpoints in the following format:

```
InterfaceType,Address,Credentials,Name,Description
vCloud,https://abxpoint2vco,svc-admin,abxpoint2vco,abxpoint
```

- 6 Click **Import**.

You can edit and manage your endpoints through the vRealize Automation console.

Create a Fabric Group

An IaaS administrator can organize virtualization compute resources and cloud endpoints into fabric groups by type and intent. An IaaS administrator also assigns one or more fabric administrators to manage the resources in the fabric group.

You can grant the Fabric Administrator role to multiple users by either adding multiple users one at a time or by choosing an identity store group or custom group as your fabric administrator.

Prerequisites

- Log in to the vRealize Automation console as an **IaaS administrator**.
- Create at least one endpoint.

Procedure

- 1 Select **Infrastructure > Groups > Fabric Groups**.
- 2 Click **New Fabric Group**.
- 3 Enter a name in the **Name** text box.
- 4 (Optional) Enter a description in the **Description** text box.
- 5 Type a user name or group name in the **Fabric administrators** text box and press Enter.

Repeat this step to add multiple users or groups to the role.

- 6 Click one or more **Compute resources** to include in your fabric group.

Only templates that exist on the clusters you select for your fabric group are discovered during data collection and available for cloning on reservations you create for business groups.

- 7 Click **OK**.

Fabric administrators can now configure machine prefixes.

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 Machine Prefixes

Fabric administrators create machine prefixes that are used to create names for machines provisioned through vRealize Automation. Tenant administrators and business group managers select these machine prefixes and assign them to provisioned machines through blueprints and business group defaults.

A prefix is a base name to be followed by a counter of a specified number of digits. When the digits are all used, vRealize Automation rolls back to the first number.

Machine prefixes must conform to the following limitations:

- Contain only the case-insensitive ASCII letters a through z, the digits 0 through 9, and the hyphen (-).
- Not begin with a hyphen.
- No other symbols, punctuation characters, or blank spaces can be used.
- No longer than 15 characters, including the digits, to conform to the Windows limit of 15 characters in host names.



Longer host names are truncated when a machine is provisioned, and updated the next time data collection is run. However, for WIM provisioning names are not truncated and provisioning fails when the specified name is longer than 15 characters.

- vRealize Automation does not support multiple virtual machines of the same name in a single instance. If you choose a naming convention that causes an overlap in machine names, vRealize Automation does not provision a machine with the redundant name. If possible, vRealize Automation skips the name that is already in use and generates a new machine name using the specified machine prefix. If a unique name cannot be generated, provisioning fails.

Prerequisites

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

Procedure

- 1 Select **Infrastructure > Blueprints > Machine Prefixes**.
- 2 Click the **Add** icon (.
- 3 Enter the machine prefix in the **Machine Prefix** text box.
- 4 Enter the number of counter digits in the **Number of Digits** text box.
- 5 Enter the counter start number in the **Next Number** text box.
- 6 Click the **Save** icon (.

Tenant administrators can create business groups so that users can access vRealize Automation to request machines.

Create a Business Group


Tenant administrators create business groups to associate a set of services and resources to a set of users, often corresponding to a line of business, department, or other organizational unit. Users must belong to a business group to request machines.

To add multiple users to a business group role, you can add multiple individual users, or you can add multiple users at the same time by adding an identity store group or a custom group to a role. For example, you can create a custom group Sales Support Team and add that group to the support role. For information about creating custom groups, see *Tenant Administration*. You can also use existing identity store user groups. The users and groups you choose must be valid in the identity store.

Prerequisites

- Log in to the vRealize Automation console as a **tenant administrator**.
- Request a machine prefix from a fabric administrator. See [“Configure Machine Prefixes,”](#) on page 18.

Procedure

- 1 Select **Infrastructure > Groups > Business Groups**.
- 2 Click the **Add** icon ()
- 3 (Optional) Select an existing business group from the **Copy from existing group** drop-down box.
Data from the group you chose appears.
- 4 Enter a name in the **Name** text box.
- 5 (Optional) Enter a description in the **Description** text box.
- 6 Select a **Default machine prefix**.
- 7 (Optional) Type a default **Active directory container** for machines provisioned in this group and press Enter.

The Active Directory container is used only for WIM provisioning. Other provisioning methods require additional configuration to join provisioned machines to an AD container.
- 8 Type a user name or group name in the **Group manager role** search box and press Enter.

Repeat this step to add more than one user or group to the role. You do not have to specify users at this time. You can create empty business groups to populate later.
- 9 Type one or more user names or group names in the **Send manager emails to** text box and press Enter.

Multiple entries must be separated with commas. For example,
JoeAdmin@mycompany.com,WeiMgr@mycompany.com.
- 10 Add users to your business group.

Multiple entries must be separated by line breaks. You do not have to specify users at this time. You can create empty business groups to populate later.
 - a Type a user name or group name in the **Support role** search box and press Enter.
Repeat this step to add more than one user or group to the role.
 - b Type a user name or group name in the **User role** search box and press Enter.
Repeat this step to add more than one user or group to the role.
- 11 (Optional) Add any custom properties.
- 12 Click **OK**.

Fabric administrators can now allocate resources to your business group by creating a reservation.

Managing Key Pairs

Key pairs are used to provision and connect to a cloud instance. A key pair is used to decrypt Windows passwords or to log in to a Linux machine.

Key pairs are required for provisioning with Amazon AWS. For Red Hat OpenStack, key pairs are optional.

Existing key pairs are imported as part of data collection when you add a cloud endpoint. A fabric administrator can also create and manage key pairs by using the vRealize Automation console. If you delete a key pair from the vRealize Automation console, it is also deleted from the cloud service account.

In addition to managing key pairs manually, you can configure vRealize Automation to generate key pairs automatically per machine or per business group.

- A fabric administrator can configure the automatic generation of key pairs at a reservation level.
- If the key pair is going to be controlled at the blueprint level, the fabric administrator must select **Not Specified** on the reservation.
- A tenant administrator or business group manager can configure the automatic generation of key pairs at a blueprint level.
- If key pair generation is configured at both the reservation and blueprint level, the reservation setting overrides the blueprint setting.


Create a Key Pair

A fabric administrator can create key pairs for use with cloud endpoints by using the vRealize Automation console.

Prerequisites

- Log in to the vRealize Automation console as a **fabric administrator**.
- An IaaS administrator must have created a cloud endpoint and added cloud compute resources to a fabric group.

Procedure

- 1 Select **Infrastructure > Reservations > Key Pairs**.
- 2 Click **New Key Pair**.
- 3 Enter a name in the **Name** text box.
- 4 Select a cloud region from the **Compute resource** drop-down menu.
- 5 Click the **Save** icon ().

The key pair is ready to use when the Secret Key column has the value *****.


Upload the Private Key for a Key Pair

A fabric administrator can upload the private key for a key pair in PEM format.

Prerequisites

- Log in to the vRealize Automation console as a **fabric administrator**.
- You must already have a key pair. See [“Create a Key Pair,”](#) on page 20.

Procedure

- 1 Select **Infrastructure > Reservations > Key Pairs**.
- 2 Locate the key pair for which you want to upload a private key.
- 3 Click the **Edit** icon (.
- 4 Use one of the following methods to upload the key.
 - Browse for a PEM-encoded file and click **Upload**.
 - Paste the text of the private key, beginning with -----BEGIN RSA PRIVATE KEY----- and ending with -----END RSA PRIVATE KEY-----.

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


Export the Private Key from a Key Pair

A fabric administrator can export the private key from a key pair to a PEM-encoded file.

Prerequisites

- Log in to the vRealize Automation console as a **fabric administrator**.
- A key pair with a private key must exist. See [“Upload the Private Key for a Key Pair,”](#) on page 20.

Procedure

- 1 Select **Infrastructure > Reservations > Key Pairs**.
- 2 Locate the key pair from which you want to export the private key.
- 3 Click the **Export** icon ()
- 4 Browse to the location where you want to save the file and click **Save**.

Cloud Reservations

A cloud reservation provides access to the provisioning services of a cloud service account for a particular business group.

A group can have multiple reservations on one endpoint or reservations on multiple endpoints.

A reservation may also define policies, priorities, and quotas that determine machine placement.

Cloud Reservation Selection

When a member of a business group requests a cloud machine, a reservation must be selected from the reservations that belong to the business group.

The reservation on which a machine is provisioned must satisfy the following criteria:

- The reservation must be of the same platform type as the blueprint from which the machine was requested.
- The reservation must be enabled.
- The reservation must have capacity remaining in its machine quota or have an unlimited quota.
The allocated machine quota includes only machines that are powered on. For example, if a reservation has a quota of 50, and 40 machines have been provisioned but only 20 of them are powered on, the reservation’s quota is 40 percent allocated, not 80 percent.
- The reservation must have the security groups specified in the machine request.
- The reservation must be associated with a region that has the machine image specified in the blueprint.
- For Amazon machines, the request specifies an availability zone and whether the machine is to be provisioned a subnet in a Virtual Private Cloud (VPC) or in a non-VPC location. The reservation must match the network type (VPC or non-VPC).
- If the cloud provider supports network selection and the blueprint has specific network settings, the reservation must have the same networks.
If the blueprint or reservation specifies a network profile for static IP address assignment, an IP address must be available to assign to the new machine.
- If the blueprint specifies a reservation policy, the reservation must belong to that reservation policy.

Reservation policies are a way to guarantee that the selected reservation satisfies any additional requirements for provisioning machines from a specific blueprint. For example, if a blueprint uses a specific machine image, you can use reservation policies to limit provisioning to reservations associated with the regions that have the required image.

If no reservation is available that meets all of the selection criteria, provisioning fails.

If multiple reservations meet all of the criteria, the reservation from which to provision a requested machine is determined by the following logic:

- Reservations with higher priority are selected over reservations with lower priority.
- If multiple reservations have the same priority, the reservation with the lowest percentage of its machine quota allocated is selected.
- If multiple reservations have the same priority and quota usage, machines are distributed among reservations in round-robin fashion.

Choosing a Reservation Scenario

A fabric administrator creates reservations to allocate resources to business groups. Depending on your scenario, the procedure to create an endpoint differs.

Each business group must have at least one reservation for its members to provision machines of that type. For example, a business group with a Red Hat OpenStack reservation, but not an Amazon AWS reservation, cannot request a machine from Amazon AWS. The group must be allocated a reservation specifically for Amazon AWS resources.

Table 1-5. Choosing a Reservation Scenario

Scenario	Procedure
Create a reservation to allocate resources on Amazon AWS (without using Virtual Private Cloud)	“Create an Amazon AWS Reservation (non-VPC),” on page 22
Create a reservation to allocate resources on Amazon AWS using Virtual Private Cloud	“Create an Amazon AWS Reservation (VPC),” on page 25
Create a reservation to allocate resources on Red Hat OpenStack	“Create a Red Hat OpenStack Reservation,” on page 28

Create an Amazon AWS Reservation (non-VPC)

Fabric administrators must allocate resources to machines by creating a reservation before members of a business group can request machines.

Prerequisites

- Log in to the vRealize Automation console as a **fabric administrator**.
- A tenant administrator must create at least one business group. See [“Create a Business Group,”](#) on page 18.

Procedure

- 1 [Specify Amazon AWS Reservation Information \(non-VPC\)](#) on page 23
Each reservation is configured for a specific business group to grant them access to request machines on a specified compute resource.
- 2 [Specify Amazon AWS Key Pairs and Network Settings \(non-VPC\)](#) on page 23
Configure the resources and network paths available to machines that are provisioned by using this reservation.

3 [Configure Amazon AWS Alerts \(non-VPC\)](#) on page 24

You can configure alerts to send email notifications whenever reservation resources are low.

Specify Amazon AWS Reservation Information (non-VPC)

Each reservation is configured for a specific business group to grant them access to request machines on a specified compute resource.

NOTE After you create a reservation, you cannot change the business group or the compute resource.

Prerequisites

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

Procedure

- 1 Select **Infrastructure > Reservations > Reservations**.
- 2 Select **New Reservation > Cloud** and select the type of reservation you are creating.
Select **Amazon EC2**.
- 3 (Optional) Select an existing reservation from the **Copy from existing reservation** drop-down menu.
Data from the reservation you chose appears, and you can make changes as required for your new reservation.
- 4 Select a compute resource on which to provision machines from the **Compute resource** drop-down menu.
When you select an available compute resource, the **Name** field automatically populates.
- 5 Select a tenant from the **Tenant** drop-down menu.
- 6 Select a business group from the **Business group** drop-down menu.
Only users in this business group can provision machines by using this reservation.
- 7 Select a reservation policy from the **Reservation policy** drop-down menu.
- 8 (Optional) Type a number in the **Machine quota** text box to set the maximum number of machines that can be provisioned on this reservation.
Only machines that are powered on are counted towards the quota. Leave blank to make the reservation unlimited.
- 9 Type a number in the **Priority** text box to set the priority for the reservation.
The priority is used when a business group has more than one reservation. A reservation with priority 1 is used for provisioning over a reservation with priority 2.
- 10 (Optional) Deselect the **Enable this reservation** check box if you do not want this reservation active.
- 11 (Optional) Add any custom properties.

Do not navigate away from this page. Your reservation is not complete.

Specify Amazon AWS Key Pairs and Network Settings (non-VPC)

Configure the resources and network paths available to machines that are provisioned by using this reservation.

Prerequisites

[“Specify Amazon AWS Reservation Information \(non-VPC\),”](#) on page 23.

Procedure

- 1 Click the **Resources** tab.
- 2 Select a method of assigning key pairs to compute instances from the **Key pair** drop-down menu.

Option	Description
Not Specified	Controls key pair behavior at the blueprint level rather than the reservation level.
Auto-Generated per Business Group	Every machine provisioned in the same business group has the same key pair, including machines provisioned on other reservations when the machine has the same compute resource and business group. Because key pairs generated this way are associated with a business group, the key pairs are deleted when the business group is deleted.
Auto-Generated per Machine	Each machine has a unique key pair. This is the most secure method because no key pairs are shared among machines.
Specific Key Pair	Every machine provisioned on this reservation has the same key pair. Browse for a key pair to use for this reservation.

- 3 Select one or more available zones in the **Locations** list to make them available for this reservation.
Because Amazon machine images are region-specific, the Amazon machine images specified in a blueprint require that the requesting user select a location in the corresponding region. This allows vRealize Automation to select an appropriate reservation during provisioning.
- 4 Select one or more security groups that can be assigned to a machine during provisioning from the security groups list.
- 5 If Amazon elastic load balancing is enabled, select from the **Load balancers** list to apply to the selected locations.

Do not navigate away from this page. Your reservation is not complete.

Configure Amazon AWS Alerts (non-VPC)

You can configure alerts to send email notifications whenever reservation resources are low.

Alerts are an optional step in the reservation configuration. If you do not want to set alerts, click **OK** to save your reservation.

If configured, alerts are generated daily, rather than when the specified thresholds are reached.

IMPORTANT Notifications are only sent if emails are configured and notifications are enabled. See *Tenant Administration*.

Prerequisites

[“Specify Amazon AWS Key Pairs and Network Settings \(non-VPC\),”](#) on page 23.

Procedure

- 1 Click the **Alerts** tab.
- 2 Set capacity alerts to On.
- 3 Use the sliders to set thresholds for resource allocation.
- 4 Type one or more user email addresses or group names to receive alert notifications in the **Recipients** text box.
Press Enter to separate multiple entries.
- 5 Select **Send alerts to group manager** to include group managers in the email alerts.

- 6 Choose a reminder frequency (days).
- 7 Click **OK**.

Tenant administrators and business group managers can now create blueprints. You can configure optional policies such as reservation policies.

Create an Amazon AWS Reservation (VPC)

Fabric administrators must allocate resources to machines by creating a reservation before members of a business group can request machines them .

Amazon Web Services users can create a Amazon Virtual Private Cloud to design a virtual network topology according to your specifications. If you plan to use Amazon VPC, you must assign an Amazon VPC to a vRealize Automation reservation.

NOTE After you create a reservation, you cannot change the business group or the compute resource.

For information about creating an Amazon VPC by using the AWS Management Console, see Amazon Web Services documentation.

Prerequisites

- Log in to the vRealize Automation console as a **fabric administrator**.
- A tenant administrator must create at least one business group. See [“Create a Business Group,”](#) on page 18.
- Create an Amazon Virtual Private Cloud environment for use with vRealize Automation. See [“Using Virtual Private Cloud,”](#) on page 40.

Procedure

- 1 [Specify Amazon AWS with Amazon VPC Reservation Information](#) on page 25
Fabric administrators configure each reservation for a specific business group to grant them access to request machines on a specified compute resource.
- 2 [Specify Amazon AWS with Amazon VPC Key Pairs and Network Settings](#) on page 26
Configure the resources and network paths that are available to machines provisioned by using this reservation.
- 3 [Configure Amazon AWS with Amazon VPC Alerts](#) on page 27
Optionally, you can configure alerts to send email notifications whenever reservation resources are low.

Specify Amazon AWS with Amazon VPC Reservation Information

Fabric administrators configure each reservation for a specific business group to grant them access to request machines on a specified compute resource.

NOTE After you create a reservation, you cannot change the business group or the compute resource.

Prerequisites

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

Procedure

- 1 Select **Infrastructure > Reservations > Reservations**.

- 2 Select **New Reservation > Cloud** and select the type of reservation you are creating.
Select **Amazon EC2**.
- 3 (Optional) Select an existing reservation from the **Copy from existing reservation** drop-down menu.
Data from the reservation you chose appears, and you can make changes as required for your new reservation.
- 4 Select a compute resource on which to provision machines from the **Compute resource** drop-down menu.
When you select an available compute resource, the **Name** field automatically populates.
- 5 Select a tenant from the **Tenant** drop-down menu.
- 6 Select a business group from the **Business group** drop-down menu.
Only users in this business group can provision machines by using this reservation.
- 7 (Optional) Select a reservation policy from the **Reservation policy** drop-down menu.
This option requires additional configuration. You must create a reservation policy.
- 8 (Optional) Type a number in the **Machine quota** text box to set the maximum number of machines that can be provisioned on this reservation.
Only machines that are powered on are counted towards the quota. Leave blank to make the reservation unlimited.
- 9 Type a number in the **Priority** text box to set the priority for the reservation.
The priority is used when a business group has more than one reservation. A reservation with priority 1 is used for provisioning over a reservation with priority 2.
- 10 (Optional) Deselect the **Enable this reservation** check box if you do not want this reservation active.
- 11 (Optional) Add any custom properties.

Do not navigate away from this page. Your reservation is not complete.

Specify Amazon AWS with Amazon VPC Key Pairs and Network Settings

Configure the resources and network paths that are available to machines provisioned by using this reservation.

Prerequisites

[“Specify Amazon AWS with Amazon VPC Reservation Information,”](#) on page 25.

Procedure

- 1 Click the **Resources** tab.
- 2 Select a method of assigning key pairs to compute instances from the **Key pair** drop-down menu.

Option	Description
Not Specified	Controls key pair behavior at the blueprint level rather than the reservation level.
Auto-Generated per Business Group	Every machine provisioned in the same business group has the same key pair, including machines provisioned on other reservations when the machine has the same compute resource and business group. Because key pairs generated this way are associated with a business group, the key pairs are deleted when the business group is deleted.

Option	Description
Auto-Generated per Machine	Each machine has a unique key pair. This is the most secure method because no key pairs are shared among machines.
Specific Key Pair	Every machine provisioned on this reservation has the same key pair. Browse for a key pair to use for this reservation.

- 3 Select the **Assign to a subnet in a VPC** check box to open the Amazon VPC list.
- 4 Locate the Amazon VPC to assign and click the **Edit** icon (✎) to open the Edit VPC page.
 - a In the **Subnets** list, select each subnet in the Amazon VPC that you want to be available for provisioning.
 - b In the **Security Groups** list, select each group that can be assigned to a machine during provisioning.
 - c If you are using the elastic load balancer feature, select from the list of **Load Balancers** that apply to the selected subnets to use in the Amazon VPC.
 - d Click the **Save** icon.

Do not navigate away from this page. Your reservation is not complete.

For related information about security groups, see [“Using Security Groups for Amazon Web Services,”](#) on page 39.

For related information about load balancers, see [“Using Elastic Load Balancers,”](#) on page 41.

Configure Amazon AWS with Amazon VPC Alerts

Optionally, you can configure alerts to send email notifications whenever reservation resources are low.

Alerts are an optional step in the reservation configuration. If you do not want to set alerts, click **OK** to save your reservation.

If configured, alerts are generated daily, rather than when the specified thresholds are reached.

IMPORTANT Notifications are only sent if emails are configured and notifications are enabled. See *Tenant Administration*.

Prerequisites

[“Specify Amazon AWS with Amazon VPC Key Pairs and Network Settings,”](#) on page 26.

Procedure

- 1 Click the **Alerts** tab.
- 2 Set capacity alerts to On.
- 3 Use the sliders to set thresholds for resource allocation.
- 4 Type one or more user email addresses or group names to receive alert notifications in the **Recipients** text box.
Press Enter to separate multiple entries.
- 5 Select **Send alerts to group manager** to include group managers in the email alerts.
- 6 Choose a reminder frequency (days).
- 7 Click **OK**.

Tenant administrators and business group managers can create blueprints. You can configure optional policies such as reservation policies.

Create a Red Hat OpenStack Reservation

Before members of a business group can request machines, fabric administrators must allocate resources to them by creating a reservation.

Prerequisites

- Log in to the vRealize Automation console as a **fabric administrator**.
- A tenant administrator must create at least one business group. See [“Create a Business Group,”](#) on page 18.

Procedure

- 1 [Specify Reservation Information](#) on page 28
Each reservation is configured for a specific business group to grant them access to request machines on a specified compute resource.
- 2 [Specify Key Pairs and Network Settings](#) on page 29
Configure the key pairs, security groups, and networks available to machines provisioned through this reservation.
- 3 [Configure Alerts](#) on page 30
Optionally, you can configure alerts to send notifications whenever reservation resources are low.

Specify Reservation Information

Each reservation is configured for a specific business group to grant them access to request machines on a specified compute resource.

NOTE Once you create a reservation, you cannot change the business group or the compute resource.

Prerequisites

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

Procedure

- 1 Select **Infrastructure > Reservations > Reservations**.
- 2 Select **New Reservation > Cloud > OpenStack**.
- 3 (Optional) Select an existing reservation from the **Copy from existing reservation** drop-down menu.
Data from the reservation you chose appears, and you can make changes as required for your new reservation.
- 4 Select a compute resource on which to provision machines from the **Compute resource** drop-down menu.
The reservation name appears in the **Name** text box.
- 5 Enter a name in the **Name** text box.
- 6 Select a tenant from the **Tenant** drop-down menu.
- 7 Select a business group from the **Business group** drop-down menu.
Only users in this business group can provision machines by using this reservation.

- 8 (Optional) Select a reservation policy from the **Reservation policy** drop-down menu.
This option requires additional configuration. You must create a reservation policy.
- 9 (Optional) Type a number in the **Machine quota** text box to set the maximum number of machines that can be provisioned on this reservation.
Only machines that are powered on are counted towards the quota. Leave blank to make the reservation unlimited.
- 10 Type a number in the **Priority** text box to set the priority for the reservation.
The priority is used when a business group has more than one reservation. A reservation with priority 1 is used for provisioning over a reservation with priority 2.
- 11 (Optional) Deselect the **Enable this reservation** check box if you do not want this reservation active.
- 12 (Optional) Add any custom properties.

Do not navigate away from this page. Your reservation is not complete.

Specify Key Pairs and Network Settings

Configure the key pairs, security groups, and networks available to machines provisioned through this reservation.

Prerequisites

[“Specify Reservation Information,”](#) on page 28.

Procedure

- 1 Click the **Resources** tab.
- 2 Select a method of assigning key pairs to compute instances from the **Key pair** drop-down menu.

Option	Description
Not Specified	Controls key pair behavior at the blueprint level rather than the reservation level.
Auto-Generated per Business Group	Every machine provisioned in the same business group has the same key pair, including machines provisioned on other reservations when the machine has the same compute resource and business group. Because key pairs generated this way are associated with a business group, the key pairs are deleted when the business group is deleted.
Auto-Generated per Machine	Each machine has a unique key pair. This is the most secure method because no key pairs are shared among machines.
Specific Key Pair	Every machine provisioned on this reservation has the same key pair. Browse for a key pair to use for this reservation.

- 3 Select one or more security groups that can be assigned to a machine during provisioning from the security groups list.
- 4 Click the **Network** tab.
- 5 Configure a network path for machines provisioned by using this reservation.
 - a Select a network path for machines provisioned on this reservation from the Network table.
 - b (Optional) Select a network profile from the **Network Profile** drop-down menu.
This option requires additional configuration to configure network profiles.

You can select more than one network path on a reservation, but only one network is selected when provisioning a machine.

At this point, you can save the reservation by clicking **OK**. Optionally, you can configure email notifications to send alerts out when resources allocated to this reservation become low.

Configure Alerts

Optionally, you can configure alerts to send notifications whenever reservation resources are low.

Alerts are an optional step in the reservation configuration. If you do not want to set alerts, click **OK** to save your reservation.

If configured, alerts are generated daily, rather than when the specified thresholds are reached.

IMPORTANT Notifications are only sent if emails are configured and notifications are enabled. See *Tenant Administration*.

Prerequisites

[“Specify Key Pairs and Network Settings,”](#) on page 29.

Procedure

- 1 Click the **Alerts** tab.
- 2 Set capacity alerts to On.
- 3 Use the sliders to set thresholds for resource allocation.
- 4 Type one or more user email addresses or group names to receive alert notifications in the **Recipients** text box.
Press Enter to separate multiple entries.
- 5 Select **Send alerts to group manager** to include group managers in the email alerts.
- 6 Choose a reminder frequency (days).
- 7 Click **OK**.

Tenant administrators and business group managers can create blueprints. You can configure optional policies such as reservation policies.

Optional Configurations

You can configure optional policies and settings to give you more control over the resource usage of provisioned machines.

This chapter includes the following topics:

- [“Build Profiles and Property Sets,”](#) on page 31
- [“Configuring Reservation Policies,”](#) on page 34
- [“Static IP Address Assignment,”](#) on page 36
- [“Cost Information for Cloud Machines,”](#) on page 39
- [“Using Optional Amazon Features,”](#) on page 39
- [“Using Optional Red Hat OpenStack Features,”](#) on page 42

Build Profiles and Property Sets

vRealize Automation contains property sets that fabric administrators can use when they create build profiles.

Property sets are groups of related properties that are commonly used together in build profiles and machine blueprints. Instead of adding custom properties to a build profile or a machine blueprint individually, you can add an entire set to a build profile and provide the values.

Many commonly used property sets are included in vRealize Automation.

For example, the `WimImagingProperties` property set contains custom properties commonly used for WIM-based provisioning:

- `Image.ISO.Location`
- `Image.ISO.Name`
- `Image.Network.Password`
- `Image.Network.User`
- `Image.WIM.Index`
- `Image.WIM.Name`
- `Image.WIM.Path`

As another example, the `vApp` property set contains the following custom properties that can be used for vApp provisioning:

- `VirtualMachine.NetworkN.Name`

- VCloud.Template.MakeIdenticalCopy
- VMware.SCSI.Type
- Sysprep.Identification.DomainAdmin
- Sysprep.Identification.DomainAdminPassword
- Sysprep.Identification.JoinDomain

Fabric administrators can create their own property sets and add them to vRealize Automation.

Create a Property Set

Fabric administrators can create their own groupings of related custom properties and add them to vRealize Automation for use in build profiles.

Property sets are available to fabric administrators of all tenants.

Procedure

- 1 [Create a Property Set XML File](#) on page 32
A fabric administrator defines property sets in an XML file and uploads them to vRealize Automation.
- 2 [Add a Property Set to vRealize Automation](#) on page 33
After you create a property set XML file, a fabric administrator can upload the property set to vRealize Automation.

Create a Property Set XML File

A fabric administrator defines property sets in an XML file and uploads them to vRealize Automation.

If you edit a property set that is already in use in a build profile, vRealize Automation does not automatically update the build profile. A fabric administrator must reload the property set to the build profile.

Procedure

- 1 Create an XML file.
- 2 Insert the following version and encoding values into the schema declaration.
version="1.0" encoding="UTF-16".
- 3 Insert a <Doc> element.

```
<Doc>
</Doc>
```
- 4 Insert a <CustomProperties> element in the <Doc> element.

```
<Doc>
  <CustomProperties>
  </CustomProperties>
</Doc>
```


- 5 Define the attributes of the custom property to include in the property set.

```
<Doc>
  <CustomProperties>
    <Property Name="property_name" DefaultValue="property_value" Encrypted="true_or_false"
    PromptUser="true_or_false"/>
  </CustomProperties>
</Doc>
```

If you do not include the DefaultValue attribute, no default value is stored. If you do not include the Encrypted or PromptUser attributes, they default to false.

- 6 Repeat [Step 5](#) step for each property to include in the property set.
- 7 Save and close the file.

A fabric administrator can now upload your property set XML file to vRealize Automation.

Add a Property Set to vRealize Automation

After you create a property set XML file, a fabric administrator can upload the property set to vRealize Automation.

Prerequisites

- Log in to the vRealize Automation console as a **fabric administrator**.
- Create a property set XML file.

Procedure

- 1 Select **Infrastructure > Blueprints > Build Profiles**.
- 2 Click **Manage Property Sets**.
- 3 Click **Browse** to select the **Property set XML file name**.
- 4 Enter a name and, optionally, a description.
- 5 Click **OK**.

You can now include your property set in build profiles.

Create a Build Profile


Fabric administrators can organize commonly used custom properties into build profiles so that tenant administrators and business group managers can include these custom property sets in blueprints.

Build profiles are available to tenant administrators and business group managers of all tenants. You can create your build profile by adding custom properties individually, by loading property sets that contain multiple custom properties, or by using a combination of the two methods.

Prerequisites



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

Procedure

- 1 Select **Infrastructure > Blueprints > Build Profiles**.
- 2 Click the **Add** icon ()
- 3 Enter a name and, optionally, a description.


- 4 (Optional) Add custom properties individually to your build profile.
 - a Click **New Property**.
 - b Enter the custom property in the **Name** text box.
 - c Enter the value of the custom property in the **Value** text box.
 - d (Optional) Select the **Encrypted** check box to encrypt the custom property in the database.
 - e (Optional) Select the **Prompt user** check box to require the user to provide a value when they request a machine.

If you choose to prompt users for a value, any value you provide for the custom property is presented to them as the default. If you do not provide a default, users cannot continue with the machine request until they provide a value for the custom property.

- f Click the **Save** icon ()
- 5 (Optional) Select a property set to load into your build profile.
 - a Select a property set from the **Add from property set** drop-down menu.
 - b Click **Load**.
- 6 (Optional) Click the **Edit** icon () to configure a custom property loaded from a property set.

- a Enter the value of the custom property in the **Value** text box.
 - b Select the **Encrypted** check box to encrypt the custom property in the database.
 - c Select the **Prompt user** check box to require the user to provide a value when they request a machine.

If you choose to prompt users for a value, any value you provide for the custom property is presented to them as the default. If you do not provide a default, users cannot continue with the machine request until they provide a value for the custom property.

- d Click the **Save** icon ()
- 7 Click **OK**.

Tenant administrators and business group managers can now select your build profile and include it in their blueprints.

Configuring Reservation Policies

When a user requests a machine, it can be provisioned on any reservation of the appropriate type that has sufficient capacity for the machine. You can apply a reservation policy to a blueprint to restrict the machines provisioned from that blueprint to a subset of available reservations.

Reservation policies provide an optional and helpful means of controlling how reservation requests are processed. A reservation policy is often used to collect resources into groups for different service levels, or to make a specific type of resource easily available for a particular purpose. The following scenarios provide a few examples of possible uses for reservation policies:

- To ensure that machines provisioned from a virtual blueprint are placed on reservations with storage devices that support NetApp FlexClone
- To restrict provisioning of cloud machines to a specific region containing a machine image that is required for a specific blueprint
- To restrict provisioning of Cisco UCS physical machines to reservations on endpoints on which the selected service profile template and boot policy are available

- As an additional means of using a Pay As You Go allocation model for vApps

You can add multiple reservations to a reservation policy, but a reservation can belong to only one policy. You can assign a single reservation policy to more than one blueprint. A blueprint can have only one reservation policy.

A reservation policy can include reservations of different types, but only reservations that match the blueprint type are considered when selecting a reservation for a particular request. For more information about how reservations are selected for provisioning a machine, see *IaaS Configuration for Cloud Platforms*, *IaaS Configuration for Physical Machines*, or *IaaS Configuration for Virtual Platforms*.

Configure a Reservation Policy

Fabric administrators create reservation policies to collect resources into groups for different service levels, or to make a specific type of resource easily available for a particular purpose. After you create the reservation policy, you then must populate it with reservations before tenant administrators and business group managers can use the policy effectively in a blueprint.

Prerequisites

- Log in to the vRealize Automation console as a **fabric administrator**.
- Create at least one reservation.

Procedure

- 1 [Create a Reservation Policy](#) on page 35

Fabric administrators use reservation policies to group similar reservations together. Create the reservation policy tag first, then add the policy to reservations to allow a tenant administrator or business group manager to use the reservation policy in a blueprint.

- 2 [Assign a Reservation Policy to a Reservation](#) on page 36

When fabric administrators create a reservation, the option to assign that reservation to a reservation policy appears. To assign an existing reservation to a reservation policy, you edit the reservation.


Create a Reservation Policy

Fabric administrators use reservation policies to group similar reservations together. Create the reservation policy tag first, then add the policy to reservations to allow a tenant administrator or business group manager to use the reservation policy in a blueprint.

Prerequisites

- Log in to the vRealize Automation console as a **fabric administrator**.
- Create at least one reservation.

Procedure

- 1 Select **Infrastructure > Reservations > Reservation Policies**.
- 2 Click **New Reservation Policy**.
- 3 Enter a name and, optionally, a description.
- 4 Click the **Save** icon ()

Assign a Reservation Policy to a Reservation

When fabric administrators create a reservation, the option to assign that reservation to a reservation policy appears. To assign an existing reservation to a reservation policy, you edit the reservation.

Prerequisites

[“Create a Reservation Policy,”](#) on page 35.

Procedure

- 1 Select **Infrastructure > Reservations > Reservations**.
- 2 Point to a reservation and click **Edit**.
- 3 Select your reservation policy from the **Reservation Policy** drop-down menu.
- 4 Click **OK**.

Tenant administrators and business group managers can now use your reservation policy in blueprints. See [“Add a Reservation Policy to a Blueprint,”](#) on page 80.

Static IP Address Assignment

For virtual machines provisioned by using cloning or Linux kickstart/autoYaST provisioning and cloud machines provisioned in Red Hat OpenStack by using kickstart, it is possible to assign static IP addresses from a predefined range.

By default, vRealize Automation uses Dynamic Host Configuration Protocol (DHCP) to assign IP addresses to provisioned machines.

Fabric administrators can create network profiles to define a range of static IP addresses that can be assigned to machines. Network profiles can be assigned to specific network paths on a reservation. Any cloud machine or virtual machine provisioned by cloning or kickstart/autoYaST that is attached to a network path that has an associated network profile is provisioned using static IP address assignment.

Tenant administrators or business group managers can also assign network profiles to blueprints by using the custom property *VirtualMachine.NetworkN.ProfileName*. If a network profile is specified in both the blueprint and the reservation, the profile specified in the blueprint takes precedence.

When a machine that has a static IP address is destroyed, its IP address is made available for use by other machines. The process to reclaim static IP addresses runs every 30 minutes, so unused addresses may not be available immediately after the machines using them are destroyed. If there are not available IP addresses in the network profile, machines cannot be provisioned with static IP assignment on the associated network path.

Create a Network Profile for Static IP Address Assignment

Fabric administrators can create network profiles to define a range of static IP addresses that can be assigned to machines.

Prerequisites

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

Procedure

- 1 [Specify Network Profile Information](#) on page 37
The network profile information identifies the external network profile and specifies settings for an existing network.
- 2 [Configure IP Ranges](#) on page 37
A fabric administrator can define one or more ranges of static IP addresses for use in provisioning a machine.

What to do next

You can assign a network profile to a network path in a reservation, or a tenant admin or business group manager can specify the network profile in a blueprint.

Specify Network Profile Information

The network profile information identifies the external network profile and specifies settings for an existing network.

Prerequisites

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

Procedure

- 1 Select **Infrastructure > Reservations > Network Profiles**.
- 2 Select **New Network Profile > External**.
- 3 Enter a name and, optionally, a description.
- 4 Type a mask address in the **Subnet mask** text box.
For example, **255.255.0.0**.
- 5 (Optional) Type an IP address in the **Gateway** text box.
The gateway address is required for a one-to-one NAT network profile.
- 6 (Optional) In the DNS/WINS group, type values as needed.
The external network profile provides these values, which you can edit.

What to do next

The network profile is not finished. Do not navigate away from this page.

Configure IP Ranges

A fabric administrator can define one or more ranges of static IP addresses for use in provisioning a machine.

Prerequisites

[“Specify Network Profile Information,”](#) on page 37.

Procedure

- 1 Click the **IP Ranges** tab.
- 2 Click **New Network Range**.
The New Network Range dialog box appears.
- 3 Enter a name and, optionally, a description.

- 4 Enter an IP address in the **Starting IP address** text box.
- 5 Enter an IP address in the **Ending IP address** text box.
- 6 Click **OK**.

The newly defined IP address range appears in the Defined Ranges list. The IP addresses in the range appear in the Defined IP Addresses list.

- 7 (Optional) Upload one or more IP addresses from a CSV file.

A row in the CSV file has the format *ip_address,mname,status*.

CSV Field	Description
<i>ip_address</i>	An IP address
<i>mname</i>	Name of a managed machine in vRealize Automation. If the field is empty, defaults to no name.
<i>status</i>	Allocated or Unallocated, case-sensitive. If the field is empty, defaults to Unallocated.

- a Click **Browse** next to the **Upload CSV** text box.
- b Navigate to the CSV file and click **Open**.
- c Click **Process CSV File**.

The uploaded IP addresses appear in the Defined IP Addresses list. If the upload fails, diagnostic messages appear that identify the problems.

- 8 (Optional) Filter IP address entries to only those that match.
 - a Click in the **Defined IP Addresses** text boxes.
 - b Type a partial IP address or machine name, or select a date from the Last Modified drop-down calendar.

The IP addresses that match the filter criteria appear.

- 9 Click **OK**.

What to do next

You can assign a network profile to a network path in a reservation, or a tenant admin or business group manager can specify the network profile in a blueprint.

Assign a Network Profile to a Reservation

Fabric administrators can assign a network profile to a reservation to enable static IP assignment to machines provisioned on that reservation.

Network profiles can also be assigned to blueprints by using the custom property *VirtualMachine.NetworkN.ProfileName*. If a network profile is specified in both the blueprint and the reservation, the profile specified in the blueprint takes precedence.

NOTE This information does not apply to Amazon Web Services.

Prerequisites

- Log in to the vRealize Automation console as a **fabric administrator**.
- [“Create a Network Profile for Static IP Address Assignment,”](#) on page 36.

Procedure

- 1 Select **Infrastructure > Reservations > Reservations**.
- 2 Point to a reservation and click **Edit**.

- 3 Click the **Network** tab.
- 4 Assign a network profile to a network path.
 - a Locate a network path on which you want to enable static IP addresses.
 - b Select a network profile from the **Network Profile** drop-down menu.
 - c (Optional) Repeat this step to assign network profiles to additional network paths on this reservation.
- 5 Click **OK**.

What to do next

Repeat this procedure for all reservations for which you want to enable static IP addresses.

Cost Information for Cloud Machines

Cost profiles do not apply to cloud machines. The only factor in calculating cost for cloud machines is the blueprint cost; this figure should reflect all cloud service charges.

Blueprints with more than one instance type or flavors allow provisioning of instances with different costs from the same blueprint. When there is more than one instance type or flavor available, ensure that the daily blueprint cost reflects charges for the most expensive instance.

Lease cost is calculated as daily cost multiplied by the total number of days in the lease period (if applicable).

Cost-to-date is calculated as daily cost multiplied by the number of days a machine has been provisioned.

Machine cost is displayed at various stages of the request and provisioning life cycle and is updated according to the current information in the request or on the provisioned item.

Table 2-1. Cost Displayed During the Request and Provisioning Life Cycle

Life Cycle Stage	Value Displayed for Cost
Viewing the catalog item details prior to request	Daily cost based on the blueprint cost and projected lease cost based on the lease duration specified in the blueprint. The lease cost may vary if the blueprint specifies a range for the lease duration.
Completing the request form, viewing details of a submitted request, or approving a request	Daily cost based on the blueprint cost and projected lease cost based on the lease duration specified in the request.
Viewing the details of a provisioned machine	Daily cost based on the blueprint cost, lease cost based on the lease duration, and cost-to-date based on the number of days since the machine was provisioned.

Using Optional Amazon Features

vRealize Automation supports several Amazon features, including Amazon Virtual Private Cloud, elastic load balancers, elastic IP addresses, and elastic block storage.

Using Security Groups for Amazon Web Services

A security group acts as a firewall to control access to a cloud machine.

A fabric administrator specifies at least one security group when creating an Amazon EC2 reservation. Each available region requires at least one specified security group. The reservation form displays the security groups that are available to an Amazon Web Services account region. Every region includes at least the default security group.

Administrators can use the Amazon Web Services Management Console to create additional security groups, configure ports for Microsoft Remote Desktop Protocol or SSH, and set up a virtual private network for an Amazon VPN.

For information about creating and using security groups in Amazon Web Services, see Amazon documentation.

Understanding Amazon Web Service Regions

Each Amazon Web Services account is represented by a cloud endpoint. When you create an Amazon Elastic Cloud Computing endpoint in vRealize Automation, regions are collected as compute resources. After the IaaS administrator selects compute resources for a business group, inventory and state data collections occur automatically.

Inventory data collection, which occurs automatically once a day, collects data about what is on a compute resource, such as the following data:

- Elastic IP addresses
- Elastic load balancers
- Elastic block storage volumes

State data collection occurs automatically every 15 minutes by default. It gathers information about the state of managed instances, which are instances that vRealize Automation creates. The following are examples of state data:

- Windows passwords
- State of machines in load balancers
- Elastic IP addresses

A fabric administrator can initiate inventory and state data collection and disable or change the frequency of inventory and state data collection.

Using Virtual Private Cloud

Amazon Virtual Private Cloud allows you to provision Amazon machine instances in a private section of the Amazon Web Services cloud.

Amazon AWS users can create a Amazon VPC to design a virtual network topology according to your specifications. You can assign a Amazon VPC in vRealize Automation. However, vRealize Automation does not track the cost of using the Amazon VPC.

Use the AWS Management Console to create the following elements:

- An Amazon VPC, which includes Internet gateways, routing table, security groups and subnets, and available IP addresses.
- An Amazon Virtual Private Network if users will need to log in to Amazon machines instances outside of the AWS Management Console.

vRealize Automation users can perform the following tasks when working with an Amazon VPC:

- A fabric administrator can assign an Amazon VPC to a cloud reservation. See [“Create an Amazon AWS Reservation \(VPC\),”](#) on page 25.
- A machine owner can assign an Amazon machine instance to an Amazon VPC.

For more information about creating an Amazon VPC, see Amazon Web Services documentation.

Using Elastic Load Balancers

Elastic load balancers distribute incoming application traffic across Amazon Web Services instances. Amazon load balancing enables improved fault tolerance and performance.

Amazon makes elastic load balancing available for machines provisioned using Amazon EC2 blueprints.

The elastic load balancer must be available in the Amazon Web Services, Amazon Virtual Private Network and at the provisioning location. For example, if a load balancer is available in us-east1c and a machine location is us-east1b, the machine cannot use the available load balancer.

vRealize Automation does not create, manage, or monitor the elastic load balancers.

For information about creating Amazon elastic load balancers by using the Amazon Web Services Management Console, see Amazon Web Services documentation.

Using Elastic IP Addresses

An elastic IP address is designed for quick fail-over in a dynamic Amazon Web Services cloud environment.

Using an elastic IP address allows you to rapidly fail over to another machine. In vRealize Automation, the elastic IP address is available to all business groups that have rights to the region.

The elastic IP address is associated with your Amazon Web Services account, not a particular machine. The address remains associated with your Amazon Web Services account until you choose to release it. You can release it to map it to a specific machine instance.

An administrator can allocate elastic IP addresses to your Amazon Web Services account by using the AWS Management Console. One range of addresses is allocated to a non-Amazon VPC region and another range to a Amazon VPC region.

A tenant administrator or business group manager can use a custom property to assign an elastic IP address to a machine during provisioning. A machine owner or an administrator with rights to edit the machine can assign an elastic IP addresses to a machines after it is provisioned. However, if the address is already associated to a machine instance, and the instance is part of the Amazon Virtual Private Cloud deployment, Amazon does not assign the address.

A machine owner can view the elastic IP address assigned to that machine. Only one Amazon machine at a time can use the elastic IP address.

vRealize Automation does not track the cost of using the elastic IP address.

There are two groups of elastic IP addresses in any given a region, one group for non-Amazon VPC instances and another group for Amazon VPCs. If you allocate addresses in a non-Amazon VPC region only, the addresses will not be available in an Amazon VPC. The reverse is also true. If you allocate addresses in an Amazon VPC only, the addresses will not be available in a non-Amazon VPC region.

For more information about creating and using Amazon elastic IP addresses, see Amazon Web Services documentation.

Using Elastic Block Storage

Amazon elastic block storage provides block level storage volumes to use with an Amazon AWS machine instance and Amazon Virtual Private Cloud. The storage volume can persist past the life of its associated Amazon machine instance in the Amazon Web Services cloud environment.

When you use an Amazon elastic block storage volume in conjunction with vRealize Automation, the following caveats apply:

- You cannot attach an existing elastic block storage volume when you provision a machine instance. However, if you create a new volume and request more than one machine at a time, the volume will be created and attached to each instance. For example, if you create one volume named `volume_1` and request three machines, a volume will be created for each machine. Three volumes named `volume_1` will be created and attached to each machine. Each volume will have a unique volume ID. Each volume will be the same size and in the same location.
- The volume must be of the same operating system and in the same location as the machine to which you attach it.
- vRealize Automation does not track the cost of using an existing elastic block storage volume.
- vRealize Automation does not manage the primary volume of an elastic block storage-backed instance.

For more information about Amazon elastic block storage, and details on how to enable it by using Amazon Web Services Management Console, see Amazon Web Services documentation.

Using Optional Red Hat OpenStack Features

vRealize Automation supports several features in Red Hat OpenStack including security groups and floating IP addresses.

Using Security Groups

Security groups allow you to specify rules to control network traffic over specific ports.

A fabric administrator specifies security groups when creating a Red Hat OpenStack reservation. A business group manager can specify security groups when requesting a machine.

Each region includes at least the default security group. Additional security groups must be managed in Red Hat OpenStack. Security groups for each region are imported during data collection. For more information about managing security groups, see the Red Hat OpenStack documentation.

Using Floating IP Addresses

Floating IP addresses can be assigned to a running virtual instance in Red Hat OpenStack.

To enable assignment of floating IP addresses, you must configure IP forwarding and create a floating IP pool in Red Hat OpenStack. For more information, see the Red Hat OpenStack documentation.

A tenant administrator or business group manager must entitle the Associate Floating IP and Disassociate Floating IP actions to machine owners. The entitled users can then associate a floating IP address to a provisioned machine from the external networks attached to the machine by selecting an available address from the floating IP address pool. After a floating IP address has been associated with a machine, the user can select Disassociate Floating IP to view the currently assigned floating IP addresses and disassociate an address from a machine.

Preparing for Provisioning

There is some configuration required outside of vRealize Automation to prepare for machine provisioning.

This chapter includes the following topics:

- [“Choosing a Cloud Provisioning Scenario,”](#) on page 43
- [“Preparing for Amazon Provisioning,”](#) on page 44
- [“Preparing for OpenStack Provisioning,”](#) on page 46

Choosing a Cloud Provisioning Scenario

There are several methods you can use with vRealize Automation to provision cloud machines.

Most of these options require some additional configuration outside of vRealize Automation.

Table 3-1. Choosing a Cloud Provisioning Method

Scenario	Supported Platforms	Provisioning Method
Launch an instance from an Amazon Machine Image.	Amazon AWS	Amazon Machine Image See “Preparing for Amazon Provisioning,” on page 44.
Launch an instance from a virtual machine image.	Red Hat OpenStack	virtual machine image See “Preparing for Virtual Machine Image Provisioning,” on page 46.
Provision a machine by booting from an ISO image, using a kickstart or autoYaSt configuration file and a Linux distribution image to install the operating system on the machine.	Red Hat OpenStack	Linux Kickstart See “Preparing for Linux Kickstart Provisioning,” on page 47.
Provision a machine by booting into a WinPE environment and installing an operating system using a Windows Imaging File Format (WIM) image of an existing Windows reference machine.	Red Hat OpenStack	WIM Image See “Preparing for WIM Provisioning,” on page 49.

Preparing for Amazon Provisioning

Before you provision instances with Amazon AWS, you must have Amazon machine images and instance types associated with your Amazon AWS account.

The vRealize Automation access rights that you need to perform tasks such as creating endpoints, reservations, and blueprints, and requesting machine provisioning are described for the applicable vRealize Automation tasks in this guide. However, you must also be configured in Amazon Web Services (AWS) to perform related endpoint tasks.

- Role and Permission Authorization in Amazon Web Services

The Power User role in AWS provides an AWS Directory Service user or group with full access to AWS services and resources.

You do not need any AWS credentials to create an AWS endpoint in vRealize Automation. However, the AWS user who creates an Amazon machine image is expected by vRealize Automation to have the Power User role.

- Authentication Credentials in Amazon Web Services

The AWS Power User role does not allow management of AWS Identity and Access Management (IAM) users and groups. For management of IAM users and groups, you must be configured with AWS Full Access Administrator credentials.

vRealize Automation requires access keys for endpoint credentials and does not support user names and passwords. To obtain the access key needed to create the Amazon endpoint, the Power User must either request a key from a user who has AWS Full Access Administrator credentials or be additionally configured with the AWS Full Access Administrator policy.

For information about enabling policies and roles, see the *AWS Identity and Access Management (IAM)* section of Amazon Web Services product documentation.

Understanding Amazon Machine Images

A tenant administrator or business group manager selects an Amazon machine image from a list of available images when creating Amazon EC2 blueprints.

An Amazon machine image is a template that contains a software configuration, including an operating system. They are managed by Amazon Web Services accounts. vRealize Automation manages the instance types that are available for provisioning.

The Amazon machine image and instance type must be available in an Amazon region. Not all instance types are available in all regions.

You can select an Amazon machine image provided by Amazon Web Services, a user community, or the AWS Marketplace site. You can also create and optionally share your own Amazon machine images. A single Amazon machine image can be used to launch one or many instances.

The following considerations apply to Amazon machine images in the Amazon Web Services accounts from which you will provision cloud machines:

- Each cloud blueprint must specify an Amazon machine image.

An Amazon machine image can be enabled for some locations in region. A private Amazon machine image is available to a specific account and all its regions. A public Amazon machine image is available to all accounts, but only to a specific region in each account.

- When a cloud blueprint is created, the specified Amazon machine image is selected from regions that have been data-collected. If multiple Amazon Web Services accounts are available, the business group manager must have rights to any private Amazon machine images. The Amazon machine image region and the specified user location restrict provisioning request to reservations that match the corresponding region and location.
- Use reservations and policies to distribute Amazon machine images in your Amazon Web Services accounts. Use policies to restrict provisioning from a blueprint to a particular set of reservations.
- vRealize Automation cannot create user accounts on a cloud machine. The first time a machine owner connects to a cloud machine, she must log in as an administrator and add her vRealize Automation user credentials or an administrator must do that for her. She can then log in using her vRealize Automation user credentials.

If the Amazon machine image generates the administrator password on every boot, the Edit Machine Record page displays the password. If it does not, you can find the password in the Amazon Web Services account. You can configure all Amazon machine images to generate the administrator password on every boot. You can also provide administrator password information to support users who provision machines for other users.

- To allow remote Microsoft Windows Management Instrumentation (WMI) requests on cloud machines provisioned in Amazon Web Services accounts, enable a Microsoft Windows Remote Management (WinRM) agent to collect data from Windows machines managed by vRealize Automation. See *Installation and Configuration*.
- A private Amazon machine image can be seen across tenants.

For related information, see *Amazon Machine Images (AMI)* topics in Amazon documentation.

Understanding Amazon Instance Types

A tenant administrator or business group manager selects one or more Amazon instance types when creating Amazon EC2 blueprints. An IaaS administrator can add or remove them.

An Amazon EC2 instance is a virtual server that can run applications in Amazon Web Services. Instances are created from an Amazon machine image and by choosing an appropriate instance type.

To provision a machine in an Amazon Web Services account, an instance type is applied to the specified Amazon machine image. The available instance types are listed on the Amazon EC2 blueprint.

When creating an Amazon EC2 blueprint, you identify one or more instance types to be available for selection by the user when they provision a machine. The instance types must be supported in the designated region.

For related information, see *Selecting Instance Types* and *Amazon EC2 Instance Details* topics in Amazon documentation.

Add an Amazon Instance Type


Several instance types are supplied with vRealize Automation for use with Amazon EC2 blueprints. An IaaS administrator can add and remove instance types.

The machine instance types managed by IaaS administrators on the Instance Types page are available to tenant administrators and business group manager when they create or edit an Amazon EC2 blueprint. Amazon machine images and instance types are made available through the Amazon Web Services product.

Prerequisites

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

Procedure

- 1 Click **Infrastructure > Blueprints > Instance Types**.
- 2 Click **New Instance Type**.
- 3 Add a new instance type, specifying the following parameters and referencing *Selecting Instance Types* and *Amazon EC2 Instance Details* topics in Amazon documentation.
 - Name
 - API name
 - Type Name
 - IO Performance Name
 - # CPUs
 - Memory (GB)
 - Storage (GB)
 - Compute Units
- 4 Click the **Save** icon ()

Preparing for OpenStack Provisioning

Before you provision instances with Red Hat OpenStack using virtual machine images, Linux kickstart, or WIM provisioning you must prepare the appropriate machine images and flavors.

Preparing for Virtual Machine Image Provisioning

Before you provision instances with Red Hat OpenStack, you must have virtual machine images and flavors configured in the Red Hat OpenStack provider.

Virtual Machine Images

A tenant administrator or business group manager selects an virtual machine image from a list of available images when creating Red Hat OpenStack blueprints.

A virtual machine image is a template that contains a software configuration, including an operating system. Virtual machine images are managed within the Red Hat OpenStack provider and are imported during data collection.

If an image that is used in a blueprint is later deleted from the Red Hat OpenStack provider, it is also removed from the blueprint. If all the images have been removed from a blueprint, the blueprint is disabled and cannot be used for machine requests until it is edited to add at least one image.

OpenStack Flavors

A tenant administrator or business group manager selects one or more flavors when creating Red Hat OpenStack blueprints.

Flavors are virtual hardware templates that define the machine resource specifications for instances provisioned in Red Hat OpenStack. Flavors are managed within the Red Hat OpenStack provider and are imported during data collection.

If a flavor that is used in a blueprint is later deleted from the Red Hat OpenStack provider, it is also removed from the blueprint. If all the flavors have been removed from a blueprint, the blueprint is disabled and cannot be used for machine requests until it is edited to add at least one flavor.

Preparing for Linux Kickstart Provisioning

Linux Kickstart provisioning uses a configuration file to automate a Linux installation on a newly provisioned machine. To prepare for provisioning you must create a bootable ISO image and a kickstart or autoYaST configuration file.

The following is a high-level overview of the steps required to prepare for Linux Kickstart provisioning:

- 1 Ensure that a DHCP server is available on the network. vRealize Automation cannot provision machines by using Linux Kickstart provisioning unless DHCP is available.
- 2 Prepare the configuration file. In the configuration file, you must specify the locations of the vRealize Automation server and the Linux agent installation package. See [“Prepare the Linux Kickstart Configuration Sample File,”](#) on page 47.
- 3 Edit the `isolinux/isolinux.cfg` or `loader/isolinux.cfg` to specify the name and location of the configuration file and the appropriate Linux distribution source.
- 4 Create the boot ISO image and save it to the location required by your virtualization platform. Please see the documentation provided by your hypervisor if you do not know the required location.
- 5 (Optional) Add customization scripts.
 - a To specify post-installation customization scripts in the configuration file, see [“Specify Custom Scripts in a kickstart/autoYaST Configuration File,”](#) on page 48.
 - b To call Visual Basic scripts in blueprint, see [“Enabling Visual Basic Scripts in Provisioning,”](#) on page 84.
- 6 Gather the following information so that tenant administrators and business group managers can include it in their blueprints:
 - a The name and location of the ISO image.
 - b For vCenter Server integrations, the vCenter Server guest operating system version with which vCenter Server is to create the machine.

NOTE You can create a build profile with the property set `BootIsoProperties` to include the required ISO information. This makes it easier for tenant administrators and business group managers to include this information correctly in their blueprints.

Prepare the Linux Kickstart Configuration Sample File

vRealize Automation provides sample configuration files that you can modify and edit to suit your needs. There are several changes required to make the files usable.

Procedure

- 1 Navigate to the vCloud Automation Center Appliance management console installation page.
For example: `https://vcac-hostname.domain.name:5480/installer/`.
- 2 Download and save the Linux Guest Agent Packages.
- 3 Unpack the `LinuxGuestAgentPkgs` file.
- 4 Navigate to the `LinuxGuestAgentPkgs` file and locate the subdirectory that corresponds to the guest operating system that you are deploying during provisioning.
- 5 Open the `sample-https.cfg` file.

- 6 Replace all instances of the string `host=dcac.example.net` with the IP address or fully qualified domain name and port number for the vRealize Automation server host.

Platform	Required Format
vSphere ESXi	IP Address, for example: <code>--host=172.20.9.59</code>
vSphere ESX	IP Address, for example: <code>--host=172.20.9.58</code>
SUSE 10	IP Address, for example: <code>--host=172.20.9.57</code>
All others	FQDN, for example: <code>--host=mycompany-host1.mycompany.local:443</code>

- 7 Locate each instance of `gugent.rpm` or `gugent.tar.gz` and replace the URL `rpm.example.net` with the location of the guest agent package.

For example:

```
rpm -i nfs:172.20.9.59/suseagent/gugent.rpm
```

- 8 Save the file to a location accessible to newly provisioned machines.

Specify Custom Scripts in a kickstart/autoYaST Configuration File

You can modify the configuration file to copy or install custom scripts onto newly provisioned machines. The Linux agent runs the scripts at the specified point in the workflow.

Your script can reference any of the `./properties.xml` files in the `/usr/share/gugent/site/workitem` directories.

Prerequisites

- Prepare a kickstart or autoYaST configuration file. See [“Prepare the Linux Kickstart Configuration Sample File,”](#) on page 47.
- Your script must return a non-zero value on failure to prevent machine provisioning failure.

Procedure

- 1 Create or identify the script you want to use.
- 2 Save the script as `NN_scriptname`.
NN is a two digit number. Scripts are executed in order from lowest to highest. If two scripts have the same number, the order is alphabetical based on *scriptname*.
- 3 Make your script executable.
- 4 Locate the post-installation section of your kickstart or autoYaST configuration file.
In kickstart, this is indicated by `%post`. In autoYaST, this is indicated by `post-scripts`.
- 5 Modify the post-installation section of the configuration file to copy or install your script into the `/usr/share/gugent/site/workitem` directory of your choice.

Custom scripts are most commonly run for virtual kickstart/autoYaST with the work items SetupOS (for create provisioning) and CustomizeOS (for clone provisioning), but you can run scripts at any point in the workflow.

For example, you can modify the configuration file to copy the script `11_addusers.sh` to the `/usr/share/gugent/site/SetupOS` directory on a newly provisioned machine by using the following command:

```
cp nfs:172.20.9.59/linuxscripts/11_addusers.sh /usr/share/gugent/site/SetupOS
```

The Linux agent runs the script in the order specified by the work item directory and the script file name.

Preparing for WIM Provisioning

Provision a machine by booting into a WinPE environment and then install an operating system using a Windows Imaging File Format (WIM) image of an existing Windows reference machine.

The following is a high-level overview of the steps required to prepare for WIM provisioning:

- 1 Identify or create the staging area. This should be a network directory that can be specified as a UNC path or mounted as a network drive by the reference machine, the system on which you build the WinPE image, and the virtualization host on which machines are provisioned.
- 2 Ensure that a DHCP server is available on the network. vRealize Automation cannot provision machines by using a WIM image unless DHCP is available.
- 3 Identify or create the reference machine within the virtualization platform you intend to use for provisioning. For vRealize Automation requirements, see [“Reference Machine Requirements for WIM Provisioning,”](#) on page 50. For information about creating a reference machine, please see the documentation provided by your hypervisor.
- 4 Using the System Preparation Utility for Windows, prepare the reference machine's operating system for deployment. See [“SysPrep Requirements for the Reference Machine,”](#) on page 50.
- 5 Create the WIM image of the reference machine. Do not include any spaces in the WIM image file name or provisioning fails.
- 6 (Optional) Create any custom scripts you want to use to customize provisioned machines and place them in the appropriate work item directory of your PEBuilder installation. See [“Specify Custom Scripts in a PEBuilder WinPE,”](#) on page 52.
- 7 Create a WinPE image that contains the vRealize Automation guest agent. You can use the vRealize Automation PEBuilder to create the WinPE image. See [“Create a WinPE Image by Using PEBuilder,”](#) on page 52. You can create the WinPE image by using another method, but you must manually insert the vRealize Automation guest agent. See [“Manually Insert the Guest Agent into a WinPE Image,”](#) on page 53.
- 8 Place the WinPE image in the location required by your virtualization platform. If you do not know the location, please see the documentation provided by your hypervisor.
- 9 Gather the following information so that tenant administrators and business group managers can include it in their blueprints:
 - a The name and location of the WinPE ISO image.
 - b The name of the WIM file, the UNC path to the WIM file, and the index used to extract the desired image from the WIM file.
 - c The user name and password under which to map the WIM image path to a network drive on the provisioned machine.
 - d (Optional) If you do not want to accept the default, K, the drive letter to which the WIM image path is mapped on the provisioned machine.
 - e For vCenter Server integrations, the vCenter Server guest operating system version with which vCenter Server is to create the machine.
 - f (Optional) For SCVMM integrations, the ISO, virtual hard disk, or hardware profile to attach to provisioned machines.

NOTE You can create a build profile with the property set `WimImagingProperties` to include all of this required information. This makes it easier for tenant administrators and business group managers to include this information correctly in their blueprints. Optionally, you can also add the property set `SysprepProperties` and prompt tenant administrators and business group managers to specify or overwrite settings in the SysPrep answer file.

Reference Machine Requirements for WIM Provisioning

WIM provisioning involves creating a WIM image from a reference machine. The reference machine must meet basic requirements for the WIM image to work for provisioning in vRealize Automation.

The following is a high-level overview of the steps to prepare a reference machine:

- 1 If the operating system on your reference machine is Windows Server 2008 R2, Windows Server 2012, Windows 7, or Windows 8, the default installation creates a small partition on the system's hard disk in addition to the main partition. vRealize Automation does not support the use of WIM images created on such multi-partitioned reference machines. You must delete this partition during the installation process.
- 2 Install NET 4.5 and Windows Automated Installation Kit (AIK) for Windows 7 (including WinPE 3.0) on the reference machine.
- 3 If the reference machine operating system is Windows Server 2003 or Windows XP, reset the administrator password to be blank. (There is no password.)
- 4 (Optional) If you want to enable XenDesktop integration, install and configure a Citrix Virtual Desktop Agent.
- 5 (Optional) A Windows Management Instrumentation (WMI) agent is required to collect certain data from a Windows machine managed by vRealize Automation, for example the Active Directory status of a machine's owner. To ensure successful management of Windows machines, you must install a WMI agent (typically on the Manager Service host) and enable the agent to collect data from Windows machines. See *Installation and Configuration*.

SysPrep Requirements for the Reference Machine

A SysPrep answer file is required for WIM provisioning and there are a number of required settings.

Table 3-2. Windows Server 2003 or Windows XP reference machine SysPrep required settings:

GuiUnattended Settings	Value
AutoLogon	Yes
AutoLogonCount	1
AutoLogonUsername	<i>username</i> (<i>username</i> and <i>password</i> are the credentials used for auto logon when the newly provisioned machine boots into the guest operating system. Administrator is typically used.)
AutoLogonPassword	<i>password</i> corresponding to the AutoLogonUsername.

Table 3-3. Required SysPrep Settings for reference machine that are not using Windows Server 2003 or Windows XP:

AutoLogon Settings	Value
Enabled	Yes
LogonCount	1

Table 3-3. Required SysPrep Settings for reference machine that are not using Windows Server 2003 or Windows XP: (Continued)

AutoLogon Settings	Value
Username	<i>username</i> (<i>username</i> and <i>password</i> are the credentials used for auto logon when the newly provisioned machine boots into the guest operating system. Administrator is typically used.)
Password	<i>password</i> (<i>username</i> and <i>password</i> are the credentials used for auto logon when the newly provisioned machine boots into the guest operating system. Administrator is typically used.) NOTE For reference machines that use a Windows platform newer than Windows Server 2003/Windows XP, you must set the autologon password by using the custom property <code>Sysprep.GuiUnattended.AdminPassword</code> . A convenient way to ensure this is done is to create a build profile that includes this custom property so that tenant administrators and business group managers can include this information correctly in their blueprints.

Install PEBuilder

The PEBuilder tool provided by vRealize Automation provides a simple way to include the vRealize Automation guest agent in your WinPE images.

PEBuilder has a 32 bit guest agent. If you need to run commands specific to 64 bit, install PEBuilder and then get the 64 bit files from the `GugentZipx64.zip` file.

Install PEBuilder in a location where you can access your staging environment.

Prerequisites

- Install NET Framework 4.5.
- Windows Automated Installation Kit (AIK) for Windows 7 (including WinPE 3.0) is installed.

Procedure

- 1 Navigate to the vCloud Automation Center Appliance management console installation page.
For example: `https://vcac-hostname.domain.name:5480/installer/`.
- 2 Download the PEBuilder.
- 3 (Optional) Download the Windows 64-bit guest agent package if you want to include the Windows 64-bit guest agent in your WinPE instead of the Windows 32-bit guest agent.
- 4 Run `vcac-WinPEBuilder-Setup.exe`.
- 5 Follow the prompts to install PEBuilder.
- 6 (Optional) Replace the Windows 32-bit guest agent files located in `\PE Builder\Plugins\VRM Agent\VRMGuestAgent` with the 64-bit files to include the 64-bit agent in your WinPE.

You can use PEBuilder to create a WinPE for use in WIM provisioning.

Specify Custom Scripts in a PEBuilder WinPE

You can use PEBuilder to customize machines by running custom bat scripts at specified points in the provisioning workflow.

Prerequisites

[“Install PEBuilder,”](#) on page 51.

Procedure

- 1 Create or identify the bat script you want to use.
Your script must return a non-zero value on failure to prevent machine provisioning failure.
- 2 Save the script as *NN_scriptname*.
NN is a two digit number. Scripts are executed in order from lowest to highest. If two scripts have the same number, the order is alphabetical based on *scriptname*.
- 3 Make your script executable.
- 4 Place the scripts in the work item subdirectory that corresponds to the point in the provisioning workflow you want the script to run.

For example, C:\Program Files (x86)\VMware\VRA\PE Builder\Plugins\VRM Agent\VRMGuestAgent\site\SetupOS.

The agent runs the script in the order specified by the work item directory and the script file name.

Create a WinPE Image by Using PEBuilder

Use the PEBuilder tool provided by vRealize Automation to create a WinPE ISO file that includes the vRealize Automation guest agent.

Prerequisites

- [“Install PEBuilder,”](#) on page 51.
- (Optional) Configure PEBuilder to include the Windows 64-bit guest agent in your WinPE instead of the Windows 32-bit guest agent. See [“Install PEBuilder,”](#) on page 51.
- (Optional) Add any third party plugins you want to add to the WinPE image to the PlugIns subdirectory of the PEBuilder installation directory.
- (Optional) [“Specify Custom Scripts in a PEBuilder WinPE,”](#) on page 52.

Procedure

- 1 Run PEBuilder.

- 2 Enter the IaaS Manager Service host information.

Option	Description
If you are using a load balancer	<ol style="list-style-type: none"> a Enter the fully qualified domain name of the load balancer for the IaaS Manager Service in the vCAC Hostname text box. For example, manager_service_LB.mycompany.com. b Enter the port number for the IaaS Manager Service load balancer in the vCAC Port text box. For example, 443.
With no load balancer	<ol style="list-style-type: none"> a Enter the fully qualified domain name of the IaaS Manager Service machine in the vCAC Hostname text box. For example, manager_service.mycompany.com. b Enter the port number for the IaaS Manager Service machine in the vCAC Port text box. For example, 443.

- 3 Enter the path to the PEBuilder plugins directory.

This depends on the installation directory specified during installation. The default is C:\Program Files (x86)\VMware\VCAC\PE Builder\PlugIns.

- 4 Enter the output path for the ISO file you are creating in the **ISO Output Path** text box.

This location should be on the staging area you prepared.

- 5 Click **File > Advanced**.

NOTE Do not change the **WinPE Architecture** or **Protocol** settings.

- 6 Select the **Include vCAC Guest Agent in WinPE ISO** check box.

- 7 Click **OK**.

- 8 Click **Build**.

What to do next

Place the WinPE image in the location required by your integration platform. If you do not know the location, please see the documentation provided by your platform.

Manually Insert the Guest Agent into a WinPE Image

You do not have to use the vRealize Automation PEBuilder to create your WinPE. However, if you do not use the PEBuilder you must manually insert the vRealize Automation guest agent into your WinPE image.

Prerequisites

- Select a Windows system from which the staging area you prepared is accessible and on which .NET 4.5 and Windows Automated Installation Kit (AIK) for Windows 7 (including WinPE 3.0) are installed.
- Create a WinPE.

Procedure

- 1 [Install the Guest Agent in a WinPE](#) on page 54

If you choose not to use the vRealize Automation PEBuilder to create your WinPE, you must install PEBuilder to manually copy the guest agent files to your WinPE image.

- 2 [Configure the doagent.bat File](#) on page 54

If you choose not to use the vRealize Automation PEBuilder, you must manually configure the `doagent.bat` file.

3 [Configure the doagentc.bat File](#) on page 55

If you choose not to use the vRealize Automation PEBuilder, you must manually configure the `doagentc.bat` file.

4 [Configure the Guest Agent Properties Files](#) on page 56

If you choose not to use the vRealize Automation PEBuilder, you must manually configure the guest agent properties files.

Install the Guest Agent in a WinPE

If you choose not to use the vRealize Automation PEBuilder to create your WinPE, you must install PEBuilder to manually copy the guest agent files to your WinPE image.

PEBuilder has a 32-bit guest agent. If you need to run commands specific to 64-bit, install PEBuilder and then get the 64-bit files from the `GugentZipx64.zip` file.

Prerequisites

- Select a Windows system from which the staging area you prepared is accessible and on which .NET 4.5 and Windows Automated Installation Kit (AIK) for Windows 7 (including WinPE 3.0) are installed.
- Create a WinPE.

Procedure

- 1 Navigate to the vCloud Automation Center Appliance management console installation page.
For example: `https://vcac-hostname.domain.name:5480/installer/`.
- 2 Download the PEBuilder.
- 3 (Optional) Download the Windows 64-bit guest agent package if you want to include the Windows 64-bit guest agent in your WinPE instead of the Windows 32-bit guest agent.
- 4 Execute `vcac-WinPEBuilder-Setup.exe`.
- 5 Deselect both **Plugins** and **PEBuilder**.
- 6 Expand **Plugins** and select **VRMAgent**.
- 7 Follow the prompts to complete the installation.
- 8 (Optional) After installation is complete, replace the Windows 32-bit guest agent files located in `\PE Builder\Plugins\VRM Agent\VRMGuestAgent` with the 64-bit files to include the 64-bit agent in your WinPE.
- 9 Copy the contents of `%SystemDrive%\Program Files (x86)\VMware\PE Builder\Plugins\VRM Agent\VRMGuestAgent` to a new location within your WinPE Image.
For example: `C:\Program Files (x86)\VMware\PE Builder\Plugins\VRM Agent\VRMGuestAgent`.

Configure the doagent.bat File

If you choose not to use the vRealize Automation PEBuilder, you must manually configure the `doagent.bat` file.

Prerequisites

[“Install the Guest Agent in a WinPE,”](#) on page 54.

Procedure

- 1 Navigate to the `VRMGuestAgent` directory within your WinPE Image.
For example: `C:\Program Files (x86)\VMware\PE Builder\Plugins\VRM Agent\VRMGuestAgent`.
- 2 Make a copy of the file `doagent-template.bat` and name it `doagent.bat`.

- 3 Open `doagent.bat` in a text editor.
- 4 Replace all instances of the string `#Dcac Hostname#` with the fully qualified domain name and port number of the IaaS Manager Service host.

Option	Description
If you are using a load balancer	Enter the fully qualified domain name and port of the load balancer for the IaaS Manager Service. For example, <code>manager_service_LB.mycompany.com:443</code>
With no load balancer	Enter the fully qualified domain name and port of the machine on which the IaaS Manager Service is installed. For example, <code>manager_service.mycompany.com:443</code>

- 5 Replace all instances of the string `#Protocol#` with the string `/ssl`.
- 6 Replace all instances of the string `#Comment#` with `REM` (`REM` must be followed by a trailing space).
- 7 (Optional) If you are using self-signed certificates, uncomment the `openssl` command.

```
echo QUIT | c:\VRMGuestAgent\bin\openssl s_client -connect
```
- 8 Save and close the file.
- 9 Edit the `Startnet.cmd` script for your WinPE to include the `doagent.bat` as a custom script.

Configure the `doagentc.bat` File

If you choose not to use the vRealize Automation PEBuilder, you must manually configure the `doagentc.bat` file.

Prerequisites

[“Configure the `doagentc.bat` File,”](#) on page 54.

Procedure

- 1 Navigate to the `VRMGuestAgent` directory within your WinPE Image.
 For example: `C:\Program Files (x86)\VMware\PE Builder\Plugins\VRM Agent\VRMGuestAgent`.
- 2 Make a copy of the file `doagentsvc-template.bat` and name it `doagentc.bat`.
- 3 Open `doagentc.bat` in a text editor.
- 4 Remove all instance of the string `#Comment#`.
- 5 Replace all instances of the string `#Dcac Hostname#` with the fully qualified domain name and port number of the Manager Service host.

The default port for the Manager Service is 443.

Option	Description
If you are using a load balancer	Enter the fully qualified domain name and port of the load balancer for the Manager Service. For example, <code>load_balancer_manager_service.mycompany.com:443</code>
With no load balancer	Enter the fully qualified domain name and port of the Manager Service. For example, <code>manager_service.mycompany.com:443</code>

- 6 Replace all instances of the string `#errorlevel#` with the character `1`.
- 7 Replace all instances of the string `#Protocol#` with the string `/ssl`.
- 8 Save and close the file.

Configure the Guest Agent Properties Files

If you choose not to use the vRealize Automation PEBuilder, you must manually configure the guest agent properties files.

Prerequisites

[“Configure the doagentc.bat File,”](#) on page 55.

Procedure

- 1 Navigate to the VRMGuestAgent directory within your WinPE Image.
For example: C:\Program Files (x86)\VMware\PE Builder\Plugins\VRM Agent\VRMGuestAgent.
- 2 Make a copy of the file gument.properties and name it gument.properties.template.
- 3 Make a copy of the file gument.properties.template and name it gumentc.properties.
- 4 Open gument.properties in a text editor.
- 5 Replace all instances of the string GuestAgent.log the string X:/VRMGuestAgent/GuestAgent.log.
- 6 Save and close the file.
- 7 Open gumentc.properties in a text editor.
- 8 Replace all instances of the string GuestAgent.log the string C:/VRMGuestAgent/GuestAgent.log.
- 9 Save and close the file.

Creating a Cloud Blueprint

Machine blueprints determine the attributes of the machine, the manner in which it is provisioned, and its policy and management settings. A tenant administrator or business group manager creates blueprints for provisioning machines.

This chapter includes the following topics:

- [“Choosing a Blueprint Scenario,”](#) on page 57
- [“Create an Amazon AWS Blueprint,”](#) on page 58
- [“Create a Basic Red Hat OpenStack Blueprint,”](#) on page 62
- [“Create a Red Hat OpenStack Blueprint for Linux Kickstart Provisioning,”](#) on page 65
- [“Create an OpenStack Blueprint for WIM Provisioning,”](#) on page 71
- [“Publish a Blueprint,”](#) on page 78

Choosing a Blueprint Scenario

Depending on your environment and the methods of provisioning your fabric administrators have prepared, there are several procedures available to create the blueprint for your needs.

Unless otherwise noted in the Choosing a Blueprint Scenario table, all of these provisioning methods require preparation by your fabric administrators before you can create a blueprint and provision machines. Please contact your fabric administrators to determine which provisioning methods they have prepared for you.

Table 4-1. Choosing a Blueprint Scenario

Scenario	Supported Platforms	Procedure	Custom Properties
Launch an instance from an Amazon Machine Image.	Amazon AWS	“Create an Amazon AWS Blueprint,” on page 58	No custom properties are required. A load balancer can be configured before or after install.
Launch an instance from a virtual machine image.	Red Hat OpenStack	“Create a Basic Red Hat OpenStack Blueprint,” on page 62	No custom properties are required.

Table 4-1. Choosing a Blueprint Scenario (Continued)

Scenario	Supported Platforms	Procedure	Custom Properties
Provision a machine by booting from an ISO image, using a kickstart or autoYaSt configuration file and a Linux distribution image to install the operating system on the machine.	Red Hat OpenStack	“Create a Red Hat OpenStack Blueprint for Linux Kickstart Provisioning,” on page 65	“Custom Properties for Linux Kickstart Blueprints,” on page 69
Provision a machine by booting into a WinPE environment and installing an operating system using a Windows Imaging File Format (WIM) image of an existing Windows reference machine.	Red Hat OpenStack	“Create an OpenStack Blueprint for WIM Provisioning,” on page 71	“Custom Properties for WIM Blueprints,” on page 75

Create an Amazon AWS Blueprint

Machine blueprints determine the attributes of the machine, the manner in which it is provisioned, and its policy and management settings. A tenant administrator or business group manager creates a blueprint for provisioning machines.

Prerequisites

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

Procedure

- 1 [Specify Amazon AWS Blueprint Information](#) on page 58
Tenant administrators or business group managers use blueprint information settings to control who can access a blueprint, how many machines they can provision with it, and daily costs.
- 2 [Specify Amazon AWS Blueprint Build Information](#) on page 59
Tenant administrators or business group managers use build information settings to determine how a machine is provisioned.
- 3 [Configure Amazon Machine Resources](#) on page 60
Machine resource settings control the resources consumed by a machine provisioned by using your blueprint.
- 4 [Add Amazon AWS Blueprint Custom Properties](#) on page 61
Adding custom properties to a blueprint gives you detailed control over provisioned machines.
- 5 [Configure Amazon AWS Blueprint Actions](#) on page 61
Tenant administrators or business group managers use blueprint actions and entitlements to maintain detailed control over provisioned machines.

What to do next

The blueprint is not available as a catalog item until you publish it. See [“Publish a Blueprint,”](#) on page 78.

Specify Amazon AWS Blueprint Information

Tenant administrators or business group managers use blueprint information settings to control who can access a blueprint, how many machines they can provision with it, and daily costs.

Prerequisites

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

Procedure

- 1 Select **Infrastructure > Blueprints > Blueprints**.
- 2 Select **New Blueprint > Cloud > Amazon EC2**.
- 3 Enter a name and, optionally, a description.
- 4 (Optional) Select the **Master** check box to allow users to copy your blueprint.
- 5 Select the **Display location on request** check box to prompt users to select a location when they submit a machine request.

This option enables users to select a region into which to provision the requested machine.

- 6 Select who can provision machines with this blueprint.

Roles	Who Can Provision
If you are both a business group manager and a tenant administrator	<ul style="list-style-type: none"> ■ Select the Shared blueprint check box to allow the blueprint to be entitled to users in any business group. ■ Deselect the Shared blueprint check box to create a local blueprint, and select a business group from the Business group drop-down menu.
Business group manager	Select a business group from the Business group drop-down menu.
Tenant administrator	Your blueprints are always shared. You cannot choose who can use them to provision machines.

- 7 (Optional) Select a reservation policy from the **Reservation policy** drop-down menu.
This option requires additional configuration by a fabric administrator to create a reservation policy.
- 8 Select a machine prefix from the **Machine prefix** drop-down menu.
You can select **Use group default** to accept the default machine prefix business group for the user.
- 9 (Optional) Enter a number in the **Maximum per user** text box to limit the number of machines that a single user can provision with this blueprint.
- 10 Set the daily cost of the machine by typing the amount in the **Cost (daily)** text box.

Your blueprint is not finished. Do not navigate away from this page.

Specify Amazon AWS Blueprint Build Information

Tenant administrators or business group managers use build information settings to determine how a machine is provisioned.

Prerequisites

[“Specify Amazon AWS Blueprint Information,”](#) on page 58.

Procedure

- 1 Click the **Build Information** tab.
- 2 Select whether machines provisioned from this blueprint are classified as Desktop or Server from the **Blueprint type** drop-down menu.
This information is for record-keeping and licensing purposes only.
- 3 Select **CloudProvisioningWorkflow** from the **Provisioning workflow** drop-down menu.
- 4 Click the **Browse** icon next to the **Amazon machine image** text box.
- 5 Select an Amazon machine image from the list.

- 6 Click **OK**.
- 7 Select a key pair generation method from the **Key pair** drop-down menu.

Option	Description
Not Specified	Select to control key pair behavior at the reservation level rather than the blueprint level.
Auto Generated per Business Group	Each machine provisioned in the same business group has the same key pair. If you delete the business group, its key pair is also deleted.
Auto Generated per Machine	Each machine has a unique key pair.

- 8 (Optional) Select the **Enable Amazon network options on machine** check box to allow users to choose whether to provision a machine in a virtual private cloud (VPC) or non-VPC location when they submit the request.

Your blueprint is not finished. Do not navigate away from this page.

Configure Amazon Machine Resources

Machine resource settings control the resources consumed by a machine provisioned by using your blueprint.

You cannot add or edit Amazon instance types after you create the blueprint.

Prerequisites

[“Specify Amazon AWS Blueprint Build Information,”](#) on page 59.

Procedure

- 1 Select one or more **Instance types** check boxes to select one or more Amazon instance types that users can select from when they submit a machine request.

Selecting one or more instance types updates values in the **# CPUs**, **Memory**, and **Storage (GB)** text boxes. Review the minimum and maximum machine resource values and adjust your instance type selections as desired.
- 2 (Optional) Specify the EBS storage settings for provisioned machines.
 - a Enter the minimum amount of storage in the **Minimum** text box.

If you only provide a minimum, this number becomes the value for all machines provisioned from this blueprint.
 - b (Optional) Enter the maximum amount of storage in the **Maximum** text box to allow users to select their own settings within the range that you provide.
- 3 (Optional) Specify the lease settings for provisioned machines, or leave blank for no expiration date.
 - a Enter the minimum number of lease days in the **Minimum** text box.

If you only provide a minimum, this number becomes the value for all machines provisioned from this blueprint.
 - b (Optional) Enter the maximum number of lease days in the **Maximum** text box to allow users to select their own settings within the range that you provide.

Your blueprint is not finished. Do not navigate away from this page.

Add Amazon AWS Blueprint Custom Properties

Adding custom properties to a blueprint gives you detailed control over provisioned machines.

For information about custom properties for Amazon, see *Custom Properties Reference*.

Prerequisites

[“Configure Amazon Machine Resources,”](#) on page 60.

Procedure

- 1 Click the **Properties** tab.
- 2 (Optional) Select one or more build profiles from the **Build profiles** menu.
Build profiles contain groups of custom properties. Fabric administrators can create build profiles.
- 3 (Optional) Add any custom properties to your blueprint.

- a Click **New Property**.
- b Enter the custom property in the **Name** text box.
- c Enter the value of the custom property in the **Value** text box.
- d (Optional) Select the **Encrypted** check box to encrypt the custom property in the database.
- e (Optional) Select the **Prompt user** check box to require the user to provide a value when they request a machine.

If you choose to prompt users for a value, any value you provide for the custom property is presented to them as the default. If you do not provide a default, users cannot continue with the machine request until they provide a value for the custom property.

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

Your blueprint is not finished. Do not navigate away from this page.

Configure Amazon AWS Blueprint Actions

Tenant administrators or business group managers use blueprint actions and entitlements to maintain detailed control over provisioned machines.

Entitlements control which machine operations specific users can access. Blueprint actions control which machine operations can be performed on machines provisioned from a blueprint. For example, if you disable the option to reprovision machines created from a blueprint, then the option to reprovision a machine created from the blueprint does not appear for any users. If you enable the reprovision machine operation, then users who are entitled to reprovision machines can reprovision machines created from this blueprint.

NOTE The options that appear on the **Actions** tab depend on your platform and any customizations made to your vRealize Automation instance.

Prerequisites

[“Add Amazon AWS Blueprint Custom Properties,”](#) on page 61.

Procedure

- 1 Click the **Actions** tab.

- 2 (Optional) Select the check boxes for each machine option to enable for machines provisioned from this blueprint.
- 3 Click **OK**.

Your blueprint is saved in draft state.

What to do next

Publish your blueprint to make it available as a catalog item. See [“Publish a Blueprint,”](#) on page 78.

Create a Basic Red Hat OpenStack Blueprint

Machine blueprints determine the attributes of the machine, the manner in which it is provisioned, and its policy and management settings. A tenant administrator or business group manager creates a blueprint for provisioning machines.

Prerequisites

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

Procedure

- 1 [Specify Basic Red Hat OpenStack Blueprint Information](#) on page 62
Blueprint information settings control who can access a blueprint, how many machines they can provision with it, and the daily cost.
- 2 [Specify Basic Red Hat OpenStack Blueprint Build Information](#) on page 63
Build information settings determine how a machine is provisioned and control the resources consumed by a machine provisioned from your blueprint.
- 3 [Add Basic Red Hat OpenStack Blueprint Custom Properties](#) on page 64
Adding custom properties to a blueprint gives you detailed control over provisioned machines.
- 4 [Configure Basic Red Hat OpenStack Blueprint Actions](#) on page 65
Use blueprint actions and entitlements together to maintain detailed control over provisioned machines.

What to do next

The blueprint is not available as a catalog item until you publish it. See [“Publish a Blueprint,”](#) on page 78.

Specify Basic Red Hat OpenStack Blueprint Information

Blueprint information settings control who can access a blueprint, how many machines they can provision with it, and the daily cost.

Prerequisites

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

Procedure

- 1 Select **Infrastructure > Blueprints > Blueprints**.
- 2 Select **New Blueprint > Cloud > OpenStack**.
- 3 Enter a name and, optionally, a description.
- 4 (Optional) Select the **Master** check box to allow users to copy your blueprint.

- 5 Select the **Display location on request** check box to prompt users to select a location when they submit a machine request.

This option enables users to select a region into which to provision the requested machine.

- 6 Select who can provision machines with this blueprint.

Roles	Who Can Provision
If you are both a business group manager and a tenant administrator	<ul style="list-style-type: none"> ■ Select the Shared blueprint check box to allow the blueprint to be entitled to users in any business group. ■ Deselect the Shared blueprint check box to create a local blueprint, and select a business group from the Business group drop-down menu.
Business group manager	Select a business group from the Business group drop-down menu.
Tenant administrator	Your blueprints are always shared. You cannot choose who can use them to provision machines.

- 7 (Optional) Select a reservation policy from the **Reservation policy** drop-down menu.
This option requires additional configuration by a fabric administrator to create a reservation policy.
- 8 Select a machine prefix from the **Machine prefix** drop-down menu.
You can select **Use group default** to accept the default machine prefix business group for the user.
- 9 (Optional) Enter a number in the **Maximum per user** text box to limit the number of machines that a single user can provision with this blueprint.
- 10 Set the daily cost of the machine by typing the amount in the **Cost (daily)** text box.

Your blueprint is not finished. Do not navigate away from this page.

Specify Basic Red Hat OpenStack Blueprint Build Information

Build information settings determine how a machine is provisioned and control the resources consumed by a machine provisioned from your blueprint.

Prerequisites

[“Specify Basic Red Hat OpenStack Blueprint Information,”](#) on page 62.

Procedure

- 1 Click the **Build Information** tab.
- 2 Select whether machines provisioned from this blueprint are classified as Desktop or Server from the **Blueprint type** drop-down menu.
This information is for record-keeping and licensing purposes only.
- 3 Select **CloudProvisioningWorkflow** from the **Provisioning workflow** drop-down menu.
- 4 Click the **Browse** icon next to the **OpenStack image** text box.
- 5 Select a virtual machine image from the list.

IMPORTANT Ensure that you select an image that is appropriate for the workflow that you selected. For example, only select kickstart images when you use the `CloudLinuxKickstartWorkflow`, and only select WIM images when you use the `CloudWIMImageWorkflow`. Using an image with the wrong workflow can result in a machine that appears to be successfully provisioned and ready to use when it is actually waiting for input in the Red Hat OpenStack provider.

- 6 Click **OK**.

- 7 Select a key pair generation method from the **Key pair** drop-down menu.

Option	Description
Not Specified	Select to control key pair behavior at the reservation level rather than the blueprint level.
Auto Generated per Business Group	Each machine provisioned in the same business group has the same key pair. If you delete the business group, its key pair is also deleted.
Auto Generated per Machine	Each machine has a unique key pair.

If you select **Not Specified** in both the reservation and the blueprint, machines provisioned from this blueprint have no key pair.

- 8 Select one or more **Flavors** check boxes to select one or more flavors that users can select from when they submit a machine request.

Selecting one or more flavors updates values in the **# CPUs**, **Memory**, and **Storage (GB)** text boxes. Review the minimum and maximum machine resource values and adjust your instance type selections as desired.

- 9 (Optional) Specify the lease settings for provisioned machines, or leave blank for no expiration date.
- Enter the minimum number of lease days in the **Minimum** text box.
If you only provide a minimum, this number becomes the value for all machines provisioned from this blueprint.
 - (Optional) Enter the maximum number of lease days in the **Maximum** text box to allow users to select their own settings within the range that you provide.

Your blueprint is not finished. Do not navigate away from this page.

Add Basic Red Hat OpenStack Blueprint Custom Properties

Adding custom properties to a blueprint gives you detailed control over provisioned machines.

Prerequisites

[“Specify Basic Red Hat OpenStack Blueprint Build Information,”](#) on page 63.

Procedure

- Click the **Properties** tab.
- (Optional) Select one or more build profiles from the **Build profiles** menu.
Build profiles contain groups of custom properties. Fabric administrators can create build profiles.
- (Optional) Add any custom properties to your blueprint.
 - Click **New Property**.
 - Enter the custom property in the **Name** text box.
 - Enter the value of the custom property in the **Value** text box.
 - (Optional) Select the **Encrypted** check box to encrypt the custom property in the database.

- e (Optional) Select the **Prompt user** check box to require the user to provide a value when they request a machine.

If you choose to prompt users for a value, any value you provide for the custom property is presented to them as the default. If you do not provide a default, users cannot continue with the machine request until they provide a value for the custom property.

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

Your blueprint is not finished. Do not navigate away from this page.

Configure Basic Red Hat OpenStack Blueprint Actions

Use blueprint actions and entitlements together to maintain detailed control over provisioned machines.

Entitlements control which machine operations specific users can access. Blueprint actions control which machine operations can be performed on machines provisioned from a blueprint. For example, if you disable the option to reprovision machines created from a blueprint, then the option to reprovision a machine created from the blueprint does not appear for any users. If you enable the reprovision machine operation, then users who are entitled to reprovision machines can reprovision machines created from this blueprint.

NOTE The options that appear on the **Actions** tab depend on your platform and any customizations made to your vRealize Automation instance.

Prerequisites

[“Add Basic Red Hat OpenStack Blueprint Custom Properties,”](#) on page 64.

Procedure

- 1 Click the **Actions** tab.
- 2 (Optional) Select the check boxes for each machine option to enable for machines provisioned from this blueprint.
- 3 Click **OK**.

Your blueprint is saved in draft state.

What to do next

Publish your blueprint to make it available as a catalog item. See [“Publish a Blueprint,”](#) on page 78.

Create a Red Hat OpenStack Blueprint for Linux Kickstart Provisioning

You can provision a machine by booting from an ISO image, then using a kickstart or autoYaSt configuration file and a Linux distribution image to install the operating system on the machine.

Prerequisites

- Log in to the vRealize Automation console as a **tenant administrator** or **business group manager**.
- Gather the following information from your fabric administrator:
 - a The name and location of the ISO image.
 - b For vCenter Server integrations, the vCenter Server guest operating system version with which vCenter Server is to create the machine.

- c (Optional) For SCVMM integrations, the ISO, virtual hard disk, or hardware profile to attach to provisioned machines.

NOTE Your fabric administrator might have provided this information in a build profile.

Procedure

- 1 [Specify Linux Kickstart Blueprint Information](#) on page 66
Blueprint information settings control who can access a blueprint, how many machines they can provision with it, and how long to archive a machine after the lease period is over.
- 2 [Specify Linux Kickstart Blueprint Build Information](#) on page 67
The build information settings determine how a machine is provisioned.
- 3 [Add Required Properties to a Linux Kickstart Blueprint](#) on page 68
There are a number of custom properties that are required for Linux Kickstart provisioning.
- 4 [Configure Linux Kickstart Blueprint Actions](#) on page 71
Use blueprint actions and entitlements together to maintain detailed control over provisioned machines.

What to do next

The blueprint is not available as a catalog item until you publish it. See [“Publish a Blueprint,”](#) on page 78.

Specify Linux Kickstart Blueprint Information

Blueprint information settings control who can access a blueprint, how many machines they can provision with it, and how long to archive a machine after the lease period is over.

Prerequisites

- Log in to the vRealize Automation console as a **tenant administrator** or **business group manager**.
- Gather the following information from your fabric administrator:
 - a The name and location of the ISO image.
 - b For vCenter Server integrations, the vCenter Server guest operating system version with which vCenter Server is to create the machine.
 - c (Optional) For SCVMM integrations, the ISO, virtual hard disk, or hardware profile to attach to provisioned machines.

NOTE Your fabric administrator might have provided this information in a build profile.

Procedure

- 1 Select **Infrastructure > Blueprints > Blueprints**.
- 2 Select **New Blueprint > Cloud > OpenStack**.
- 3 Enter a name and, optionally, a description.
- 4 (Optional) Select the **Master** check box to allow users to copy your blueprint.
- 5 Select the **Display location on request** check box to prompt users to select a location when they submit a machine request.

This option enables users to select a region into which to provision the requested machine.

- 6 Select who can provision machines with this blueprint.

Roles	Who Can Provision
If you are both a business group manager and a tenant administrator	<ul style="list-style-type: none"> ■ Select the Shared blueprint check box to allow the blueprint to be entitled to users in any business group. ■ Deselect the Shared blueprint check box to create a local blueprint, and select a business group from the Business group drop-down menu.
Business group manager	Select a business group from the Business group drop-down menu.
Tenant administrator	Your blueprints are always shared. You cannot choose who can use them to provision machines.

- 7 (Optional) Select a reservation policy from the **Reservation policy** drop-down menu.
This option requires additional configuration by a fabric administrator to create a reservation policy.
- 8 Select a machine prefix from the **Machine prefix** drop-down menu.
You can select **Use group default** to accept the default machine prefix business group for the user.
- 9 (Optional) Enter a number in the **Maximum per user** text box to limit the number of machines that a single user can provision with this blueprint.
- 10 Set the daily cost of the machine by typing the amount in the **Cost (daily)** text box.

Your blueprint is not finished. Do not navigate away from this page.

Specify Linux Kickstart Blueprint Build Information

The build information settings determine how a machine is provisioned.

Prerequisites

[“Specify Linux Kickstart Blueprint Information,”](#) on page 66.

Procedure

- 1 Click the **Build Information** tab.
- 2 Select whether machines provisioned from this blueprint are classified as Desktop or Server from the **Blueprint type** drop-down menu.
This information is for record-keeping and licensing purposes only.
- 3 Select **CloudLinuxKickstartWorkflow** from the **Provisioning workflow** drop-down menu.
- 4 Click the **Browse** icon next to the **OpenStack image** text box.
- 5 Select a virtual machine image from the list.

IMPORTANT Ensure that you select an image that is appropriate for the workflow that you selected. For example, only select kickstart images when you use the CloudLinuxKickstartWorkflow, and only select WIM images when you use the CloudWIMImageWorkflow. Using an image with the wrong workflow can result in a machine that appears to be successfully provisioned and ready to use when it is actually waiting for input in the Red Hat OpenStack provider.

- 6 Click **OK**.

- 7 Select a key pair generation method from the **Key pair** drop-down menu.

Option	Description
Not Specified	Select to control key pair behavior at the reservation level rather than the blueprint level.
Auto Generated per Business Group	Each machine provisioned in the same business group has the same key pair. If you delete the business group, its key pair is also deleted.
Auto Generated per Machine	Each machine has a unique key pair.

If you select **Not Specified** in both the reservation and the blueprint, machines provisioned from this blueprint have no key pair.

- 8 Select one or more **Flavors** check boxes to select one or more flavors that users can select from when they submit a machine request.

Selecting one or more flavors updates values in the **# CPUs**, **Memory**, and **Storage (GB)** text boxes. Review the minimum and maximum machine resource values and adjust your instance type selections as desired.

- 9 (Optional) Specify the lease settings for provisioned machines, or leave blank for no expiration date.

- a Enter the minimum number of lease days in the **Minimum** text box.

If you only provide a minimum, this number becomes the value for all machines provisioned from this blueprint.

- b (Optional) Enter the maximum number of lease days in the **Maximum** text box to allow users to select their own settings within the range that you provide.

Your blueprint is not finished. Do not navigate away from this page.

Add Required Properties to a Linux Kickstart Blueprint

There are a number of custom properties that are required for Linux Kickstart provisioning.

NOTE If your fabric administrator creates a build profile that contains the required custom properties and you include it in your blueprint, you do not need to individually add the required custom properties to the blueprint.

For a list of commonly used custom properties for Linux Kickstart provisioning, see [“Custom Properties for Linux Kickstart Blueprints,”](#) on page 69.

Prerequisites

[“Specify Linux Kickstart Blueprint Build Information,”](#) on page 67.

Procedure



- 1 Click the **Properties** tab.
- 2 (Optional) Select one or more build profiles from the **Build profiles** menu.
Build profiles contain groups of custom properties. Fabric administrators can create build profiles.
- 3 Add the ISO name custom property.
 - a Click **New Property**.
 - b Type `Image.ISO.Name` in the **Name** text box.

- c Type the name or location of the ISO image that your fabric administrator provided for you in the **Value** text box.

Option	Description
vCenter Server	The path to the ISO, including the name. The value must use forward slashes.
Hyper-V	The full local path to the ISO file, including the file name.
XenServer	The name of the ISO file.

- d (Optional) Select the **Encrypted** check box to encrypt the custom property in the database.
- e (Optional) Select the **Prompt user** check box to require the user to provide a value when they request a machine.

If you choose to prompt users for a value, any value you provide for the custom property is presented to them as the default. If you do not provide a default, users cannot continue with the machine request until they provide a value for the custom property.

- f Click the **Save** icon (.
- 4 Add the ISO location custom property.
- This custom property is not required for Hyper-V integrations. The name and location are both specified in the ISO name custom property.
- a Click **New Property**.
- b Type `Image.ISO.Location` in the **Name** text box.
- c Type the location of the ISO image that your fabric administrator provided for you in the **Value** text box.
- d Click the **Save** icon (.

Your blueprint is not finished. Do not navigate away from this page.

Custom Properties for Linux Kickstart Blueprints

vRealize Automation includes custom properties that you can use to provide additional controls for Linux Kickstart blueprints.

Certain vRealize Automation custom properties are required to use with Linux Kickstart blueprints.

Table 4-2. Required Custom Properties for Linux Kickstart Blueprints

Custom Property	Description
<code>VMware.VirtualCenter.OperatingSystem</code>	Specifies the vCenter Server guest operating system version (<code>VirtualMachineGuestOsIdentifier</code>) with which vCenter Server creates the machine. This operating system version must match the operating system version to be installed on the provisioned machine. Administrators can create property groups using one of several property sets, for example, <code>VMware[OS_Version]Properties</code> , that are predefined to include the correct <code>VMware.VirtualCenter.OperatingSystem</code> values. This property is for virtual provisioning. For related information, see the enumeration type <code>VirtualMachineGuestOsIdentifier</code> in vSphere API/SDK Documentation. For a list of currently accepted values, see the VMware vCenter Server™ documentation.
<code>Image.ISO.Location</code>	Values for this property are case sensitive. Specifies the location of the ISO image from which to boot, for example <code>http://192.168.2.100/site2/winpe.iso</code> . The format of this value depends on your platform. For details, see documentation provided for your platform. This property is required for WIM-based provisioning, Linux Kickstart and autoYaST provisioning, and SCCM-based provisioning.
<code>Image.ISO.Name</code>	Values for this property are case sensitive. Specifies the name of the ISO image from which to boot, for example <code>/ISO/Microsoft/WinPE.iso</code> . The format of this value depends on your platform. For details, see documentation provided for your platform. This property is required for WIM-based provisioning, Linux Kickstart and autoYaST provisioning, and SCCM-based provisioning.
<code>Image.ISO.UserName</code>	Specifies the user name to access the CIFS share in the format <code>username@domain</code> . For Dell iDRAC integrations where the image is located on a CIFS share that requires authentication to access.
<code>Image.ISO.Password</code>	Specifies the password associated with the <code>Image.ISO.UserName</code> property. For Dell iDRAC integrations where the image is located on a CIFS share that requires authentication to access.

Optional custom properties are available.

Table 4-3. Optional Custom Properties for Linux Kickstart Blueprints

Custom Property	Description
<code>VirtualMachine.Admin.ThinProvision</code>	Determines whether thin provisioning is used on ESX compute resources using local or iSCSI storage. Set to True to use thin provisioning. Set to False to use standard provisioning. This property is for virtual provisioning.
<code>Machine.SSH</code>	Set to True to enable the Connect Using SSH option, on the vRealize Automation Items page, for Linux machines provisioned from this blueprint. If set to True and the Connect using RDP or SSH machine operation is enabled in the blueprint, all Linux machines that are provisioned from the blueprint display the Connect Using SSH option to entitled users.

Configure Linux Kickstart Blueprint Actions

Use blueprint actions and entitlements together to maintain detailed control over provisioned machines.

Entitlements control which machine operations specific users can access. Blueprint actions control which machine operations can be performed on machines provisioned from a blueprint. For example, if you disable the option to reprovision machines created from a blueprint, then the option to reprovision a machine created from the blueprint does not appear for any users. If you enable the reprovision machine operation, then users who are entitled to reprovision machines can reprovision machines created from this blueprint.

NOTE The options that appear on the **Actions** tab depend on your platform and any customizations made to your vRealize Automation instance.

Prerequisites

[“Add Required Properties to a Linux Kickstart Blueprint,”](#) on page 68.

Procedure

- 1 Click the **Actions** tab.
- 2 (Optional) Select the check boxes for each machine option to enable for machines provisioned from this blueprint.
- 3 Click **OK**.

Your blueprint is saved in draft state.

What to do next

Publish your blueprint to make it available as a catalog item. See [“Publish a Blueprint,”](#) on page 78.

Create an OpenStack Blueprint for WIM Provisioning

You can provision a machine by booting into a WinPE environment and then installing an operating system using a Windows Imaging File Format (WIM) image of an existing Windows reference machine.

Prerequisites

- Log in to the vRealize Automation console as a **tenant administrator** or **business group manager**.
- Gather the following information from your fabric administrator:
 - a The name and location of the WinPE ISO image.
 - b The name of the WIM file, the UNC path to the WIM file, and the index used to extract the desired image from the WIM file.
 - c The user name and password under which to map the WIM image path to a network drive on the provisioned machine.
 - d (Optional) If you do not want to accept the default, K, the drive letter to which the WIM image path is mapped on the provisioned machine.
 - e For vCenter Server integrations, the vCenter Server guest operating system version with which vCenter Server is to create the machine.
 - f (Optional) For SCVMM integrations, the ISO, virtual hard disk, or hardware profile to attach to provisioned machines.

NOTE Your fabric administrator might have provided this information in a build profile.

Procedure

- 1 [Specify WIM Blueprint Information](#) on page 72
Blueprint information settings control who can access a blueprint, how many machines they can provision with it, and how long to archive a machine after the lease period is over.
- 2 [Specify WIM Blueprint Build Information](#) on page 73
The build information settings determine how a machine is provisioned.
- 3 [Add Required Properties to a WIM Blueprint](#) on page 74
There are a number of custom properties that are required for WIM provisioning.
- 4 [Configure WIM Blueprint Actions](#) on page 78
Use blueprint actions and entitlements together to maintain detailed control over provisioned machines.

What to do next

The blueprint is not available as a catalog item until you publish it. See [“Publish a Blueprint,”](#) on page 78.

Specify WIM Blueprint Information

Blueprint information settings control who can access a blueprint, how many machines they can provision with it, and how long to archive a machine after the lease period is over.

Prerequisites

- Log in to the vRealize Automation console as a **tenant administrator** or **business group manager**.
- Gather the following information from your fabric administrator:
 - a The name and location of the WinPE ISO image.
 - b The name of the WIM file, the UNC path to the WIM file, and the index used to extract the desired image from the WIM file.
 - c The user name and password under which to map the WIM image path to a network drive on the provisioned machine.
 - d (Optional) If you do not want to accept the default, K, the drive letter to which the WIM image path is mapped on the provisioned machine.
 - e For vCenter Server integrations, the vCenter Server guest operating system version with which vCenter Server is to create the machine.
 - f (Optional) For SCVMM integrations, the ISO, virtual hard disk, or hardware profile to attach to provisioned machines.

NOTE Your fabric administrator might have provided this information in a build profile.

Procedure

- 1 Select **Infrastructure > Blueprints > Blueprints**.
- 2 Select **New Blueprint > Cloud > OpenStack**.
- 3 Enter a name and, optionally, a description.
- 4 (Optional) Select the **Master** check box to allow users to copy your blueprint.
- 5 (Optional) Select the **Display location on request** check box to prompt users to choose a datacenter location when they submit a machine request.

This option requires additional configuration to add datacenter locations and associate compute resources with those locations.

- 6 Select who can provision machines with this blueprint.

Roles	Who Can Provision
If you are both a business group manager and a tenant administrator	<ul style="list-style-type: none"> ■ Select the Shared blueprint check box to allow the blueprint to be entitled to users in any business group. ■ Deselect the Shared blueprint check box to create a local blueprint, and select a business group from the Business group drop-down menu.
Business group manager	Select a business group from the Business group drop-down menu.
Tenant administrator	Your blueprints are always shared. You cannot choose who can use them to provision machines.

- 7 (Optional) Select a reservation policy from the **Reservation policy** drop-down menu.
This option requires additional configuration by a fabric administrator to create a reservation policy.
- 8 Select a machine prefix from the **Machine prefix** drop-down menu.
You can select **Use group default** to accept the default machine prefix business group for the user.
- 9 (Optional) Enter a number in the **Maximum per user** text box to limit the number of machines that a single user can provision with this blueprint.
- 10 Set the daily cost of the machine by typing the amount in the **Cost (daily)** text box.

Your blueprint is not finished. Do not navigate away from this page.

Specify WIM Blueprint Build Information

The build information settings determine how a machine is provisioned.

Prerequisites

[“Specify WIM Blueprint Information,”](#) on page 72.

Procedure

- 1 Click the **Build Information** tab.
- 2 Select whether machines provisioned from this blueprint are classified as Desktop or Server from the **Blueprint type** drop-down menu.
This information is for record-keeping and licensing purposes only.
- 3 Select **CloudWIMImageWorkflow** from the **Provisioning workflow** drop-down menu.
- 4 Click the **Browse** icon next to the **OpenStack image** text box.
- 5 Select a virtual machine image from the list.

IMPORTANT Ensure that you select an image that is appropriate for the workflow that you selected. For example, only select kickstart images when you use the `CloudLinuxKickstartWorkflow`, and only select WIM images when you use the `CloudWIMImageWorkflow`. Using an image with the wrong workflow can result in a machine that appears to be successfully provisioned and ready to use when it is actually waiting for input in the Red Hat OpenStack provider.

- 6 Click **OK**.

- 7 Select a key pair generation method from the **Key pair** drop-down menu.

Option	Description
Not Specified	Select to control key pair behavior at the reservation level rather than the blueprint level.
Auto Generated per Business Group	Each machine provisioned in the same business group has the same key pair. If you delete the business group, its key pair is also deleted.
Auto Generated per Machine	Each machine has a unique key pair.

If you select **Not Specified** in both the reservation and the blueprint, machines provisioned from this blueprint have no key pair.

- 8 Select one or more **Flavors** check boxes to select one or more flavors that users can select from when they submit a machine request.

Selecting one or more flavors updates values in the **# CPUs**, **Memory**, and **Storage (GB)** text boxes. Review the minimum and maximum machine resource values and adjust your instance type selections as desired.

- 9 (Optional) Specify the lease settings for provisioned machines, or leave blank for no expiration date.
- Enter the minimum number of lease days in the **Minimum** text box.
If you only provide a minimum, this number becomes the value for all machines provisioned from this blueprint.
 - (Optional) Enter the maximum number of lease days in the **Maximum** text box to allow users to select their own settings within the range that you provide.

Your blueprint is not finished. Do not navigate away from this page.

Add Required Properties to a WIM Blueprint

There are a number of custom properties that are required for WIM provisioning.

NOTE If your fabric administrator creates a build profile that contains the required custom properties and you include it in your blueprint, you do not need to individually add the required custom properties to the blueprint.

For a list of custom properties commonly used with WIM provisioning, see [“Custom Properties for WIM Blueprints,”](#) on page 75.

Prerequisites

[“Specify WIM Blueprint Build Information,”](#) on page 73.

Procedure

- Click the **Properties** tab.
- (Optional) Select one or more build profiles from the **Build profiles** menu.
Build profiles contain groups of custom properties. Fabric administrators can create build profiles.
- Add the ISO name custom property.
 - Click **New Property**.
 - Enter `Image.ISO.Name` in the **Name** text box.

- c Enter the location of the ISO image that your fabric administrator provided for you in the **Value** text box.

The value of `Image.ISO.Name` is case sensitive.

Option	Description
vCenter Server	The path to the WinPE ISO, including the name. The value must use forward slashes, for example: <code>/MyISOs/Microsoft/MSDN/win2003.iso</code> .
Hyper-V	The full local path to the WinPE ISO file, including the file name.
XenServer	The name of the WinPE ISO file.

- d (Optional) Select the **Encrypted** check box to encrypt the custom property in the database.
- e (Optional) Select the **Prompt user** check box to require the user to provide a value when they request a machine.

If you choose to prompt users for a value, any value you provide for the custom property is presented to them as the default. If you do not provide a default, users cannot continue with the machine request until they provide a value for the custom property.

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

- 4 Add the remaining required custom properties for WIM provisioning.

Option	Description
Image.ISO.Location	Enter the location of the ISO image that your fabric administrator provided for you in the Value text box. This custom property is case sensitive. For Hyper-V integrations, this custom property is not required because the name and location are both specified in the ISO name custom property.
Image.WIM.Path	Enter the UNC path to the WIM file that your fabric administrator provided for you in the Value text box.
Image.WIM.Name	Enter the name of the WIM file that your fabric administrator provided for you in the Value text box.
Image.WIM.Index	Enter the index to be used to extract the desired image from the WIM file that your fabric administrator provided for you in the Value text box.
Image.Network.User	Enter the user name under which to map the WIM image path (<code>Image.WIM.Path</code>) to a network drive on the machine in the Value text box.
Image.Network.Password	Enter the associated password for the network user (<code>Image.Network.User</code>) in the Value text box.
VirtualMachine.Admin.Password	Enter the password to use for the administrator account on the provisioned machine.

Your blueprint is not finished. Do not navigate away from this page.

Custom Properties for WIM Blueprints

vRealize Automation includes custom properties that provide additional controls for WIM blueprints.

Certain vRealize Automation custom properties are required for WIM blueprints.

Table 4-4. Required Custom Properties for WIM Blueprints

Custom Property	Description
<code>Image.ISO.Location</code>	Values for this property are case sensitive. Specifies the location of the ISO image from which to boot, for example <code>http://192.168.2.100/site2/winpe.iso</code> . The format of this value depends on your platform. For details, see documentation provided for your platform. This property is required for WIM-based provisioning, Linux Kickstart and autoYaST provisioning, and SCCM-based provisioning.
<code>Image.ISO.Name</code>	Values for this property are case sensitive. Specifies the name of the ISO image from which to boot, for example <code>/ISO/Microsoft/WinPE.iso</code> . The format of this value depends on your platform. For details, see documentation provided for your platform. This property is required for WIM-based provisioning, Linux Kickstart and autoYaST provisioning, and SCCM-based provisioning.
<code>Image.ISO.UserName</code>	Specifies the user name to access the CIFS share in the format <code>username@domain</code> . For Dell iDRAC integrations where the image is located on a CIFS share that requires authentication to access.
<code>Image.ISO.Password</code>	Specifies the password associated with the <code>Image.ISO.UserName</code> property. For Dell iDRAC integrations where the image is located on a CIFS share that requires authentication to access.
<code>Image.Network.Letter</code>	Specifies the drive letter to which the WIM image path is mapped on the provisioned machine. The default value is K.
<code>Image.WIM.Path</code>	Specifies the UNC path to the WIM file from which an image is extracted during WIM-based provisioning. The path format is <code>\\server\share\$</code> format, for example <code>\\lab-ad\dfs\$</code> .
<code>Image.WIM.Name</code>	Specifies the name of the WIM file, for example <code>win2k8.wim</code> , as located by the <code>Image.WIM.Path</code> property.
<code>Image.WIM.Index</code>	Specifies the index used to extract the correct image from the WIM file.
<code>Image.Network.User</code>	Specifies the user name with which to map the WIM image path (<code>Image.WIM.Path</code>) to a network drive on the provisioned machine. This is typically a domain account with access to the network share.
<code>Image.Network.Password</code>	Specifies the password associated with the <code>Image.Network.User</code> property.
<code>VMware.VirtualCenter.OperatingSystem</code>	Specifies the vCenter Server guest operating system version (<code>VirtualMachineGuestOsIdentifier</code>) with which vCenter Server creates the machine. This operating system version must match the operating system version to be installed on the provisioned machine. Administrators can create property groups using one of several property sets, for example, <code>VMware[OS_Version]Properties</code> , that are predefined to include the correct <code>VMware.VirtualCenter.OperatingSystem</code> values. This property is for virtual provisioning. For related information, see the enumeration type <code>VirtualMachineGuestOsIdentifier</code> in vSphere API/SDK Documentation. For a list of currently accepted values, see the VMware vCenter Server™ documentation.

Optional custom properties are also available for WIM blueprints.

Table 4-5. Common Custom Properties for WIM Blueprints

Custom Property	Description
SysPrep. <i>Section.Key</i> ■ SysPrep.GuiUnattended.AdminPassword ■ SysPrep.GuiUnattended.EncryptedAdminPassword ■ SysPrep.GuiUnattended.TimeZone	<p>Specifies information to be added to the SysPrep answer file on machines during the WinPE stage of provisioning. Information that already exists in the SysPrep answer file is overwritten by these custom properties. <i>Section</i> represents the name of the section of the SysPrep answer file, for example GuiUnattended or UserData. <i>Key</i> represents a key name in the section. For example, to set the time zone of a provisioned machine to West Pacific Standard Time, define the custom property GuiUnattended.UserData.TimeZone and set the value to 275.</p> <p>For a full list of sections, keys, and accepted values, see the System Preparation Utility for Windows documentation.</p> <p>The following <i>Section.Key</i> combinations can be specified for WIM-based provisioning:</p> <ul style="list-style-type: none"> ■ GuiUnattended <ul style="list-style-type: none"> ■ AdminPassword ■ EncryptedAdminPassword ■ TimeZone ■ UserData <ul style="list-style-type: none"> ■ ProductKey ■ FullName ■ ComputerName ■ OrgName ■ Identification <ul style="list-style-type: none"> ■ DomainAdmin ■ DomainAdminPassword ■ JoinDomain ■ JoinWorkgroup
Sysprep.Identification.DomainAdmin	Specifies a user name with administrator-level access to the target domain in Active Directory. Do not include the user domain in the credentials that you send to vCloud Director.
Sysprep.Identification.DomainAdminPassword	Specifies the password to associate with the Sysprep.Identification.DomainAdmin property.
Sysprep.Identification.JoinDomain	Specifies the name of the domain to join in Active Directory.
Sysprep.Identification.JoinWorkgroup	Specifies the name of the workgroup to join if not using a domain.
SysPrep.UserData.ComputerName	Specifies a machine name, for example lab-client005.
SysPrep.UserData.FullName	Specifies the full name of a user.
SysPrep.UserData.OrgName	Specifies the organization name of the user.
SysPrep.UserData.ProductKey	Specifies the Windows product key.
VirtualMachine.Admin.ThinProvision	Determines whether thin provisioning is used on ESX compute resources using local or iSCSI storage. Set to True to use thin provisioning. Set to False to use standard provisioning. This property is for virtual provisioning.

Configure WIM Blueprint Actions

Use blueprint actions and entitlements together to maintain detailed control over provisioned machines.

Entitlements control which machine operations specific users can access. Blueprint actions control which machine operations can be performed on machines provisioned from a blueprint. For example, if you disable the option to reprovision machines created from a blueprint, then the option to reprovision a machine created from the blueprint does not appear for any users. If you enable the reprovision machine operation, then users who are entitled to reprovision machines can reprovision machines created from this blueprint.

NOTE The options that appear on the **Actions** tab depend on your platform and any customizations made to your vRealize Automation instance.

Prerequisites

[“Add Required Properties to a WIM Blueprint,”](#) on page 74.

Procedure

- 1 Click the **Actions** tab.
- 2 (Optional) Select the check boxes for each machine option to enable for machines provisioned from this blueprint.
- 3 Click **OK**.

Your blueprint is saved in draft state.

What to do next

Publish your blueprint to make it available as a catalog item. See [“Publish a Blueprint,”](#) on page 78.

Publish a Blueprint

Blueprints are saved in the draft state and must be manually published before you can configure them as catalog items.

You need to publish a blueprint only once. Any changes you make to a published blueprint are automatically reflected in the catalog.

Prerequisites

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

Procedure

- 1 Select **Infrastructure > Blueprints > Blueprints**.
- 2 Point to the blueprint to publish and click **Publish** from the drop-down menu.
- 3 Click **OK**.

Your blueprint is now ready for tenant administrators, business group managers, and service architects to associate it with a catalog service and entitle users to request it from the catalog.

What to do next

For information about how to configure and manage the catalog, see *Tenant Administration*.

Configuring Advanced Blueprint Settings

5

Tenant administrators and business group managers can use custom properties and optional policies to configure advanced blueprint settings such as reservation policies, Visual Basic scripts, and Active Directory cleanup.

This chapter includes the following topics:

- [“Reservation Policies,”](#) on page 79
- [“Configuring Network Settings,”](#) on page 80
- [“Enabling Visual Basic Scripts in Provisioning,”](#) on page 84
- [“Add Active Directory Cleanup to a Blueprint,”](#) on page 86
- [“Enabling Remote Desktop Connections,”](#) on page 87
- [“Enable Connections Using SSH,”](#) on page 89

Reservation Policies

When a user requests a machine, it can be provisioned on any reservation of the appropriate type that has sufficient capacity for the machine. You can apply a reservation policy to a blueprint to restrict the machines provisioned from a that blueprint to a subset of available reservations.

Reservation policies provide an optional and helpful means of controlling how reservation requests are processed. A reservation policy is often used to collect resources into groups for different service levels, or to make a specific type of resource easily available for a particular purpose. The following scenarios provide a few examples of possible uses for reservation policies:

- To ensure that machines provisioned from a virtual blueprint are placed on reservations with storage devices that support NetApp FlexClone
- To restrict provisioning of cloud machines to a specific region containing a machine image that is required for a specific blueprint
- To restrict provisioning of Cisco UCS physical machines to reservations on endpoints on which the selected service profile template and boot policy are available
- As an additional means of using a Pay As You Go allocation model for vApps

You can add multiple reservations to a reservation policy, but a reservation can belong to only one policy. You can assign a single reservation policy to more than one blueprint. A blueprint can have only one reservation policy.

A reservation policy can include reservations of different types, but only reservations that match the blueprint type are considered when selecting a reservation for a particular request. For more information about how reservations are selected for provisioning a machine, see *IaaS Configuration for Cloud Platforms*, *IaaS Configuration for Physical Machines*, or *IaaS Configuration for Virtual Platforms*.

NOTE If you have SDRS enabled on your platform, you can allow SDRS to load balance storage for individual virtual machine disks, or all storage for the virtual machine. If you are working with SDRS datastore clusters, conflicts can occur when you use reservation policies and storage reservation policies. For example, if a standalone datastore or a datastore within an SDRS cluster is selected on one of the reservations in a policy or storage policy, your virtual machine storage might be frozen instead of driven by SDRS.


Add a Reservation Policy to a Blueprint

When tenant administrators and business group managers create a new blueprint, the option to add a reservation policy appears. To add a reservation policy to an existing blueprint, you edit the blueprint.

Prerequisites

- Log in to the vRealize Automation console as a **tenant administrator** or **business group manager**.
- Verify that a fabric administrator configured a reservation policy. See [“Configure a Reservation Policy,”](#) on page 35.

Procedure

- 1 Select **Infrastructure > Blueprints > Blueprints**.
- 2 Point to the blueprint to edit.
- 3 Click the **Edit** icon ().
- 4 Select a **Reservation policy** from the **Reservation policy** drop-down menu.
- 5 Click **OK**.

The machines provisioned from your blueprint are now restricted to the resources specified in your reservation policy.

Configuring Network Settings

You can also specify aspects of network configuration such as static IP address assignment, the networks to which machines should be attached, and other networking details.

By default, vCloud Automation Center uses DHCP to assign IP addresses to provisioned machines. For cloud and virtual machines provisioned by using cloning or Linux kickstart/autoYaST provisioning, it is possible to assign static IP addresses from a predefined range. Static IP address assignment can be configured either at the reservation level or at the blueprint level.

NOTE This information does not apply to Amazon Web Services.

Add a Network Profile for Static IP Assignment to a Blueprint

Tenant administrators and business group managers can configure static IP assignment at the blueprint level.

NOTE This information does not apply to Amazon Web Services.

For a list of custom properties related to networking, see [“Custom Properties for Networking,”](#) on page 81.


Prerequisites

- Log in to the vRealize Automation console as a **tenant administrator** or **business group manager**.
- A fabric administrator must “[Create a Network Profile for Static IP Address Assignment](#),” on page 36.
- Create at least one blueprint that is eligible for static IP assignment. Static IP is supported with cloud and virtual machines provisioned by cloning or Linux kickstart/autoYaST. For SCVMM, you can only use static IP addresses if you are cloning Linux machines and you installed the guest agent on your template.
- For vSphere clone blueprints, you must specify a valid customization specification on the Build Information tab of the blueprint.

Procedure

- 1 Select **Infrastructure > Blueprints > Blueprints**.
- 2 Point to the blueprint for which you want to enable static IP assignment and click **Edit**.
- 3 Click the **Properties** tab.
- 4 Specify a network profile to use for assigning static IP addresses to a network device on machines provisioned from this blueprint.
 - a Click **New Property**.
 - b Type **VirtualMachine.NetworkN.ProfileName** in the **Name** text box, where *N* is the number of the network device for which to enable static IP assignment.

For example: **VirtualMachine.Network0.ProfileName**.
 - c Type the name of the network profile that defines the static IP addresses that can be assigned to the network device.
 - d (Optional) Select the **Prompt user** check box to require the user to provide a value when they request a machine.

If you choose to prompt users for a value, any value you provide for the custom property is presented to them as the default. If you do not provide a default, users cannot continue with the machine request until they provide a value for the custom property.
 - e Click the **Save** icon ().

Repeat this step to configure static IP assignment for additional network devices in this blueprint.
- 5 Click **OK**.

Custom Properties for Networking

The vRealize Automation custom properties for networking specify configuration for a specific network device on a machine.

NOTE This information does not apply to Amazon Web Services.

Network assignments are performed during machine allocation. vRealize Automation retrieves network information from the blueprint. If you want to assign more than one network, use the `VirtualMachine.NetworkN.Name` custom property on your machine blueprint. If you do not provide custom properties, allocation only assigns one network, which is picked using a round robin method in conjunction with the selected reservation.

By default, a machine has one network device configured with the `VirtualMachine.Network0.Name` property. You can configure additional network devices by using the `VirtualMachine.NetworkN.Name` custom property.

The numbering of network properties must be sequential, starting with 0. For example, if you specify custom properties for only `VirtualMachine.Network0` and `VirtualMachine.Network2`, the properties for `VirtualMachine.Network2` are ignored, because the preceding network, `VirtualMachine.Network1`, was not specified.

Table 5-1. Custom Properties for Networking Configuration

Custom Property	Description
<code>VirtualMachine.NetworkN.Address</code>	Specifies the IP address of network device <i>N</i> in a machine provisioned with a static IP address.
<code>VirtualMachine.NetworkN.MacAddressType</code>	<p>Indicates whether the MAC address of network device <i>N</i> is generated or user-defined (static). This property is available for cloning.</p> <p>The default value is generated. If the value is static, you must also use <code>VirtualMachine.NetworkN.MacAddress</code> to specify the MAC address.</p> <p><code>VirtualMachine.NetworkN</code> custom properties are designed to be specific to individual blueprints and machines. When a machine is requested, network and IP address allocation is performed before the machine is assigned to a reservation. Because blueprints are not guaranteed to be allocated to a specific reservation, do not use this property on a reservation.</p>
<code>VirtualMachine.NetworkN.MacAddress</code>	<p>Specifies the MAC address of a network device <i>N</i>. This property is available for cloning.</p> <p>If the value of <code>VirtualMachine.NetworkN.MacAddressType</code> is generated, this property contains the generated address.</p> <p>If the value of <code>VirtualMachine.Network.N.MacAddressType</code> is static, this property specifies the MAC address. For virtual machines provisioned on ESX server hosts, the address must be in the range specified by VMware. For details, see vSphere documentation.</p> <p><code>VirtualMachine.NetworkN</code> custom properties are designed to be specific to individual blueprints and machines. When a machine is requested, network and IP address allocation is performed before the machine is assigned to a reservation. Because blueprints are not guaranteed to be allocated to a specific reservation, do not use this property on a reservation.</p>
<code>VirtualMachine.NetworkN.Name</code>	<p>Specifies the name of the network to connect to, for example the network device <i>N</i> to which a machine is attached.</p> <p>By default, a network is assigned from the network paths available on the reservation on which the machine is provisioned.</p> <p>You can ensure that a network device is connected to a specific network by setting the value of this property to the name of a network on an available reservation.</p> <p><code>VirtualMachine.NetworkN</code> custom properties are designed to be specific to individual blueprints and machines. When a machine is requested, network and IP address allocation is performed before the machine is assigned to a reservation. Because blueprints are not guaranteed to be allocated to a specific reservation, do not use this property on a reservation.</p> <p>You can add this property to a vApp (vCloud) blueprint or a vApp (vCloud) Component blueprint.</p>

Table 5-1. Custom Properties for Networking Configuration (Continued)

Custom Property	Description
VirtualMachine.NetworkN.PortID	<p>Specifies the port ID to use for network device <i>N</i> when using a dvPort group with a vSphere distributed switch.</p> <p>VirtualMachine.NetworkN custom properties are designed to be specific to individual blueprints and machines. When a machine is requested, network and IP address allocation is performed before the machine is assigned to a reservation. Because blueprints are not guaranteed to be allocated to a specific reservation, do not use this property on a reservation.</p>
VirtualMachine.NetworkN.ProfileName	<p>Specifies the name of a network profile from which to assign a static IP address to network device <i>N</i> or from which to obtain the range of static IP addresses that can be assigned to network device <i>N</i> of a cloned machine, where <i>N</i>=0 for the first device, 1 for the second, and so on.</p> <p>If a network profile is specified in the network path in the reservation on which the machine is provisioned, a static IP address is assigned from that network profile. You can ensure that a static IP address is assigned from a specific profile by setting the value of this property to the name of a network profile.</p>
<ul style="list-style-type: none"> ■ VirtualMachine.NetworkN.SubnetMask ■ VirtualMachine.NetworkN.Gateway ■ VirtualMachine.NetworkN.PrimaryDns ■ VirtualMachine.NetworkN.SecondaryDns ■ VirtualMachine.NetworkN.PrimaryWins ■ VirtualMachine.NetworkN.SecondaryWins ■ VirtualMachine.NetworkN.DnsSuffix ■ VirtualMachine.NetworkN.DnsSearchSuffixes 	<p>Appending a name allows you to create multiple versions of a custom property. For example, the following properties might list load balancing pools set up for general use and machines with high, moderate, and low performance requirements:</p> <ul style="list-style-type: none"> ■ VCNS.LoadBalancerEdgePool.Names ■ VCNS.LoadBalancerEdgePool.Names.moderate ■ VCNS.LoadBalancerEdgePool.Names.high ■ VCNS.LoadBalancerEdgePool.Names.low <p>NOTE In a multi-machine service, this custom property is supported only for VM components with Routed Virtual Wire. It is also not supported for pre-defined static IP addresses for VM components in the context of a multi-machine blueprint.</p> <p>Configures attributes of the network profile specified in VirtualMachine.NetworkN.ProfileName.</p>
VCNS.LoadBalancerEdgePool.Names.name	<p>Specifies the vCloud Networking and Security load balancing pools to which the virtual machine is assigned during provisioning. The virtual machine is assigned to all service ports of all specified pools. The value is an <i>edge/pool</i> name or a list of <i>edge/pool</i> names separated by commas. Names are case-sensitive.</p> <p>Appending a name allows you to create multiple versions of a custom property. For example, the following properties might list load balancing pools set up for general use and machines with high, moderate, and low performance requirements:</p> <ul style="list-style-type: none"> ■ VCNS.LoadBalancerEdgePool.Names ■ VCNS.LoadBalancerEdgePool.Names.moderate ■ VCNS.LoadBalancerEdgePool.Names.high ■ VCNS.LoadBalancerEdgePool.Names.low <p>NOTE In a multi-machine service, this custom property is supported only for VM components with Routed Virtual Wire. It is also not supported for pre-defined static IP addresses for VM components in the context of a multi-machine blueprint.</p>

Table 5-1. Custom Properties for Networking Configuration (Continued)

Custom Property	Description
<code>VCNS.SecurityGroup.Names.name</code>	<p>Specifies the vCloud Networking and Security security group or groups to which the virtual machine is assigned during provisioning. The value is a security group name or a list of names separated by commas. Names are case-sensitive.</p> <p>Appending a name allows you to create multiple versions of the property, which can be used separately or in combination. For example, the following properties can list security groups intended for general use, for the sales force, and for support:</p> <ul style="list-style-type: none"> ■ <code>VCNS.SecurityGroup.Names</code> ■ <code>VCNS.SecurityGroup.Names.sales</code> ■ <code>VCNS.SecurityGroup.Names.support</code>
<code>VCNS.SecurityTag.Names.name</code>	<p>Specifies the vCloud Networking and Security security tag or tags to which the virtual machine is associated during provisioning. The value is a security tag name or a list of names separated by commas. Names are case-sensitive.</p> <p>Appending a name allows you to create multiple versions of the property, which can be used separately or in combination. For example, the following properties can list security tags intended for general use, for the sales force, and for support:</p> <ul style="list-style-type: none"> ■ <code>VCNS.SecurityTag.Names</code> ■ <code>VCNS.SecurityTag.Names.sales</code> ■ <code>VCNS.SecurityTag.Names.support</code>

Enabling Visual Basic Scripts in Provisioning

Visual Basic scripts are run outside of vRealize Automation as additional steps in the machine life cycle and can be used to update the custom property values of machines. Visual Basic scripts can be used with any provisioning method.

For example, you could use a script to generate certificates or security tokens before provisioning and then use those certificates and tokens in provisioning a machine.

NOTE This information does not apply to Amazon Web Services.

When executing a Visual Basic script, the EPI agent passes all machine custom properties as arguments to the script. To return updated property values to vRealize Automation, you must place these properties in a dictionary and call a function provided by vRealize Automation.

The sample Visual Basic script `PrePostProvisioningExample.vbs` is included in the `Scripts` subdirectory of the EPI agent installation directory. This script contains a header to load all arguments into a dictionary, a body in which you can include your functions, and a footer to return updated custom properties to vRealize Automation.

The following is a high-level overview of the steps required to use Visual Basic scripts in provisioning:

- 1 A system administrator installs and configures an EPI agent for Visual Basic scripts. See *Installation and Configuration*.
- 2 A system administrator creates Visual Basic scripts and places them on the system where the EPI agent is installed.

- 3 Gather the following information for tenant administrators and business group managers for each Visual Basic script:
 - The complete path to the Visual Basic script, including the filename and extension. For example, `%System Drive%Program Files (x86)\VMware\VCAC Agents\EPI_Agents\Scripts\SendEmail.vbs`.

NOTE A fabric administrator can create a build profile by using the property sets `ExternalPreProvisioningVbScript` and `ExternalPostProvisioningVbScript` to provide this required information. Doing so makes it easier for tenant administrators and business group managers to include this information correctly in their blueprints.

- 4 Tenant administrators and business group managers use custom properties in their blueprints to call the Visual Basic scripts.

Call a Visual Basic Script from a Blueprint

Tenant administrators and business group managers can call Visual Basic scripts to run outside of vRealize Automation as additional steps in the machine life cycle and to update custom property values of machines.

NOTE If your fabric administrator creates a build profile that contains the required custom properties and you include it in your blueprint, you do not need to individually add the required custom properties to the blueprint.

Prerequisites

- Log in to the vRealize Automation console as a **tenant administrator** or **business group manager**.
- Gather the following information from your fabric administrator for each Visual Basic script:
 - The complete path to the Visual Basic script, including the filename and extension. For example, `%System Drive%Program Files (x86)\VMware\VCAC Agents\EPI_Agents\Scripts\SendEmail.vbs`.

NOTE Your fabric administrator might have provided this information in a build profile.

- Create at least one blueprint.


Procedure

- 1 Select **Infrastructure > Blueprints > Blueprints**.
- 2 Point to the blueprint to which you want to add Visual Basic scripts and click **Edit**.
- 3 Click the **Properties** tab.
- 4 (Optional) Select one or more build profiles from the **Build profiles** menu.
Build profiles contain groups of custom properties. Fabric administrators can create build profiles.
- 5 Add the Visual Basic script custom properties.
 - a Click **New Property**.
 - b Type the Visual Basic script custom property name in the **Name** text box.

Option	Description
<code>VbScript.PreProvisioning.Name</code>	Runs the Visual Basic script before a machine is provisioned.
<code>VbScript.PostProvisioning.Name</code>	Runs the Visual Basic script after a machine is provisioned.
<code>VbScript.UnProvisioning.Name</code>	Runs the Visual Basic script when a provisioned machine is destroyed.

- c Type the complete pathname, including filename and extension, to the Visual Basic script in the **Value** text box.

For example, %SystemDrive%\Program Files (x86)\VMware\VCAC Agents\EPI_Agent\Scripts\SendEmail.vbs.
 - d (Optional) Select the **Encrypted** check box to encrypt the custom property in the database.
 - e (Optional) Select the **Prompt user** check box to require the user to provide a value when they request a machine.

If you choose to prompt users for a value, any value you provide for the custom property is presented to them as the default. If you do not provide a default, users cannot continue with the machine request until they provide a value for the custom property.
 - f Click the **Save** icon ()
- 6 Repeat to add multiple Visual Basic scripts to your blueprint.
 - 7 Click **OK**.

Add Active Directory Cleanup to a Blueprint

Tenant administrators and business group managers configure the Active Directory Cleanup Plugin by using a set of custom properties to specify actions to take in Active Directory when a machine is deleted from a hypervisor.

NOTE If your fabric administrator creates a build profile that contains the required custom properties and you include it in your blueprint, you do not need to individually add the required custom properties to the blueprint.

Prerequisites

NOTE This information does not apply to Amazon Web Services.

- Log in to the vRealize Automation console as a **tenant administrator** or **business group manager**.
- Gather the following information from your fabric administrator:
 - An Active Directory account user name and password with sufficient rights to delete, disable, rename, or move AD accounts. The user name must be in domain\username format.
 - (Optional) The name of the OU to which to move destroyed machines.
 - (Optional) The prefix to attach to destroyed machines.

NOTE Your fabric administrator might have provided this information in a build profile.

- Create at least one blueprint.

Procedure

- 1 Select **Infrastructure > Blueprints > Blueprints**.
- 2 Point to the blueprint to which you want to add the Active Directory Cleanup Plugin and click **Edit**.
- 3 Click the **Properties** tab.
- 4 (Optional) Select one or more build profiles from the **Build profiles** menu.

Build profiles contain groups of custom properties. Fabric administrators can create build profiles.

- 5 Enable the Active Directory Cleanup Plugin.
 - a Click **New Property**.
 - b Type `Plugin.AdMachineCleanup.Execute` in the **Name** text box.
 - c Type **true** in the **Value** text box.
 - d (Optional) Select the **Encrypted** check box to encrypt the custom property in the database.
 - e (Optional) Select the **Prompt user** check box to require the user to provide a value when they request a machine.

If you choose to prompt users for a value, any value you provide for the custom property is presented to them as the default. If you do not provide a default, users cannot continue with the machine request until they provide a value for the custom property.

- f Click the **Save** icon ().
- 6 Add the remaining Active Directory Cleanup Plugin custom properties.

Option	Description and Value
<code>Plugin.AdMachineCleanup.UserName</code>	Type the Active Directory account user name in the Value text box. This user must have sufficient privileges to delete, disable, move, and rename Active Directory accounts. The user name must be in the format <code>domain\username</code> .
<code>Plugin.AdMachineCleanup.Password</code>	Type the password for the Active Directory account user name in the Value text box.
<code>Plugin.AdMachineCleanup.Delete</code>	Set to True to delete the accounts of destroyed machines, instead of disabling them.
<code>Plugin.AdMachineCleanup.MoveToOU</code>	Moves the account of destroyed machines to a new Active Directory organizational unit. The value is the organization unit to which you are moving the account. This value must be in <code>ou=OU, dc=dc</code> format, for example <code>ou=trash,cn=computers,dc=lab,dc=local</code> .
<code>Plugin.AdMachineCleanup.RenamePrefix</code>	Renames the accounts of destroyed machines by adding a prefix. The value is the prefix string to prepend, for example <code>destroyed_</code> .

- 7 Click **OK**.

Enabling Remote Desktop Connections

A system administrator can create a custom remote desktop protocol file that tenant administrators and business group managers use in blueprints to configure RDP settings.

The following high-level overview is the sequence of tasks required to enable machine users to connect using RDP.

- 1 A system administrator creates a custom RDP file and places it in the `Website\Rdp` subdirectory of the vRealize Automation installation directory. Provide fabric administrators, tenant administrators, and business group managers with the full pathname for the custom RDP file so that it can be included in blueprints.
- 2 (Optional) A fabric administrator creates a build profile using the property set `RemoteDesktopProtocolProperties` to compile RDP custom properties and values for tenant administrators and business group managers to include in their blueprints.
- 3 A tenant administrator or business group manager adds the RDP custom properties to a blueprint to configure the RDP settings of machines provisioned from the blueprint.
- 4 A tenant administrator or business group manager enables the **Connect using RDP or SSH** option in a blueprint.

- 5 A tenant administrator or business group manager entitles users or groups to use the **Connect using RDP or SSH** option. See *Tenant Administration*.

Configure Connections Using RDP

Tenant administrators and business group managers can use custom properties to configure RDP settings.

NOTE If your fabric administrator creates a build profile that contains the required custom properties and you include it in your blueprint, you do not need to individually add the required custom properties to the blueprint.

Prerequisites

- Log in to the vRealize Automation console as a **tenant administrator** or **business group manager**.
- Create at least one blueprint.

Procedure

- 1 Select **Infrastructure > Blueprints > Blueprints**.
- 2 Point to the blueprint to change and click **Edit**.
- 3 Click the **Properties** tab.
- 4 (Optional) Select one or more build profiles from the **Build profiles** menu.
Build profiles contain groups of custom properties. Fabric administrators can create build profiles.
- 5 Configure RDP settings.
 - a Click **New Property**.
 - b Enter the RDP custom property names in the **Name** text box and the corresponding values in the **Value** text box.

Option	Description and Value
(Required for RDP customization) RDP.File.Name	Specifies an RDP file from which to obtain settings, for example <code>My_RDP_Settings.rdp</code> . The file must reside in the <code>Website\Rdp</code> subdirectory of the vRealize Automation installation directory.
(Required for RDP customization) VirtualMachine.Rdp.Setting<i>N</i>	Configures specific RDP settings. <i>N</i> is a unique number used to distinguish one RDP setting from another. For example, to specify the Authentication Level so that no authentication requirement is specified, define the custom property <code>VirtualMachine.Rdp.Setting1</code> and set the value to <code>authentication level:i:3</code> . Use to open an RDP link to specify settings. For a list of available settings and correct syntax, see the Microsoft Windows RDP documentation.
VirtualMachine.Admin.NameCompletion	Specifies the domain name to include in the fully qualified domain name of the machine that the RDP or SSH files generate for the user interface options Connect Using RDP or Connect Using SSH option. For example, set the value to <code>myCompany.com</code> to generate the fully qualified domain name <code>my-machine-name.myCompany.com</code> in the RDP or SSH file.

- c (Optional) Select the **Encrypted** check box to encrypt the custom property in the database.

- d (Optional) Select the **Prompt user** check box to require the user to provide a value when they request a machine.

If you choose to prompt users for a value, any value you provide for the custom property is presented to them as the default. If you do not provide a default, users cannot continue with the machine request until they provide a value for the custom property.

- e Click the **Save** icon ().

- 6 Click **OK**.

What to do next

[“Enable Connections Using RDP,”](#) on page 89.

Enable Connections Using RDP

Tenant administrators and business group managers can enable RDP connections on a blueprint so that entitled users can remotely access machines provisioned from that blueprint.

Prerequisites

- Log in to the vRealize Automation console as a **tenant administrator** or **business group manager**.
- Create at least one blueprint.
- (Optional) [“Configure Connections Using RDP,”](#) on page 88.

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.

Procedure

- 1 Select **Infrastructure > Blueprints > Blueprints**.
- 2 Point to the blueprint to change and click **Edit**.
- 3 Click the **Actions** tab.
- 4 Select the **Connect using RDP or SSH** check box to allow users to connect to their machine by using an RDP or SSH connection.
- 5 Click **OK**.

All users who are entitled to connect using RDP now see the **Connect Using RDP** option for all machines they request from your blueprint. For information about entitlements, see *Tenant Administration*.

Enable Connections Using SSH

Tenant administrators and business group managers can use a custom property to enable users to connect remotely to their Linux machines using SSH.

NOTE If your fabric administrator creates a build profile that contains the required custom properties and you include it in your blueprint, you do not need to individually add the required custom properties to the blueprint.

Prerequisites

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

- Create at least one blueprint.


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.

The SSH daemon must also be running on the target machine.

Procedure

- 1 Select **Infrastructure > Blueprints > Blueprints**.
- 2 Point to the blueprint to change and click **Edit**.
- 3 Click the **Properties** tab.
- 4 (Optional) Select one or more build profiles from the **Build profiles** menu.
Build profiles contain groups of custom properties. Fabric administrators can create build profiles.
- 5 Click **New Property**.
- 6 Type **Machine.SSH** in the **Name** text box.
You may also need to add the **VirtualMachine.Admin.NameCompletion** custom property.
- 7 Type **true** in the **Value** text box.
- 8 (Optional) Select the **Encrypted** check box to encrypt the custom property in the database.
- 9 (Optional) Select the **Prompt user** check box to require the user to provide a value when they request a machine.

If you choose to prompt users for a value, any value you provide for the custom property is presented to them as the default. If you do not provide a default, users cannot continue with the machine request until they provide a value for the custom property.
- 10 (Optional) Add the **VirtualMachine.Admin.NameCompletion** custom property.

Specifies the domain name to include in the fully qualified domain name of the machine that the RDP or SSH files generate for the user interface options **Connect Using RDP** or **Connect Using SSH** option. For example, set the value to `myCompany.com` to generate the fully qualified domain name `my-machine-name.myCompany.com` in the RDP or SSH file.
- 11 Click the **Save** icon ().
- 12 Click the **Actions** tab.
- 13 Select the **Connect using RDP or SSH** check box to allow users to connect to their machine by using an RDP or SSH connection.
- 14 Click **OK**.

All users who are entitled to connect using SSH see the **Connect Using SSH** option for all machines they request from your blueprint. For information about entitlements, see *Tenant Administration*.

Managing Cloud Infrastructure

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

This chapter includes the following topics:

- [“Managing Resource Usage,”](#) on page 91
- [“Data Collection,”](#) on page 93
- [“Manage Amazon EBS Volumes,”](#) on page 95
- [“Connecting to a Cloud Machine,”](#) on page 96

Managing Resource Usage

Fabric administrators manage the resource usage of the compute resources and reservations they own, and tenant administrators and business group managers manage the resource usage of the business groups they own.

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 6-1. 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. For physical machines, this term indicates machines that are reserved for a reservation but not yet provisioned.
Managed	For physical machines, this 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.

Table 6-1. Resource Usage Terminology (Continued)

Term	Description
Used	For virtual provisioning this indicates the amount of storage in use by provisioned machines. When standard storage is employed, storage used and storage allocated are the same. However, when space-efficient storage technology is employed (for example, FlexClone or thin provisioning) storage used is typically less than storage allocated because machines only consume exactly the amount of storage they require.
Free	For virtual provisioning, this is the unused physical capacity on a storage path.

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 6-2. 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 physical machines that are reserved for use but not yet provisioned.	Fabric Administrator	Infrastructure > Machines > Reserved Machines
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) 	Infrastructure > Groups > Business Groups

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

Data Collection

vRealize Automation collects data from both 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. IaaS administrators can manually initiate data collection for infrastructure source endpoints and fabric administrators can manually initiate data collection for compute resources.

Table 6-3. Data Collection Types

Data Collection Type	Description
Infrastructure Source Endpoint Data Collection	Updates information about virtualization hosts, templates, and ISO images for virtualization environments. Updates virtual datacenters and templates for vCloud Director. Updates regions and machines provisioned on them for Amazon. Updates installed memory and CPU count for physical management interfaces.
Inventory Data Collection	Updates the record of the virtual machines whose resource use is tied to a specific compute resource, including detailed information about the networks, storage, and virtual machines. This record also includes information about unmanaged virtual machines, which are machines provisioned outside of vRealize Automation.
State Data Collection	Updates the record of the power state of each machine discovered through inventory data collection. State data collection also records missing machines that vRealize Automation manages but cannot be detected on the virtualization compute resource or cloud endpoint.
Performance Data Collection (vSphere compute resources only)	Updates the record of the average CPU, storage, memory, and network usage for each virtual machine discovered through inventory data collection.
vCNS inventory data collection (vSphere compute resources only)	Updates the record of network and security data related to vCloud Networking and Security and NSX, particularly information about security groups and load balancing, for each machine following inventory data collection.
WMI data collection (Windows compute resources only)	Updates the record of the management data for each Windows machine. A WMI agent must be installed, typically on the Manager Service host, and enabled to collect data from Windows machines.
Cost data collection (compute resources managed by vRealize Business Standard Edition only)	Updates the CPU, memory, and storage costs for each compute resource managed by vRealize Business Standard Edition. The costs of catalog items that can be provisioned by using the compute resources are updated.

Start Endpoint Data Collection Manually

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

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

Prerequisites

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

Procedure

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

Configure Compute Resource Data Collection

Fabric administrators can enable or disable data collections, configure the frequency of data collections, or manually request new data collections.

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

Prerequisites

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

Procedure

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

This is available only for vSphere integrations.

 - Select **On** to enable data collection.
 - Select **Off** to disable data collection.

- Type a number in the **Frequency** text box to configure the time interval (in hours) between performance data collections.
 - Click **Request Now** to manually start data collection.
- 7 Configure **vCNS Inventory** data collection.
- This is only available for vSphere integrations configured to use NSX or vCloud Networking and Security.
- Select **On** to enable data collection.
 - Select **Off** to disable data collection.
 - Type a number in the **Frequency** text box to configure the time interval (in hours) between vCNS Inventory data collections.
 - Click **Request Now** to manually start data collection.
- 8 Configure **Cost** data collection.
- This is only available for compute resources managed by vRealize Business Standard Edition.
- Select **On** to enable data collection.
 - Select **Off** to disable data collection.
 - Type a number in the **Frequency** text box to configure the time interval (in hours) between cost data collections.
 - Click **Request Now** to manually start data collection.
- 9 Click **OK**.



Manage Amazon EBS Volumes

A fabric administrator can destroy Amazon Elastic Block Store volumes that belong to an Amazon region compute resource in a fabric group. If you are also a business group manager, you can change the owner of an Amazon EBS volume.

Prerequisites

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

Procedure

- 1 Select **Infrastructure > Compute Resources > EBS Volumes**.
 - 2 Destroy or change the owner for a specific Amazon EBS volume.
 - Locate the Amazon EBS volume to delete and click the **Delete** icon ()
 - Locate the Amazon EBS volume to edit and click the **Edit** icon ().
Type or select a name in the **Owner** text box.
- For related information about Amazon EBS, see Amazon documentation.

Connecting to a Cloud Machine

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

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

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

Collect User Credentials for an Amazon Machine

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

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

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

Prerequisites

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

Procedure

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

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

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

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

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

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

Monitoring Workflows and Viewing Logs

7

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

Table 7-1. 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 Logs . The audit log provides details about the status of managed virtual machines and activities performed on these machines during reconfiguration. It also displays information about Amazon machine provisioning. The log includes Amazon machine provisioning, multi-machine, vCloud Networking and Security, 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 > Distributed Execution 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 > Event Logs .
View information about recent events.	IaaS administrator or Tenant administrator	Display recent events for the currently logged in user. Select Infrastructure > Recent Events

Machine Life Cycle and Workflow States for Cloud Platforms

8

Several machine workflows and workflow states are available in vRealize Automation.

Provisioning workflows are selected by the tenant administrator or business group manager when they create a blueprint. All of the available options are listed.

Table 8-1. MasterWorkflow States

Workflow State	Work Item	Description
Requested		A new machine is requested. The machine is created or registered.
AwaitingApproval		The approval process starts. The machine build is on hold until it is approved.
RegisterMachine	RegisterVM	The existing machine is registered. Attributes are set in the hypervisor. For information about registering machines, see vSphere product documentation.
BuildingMachine		The machine build is about to start. The provisioning workflow that is specified in the blueprint is being created.
MachineProvisioned	SetMachineOperations	The machine build completed successfully. Operations are being performed on the machine before it is made available for use. The allowed machine operations that are defined in the blueprint are configured on the machine.
MachineActivated		The requested machine is activated.
InstallTools	InstallTools	Hypervisor guest operating system tools are installed on the machine by the hypervisor.
Expired		The machine has expired and is turned off. After the archive period elapses, it is disposed of. The machine cannot be used. The machine can only be reactivated or disposed of in this state.
DeactivateMachine		The machine disposal process has started.
UnprovisionMachine		The machine unprovisioning process has started.

Table 8-1. MasterWorkflow States (Continued)

Workflow State	Work Item	Description
Disposing	DisposeVM	The hypervisor is disposing of the machine.
Finalized		The machine was disposed of and is about to be removed from management. The master workflow is about to stop.

Table 8-2. CloudProvisioningWorkflow States

Workflow State	Work Item	Description
CloneMachine	CloneVM	The hypervisor is cloning the machine.
CustomizeMachine	CustomizeVM	The hypervisor is configuring the machine.
InitialPowerOn	PowerOn	The machine is being powered on for the first time after the clone process is finished.
BuildComplete		The machine build process is finished.

Index

A

- Active Directory, cleanup plugin **86**
- AD, *See* Active Directory
- alerts, configuring for Amazon reservations **24**
- Amazon
 - configuring reservation alerts **24**
 - configuring VPC alerts **27**
 - creating endpoints **13**
 - creating a reservation **22**
 - creating a VPC reservation **25**
 - specifying network settings **23**
 - specifying key pairs **23**
 - specifying key pairs for VPC **26**
 - specifying reservation information **23**
 - specifying reservations for VPC **25**
 - using machine images and instance types **44**
- Amazon blueprints
 - adding custom properties **61**
 - configuring options **58**
 - specifying actions **61**
 - specifying build information **59**
- Amazon machine resources, configuring
 - instance types **45, 60**
- Amazon tools
 - managing EBS volumes **95**
 - using an elastic load balancer **41**
 - using elastic block storage **42**
 - using elastic IP addresses **41**
 - using security groups **39**
 - using virtual private cloud **40**
 - using web service regions **40**
- autoYaST
 - configuration file **47**
 - preparing for provisioning **47**
 - specifying custom scripts **48**

B

- blueprints
 - adding reservation policies **80**
 - choosing scenario **57**
 - configuring Amazon machine resources **45, 60**
 - creating Amazon blueprints) **58**
 - creating basic Red Hat OpenStack blueprints **62**
 - creating advanced blueprints **79**

- creating Amazon blueprints **58**
- creating for Linux Kickstart **68**
- creating Linux kickstart blueprints for Red Hat OpenStack **65–67, 71–74**
- publishing **78**
- specifying Amazon build information **59**
- specifying an Amazon machine image **44, 45, 59, 60**
- specifying Amazon actions **61**
- specifying an Amazon instance type **45, 59, 60**
- build profiles
 - creating **33**
 - creating property sets **32, 33**
 - using property sets **31**
- business groups
 - create **18**
 - monitoring resource usage **91**

C

- catalog items, publishing **78**
- cloud reservations **21**
- cloud reservations, endpoints **21**
- compute resources
 - configure data collection **94**
 - monitoring resource usage **91**
- costs
 - calculation **39**
 - during request and provisioning life cycle **39**
- credentials
 - storing Amazon AWS credentials **12**
 - storing RedHat OpenStack credentials **12**
- custom **10**
- custom properties
 - Active Directory cleanup **86**
 - creating build profiles **33**
 - creating property sets **32, 33**
 - Linux Kickstart provisioning **69**
 - machine provisioning **10**
 - machine.ssh **89**
 - networking **81**
 - Openstack endpoints **15**
 - order of precedence **11**
 - WIM provisioning **75**

- custom
 - properties, VirtualMachine.Admin.Name Completion **89**
- custom scripts, specifying for Linux
 - Kickstart/autoYaST **48**

D

- data collection
 - configuring **94**
 - overview **93**
 - starting endpoint data collection manually **93**

E

- elastic block storage, creating for Amazon
 - deployments **42**
- elastic IP addresses, creating for Amazon
 - deployments **41**
- endpoints
 - creating **12, 13**
 - creating for Amazon AWS **13**
 - import **16**
 - importing a CSV file **16**
 - OpenStack **14**
 - PowerVC **14**
 - starting data collection manually **93**
 - storing credentials **12**
- ESX, *See* VMware ESX
- ESXi, *See* VMware ESXi

F

- fabric groups, creating **12**
- flavors, overview **46**

G

- goal navigator, using **8**
- guest agent, inserting in a WinPE **53–56**

I

- IaaS
 - configuring **9**
 - configuring overview **9**
- infrastructure, managing **91**
- inventory data collection, configure data
 - collection **94**
- IP ranges, configuring **37**

K

- key pairs
 - creating **20**
 - exporting a private key **21**
 - overview **19**
 - uploading a private key **20**
- kickstart, specifying custom scripts **48**
- Kickstart
 - configuration file **47**
 - preparing for provisioning **47**

L

- Linux
 - configuration file **47**
 - preparing for provisioning **47**
- Linux kickstart, specifying custom scripts **48**
- Linux kickstart provisioning, creating blueprints
 - for Red Hat OpenStack **65–67, 71**
- Linux Kickstart provisioning
 - creating blueprints **68**
 - custom properties **69**
 - logs, viewing activity logs **99**

M

- machine prefixes, configuring **18**
- machine life cycles, understanding **101**
- machine provisioning, custom properties **10**
- machines
 - configuring optional Amazon settings **39**
 - connecting to Amazon machines **96**
 - connecting to cloud machines **96**
 - logging in as an Amazon administrator **96**
 - provisioning with optional configurations **31**

N

- network profiles
 - assigning to a reservation **38**
 - configuring IP ranges **37**
 - creating **36**
 - overview **36**
 - specifying network information **37**
- network settings, overview of blueprint
 - settings **80**

O

- Openstack, custom properties for endpoints **15**
- optional policies **31**

P

- PEBuilder
 - using custom scripts **52**
 - See also* VMware PEBuilder
- performance data collection, configure data
 - collection **94**
- policies, creating reservation policies **35**
- PrePostProvisioningExample.vbs, sample
 - script **84**
- property sets
 - creating **32, 33**
 - creating the XML file **32**
 - using with build profiles **31**
- provisioning
 - preparations **43**
 - WIM provisioning preparations **49**
- provisioning methods, choosing a scenario **43**
- provisioning workflow, specifying Amazon
 - blueprints **59**

R

- RDP, *See* Remote Desktop Connection
- Red Hat OpenStack
 - flavors **46**
 - support for security groups **42**
 - virtual machine image provisioning **46**
 - virtual machine images **46**
- Red Hat OpenStack blueprints
 - configuring options **62**
 - creating for Linux kickstart **65–67, 71**
 - creating for WIM **71–74**
 - specifying actions **65, 78**
 - specifying build information **63**
 - specifying custom properties **64**
- remote connections
 - configuring connect using RDP **87, 88**
 - enabling connect using RDP **89**
 - enabling connect using SSH **89**
- Remote Desktop Connection
 - configuring connect using RDP **87, 88**
 - enabling connect using RDP **89**
- requests, monitoring status **99**
- reservation
 - alerts **30**
 - create **28**
- reservation policies
 - about **79**
 - assigning to a reservation **36**
 - configuring **34**
 - create **35**
- reservations
 - allocating Cloud capacity **91**
 - assigning reservation policies **36**
 - cloud selection logic **21**
 - configuring Amazon alerts **24**
 - configuring Amazon with VPC alerts **27**
 - creating **22**
 - creating an Amazon reservation **22**
 - creating an Amazon reservation for VPC **25**
 - creating reservation policies **35**
 - creating reservations **28**
 - creating reservations for Amazon **39**
 - monitoring resource usage **91**
 - specifying Amazon reservation information **23**
 - specifying Amazon reservation information with VPC **25**
- resource monitoring
 - choosing scenario **92**
 - terminology **91**
- resource usage, monitoring **91**

S

- security groups
 - specifying in an Amazon reservation **39**
 - with Red Hat OpenStack **42**
- selection logic, cloud reservations **21**
- SSH, enabling connect using SSH **89**
- state data collection, configure data collection **94**
- static IP, *See* network profiles
- static IP addresses
 - assigning a network profile to a blueprint **80**
 - overview **36**
- Sysprep, for WIM provisioning **50**

U

- updated information **7**
- user credentials, storing **12**

V

- VB scripts, *See* Visual Basic scripts
- virtual machine image provisioning, overview **46**
- virtual machine images, overview **46**
- virtual private cloud, configuring for use with Amazon **40**
- VirtualMachine.Network, specifying a network device **81**
- VirtualMachine.NetworkN.ProfileName
 - assigning a network profile to a reservation **38**
 - using a static IP address **36**
- Visual Basic scripts
 - enabling in provisioning **84**
 - in blueprints **85**
- VMware ESXi, preparing for provisioning **47**
- VMware PEBuilder
 - creating a WinPE **52**
 - installing **51**

W

- WIM Provisioning, creating blueprints for Red Hat OpenStack **71–74**
- WIM provisioning
 - creating a PEBuilder WinPE **52**
 - custom properties **75**
 - inserting the guest agent **53–56**
 - installing PEBuilder **51**
 - preparing **49**
 - reference machine requirements **50**
 - Sysprep the reference machine **50**
- workflows, monitoring **99**

