

IaaS Configuration for Virtual Platforms

vRealize Automation 6.2



vmware®

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

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

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

docfeedback@vmware.com

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

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

Contents

IaaS Configuration for Virtual Platforms	8
Updated Information	8
Using the Goal Navigator	9
1 Configuring IaaS	11
Configuring IaaS Checklist	11
Custom Properties in Machine Provisioning, Blueprints, and Build Profiles	12
Order of Precedence for Custom Properties	13
Bringing Resources Under vRealize Automation Management	14
Choosing an Endpoint Scenario	15
Store User Credentials for Endpoints	15
Create a vSphere Endpoint	16
Create a vSphere Endpoint with Networking Integration	19
Create a NetApp ONTAP Endpoint	21
Create a KVM (RHEV) Endpoint	22
Create a Standalone Hyper-V Endpoint	22
Create an SCVMM Endpoint	23
Create a Xen Pool Endpoint	24
Create a XenServer Endpoint	25
Preparing an Endpoint CSV File for Import	25
Import a List of Endpoints	26
Create a Fabric Group	27
Configure Machine Prefixes	28
Create a Business Group	29
Virtual Reservations	30
Virtual Reservation Selection	30
Create a Reservation	32
2 Optional Configurations	36
Build Profiles and Property Sets	36
Create a Property Set	37
Create a Build Profile	39
Reservation Policies	40
Configure a Reservation Policy	41
Storage Reservation Policies	42
Configure a Storage Reservation Policy	43
Static IP Address Assignment	44
Create a Network Profile for Static IP Address Assignment	45

- Assign a Network Profile to a Reservation 47
- Cost Profiles 48
 - Cost Calculation for Virtual Machines 48
 - How Cost Is Displayed 49
 - Configure a Cost Profile 50
 - Configure a Storage Cost Profile 52
 - Cost Profiles with vRealize Business Standard Edition 53
- Enabling Users to Select Datacenter Locations 55
 - Associate a Compute Resource with a Location 56
 - Removing Datacenter Locations 56
 - Remove a Datacenter Location 57
- Active Directory Cleanup Plugin 57
- Enabling Remote Desktop Connections 58
- Enabling Connections to Machines Using SSH 58

- 3 Preparing for Provisioning 59**
 - Space-Efficient Storage for Virtual Provisioning 60
 - Preparing Your SCVMM Environment 61
 - Enabling Visual Basic Scripts in Provisioning 62
 - Choosing a Provisioning Scenario 62
 - Checklist for Virtual Provisioning by Cloning Preparation 63
 - Using vRealize Automation Guest Agent in Provisioning 66
 - Install the Guest Agent on a Linux Machine 67
 - Install the Guest Agent on a Windows Machine 69
 - Configuring the Windows Guest Agent to Trust a Server 70
 - Choosing Your Linux Kickstart Provisioning Preparation Method 71
 - Checklist for Preparing for Linux Kickstart Provisioning by using Solaris Automated Installer 72
 - Checklist for Preparing for Linux Kickstart Provisioning by using Configuration Files 78
 - Preparing for SCCM Provisioning 80
 - Create a Software Package for SCCM Provisioning 81
 - Preparing for WIM Provisioning 81
 - Reference Machine Requirements for WIM Provisioning 83
 - SysPrep Requirements for the Reference Machine 83
 - Install PEBuilder 84
 - Specify Custom Scripts in a PEBuilder WinPE 85
 - Preparing for WIM Provisioning with VirtIO Drivers 85
 - Create a WinPE Image by Using PEBuilder 86
 - Manually Insert the Guest Agent into a WinPE Image 87

- 4 Creating Blueprints 91**
 - Space-Efficient Storage for Virtual Provisioning 91
 - Choosing a Blueprint Scenario 92

- Create a Blueprint for the Basic Workflow 93
 - Specify Basic Workflow Blueprint Information 94
 - Specify Basic Workflow Build Information 95
 - Configure Basic Workflow Blueprint Machine Resources 96
 - Add Blueprint Custom Properties 97
 - Configure Blueprint Actions 99
- Create a Blueprint for Cloning 100
 - Specify Clone Blueprint Information 101
 - Specify Clone Blueprint Build Information 102
 - Configure Clone Blueprint Machine Resources 103
 - Add Custom Properties to a Blueprint 105
 - Configure Blueprint Actions 109
- Create a Linked Clone Blueprint 110
 - Specify Linked Clone Blueprint Information 111
 - Specify Linked Clone Build Information 112
 - Configure Linked Clone Machine Resources 113
 - Add Custom Properties to a Blueprint 114
 - Configure Blueprint Actions 118
- Create a Blueprint for Net App FlexClone Provisioning 119
 - Specify FlexClone Blueprint Information 120
 - Specify FlexClone Blueprint Build Information 121
 - Configure FlexClone Blueprint Machine Resources 122
 - Add Custom Properties to a Blueprint 124
 - Configure Blueprint Actions 127
- Create a Blueprint for WIM Provisioning 128
 - Specify WIM Blueprint Information 129
 - Specify WIM Blueprint Build Information 131
 - Configure WIM Blueprint Machine Resources 131
 - Add Required Properties to a WIM Blueprint 133
 - Configure Blueprint Actions 138
- Create a Blueprint for Linux Kickstart Provisioning 139
 - Specify Linux Kickstart Blueprint Information 140
 - Specify Linux Kickstart Blueprint Build Information 141
 - Configure Linux Kickstart Blueprint Machine Resources 142
 - Add Required Properties to a Linux Kickstart Blueprint 143
 - Configure Blueprint Actions 146
- Create a Blueprint for SCCM Provisioning 147
 - Specify SCCM Blueprint Information 148
 - Specify SCCM Build Information 149
 - Configure SCCM Machine Resources 150
 - Add Required Properties to a SCCM Blueprint 152
 - Configure Blueprint Actions 154

- Troubleshooting Blueprints for Clone and Linked Clone 155
- Publish a Blueprint 156

5 Configuring Advanced Blueprint Settings 158

- Assigning Datastores to Machine Volumes 158
 - Choosing a Storage Scenario 159
 - Add a Datastore to a Create, Clone, or FlexClone Blueprint 160
 - Add a Datastore to a Linked Clone Blueprint 161
 - Add a Storage Reservation Policy to a Create, Clone, or FlexClone Blueprint 162
 - Add a Storage Reservation Policy to a Linked Clone Blueprint 162
 - Allow Alternative Datastores at Provisioning Time 163
- Reservation Policies 164
 - Add a Reservation Policy to a Blueprint 165
- Configuring Network Settings 165
 - Add a Network Profile for Static IP Assignment to a Blueprint 166
 - Custom Properties for Networking 167
- Enable Users to Select Datacenter Locations 170
- Allow Requesters to Specify Machine Host Name 171
- Enabling Visual Basic Scripts in Provisioning 172
 - Call a Visual Basic Script from a Blueprint 173
- Add Active Directory Cleanup to a Blueprint 174
- Enable Machine Users to Create Snapshots 176
- RDP Connections 176
 - Configure Connections Using RDP 177
 - Enable Connections Using RDP 178
- Enable Connections Using SSH 179

6 Infrastructure Organizer 181

- Creating Blueprints for Importing Virtual Machines 181
- Choosing an Infrastructure Organizer Scenario 181
- Organize Compute Resources and Import Existing Virtual Machines 182
 - Create the Infrastructure Organizer Fabric Group 183
 - Configure Compute Resources 183
 - Bring Virtual Machines under vRealize Automation Management 184
- Import Existing Virtual Machines 185
 - Choose Compute Resources 186
 - Bring Virtual Machines under vRealize Automation Management 187

7 Managing Infrastructure 189

- Managing Resource Usage 189
 - Resource Usage Terminology 189
 - Choosing a Resource Monitoring Scenario 190

Reducing Reservation Usage by Attrition	191
Data Collection	191
Start Endpoint Data Collection Manually	192
Configure Compute Resource Data Collection	193
Update Cost Data for All Compute Resources	194
vSphere Host Maintenance	195
Decommissioning a Storage Path	195
Disable a Storage Path	195
8 Monitoring Workflows and Viewing Logs	197
9 Appendix A: Machine Life Cycle and Workflow States	198

IaaS Configuration for Virtual Platforms

The Infrastructure as a service (IaaS) feature of VMware vRealize™ Automation enables the rapid modeling and provisioning of servers and desktops across virtual and physical, private and public, or hybrid cloud infrastructure. This documentation is specifically designed for virtual platform integrations with vRealize Automation.

This documentation guides you through the following processes:

- Configuring IaaS features and policies
- Preparing for provisioning
- Bringing resources under vRealize Automation management
- 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 Virtual Platforms* is updated with each release of the product or when necessary.

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

Revision	Description
001645-06	Updated Create an SCVMM Endpoint .
001645-05	Added information about installing Solaris operating systems as part of the provisioning process. See Choosing Your Linux Kickstart Provisioning Preparation Method .
001645-04	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 . Added information about custom specification settings to the following topics: <ul style="list-style-type: none"> ■ Specify Clone Blueprint Build Information ■ Specify FlexClone Blueprint Build Information ■ Specify Linked Clone Build Information
001645-03	<ul style="list-style-type: none"> ■ Added vSphere Host Maintenance. ■ Added Preparing Your SCVMM Environment. ■ Updated Add Required Properties to a WIM Blueprint and Custom Properties for WIM Blueprints.
001645-02	<ul style="list-style-type: none"> ■ Added Using vRealize Automation Guest Agent in Provisioning. ■ Updated Install PEBuilder. ■ Updated Create a WinPE Image by Using PEBuilder. ■ Updated Cost Profiles with vRealize Business Standard Edition. ■ Updated Custom Properties for SCCM Blueprints. ■ Editorial updates.
001645-01	<ul style="list-style-type: none"> ■ Added Understanding vSwap Allocation Checking for vCenter Server Endpoints. ■ Updated Custom Properties for Basic Workflow Blueprints.
001645-00	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 Checklist](#)
- [Custom Properties in Machine Provisioning, Blueprints, and Build Profiles](#)
- [Bringing Resources Under vRealize Automation Management](#)
- [Configure Machine Prefixes](#)
- [Create a Business Group](#)
- [Virtual Reservations](#)

Configuring IaaS Checklist

Required and optional configurations are performed by IaaS administrators, tenant administrators, and fabric administrators to fully implement and customize the IaaS feature of 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 .	IaaS administrator
<input type="checkbox"/> Configure the machine prefixes used to create names for machines provisioned through vRealize Automation. See Configure Machine Prefixes .	Fabric administrator
<input type="checkbox"/> Create at least one business group of users who need to request machines. See Create a Business Group .	Tenant administrator
<input type="checkbox"/> Create at least one reservation to allocate resources to a business group. See Create a Reservation .	Fabric administrator

Table 1-1. Configuring IaaS checklist (Continued)

Task	Required Role
<input type="checkbox"/> Configure optional policies and settings. See Chapter 2 Optional Configurations .	Fabric administrator
<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, if you intend to provision by cloning, you need to create the template and the customization object. See Choosing a Provisioning Scenario .	Outside of vRealize Automation
<input type="checkbox"/> Create and publish machine blueprints. See Choosing a Blueprint Scenario and Chapter 5 Configuring Advanced Blueprint Settings .	<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 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

To allow vRealize Automation to manage your infrastructure, IaaS administrators create endpoints, store administrator-level user credentials for those endpoints, and add compute resources to a fabric group. The procedure for creating endpoints differs depending on your environment.

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 procedures to create endpoints and store the corresponding credentials differ slightly.

Table 1-2. Choosing an Endpoint Scenario

Scenario	Store Credentials	Create Endpoint
Connect to a vSphere environment.	Store User Credentials for Endpoints	Create a vSphere Endpoint
Integrate with vCloud Networking and Security or NSX (vSphere only).	Store User Credentials for Endpoints	Create a vSphere Endpoint with Networking Integration
Use Net App FlexClone technology for storage (vSphere only).	Store User Credentials for Endpoints	Create a NetApp ONTAP Endpoint
Connect to a KVM (RHEV) (Kernel-based Virtual Machine with Red Hat Enterprise Virtualization) environment.	Store User Credentials for Endpoints	Create a KVM (RHEV) Endpoint
Connect to a standalone Hyper-V environment.	A system administrator must install a proxy agent with stored credentials that correspond to your endpoint. See <i>Installation and Configuration</i> .	Create a Standalone Hyper-V Endpoint
Connect to a Hyper-V and SCVMM (Microsoft Center Virtual Machine Manager) environment.	Store User Credentials for Endpoints	Create an SCVMM Endpoint
Connect to a Xen pool on a XenServer.	A system administrator must install a proxy agent with stored credentials that correspond to your endpoint. See <i>Installation and Configuration</i> .	Create a Xen Pool Endpoint
Connect to a XenServer environment.	A system administrator must install a proxy agent with stored credentials that correspond to your endpoint. See <i>Installation and Configuration</i> .	Create a XenServer Endpoint
Import a list of endpoints.	Store User Credentials for Endpoints	<ul style="list-style-type: none"> ■ Preparing an Endpoint CSV File for Import ■ Import a List of Endpoints

Store User Credentials for Endpoints

You must store administrator-level credentials for your environment so that vRealize Automation can communicate with your 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.

Credentials used to manage endpoints must have administrator privileges.

Note If you are working in a vSphere environment, your system administrator might have configured the vSphere proxy agent to use integrated credentials. If this is the case, you might not need to store additional credentials for your vSphere endpoint.

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 in the **Name** text box.
- 4 (Optional) Enter a description in the **Description** text box.
- 5 Type the user name in the **User name** text box.

Option	Description
vSphere	Must be in domain\username format, for example mycompany\admin . The credentials must have permission to modify custom attributes.
vCloud Networking and Security (vSphere only)	Must be in domain\username format, for example mycompany\admin .
NSX (vSphere only)	Does not have to include the domain name, for example admin .
Hyper-V (SCVMM)	Must be in domain\username format, for example mycompany\admin .
KVM (RHEV)	Must be in username@domain format, for example admin@mycompany .
NetApp ONTAP	Does not have to include the domain name, for example admin .

- 6 Type the password in the **Password** text boxes.
- 7 Click the **Save** icon (✔).

What to do next

Now that your credentials are stored, you are ready to create an endpoint. See [Choosing an Endpoint Scenario](#).

Create a vSphere Endpoint

IaaS administrators create endpoints that allow vRealize Automation to communicate with the vSphere environment and discover compute resources, collect data, and provision machines.

Note You do not have to create endpoints one at a time. You can also create endpoints by importing them from a CSV file.

Prerequisites

- Log in to the vRealize Automation console as an **IaaS administrator**.
- A system administrator must install a vSphere agent to correspond to your vSphere endpoints and credentials. See *Installation and Configuration*.
- The endpoint name you configure in vRealize Automation must match the endpoint name provided to the vSphere proxy agent during installation. If you do not know the endpoint name your system administrator provided to the proxy agent, see [Attached vSphere Endpoint Cannot be Found](#).
- If your system administrator did not configure the proxy to use integrated credentials, you must store administrator-level credentials for your endpoint. See [Store User Credentials for Endpoints](#).

Procedure

1 Select **Infrastructure > Endpoints > Endpoints**.

2 Select **New Endpoint > Virtual > vSphere**.

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

This must match the endpoint name provided to the vSphere proxy agent during installation or data collection fails.

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

5 Enter the URL for the vCenter Server instance in the **Address** text box.

The URL must be of the type: **https://hostname/sdk** or **https://IP_address/sdk**.

For example, **https://vsphereA/sdk**.

6 Select the **Credentials** for the endpoint.

If your system administrator configured the vSphere proxy agent to use integrated credentials, you can select the **Integrated** credentials.

7 Do not select **Specify manager for network and security platform**.

This is for implementations that use vCloud Networking and Security or NSX and requires additional configuration.

8 (Optional) Add any custom properties.

9 Click **OK**.

vRealize Automation can now discover your compute resources.

Important Renaming vSphere assets after discovery can cause provisioning to fail.

What to do next

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

See [Create a Fabric Group](#).

Understanding vSwap Allocation Checking for vCenter Server Endpoints

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

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

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

Consider the following behaviors for vSwap allocation checks:

- The swap file is located on the datastore that contains the virtual machine. Alternate vCenter Server configurations for locating swap files on a dedicated or different datastore are not supported.
- Swap size is considered when creating or reconfiguring a virtual machine. The maximum swap size is the size of the virtual machine's memory.
- Reserved values for vRealize Automation storage reservations in a host must not exceed the physical capacity of the compute resource.
- When creating a reservation, the sum of the reserved values must not exceed the available storage space.
- Resource pool or host level or virtual machine level memory reservations on vSphere are not collected from the vSphere endpoint and not considered during the calculations on vRealize Automation.
- vSwap does not validate the swap space that is available during power on operations for existing machines.
- You must re-run data collection to capture any changes made to the vSphere endpoint relative to vSwap.

Attached vSphere Endpoint Cannot be Found

When data collection fails for a vSphere endpoint, it is often due to a mismatch between the proxy name and the endpoint name.

Problem

Data collection fails for a vSphere endpoint. The log messages return an error similar to the following:

```
This exception was caught: The attached endpoint
    'vCenter' cannot be found.
```

Cause

The endpoint name you configure in vRealize Automation must match the endpoint name provided to the vSphere proxy agent during installation. Data collection fails for a vSphere endpoint if there is a mismatch between the endpoint name and the proxy agent name. Until an endpoint with a matching name is configured, the log messages return an error similar to the following:

```
This exception was caught: The attached endpoint
    'expected endpoint name' cannot be found.
```

Solution

- 1 Select **Infrastructure > Monitoring > Log**.
- 2 Look for an Attached Endpoint Cannot be Found error message.

For example,

```
This exception was caught: The attached endpoint
    'expected endpoint name' cannot be found.
```

- 3 Edit your vSphere endpoint to match the expected endpoint name shown in the log message.
 - a Select **Infrastructure > Endpoints > Endpoints**.
 - b Click the name of the endpoint to edit.
 - c Enter the expected endpoint name in the **Name** text box.
 - d Click **OK**.

The proxy agent can communicate with the expected endpoint and data collection is successful.

Create a vSphere Endpoint with Networking Integration

IaaS administrators create endpoints that allow vRealize Automation to communicate with the vSphere environment, a vCloud Networking and Security or an NSX instance.

Note You do not have to create endpoints one at a time. You can also create endpoints by importing them from a CSV file.

Prerequisites

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

- A system administrator must install a vSphere agent to correspond to your vSphere endpoints and credentials. See *Installation and Configuration*.
- The endpoint name you configure in vRealize Automation must match the endpoint name provided to the vSphere proxy agent during installation. If you do not know the endpoint name your system administrator provided to the proxy agent, see [Attached vSphere Endpoint Cannot be Found](#).
- If your system administrator did not configure the proxy to use integrated credentials, you must store administrator-level credentials for your endpoint. See [Store User Credentials for Endpoints](#).

Procedure

1 Select **Infrastructure > Endpoints > Endpoints**.

2 Select **New Endpoint > Virtual > vSphere**.

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

This must match the endpoint name provided to the vSphere proxy agent during installation or data collection fails.

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

5 Type the URL for the vCenter Server instance in the **Address** text box.

The URL must be of the type: **https://hostname/sdk** or **https://IP_address/sdk**.

For example, **https://vsphereA/sdk**.

6 Select the **Credentials** for the endpoint.

If your system administrator configured the vSphere proxy agent to use integrated credentials, you can select the **Integrated** credentials.

7 Configure a networking solution platform.

a Select **Specify manager for network and security platform**.

b Type the URL for the vCloud Networking and Security or NSX instance in the **Address** text box.

The URL must be of the type: **https://hostname** or **https://IP_address**.

For example, **https://vCNSa**.

c Select the **Credentials** for the endpoint.

8 (Optional) Add any custom properties.

9 Click **OK**.

vRealize Automation can now discover your compute resources.

Important Renaming vSphere assets after discovery can cause provisioning to fail.

What to do next

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

See [Create a Fabric Group](#).

Create a NetApp ONTAP Endpoint

IaaS administrators create endpoints to allow vRealize Automation to communicate with storage devices that use Net App FlexClone technology.

Note You do not have to create endpoints one at a time. You can also create endpoints by importing them from a CSV file.

Prerequisites

- Log in to the vRealize Automation console as an **IaaS administrator**.
- [Store User Credentials for Endpoints](#).

Procedure

- 1 Select **Infrastructure > Endpoints > Endpoints**.
- 2 Select **New Endpoint > Storage > NetApp ONTAP**.
- 3 Enter a name in the **Name** text box.
- 4 (Optional) Enter a description in the **Description** text box.
- 5 Type the URL for the endpoint in the **Address** text box.
The URL must be of the type: *FQDN* or *IP_address*.
For example: **netapp-1.mycompany.local**.
- 6 Select the **Credentials** for the endpoint.
If you did not already store the credentials, you can do so now.
- 7 (Optional) Add any custom properties.
- 8 Click **OK**.

vRealize Automation can now discover your compute resources.

Important Renaming vSphere assets after discovery can cause provisioning to fail.

What to do next

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

See [Create a Fabric Group](#).

Create a KVM (RHEV) Endpoint

IaaS administrators create endpoints to allow vRealize Automation to communicate with the KVM (RHEV) environment and discover compute resources, collect data, and provision machines.

Note You do not have to create endpoints one at a time. You can also create endpoints by importing them from a CSV file.

Prerequisites

- Log in to the vRealize Automation console as an **IaaS administrator**.
- [Store User Credentials for Endpoints](#).

Procedure

- 1 Select **Infrastructure > Endpoints > Endpoints**.
- 2 Select **New Endpoint > Virtual > KVM (RHEV)**.
- 3 Enter a name in the **Name** text box.
- 4 (Optional) Enter a description in the **Description** text box.
- 5 Type the URL for the endpoint in the **Address** text box.
The URL must be of the type: **https://FQDN** or **https://IP_address**
For example, **https://mycompany-kvmrhev1.mycompany.local**.
- 6 Select the **Credentials** for the endpoint.
If you did not already store the credentials, you can do so now.
- 7 (Optional) Add any custom properties.
- 8 Click **OK**.

vRealize Automation can now discover your compute resources.

What to do next

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

See [Create a Fabric Group](#).

Create a Standalone Hyper-V Endpoint

IaaS administrators create endpoints to allow vRealize Automation to communicate with the Hyper-V environment and discover compute resources, collect data, and provision machines.

Prerequisites

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

- A system administrator must install a proxy agent with stored credentials that correspond to your endpoint. See *Installation and Configuration*.

Procedure

- 1 Select **Infrastructure > Endpoints > Agents**.
- 2 Enter the fully qualified DNS name of your Hyper-V server in the **Compute resource** text box.
- 3 Select the proxy agent that your system administrator installed for this endpoint from the **Proxy agent name** drop-down menu.
- 4 (Optional) Enter a description in the **Description** text box.
- 5 Click **OK**.

vRealize Automation can now discover your compute resources.

What to do next

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

See [Create a Fabric Group](#).

Create an SCVMM Endpoint

IaaS administrators create endpoints to allow vRealize Automation to communicate with the SCVMM environment and discover compute resources, collect data, and provision machines.

For related information, see [Preparing Your SCVMM Environment](#).

Note You do not have to create endpoints one at a time. You can also create endpoints by importing them from a CSV file.

Prerequisites

- Log in to the vRealize Automation console as an **IaaS administrator**.
- You must install and configure a DEM agent to manage your Hyper-V (SCVMM) endpoint. For information, see *SCVMM Requirements* in *Installation and Configuration*.
- [Store User Credentials for Endpoints](#).

Procedure

- 1 Select **Infrastructure > Endpoints > Endpoints**.
- 2 Select **New Endpoint > Virtual > Hyper-V (SCVMM)**.
- 3 Enter a name in the **Name** text box.
- 4 (Optional) Enter a description in the **Description** text box.

- 5 Enter the URL for the endpoint in the **Address** text box.
The URL must be of the type: *FQDN* or *IP_address*.
For example: `mycompany-scvmm1.mycompany.local`.
- 6 Select the **Credentials** for the endpoint.
If you did not already store the credentials, you can do so now.
- 7 (Optional) Add any custom properties.
- 8 Click **OK**.

vRealize Automation can now discover your compute resources.

What to do next

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

See [Create a Fabric Group](#).

Create a Xen Pool Endpoint

IaaS administrators create endpoints to allow vRealize Automation to communicate with the Xen pool master and discover compute resources, collect data, and provision machines.

Prerequisites

- Log in to the vRealize Automation console as an **IaaS administrator**.
- A system administrator must install a proxy agent with stored credentials that correspond to your endpoint. See *Installation and Configuration*.

Procedure

- 1 Select **Infrastructure > Endpoints > Agents**.
- 2 Enter the fully qualified DNS name of your Xen pool master in the **Compute resource** text box.

Note Do not enter the name of the Xen pool. You must enter the name of the pool master.

- 3 Select the proxy agent that your system administrator installed for this endpoint from the **Proxy agent name** drop-down menu.
- 4 (Optional) Enter a description in the **Description** text box.
- 5 Click **OK**.

vRealize Automation can now discover your compute resources.

What to do next

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

See [Create a Fabric Group](#).

Create a XenServer Endpoint

IaaS administrators create endpoints to allow vRealize Automation to communicate with the XenServer environment and discover compute resources, collect data, and provision machines.

Prerequisites

- Log in to the vRealize Automation console as an **IaaS administrator**.
- A system administrator must install a proxy agent with stored credentials that correspond to your endpoint. See *Installation and Configuration*.

Procedure

- 1 Select **Infrastructure > Endpoints > Agents**.
- 2 Enter the fully qualified DNS name of your XenServer server in the **Compute resource** text box.
- 3 Select the proxy agent that your system administrator installed for this endpoint from the **Proxy agent name** drop-down menu.
- 4 (Optional) Enter a description in the **Description** text box.
- 5 Click **OK**.

vRealize Automation can now discover your compute resources.

What to do next

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

See [Create a Fabric Group](#).

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 1-3. 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](#).

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.

Virtual Reservations

A virtual reservation is a share of the memory, CPU, networking, and storage resources of one compute resource allocated to a particular business group.

To provision virtual machines, a business group must have at least one reservation on a virtual compute resource. Each reservation is for one business group only, but a business group can have multiple reservations on a single compute resource, or multiple reservations on compute resources of different types.

In addition to the share of fabric resources allocated to the business group, a reservation can define policies, priorities, and quotas that determine machine placement.

Virtual Reservation Selection

When a member of a business group requests a virtual machine, resources must be allocated on one of the group's reservations to provision the machine.

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.

A generic virtual blueprint can be provisioned on any type of virtual reservation.

- The reservation must be enabled.
- The compute resource must be accessible and not in maintenance mode.
- 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 sufficient unallocated memory and storage resources to provision the machine.

When a virtual reservation's machine quota, memory, or storage is fully allocated, no further virtual machines can be provisioned from it. Resources may be reserved beyond the physical capacity of a virtualization compute resource (overcommitted), but when the physical capacity of a compute resource is 100% allocated, no further machines can be provisioned on any reservations with that compute resource until the resources are reclaimed.

- If 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 or request specifies a location, the compute resource must be associated with that location.

If the value of the custom property *VRM.Datacenter.Policy* is **Exact** and there is no reservation for a compute resource associated with that location that satisfies all the other criteria, then provisioning fails.

If the value of *VRM.Datacenter.Policy* is **NotExact** and there is no reservation for a compute resource associated with that location that satisfies all the other criteria, provisioning can proceed on another reservation regardless of location. This option is the default.

- If the blueprint or request specifies the custom property *VirtualMachine.Host.TpmEnabled*, trusted hardware must be installed on the compute resource for the reservation.
- 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, you can use reservation policies to limit provisioning to compute resources with a specific template for cloning.

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.

If multiple storage paths are available on a reservation with sufficient capacity to provision the machine volumes, storage paths are selected according to the following logic:

- If the blueprint or request specifies a storage reservation policy, the storage path must belong to that storage reservation policy.

If the value of the custom property *VirtualMachine.DiskN.StorageReservationPolicyMode* is **NotExact** and there is no storage path with sufficient capacity within the storage reservation policy, then provisioning can proceed with a storage path outside the specified storage reservation policy. The default value of *VirtualMachine.DiskN.StorageReservationPolicyMode* is **Exact**.

- Storage paths with higher priority are selected over storage paths with lower priority.
- If multiple storage paths have the same priority, machines are distributed among storage paths in round-robin fashion.

Create a Reservation

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

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 vSphere reservation, but not a KVM (RHEV) reservation, cannot request a KVM (RHEV) virtual machine. The group must be allocated a reservation specifically for KVM (RHEV) resources.

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](#).

Procedure

1 [Specify Reservation Information](#)

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

2 [Specify Memory, Storage, and Network Resources](#)

Configure the memory, storage, and network paths available to machines provisioned through this reservation.

3 [Configure Alerts](#)

Optionally, you can configure alerts to send email notifications whenever reserved resources run 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 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 > Virtual** and select the type of reservation you are creating.

- 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.
Only templates located on the cluster you select are available for cloning with this reservation.
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 Memory, Storage, and Network Resources

Configure the memory, storage, and network paths available to machines provisioned through this reservation.

You can select a FlexClone datastore in your reservation if you have a vSphere environment and storage devices that use Net App FlexClone technology. SDRS is not supported for FlexClone storage devices.

Prerequisites

[Specify Reservation Information.](#)

Procedure

- 1 Click the **Resources** tab.
- 2 Specify the amount of memory, in GB, to be allocated to this reservation from the Memory table.

3 Configure a storage path in the **Storage** table.

- a Select a storage path from the **Storage Path** column.

Option	Description
Storage Distributed Resource Scheduler (SDRS) Storage	For integrations that use SDRS, you can select a storage cluster to allow SDRS to automatically handle storage placement and load balancing for machines provisioned from this reservation. The SDRS Automation Mode must be set to Automatic. Select a datastore within the cluster for standalone datastore behavior.
Individual Datastore Selection	Select a standalone datastore.

- b (Optional) Select a storage endpoint from the **Endpoint** drop-down menu to specify a storage device that uses FlexClone technology.

SDRS is not supported for FlexClone storage devices.

The FlexClone option is visible in the endpoint column if a NetApp ONTAP endpoint exists and if the host is virtual. If there is a NetApp ONTAP endpoint, the reservation page displays the endpoint assigned to the storage path. When you add, update, or delete an endpoint for a storage path, the change is visible in all the applicable reservations.

When you add, update, or delete an endpoint for a storage path, the change is visible in the reservation page.

- c Type a value in **This Reservation Reserved** to specify how much storage to allocate to this reservation.
- d Specify the **Priority** for the storage path.

The priority is used for multiple storage paths. A storage path with priority 0 is used before a path with priority 1.

- e Repeat this step to configure clusters and datastores as needed.

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 email notifications whenever reserved resources run 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 Memory, Storage, and Network Resources.](#)

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 storage reservation policies and 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](#)
- [Reservation Policies](#)
- [Storage Reservation Policies](#)
- [Static IP Address Assignment](#)
- [Cost Profiles](#)
- [Enabling Users to Select Datacenter Locations](#)
- [Active Directory Cleanup Plugin](#)
- [Enabling Remote Desktop Connections](#)
- [Enabling Connections to Machines Using SSH](#)

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](#)

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

2 [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.

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.

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.

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](#)

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](#)

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](#).

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](#).

Storage Reservation Policies

Fabric administrators create storage reservation policies to allow tenant administrators and business group managers to assign the volumes of a virtual machine to different datastores for the vSphere, KVM (RHEV), and SCVMM platform types.

Assigning the volumes of a virtual machine to different datastores allows tenant administrators and business group managers to control and use storage space more effectively. For example, they might deploy the operating system volume to a slower, less expensive datastore and the database volume to a faster datastore.

Tenant administrators and business group managers can assign a single datastore or a storage reservation policy that represents multiple datastores to a volume. When they assign a single datastore to a volume, vRealize Automation uses that datastore at provisioning time, if possible. When they assign a storage reservation policy to a volume, vRealize Automation uses one of its datastores at provisioning time, if possible.

A storage reservation policy is essentially a tag applied to one or more datastores by a fabric administrator to group datastores that have similar characteristics, such as speed or price. A datastore can be assigned to only one storage reservation policy at a time, but a storage reservation policy can have many different datastores. A fabric administrator creates a storage reservation policy and assigns it to one or more datastores. A tenant administrator or business group manager then assigns the storage reservation policy to a volume in a virtual blueprint. When a user requests a virtual machine that uses the blueprint, vRealize Automation uses the storage reservation policy specified in the blueprint to select a datastore for the machine's volume.

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.

Configure a Storage Reservation Policy

Fabric administrators create vRealize Automation storage reservation policies to group datastores that have similar characteristics, such as speed or price. After you create the storage reservation policy, you then need to populate it with datastores before tenant administrators and business group managers can use the policy in a blueprint.

vCloud Director 5.5 does not allow a virtual machine to have disks in different storage tiers. For example, if you change the storage reservation policy for one volume in a vApp, vRealize Automation changes the storage reservation policy for all the volumes in that vApp. You cannot provision a virtual machine with volumes on different types of storage.

Prerequisites

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

Procedure

1 [Create a Storage Reservation Policy](#)

When you create a storage reservation policy, you are creating an empty container that you use to group datastores that have similar characteristics, such as speed or price.

2 [Assign a Storage Reservation Policy to a Datastore](#)

After the storage reservation policy is created, populate it with datastores. A datastore can belong to only one storage reservation policy. Add multiple datastores to create a group of datastores for use with a blueprint.

Create a Storage Reservation Policy

When you create a storage reservation policy, you are creating an empty container that you use to group datastores that have similar characteristics, such as speed or price.

Prerequisites

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

Procedure

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

Assign a Storage Reservation Policy to a Datastore

After the storage reservation policy is created, populate it with datastores. A datastore can belong to only one storage reservation policy. Add multiple datastores to create a group of datastores for use with a blueprint.

Prerequisites

Create a storage reservation policy.

Procedure

- 1 Select **Infrastructure > Compute Resources > Compute Resources**.
- 2 Point to a compute resource and click **Edit**.
- 3 Click the **Configuration** tab.
- 4 Locate the datastore to add to your storage reservation policy in the Storage table.
- 5 Click the **Edit** icon (✎).
- 6 Select your storage reservation policy from the **Storage Reservation Policy** column drop-down menu.
- 7 Click the **Save** icon (✔).
- 8 Click **OK**.
- 9 (Optional) Assign additional datastores to your storage reservation policy.

Tenant administrators and business group managers can now use your storage reservation policy in blueprints.

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

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

2 Configure IP Ranges

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](#).

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
ip_address	An IP address
mname	Name of a managed machine in vRealize Automation. If the field is empty, defaults to no name.
status	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](#).

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 Profiles

Fabric administrators can associate compute resources and physical machines with cost profiles to enable calculation of a machine's cost. The cost is displayed to machine owners, requesters, approvers, and administrators at various points in the request and provisioning life cycle.

A cost profile includes the following values for daily cost:

- Cost per GB of memory capacity specified in the virtual blueprint or installed in the physical machine
- Cost per CPU specified in the virtual blueprint or installed in the physical machine
- Cost per GB of storage capacity as specified in the virtual blueprint (not used for physical machines, because storage attached to physical machines is not discovered or tracked)

For finer definition of storage cost for virtual machines, you can also associate each known datastore on a compute resource with a storage cost profile. A storage cost profile contains only a daily cost per GB of storage. If you assign a storage cost profile to a datastore, this storage cost overrides the storage cost in the cost profile assigned to the compute resource.

For virtual machines, the machine cost is calculated from the cost profile and storage cost profile on the compute resource, the resources it consumes, and the daily blueprint cost. You can use the blueprint cost to represent a markup for using the machine in addition to the resources that the machine consumes, for example to account for the cost of specific software deployed with that blueprint.

For physical machines, the machine cost is calculated from the cost profile on the machine, the CPU and memory on the machine, and the daily blueprint cost. You can use the blueprint cost to represent such factors as storage cost or additional costs for using the machine.

You cannot apply cost profiles to machines provisioned on Amazon Web Services or Red Hat OpenStack. For machines provisioned on these cloud platforms, the only cost factor is the daily cost in the blueprint from which it was provisioned. The cost for vCloud Director vApps includes any cost profile and storage cost profile on the virtual datacenter and the blueprint cost.

Cost Calculation for Virtual Machines

Several factors can contribute to the calculation of cost for virtual machines.

The daily cost of a machine is based on the cost of the memory, CPU, and storage resources that it consumes, and the blueprint cost.

Table 2-1. Daily Cost Drivers

Cost Driver	Calculated Cost
Memory	Allocated memory (GB) multiplied by the daily memory cost per gigabyte specified in the cost profile associated with the compute resource on which the machine is provisioned
CPUs	Number of allocated CPUs multiplied by the daily CPU cost specified in the cost profile associated with the compute resource on which the machine is provisioned
Storage	<p>Allocated storage (GB) multiplied by the storage cost per gigabyte specified in one of the following places:</p> <ul style="list-style-type: none"> ■ The cost profile associated with the compute resource on which the machine is provisioned ■ The storage cost profile on the datastore on which the machine is provisioned <p>The cost specified in a storage cost profile for a specific datastore overrides the storage cost specified in a cost profile on the compute resource.</p>
Blueprint cost	<p>The value for daily cost specified in the machine blueprint is added to the total cost of the machine.</p> <p>This value can represent a markup for using the machine in addition to the resources consumed by the machine.</p>

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 is provisioned.

How Cost Is Displayed

The 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.

Daily cost and lease cost appear at all stages of the lifecycle. After the machine is provisioned, the cost-to-date also appears.

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

Life Cycle Stage	Value Displayed for Cost
Viewing the catalog item details before a request	<p>Projected costs based on the cost profile, the values for machine resources and lease duration specified in the blueprint, and the daily blueprint cost.</p> <p>The values of some of the cost drivers can be unknown before a machine is requested:</p> <ul style="list-style-type: none"> ■ The blueprint can specify a range for machine resources. The requester can specify any value in the blueprint range. ■ More than one reservation can be available for provisioning the machine. If these reservations have different compute resources, different cost profiles can apply. ■ More than one datastore can be available for provisioning the machine, either on the same reservation or on different reservations, that have different storage cost profiles. ■ If more than one reservation or datastore is available, the daily cost appears as a range with the lower bound being the minimum resource use multiplied by the lowest resource cost for memory, CPU, and storage, and the upper bound being the maximum resource use multiplied by the highest cost. <p>If the blueprint specifies a range for the lease duration or the daily cost is a range, the lease cost appears as a range with the lower bound being the minimum lease duration multiplied by the minimum cost, and the upper bound being the maximum lease duration multiplied by the maximum cost.</p>
Completing the request form	<p>Projected costs based on the cost profile, the values for machine resources and lease duration specified in the request form and the blueprint cost.</p> <p>The daily cost can be a range if more than one reservation is available for provisioning the machine and these reservations have different compute resources that vary in cost. Alternatively, more than one datastore can be available for provisioning the machine, either on the same reservation or on different reservations, that vary in cost.</p> <p>By default, the request form is populated with the minimum values for machine resources and lease duration, and the values for daily and lease costs are calculated accordingly. The costs are updated as the user edits the relevant values in the request form.</p>
Viewing details of a submitted request or approving a request	<p>Projected costs based on the requested machine resources, lease duration, and blueprint cost.</p> <p>The daily cost can be a range if more than one reservation is available for provisioning the machine and these reservations have different compute resources that vary in cost. Alternatively, more than one datastore can be available for provisioning the machine, either on the same reservation or on different reservations, that vary in cost.</p> <p>Depending on the approval policy, an approver can edit values for machine resources or lease duration that affect the cost, but the cost information is not updated in the request.</p>
Viewing the details of a provisioned machine	<p>Actual daily cost, lease cost, and cost-to-date based on the resource use of the provisioned machine, lease duration, blueprint cost, and number of days since the machine was provisioned.</p>

Configure a Cost Profile

Fabric administrators can create cost profiles and associate them with compute resources to enable calculation of a machine's cost.

Prerequisites

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

Procedure

1 [Create a Cost Profile](#)

Fabric administrators create cost profiles to define daily costs for memory, storage, and CPU for compute resources.

2 [Assign a Cost Profile to a Compute Resource](#)

Fabric administrators can assign a cost profile to a compute resource to apply those costs to all machines provisioned on that compute resource.

Create a Cost Profile

Fabric administrators create cost profiles to define daily costs for memory, storage, and CPU for compute resources.

The values for daily cost specified in a cost profile can be up to four decimal places, although the values are rounded to two decimal places when displayed to end users in the catalog.

Prerequisites

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

Procedure

1 Select **Infrastructure > Compute Resources > Cost Profiles**.

2 Click **New Cost Profile**.

3 Enter a name and, optionally, a description.

4 Specify the costs for the machine resources on a compute resource.

a Type a value for the daily cost per gigabyte of memory in the **Memory Cost (per GB)** text box.

b Type a value for the daily cost per gigabyte of storage in the **Storage Cost (per GB)** text box.

c Type a value for the daily cost per CPU in the **CPU Cost** text box.

It is not necessary to specify a cost value for all the resources, only the resources that you want to include in cost calculations.

5 Click the **Save** icon (✔).

Assign a Cost Profile to a Compute Resource

Fabric administrators can assign a cost profile to a compute resource to apply those costs to all machines provisioned on that compute resource.

Prerequisites

[Create a Cost Profile](#).

Procedure

- 1 Select **Infrastructure > Compute Resources > Compute Resources**.
- 2 Point to a compute resource and click **Edit**.
- 3 Click the **Configuration** tab.
- 4 Select a cost profile from the **Cost Profile** drop-down menu.

The costs for memory, storage, and CPU appear, and the cost for each datastore is populated with the storage cost from the cost profile, unless a storage cost profile is applied to that datastore.

- 5 Click **OK**.

If reservations for this compute resource exist, the costs for each datastore are also populated based on the cost profile.

Configure a Storage Cost Profile

Fabric administrators can create storage cost profiles and associate them with storage paths to provide more detailed storage cost information. Storage cost profiles override the storage cost in any cost profile assigned to the compute resource.

Prerequisites

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

Procedure

1 [Create a Storage Cost Profile](#)

Fabric administrators create storage cost profiles to define daily costs for specific datastores.

2 [Assign a Storage Cost Profile to a Datastore](#)

Fabric administrators can assign a storage cost profile to a datastore to apply those costs to all machines provisioned on that datastore.

Create a Storage Cost Profile

Fabric administrators create storage cost profiles to define daily costs for specific datastores.

The values for daily cost specified in a storage cost profile can be up to four decimal places. The values are rounded to two decimal places when displayed to end users in the catalog.

Prerequisites

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

Procedure

- 1 Select **Infrastructure > Compute Resources > Cost Profiles**.
- 2 Click **New Storage Cost Profile**.
- 3 Enter a name and, optionally, a description.

- 4 Type a value for the daily cost per gigabyte of storage in the **Storage Cost (per GB)** text box.
- 5 Click the **Save** icon (✔).

Assign a Storage Cost Profile to a Datastore

Fabric administrators can assign a storage cost profile to a datastore to apply those costs to all machines provisioned on that datastore.

Prerequisites

[Create a Storage Cost Profile.](#)

Procedure

- 1 Select **Infrastructure > Compute Resources > Compute Resources**.
- 2 Point to a compute resource and click **Edit**.
- 3 Click the **Configuration** tab.
- 4 Locate the datastore to which to assign the storage cost profile in the Storage table.
- 5 Click the **Edit** icon (✎).
- 6 Select a storage cost profile from the **Storage Cost Profile** drop-down menu.
The cost for the datastore is populated with the value from the storage cost profile.
- 7 Click the **Save** icon (✔).
- 8 Click **OK**.

If there are any reservations that include this datastore, its cost is also updated based on the cost in the storage cost profile.

Cost Profiles with vRealize Business Standard Edition

vRealize Business Standard Edition can be configured to provide cost information for vCenter Server.

The following is a high-level overview of the steps required to use cost information from vRealize Business Standard Edition in IaaS and the service catalog:

- 1 Install and configure vRealize Business Standard Edition.
 - a Install vRealize Business Standard Edition.
 - b As a business management administrator, configure vRealize Business Standard Edition to manage costs for one or more vCenter Server instances.

- c (Optional) Edit cost drivers in vRealize Business Standard Edition. This can include a markup for using the resources in addition to the raw resource cost, which contributes to the overall price.

Note vRealize Business Standard Edition supports configuring pricing policies based on vRealize Automation reservations and reservation policies. However, costs configured at this level in vRealize Business Standard Edition are not reflected in vRealize Automation.

For more information, see the vRealize Business Standard Edition documentation.

- 2 Create vSphere endpoints in IaaS for the same vCenter Server instances that you configured in vRealize Business Standard Edition.

Cost profiles are automatically generated for any hosts and clusters in the vCenter Server instances that are configured in both IaaS and vRealize Business Standard Edition. Values for monthly price in vRealize Business Standard Edition are converted to daily cost in IaaS.

Cost information is refreshed in IaaS every 24 hours. A fabric administrator can also manually initiate an update of cost information.

The automatically generated cost profiles cannot be viewed on the Cost Profiles page in IaaS. You can view the cost information on a compute resource that is managed by vRealize Business Standard Edition and update cost profiles for compute resources that are not managed by vRealize Business Standard Edition.

Note Although vRealize Business Standard Edition supports cost management for vCloud Director and Amazon AWS, costs from vRealize Business Standard Edition for these platforms are not integrated with IaaS or the service catalog.

View Cost Information from vRealize Business Standard Edition

Fabric administrators can view the costs for compute resources automatically populated from vRealize Business Standard Edition.

Prerequisites

- To view cost information from vRealize Business Standard Edition for a compute resource in IaaS, the vCenter Server that manages the host or cluster must be configured in both IaaS and vRealize Business Standard Edition.
- Log in to the vRealize Automation console as a **fabric administrator**.

Procedure

- 1 Select **Infrastructure > Compute Resources > Compute Resources**.
- 2 Point to a compute resource and click **Edit**.

Compute resources that have cost information from vRealize Business Standard Edition have a value of `Automatic` in the Cost Profile column.

- 3 Click the **Configuration** tab.

The daily costs for memory and CPU for the compute resource and the storage costs for each datastore are displayed. Storage cost for a specific datastore can also be viewed in reservations that include that datastore.

Update Cost Information from vRealize Business Standard Edition

Cost information from vRealize Business Standard Edition for compute resources is updated every 24 hours. A fabric administrator can also trigger a manual update of cost information.

Prerequisites

- Verify that a tenant administrator or business management administrator has updated cost information in vRealize Business Standard Edition since the last time that cost information was updated. For more information, see the vRealize Business Standard Edition documentation.
- Log in to the vRealize Automation console as a **fabric administrator**.

Procedure

1 Select **Infrastructure > Compute Resources > Compute Resources**.

2 Click **Update Costs**.

The Confirm Update Costs page appears with the cost status, including the last time costs were updated and the next scheduled update.

3 Click **OK**.

4 (Optional) Confirm that the cost update has completed successfully.

a Click **Update Costs**.

If the updated costs are still being applied to compute resources, the status appears as In Progress. After the update is complete, the page automatically refreshes with the new status and last updated time.

b Click **Cancel**.

The cost information is updated for all compute resources that are configured with vRealize Business Standard Edition.

Enabling Users to Select Datacenter Locations

The **Display location on request** check box on the Blueprint Information tab allows users to select a particular datacenter location at which to provision a requested virtual or cloud machine.

For example, if you have an office in London and an office in Boston, you might have compute resources and business groups in both locations. By enabling the **Display location on request** check box, your business group users can choose to provision their machines with the resources that are local, for example.

The following is a high-level overview of the sequence of steps required to enable users to select datacenter locations:

- 1 A system administrator adds datacenter location information to a locations file.
- 2 A fabric administrator edits a compute resource to associate it with a location.
- 3 A tenant administrator or business group manager creates a blueprint that prompts users to choose a datacenter location when submitting a machine request.

Associate a Compute Resource with a Location

Fabric administrators associate compute resources to datacenter locations before tenant administrators and business group managers can prompt machine requesters to select a datacenter location at provisioning time.

Prerequisites

- Log in to the vRealize Automation console as a **fabric administrator**.
- Your system administrator must add datacenter locations to a locations file. See *System Administration*.

Procedure

- 1 Select **Infrastructure > Compute Resources > Compute Resources**.
- 2 Point to the compute resource to which you want to add a location and click **Edit**.
- 3 Select a datacenter location from the **Locations** drop-down menu.
- 4 Click **OK**.
- 5 Repeat this procedure as necessary to associate your compute resources to locations.

Tenant administrators and business group managers can set the Display location on request option on blueprints to prompt users to select a data location when requesting a machine.

Removing Datacenter Locations

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

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

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

- 1 A system administrator removes the datacenter location information from the locations file.

- 2 A fabric administrator removes all the compute resource associations to the location by editing the locations of each associated compute resource.

Remove a Datacenter Location

Fabric administrators must remove all associations to a datacenter location before it stops displaying to machine requesters as a datacenter location option.

Prerequisites

- Log in to the vRealize Automation console as a **fabric administrator**.
- Your system administrator removes the datacenter from the locations file. See *System Administration*.

Procedure

- 1 Select **Infrastructure > Compute Resources > Compute Resources**.
- 2 Point to a compute resource associated with the location you want to remove and click **Edit**.
- 3 Remove the datacenter location by selecting the blank option in the **Locations** drop-down menu.
- 4 Click **OK**.
- 5 Repeat this procedure to remove the location from all of your compute resources.

After the datacenter location is removed from all compute resources, it no longer appears to machine requesters in the locations drop-down menu.

Active Directory Cleanup Plugin

The Active Directory Cleanup Plugin is a plugin that tenant administrators and business group managers configure with a set of custom properties to specify actions to take in Active Directory when a machine is deleted from a hypervisor.

Using the Active Directory Cleanup Plugin, you can specify the following Active Directory account actions to occur when a machine is deleted from a hypervisor:

- The AD account is deleted
- The AD account is disabled
- The AD account is renamed
- The AD account is moved to another AD organizational unit (OU)

Gather the following information for tenant administrators and business group managers to include in their blueprint:

- 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 A fabric administrator can create a build profile by using the property set `ActiveDirectoryCleanupPlugin` 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.

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*.

Enabling Connections to Machines Using SSH

With a combination of custom properties and blueprint settings, you can enable the **Connect Using SSH** option for Linux machines provisioned through vRealize Automation.

To enable machine users to connect using SSH:

- 1 A blueprint creator adds the custom property `Machine.SSH` to add the ability to connect using SSH to a blueprint. Optionally, a fabric administrator can include this custom property in a property group.
- 2 A blueprint creator enables the **Connect using RDP or SSH** machine operation for machines provisioned from the blueprint.
- 3 A tenant administrator or business group manager entitles users or groups to use the **Connect using RDP** or **Connect Using SSH** options.

If the custom property `Machine.SSH` is set to **true**, and the **Connect using RDP or SSH** machine operation is enabled in the blueprint, all Linux machines provisioned from the blueprint display the **Connect Using SSH** option to entitled users.

Preparing for Provisioning

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

- [Space-Efficient Storage for Virtual Provisioning](#)

Space-efficient storage technology eliminates the inefficiencies of traditional storage methods by using only the storage actually required for a machine's operations. Typically, this is only a fraction of the storage actually allocated to machines. vRealize Automation supports two methods of provisioning with space-efficient technology, thin provisioning and FlexClone provisioning.

- [Preparing Your SCVMM Environment](#)

Before you begin creating SCVMM templates and hardware profiles for use in vRealize Automation machine provisioning, you must understand the naming restrictions on template and hardware profile names, and configure SCVMM network and storage settings.

- [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.

- [Choosing a Provisioning Scenario](#)

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

- [Checklist for Virtual Provisioning by Cloning Preparation](#)

With vRealize Automation, you can clone both Windows and Linux machines. Some preparation outside of vRealize Automation is required to create the template and the customization objects used to clone machines.

- [Choosing Your Linux Kickstart Provisioning Preparation Method](#)

You can automate a Linux installation on newly provisioned machines. The procedures for preparing for this type of provisioning differ depending on whether you are using Solaris AI or Kickstart or autoYaST configuration files to perform your machine provisioning.

- [Preparing for SCCM Provisioning](#)

vRealize Automation boots a newly provisioned machine from an ISO image, and then passes control to the specified SCCM task sequence.

- [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.

Space-Efficient Storage for Virtual Provisioning

Space-efficient storage technology eliminates the inefficiencies of traditional storage methods by using only the storage actually required for a machine's operations. Typically, this is only a fraction of the storage actually allocated to machines. vRealize Automation supports two methods of provisioning with space-efficient technology, thin provisioning and FlexClone provisioning.

When standard storage is used, the storage allocated to a provisioned machine is fully committed to that machine, even when it is powered off. This can be a significant waste of storage resources because few virtual machines actually use all of the storage allocated to them, just as few physical machines operate with a 100% full disk. When a space-efficient storage technology is used, the storage allocated and the storage used are tracked separately and only the storage used is fully committed to the provisioned machine.

Thin Provisioning

Thin provisioning is supported for all virtual provisioning methods. Depending on your virtualization platform, storage type, and default storage configuration, thin provisioning might always be used during machine provisioning. For example, for vSphere ESX Server integrations using NFS storage, thin provisioning is always employed. However, for vSphere ESX Server integrations that use local or iSCSI storage, thin provisioning is only used to provision machines if the custom property `VirtualMachine.Admin.ThinProvision` is specified in the blueprint. For more information about thin provisioning, please see the documentation provided by your virtualization platform.

Net App FlexClone Provisioning

You can create a blueprint for Net App FlexClone provisioning if you are working in a vSphere environment that uses Network File System (NFS) storage and FlexClone technology.

You can only use NFS storage, or machine provisioning fails. You can specify a FlexClone storage path for other types of machine provisioning, but the FlexClone storage path behaves like standard storage.

The following is a high-level overview of the sequence of steps required to provision machines that use FlexClone technology:

- 1 An IaaS administrator creates a NetApp ONTAP endpoint. See [Create a NetApp ONTAP Endpoint](#).
- 2 An IaaS administrator runs data collection on the endpoint to enable the endpoint to be visible on the compute resource and reservation pages.

The FlexClone option is visible on a reservation page in the endpoint column if a NetApp ONTAP endpoint exists and if the host is virtual. If there is a NetApp ONTAP endpoint, the reservation page displays the endpoint assigned to the storage path.

- 3 A fabric administrator creates a vSphere reservation, enables FlexClone storage, and specifies an NFS storage path that uses FlexClone technology. See [Create a Reservation](#).
- 4 A tenant administrator or business group manager creates a FlexClone blueprint. See [Create a Blueprint for Net App FlexClone Provisioning](#).

Preparing Your SCVMM Environment

Before you begin creating SCVMM templates and hardware profiles for use in vRealize Automation machine provisioning, you must understand the naming restrictions on template and hardware profile names, and configure SCVMM network and storage settings.

For related information about preparing your environment, see *SCVMM Requirements* in *Installation and Configuration*.

For related information about machine provisioning, see [Create an SCVMM Endpoint](#).

Template and Hardware Profile Naming

Because of naming conventions that SCVMM and vRealize Automation use for templates and hardware profiles, do not start your template or hardware profile names with the words temporary or profile. For example, the following words are ignored during data collection:

- TemporaryTemplate
- Temporary Template
- TemporaryProfile
- Temporary Profile
- Profile

Required Network Configuration for SCVMM Clusters

SCVMM clusters only expose virtual networks to vRealize Automation, so you must have a 1:1 relationship between your virtual and logical networks. Using the SCVMM console, map each logical network to a virtual network and configure your SCVMM cluster to access machines through the virtual network.

Required Storage Configuration for SCVMM Clusters

On SCVMM Hyper-V clusters, vRealize Automation collects data and provisions on shared volumes only. Using the SCVMM console, configure your clusters to use shared resource volumes for storage.

Required Storage Configuration for Standalone SCVMM Hosts

For standalone SCVMM hosts, vRealize Automation collects data and provisions on the default virtual machine path. Using the SCVMM console, configure default virtual machine paths for your standalone hosts.

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.

Choosing a Provisioning Scenario

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

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

Table 3-1. Choosing a Provisioning Method

Scenario	Supported Platforms	Provisioning Method	Pre-provisioning Preparations
Provision machines with no guest operating system. You can install an operating system after provisioning.	All	Basic	No required pre-provisioning preparations outside of vRealize Automation.
Provision machines by cloning from a template object created from an existing Windows or Linux machine, called the reference machine, and a customization specification.	<ul style="list-style-type: none"> ■ vSphere ■ KVM (RHEV) ■ SCVMM 	Clone	Checklist for Virtual Provisioning by Cloning Preparation
Provision a space-efficient copy of a virtual machine called a linked clone. Linked clones are based on a snapshot of a VM and use a chain of delta disks to track differences from a parent machine.	vSphere	Linked clone	No required pre-provisioning preparations outside of vRealize Automation.
Provision a space-efficient copy of a virtual machine by using Net App FlexClone technology.	vSphere	FlexClone	Checklist for Virtual Provisioning by Cloning Preparation
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.	All	Linux Kickstart	Choosing Your Linux Kickstart Provisioning Preparation Method
Provision a machine and pass control to an SCCM task sequence to boot from an ISO image, deploy a Windows operating system, and install the vRealize Automation guest agent.	All	Virtual SCCM	Preparing for SCCM Provisioning
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.	All	WIM Image	Preparing for WIM Provisioning

Checklist for Virtual Provisioning by Cloning Preparation

With vRealize Automation, you can clone both Windows and Linux machines. Some preparation outside of vRealize Automation is required to create the template and the customization objects used to clone machines.

Cloning requires a template to clone from, created from a reference machine. You are not required to install the guest agent or the WMI agent, and you are not required to specify Visual Basic scripts. These options provide you with the ability to further configure machines provisioned from a clone template.

If you are provisioning a Windows machine by cloning, the only way to join the provisioned machine to an Active Directory domain is by using the customization specification from vCenter Server or including a guest operating system profile with your SCVMM template. Machines provisioned by cloning cannot be placed in an Active Directory container during provisioning. This must be done manually after provisioning.

Table 3-2. Preparing for Virtual Provisioning by Cloning Checklist

Task	Location	Details
<input type="checkbox"/> Identify or create the reference machine.	Hypervisor	See the documentation provided by your hypervisor.
<input type="checkbox"/> (Optional) Install the vRealize Automation guest agent on your reference machine to customize deployed machines.	Reference machine	See Using vRealize Automation Guest Agent in Provisioning .
<input type="checkbox"/> (Optional) Install VMware Tools on the reference machine.	vCenter Server	See the VMware Tools documentation.
<input type="checkbox"/> Use the reference machine to create a template for cloning.	Hypervisor	The reference machine may be powered on or off. If you are cloning in vCenter Server, you can use a reference machine directly without creating a template. See the documentation provided by your hypervisor.
<input type="checkbox"/> Create the customization object to configure cloned machines by applying System Preparation Utility information or a Linux customization.	Hypervisor	If you are cloning for Linux you can install the Linux guest agent and provide external customization scripts instead of creating a customization object. If you are cloning with vCenter Server, you must provide the customization specification as the customization object. See the documentation provided by your hypervisor.
<input type="checkbox"/> (Optional) If you installed and configured an EPI agent, specify Visual Basic scripts to run outside vRealize Automation as additional steps in the machine life cycle.	Machine where the EPI agent is installed	See Enabling Visual Basic Scripts in Provisioning .

Table 3-2. Preparing for Virtual Provisioning by Cloning Checklist (Continued)

Task	Location	Details
<input type="checkbox"/> (Optional) Install and enable a Windows Management Instrumentation (WMI) agent to collect certain data from a Windows machine managed by vRealize Automation, for example the Active Directory status of a machine's owner.	Manager Service host	See <i>Installation and Configuration</i> .
<input type="checkbox"/> (Optional) Run endpoint data collection to collect your templates immediately.	vRealize Automation	See Start Endpoint Data Collection Manually .

Gather information about your template and customizations and transfer it to the tenant administrators and business group managers who create blueprints.

- Gather the following information so that tenant administrators and business group managers can include it in their blueprints:
 - a The name of the template.
 - b The name of the customization specification.
 - c The amount of total storage specified for the template.
 - d The reservations on which the template is available. Business group managers and tenant administrators are provided with a list of all templates on all reservations to choose from when creating a blueprint. vRealize Automation does not validate that the selected template exists on the reservation selected in the blueprint. To avoid errors during provisioning, ensure that the template is available on all reservations or create reservation policies that tenant administrators and business group managers can use to restrict the blueprint to reservations where the template is available.
 - 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.
- (Optional) If you have customization scripts for the Linux guest agent to run during provisioning, gather the following information for your tenant administrators and business group managers to include in their blueprints:
 - a The name of the customization script.
 - b Whether the location of the customization script is nfs or local.
 - c The name of the NFS server on which the customization script is located if it is not local.
 - d The local path to the customization path or the export path if the script is on an NFS server.

Note You can create build profiles by using the property set `LinuxExternalScriptProperties` to provide the information for Linux customization scripts. Doing so makes it easier for tenant administrators and business group managers to include this information correctly in their blueprints.

Using vRealize Automation Guest Agent in Provisioning

You can install the guest agent on your templates for cloning to further customize a machine after deployment. You can use the reserved guest agent custom properties to perform basic customizations such as adding and formatting disks, or you can create your own custom scripts for the guest agent to run within the guest operating system of a provisioned machine.

After the deployment is completed and the customization specification is run (if you provided one), the guest agent creates an XML file that contains all of the deployed machine's custom properties `c:\VRMGuestAgent\site\workitem.xml`, completes any tasks assigned to it with the guest agent custom properties, and then deletes itself from the provisioned machine.

You can write your own custom scripts for the guest agent to run on deployed machines, and use custom properties on the machine blueprint to specify the location of those scripts and the order in which to run them. You can also use custom properties on the machine blueprint to pass custom property values to your scripts as parameters.

For example, you could use the guest agent to make the following customizations on deployed machines:

- Change the IP address
- Add or format drives
- Run security scripts
- Initialize another agent, for example Puppet or Chef

Your custom scripts do not have to be locally installed on the machine. As long as the provisioned machine has network access to the script location, the guest agent can access and run the scripts. This lowers maintenance costs because you can update your scripts without having to rebuild all of your templates .

If you choose to install the guest agent to run custom scripts on provisioned machines, your blueprints must include the appropriate guest agent custom properties. For example, if you install the guest agent on a template for cloning, create a custom script that changes the provisioned machine's IP address, and place the script in a shared location, you need to include a number of custom properties in your blueprint.

Table 3-3. Custom Properties for Changing IP Address of a Provisioned Machine with a Guest Agent

Custom Property	Description
<code>VirtualMachine.Admin.UseGuestAgent</code>	Set to true to initialize the guest agent when the provisioned machine is started.
<code>VirtualMachine.Customize.WaitComplete</code>	Set to True to prevent the provisioning workflow from sending work items to the guest agent until all customizations have been completed.
<code>VirtualMachine.SoftwareN.ScriptPath</code>	Specifies the full path to an application's install script. The path must be a valid absolute path as seen by the guest operating system and must include the name of the script file name. You can pass custom property values as parameters to the script by inserting <code>{YourCustomProperty}</code> in the path string. For example, entering the value <code>\\vra-scripts.mycompany.com\scripts\changeIP.bat</code> runs the <code>changeIP.bat</code> script from a shared location, but entering the value <code>\\vra-scripts.mycompany.com\scripts\changeIP.bat {VirtualMachine.Network0.Address}</code> runs the <code>changeIP</code> script but also passes the value of the <code>VirtualMachine.Network0.Address</code> property to the script as a parameter.

For more information about custom properties you can use with the guest agent, see *Custom Properties Reference*.

Install the Guest Agent on a Linux Machine

Install the Linux guest agent on your reference machines to further customize machines after deployment.

Prerequisites

- Identify or create the reference machine.
- The guest agent files you download contain both `tar .gz` and RPM package formats. If your operating system cannot install `tar .gz` or RPM files, use a conversion tool to convert the installation files to your preferred package format.

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 Install the guest agent package that corresponds to the guest operating system you are deploying during provisioning.

- a Navigate to the `LinuxGuestAgentPkgs` subdirectory for your guest operating system.
- b Locate your preferred package format or convert a package to your preferred package format.
- c Install the guest agent package on your reference machine.

For example, to install the files from the RPM package, run `rpm -i gugent-7.0.0-012715.x86_64.rpm`.

- 5 Configure the guest agent to communicate with the Manager Service by running `installgugent.sh Manager_Service_Hostname_fdqn:portnumber ssl`.

Option	Description
If you are using a load balancer	<p>Enter the fully qualified domain name and port number of your Manager Service load balancer. For example:</p> <pre>cd /usr/share/gugent ./installgugent.sh load_balancer_manager_service.mycompany.com:443 ssl</pre>
With no load balancer	<p>Enter the fully qualified domain name and port number of your Manager Service machine. For example:</p> <pre>cd /usr/share/gugent ./installgugent.sh manager_service_machine.mycompany.com:443 ssl</pre>

The default port number for the Manager Service is 443.

- 6 If deployed machines are not already configured to trust the Manager Service SSL certificate, you must install the `cert.pem` file on your reference machine to establish trust.
- For the most secure approach, obtain the `cert.pem` certificate and manually install the file on the reference machine.
 - For a more convenient approach, you can connect to the manager service load balancer or manager service machine and download the `cert.pem` certificate.

Option	Description
If you are using a load balancer	<p>As the root user on the reference machine, run the following command:</p> <pre>echo openssl s_client -connect manager_service_load_balancer.mycompany.com:443 sed -ne '/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p' > cert.pem</pre>
With no load balancer	<p>As the root user on the reference machine, run the following command:</p> <pre>echo openssl s_client -connect manager_service_machine.mycompany.com:443 sed -ne '/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p' > cert.pem</pre>

- 7 If you are installing the guest agent on a Ubuntu operating system, create symbolic links for shared objects by running one of the following command sets.

Option	Description
64-bit systems	<pre>cd /lib/x86_64-linux-gnu sudo ln -s libssl.so.1.0.0 libssl.so.10 sudo ln -s libcrypto.so.1.0.0 libcrypto.so.10</pre>
32-bit systems	<pre>cd /lib/i386-linux-gnu sudo ln -s libssl.so.1.0.0 libssl.so.10 sudo ln -s libcrypto.so.1.0.0 libcrypto.so.10</pre>

Install the Guest Agent on a Windows Machine

Install the Windows guest agent on a Windows reference machines to run as a Windows service and enable further customization of machines provisioned by cloning.

Prerequisites

Identify or create the reference machine.

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 Windows guest agent files.
 - Windows guest agent files (**32-bit**.)
 - Windows guest agent files (**64-bit**.)
- 3 Install the guest agent on the reference machine.
 - a Download `GugentZip_version` to the C drive on the reference machine.
 - b Right-click the file and select **Properties**.
 - c Click **General**.
 - d Click **Unblock**.
 - e Extract the files.

This produces the directory `C:\VRMGuestAgent`. Do not rename this directory.

- 4 Configure the guest agent to communicate with the Manager Service.
 - a Open an elevated command prompt.
 - b Navigate to C:\VRMGuestAgent.
 - c Run `winservice -i -h Manager_Service_Hostname_fdn:portnumber -p ssl`.

Option	Description
If you are using a load balancer	Enter the fully qualified domain name and port number of your Manager Service load balancer. For example, <code>winservice -i -h <i>load_balancer_manager_service.mycompany.com:443</i> -p ssl</code> .
With no load balancer	Enter the fully qualified domain name and port number of your Manager Service machine. For example, <code>winservice -i -h <i>manager_service_machine.mycompany.com:443</i> -p ssl</code> .

The default port number for the Manager Service is 443.

The name of the Windows service is VCACGuestAgentService. You can find the installation log VCAC-GuestAgentService.log in C:\VRMGuestAgent.

What to do next

Configure the guest agent to trust a server. See [Configuring the Windows Guest Agent to Trust a Server](#).

Configuring the Windows Guest Agent to Trust a Server

The most secure approach is to install the trusted PEM file manually on each template that uses the guest agent, but you can also allow the guest agent to trust the first machine to which it connects.

Installing the PEM file for the trusted server on each template along with the guest agent is the most secure approach. For security, the guest agent does not check for a certificate if a PEM file already exists in the VRMGuestAgent directory. If the server certificates change, you must manually rebuild your templates with the new PEM files.

You can also configure the guest agent to populate the trusted PEM file on first use. This is less secure than manually installing the PEM files on each template, but is more flexible for environments where you might use a single template for multiple servers. To allow the guest agent to trust the first server it connects to, you create a template with no PEM files in the VRMGuestAgent directory. The guest agent populates the PEM file the first time it connect to a server. The template always trusts the first system to which it connects. For security, the guest agent does not check for a certificate if a PEM file already exists in the VRMGuestAgent directory. If the server certificate changes, you must remove the PEM file from your VRMGuestAgent directory. The guest agent installs the new PEM file the next time it connects to the server.

Choosing Your Linux Kickstart Provisioning Preparation Method

You can automate a Linux installation on newly provisioned machines. The procedures for preparing for this type of provisioning differ depending on whether you are using Solaris AI or Kickstart or autoYaST configuration files to perform your machine provisioning.

Figure 3-1. Workflow for Preparing for Linux Machine Provisioning

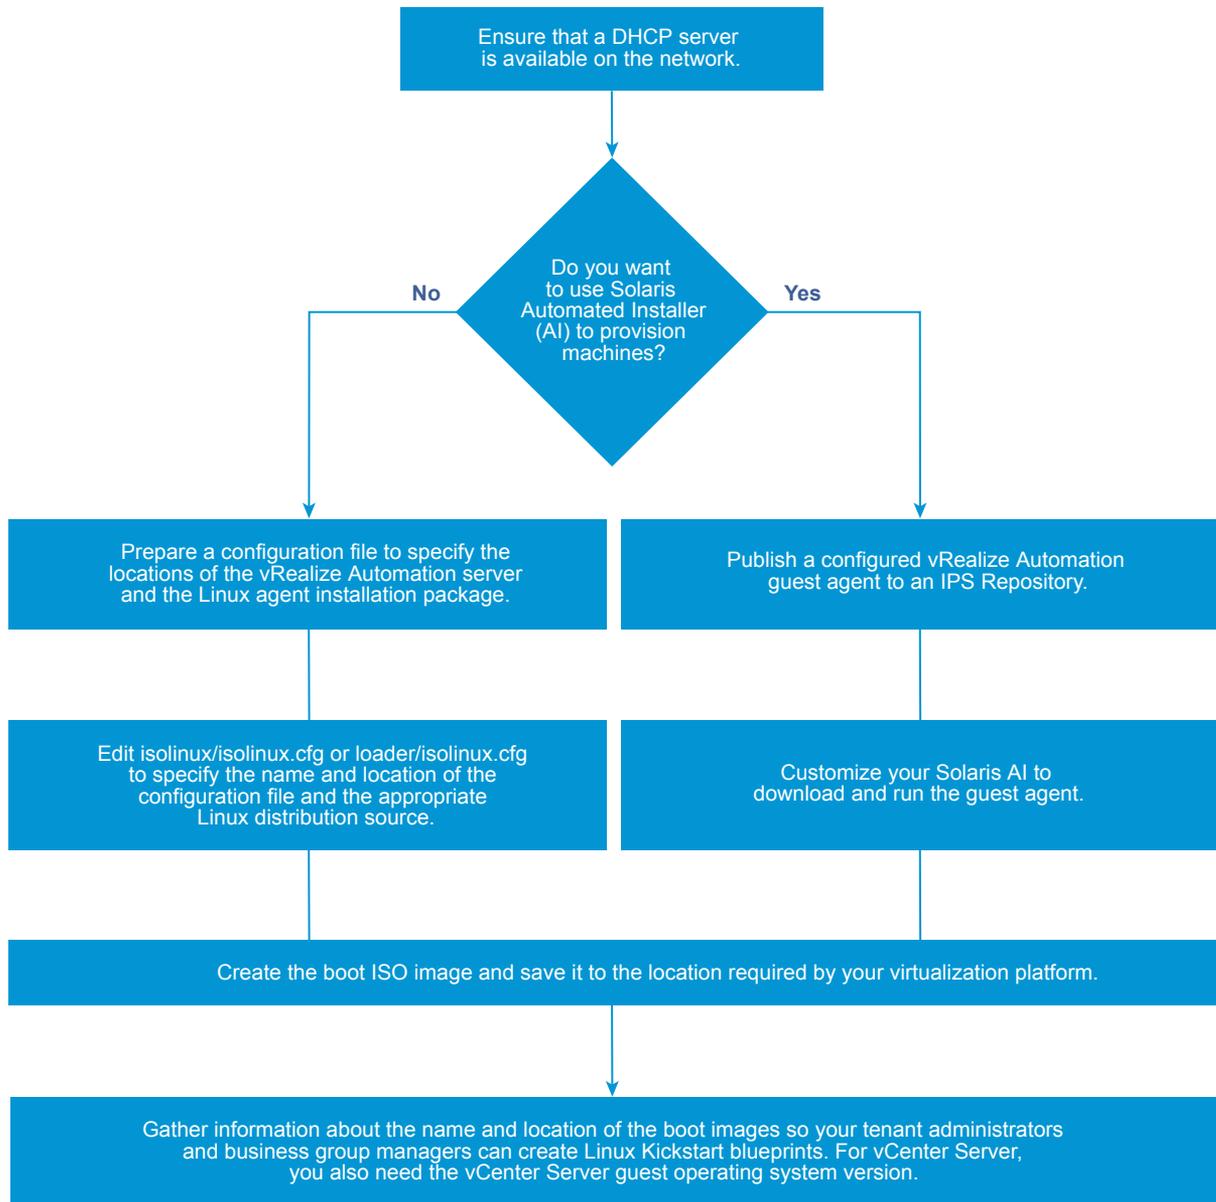


Table 3-4. Choosing Your Linux Kickstart Provisioning Preparation Method

Preparation Method	More Information
If you are using Oracle Solaris Automated Installer (AI) service and Image Packaging System (IPS) to install and configure operating systems on provisioned machines, you must configure and publish the vRealize Automation guest agent as an IPS software package and then include that package in your AI manifest.	See Checklist for Preparing for Linux Kickstart Provisioning by using Solaris Automated Installer .
If you are not using Solaris Automated Installer to provision machines, you can create configuration files that automate the installation of the vRealize Automation guest agent and a Linux operating system on newly provisioned machines.	See Checklist for Preparing for Linux Kickstart Provisioning by using Configuration Files .

Checklist for Preparing for Linux Kickstart Provisioning by using Solaris Automated Installer

If you are using Oracle Solaris Automated Installer (AI) service and Image Packaging System (IPS) to install and configure operating systems on provisioned machines, you must configure and publish the vRealize Automation guest agent as an IPS software package and then include that package in your AI manifest.

Table 3-5. Checklist for Preparing for Linux Kickstart Provisioning by using Solaris AI

Linux Kickstart by using Solaris AI
<input type="checkbox"/> Ensure that a DHCP server is available on the network. vRealize Automation cannot provision machines by using Linux Kickstart provisioning unless DHCP is available.
<input type="checkbox"/> Configure and publish the vRealize Automation guest agent as an IPS software package. See Configure and Publish the Guest Agent to your Oracle Solaris IPS Repository .
<input type="checkbox"/> Include your configured guest agent package in your automated installer manifest. You can use the sample automated installer manifest included with your Linux guest agent files and edit it for your own specific environment, or you can include your guest agent package in your existing manifests. See Customize your Solaris Automated Installation for Linux Kickstart Provisioning .
<input type="checkbox"/> Gather the following information so that tenant administrators and business group managers can include it in their blueprints: <ol style="list-style-type: none"> 1 The name and location of the ISO image. 2 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 <code>BootIsoProperties</code> 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.

Configure and Publish the Guest Agent to your Oracle Solaris IPS Repository

If you are using Solaris AI to provision Linux machines, you must configure and publish the vRealize Automation guest agent as an IPS software package to include in your AI manifest.

Prerequisites

Create or identify an Oracle Solaris IPS Repository. For an example of creating an IPS repository, see [Example: Create a Local IPS Repository and Publish the Guest Agent](#).

Procedure

- 1 Log in to the Solaris machine hosting your IPS Repository.
- 2 Download the guest agent files from your vRealize Appliance.

```
wget https://vRealize_VA_Hostname_fqdn:5480/installer/LinuxGuestAgentPkgs.zip
```

If you are using self-signed certificates, you might have to use the wget option `--no-check-certificate`. For example:

```
wget --no-check-certificate https://vRealize_VA_Hostname_fqdn:5480/installer/LinuxGuestAgentPkgs.zip
```

- 3 Unpack the guest agent files.
- 4 Navigate to `LinuxGuestAgentPkgs/solaris11-amd64`.
- 5 Run `create_pkg_manifest.sh` to establish communication with your Manager Service machine and create a package manifest file `gugent-pkg.p5m`.

```
/bin/bash ./create_pkg_manifest.sh Manager_Service_Hostname_fqdn:port
```

The default port for the Manager Service is 443.

- 6 Publish your configured guest agent package to your IPS repository.

For example,

```
pkgsend -s /export/my-repository -d gugent-proto gugent-pkg.p5m
```

You can verify that your configured guest agent package is published in your repository by running the `list` command for your repository. For example

```
pkgrepo list -s http://localhost:9000
```

Example: Create a Local IPS Repository and Publish the Guest Agent

You configure a Solaris 11.3 virtual machine to serve as your IPS repository and enable remote access to the repository over HTTP connections. You run a script to configure your guest agent to communicate with your manager service, then publish the configured guest agent as an IPS software package on the local repository.

- 1 Create or identify a Solaris 11.3 virtual machine with SSH root login enabled.
- 2 Log in to the virtual machine as root by using SSH.

- 3 Create an IPS Repository called gugent-repo and set the default publisher as VMware.

```
root@solaris:~# cd /export
root@solaris:/export# pkgrepo create gugent-repo
root@solaris:/export# pkgrepo -s gugent-repo set publisher/prefix=vmware
```

- 4 Configure the application/pkg/server service to enable clients to access the repository over HTTP connections.

```
root@solaris:~# svccfg -s application/pkg/server setprop pkg/inst_root=/export/gugent-repo
root@solaris:~# svccfg -s application/pkg/server setprop pkg/readonly=false
root@solaris:~# svccfg -s application/pkg/server setprop pkg/port=9000
root@solaris:~# svcadm refresh application/pkg/server
root@solaris:~# svcadm enable application/pkg/server
```

- 5 Download a copy of the guest agent from your vRealize Appliance without checking certificates.

```
wget --no-check-certificate https://vra01svr01.rainpole.local:5480/installer/LinuxGuestAgentPkgs.zip
```

- 6 Run `create_pkg_manifest.sh` to establish communication with your Manager Service machine and create a package manifest file called `gugent-pkg.p5m`.

```
/bin/bash ./create_pkg_manifest.sh vra01svr01.rainpole.local:443
```

- 7 Publish your guest agent package.

```
pkgsend -s /export/gugent-repo -d gugent-proto gugent-pkg.p5m
```

- 8 Verify that your configured guest agent package is published in your repository.

```
pkgrepo list -s http://localhost:9000
```

What to do next

You must include your configured guest agent software package in your Automated Installer manifest. See [Customize your Solaris Automated Installation for Linux Kickstart Provisioning](#).

Customize your Solaris Automated Installation for Linux Kickstart Provisioning

You must include your configured guest agent package in your automated installer manifest. You can use the sample automated installer manifest included with your Linux guest agent files and edit it for your own specific environment, or you can include your guest agent package in your existing manifests.

Prerequisites

- [Configure and Publish the Guest Agent to your Oracle Solaris IPS Repository.](#)

- Create or identify the Solaris Automated Install Boot Image and AI manifest that you want to use for provisioning machines. For an example of how to create your own, see [Example: Customize your Solaris Automated Installation](#).

Procedure

- 1 Edit your AI manifest to point to the IPS Repository where you published the configured guest agent.

For example,

```
<publisher name="vmware">
  <origin name="http://10.118.192.254:9000/" />
</publisher>
```

- 2 Edit the software packages and package groups.

You must include your configured guest agent package in the software to install list.

```
<software_data action="install">
  <name>pkg:/entire@0.5.11-0.175.3</name>
  <name>pkg:/group/system/solaris-large-server</name>
  <name>pkg:/mygugent-pkg</name>
</software_data>
</software>
```

- 3 Save and exit.

Your AI manifest includes the vRealize Automation guest agent software package, which is required for Linux Kickstart provisioning.

Example: Customize your Solaris Automated Installation

You download the existing Solaris 11.3x86 AI Boot image and edit the boot grub menu and system configuration profile for your own specifications. You replace the AI manifest with the sample manifest included in your Linux guest agent packages. When you're satisfied with all your changes, you generate a new AI boot image.

- 1 Download the Solaris 11.3 x86 Automated Installer Boot Image onto your Solaris 11.3 virtual machine.
- 2 Mount the image and copy the files to a new location so you can modify them.

```
root@solaris:~# mkdir /root/ai /root/ai/image /mnt/ai
root@solaris:~# mount -F hsfs $(lofiadm -a /root/ai/sol-11_3-ai-x86.iso) /mnt/ai
root@solaris:~# cd /mnt/ai
root@solaris:/mnt/ai# find . -depth -print | cpio -pdm /root/ai/image
```

- 3 Edit the `set default` configuration in the `/root/ai/image/boot/grub/grub.cfg` file to point to the "Oracle Solaris 11.3 Automated Install" grub boot menu entry. The first entry in the list is "0", and the second entry is "1."

```
# GRUB2 configuration file
load_video_$target
terminal_output console
set timeout=30
set default="1"
menuentry "Oracle Solaris 11.3 Automated Install custom" {
    search --no-floppy --file --set=root /.drought.us.oracle.com-2015-10-06T08:49:49.442166
    set kern=/platform/i86pc/kernel/amd64/unix
    echo -n "Loading ${root}$kern: "
    $multiboot $kern $kern -B install=true,aimanifest=prompt
    set gfxpayload="1024x768x32;1024x768x16;800x600x16;640x480x16;640x480x15;640x480x32"
    insmod gzio
    echo -n "Loading ${root}/platform/i86pc/amd64/boot_archive: "
    $module /platform/i86pc/amd64/boot_archive
}
menuentry "Oracle Solaris 11.3 Automated Install" {
    search --no-floppy --file --set=root /.drought.us.oracle.com-2015-10-06T08:49:49.442166
    set kern=/platform/i86pc/kernel/amd64/unix
    echo -n "Loading ${root}$kern: "
    $multiboot $kern $kern -B install=true
    set gfxpayload="1024x768x32;1024x768x16;800x600x16;640x480x16;640x480x15;640x480x32"
    insmod gzio
    echo -n "Loading ${root}/platform/i86pc/amd64/boot_archive: "
    $module /platform/i86pc/amd64/boot_archive
}
...
```

- 4 Generate a new grub boot image.

```
root@solaris:/mnt/ai# cat /root/ai/image/boot/grub/i386-
pc/cdboot.img /root/ai/image/boot/grub/i386-pc/core.img > /root/ai/myboot.img
root@solaris:/mnt/ai# cp /root/ai/myboot.img /root/ai/image/boot/grub/i386-pc/myboot.img
```

- 5 Mount the `solaris.zlib` image and copy the files to a new location so you can modify them.

```
root@solaris:/mnt/ai# mkdir /root/ai/zlib /mnt/zlib
root@solaris:/mnt/ai# mount -F hsfs $(lofiadm -a /root/ai/image/solaris.zlib) /mnt/zlib
root@solaris:/mnt/ai# cd /mnt/zlib
root@solaris:/mnt/zlib# find . -depth -print | cpio -pdm /root/ai/zlib
```

- 6 Run the `sysconfig create-profile` command to create a new system profile.

```
root@solaris:/mnt/zlib# sysconfig create-profile -o /root/ai
```

- 7 Generate a new copy of `solaris.zlib` with your changes included and replace the old `solaris.zlib`.

```
root@solaris:/mnt/zlib#
cp /root/ai/sc_profile.xml /root/ai/zlib/share/auto_install/sc_profiles/enable_sci.xml
root@solaris:/mnt/zlib# cd /root/ai/zlib
root@solaris:~/ai/zlib# mkisofs -o /root/ai/mysolaris.zlib -quiet -N -l -R -U -allow-multidot -no-
iso-translate -cache-inodes -d -D -V "compress" /root/ai/zlib/
root@solaris:~/ai/zlib# lofiadm -C lzma /root/ai/mysolaris.zlib
root@solaris:~/ai/zlib# cp /root/ai/mysolaris.zlib /root/ai/image/solaris.zlib
```

- 8 Create a copy of the sample manifest file included in your guest agent package and open it for editing.

```
root@solaris:/mnt/ai# cp /root/LinuxGuestAgentPkgs/solaris11-
amd64/sample_ai_manifest.xml /root/ai/default.xml
root@solaris:/mnt/ai# vi /root/ai/default.xml
```

- 9 Edit the publisher origin name in the manifest file to point to your Solaris 11.3 virtual machine, which you configured to run as your local IPS repository.

```
<publisher name="vmware">
  <origin name="http://10.118.192.254:9000/" />
</publisher>
```

- 10 Replace the AI boot image manifest file with your own.

```
root@solaris:/mnt/ai# cp /root/ai/default.xml /root/ai/image/auto_install/default.xml
```

- 11 Generate a new AI boot image called `myai.iso` with all of your changes.

```
root@solaris:~/ai/zlib# cd /root/ai/image
root@solaris:~/ai/image# /usr/bin/mkisofs -d -D -J -l -r -U -relaxed-filenames -b boot/grub/i386-
pc/myboot.img -no-emul-boot -boot-load-size 4 -boot-info-table -c .catalog -V "myai.iso" -
o /root/ai/myai.iso "/root/ai/image/"
```

- 12 Save the resulting `myai.iso` file to the location required by your hypervisor.

What to do next

Gather the following information so that your IaaS architects can include it in their blueprints:

- 1 The name and location of the ISO image.
- 2 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 property group with the property set `BootIsoProperties` to include the required ISO information. This makes it easier to include this information correctly on blueprints.

Checklist for Preparing for Linux Kickstart Provisioning by using Configuration Files

If you are not using Solaris Automated Installer to provision machines, you can create Kickstart or autoYaST configuration files that automate the installation of the vRealize Automation guest agent and a Linux operating system on newly provisioned machines.

Table 3-6. Checklist for Preparing for Linux Kickstart Provisioning by using Configuration Files

Linux Kickstart by using a configuration file

- Ensure that a DHCP server is available on the network. vRealize Automation cannot provision machines by using Linux Kickstart provisioning unless DHCP is available.

- 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](#).

- 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.

- (Optional) Add customization scripts.
 - 1 To specify post-installation customization scripts in the configuration file, see [Specify Custom Scripts in a kickstart/autoYaST Configuration File](#).
 - 2 To call Visual Basic scripts in blueprint, see [Enabling Visual Basic Scripts in Provisioning](#).

- 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.

- Gather the following information so that tenant administrators and business group managers can include it in their blueprints:
 - 1 The name and location of the ISO image.
 - 2 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](#).
- 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 SCCM Provisioning

vRealize Automation boots a newly provisioned machine from an ISO image, and then passes control to the specified SCCM task sequence.

SCCM provisioning is supported for the deployment of Windows operating systems. Linux is not supported. Software distribution and updates are not supported.

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

- 1 Consult with your network administrator to ensure that the following network requirements are met:
 - Communication with SCCM requires the NetBios name of the SCCM server. At least one Distributed Execution Manager must be able to resolve the fully qualified name of the SCCM server to its NetBios name.
 - The SCCM server and the vRealize Automation server must be on the same network and available to each other.
- 2 Create a software package that includes the vRealize Automation guest agent. See [Create a Software Package for SCCM Provisioning](#).
- 3 In SCCM, create the desired task sequence for provisioning the machine. The final step must be to install the software package you created that contains the vRealize Automation guest agent. For information about creating task sequences and installing software packages, please see the SCCM documentation.
- 4 Create a zero touch boot ISO image for the task sequence. By default, SCCM creates a light touch boot ISO image. For information about configuring SCCM for zero touch ISO images, see the SCCM documentation.
- 5 Copy the ISO image to the location required by your virtualization platform. If you do not know the appropriate location, please refer to the documentation provided by your hypervisor.

- 6 Gather the following information so that tenant administrators and business group managers can include it in their blueprints:
 - a The name of the collection containing the task sequence.
 - b The fully qualified domain name of the SCCM server on which the collection containing the sequence resides.
 - c The site code of the SCCM server.
 - d Administrator-level credentials for the SCCM server.
 - e (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 `SCCMProvisioningProperties` 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.

Create a Software Package for SCCM Provisioning

The final step in your SCCM task sequence must be to install a software package that includes the vRealize Automation guest agent.

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 Windows guest agent files.
 - Windows guest agent files (**32-bit**.)
 - Windows guest agent files (**64-bit**.)
- 3 Extract the Windows guest agent files to a location available to SCCM.
- 4 Create a software package from the definition file `SCCMPackageDefinitionFile.sms`.
- 5 Make the software package available to your distribution point.
- 6 Select the contents of the extracted Windows guest agent files as your source files.

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](#). 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](#).
- 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](#).
- 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](#). 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](#).
- 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 3-7. 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 3-8. Required SysPrep Settings for reference machine that are not using Windows Server 2003 or Windows XP:

AutoLogon Settings	Value
Enabled	Yes
LogonCount	1
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 Sysprep.GuiUnattended.AdminPassword. 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.](#)

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.

Preparing for WIM Provisioning with VirtIO Drivers

If you are using VirtIO for network or storage interfaces, you must ensure that the necessary drivers are included in your WinPE image and WIM image. VirtIO generally offers better performance when provisioning with KVM (RHEV).

Windows drivers for VirtIO are included as part of the Red Hat Enterprise Virtualization and are located in the `/usr/share/virtio-win` directory on the file system of the Red Hat Enterprise Virtualization Manager. The drivers are also included in the Red Hat Enterprise Virtualization Guest Tools located `/usr/share/rhev-guest-tools-iso/rhev-tools-setup.iso`.

The high-level process for enabling WIM-based provisioning with VirtIO drivers is as follows:

- 1 Create a WIM image from a Windows reference machine with the VirtIO drivers installed or insert the drivers into an existing WIM image.
- 2 Copy the VirtIO driver files to the `Plugins` subdirectory of your PEBuilder installation directory before creating a WinPE image, or insert the drivers into a WinPE image created using other means.

- 3 Upload the WinPE image ISO to the Red Hat Enterprise Virtualization ISO storage domains using the `rhevms-iso-uploader` command. For more information about managing ISO images in RHEV refer to the Red Hat documentation.
- 4 A tenant administrator or business group manager creates a KVM (RHEV) blueprint for WIM provisioning and selects the WinPE ISO on the **Build** information tab. The custom property `VirtualMachine.Admin.DiskInterfaceType` must be included with the value **VirtIO**. A fabric administrator can include this information in a build profile to make it easier for tenant administrators and business group managers to include in their blueprints.

The custom properties `Image.ISO.Location` and `Image.ISO.Name` are not used for KVM (RHEV) blueprints.

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](#).
- (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](#).
- (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](#).

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, <code>manager_service_LB.mycompany.com</code>. b Enter the port number for the IaaS Manager Service load balancer in the vCAC Port text box. For example, <code>443</code>.
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, <code>manager_service.mycompany.com</code>. b Enter the port number for the IaaS Manager Service machine in the vCAC Port text box. For example, <code>443</code>.

- 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](#)

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](#)

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

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

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](#)

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.](#)

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 `doagent.bat` File.](#)

Procedure

- 1 Navigate to the `VRMGuestAgent` directory within your WinPE Image.
For example: `C:\Program Files (x86)\VMware\PE Builder\Plugins\VRMGuestAgent\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.](#)

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 `gugent.properties` and name it `gugent.properties.template`.
- 3 Make a copy of the file `gugent.properties.template` and name it `gugentc.properties`.
- 4 Open `gugent.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 `gugentc.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 Blueprints

Machine blueprints are the complete specification for a virtual, cloud, or physical machine. Blueprints determine a machine's attributes, the manner in which it is provisioned, and its policy and management settings.

This chapter includes the following topics:

- [Space-Efficient Storage for Virtual Provisioning](#)
- [Choosing a Blueprint Scenario](#)
- [Create a Blueprint for the Basic Workflow](#)
- [Create a Blueprint for Cloning](#)
- [Create a Linked Clone Blueprint](#)
- [Create a Blueprint for Net App FlexClone Provisioning](#)
- [Create a Blueprint for WIM Provisioning](#)
- [Create a Blueprint for Linux Kickstart Provisioning](#)
- [Create a Blueprint for SCCM Provisioning](#)
- [Troubleshooting Blueprints for Clone and Linked Clone](#)
- [Publish a Blueprint](#)

Space-Efficient Storage for Virtual Provisioning

Space-efficient storage technology eliminates the inefficiencies of traditional storage methods by using only the storage actually required for a machine's operations. Typically, this is only a fraction of the storage actually allocated to machines. vRealize Automation supports two methods of provisioning with space-efficient technology, thin provisioning and FlexClone provisioning.

When standard storage is used, the storage allocated to a provisioned machine is fully committed to that machine, even when it is powered off. This can be a significant waste of storage resources because few virtual machines actually use all of the storage allocated to them, just as few physical machines operate with a 100% full disk. When a space-efficient storage technology is used, the storage allocated and the storage used are tracked separately and only the storage used is fully committed to the provisioned machine.

Thin Provisioning

Thin provisioning is supported for all virtual provisioning methods. Depending on your virtualization platform, storage type, and default storage configuration, thin provisioning might always be used during machine provisioning. For example, for vSphere ESX Server integrations using NFS storage, thin provisioning is always employed. However, for vSphere ESX Server integrations that use local or iSCSI storage, thin provisioning is only used to provision machines if the custom property `VirtualMachine.Admin.ThinProvision` is specified in the blueprint. For more information about thin provisioning, please see the documentation provided by your virtualization platform.

Net App FlexClone Provisioning

You can create a blueprint for Net App FlexClone provisioning if you are working in a vSphere environment that uses Network File System (NFS) storage and FlexClone technology.

You can only use NFS storage, or machine provisioning fails. You can specify a FlexClone storage path for other types of machine provisioning, but the FlexClone storage path behaves like standard storage.

The following is a high-level overview of the sequence of steps required to provision machines that use FlexClone technology:

- 1 An IaaS administrator creates a NetApp ONTAP endpoint. See [Create a NetApp ONTAP Endpoint](#).
- 2 An IaaS administrator runs data collection on the endpoint to enable the endpoint to be visible on the compute resource and reservation pages.

The FlexClone option is visible on a reservation page in the endpoint column if a NetApp ONTAP endpoint exists and if the host is virtual. If there is a NetApp ONTAP endpoint, the reservation page displays the endpoint assigned to the storage path.

- 3 A fabric administrator creates a vSphere reservation, enables FlexClone storage, and specifies an NFS storage path that uses FlexClone technology. See [Create a Reservation](#).
- 4 A tenant administrator or business group manager creates a FlexClone blueprint. See [Create a Blueprint for Net App FlexClone Provisioning](#).

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
Provision machines with no guest operating system. You can install an operating system after provisioning. This option does not require any pre-provisioning preparations.	All	Create a Blueprint for the Basic Workflow	Custom Properties for Basic Workflow Blueprints
Provision machines by cloning from a template object created from an existing Windows or Linux machine, called the reference machine, and a customization specification.	<ul style="list-style-type: none"> ■ vSphere ■ KVM (RHEV) ■ SCVMM 	Create a Blueprint for Cloning	Custom Properties for Clone Blueprints
Provision a space-efficient copy of a virtual machine called a linked clone. Linked clones are based on a snapshot of a VM and use a chain of delta disks to track differences from a parent machine.	vSphere	Create a Linked Clone Blueprint	Custom Properties for Linked Clone Blueprints
Provision a space-efficient copy of a virtual machine by using Net App FlexClone technology.	vSphere	Create a Blueprint for Net App FlexClone Provisioning	Custom Properties for FlexClone Blueprints
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.	All	Create a Blueprint for Linux Kickstart Provisioning	Custom Properties for Linux Kickstart Blueprints
Provision a machine and pass control to an SCCM task sequence to boot from an ISO image, deploy a Windows operating system, and install the vRealize Automation guest agent.	All	Create a Blueprint for SCCM Provisioning	Custom Properties for SCCM Blueprints
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.	All	Create a Blueprint for WIM Provisioning	Custom Properties for WIM Blueprints

Create a Blueprint for the Basic Workflow

A basic workflow blueprint provisions a virtual machine with no guest operating system. You can install the operating system after the machine is provisioned.

Prerequisites

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

Procedure

1 [Specify Basic Workflow 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.

2 [Specify Basic Workflow Build Information](#)

The build information settings determine how a machine is provisioned.

3 [Configure Basic Workflow Blueprint Machine Resources](#)

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

4 [Add Blueprint Custom Properties](#)

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

5 [Configure Blueprint Actions](#)

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

Specify Basic Workflow 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**.

Procedure

1 Select **Infrastructure > Blueprints > Blueprints**.

2 Select **New Blueprint > Virtual** and select the type of blueprint you are creating.

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

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

5 (Optional) Select the **Master** check box to allow users to copy your blueprint.

6 (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.

7 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.

8 (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.

9 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.

10 (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.

11 Specify the number of days to archive machines provisioned from this blueprint in the **Archive (days)** text box.

Enter **0** if you do not want to archive machines.

12 (Optional) Set the daily cost of the machine by typing the amount in the **Cost (daily)** text box.

This cost is added to any cost profiles that your fabric administrator sets up.

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

Specify Basic Workflow Build Information

The build information settings determine how a machine is provisioned.

Prerequisites

[Specify Basic Workflow Blueprint Information.](#)

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 **Create** from the **Action** drop-down menu.
- 4 Select **BasicVmWorkflow** from the **Provisioning workflow** drop-down menu.

- 5 (Optional) If you are working with a KVM (RHEV) or Hyper-V (SCVMM) blueprint, you can provide an ISO for provisioning. For Hyper-V (SCVMM), you can also select additional profiles to attach to machines provisioned with your blueprint.
 - a Select a virtual hard disk from the **Virtual hard disk** drop-down menu.
 - b Select a profile from the **Hardware profile** drop-down menu.
 - c Select a bootable ISO image from the **ISO** drop-down menu.

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

Configure Basic Workflow Blueprint Machine Resources

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

Prerequisites

[Specify Basic Workflow Build Information.](#)

Procedure

- 1 Specify CPU settings for provisioned machines.
 - a Enter the minimum number of CPUs 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 CPUs in the **Maximum** text box to allow users to select their own settings within the range that you provide.
- 2 Specify memory settings for provisioned machines.
 - a Enter the minimum amount of memory 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 memory in the **Maximum** text box to allow users to select their own settings within the range that you provide.
- 3 Specify 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.

- 4 (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.

5 Add **Storage volumes**.

- a Click **New Volume**.
- b Enter the capacity in the **Capacity (GB)** text box.
- c (Optional) Enter a drive letter and mount path in the **Drive Letter/Mount Path** text box.
- d (Optional) Enter a name in the **Label** text box.
- e (Optional) Select a **Storage Reservation Policy** from the drop-down menu.

This option requires additional configuration by a fabric administrator to create storage reservation policies.

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

6 (Optional) Select **Allow user to see and change storage reservation policies**.

If selected, users who are provisioning machines can see and change the storage reservation policies that you apply to blueprints.

7 Enter the maximum number of volumes a user can select for provisioned machines in the **Max # of volumes** text box.

Leave blank to accept the default of 15, or enter 0 to disable adding volumes.

8 Enter the maximum number of network adapters a user can select for provisioned machines in the **Max # of network adapters** text box.

For example, if your integration platform has a limit of 10 network adapters, set your maximum at or below this threshold so that users cannot cause provisioning to fail by selecting more network adapters than your platform supports.

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

Add Blueprint Custom Properties

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

Note You might not need to add custom properties one at a time through the vRealize Automation console, instead a fabric administrator can create build profiles that contain common groupings of custom properties.

Prerequisites

[Configure Basic Workflow Blueprint Machine Resources.](#)

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.

Custom Properties for Basic Workflow Blueprints

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

Table 4-2. Custom Properties for Basic Workflow Blueprints

Custom Property	Description
VirtualMachine.CDRom.Attach	Set to False to provision the machine without a CD-ROM device. The default is True.
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.

Table 4-2. Custom Properties for Basic Workflow Blueprints (Continued)

Custom Property	Description
<code>VirtualMachine.DiskN.StorageReservationPolicy</code>	Specifies the storage reservation policy to use to find storage for disk <i>N</i> . Also assigns the named storage reservation policy to a volume. To use this property, substitute the volume number for <i>N</i> in the property name and specify a storage reservation policy name as the value. This property is equivalent to the storage reservation policy name specified on the blueprint. Disk numbering must be sequential. This property is valid for all Virtual and vCloud reservations. This property is not valid for Physical or Cloud reservations, other than vCloud reservations.
<code>VirtualMachine.Storage.AllocationType</code>	Stores collected groups to a single datastore. A distributed environment stores disks round-robin style.
<code>VirtualMachine.Storage.Name</code>	Identifies the storage path on which the machine resides. The default is the value specified in the reservation that was used to provision the machine.
<code>VirtualMachine.Storage.ReserveMemory</code>	Set to True to manage vSwap storage allocation to ensure availability and set allocation in the reservation. vSwap allocation is considered when you create or reconfigure a virtual machine. vSwap allocation checking is only available for vCenter Server endpoints. For multi-machine blueprints, add the property to the component blueprints. Note If you do not specify the <code>VirtualMachine.Storage.ReserveMemory</code> custom property when you create or provision the machine from vRealize Automation, swap space availability is not ensured. If you add the property for an already provisioned machine, and the allocated reservation is full, the storage allocated in the reservation might exceed the actual allocated storage.
<code>VMware.Hardware.Version</code>	Specifies the VM hardware version to be used for vSphere settings. Supported values are currently vmx-04, vmx-07, vmx-08, vmx-09 and vmx-10. This property is applicable for VM Create and VM Update workflows and is available only for basic workflow blueprints.

Configure 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 Blueprint Custom Properties.](#)

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 (Optional) Specify whether to allow users to use Snapshot Manager to create, apply, and delete snapshots of their machine.

Depending on your vRealize Automation configuration, this option might not be available.

Option	Description
Enable Snapshot Manager	Click Yes .
Disable Snapshot Manager	Click No .

- 4 Specify reconfigure settings.
 - a Specify whether to allow machines provisioned from this blueprint to be reconfigured.

Option	Description
Allow entitled users to reconfigure machines provisioned from this blueprint.	Click Yes .
Do not allow any users to reconfigure machines provisioned from this blueprint.	Click No .

- b Select **Enable execution selector** to allow users to schedule a machine reconfiguration.
 - c Select **Enable power action selector** to allow users to choose whether to shut down their machine before reconfiguration.

If you disable this option, machines are shut down before reconfiguration.

- 5 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](#).

Create a Blueprint for Cloning

You can use vRealize Automation to provision machines by cloning from a template object created from a Windows or Linux reference machine.

Prerequisites

- Log in to the vRealize Automation console as a **tenant administrator** or **business group manager**.
- Contact your fabric administrator for information about the templates and customization objects you should use.

Procedure

1 [Specify Clone 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 once the lease period is over.

2 [Specify Clone Blueprint Build Information](#)

The build information settings determine how a machine is provisioned.

3 [Configure Clone Blueprint Machine Resources](#)

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

4 [Add Custom Properties to a Blueprint](#)

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

5 [Configure Blueprint Actions](#)

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

Specify Clone 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 once the lease period is over.

Prerequisites

- Log in to the vRealize Automation console as a **tenant administrator** or **business group manager**.
- Contact your fabric administrator for information about the templates and customization objects you should use.

Procedure

- 1 Select **Infrastructure > Blueprints > Blueprints**.
- 2 Select **New Blueprint > Virtual** and select the type of blueprint you are creating.
- 3 Enter a name in the **Name** text box.
- 4 (Optional) Enter a description in the **Description** text box.
- 5 (Optional) Select the **Master** check box to allow users to copy your blueprint.

- 6 (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.

- 7 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.

- 8 (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.

- 9 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.

- 10 (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.

- 11 Specify the number of days to archive machines provisioned from this blueprint in the **Archive (days)** text box.

Enter **0** if you do not want to archive machines.

- 12 (Optional) Set the daily cost of the machine by typing the amount in the **Cost (daily)** text box.

This cost is added to any cost profiles that your fabric administrator sets up.

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

Specify Clone Blueprint Build Information

The build information settings determine how a machine is provisioned.

Prerequisites

[Specify Clone Blueprint Information.](#)

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 **Clone** from the **Action** drop-down menu.
- 4 Select **CloneWorkflow** from the **Provisioning workflow** drop-down menu.
- 5 Click the **Browse** icon next to the **Clone from** text box.
- 6 Select a template to clone.

Blueprint Type	Available Templates
Shared	The list includes all templates for all reservations. No validation is available to ensure that the template you select exists in the reservation used to provision a machine from this blueprint.
Local (you selected a business group for the blueprint)	You only see templates that exist on reservations for the business group you selected.

- 7 Click **OK**.
- 8 (Optional) Enter the name of the customization specification in the **Customization spec** text box.

A customization specification is required only if you are cloning with static IP addresses. However, you cannot perform customizations of Windows machines without a customization specification object. For Linux clone machines, you can use a customization specification, an external script, or both to perform customizations.

When you create a customization specification in the vSphere Client for DHCP, use the **Typical settings** option on the **Customization Specification > Network** tab page. The **Typical settings** option is related to network settings in vRealize Automation and is indirectly related to network profiles that have the DHCP private network profile and NAT one-to-many settings enabled. If the clone machine is provisioned by using the custom specification, and no profiles are enabled on the vRealize Automation virtual appliance, the machine IP address is selected from the vSphere Client DHCP.

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

Configure Clone Blueprint Machine Resources

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

Depending on your platform type, not all machine resource settings are available.

Prerequisites

[Specify Clone Blueprint Build Information.](#)

Procedure

- 1 Specify CPU settings for provisioned machines.
 - a Enter the minimum number of CPUs 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 CPUs in the **Maximum** text box to allow users to select their own settings within the range that you provide.
- 2 Specify memory settings for provisioned machines.
 - a Enter the minimum amount of memory 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 memory in the **Maximum** text box to allow users to select their own settings within the range that you provide.
- 3 Specify 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.
- 4 (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.
- 5 Add **Storage volumes**.
 - a Click **New Volume**.
 - b Enter the capacity in the **Capacity (GB)** text box.
 - c (Optional) Enter a drive letter and mount path in the **Drive Letter/Mount Path** text box.
 - d (Optional) Enter a name in the **Label** text box.

- e (Optional) Select a **Storage Reservation Policy** from the drop-down menu.

This option requires additional configuration by a fabric administrator to create storage reservation policies.

- f Click the **Save** icon (✔).

- 6 (Optional) Select **Allow user to see and change storage reservation policies**.

If selected, users who are provisioning machines can see and change the storage reservation policies that you apply to blueprints.

- 7 Enter the maximum number of volumes a user can select for provisioned machines in the **Max # of volumes** text box.

Leave blank to accept the default of **15**, or enter **0** to disable adding volumes.

- 8 Enter the maximum number of network adapters a user can select for provisioned machines in the **Max # of network adapters** text box.

For example, if your integration platform has a limit of 10 network adapters, set your maximum at or below this threshold so that users cannot cause provisioning to fail by selecting more network adapters than your platform supports.

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

Add Custom Properties to a Blueprint

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

Note You might not need to add custom properties one at a time through the vRealize Automation console, instead a fabric administrator can create build profiles that contain common groupings of custom properties.

Prerequisites

[Configure Clone Blueprint Machine Resources.](#)

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.

Custom Properties for Clone Blueprints

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

Table 4-3. Custom Properties for Clone Blueprints

Custom Property	Description
<code>VirtualMachine.DiskN.StorageReservationPolicy</code>	Specifies the storage reservation policy to use to find storage for disk <i>N</i> . Also assigns the named storage reservation policy to a volume. To use this property, substitute the volume number for <i>N</i> in the property name and specify a storage reservation policy name as the value. This property is equivalent to the storage reservation policy name specified on the blueprint. Disk numbering must be sequential. This property is valid for all Virtual and vCloud reservations. This property is not valid for Physical or Cloud reservations, other than vCloud reservations.
<code>VirtualMachine.NetworkN.ProfileName</code>	<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> <p>With WIM-based provisioning for virtual machines, you can use this property to specify a network profile and network interface or you can use the Network section of the Virtual Reservation page. You can also assign the network interface to a virtual network using the <code>VirtualMachine.NetworkN.Name</code> custom property.</p>

Table 4-3. Custom Properties for Clone Blueprints (Continued)

Custom Property	Description
<code>Linux.ExternalScript.Name</code>	Specifies the name of an optional customization script, for example <code>config.sh</code> , that the Linux guest agent runs after the operating system is installed. This property is available for Linux machines cloned from templates on which the Linux agent is installed. If you specify an external script, you must also define its location by using the <code>Linux.ExternalScript.LocationType</code> and <code>Linux.ExternalScript.Path</code> properties.
<code>Linux.ExternalScript.LocationType</code>	Specifies the location type of the customization script named in the <code>Linux.ExternalScript.Name</code> property. This can be either <code>local</code> or <code>nfs</code> . You must also specify the script location using the <code>Linux.ExternalScript.Path</code> property. If the location type is <code>nfs</code> , also use the <code>Linux.ExternalScript.Server</code> property.
<code>Linux.ExternalScript.Server</code>	Specifies the name of the NFS server, for example <code>lab-ad.lab.local</code> , on which the Linux external customization script named in <code>Linux.ExternalScript.Name</code> is located.
<code>Linux.ExternalScript.Path</code>	Specifies the local path to the Linux customization script or the export path to the Linux customization on the NFS server. The value must begin with a forward slash and not include the file name, for example <code>/scripts/linux/config.sh</code> .

If your administrators installed the guest agent to run scripts that accept custom properties and customize provisioned machines, you can use custom properties to further customize cloned machines that use the guest agent.

Table 4-4. Custom Properties for Customizing Cloned Machines with a Guest Agent

Custom Property	Description
<code>VirtualMachine.Admin.AddOwnerToAdmins</code>	Set to <code>True</code> (default) to add the machine's owner, as specified by the <code>VirtualMachine.Admin.Owner</code> property, to the local administrators group on the machine.
<code>VirtualMachine.Admin.AllowLogin</code>	Set to <code>True</code> (default) to add the machine owner to the local remote desktop users group, as specified by the <code>VirtualMachine.Admin.Owner</code> property.
<code>VirtualMachine.Admin.UseGuestAgent</code>	If the guest agent is installed as a service on a template for cloning, set to <code>True</code> on the machine blueprint to enable the guest agent service on machines cloned from that template. When the machine is started, the guest agent service is started. Set to <code>False</code> to disable the guest agent. If set to <code>False</code> , the enhanced clone workflow will not use the guest agent for guest operating system tasks, reducing its functionality to <code>VMwareCloneWorkflow</code> . If not specified or set to anything other than <code>False</code> , the enhanced clone workflow will send work items to the guest agent.

Table 4-4. Custom Properties for Customizing Cloned Machines with a Guest Agent (Continued)

Custom Property	Description
VirtualMachine.DiskN.Active	Set to True (default) to specify that the machine's disk <i>N</i> is active. Set to False to specify that the machine's disk <i>N</i> is not active.
VirtualMachine.DiskN.Size	Defines the size in GB of disk <i>N</i> . For example, to give a size of 150 GB to a disk G, define the custom property VirtualMachine.Disk0.Size and enter a value of 150. Disk numbering must be sequential. By default a machine has one disk referred to by VirtualMachine.Disk0.Size, where size is specified by the storage value on the blueprint from which the machine is provisioned. The storage value on the blueprint user interface overwrites the value in the VirtualMachine.Disk0.Size property. The VirtualMachine.Disk0.Size property is not available as a custom property because of its relationship with the storage option on the blueprint. More disks can be added by specifying VirtualMachine.Disk1.Size, VirtualMachine.Disk2.Size and so on. VirtualMachine.Admin.TotalDiskUsage always represents the total of the .DiskN.Size properties plus the VMware.Memory.Reservation size allocation.
VirtualMachine.DiskN.Label	Specifies the label for a machine's disk. The disk label maximum is 32 characters. Disk numbering must be sequential. When used in conjunction with a guest agent, specifies the label of a machine's disk <i>N</i> inside the guest operating system.
VirtualMachine.DiskN.Letter	Specifies the drive letter or mount point of a machine's disk <i>N</i> . The default is C. For example, to specify the letter D for Disk 1, define the custom property as VirtualMachine.Disk1.Letter and enter the value D. Disk numbering must be sequential. When used in conjunction with a guest agent, this value specifies the drive letter or mount point under which an additional disk <i>N</i> is mounted by the guest agent in the guest operating system.
VirtualMachine.Admin.CustomizeGuestOSDelay	Specifies the time to wait after customization is complete and before starting the guest operating system customization. The value must be in HH:MM:SS format. If the value is not set, the default value is one minute (00:01:00). If you choose not to include this custom property, provisioning can fail if the virtual machine reboots before guest agent work items are completed.
VirtualMachine.Customize.WaitComplete	Set to True to prevent the provisioning workflow from sending work items to the guest agent until all customizations have been completed.
VirtualMachine.SoftwareN.Name	Specifies the descriptive name of a software application or script to install or run during provisioning. This is an optional and information-only property. It serves no real function for the enhanced clone workflow or the guest agent but it is useful for a custom software selection in a user interface or for software usage reporting.

Table 4-4. Custom Properties for Customizing Cloned Machines with a Guest Agent (Continued)

Custom Property	Description
<code>VirtualMachine.SoftwareN.ScriptPath</code>	Specifies the full path to an application's install script. The path must be a valid absolute path as seen by the guest operating system and must include the name of the script file name. You can pass custom property values as parameters to the script by inserting <code>{CustomPropertyName}</code> in the path string. For example, if you have a custom property named <code>ActivationKey</code> whose value is 1234, the script path is <code>D:\InstallApp.bat -key {ActivationKey}</code> . The guest agent runs the command <code>D:\InstallApp.bat -key 1234</code> . Your script file can then be programmed to accept and use this value.
<code>VirtualMachine.SoftwareN.ISOName</code>	Specifies the path and file name of the ISO file relative to the data store root. The format is <code>/folder_name/subfolder_name/file_name.iso</code> . If a value is not specified, the ISO is not mounted.
<code>VirtualMachine.SoftwareN.ISOLocation</code>	Specifies the storage path that contains the ISO image file to be used by the application or script. Format the path as it appears on the host reservation, for example <code>netapp-1:it_nfs_1</code> . If a value is not specified, the ISO is not mounted.

Configure 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 Custom Properties to a Blueprint.](#)

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 (Optional) Specify whether to allow users to use Snapshot Manager to create, apply, and delete snapshots of their machine.

Depending on your vRealize Automation configuration, this option might not be available.

Option	Description
Enable Snapshot Manager	Click Yes .
Disable Snapshot Manager	Click No .

- 4 Specify reconfigure settings.

- a Specify whether to allow machines provisioned from this blueprint to be reconfigured.

Option	Description
Allow entitled users to reconfigure machines provisioned from this blueprint.	Click Yes .
Do not allow any users to reconfigure machines provisioned from this blueprint.	Click No .

- b Select **Enable execution selector** to allow users to schedule a machine reconfiguration.
- c Select **Enable power action selector** to allow users to choose whether to shut down their machine before reconfiguration.

If you disable this option, machines are shut down before reconfiguration.

- 5 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](#).

Create a Linked Clone Blueprint

If you are working with a vSphere environment, you can provision a space-efficient copy of a virtual machine called a linked clone. Linked clones are based on a snapshot of a VM and use a chain of delta disks to track differences from a parent machine. Linked clones are provisioned quickly, reduce storage cost, and are ideal to use when performance is not a high priority.

You can create a vSphere linked clone blueprint by using any vSphere virtual machine you manage as a parent machine. Tenant administrators can create linked clone blueprints from any of the virtual machines in their infrastructure, and business group managers can create linked clone blueprints from any of the virtual machines requested by users in their business group.

Prerequisites

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

- A snapshot of a virtual machine. These do not have to be created in vRealize Automation. Snapshots created in vCenter Server are discovered during inventory data collection.

Procedure

1 Specify Linked Clone 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 once the lease period is over.

2 Specify Linked Clone Build Information

The build information settings determine how a machine is provisioned.

3 Configure Linked Clone Machine Resources

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

4 Add Custom Properties to a Blueprint

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

5 Configure Blueprint Actions

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

What to do next

Publish your blueprint to make it available as a catalog item. See [Publish a Blueprint](#).

Optionally, you can add a storage reservation policy or a single datastore to your linked clone blueprint. See [Choosing a Storage Scenario](#).

Specify Linked Clone 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 once the lease period is over.

Prerequisites

- Log in to the vRealize Automation console as a **tenant administrator** or **business group manager**.
- A snapshot of a virtual machine. These do not have to be created in vRealize Automation. Snapshots created in vSphere are discovered during inventory data collection.

Procedure

- 1 Select **Infrastructure > Blueprints > Blueprints**.
- 2 Click **New Blueprint > Virtual > vSphere (vCenter)**.
- 3 Enter a name in the **Name** text box.
- 4 (Optional) Enter a description in the **Description** text box.
- 5 (Optional) Select the **Master** check box to allow users to copy your blueprint.

- 6 (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.

- 7 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.

- 8 (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.

- 9 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.

- 10 (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.

- 11 Specify the number of days to archive machines provisioned from this blueprint in the **Archive (days)** text box.

Enter **0** if you do not want to archive machines.

- 12 (Optional) Set the daily cost of the machine by typing the amount in the **Cost (daily)** text box.

This cost is added to any cost profiles that your fabric administrator sets up.

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

Specify Linked Clone Build Information

The build information settings determine how a machine is provisioned.

Prerequisites

[Specify Linked Clone Blueprint Information.](#)

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 **Linked Clone** from the **Action** drop-down menu.
- 4 Select **CloneWorkflow** from the **Provisioning workflow** drop-down menu.
- 5 Click the **Browse** icon next to the **Clone from** text box.
- 6 Select a machine to clone.

You can only clone from machines that you manage as a business group manager or tenant administrator.
- 7 Click **Next**.
- 8 Select a snapshot to clone.

You can create a new snapshot from this menu.
- 9 Click **Finish**.
- 10 Click **OK**.
- 11 (Optional) Select **Delete snapshot when the blueprint is deleted** if you do not want to keep the snapshot when the blueprint is deleted.
- 12 (Optional) Enter the name of the customization specification in the **Customization spec** text box.

A customization specification is required only if you are cloning with static IP addresses. However, you cannot perform customizations of Windows machines without a customization specification object. For Linux clone machines, you can use a customization specification, an external script, or both to perform customizations.

When you create a customization specification in the vSphere Client for DHCP, use the **Typical settings** option on the **Customization Specification > Network** tab page. The **Typical settings** option is related to network settings in vRealize Automation and is indirectly related to network profiles that have the DHCP private network profile and NAT one-to-many settings enabled. If the clone machine is provisioned by using the custom specification, and no profiles are enabled on the vRealize Automation virtual appliance, the machine IP address is selected from the vSphere Client DHCP.

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

What to do next

[Configure Clone Blueprint Machine Resources.](#)

Configure Linked Clone Machine Resources

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

Prerequisites

[Specify Linked Clone Build Information.](#)

Procedure

- 1 Specify CPU settings for provisioned machines.
 - a Enter the minimum number of CPUs 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 CPUs in the **Maximum** text box to allow users to select their own settings within the range that you provide.
- 2 Specify memory settings for provisioned machines.
 - a Enter the minimum amount of memory 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 memory in the **Maximum** text box to allow users to select their own settings within the range that you provide.
- 3 Specify 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.
- 4 (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 Custom Properties to a Blueprint

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

Note You might not need to add custom properties one at a time through the vRealize Automation console, instead a fabric administrator can create build profiles that contain common groupings of custom properties.

Prerequisites

[Configure Linked Clone Machine Resources.](#)

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.

Custom Properties for Linked Clone Blueprints

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

Certain vRealize Automation custom properties are required to use with linked clone blueprints.

Table 4-5. Custom Properties for Linked Clone Blueprints

Custom Property	Description
<code>VirtualMachine.DiskN.Storage</code>	Specifies the datastore on which to place the machine disk <i>N</i> , for example DATASTORE01. This property is also used to add a single datastore to a linked clone blueprint. <i>N</i> is the index (starting at 0) of the volume to assign. Enter the name of the datastore to assign to the volume. This is the datastore name as it appears in the Storage Path on the Edit Compute Resource page. Disk numbering must be sequential.
<code>VirtualMachine.DiskN.StorageReservationPolicy</code>	Specifies the storage reservation policy to use to find storage for disk <i>N</i> . Also assigns the named storage reservation policy to a volume. To use this property, substitute the volume number for <i>N</i> in the property name and specify a storage reservation policy name as the value. This property is equivalent to the storage reservation policy name specified on the blueprint. Disk numbering must be sequential. This property is valid for all Virtual and vCloud reservations. This property is not valid for Physical or Cloud reservations, other than vCloud reservations.
<code>VirtualMachine.DiskN.Size</code>	Defines the size in GB of disk <i>N</i> . For example, to give a size of 150 GB to a disk <i>G</i> , define the custom property <code>VirtualMachine.Disk0.Size</code> and enter a value of 150. Disk numbering must be sequential. By default a machine has one disk referred to by <code>VirtualMachine.Disk0.Size</code> , where size is specified by the storage value on the blueprint from which the machine is provisioned. The storage value on the blueprint user interface overwrites the value in the <code>VirtualMachine.Disk0.Size</code> property. The <code>VirtualMachine.Disk0.Size</code> property is not available as a custom property because of its relationship with the storage option on the blueprint. More disks can be added by specifying <code>VirtualMachine.Disk1.Size</code> , <code>VirtualMachine.Disk2.Size</code> and so on. <code>VirtualMachine.Admin.TotalDiskUsage</code> always represents the total of the <code>.DiskN.Size</code> properties plus the <code>VMware.Memory.Reservation</code> size allocation.
<code>VirtualMachine.DiskN.Label</code>	Specifies the label for a machine's disk. The disk label maximum is 32 characters. Disk numbering must be sequential. When used in conjunction with a guest agent, specifies the label of a machine's disk <i>N</i> inside the guest operating system.
<code>VirtualMachine.DiskN.Letter</code>	Specifies the drive letter or mount point of a machine's disk <i>N</i> . The default is C. For example, to specify the letter D for Disk 1, define the custom property as <code>VirtualMachine.Disk1.Letter</code> and enter the value D. Disk numbering must be sequential. When used in conjunction with a guest agent, this value specifies the drive letter or mount point under which an additional disk <i>N</i> is mounted by the guest agent in the guest operating system.
<code>MaximumProvisionedMachines</code>	Specifies the maximum number of linked clones for one machine snapshot. The default is unlimited.

Table 4-5. Custom Properties for Linked Clone Blueprints (Continued)

Custom Property	Description
<code>Linux.ExternalScript.Name</code>	<p>Specifies the name of an optional customization script, for example <code>config.sh</code>, that the Linux guest agent runs after the operating system is installed. This property is available for Linux machines cloned from templates on which the Linux agent is installed.</p> <p>If you specify an external script, you must also define its location by using the <code>Linux.ExternalScript.LocationType</code> and <code>Linux.ExternalScript.Path</code> properties.</p>
<code>Linux.ExternalScript.LocationType</code>	<p>Specifies the location type of the customization script named in the <code>Linux.ExternalScript.Name</code> property. This can be either <code>local</code> or <code>nfs</code>.</p> <p>You must also specify the script location using the <code>Linux.ExternalScript.Path</code> property. If the location type is <code>nfs</code>, also use the <code>Linux.ExternalScript.Server</code> property.</p>
<code>Linux.ExternalScript.Server</code>	<p>Specifies the name of the NFS server, for example <code>lab-ad.lab.local</code>, on which the Linux external customization script named in <code>Linux.ExternalScript.Name</code> is located.</p>
<code>Linux.ExternalScript.Path</code>	<p>Specifies the local path to the Linux customization script or the export path to the Linux customization on the NFS server. The value must begin with a forward slash and not include the file name, for example <code>/scripts/linux/config.sh</code>.</p>

If you installed the guest agent to customize cloned machines, you use some custom properties more often than others.

Table 4-6. Custom Properties for Customizing Cloned Machines with a Guest Agent

Custom Property	Description
<code>VirtualMachine.Admin.UseGuestAgent</code>	<p>If the guest agent is installed as a service on a template for cloning, set to <code>True</code> on the machine blueprint to enable the guest agent service on machines cloned from that template. When the machine is started, the guest agent service is started. Set to <code>False</code> to disable the guest agent. If set to <code>False</code>, the enhanced clone workflow will not use the guest agent for guest operating system tasks, reducing its functionality to <code>VMwareCloneWorkflow</code>. If not specified or set to anything other than <code>False</code>, the enhanced clone workflow will send work items to the guest agent.</p>
<code>VirtualMachine.Admin.CustomizeGuestOSDelay</code>	<p>Specifies the time to wait after customization is complete and before starting the guest operating system customization. The value must be in <code>HH:MM:SS</code> format. If the value is not set, the default value is one minute (<code>00:01:00</code>). If you choose not to include this custom property, provisioning can fail if the virtual machine reboots before guest agent work items are completed.</p>

Table 4-6. Custom Properties for Customizing Cloned Machines with a Guest Agent (Continued)

Custom Property	Description
<code>VirtualMachine.Customize.WaitComplete</code>	Set to True to prevent the provisioning workflow from sending work items to the guest agent until all customizations have been completed.
<code>VirtualMachine.SoftwareN.ScriptPath</code>	Specifies the full path to an application's install script. The path must be a valid absolute path as seen by the guest operating system and must include the name of the script file name. You can pass custom property values as parameters to the script by inserting <code>{CustomPropertyName}</code> in the path string. For example, if you have a custom property named <code>ActivationKey</code> whose value is 1234, the script path is <code>D:\InstallApp.bat -key {ActivationKey}</code> . The guest agent runs the command <code>D:\InstallApp.bat -key 1234</code> . Your script file can then be programmed to accept and use this value.

Configure 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 Custom Properties to a Blueprint.](#)

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 (Optional) Specify whether to allow users to use Snapshot Manager to create, apply, and delete snapshots of their machine.

Depending on your vRealize Automation configuration, this option might not be available.

Option	Description
Enable Snapshot Manager	Click Yes .
Disable Snapshot Manager	Click No .

- 4 Specify reconfigure settings.

- a Specify whether to allow machines provisioned from this blueprint to be reconfigured.

Option	Description
Allow entitled users to reconfigure machines provisioned from this blueprint.	Click Yes .
Do not allow any users to reconfigure machines provisioned from this blueprint.	Click No .

- b Select **Enable execution selector** to allow users to schedule a machine reconfiguration.
- c Select **Enable power action selector** to allow users to choose whether to shut down their machine before reconfiguration.

If you disable this option, machines are shut down before reconfiguration.

- 5 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](#).

Optionally, you can add a storage reservation policy or a single datastore to your linked clone blueprint. See [Choosing a Storage Scenario](#).

Create a Blueprint for Net App FlexClone Provisioning

You can use vRealize Automation to provision clone machines using Net App FlexClone technology if you are working with a vSphere integration.

Prerequisites

- Log in to the vRealize Automation console as a **tenant administrator** or **business group manager**.
- Contact your fabric administrator for information about the templates and customization objects you should use.

Procedure

1 Specify FlexClone 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 once the lease period is over.

2 Specify FlexClone Blueprint Build Information

The build information settings determine how a machine is provisioned.

3 Configure FlexClone Blueprint Machine Resources

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

4 Add Custom Properties to a Blueprint

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

5 Configure Blueprint Actions

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

Specify FlexClone 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 once the lease period is over.

Prerequisites

- Log in to the vRealize Automation console as a **tenant administrator** or **business group manager**.
- Contact your fabric administrator for information about the templates and customization objects you should use.

Procedure

1 Select **Infrastructure > Blueprints > Blueprints**.

2 Click **New Blueprint > Virtual > vSphere (vCenter)**.

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

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

5 (Optional) Select the **Master** check box to allow users to copy your blueprint.

6 (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.

7 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.

8 (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.

9 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.

10 (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.

11 Specify the number of days to archive machines provisioned from this blueprint in the **Archive (days)** text box.

Enter **0** if you do not want to archive machines.

12 (Optional) Set the daily cost of the machine by typing the amount in the **Cost (daily)** text box.

This cost is added to any cost profiles that your fabric administrator sets up.

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

Specify FlexClone Blueprint Build Information

The build information settings determine how a machine is provisioned.

Prerequisites

[Specify FlexClone Blueprint Information.](#)

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 **NetApp FlexClone** from the **Action** drop-down menu.
- 4 Select **CloneWorkflow** from the **Provisioning workflow** drop-down menu.
- 5 Click the **Browse** icon next to the **Clone from** text box.

6 Select a template to clone.

Blueprint Type	Available Templates
Shared	The list includes all templates for all reservations. No validation is available to ensure that the template you select exists in the reservation used to provision a machine from this blueprint.
Local (you selected a business group for the blueprint)	You only see templates that exist on reservations for the business group you selected.

7 Click **OK**.

8 (Optional) Enter the name of the customization specification in the **Customization spec** text box.

A customization specification is required only if you are cloning with static IP addresses. However, you cannot perform customizations of Windows machines without a customization specification object. For Linux clone machines, you can use a customization specification, an external script, or both to perform customizations.

When you create a customization specification in the vSphere Client for DHCP, use the **Typical settings** option on the **Customization Specification > Network** tab page. The **Typical settings** option is related to network settings in vRealize Automation and is indirectly related to network profiles that have the DHCP private network profile and NAT one-to-many settings enabled. If the clone machine is provisioned by using the custom specification, and no profiles are enabled on the vRealize Automation virtual appliance, the machine IP address is selected from the vSphere Client DHCP.

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

Configure FlexClone Blueprint Machine Resources

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

Prerequisites

[Specify FlexClone Blueprint Build Information.](#)

Procedure

- 1 Specify CPU settings for provisioned machines.
 - a Enter the minimum number of CPUs 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 CPUs in the **Maximum** text box to allow users to select their own settings within the range that you provide.

2 Specify memory settings for provisioned machines.

- a Enter the minimum amount of memory 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 memory in the **Maximum** text box to allow users to select their own settings within the range that you provide.

3 Specify 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.

4 (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.

5 Add **Storage volumes**.

- a Click **New Volume**.

- b Enter the capacity in the **Capacity (GB)** text box.

- c (Optional) Enter a drive letter and mount path in the **Drive Letter/Mount Path** text box.

- d (Optional) Enter a name in the **Label** text box.

- e (Optional) Select a **Storage Reservation Policy** from the drop-down menu.

This option requires additional configuration by a fabric administrator to create storage reservation policies.

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

6 Enter the maximum number of volumes a user can select for provisioned machines in the **Max # of volumes** text box.

Leave blank to accept the default of **15**, or enter **0** to disable adding volumes.

- 7 Enter the maximum number of network adapters a user can select for provisioned machines in the **Max # of network adapters** text box.

For example, if your integration platform has a limit of 10 network adapters, set your maximum at or below this threshold so that users cannot cause provisioning to fail by selecting more network adapters than your platform supports.

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

Add Custom Properties to a Blueprint

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

Note You might not need to add custom properties one at a time through the vRealize Automation console, instead a fabric administrator can create build profiles that contain common groupings of custom properties.

Prerequisites

[Configure FlexClone Blueprint Machine Resources.](#)

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.

Custom Properties for FlexClone Blueprints

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

Table 4-7. Custom Properties for FlexClone Blueprints

Custom Property	Description
VirtualMachine.NetworkN.ProfileName	<p>Specifies the name of a network profile from which to assign a static IP address to network device <i>Nor</i> 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=0</i> 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> <p>With WIM-based provisioning for virtual machines, you can use this property to specify a network profile and network interface or you can use the Network section of the Virtual Reservation page. You can also assign the network interface to a virtual network using the <code>VirtualMachine.NetworkN.Name</code> custom property.</p>
Linux.ExternalScript.Name	<p>Specifies the name of an optional customization script, for example <code>config.sh</code>, that the Linux guest agent runs after the operating system is installed. This property is available for Linux machines cloned from templates on which the Linux agent is installed.</p> <p>If you specify an external script, you must also define its location by using the <code>Linux.ExternalScript.LocationType</code> and <code>Linux.ExternalScript.Path</code> properties.</p>
Linux.ExternalScript.LocationType	<p>Specifies the location type of the customization script named in the <code>Linux.ExternalScript.Name</code> property. This can be either <code>local</code> or <code>nfs</code>.</p> <p>You must also specify the script location using the <code>Linux.ExternalScript.Path</code> property. If the location type is <code>nfs</code>, also use the <code>Linux.ExternalScript.Server</code> property.</p>
Linux.ExternalScript.Server	<p>Specifies the name of the NFS server, for example <code>lab-ad.lab.local</code>, on which the Linux external customization script named in <code>Linux.ExternalScript.Name</code> is located.</p>
Linux.ExternalScript.Path	<p>Specifies the local path to the Linux customization script or the export path to the Linux customization on the NFS server. The value must begin with a forward slash and not include the file name, for example <code>/scripts/linux/config.sh</code>.</p>

If you installed the guest agent to customize cloned machines, the Custom Properties for Customizing FlexClone Machines with a Guest Agent table describes the most commonly used custom properties for your situation.

Table 4-8. Custom Properties for Customizing FlexClone Machines with a Guest Agent

Custom Property	Description
<code>VirtualMachine.Admin.UseGuestAgent</code>	If the guest agent is installed as a service on a template for cloning, set to True on the machine blueprint to enable the guest agent service on machines cloned from that template. When the machine is started, the guest agent service is started. Set to False to disable the guest agent. If set to False, the enhanced clone workflow will not use the guest agent for guest operating system tasks, reducing its functionality to <code>VMwareCloneWorkflow</code> . If not specified or set to anything other than False, the enhanced clone workflow will send work items to the guest agent.
<code>VirtualMachine.DiskN.Size</code>	Defines the size in GB of disk <i>N</i> . For example, to give a size of 150 GB to a disk G, define the custom property <code>VirtualMachine.Disk0.Size</code> and enter a value of 150. Disk numbering must be sequential. By default a machine has one disk referred to by <code>VirtualMachine.Disk0.Size</code> , where size is specified by the storage value on the blueprint from which the machine is provisioned. The storage value on the blueprint user interface overwrites the value in the <code>VirtualMachine.Disk0.Size</code> property. The <code>VirtualMachine.Disk0.Size</code> property is not available as a custom property because of its relationship with the storage option on the blueprint. More disks can be added by specifying <code>VirtualMachine.Disk1.Size</code> , <code>VirtualMachine.Disk2.Size</code> and so on. <code>VirtualMachine.Admin.TotalDiskUsage</code> always represents the total of the <code>.DiskN.Size</code> properties plus the <code>VMware.Memory.Reservation</code> size allocation.
<code>VirtualMachine.DiskN.Label</code>	Specifies the label for a machine's disk. The disk label maximum is 32 characters. Disk numbering must be sequential. When used in conjunction with a guest agent, specifies the label of a machine's disk <i>N</i> inside the guest operating system.
<code>VirtualMachine.DiskN.Letter</code>	Specifies the drive letter or mount point of a machine's disk <i>N</i> . The default is C. For example, to specify the letter D for Disk 1, define the custom property as <code>VirtualMachine.Disk1.Letter</code> and enter the value D. Disk numbering must be sequential. When used in conjunction with a guest agent, this value specifies the drive letter or mount point under which an additional disk <i>N</i> is mounted by the guest agent in the guest operating system.
<code>VirtualMachine.Admin.CustomizeGuestOSDelay</code>	Specifies the time to wait after customization is complete and before starting the guest operating system customization. The value must be in HH:MM:SS format. If the value is not set, the default value is one minute (00:01:00). If you choose not to include this custom property, provisioning can fail if the virtual machine reboots before guest agent work items are completed.

Table 4-8. Custom Properties for Customizing FlexClone Machines with a Guest Agent (Continued)

Custom Property	Description
<code>VirtualMachine.Customize.WaitComplete</code>	Set to True to prevent the provisioning workflow from sending work items to the guest agent until all customizations have been completed.
<code>VirtualMachine.SoftwareN.ScriptPath</code>	Specifies the full path to an application's install script. The path must be a valid absolute path as seen by the guest operating system and must include the name of the script file name. You can pass custom property values as parameters to the script by inserting <code>{CustomPropertyName}</code> in the path string. For example, if you have a custom property named <code>ActivationKey</code> whose value is 1234, the script path is <code>D:\InstallApp.bat -key {ActivationKey}</code> . The guest agent runs the command <code>D:\InstallApp.bat -key 1234</code> . Your script file can then be programmed to accept and use this value.

Configure 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 Custom Properties to a Blueprint.](#)

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 (Optional) Specify whether to allow users to use Snapshot Manager to create, apply, and delete snapshots of their machine.

Depending on your vRealize Automation configuration, this option might not be available.

Option	Description
Enable Snapshot Manager	Click Yes .
Disable Snapshot Manager	Click No .

- 4 Specify reconfigure settings.

- a Specify whether to allow machines provisioned from this blueprint to be reconfigured.

Option	Description
Allow entitled users to reconfigure machines provisioned from this blueprint.	Click Yes .
Do not allow any users to reconfigure machines provisioned from this blueprint.	Click No .

- b Select **Enable execution selector** to allow users to schedule a machine reconfiguration.
- c Select **Enable power action selector** to allow users to choose whether to shut down their machine before reconfiguration.

If you disable this option, machines are shut down before reconfiguration.

- 5 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](#).

Create a 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

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

The build information settings determine how a machine is provisioned.

3 Configure WIM Blueprint Machine Resources

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

4 Add Required Properties to a WIM Blueprint

WIM provisioning requires a number of custom properties to locate the WIM image and WinPE.

5 Configure Blueprint Actions

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

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 > Virtual** and select the type of blueprint you are creating.
- 3 Enter a name in the **Name** text box.
- 4 (Optional) Enter a description in the **Description** text box.
- 5 (Optional) Select the **Master** check box to allow users to copy your blueprint.
- 6 (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.

- 7 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.

- 8 (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.
- 9 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.
- 10 (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.
- 11 Specify the number of days to archive machines provisioned from this blueprint in the **Archive (days)** text box.
Enter **0** if you do not want to archive machines.

12 (Optional) Set the daily cost of the machine by typing the amount in the **Cost (daily)** text box.

This cost is added to any cost profiles that your fabric administrator sets up.

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.](#)

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 **Create** from the **Action** drop-down menu.
- 4 Select **WIMImageWorkflow** from the **Provisioning workflow** drop-down menu.
- 5 (Optional) If you are working with a KVM (RHEV) or Hyper-V (SCVMM) blueprint, you can provide an ISO for provisioning. For Hyper-V (SCVMM), you can also select additional profiles to attach to machines provisioned with your blueprint.
 - a Select a virtual hard disk from the **Virtual hard disk** drop-down menu.
 - b Select a profile from the **Hardware profile** drop-down menu.
 - c Select a bootable ISO image from the **ISO** drop-down menu.

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

Configure WIM Blueprint Machine Resources

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

Prerequisites

[Specify WIM Blueprint Build Information.](#)

Procedure

- 1 Specify CPU settings for provisioned machines.
 - a Enter the minimum number of CPUs 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 CPUs in the **Maximum** text box to allow users to select their own settings within the range that you provide.
- 2 Specify memory settings for provisioned machines.
 - a Enter the minimum amount of memory 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 memory in the **Maximum** text box to allow users to select their own settings within the range that you provide.
- 3 Specify 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.
- 4 (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.
- 5 Add **Storage volumes**.
 - a Click **New Volume**.
 - b Enter the capacity in the **Capacity (GB)** text box.
 - c (Optional) Enter a drive letter and mount path in the **Drive Letter/Mount Path** text box.
 - d (Optional) Enter a name in the **Label** text box.

- e (Optional) Select a **Storage Reservation Policy** from the drop-down menu.

This option requires additional configuration by a fabric administrator to create storage reservation policies.

- f Click the **Save** icon (✔).

- 6 Enter the maximum number of volumes a user can select for provisioned machines in the **Max # of volumes** text box.

Leave blank to accept the default of **15**, or enter **0** to disable adding volumes.

- 7 Enter the maximum number of network adapters a user can select for provisioned machines in the **Max # of network adapters** text box.

For example, if your integration platform has a limit of 10 network adapters, set your maximum at or below this threshold so that users cannot cause provisioning to fail by selecting more network adapters than your platform supports.

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

Add Required Properties to a WIM Blueprint

WIM provisioning requires a number of custom properties to locate the WIM image and WinPE.

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 related to WIM provisioning, see [Custom Properties for WIM Blueprints](#).

Prerequisites

[Configure WIM Blueprint Machine Resources](#).

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 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	Type 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	Type the UNC path to the WIM file that your fabric administrator provided for you in the Value text box.
Image.WIM.Name	Type the name of the WIM file that your fabric administrator provided for you in the Value text box.
Image.WIM.Index	Type 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	Type 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	Type the associated password for the network user (<code>Image.Network.User</code>) in the Value text box.

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 4-9. Required Custom Properties for WIM Blueprints

Custom Property	Description
Image.ISO.Location	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.
Image.ISO.Name	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.
Image.ISO.UserName	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.
Image.ISO.Password	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.
Image.Network.Letter	Specifies the drive letter to which the WIM image path is mapped on the provisioned machine. The default value is K.
Image.WIM.Path	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> .
Image.WIM.Name	Specifies the name of the WIM file, for example <code>win2k8.wim</code> , as located by the <code>Image.WIM.Path</code> property.
Image.WIM.Index	Specifies the index used to extract the correct image from the WIM file.
Image.Network.User	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.

Table 4-9. Required Custom Properties for WIM Blueprints (Continued)

Custom Property	Description
Image.Network.Password	Specifies the password associated with the Image.Network.User property.
VMware.VirtualCenter.OperatingSystem	<p>Specifies the vCenter Server guest operating system version (VirtualMachineGuestOsIdentifier) 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, VMware[OS_Version]Properties, that are predefined to include the correct VMware.VirtualCenter.OperatingSystem values. This property is for virtual provisioning.</p> <p>For related information, see the enumeration type VirtualMachineGuestOsIdentifier in vSphere API/SDK Documentation. For a list of currently accepted values, see the VMware vCenter Server™ documentation.</p>

Optional custom properties are also available for WIM blueprints.

Table 4-10. Common Custom Properties for WIM Blueprints

Custom Property	Description
SysPrep. <i>Section.Key</i> <ul style="list-style-type: none"> ■ 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. For a full list of sections, keys, and accepted values, see the System Preparation Utility for Windows documentation. 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 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.](#)

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 (Optional) Specify whether to allow users to use Snapshot Manager to create, apply, and delete snapshots of their machine.

Depending on your vRealize Automation configuration, this option might not be available.

Option	Description
Enable Snapshot Manager	Click Yes .
Disable Snapshot Manager	Click No .

4 Specify reconfigure settings.

- a Specify whether to allow machines provisioned from this blueprint to be reconfigured.

Option	Description
Allow entitled users to reconfigure machines provisioned from this blueprint.	Click Yes .
Do not allow any users to reconfigure machines provisioned from this blueprint.	Click No .

- b Select **Enable execution selector** to allow users to schedule a machine reconfiguration.
- c Select **Enable power action selector** to allow users to choose whether to shut down their machine before reconfiguration.

If you disable this option, machines are shut down before reconfiguration.

5 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](#).

Create a 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](#)

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

The build information settings determine how a machine is provisioned.

3 Configure Linux Kickstart Blueprint Machine Resources

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

4 Add Required Properties to a Linux Kickstart Blueprint

Linux Kickstart provisioning requires custom properties to locate the boot ISO image.

5 Configure Blueprint Actions

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

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 > Virtual** and select the type of blueprint you are creating.
- 3 Enter a name in the **Name** text box.
- 4 (Optional) Enter a description in the **Description** text box.
- 5 (Optional) Select the **Master** check box to allow users to copy your blueprint.
- 6 (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.

7 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.

8 (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.

9 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.

10 (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.

11 Specify the number of days to archive machines provisioned from this blueprint in the **Archive (days)** text box.

Enter **0** if you do not want to archive machines.

12 (Optional) Set the daily cost of the machine by typing the amount in the **Cost (daily)** text box.

This cost is added to any cost profiles that your fabric administrator sets up.

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.](#)

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 **Create** from the **Action** drop-down menu.
- 4 Select **LinuxKickstartWorkflow** from the **Provisioning workflow** drop-down menu.

- 5 (Optional) If you are working with a KVM (RHEV) or Hyper-V (SCVMM) blueprint, you can provide an ISO for provisioning. For Hyper-V (SCVMM), you can also select additional profiles to attach to machines provisioned with your blueprint.
 - a Select a virtual hard disk from the **Virtual hard disk** drop-down menu.
 - b Select a profile from the **Hardware profile** drop-down menu.
 - c Select a bootable ISO image from the **ISO** drop-down menu.

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

Configure Linux Kickstart Blueprint Machine Resources

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

Prerequisites

[Specify Linux Kickstart Blueprint Build Information.](#)

Procedure

- 1 Specify CPU settings for provisioned machines.
 - a Enter the minimum number of CPUs 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 CPUs in the **Maximum** text box to allow users to select their own settings within the range that you provide.
- 2 Specify memory settings for provisioned machines.
 - a Enter the minimum amount of memory 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 memory in the **Maximum** text box to allow users to select their own settings within the range that you provide.
- 3 Specify 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.

- 4 (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.
- 5 Add **Storage volumes**.
 - a Click **New Volume**.
 - b Enter the capacity in the **Capacity (GB)** text box.
 - c (Optional) Enter a drive letter and mount path in the **Drive Letter/Mount Path** text box.
 - d (Optional) Enter a name in the **Label** text box.
 - e (Optional) Select a **Storage Reservation Policy** from the drop-down menu.
This option requires additional configuration by a fabric administrator to create storage reservation policies.
 - f Click the **Save** icon ().
- 6 (Optional) Select **Allow user to see and change storage reservation policies**.
If selected, users who are provisioning machines can see and change the storage reservation policies that you apply to blueprints.
- 7 Enter the maximum number of volumes a user can select for provisioned machines in the **Max # of volumes** text box.
Leave blank to accept the default of 15, or enter 0 to disable adding volumes.
- 8 Enter the maximum number of network adapters a user can select for provisioned machines in the **Max # of network adapters** text box.
For example, if your integration platform has a limit of 10 network adapters, set your maximum at or below this threshold so that users cannot cause provisioning to fail by selecting more network adapters than your platform supports.

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

Add Required Properties to a Linux Kickstart Blueprint

Linux Kickstart provisioning requires custom properties to locate the boot ISO image.

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](#).

Prerequisites

[Configure Linux Kickstart Blueprint Machine Resources](#).

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 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 4-11. 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 4-12. Optional Custom Properties for Linux Kickstart Blueprints

Custom Property	Description
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.
Machine.SSH	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 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.](#)

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 (Optional) Specify whether to allow users to use Snapshot Manager to create, apply, and delete snapshots of their machine.

Depending on your vRealize Automation configuration, this option might not be available.

Option	Description
Enable Snapshot Manager	Click Yes .
Disable Snapshot Manager	Click No .

4 Specify reconfigure settings.

- a Specify whether to allow machines provisioned from this blueprint to be reconfigured.

Option	Description
Allow entitled users to reconfigure machines provisioned from this blueprint.	Click Yes .
Do not allow any users to reconfigure machines provisioned from this blueprint.	Click No .

- b Select **Enable execution selector** to allow users to schedule a machine reconfiguration.
- c Select **Enable power action selector** to allow users to choose whether to shut down their machine before reconfiguration.

If you disable this option, machines are shut down before reconfiguration.

5 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](#).

Create a Blueprint for SCCM Provisioning

You can provision a machine and then pass control to an SCCM task sequence to boot from an ISO image, deploy a Windows operating system, and install the vRealize Automation guest agent.

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 of the collection containing the task sequence.
 - b The fully qualified domain name of the SCCM server on which the collection containing the sequence resides.
 - c The site code of the SCCM server.
 - d Administrator-level credentials for the SCCM server.
 - e (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 SCCM 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.

2 Specify SCCM Build Information

The build information settings determine how a machine is provisioned.

3 Configure SCCM Machine Resources

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

4 Add Required Properties to a SCCM Blueprint

SCCM provisioning requires custom properties to successfully provision machines with the desired ISO and task sequence.

5 Configure Blueprint Actions

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

Specify SCCM 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 of the collection containing the task sequence.
 - b The fully qualified domain name of the SCCM server on which the collection containing the sequence resides.
 - c The site code of the SCCM server.
 - d Administrator-level credentials for the SCCM server.
 - e (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 > Virtual** and select the type of blueprint you are creating.
- 3 Select **New Blueprint > Virtual** and select the type of blueprint you are creating.

- 4 Enter a name in the **Name** text box.
- 5 (Optional) Enter a description in the **Description** text box.
- 6 (Optional) Select the **Master** check box to allow users to copy your blueprint.
- 7 (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.

- 8 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.

- 9 (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.

- 10 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.

- 11 (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.

- 12 Specify the number of days to archive machines provisioned from this blueprint in the **Archive (days)** text box.

Enter **0** if you do not want to archive machines.

- 13 (Optional) Set the daily cost of the machine by typing the amount in the **Cost (daily)** text box.

This cost is added to any cost profiles that your fabric administrator sets up.

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

Specify SCCM Build Information

The build information settings determine how a machine is provisioned.

Prerequisites

[Specify SCCM Blueprint Information.](#)

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 **Create** from the **Action** drop-down menu.
- 4 Select **VirtualSccmProvisioningWorkflow** from the **Provisioning workflow** drop-down menu.
- 5 (Optional) If you are working with a KVM (RHEV) or Hyper-V (SCVMM) blueprint, you can provide an ISO for provisioning. For Hyper-V (SCVMM), you can also select additional profiles to attach to machines provisioned with your blueprint.
 - a Select a virtual hard disk from the **Virtual hard disk** drop-down menu.
 - b Select a profile from the **Hardware profile** drop-down menu.
 - c Select a bootable ISO image from the **ISO** drop-down menu.

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

Configure SCCM Machine Resources

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

Prerequisites

[Specify SCCM Build Information.](#)

Procedure

- 1 Specify CPU settings for provisioned machines.
 - a Enter the minimum number of CPUs 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 CPUs in the **Maximum** text box to allow users to select their own settings within the range that you provide.
- 2 Specify memory settings for provisioned machines.
 - a Enter the minimum amount of memory 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 memory in the **Maximum** text box to allow users to select their own settings within the range that you provide.

3 Specify 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.

4 (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.

5 Add **Storage volumes**.

- a Click **New Volume**.

- b Enter the capacity in the **Capacity (GB)** text box.

- c (Optional) Enter a drive letter and mount path in the **Drive Letter/Mount Path** text box.

- d (Optional) Enter a name in the **Label** text box.

- e (Optional) Select a **Storage Reservation Policy** from the drop-down menu.

This option requires additional configuration by a fabric administrator to create storage reservation policies.

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

6 (Optional) Select **Allow user to see and change storage reservation policies**.

If selected, users who are provisioning machines can see and change the storage reservation policies that you apply to blueprints.

7 Enter the maximum number of volumes a user can select for provisioned machines in the **Max # of volumes** text box.

Leave blank to accept the default of **15**, or enter **0** to disable adding volumes.

8 Enter the maximum number of network adapters a user can select for provisioned machines in the **Max # of network adapters** text box.

For example, if your integration platform has a limit of 10 network adapters, set your maximum at or below this threshold so that users cannot cause provisioning to fail by selecting more network adapters than your platform supports.

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

Add Required Properties to a SCCM Blueprint

SCCM provisioning requires custom properties to successfully provision machines with the desired ISO and task sequence.

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 for SCCM provisioning, see [Custom Properties for SCCM Blueprints](#).

Prerequisites

[Configure SCCM Machine Resources](#).

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 location custom property.
 - a Click **New Property**.
 - b Enter `Image.ISO.Location` 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.
 - 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 SCCM custom properties.

Option	Description
<code>Image.ISO.Name</code>	Enter the location of the ISO image that your fabric administrator provided for you in the Value text box.
<code>SCCM.Collection.Name</code>	Enter the SCCM collection name that your fabric administrator provided for you in the Value text box.
<code>SCCM.Server.Name</code>	Enter the name of the SCCM file that your fabric administrator provided for you in the Value text box.

Option	Description
SCCM.Server.SiteCode	Enter the SCCM site code that your fabric administrator provided for you in the Value text box.
SCCM.Server.UserName	Enter the administrator-level user name that your fabric administrator provided for you in the Value text box.
SCCM.Server.Password	Enter the associated password for the SCCM user in the Value text box.

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

Custom Properties for SCCM Blueprints

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

Certain custom properties are required to use with SCCM blueprints.

Table 4-13. Required Custom Properties for SCCM Blueprints

Custom Property	Description
Image.ISO.Location	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.
Image.ISO.Name	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.
Image.ISO.UserName	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.
Image.ISO.Password	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.
SCCM.Collection.Name	Specifies the name of the SCCM collection that contains the operating system deployment task sequence.
SCCM.Server.Name	Specifies the fully qualified domain name of the SCCM server on which the collection resides, for example <code>lab-sccm.lab.local</code> .
SCCM.Server.SiteCode	Specifies the site code of the SCCM server.

Table 4-13. Required Custom Properties for SCCM Blueprints (Continued)

Custom Property	Description
SCCM.Server.UserName	Specifies a user name with administrator-level access to the SCCM server.
SCCM.Server.Password	Specifies the password associated with the SCCM.Server.UserName property.

Certain custom properties are used most often with SCCM blueprints.

Table 4-14. Common Custom Properties for SCCM Blueprints

Custom Property	Description
SCCM.CustomVariable.Name	Specifies the value of a custom variable, where <i>Name</i> is the name of any custom variable to be made available to the SCCM task sequence after the provisioned machine is registered with the SCCM collection. The value is determined by your choice of custom variable. If your integration requires it, you can use SCCM.RemoveCustomVariablePrefix to remove the SCCM.CustomVariable. prefix from your custom variable.
SCCM.RemoveCustomVariablePrefix	Set to <i>true</i> to remove the prefix SCCM.CustomVariable. from SCCM custom variables you created by using the custom property SCCM.CustomVariable.Name.

Configure 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 SCCM Blueprint.](#)

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 (Optional) Specify whether to allow users to use Snapshot Manager to create, apply, and delete snapshots of their machine.

Depending on your vRealize Automation configuration, this option might not be available.

Option	Description
Enable Snapshot Manager	Click Yes .
Disable Snapshot Manager	Click No .

- 4 Specify reconfigure settings.

- a Specify whether to allow machines provisioned from this blueprint to be reconfigured.

Option	Description
Allow entitled users to reconfigure machines provisioned from this blueprint.	Click Yes .
Do not allow any users to reconfigure machines provisioned from this blueprint.	Click No .

- b Select **Enable execution selector** to allow users to schedule a machine reconfiguration.
- c Select **Enable power action selector** to allow users to choose whether to shut down their machine before reconfiguration.

If you disable this option, machines are shut down before reconfiguration.

- 5 Click **OK**.

Your blueprint is saved in draft state.

What to do next

Your blueprint is saved in draft state and you must manually publish it before you can associate it with a catalog service and entitle users to request it from the catalog. See [Publish a Blueprint](#).

Troubleshooting Blueprints for Clone and Linked Clone

When creating a linked clone or clone blueprint, machine or templates are missing. Using your shared clone blueprint to request machines fails to provision machines.

Problem

When working with clone or linked clone blueprints, you might encounter one of the following problems:

- When you create a linked clone blueprint, no machines appear in the list to clone, or the machine you want to clone does not appear.
- When you create a clone blueprint for a business group, no templates appear in the list of templates to clone, or the template you want does not appear.
- When machines are requested by using your shared clone blueprint, provisioning fails.

Cause

There are multiple possible causes for common clone and linked clone blueprint problems.

Table 4-15. Causes for Common Clone and Linked Clone Blueprints Problems

Problem	Cause	Solution
Machines missing	You can only create linked clone blueprints by using machines you manage as a tenant administrator or business group manager.	A user in your tenant or business group must request a vSphere machine. If you have the appropriate roles, you can do this yourself.
Templates missing	If you are creating a clone blueprint for a business group, then you only see templates that exist on compute resources on which the business group has a reservation.	<ul style="list-style-type: none"> ■ Verify that your fabric administrator created a reservation for your business group on the compute resource containing the templates. ■ If your endpoints are clustered and contain multiple compute resources, verify that your IaaS administrator added the cluster containing the templates to your fabric group. ■ For new templates, verify that IT placed the templates on the cluster included in your fabric group.
Provisioning failure with a shared blueprint	For shared blueprints, no validation is available to ensure that the template you select exists in the reservation used to provision a machine from your shared clone blueprint.	Consider using entitlements to restrict the blueprint to users who have a reservation on the compute resource where the template exists. For more information about entitlements, see <i>Tenant Administration</i> .
Provisioning failure with a guest agent	The virtual machine might be rebooting immediately after the guest operating system customization is completed, but before the guest agent work items are completed, causing provisioning to fail. You can use the custom property <code>VirtualMachine.Admin.CustomizeGuestOSDelay</code> to increase the time delay.	Verify that you have added the custom property <code>VirtualMachine.Admin.CustomizeGuestOSDelay</code> . The value must be in HH:MM:SS format. If the value is not set, the default value is one minute (00:01:00).

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:

- [Assigning Datastores to Machine Volumes](#)
- [Reservation Policies](#)
- [Configuring Network Settings](#)
- [Enable Users to Select Datacenter Locations](#)
- [Allow Requesters to Specify Machine Host Name](#)
- [Enabling Visual Basic Scripts in Provisioning](#)
- [Add Active Directory Cleanup to a Blueprint](#)
- [Enable Machine Users to Create Snapshots](#)
- [RDP Connections](#)
- [Enable Connections Using SSH](#)

Assigning Datastores to Machine Volumes

Storage assignment is determined by reservations. Tenant administrators and business group managers can further restrict storage volume assignment by specifying a single datastore or a storage reservation policy that represents multiple datastores to a volume. With the multi-storage feature you can assign the volumes of a virtual machine to different datastores for the vSphere, KVM (RHEV), and SCVMM platforms.

When you assign the volumes of a virtual machine to different datastores, you can control and use storage space more effectively. For example, you might deploy the operating system volume to a slower, less expensive datastore and the database volume to a faster datastore.

You can assign a single datastore or a storage reservation policy that represents multiple datastores to a volume. When you assign a single datastore to a volume, that datastore is used at provisioning time, if possible. When you assign a storage reservation policy to a volume, one of its datastores is used at provisioning time.

A storage reservation policy is a tag applied to one or more datastores by a fabric administrator and is used to group datastores that have similar characteristics, such as speed or price. A datastore can be assigned to only one storage reservation policy at a time, but a storage reservation policy can have many different datastores. A fabric administrator creates a storage reservation policy and assigns it to one or more datastores. A tenant administrator or business group manager then assigns the storage reservation policy to a volume in a virtual blueprint. When a user requests a virtual machine that uses the blueprint, the storage reservation policy specified in the blueprint is used to select a datastore for the machine's volume.

When you add or edit a volume in a virtual blueprint, you can make the following assignments:

- Assign a single datastore to the volume. This one-to-one relationship assigns only the specified datastore at provisioning time. If the datastore is not available or lacks sufficient space, provisioning fails.
- Assign a storage reservation policy to the volume. This assignment allows the selection of one of the datastores associated with the storage reservation policy at provisioning time. If the policy's datastores are unavailable or lack sufficient space, provisioning fails.
- Assign a storage reservation policy to the volume and include the custom property `VirtualMachine.DiskN.StorageReservationPolicyMode` with the value of **NonExact**. This configuration allows you to assign a datastore that is not included in the storage reservation policy.
- Do not assign a datastore or a storage reservation policy to a volume. Not making an assignment allows selection from all available datastores at provisioning time. Provisioning fails if a datastore with sufficient space is not available.

Datastore priority, as assigned by a datastore reservation, is also used to select a datastore at provisioning time.

Avoid specifying a storage path and a storage reservation policy for the same volume. If you must specify a storage path and a storage reservation policy for the same volume, the storage path takes precedence over the storage reservation policy.

Choosing a Storage Scenario

When creating vSphere, KVM (RHEV), or SCVMM Create, Clone, FlexClone, or Linked Clone blueprints, tenant administrators and business group managers can assign datastores to volumes, or assign storage reservation policies to volumes.

Depending on the type of blueprint you are working with and your storage needs, the procedure to assign datastores to volumes differs slightly.

Table 5-1. Choosing a Storage Scenario

Scenario	Supported Platforms	Procedure
Assign a datastore to a Create, Clone, or FlexClone blueprint. Note SDRS is not supported for FlexClone.	<ul style="list-style-type: none"> ■ vSphere ■ KVM (RHEV) ■ SCVMM 	Add a Datastore to a Create, Clone, or FlexClone Blueprint.
Assign a datastore to a Linked Clone blueprint.	vSphere	Add a Datastore to a Linked Clone Blueprint.
Assign a storage reservation policy to a Create, Clone, or FlexClone blueprint. Note SDRS is not supported for FlexClone.	<ul style="list-style-type: none"> ■ vSphere ■ KVM (RHEV) ■ SCVMM 	Add a Storage Reservation Policy to a Create, Clone, or FlexClone Blueprint.
Assign a storage reservation policy to a Linked Clone blueprint.	vSphere	Add a Storage Reservation Policy to a Linked Clone Blueprint.
Prevent provisioning from failing in the case of insufficient space on the datastores in a storage reservation policy.	<ul style="list-style-type: none"> ■ vSphere ■ KVM (RHEV) ■ SCVMM 	Allow Alternative Datastores at Provisioning Time.

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 Datastore to a Create, Clone, or FlexClone Blueprint

Tenant administrators and business group managers can assign a datastore to a volume so that only the specified datastore is used at provisioning time.

Prerequisites

- Log in to the vRealize Automation console as a **fabric administrator** or **business group manager**.
- Create a Clone, Create, or FlexClone blueprint for a vSphere, KVM (RHEV), or SCVMM environment.

Procedure

- 1 Select **Infrastructure > Blueprints > Blueprints**.
- 2 Point to the blueprint you want to edit.
- 3 Click the **Edit** icon ().
- 4 Click the **Build Information** tab.
- 5 Click **New Volume**.
- 6 Enter the capacity in the **Capacity (GB)** text box.

- 7 Enter a drive letter and mount path in the **Drive Letter/Mount Path** text box.
- 8 Enter a name in the **Label** text box.
- 9 Click the **Save** icon (✔).
- 10 (Optional) Add additional volumes.
- 11 Enter the maximum number of volumes a user can select for provisioned machines in the **Max # of volumes** text box.
Leave blank to accept the default of 15, or enter 0 to disable adding volumes.
- 12 Click **OK**.

Add a Datastore to a Linked Clone Blueprint

Tenant administrators and business group managers can assign a datastore to a volume so that only the specified datastore is used at provisioning time.

Prerequisites

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

Procedure

- 1 Select **Infrastructure > Blueprints > Blueprints**.
- 2 Point to the blueprint you want to edit.
- 3 Click the **Edit** icon (✎).
- 4 Click the **Properties** tab.
- 5 Click **New Property**.
- 6 Type `VirtualMachine.DiskN.Storage` in the **Name** text box.
N is the index (starting at 0) of the volume you want to assign.
- 7 Type the name of the datastore to assign to the volume in the **Value** text box.
This is the datastore name as it appears in the Storage Path column on the **Configuration** tab of the **Edit Compute Resource** page.
- 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 Click the **Save** icon (✔).

11 Click **OK**.

Add a Storage Reservation Policy to a Create, Clone, or FlexClone Blueprint

Tenant administrators and business group managers can assign storage reservation policies to blueprints to control and use storage space more efficiently.

Prerequisites

- Log in to the vRealize Automation console as a **fabric administrator** or **business group manager**.
- A fabric administrator must create a storage reservation policy.
- A Create, Clone, or FlexClone blueprint for a vSphere, KVM (RHEV), or SCVMM environment.

Procedure

- 1 Select **Infrastructure > Blueprints > Blueprints**.
- 2 Point to the blueprint to which you want to add a storage reservation policy and click **Edit**.
- 3 Click the **Build Information** tab.
- 4 Add **Storage volumes**.
 - a Click **New Volume**.
 - b Enter the capacity in the **Capacity (GB)** text box.
 - c (Optional) Enter a drive letter and mount path in the **Drive Letter/Mount Path** text box.
 - d (Optional) Enter a name in the **Label** text box.
 - e (Optional) Select a **Storage Reservation Policy** from the drop-down menu.

This option requires additional configuration by a fabric administrator to create storage reservation policies.
 - f Click the **Save** icon (✔).
- 5 (Optional) Select **Allow user to see and change storage reservation policies**.

If selected, users who are provisioning machines can see and change the storage reservation policies that you apply to blueprints.
- 6 Click the **Save** icon (✔).
- 7 Click **OK**.

Add a Storage Reservation Policy to a Linked Clone Blueprint

You must use custom properties to add a storage reservation policy to a linked clone blueprint.

Prerequisites

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

- A fabric administrator must create at least one storage reservation policy. See [Configure a Storage Reservation Policy](#).

Procedure

- 1 Select **Infrastructure > Blueprints > Blueprints**.
- 2 Point to the vSphere linked clone blueprint you want to edit.
- 3 Click the **Edit** icon ()
- 4 Click the **Properties** tab.
- 5 Click **New Property**.
- 6 Enter `VirtualMachine.DiskN.StorageReservationPolicy` in the **Name** text box.
N is the index, starting at 0, of the volume you want to assign to a storage reservation policy.
- 7 Enter the name of the storage reservation policy to assign to the volume 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 Click the **Save** icon ()
- 11 Click **New Property**.
- 12 Enter `VirtualMachine.DiskN.Size` in the **Name** text box.
N is the index you specified in step [Step 6](#)
- 13 Enter the amount of storage capacity (in gigabytes) to assign to the volume in the **Value** text box.
- 14 Click the **Save** icon ()
- 15 Click **OK**.

Allow Alternative Datastores at Provisioning Time

Tenant administrators and business group managers can use the custom property `VirtualMachine.DiskN.StorageReservationPolicyMode` to prevent provisioning from failing in the case of insufficient space on the datastores in a storage reservation policy. This custom property is set to allow vRealize Automation to choose a datastore outside the specified storage reservation policy in cases where there is not sufficient space remaining on the datastores in the policy.

Prerequisites

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

- Create a vSphere, KVM (RHEV), or SCVMM create, clone, or FlexClone blueprint that uses a storage reservation policy. See [Choosing a Storage Scenario](#).

Procedure

- 1 Select **Infrastructure > Blueprints > Blueprints**.
- 2 Point to the blueprint for which you want to allow alternative datastores and click **Edit**.
- 3 Click the **Properties** tab.
- 4 Click **New Property**.
- 5 Type `VirtualMachine.DiskN.StorageReservationPolicyMode` in the **Name** text box.
N is the index (starting at 0) of the volume.
- 6 Type the **Not Exact** in the **Value** text box.
Exact restricts provisioning to the datastores in the storage reservation policy. This is the default.
- 7 (Optional) Select the **Encrypted** check box to encrypt the custom property in the database.
- 8 (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.

- 9 Click the **Save** icon (✔).
- 10 Click **OK**.

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](#).

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](#).

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](#).
- 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 5-2. Custom Properties for Networking Configuration

Custom Property	Description
VirtualMachine.NetworkN.Address	Specifies the IP address of network device <i>N</i> in a machine provisioned with a static IP address.
VirtualMachine.NetworkN.MacAddressType	<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>
VirtualMachine.NetworkN.MacAddress	<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>
VirtualMachine.NetworkN.Name	<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 5-2. 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>

Table 5-2. Custom Properties for Networking Configuration (Continued)

Custom Property	Description
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 <hr/> <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>
VCNS.SecurityGroup.Names.name	<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"> ■ VCNS.SecurityGroup.Names ■ VCNS.SecurityGroup.Names.sales ■ VCNS.SecurityGroup.Names.support
VCNS.SecurityTag.Names.name	<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"> ■ VCNS.SecurityTag.Names ■ VCNS.SecurityTag.Names.sales ■ VCNS.SecurityTag.Names.support

Enable Users to Select Datacenter Locations

Tenant administrators and business group managers can create blueprints that prompt users to select datacenter locations, for example, London or Boston, when they request a machine.

Prerequisites

- Log in to the vRealize Automation console as a **fabric administrator** or **business group manager**.
- Your system administrator must add datacenter locations to a locations file. See *System Administration*.
- Your fabric administrator must associate compute resources to datacenter locations. See [Associate a Compute Resource with a Location](#).
- Create at least one cloud or virtual blueprint.

Procedure

- 1 Select **Infrastructure > Blueprints > Blueprints**.
- 2 Point to the blueprint to which you want to add datacenter location selection and click **Edit**.
- 3 Select **Display location on request**.
- 4 Click **OK**.

Business group users are now prompted to select a datacenter location when they request a machine to be provisioned from this blueprint.

Allow Requesters to Specify Machine Host Name

Tenant administrators and business group managers can use a custom property to prompt users to specify a host name for their 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**.
- 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 Click **New Property**.
- 5 Type **Hostname** in the **Name** text box.
- 6 Leave the **Value** text box blank.

- 7 Select the **Prompt user** check box to require the user to provide a value when they request a machine.

Because host names must be unique, users can only request one machine at a time from this blueprint.

- 8 Click the **Save** icon (✔).
- 9 Click **OK**.

Users who request a machine from your blueprint are required to specify a hostname for their machine. vRealize Automation validates that the specified host name is unique.

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
VbScript.PreProvisioning.Name	Runs the Visual Basic script before a machine is provisioned.
VbScript.PostProvisioning.Name	Runs the Visual Basic script after a machine is provisioned.
VbScript.UnProvisioning.Name	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**.

Enable Machine Users to Create Snapshots

Tenant administrators and business group managers can enable Snapshot Manager in machines created from a vSphere blueprint. This allows users to capture snapshots of their virtual machines and keep space-efficient copies of the original VM image.

Prerequisites

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

Procedure

- 1 Select **Infrastructure > Blueprints > Blueprints**.
- 2 Point to the blueprint you want to change and click **Edit**.
- 3 Click the **Actions** tab.
- 4 Specify whether to allow users to use Snapshot Manager to create, apply, and delete snapshots of their machine.
- 5 Configure the age limit and number of snapshots you allow users to take and apply.
 - a Click the **Properties** tab.
 - b Click **New Property**.
 - c Configure the custom properties.

Option	Description
<code>Snapshot.Policy.Limit</code>	Sets the number of snapshots allowed per machine. The default setting is one snapshot per machine. This property applies to vSphere provisioning. When set to 0, the blueprint option to create a snapshot is hidden for all users except for support and manager roles.
<code>Snapshot.Policy.AgeLimit</code>	Sets the age limit, in days, for snapshots that can be applied to machines. This property applies to vSphere provisioning.

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

6 Click **OK**.

RDP Connections

Tenant administrators and business group managers can configure RDP settings and enable connect by using RDP for machines provisioned with their blueprints.

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) <code>RDP.File.Name</code>	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) <code>VirtualMachine.Rdp.SettingN</code>	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.
<code>VirtualMachine.Admin.NameCompletion</code>	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.](#)

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.](#)

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*.

Infrastructure Organizer

IaaS administrators and fabric administrators can use the Infrastructure Organizer to quickly organize multiple compute resources into fabric groups, assign cost profiles, and bring existing virtual machines under vRealize Automation management.

This chapter includes the following topics:

- [Creating Blueprints for Importing Virtual Machines](#)
- [Choosing an Infrastructure Organizer Scenario](#)
- [Organize Compute Resources and Import Existing Virtual Machines](#)
- [Import Existing Virtual Machines](#)

Creating Blueprints for Importing Virtual Machines

A tenant administrator or business group manager must create a blueprint that matches the machines you want to import using the Infrastructure Organizer.

Make sure that the blueprint matches as closely as possible the machines you are importing. . For example, if you are importing vSphere machines that were created by using Linux Kickstart, create a vSphere blueprint for Linux Kickstart provisioning for importing those machines.

Consider which business groups need access to these imported machines, whether you want them to be reconfigured or reprovisioned, and how long to lease or archive the machines. These settings and more are all configured on the blueprint you use to import the machines.

Choosing an Infrastructure Organizer Scenario

Fabric administrators and IaaS administrators can use the Infrastructure Organizer to organize compute resources into fabric groups, assign cost profiles to compute resources, import existing virtual machines, and associate imported machines with business groups.

Depending on your role or combination of roles, you can leverage different capabilities of the Infrastructure Organizer.

Table 6-1. Choosing an Infrastructure Organizer Scenario

Scenario	Role	Procedure
Assign cost profiles to your compute resources and organize them into fabric groups. If you have business group manager privileges, you can continue with the Infrastructure Organizer Wizard to import existing virtual machines.	IaaS Administrator	IaaS administrators do not have access to the Infrastructure Organizer. However, you can create a fabric group and nominate yourself as the fabric administrator so that you can leverage the organizer to quickly place your compute resources into fabric groups. See Organize Compute Resources and Import Existing Virtual Machines .
Assign cost profiles to your compute resources and import existing virtual machines.	Fabric Administrator and Business Group Manager	Import Existing Virtual Machines .

Organize Compute Resources and Import Existing Virtual Machines

IaaS administrators can leverage the Infrastructure Organizer to quickly assign cost profiles to compute resources and organize them into fabric groups. Optionally, you can continue with the Infrastructure Organizer Wizard to import existing virtual machines and bring them under vRealize Automation management.

Prerequisites

- Log in to the vRealize Automation console as an **IaaS administrator**.
- Follow the appropriate procedures to store credentials and create endpoints.
- Create the fabric groups to which you want to assign your compute resources.
- (Optional) If you intend to assign cost profiles to your compute resources, a fabric administrators must create the cost profiles.
- (Optional) If you intend to import existing virtual machines, a tenant administrator or business group manager must create a blueprint to apply to machines you import.

Procedure

1 [Create the Infrastructure Organizer Fabric Group](#)

Create a fabric group that contains all of the virtual compute resources you want to organize and assign yourself as the fabric administrator.

2 [Configure Compute Resources](#)

IaaS administrators with fabric administrator privileges can leverage the Infrastructure Organizer to quickly organize multiple compute resources into fabric groups.

3 [Bring Virtual Machines under vRealize Automation Management](#)

You can import existing virtual machines by assigning them to business groups, reservations, blueprints, and owners.

Create the Infrastructure Organizer Fabric Group

Create a fabric group that contains all of the virtual compute resources you want to organize and assign yourself as the fabric administrator.

Prerequisites

- Log in to the vRealize Automation console as an **IaaS administrator**.
- Follow the appropriate procedures to store credentials and create endpoints.
- Create the fabric groups to which you want to assign your compute resources.
- (Optional) If you intend to assign cost profiles to your compute resources, a fabric administrators must create the cost profiles.
- (Optional) If you intend to import existing virtual machines, a tenant administrator or business group manager must create a blueprint to apply to machines you import.

Procedure

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

As a fabric administrator, you now have access to the Infrastructure Organizer and can also create the cost profiles to assign to your compute resources.

- 6 Select all of the virtual compute resources that you want to organize from the **Compute resource** list.
- 7 Click **OK**.
- 8 Refresh your browser.

Configure Compute Resources

IaaS administrators with fabric administrator privileges can leverage the Infrastructure Organizer to quickly organize multiple compute resources into fabric groups.

Note You can edit multiple elements at the same time by clicking the pin icon (📌) and then editing one of your selections. The changes you make are applied to all of the selections you pinned.

Prerequisites

[Create the Infrastructure Organizer Fabric Group.](#)

Procedure

- 1 Select **Infrastructure > Infrastructure Organizer > Infrastructure Organizer**.

- 2 Click **Next**.
- 3 Select the compute resources you want to organize into fabric groups.
- 4 Click **Next**.
- 5 Click the **Edit** icon ().
- 6 Select the fabric groups to which you want to add your compute resource from the **Fabric Groups** list.
- 7 (Optional) Select a cost profile to assign to your compute resource from the **Cost Profile** drop-down menu.
- 8 Click the **Save** icon ().
- 9 Click **Next**.

Your compute resources are organized into fabric groups. You can stop here by clicking **Next** and then **Finish**, or if you have **Business Group Manager** privileges, you can continue with the Infrastructure Organizer to import existing virtual machines from your compute resources.

Bring Virtual Machines under vRealize Automation Management

You can import existing virtual machines by assigning them to business groups, reservations, blueprints, and owners.

Note You can edit multiple elements at the same time by clicking the pin icon () and then editing one of your selections. The changes you make are applied to all of the selections you pinned.

Prerequisites

- [Configure Compute Resources](#).
- Verify that the machines you are importing have unique names. Duplicate machine names are not supported in vRealize Automation.

Procedure

- 1 Locate the machine you want to import.
- 2 Click the **Edit** icon ().
- 3 Select a business group to assign to the imported machine from the **Business Group** drop-down menu.
- 4 Click the **Save** icon ().
- 5 Click **Next**.
- 6 Click the **Edit** icon ().
- 7 Select a blueprint to apply to your existing virtual machine from the **Blueprint** drop-down menu.

8 Assign the machine you are importing to a reservation.

- Select a reservation from the **Reservation** drop-down menu.
- Type a name in the **Reservation** text box to create a new reservation specifically for the imported machines.

9 (Optional) Select **Grow Allocations**.

The memory and storage allocated to the reservation you assigned to the imported machines is increased by exactly the amount the imported machines consume.

10 (Optional) Select **Increase Quota**.

The machine quota on the reservation you assigned to the imported machines is increased by exactly the number of virtual machines you import.

11 Specify a machine owner.

- Select a machine owner from the **Owner** drop-down menu. Only business group members who already own machines in vRealize Automation appear in the drop-down menu.
- Type a user name into the **Owner** text box and press Enter. The user must belong to the business group to which you assigned the imported machines.

12 Click the **Save** icon (✔).

13 Click **Next**.

14 Click **Finish**.

Import Existing Virtual Machines

Fabric administrators can import virtual machines that currently exist on their compute resources outside of vRealize Automation.

Prerequisites

- Log in to the vRealize Automation console as a **fabric administrator** and as a **business group manager**.
- Create the business groups to which you want to assign imported machines.
- A tenant administrator or business group manager must create a blueprint to apply to machines you import.
- (Optional) If you intend to assign cost profiles to your compute resources, create the cost profiles you want to assign to your compute resources.

Procedure

1 [Choose Compute Resources](#)

Fabric administrators select compute resources from which they want to import existing virtual machines and assign cost profiles to the compute resources that they manage.

2 Bring Virtual Machines under vRealize Automation Management

A fabric administrator imports existing virtual machines and assigns them to business groups, reservations, and users.

Choose Compute Resources

Fabric administrators select compute resources from which they want to import existing virtual machines and assign cost profiles to the compute resources that they manage.

Note You can edit multiple elements at the same time by clicking the pin icon and then editing one of your selections. The changes you make are applied to all of the selections you pinned.

Prerequisites

- Log in to the vRealize Automation console as a **fabric administrator** and as a **business group manager**.
- Create the business groups to which you want to assign imported machines.
- A tenant administrator or business group manager must create a blueprint to apply to machines you import.
- (Optional) If you intend to assign cost profiles to your compute resources, create the cost profiles you want to assign to your compute resources.

Procedure

- 1 Select **Infrastructure > Infrastructure Organizer > Infrastructure Organizer**.
- 2 Click **Next**.
- 3 Select the compute resources that contain the existing virtual machines you want to import.
Only compute resources that you manage as a fabric administrator appear.
- 4 Click **Next**.
- 5 (Optional) Assign cost profiles to your compute resources.
 - a Click the **Edit** icon ().
 - b Select a cost profile to assign to your compute resource from the **Cost Profile** drop-down menu.
 - c Click the **Save** icon (.
- 6 Click **Next**.

Bring Virtual Machines under vRealize Automation Management

A fabric administrator imports existing virtual machines and assigns them to business groups, reservations, and users.

Note You can edit multiple elements at the same time by clicking the pin icon and then editing one of your selections. The changes you make are applied to all of the selections you pinned.

Prerequisites

- [Choose Compute Resources](#).
- Verify that the machines you are importing have unique names. Duplicate machine names are not supported in vRealize Automation.

Procedure

- 1 Locate the machines you want to import.
- 2 Click the **Edit** icon ()
- 3 Select a business group to assign to the imported machine from the **Business Group** drop-down menu.
- 4 Click the **Save** icon ()
- 5 Click **Next**.
- 6 Click the **Edit** icon ()
- 7 Select a blueprint to apply to your existing virtual machine from the **Blueprint** drop-down menu.
- 8 Assign the machines you are importing to a reservation.
 - Select a reservation from the **Reservation** drop-down menu.
 - Type a name in the **Reservation** text box to create a new reservation specifically for the imported machines.
- 9 (Optional) Select **Grow Allocations**.

The memory and storage allocated to the reservation you assigned to the imported machines is increased by exactly the amount the imported machines consume.
- 10 (Optional) Select **Increase Quota**.

The machine quota on the reservation you assigned to the imported machines is increased by exactly the number of virtual machines you import.

11 Specify a machine owner.

- Select a machine owner from the **Owner** drop-down menu. Only business group members who already own machines in vRealize Automation appear in the drop-down menu.
- Type a user name into the **Owner** text box and press Enter. The user must belong to the business group to which you assigned the imported machines.

12 Click the **Save** icon (.

13 Click **Next**.

14 Click **Finish**.

Managing Infrastructure

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

This chapter includes the following topics:

- [Managing Resource Usage](#)
- [Data Collection](#)
- [vSphere Host Maintenance](#)
- [Decommissioning a Storage Path](#)

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 7-1. 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.

Table 7-1. Resource Usage Terminology (Continued)

Term	Description
Allocated	Indicates the machine quota, memory, or storage resources actively being consumed by provisioned machines. For example, consider a reservation with a machine quota of 10. If there are 15 provisioned machines on it, but only 6 of them are currently powered on, the machine quota is 60 percent allocated.
Used	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 7-2. 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

Table 7-2. Choose a Resource Monitoring Scenario (Continued)

Resource Monitoring Scenario	Privileges Required	Location
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*.

Reducing Reservation Usage by Attrition

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

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

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

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

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 7-3. 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**.

Update Cost Data for All Compute Resources

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

Prerequisites

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

Procedure

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

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

vSphere Host Maintenance

Before you place your host in maintenance mode to service it, for example, to install more memory, migrate any machines managed by vRealize Automation to another host.

Any machines managed by vRealize Automation that remain on the vSphere host when you place it in maintenance mode permanently display as missing from vRealize Automation.

Decommissioning a Storage Path

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

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

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

Disable a Storage Path

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

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

Prerequisites

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

Procedure

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

- 8 Click **OK**.
- 9 Repeat this procedure for all reservations that use the storage path you are decommissioning.

Monitoring Workflows and Viewing Logs



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

Table 8-1. 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.	laaS 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.	laaS 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.	laaS administrator	Display default log information or control display content using column and filter options. Select Infrastructure > Monitoring > Log .
View the status and history of executed Distributed Execution Manager and other workflows.	laaS administrator	Display workflow history and optionally open a specific workflow to display its execution details. Select Infrastructure > Monitoring > Workflow History .
Display a list of events, including event type, time, user ID, and so on, and optionally display an event details page.	System administrator	View a list of events and their associated attributes, such as run time, event description, tenant name, target type and ID, and other characteristics. Select Administration > Event Logs .
Monitor the status of your requests and view request details.	Tenant administrator or business group manager	Display the status of requests that you are responsible for or own. Click Requests .

Appendix A: Machine Life Cycle and Workflow States

9

Depending on the method of provisioning specified in the machine blueprint, machines are provisioned using different vRealize Automation workflows.

The master workflow applies to all machines regardless of provisioning method.

Table 9-1. 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.

Table 9-1. MasterWorkflow States (Continued)

Workflow State	Work Item	Description
UnprovisionMachine		The machine unprovisioning process has started.
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 9-2. BasicVMWorkflow States

Workflow State	Work Item	Description
CreatingMachine	CreateVM	The machine is created by the hypervisor.
AddingDisks	AddDisks	Additional disks, if any, are being created and added to the machine by the hypervisor.
BuildComplete		The machine build process completed.

Table 9-3. CloneWorkflow States

Workflow State	Work Item	Description
CloneMachine	CloneVM	The hypervisor is cloning the machine.
CustomizeMachine	CustomizeMachine	The hypervisor is configuring the machine.
InitialPowerOn	InitialPowerOn	The machine is being powered on for the first time after the clone process is finished.
CustomizeOS	CustomizeOS	A guest agent is configuring the machine. Custom scripts are running.
BuildComplete		The machine build process is finished.

Table 9-4. WIMImageWorkflow States

Workflow State	Work Item	Description
CreatingMachine	CreateVM	The hypervisor is creating the machine.
AddingDisks	AddDisks	The hypervisor is creating and adding additional disks, if any, to the machine.
InitialPowerOn	PowerOn	The machine is being powered on for the first time.
SetupOS	SetupOS	The machine is booted from the CD-ROM or ISO image and the operating system WIM image is being copied to the C:\ drive from the configure location.
EjectingCD	EjectCD	The CD-ROM or ISO image is being unmounted from the machine.

Table 9-4. WIMImageWorkflow States (Continued)

Workflow State	Work Item	Description
InstallOS	InstallOS	The machine is being booted for the first time from the hard disk. The guest agent is performing system preparation and customization. Additional disks are partitioned and formatted and the owner is added to the administrator group.
BuildComplete		The machine build process is finished.

Table 9-5. LinuxKickstartWorkflow States

Workflow State	Work Item	Description
CreatingMachine	CreateVM	The machine is being created by the hypervisor.
AddingDisks	AddDisks	Additional disks, if any, are being created and added to the machine by the hypervisor.
InitialPowerOn	PowerOn	The machine is being powered on for the first time.
InstallingOS	SetupOS	The machine is booted from CD-ROM and the operating system image is being created by the kickstart process.
EjectingCD	EjectCD	The CD-ROM is being unmounted from the machine.
CustomizeOS	CustomizeOS	The machine is being configured by a guest agent. Custom scripts are being executed.
BuildComplete		The machine build process completed.

Table 9-6. VirtualSccmProvisionigWorkflow States

Workflow State	Work Item	Description
CreatingMachine	CreateVM	The machine is being created by the hypervisor.
AddingDisks	AddDisks	Additional disks, if any, are being created and added to the machine by the hypervisor.
SccmRegistration	SccmRegisterMachine	The machine is being registered in SCCM provisioning infrastructure.
InitialPowerOn	PowerOn	The machine is being powered on for the first time.
InstallingOS	Complete	The machine is booted and is being provisioned by SCCM.
EjectingCD	EjectCD	The CD-ROM is being unmounted from the machine.

Table 9-6. VirtualSccmProvisionigWorkflow States (Continued)

Workflow State	Work Item	Description
FailedProvisioning	SccmUnregisterMachine	The machine build failed. The machine is being unregistered from the SCCM provisioning infrastructure.
EndedVirtualProvisioning		The machine build process completed. The provisioning workflow is about to stop.