

Foundations and Concepts

12 April 2018

vRealize Automation 7.4



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 © 2008–2018 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

Foundations and Concepts	4
Foundations and Concepts	5
Using Scenarios	5
Using the Goal Navigator	5
Introducing vRealize Automation	6
Tenancy and User Roles	12
Service Catalog	24
Infrastructure as a Service	27
XaaS Blueprints and Resource Actions	37
Common Components	39
Life Cycle Extensibility	41

Foundations and Concepts

VMware vRealize™ Automation provides a secure portal where authorized administrators, developers, or business users can request new IT services. In addition, they can manage specific cloud and IT resources that enable IT organizations to deliver services that can be configured to their lines of business in a self-service catalog.

This documentation describes the features and capabilities of vRealize Automation. It includes information about the following subjects:

- vRealize Automation components
- Common service catalog
- Infrastructure as a Service
- XaaS
- Software

For information about cost management for VMware vRealize™ Automation, see the documentation for VMware vRealize™ Business™ for Cloud.

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 anyone who needs to familiarize themselves with the features and capabilities of vRealize Automation.

VMware Technical Publications Glossary

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

Foundations and Concepts

Before you begin working with vRealize Automation, you can familiarize yourself with basic vRealize Automation concepts.

Using Scenarios

You can use scenarios to build working samples of vRealize Automation functionality that you can learn from or customize to suit your needs.

Scenarios walk you through the most common and simplified workflow to complete a vRealize Automation task. They do not contain options or choices, and serve as introductory examples to both basic and advanced vRealize Automation functionality.

For example, you can use *Installing and Configuring vRealize Automation for the Rainpole Scenario* to install a working proof of concept vRealize Automation deployment into your existing vSphere environment.

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.

Introducing vRealize Automation

IT organizations can use VMware vRealize™ Automation to deliver services to their lines of business.

vRealize Automation provides a secure portal where authorized administrators, developers, or business users can request new IT services and manage specific cloud and IT resources, while ensuring compliance with business policies. Requests for IT services, including infrastructure, applications, desktops, and many others, are processed through a common service catalog to provide a consistent user experience.

To improve cost control, you can integrate vRealize Business for Cloud with your vRealize Automation instance to expose the month-to-date expense of cloud and virtual machine resources, and help you better manage capacity, price, and efficiency.

Note Beginning with version 7.3, vRealize Automation supports only vRealize Business for Cloud version 7.3 and later.

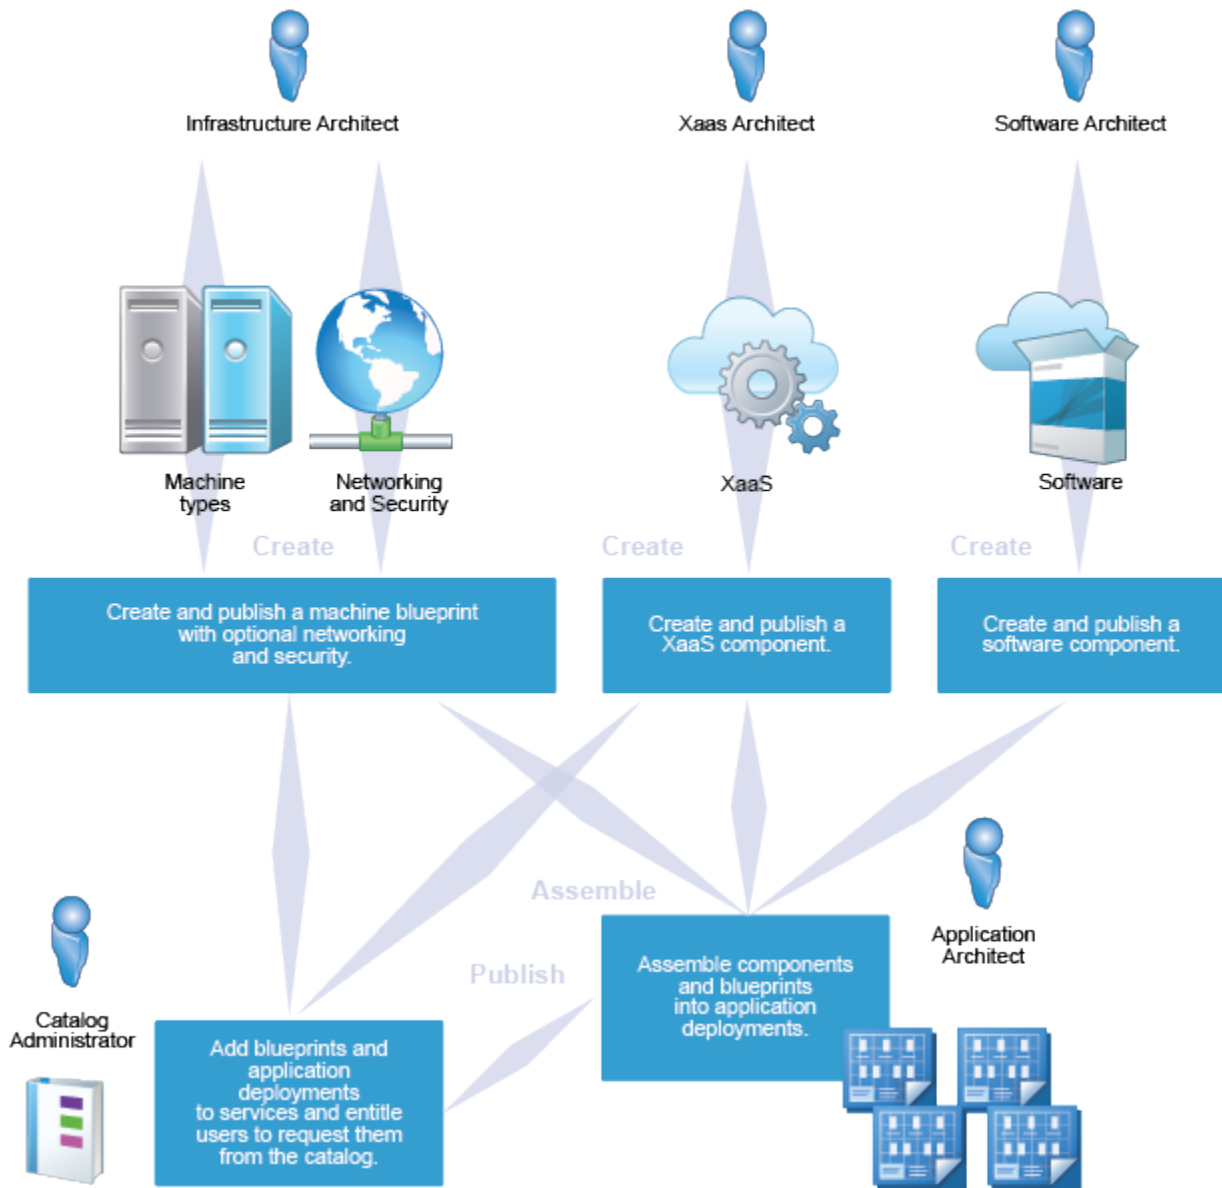
Providing On-Demand Services to Users Overview

You can use the IaaS, Software, and XaaS features of vRealize Automation to model custom on-demand IT services and deliver them to your users through the vRealize Automation common service catalog.

You use blueprints to define machine deployment settings. Published blueprints become catalog items, and are the means by which entitled users provision machine deployments. Catalog items can range in complexity from a single, simple machine with no guest operating system to complex custom application stacks delivered on multiple machines under an NSX load balancer with networking and security controls.

You can create and publish blueprints for a single machine deployment, or a single custom XaaS resource, but you can also combine machine blueprints and XaaS blueprints with other building blocks to design elaborate application blueprints that include multiple machines, networking and security, software with full life cycle support, and custom XaaS functionality. You can also control deployment settings by using a parameterized blueprint, which allows you to specify pre-configured size and image settings at request time. Because all published blueprints and blueprint components are reusable, you can create a library of these components and combine them in new nested blueprints to deliver increasingly complex on-demand services.

Published blueprints become catalog items that your service catalog administrators can deliver to your users. The service catalog provides a unified self-service portal for consuming IT services. Service catalog administrators can manage user access to catalog services, items, and actions by using entitlements and approvals, and users can browse the catalog to request items they need, track their requests, and manage their provisioned items.



- [Infrastructure as a Service Overview](#)

With Infrastructure as a Service (IaaS), you can rapidly model and provision servers and desktops across virtual and physical, private and public, or hybrid cloud infrastructures.

- [Software Components Overview](#)

Software components automate the installation, configuration, and life cycle management of middleware and application deployments in dynamic cloud environments. Applications can range from simple Web applications to complex and even packaged applications.

- [XaaS Overview](#)

With the XaaS, XaaS architects can create XaaS blueprints and resource action, and publish them as catalog items.

- [Service Catalog Overview](#)

The service catalog provides a unified self-service portal for consuming IT services. Users can browse the catalog to request items they need, track their requests, and manage their provisioned items.

- [Containers Overview](#)

You can use containers to gain access to additional instrumentation for developing and deploying applications in vRealize Automation.

Infrastructure as a Service Overview

With Infrastructure as a Service (IaaS), you can rapidly model and provision servers and desktops across virtual and physical, private and public, or hybrid cloud infrastructures.

Modeling is accomplished by creating a machine blueprint, which is a specification for a machine. Blueprints are published as catalog items in the common service catalog, and are available for reuse as components inside of application blueprints. When an entitled user requests a machine based on one of these blueprints, IaaS provisions the machine.

With IaaS, you can manage the machine life cycle from a user request and administrative approval through decommissioning and resource reclamation. Built-in configuration and extensibility features also make IaaS a highly flexible means of customizing machine configurations and integrating machine provisioning and management with other enterprise-critical systems such as load balancers, configuration management databases (CMDBs), ticketing systems, IP address management systems, or Domain Name System (DNS) servers.

Software Components Overview

Software components automate the installation, configuration, and life cycle management of middleware and application deployments in dynamic cloud environments. Applications can range from simple Web applications to complex and even packaged applications.

By using a configurable scriptable engine, software architects fully control how middleware and application deployment components are installed, configured, updated, and uninstalled on machines. Through the use of Software properties, software architects can require or allow blueprint architects and end-users to specify configuration elements such as environment variables. For repeated deployments, these blueprints standardize the structure of the application, including machine blueprints, software components, dependencies, and configurations, but can allow environment variables and property binding to be reconfigured if necessary.

To successfully add software components to the design canvas, you must also have business group member, business group administrator, or tenant administrator role access to the target catalog.

Deploying Any Application and Middleware Service

You can deploy Software components on Windows or Linux operating systems on vSphere, vCloud Director, vCloud Air, and Amazon AWS machines.

- IaaS architects create reusable machine blueprints based on templates, snapshots, or Amazon machine images that contain the guest agent and Software bootstrap agent to support Software components.
- Software architects create reusable software components that specify exactly how the software is installed, configured, updated during deployment scale operations, and uninstalled on machines.
- Software architects, IaaS architects, and application architects use a graphical interface to model application deployment topologies. Architects reconfigure Software properties and bindings as required by the software architect, and publish application blueprints that combine Software components and machine blueprints.
- Catalog administrators add the published blueprints to a catalog service, and entitle users to request the catalog item.
- Entitled users request the catalog item and provide any configuration values designed to be editable. vRealize Automation deploys the requested application, provisioning any machine(s), networking and security components, and Software component(s) defined in the application blueprint.
- Entitled users request the scale in or scale out actions to adjust their deployments to changing workload demands. vRealize Automation installs or uninstalls Software components on machines for scale, and runs update scripts for dependent Software components.

Standardization in Software

With Software, you can create reusable services using standardized configuration properties to meet strict requirements for IT compliance. Software includes the following standardized configuration properties:

- Model-driven architecture that enables adding IT certified machine blueprints and middleware services within the application blueprint.
- A delegation model for overriding configuration name value pairs between software architect, application architect, and end user to standardize configuration values for application and middleware service.

Software Extensibility and Open Architecture

You can download predefined Software components for a variety of middleware services and applications from the VMware Solution Exchange. Using either the vRealize CloudClient or vRealize Automation REST API, you can programmatically import predefined Software components into your vRealize Automation instance.

- To visit the VMware Solution Exchange, see https://solutionexchange.vmware.com/store/category_groups/cloud-management.
- For information about vRealize Automation REST API, see *Programming Guide* and *vRealize Automation API Reference*.

- For information about vRealize CloudClient, see <https://developercenter.vmware.com/tool/cloudclient>.

XaaS Overview

With the XaaS, XaaS architects can create XaaS blueprints and resource action, and publish them as catalog items.

With XaaS, you can provide anything as a service using the capabilities of VMware vRealize™ Orchestrator™. For example, you can create a blueprint that allows a user to request a backup of a database. After completing and submitting a backup request, the user receives a backup file of the database they specified.

An XaaS architect can create custom resource types mapped to vRealize Orchestrator object types and define them as items to be provisioned. A XaaS architect can then create blueprints from vRealize Orchestrator workflows and publish the blueprints as catalog items. The vRealize Orchestrator workflows can be either predefined or independently developed by workflow developers.

You can also use the XaaS to design additional actions that the consumer can perform on the provisioned items. These additional actions are connected to vRealize Orchestrator workflows and take the provisioned item as input to the workflow. To use this function for items provisioned by sources other than the XaaS, you must create resource mappings to define their resource types in vRealize Orchestrator.

For more information about vRealize Orchestrator and its capabilities, see the vRealize Orchestrator documentation.

Service Catalog Overview

The service catalog provides a unified self-service portal for consuming IT services. Users can browse the catalog to request items they need, track their requests, and manage their provisioned items.

Service architects and administrators can define new services and publish them to the common catalog. When defining a service, the architect can specify the kind of item that can be requested, and what options are available to the consumer as part of submitting the request.

Group managers or line-of-business administrators can specify business policies such as who is entitled to request specific catalog items or perform specific actions on provisioned items. They can also apply configurable approval policies to catalog requests.

Users responsible for managing the catalog, such as tenant administrators and service architects, can manage the presentation of catalog items to the consumers of IT services, for example by grouping items into service categories for easier navigation and highlighting new services to consumers on the portal home page.

Containers Overview

You can use containers to gain access to additional instrumentation for developing and deploying applications in vRealize Automation.

Containers for vRealize Automation allows vRealize Automation to support containers. You can provision an application that is built from containers or from a combination of containers and VMs.

Container administrators can use Containers to perform the following tasks:

- Model containerized applications in vRealize Automation blueprints.
- Provision container hosts from the vRealize Automation service catalog.
- Manage container hosts from within vRealize Automation.
- Create and configure hosts.
- Set resource quotas for containers.
- Work with templates, images, and registries.
- Create and edit blueprints in the vRealize Automation service catalog.
- Develop multi-container templates.

Container architects can add container components to a vRealize Automation blueprint.

The integrated Containers application uses the Docker Remote API to provision and manage containers, including retrieving information about container instances. From a deployment perspective, developers can use Docker Compose to create their application and deploy it through Containers in vRealize Automation. Because that application is ready to be promoted from development to production, developers can enhance the application to include dynamic networks or micro-segmentation.

Cloud administrators can manage the container host infrastructure, for example to govern capacity quotas and approval workflows.

Use the Containers Context-Sensitive Help

When working with Containers for vRealize Automation, you have access to a context-sensitive help system that dynamically displays content for the task that you are currently performing.

After you open the Containers help system, the page content automatically updates based on your location in the Containers user interface. You can view the Containers help system in a separate window, on a second screen, or from a mobile device in parallel with the primary interface.

You can use the Containers help system outside of the trusted network and still receive instant documentation page updates relative to where your cursor is in the Containers application.

- 1 Log in to vRealize Automation as a **container administrator**.
- 2 Click the **Containers** tab.
- 3 Click **Help** on the Containers Welcome page, next to the **Add a Host** button.

You can refresh the web browser to redisplay the Welcome page.

vRealize Business for Cloud Overview

With vRealize Business for Cloud, directors of cloud operations can monitor their expenditures and design more price-efficient cloud services.

vRealize Business for Cloud provides the following benefits:

- Drives accountability by providing visibility into the price of virtual infrastructure and public cloud providers and providing daily price and month-to-date expense updates in vRealize Automation.

- Promotes efficiencies in the virtual infrastructure by making it possible to compare the prices, efficiency, and availability of their private cloud with public cloud providers and industry benchmark data.
- Optimizes decisions about placement for virtual workloads and tradeoffs between buying new hardware and using public cloud providers.

For more information about vRealize Business for Cloud, see the [vRealize Business for Cloud documentation](#).

Tenancy and User Roles

vRealize Automation supports multiple tenants in the same installation. Users always log in and perform their tasks in a specific tenant. Some administrator roles can manage configuration that affects multiple tenants.

Tenancy Overview

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

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

Table 1. Tenant Configuration

Configuration Area	Description
Login URL	Each tenant has a unique URL to the vRealize Automation console. <ul style="list-style-type: none"> ■ The default tenant URL is in the following format: <code>https://hostname/vcac</code> ■ The URL for additional tenants is in the following format: <code>https://hostname/vcac/org/tenantURL</code>
Identity stores	Each tenant requires access to one or more directory services, such as OpenLDAP or Microsoft Active Directory servers, that are configured to authenticate users. You can use the same directory service for more than one tenant, but you must configure it separately for each tenant.
Branding	A tenant administrator can configure the branding of the vRealize Automation console including the logo, background color, and information in the header and footer. System administrators control the default branding for all tenants.
Notification providers	System administrators can configure global email servers that process email notifications. Tenant administrators can override the system default servers, or add their own servers if no global servers are specified.
Business policies	Administrators in each tenant can configure business policies such as approval workflows and entitlements. Business policies are always specific to a tenant.

Table 1. Tenant Configuration (Continued)

Configuration Area	Description
Service catalog offerings	Service architects can create and publish catalog items to the service catalog and assign them to service categories. Services and catalog items are always specific to a tenant.
Infrastructure resources	The underlying infrastructure fabric resources, for example, vCenter servers, Amazon AWS accounts, or Cisco UCS pools, are shared among all tenants. For each infrastructure source that vRealize Automation manages, a portion of its compute resources can be reserved for users in a specific tenant to use.

About the Default Tenant

When the system administrator configures an Active Directory link using Directories management during the installation of vRealize Automation, a default tenant is created with the built-in system administrator account to log in to the vRealize Automation console. The system administrator can then configure the default tenant and create additional tenants.

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

User and Group Management

All user authentication is handled by Active Directory links that are configured through Directories Management. Each tenant has one or more Active Directory links that provide authentication on a user or group level.

The root system administrator performs the initial configuration of single sign-on and basic tenant creation and setup, including designating at least one tenant administrator for each tenant. Thereafter, a tenant administrator can configure Active Directory links and assign roles to users or groups as needed from within their designated tenant.

Tenant administrators can also create custom groups within their own tenants and add users and groups to those groups. Custom groups can be assigned roles or designated as the approvers in an approval policy.

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

Comparison of Single-Tenant and Multi-tenant Deployments

vRealize Automation supports deployments with either a single tenant or multiple tenants. The configuration can vary depending on how many tenants are in your deployment. Many NSX and vSphere-related blueprint selections are tenant-specific.

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

Infrastructure configuration, including the infrastructure sources that are available for provisioning, can be configured in any tenant and is shared among all tenants. You divide your infrastructure resources, such as cloud or virtual compute resources, into fabric groups and assign an administrator to manage those resources as the fabric administrator. Fabric administrators can allocate resources in their fabric group to business groups by creating reservations.

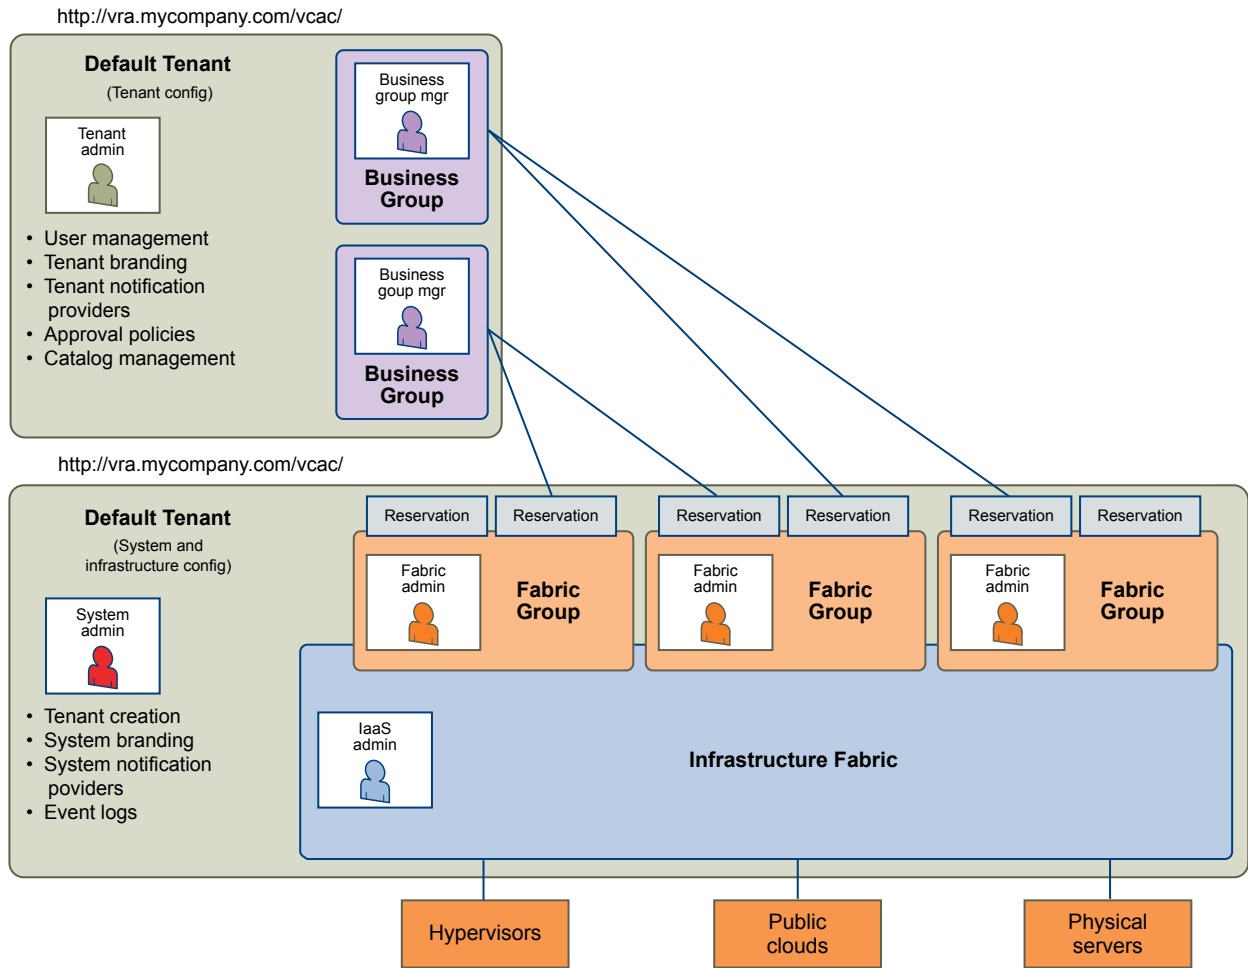
To support tenant allocation of vSphere and NSX endpoint resources, only the network profiles, reservation policies, storage policies, security groups and tags, and transport zones that are applicable to the current tenant are visible when authoring blueprints.

Single-Tenant Deployment

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

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

Figure 1. Single-Tenant Example



Note In a single-tenant scenario, it is common for the system administrator and tenant administrator roles to be assigned to the same person, but two distinct accounts exist. The system administrator account is always `administrator@vsphere.local`, and the system administrator account creates a local user account to assign the tenant administrator role.

Multi-tenant Deployment

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

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

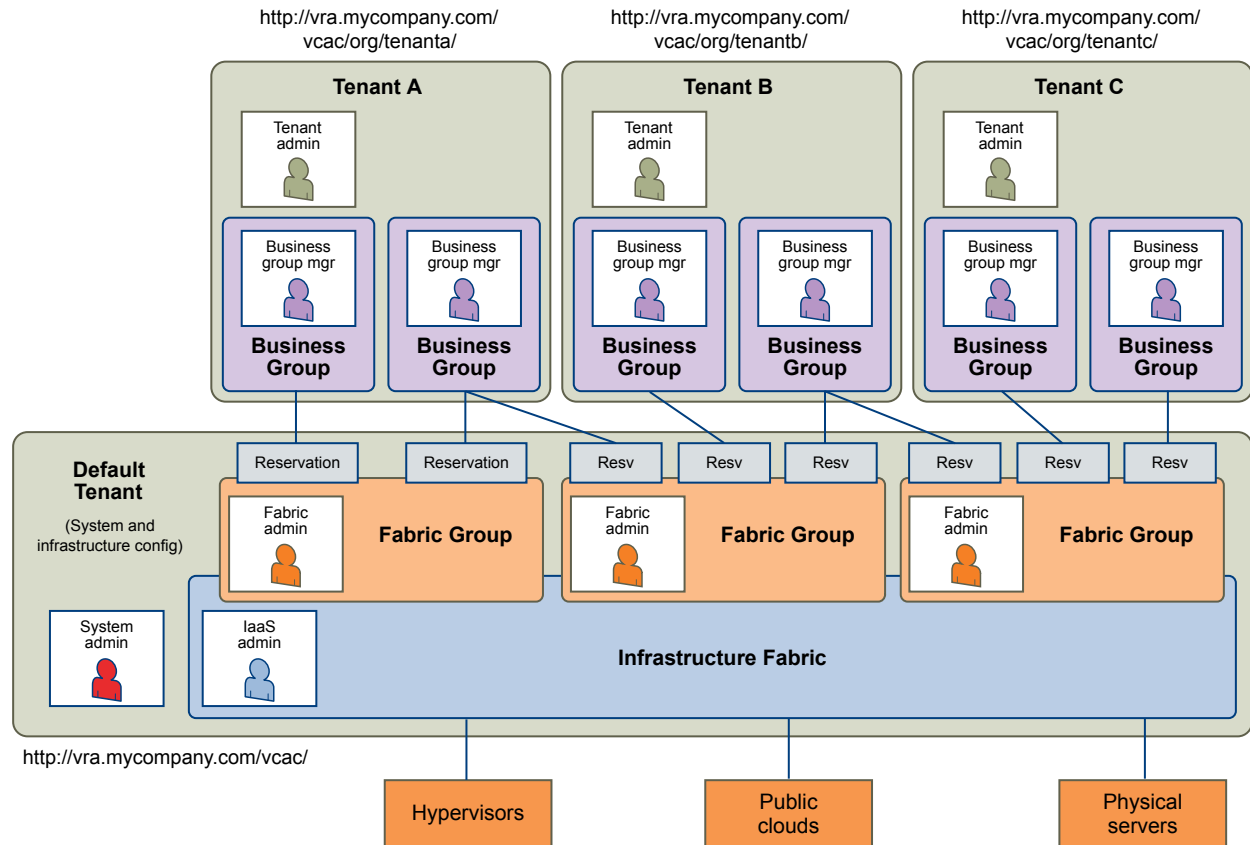
Table 2. Multi-tenant Deployment Examples

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

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

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

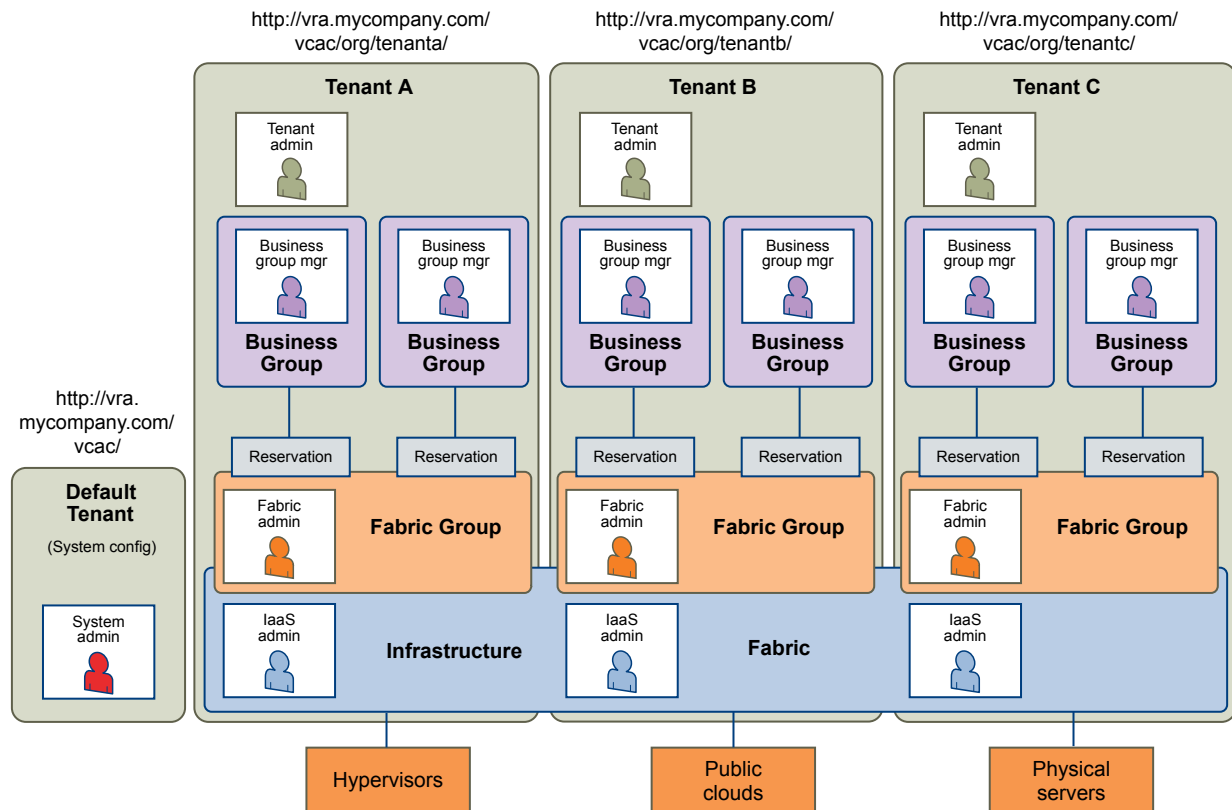
Figure 2. Multi-tenant Example with Infrastructure Configuration Only in Default Tenant



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

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

Figure 3. Multi-tenant Example with Infrastructure Configuration in Each Tenant



User Roles Overview

Roles consist of a set of privileges that can be associated with users to determine what tasks they can perform. Based on their responsibilities, individuals might have one or more roles associated with their user account.

All user roles are assigned within the context of a specific tenant. However, some roles in the default tenant can manage system-wide configuration that applies to multiple tenants.

System-Wide Role Overview

System-wide roles are typically assigned to an IT system administrator. In some organizations, the IaaS administrator role might be the responsibility of a cloud administrator.

System Administrator

The system administrator is typically the person who installs vRealize Automation and is responsible for ensuring its availability for other users. The system administrator creates tenants and manages system-wide configuration such as system defaults for branding and notification providers. This role is also responsible for monitoring system logs.

In a single-tenant deployment, the same person might also act as the tenant administrator.

IaaS Administrator

IaaS administrators manage cloud, virtual, networking, and storage infrastructure at the system level, creating and managing endpoints and credentials, and monitoring IaaS logs. IaaS administrators organize infrastructure into tenant-level fabric groups, appointing the fabric administrators who are responsible for allocating resources within each tenant through reservations and reservation, storage, and networking policies.

System-Wide Roles and Responsibilities

Users with system-wide roles manage configurations that can apply to multiple tenants. The system administrator is only present in the default tenant, but you can assign IaaS administrators to any tenant.

Table 3. System-Wide Roles and Responsibilities

Role	Responsibilities	How Assigned
System Administrator	<ul style="list-style-type: none"> ■ Create tenants. ■ Configure tenant identity stores. ■ Assign IaaS administrator role. ■ Assign tenant administrator role. ■ Configure system default branding. ■ Configure system default notification providers. ■ Monitor system event logs, not including IaaS logs. ■ Configure the vRealize Orchestrator server for use with XaaS. ■ Create and manage (view, edit, and delete) reservations across tenants if also a fabric administrator. 	Built-in administrator credentials are specified when configuring single sign-on.
IaaS Administrator	<ul style="list-style-type: none"> ■ Configure IaaS features, global properties. ■ Create and manage fabric groups. ■ Create and manage endpoints. ■ Manage endpoint credentials. ■ Configure proxy agents. ■ Manage Amazon AWS instance types. ■ Monitor IaaS-specific logs. ■ Create and manage (view, edit, and delete) reservations across tenants if also a fabric administrator. 	The system administrator designates the IaaS administrator when configuring a tenant.

Tenant Role Overview

Tenant roles typically have responsibilities that are limited to a specific tenant and cannot affect other tenants in the system.

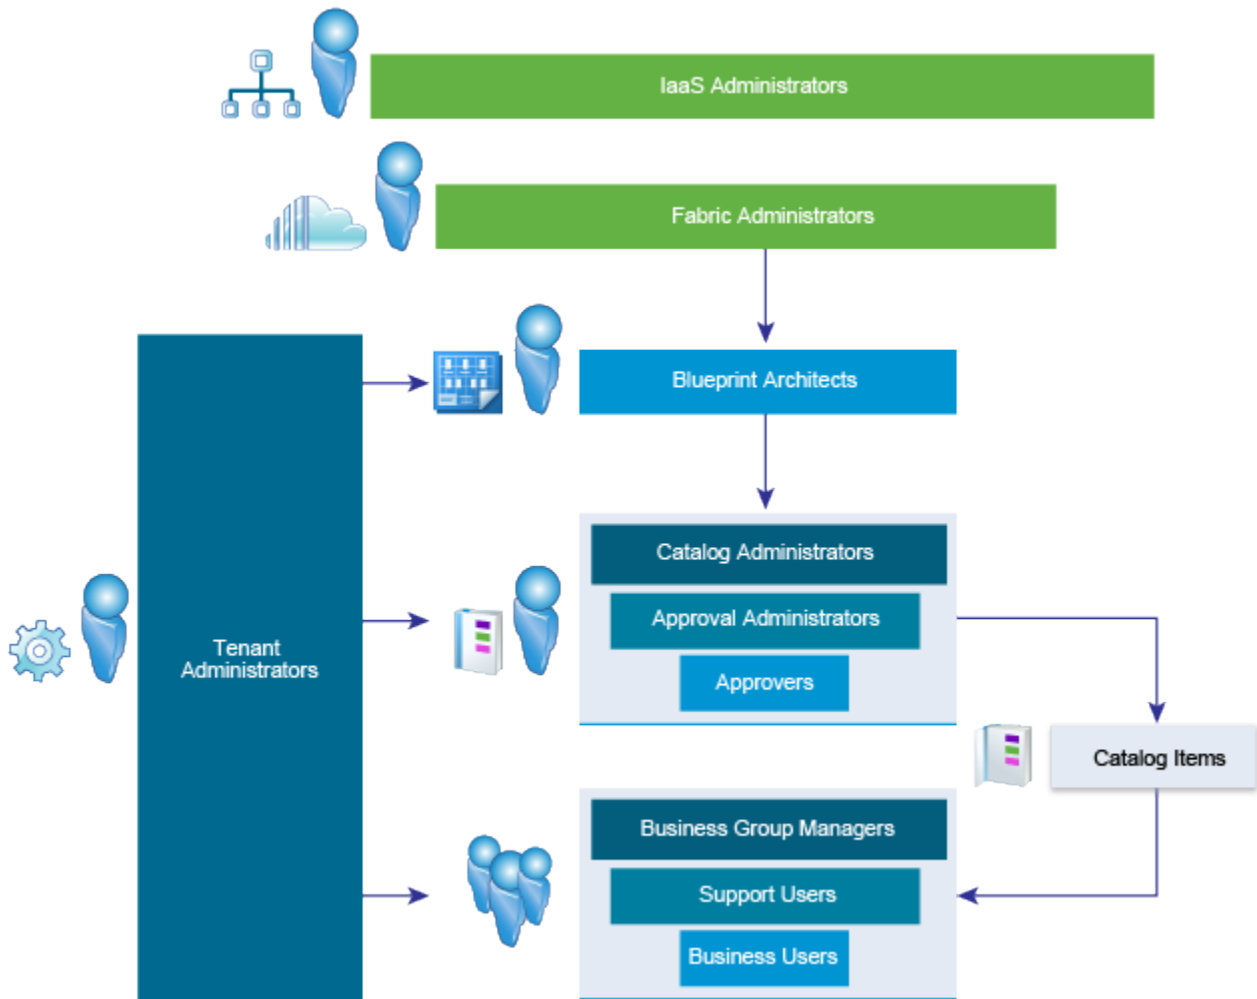


Table 4. Tenant Role Overview

Role	Description
Tenant Administrator	Typically a line-of-business administrator, business manager, or IT administrator who is responsible for a tenant. Tenant administrators configure vRealize Automation for the needs of their organizations. They are responsible for user and group management, tenant branding and notifications, and business policies such as approvals and entitlements. They also track resource usage by all users within the tenant and initiate reclamation requests for virtual machines.
Fabric Administrator	<p>Manages physical machines and compute resources assigned to their fabric groups and creates and manages the reservations and policies associated with those resources within the scope of their tenant. They also manage property groups, machine prefixes, and the property dictionary that are used across all tenants and business groups.</p> <p>Note If you add the fabric administrator role to a system-wide role such as IaaS administrator or system administrator, the fabric administrator can create reservations for any tenant, not just their own.</p>
Blueprint Architects	Umbrella term for the individuals who are responsible for creating blueprint components and assembling the blueprints that define catalog items for consumers to request from the service catalog. These roles are typically assigned to individuals in the IT department, such as architects or analysts.
Catalog Administrator	Creates and manages catalog services and manages the placement of catalog items into services.
Approval Administrator	Defines approval policies. These policies can be applied to catalog requests through entitlements that a tenant administrator or business group manager manage.
Approver	Any user of vRealize Automation, for example, a line manager, finance manager, or project manager, can be designated as an approver as part of an approval policy.
Business Group Manager	Manages one or more business groups. Typically a line manager or project manager. Business group managers entitlements for their groups in the service catalog. They can request and manage items on behalf of users in their groups.
Support User	A role in a business group. Support users can request and manage catalog items on behalf of other members of their groups.
Business User	Any user in the system can be a consumer of IT services. Users can request catalog items from the service catalog and manage their provisioned resources.
Health Consumer	Any user of vRealize Automation, for example, a line manager, finance manager, or project manager, can be designated as a Health Consumer with read-only privileges for Health Service reports.

Tenant Roles and Responsibilities in vRealize Automation

You can assign tenant roles to users in any tenant. The roles have responsibilities that are specific to that tenant.

Table 5. Tenant Roles and Responsibilities

Role	Responsibilities	How Assigned
Tenant administrator	<ul style="list-style-type: none"> ■ Customize tenant branding. ■ Manage tenant identity stores. ■ Manage user and group roles. ■ Create custom groups. ■ Manage notification providers. ■ Enable notification scenarios for tenant users. ■ Configure vRealize Orchestrator servers, plug-ins and workflows for XaaS. ■ Create and manage catalog services. ■ Manage catalog items. ■ Manage actions. ■ Create and manage entitlements. ■ Create and manage approval policies. ■ Monitor tenant machines and send reclamation requests. 	The system administrator designates a tenant administrator when creating a tenant. Tenant administrators can assign the role to other users in their tenant at any time from the Administration tab.
Fabric administrator	<ul style="list-style-type: none"> ■ Manage property groups. ■ Manage compute resources. ■ Manage network profiles. ■ Manage Amazon EBS volumes and key pairs. ■ Manage machine prefixes. ■ Manage property dictionary. ■ Create and manage reservations and reservation policies in their own tenant. ■ If this role is added to a user with IaaS administrator or system administrator privileges, the user can create and manage reservations and reservation policies in any tenant. 	The IaaS administrator designates the fabric administrator when creating or editing fabric groups.
Application architect To successfully add software components to the design canvas, you must also have business group member, business group administrator, or tenant administrator role access to the target catalog.	<ul style="list-style-type: none"> ■ Assemble and manage composite blueprints. 	Tenant administrators can assign this role to users in their tenant at any time from the Administration tab.

Table 5. Tenant Roles and Responsibilities (Continued)

Role	Responsibilities	How Assigned
<p>Infrastructure architect</p> <p>To successfully add software components to the design canvas, you must also have business group member, business group administrator, or tenant administrator role access to the target catalog.</p>	<ul style="list-style-type: none"> ■ Create and manage infrastructure blueprint components. ■ Assemble and manage composite blueprints. 	<p>Tenant administrators can assign this role to users in their tenant at any time from the Administration tab.</p>
<p>XaaS architect</p>	<ul style="list-style-type: none"> ■ Define custom resource types. ■ Create and publish XaaS blueprints. ■ Create and manage resource mappings. ■ Create and publish resource actions. 	<p>Tenant administrators can assign this role to users in their tenant at any time from the Administration tab.</p>
<p>Software architect</p> <p>To successfully add software components to the design canvas, you must also have business group member, business group administrator, or tenant administrator role access to the target catalog.</p>	<ul style="list-style-type: none"> ■ Create and manage software blueprint components. ■ Assemble and manage composite blueprints. 	<p>Tenant administrators can assign this role to users in their tenant at any time from the Administration tab.</p>
<p>Container architect</p>	<ul style="list-style-type: none"> ■ Add, edit, and remove container components in a blueprint by using options on the Design tab. ■ Add, edit, and remove container network components in a blueprint by using options on the Design tab. 	<p>Tenant administrators can assign this role to users and groups in their tenant at any time from the Administration tab.</p>
<p>Container administrator</p>	<p>Use all available options in the Containers tab, including the following tasks:</p> <ul style="list-style-type: none"> ■ Configure container hosts, placements, and registries ■ Configure container network settings ■ Create container templates 	<p>Tenant administrators can assign this role to users and groups in their tenant at any time from the Administration tab.</p>
<p>Catalog administrator</p>	<ul style="list-style-type: none"> ■ Create and manage catalog services. ■ Manage catalog items. ■ Assign icons to actions. 	<p>Tenant administrators can assign this role to users in their tenant at any time from the Administration tab.</p>

Table 5. Tenant Roles and Responsibilities (Continued)

Role	Responsibilities	How Assigned
Business group manager	<ul style="list-style-type: none"> ■ Add and delete users within the business group. ■ Assign support user roles to users in the business group. ■ Create and manage entitlements for the business group. ■ Request and manage items on behalf of a user in the business group. ■ Assign approval policies for the business group. ■ Monitor resource usage in a business group. ■ Change machine owner. 	The tenant administrator designates the business group manager when creating or editing business groups.
Shared access user	<ul style="list-style-type: none"> ■ Use and run actions on the resources that other business group members deploy. ■ Cannot initiate a provisioning request. 	The tenant administrator designates the shared access users when creating or editing business groups.
Approval administrator	<ul style="list-style-type: none"> ■ Create and manage approval policies. 	Tenant administrators can assign this role to users in their tenant at any time from the Administration tab.
Approver	<ul style="list-style-type: none"> ■ Approve service catalog requests, including provisioning requests or any resource actions. 	The tenant administrator or approval administrator creates approval policies and designates the approvers for each policy.
Support user	<ul style="list-style-type: none"> ■ Request and manage service catalog items on behalf of the other members of the business group ■ Change machine owner. 	The tenant administrator designates the support user when creating or editing business groups.
Business user	<ul style="list-style-type: none"> ■ Request service catalog items to which they are entitled. ■ Manage their provisioned resources. 	The tenant administrator designates the business users who can consume IT services when creating or editing business groups.
Health Consumer	<ul style="list-style-type: none"> ■ Can view test results. ■ Cannot configure, edit, or delete a test. 	The IaaS administrator designates privilege to any role..
Security administrator	<ul style="list-style-type: none"> ■ Create a message board whitelist. 	Tenant administrators can assign this role to users in their tenant at any time from the Administration tab.

Containers User Roles and Access Privileges

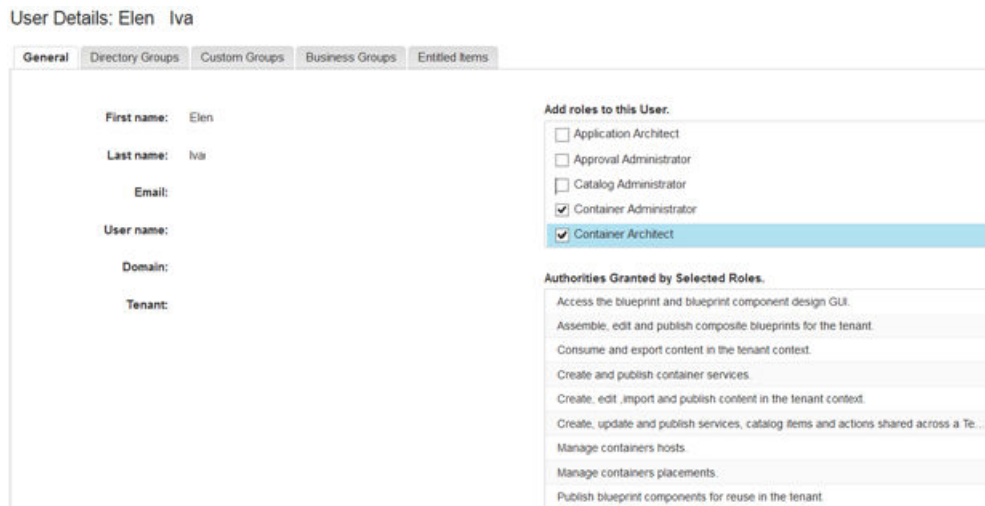
You can use container-specific roles to control who can create and configure containers by using options in the vRealize Automation Containers tab and who can add and configure container components in blueprints by using options in the **Design** tab.

When you enable Containers, two container-specific roles appear in the list of roles that a vRealize Automation tenant administrator can assign to users and groups.

User Role	Description
Container Administrator	Users and groups with this role can see the Containers tab in vRealize Automation. They can use all the Containers options, such as configuring hosts, placements, and registries. They can also create templates and provision containers and applications for configuration and validation purposes.
Container Architect	Users and groups with this role can use containers as components when creating and editing blueprints in vRealize Automation. They have permission to see the Design tab in vRealize Automation and to work with blueprints.

For information about vRealize Automation administrator and user roles, see *User Roles Overview* in the *vRealize Automation Information Center*.

Tenant administrators can assign one or both of these roles to users or groups in their tenant at any time by using options on the vRealize Automation **Administration** tab.



IaaS administrators automatically inherit the container administrator permissions to perform Containers administrative tasks.

Consumers of catalog items that involve containers inherit the necessary privileges to access the resources provided by the Containers. They can open and see the details of their container-related items and perform day-two operations on them.

vRealize Automation users authenticated through VMware Identity Manager have access to Containers.

vRealize Automation multi-tenancy and business group membership is implemented in Containers.

Service Catalog

The service catalog provides a common interface for consumers of IT services to use to request and manage the services and resources they need.

Requesting and Managing Items in the Catalog

The catalog provides a self-service portal for requesting services and also enables business users to manage their own provisioned resources.

The following example is of a typical life cycle.

Connie, the consumer of IT services, logs in to the vRealize Automation console. On the **Catalog** tab, she browses for the service offerings she needs to do her job. The items that are available in the catalog are grouped into service categories, which helps her find what she is looking for. After Connie selects a catalog item, she can view its details to confirm that it is what she wants before submitting a request.

When Connie requests a catalog item, a form appears where she can provide information such as the reason for her request, and any parameters for the request. For example, if she is requesting a virtual machine, she might be able to specify the number of CPUs or amount of storage on the machine. If Connie is not ready to submit her request, she can save it and return to it at a later time.

After Connie submits her request, it might be subject to approval. Connie can look on the **Requests** tab to track the progress of her request, including whether it is pending approval, in progress, or completed.

If the request results in an item being provisioned, it is added to Connie's list of items on the **Items** tab. Here she can view the item details or perform additional actions on her items. In the virtual machine example, she might be able to power on or power off the machine, connect to it through Remote Desktop, reconfigure it to add more resources, or dispose of it when she no longer needs it. The actions she can perform are based on entitlements and can also be made subject to approval based on flexible approval policies.

Creating and Publishing Catalog Items

Catalog administrators and tenant administrators can define new catalog items and publish them to the service catalog. Tenant administrators and business group managers can entitle the new item to consumers.

Typically, a catalog item provides a complete specification of the resource to be provisioned and the process to initiate when the item is requested. It also defines the options that are available to a requester of the item, such as virtual machine configuration or lease duration, or any additional information that the requester is prompted to provide when submitting the request.

For example, Sean has privileges to create and publish blueprints, including software components and XaaS. After the blueprint is published, Sean, or a catalog administrator or a tenant administrator responsible for managing the catalog, can then configure the catalog item, including specifying an icon and adding the item to a service.

To make the catalog item available to users, a tenant administrator or business group manager must entitle the item to the users and groups who should have access to it in the service catalog.

Services for the Service Catalog

Services are used to organize catalog items into related offerings to make it easier for service catalog users to browse for the catalog items they need.

For example, catalog offerings can be organized into Infrastructure Services, Application Services, and Desktop Services.

A tenant administrator or catalog administrator can specify information about the service such as the service hours, support team, and change window. Although the catalog does not enforce service-level agreements on services, this information is available to business users browsing the service catalog.

Catalog Items

Users can browse the service catalog for catalog items that they are entitled to request.

Some catalog items result in an item being provisioned that the user can manage through its life cycle. For example, an application developer can request storage as a service, then later add capacity, request backups, and restore previous backups.

Other catalog items do not result in provisioned items. For example, a cell phone user can submit a request for additional minutes on a mobile plan. The request initiates a workflow that adds minutes to the plan. The user can track the request as it progresses, but cannot manage the minutes after they are added.

Some catalog items are available only in a specific business group, other catalog items are shared between business groups in the same tenant.

Actions

Actions are operations that you can perform on provisioned items.

Users can manage their provisioned items on the **Items** tab. The **View Details** option is always present in the **Actions** menu. Additional options might be available depending on the type of item and the user's entitlements. For example, Power On can be available for machines but not for HR services such as provisioning a new hire.

You can perform request actions and immediate actions. Request actions initiate requests, which you can track on the **Requests** tab and which can be made subject to approval. Statuses shown on the **Requests** tab indicate the success or failure of the request, and do not indicate the successful completion of an action. Immediate actions do not create requests and are always run immediately.

Built-in actions are available to all tenants and cannot be edited, although they can be enabled or disabled. Custom actions can be created at a per-tenant level and shared across all business groups in that tenant.

Entitlements

Entitlements determine which users and groups can request specific catalog items or perform specific actions. Entitlements are specific to a business group.

Business group managers can create entitlements for the groups that they manage. Tenant administrators can create entitlements for any business group in their tenant. When you create an entitlement, you must select a business group and specify individual users and groups in the business group for the entitlement.

You can entitle an entire service category, which entitles all of the catalog items in that service, including items that are added to the service after you create the entitlement. You can also add individual catalog items in a service to an entitlement. Services do not contain actions. You must add actions to an entitlement individually.

For each service, catalog item, or action that you entitle, you can optionally specify an approval policy to apply to requests for that item. If you entitle an entire service and a specific catalog item in that service in the same entitlement, the approval policy on the catalog item overrides the policy on the service. For example, you can entitle the Cloud Infrastructure service to members of a business group and allow them to request any of its items with no approval policy. For a select number of catalog items that require more governance for their provisioning, you can entitle those in the same entitlement and apply an approval policy on just those items.

The actions that you entitle to users apply to any items that support the entitled action and they are not limited to the services and actions in the same entitlement. For example, if Connie, a consumer of infrastructure services, is entitled to Machine Blueprint 1 and the action Reconfigure in one entitlement, and she is also entitled to Machine Blueprint 2 in a different entitlement, then she is entitled to reconfigure machines provisioned from Machine Blueprint 1 and Machine Blueprint 2, as long as both blueprints allow that action to be performed.

If multiple entitlements exist for the same business group, you can prioritize the entitlements. When a user makes a catalog request, the entitlement and associated approval policy that applies is the highest priority entitlement that grants the user access to that item or action.

Approval Policies

An approval policy is used to govern whether a service catalog user needs approval from someone in your organization to provision items in your environment.

A tenant administrator or approval administrator can create approval policies. The policies can be for pre-provisioning or post-provisioning. If a pre-approval is configured, then the request must be approved before the request is provisioned. If it is a post-approval, the request must be approved before the provisioned item is released to the requesting user.

The policies are applied to items in an entitlement. You can apply them to services, catalog items, catalog item components, or actions that require an approver to approve or reject a provisioning request.

When a service catalog user requests an item that includes one or more approval policies, the approval request is sent to the approvers. If approved, the request moves forward. If rejected, the request is canceled and the service catalog user is notified regarding the rejection.

Infrastructure as a Service

With Infrastructure as a Service (IaaS), you can rapidly model and provision servers and desktops across virtual and physical, private and public, or hybrid cloud infrastructures.

- **Configuring Infrastructure Fabric**

The IaaS administrator and fabric administrator roles are responsible for configuring the fabric to enable provisioning of infrastructure services. Fabric configuration is system-wide and is shared across all tenants.

- **Infrastructure Source Endpoints**

Infrastructure sources can include a group of virtualization compute resources or a cloud service account.

- **Compute Resources**

A compute resource is an object that represents a host, host cluster, or pool in a virtualization platform, a virtual datacenter, or an Amazon region on which machines can be provisioned.

- **Data Collection**

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

- **Fabric Groups**

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

- **Business Groups**

A business group associates a set of services and resources to a set of users, often corresponding to a line of business, department, or other organizational unit.

- **Machine Prefixes**

You use machine prefixes to generate the names of provisioned machines.

- **Resource Reservations**

You can create a reservation to allocate provisioning resources in the fabric group to a specific business group.

- **Configuring Reservation Policies**

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

- **Machine Blueprints**

A blueprint that contains a machine component specifies the workflow used to provision a machine and includes information such as CPU, memory, and storage. Machine blueprints specify the workflow used to provision a machine and include additional provisioning information such as the locations of required disk images or virtualization platform objects. Blueprints also specify policies such as the lease period and can include networking and security components such as security groups, policies, or tags.

- **Machine Leases and Reclamation**

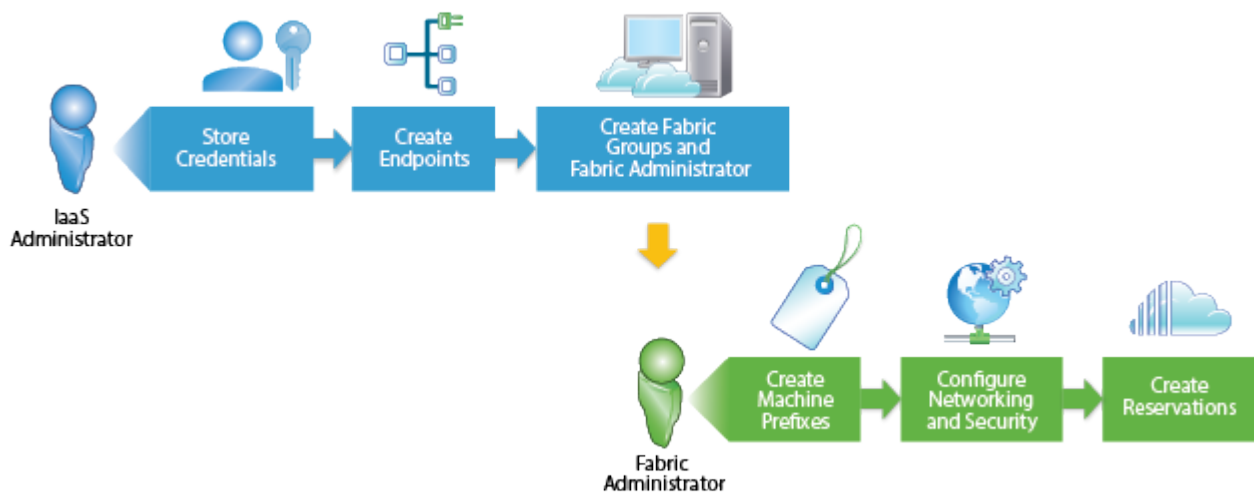
Machine lease and reclamation options provides mechanisms for controlling resource use and controlling prices.

- **Scaling and Reconfiguring Deployments**

You can scale provisioned deployments to adjust to changing workload demands. You use the scale in or scale out actions for horizontal scale, and the machine reconfigure action for vertical scale. You govern scale and reconfigure actions by using entitlements, approval policies, or by designing constraints directly into blueprints.

Configuring Infrastructure Fabric

The IaaS administrator and fabric administrator roles are responsible for configuring the fabric to enable provisioning of infrastructure services. Fabric configuration is system-wide and is shared across all tenants.



An IaaS administrator creates an endpoint to configure access to an infrastructure source. When the connection to an infrastructure source is established, vRealize Automation collects information about the compute resources available through that source. The IaaS administrator can then organize those resources into fabric groups and assign a fabric administrator to manage each group as well as cross-tenant configuration such as machine prefixes.

A fabric administrator can create reservations to allocate provisioning resources in the fabric group to specific business groups that the tenant administrator created during tenant configuration. Optionally, the fabric administrator can configure reservation, network, or storage reservation policies. For example, they can create a reservation policy to control placement of provisioned machines.

When the fabric administrator has created reservations, the IaaS architects can create and publish machine blueprints for reuse in application blueprints and for catalog administrators to make available in the service catalog.

Infrastructure Source Endpoints

Infrastructure sources can include a group of virtualization compute resources or a cloud service account.

An IaaS administrator configures an infrastructure source by specifying the endpoint details and credentials that vRealize Automation can use to communicate with the source.

vRealize Automation collects information about all configured infrastructure sources at regular intervals.

Table 6. Examples of Infrastructure Source Endpoints

Infrastructure Source	Endpoints
vSphere	vCenter Server
vCloud Air	vCloud Air OnDemand or subscription service
vCloud Director	vCloud Director server
Amazon or OpenStack	Cloud service account
Hyper-V (SCVMM)	Microsoft System Center Virtual Machine Manager server
KVM (RHEV)	Red Hat Enterprise Virtualization server

Compute Resources

A compute resource is an object that represents a host, host cluster, or pool in a virtualization platform, a virtual datacenter, or an Amazon region on which machines can be provisioned.

An IaaS administrator can add compute resources to or remove compute resources from a fabric group. A compute resource can belong to more than one fabric group, including groups that different fabric administrators manage. After a compute resource is added to a fabric group, a fabric administrator can create reservations on it for specific business groups. Users in those business groups can then be entitled to provision machines on that compute resource.

Information about the compute resources on each infrastructure source endpoint and machines provisioned on each compute resource is collected at regular intervals.

Table 7. Examples of Compute Resources for Infrastructure Sources

Infrastructure Source	Compute Resource
vSphere (vCenter)	ESX or ESXi host or cluster
Hyper-V (SCVMM)	Hyper-V host
KVM (RHEV)	KVM host
vCloud Director	virtual datacenter
Amazon AWS	Amazon region

Data Collection

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

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

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

Table 8. Data Collection Types

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

Fabric Groups

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

Fabric administrators are responsible for creating reservations on the compute resources in their groups to allocate fabric to specific business groups. Fabric groups are created in a specific tenant, but their resources can be made available to users who belong to business groups in all tenants.

Business Groups

A business group associates a set of services and resources to a set of users, often corresponding to a line of business, department, or other organizational unit.

Business groups are managed in **Administration > Users and Groups** and are used when creating reservations and entitling users to items in the service catalog.

To request catalog items, a user must belong to the business group that is entitled to request the item. A business group can have access to catalog items specific to that group and to catalog items that are shared between business groups in the same tenant. In IaaS, each business group has one or more reservations that determine on which compute resources the machines that this group requested can be provisioned.

A business group must have at least one business group manager, who monitors the resource use for the group and often is an approver for catalog requests. Business groups can include support users. Support users can request and manage machines on behalf of other group members. Business group managers can also submit requests on behalf of their users. A user can be a member of more than one business group, and can have different roles in different groups.

Machine Prefixes

You use machine prefixes to generate the names of provisioned machines.

You should assign a default machine prefix to every business group that you expect to need IaaS resources. Every blueprint must have a machine prefix or use the group default prefix.

Only the machine prefixes that are applicable to the current tenant are exposed with authoring a blueprint or editing a business group.

Fabric administrators are responsible for managing machine prefixes. A prefix is a base name to be followed by a counter of a specified number of digits. For example, a prefix of g1dw for group1 and developer workstation, with a counter of three digits produces machines named g1dw001, g1dw002, and so on. A prefix can also specify a number other than 1 to start the counter.

If a business group is not intended to provision IaaS resources, tenant administrators do not need to assign a default machine prefix when they create the business group. If the business group is intended to provision IaaS resources, tenant administrators should assign one of the existing machine prefixes as the default for the business group. This assignment does not restrict blueprint architects from choosing a different prefix when they create blueprints. A tenant administrator can change the default prefix of a business group at any time. The new default prefix is used in the future, but does not affect previously provisioned machines.

Resource Reservations

You can create a reservation to allocate provisioning resources in the fabric group to a specific business group.

A virtual reservation allocates a share of the memory, CPU and storage resources on a particular compute resource for a business group to use.

A cloud reservation provides access to the provisioning services of a cloud service account, for Amazon AWS, or to a virtual datacenter, for vCloud Director, for a business group to use.

A business group can have multiple reservations on the same compute resource or different compute resources, or any number of reservations containing any number of machines.

A compute resource can also have multiple reservations for multiple business groups. In the case of virtual reservations, you can reserve more resources across several reservations than are physically present on the compute resource. For example, if a storage path has 100 GB of storage available, a fabric administrator can create one reservation for 50 GB of storage and another reservation using the same path for 60 GB of storage. You can provision machines by using either reservation as long as sufficient resources are available on the storage host.

Configuring Reservation Policies

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

You can use a reservation policy to collect resources into groups for different service levels, or to make a specific type of resource easily available for a particular purpose. When a user requests a machine, it can be provisioned on any reservation of the appropriate type that has sufficient capacity for the machine. The following scenarios provide a few examples of possible uses for reservation policies:

- To ensure that provisioned machines are placed on reservations with specific 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.
- As an additional means of using a Pay As You Go allocation model for machine types that support that capability.

Note Reservations defined for vCloud Air endpoints and vCloud Director endpoints do not support the use of network profiles for provisioning machines.

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.

Reservation policies provide an optional means of controlling how reservation requests are processed. You can apply a reservation policy to a blueprint to restrict the machines provisioned from that blueprint to a subset of available reservations.

Machine Blueprints

A blueprint that contains a machine component specifies the workflow used to provision a machine and includes information such as CPU, memory, and storage. Machine blueprints specify the workflow used to provision a machine and include additional provisioning information such as the locations of required disk images or virtualization platform objects. Blueprints also specify policies such as the lease period and can include networking and security components such as security groups, policies, or tags.

A machine blueprint typically refers to a blueprint that contains only one machine component and the associated security and networking elements. It can be published as a standalone blueprint and made available to users in the service catalog. However, published machine blueprints also become available for reuse in your design library, and you can assemble multiple machine blueprints, along with Software components and XaaS blueprints, to design elaborate application blueprints for delivering catalog items that include multiple machines, networking and security, software with full life cycle support, and custom XaaS functionality to your users.

An example of a standalone virtual machine blueprint might be one that specifies a Windows 7 developer workstation with one CPU, 2 GB of memory, and a 30 GB hard disk. A standalone cloud machine blueprint might specify a Red Hat Linux web server image in a small instance type with one CPU, 2 GB of memory, and 160 GB of storage.

Blueprints can be specific to a business group or shared among groups in a tenant, depending on the entitlements that are configured for the published blueprint .

You can add custom properties to a machine component in a blueprint to specify attributes of a machine or to override default specifications. You can also add property groups as a convenience for specifying multiple custom properties.

Machine Leases and Reclamation

Machine lease and reclamation options provides mechanisms for controlling resource use and controlling prices.

Machine leases provide access to a machine for a limited period.

Deployment reclamation allows you to identify underused resources and reclaim them from their owners.

Machine Leases

A blueprint can optionally define a lease duration for machines provisioned from that blueprint.

If a blueprint does not specify a lease period, machines are provisioned from that blueprint with no expiration date. If a blueprint specifies a single value for lease duration, machines are provisioned from that blueprint with an expiration date based on the blueprint lease duration. The expiration date is calculated from the time of the request, not from when the machine is provisioned.

If a blueprint specifies a range of possible lease durations, a user can select the desired lease duration within that range when submitting the machine request. Machine requests can be subject to approval based on the requested lease duration.

When a machine lease expires, the machine is powered off. When the archive period expires, the machine is destroyed. You can reactivate an archived machine by setting the expiration date to a date in the future to extend its lease, and powering it back on.

You can send notification emails to alert machine owners and business group managers that a machine's lease is about to expire and again when the lease expires.

Users can be entitled to request a lease extension at any time before it expires. A business group manager or support user can also change the expiration date for a machine after it is provisioned.

Reclamation Overview

You can use metrics to identify underused machines that might be candidates for deployment reclamation.

You can use the basic metrics provided by vRealize Automation to sort and filter metrics information for all of your machines, or you can configure a vRealize Operations Manager endpoint to provide metrics and health badges for your vSphere virtual machines.

Select the candidate deployment and send a reclamation request to the owners of the machines. The machine owner has a fixed period of time to respond to the request. If machines in the deployment are still in use, the machine owner can stop the reclamation process and continue using the machine. If the machine is no longer needed, the owner can release the machine for reclamation, in which case the machine lease is ended. If the owner does not respond in a timely manner, a lease determined by the administrator is imposed. If the owner continues to take no action, the machine is powered off on the new expiration date, the machine is reclaimed, and the resources are freed.

Scaling and Reconfiguring Deployments

You can scale provisioned deployments to adjust to changing workload demands. You use the scale in or scale out actions for horizontal scale, and the machine reconfigure action for vertical scale. You govern scale and reconfigure actions by using entitlements, approval policies, or by designing constraints directly into blueprints.

Scale In or Scale Out

After you provision a deployment, you can adjust to changing workload demands by increasing or decreasing the number of instances of virtual or cloud machines in your deployment. For example, you deployed a three-tiered banking application with a clustered application server node, a database node, and a load balancer node. Demand increases, and you find that the two instances of your application server node cannot handle all the traffic. Because your blueprint supports up to ten instances of the application server, and you are entitled to scale actions, you can scale out your application. You navigate to your provisioned application item in vRealize Automation and select the scale out action to add another instance of your application server node to the deployment. vRealize Automation provisions a new machine, installs the application software component, and updates your load balancer so your application can handle the increased demands.

If demand decreases, you can scale the deployment in. The newest machines and software components are destroyed first, and your networking and security components are updated so that your deployed application isn't using any unnecessary resources.

Table 9. Support for Scalable Components

Component Type	Supported	Notes
Machine components	Yes	Scale out provisions additional instances of your machines, and scale in destroys machines in last in, first out order.
Software components	Yes	Software components are provisioned or destroyed along with machines that are scaled, and the update life cycle scripts are run for any software components that depend on the scaled machine components.
Networking and security components	Yes	<p>Networking and security components, including NSX load balancers, security groups and security tags, are updated for the new deployment configuration. Scaling impacts the network and security, including load balancer, settings for the deployment. When you scale in or scale out a deployment that contains one or more nodes, the associated NSX networking components are updated. For example, if there is an on-demand NAT networking component associated with the deployment, the NAT rules are updated in accordance with the scaling request.</p> <p>When you scale in or scale out a deployment that contains an associated load balancer, the load balancer is automatically configured to include newly added machines or to stop load balancing machines that are targeted for tear down.</p> <p>When you scale out a deployment that contains a load balancer, secondary IP addresses are added to the load balancer. Depending on whether you scale in or scale out, virtual machines are added or removed from the load balancer and saved or removed in the IaaS database.</p>
XaaS components	No	<p>XaaS components are not scalable and are not updated during a scale operation. If you are using XaaS components in your blueprint, you could create a resource action for users to run after a scale operation, which could either scale or update your XaaS components as required. Alternatively, you could disable scale by configuring exactly the number of instances you want to allow for each machine component.</p>
Nested blueprints	Yes	Supported components in nested blueprints might only update if you create explicit dependencies to scaled machine components. You create explicit dependencies by drawing dependency lines on the design canvas.

When you scale out a deployment, vRealize Automation allocates the requested resources on the current reservation before proceeding. If the scale is partially successful, and fails to provision one or more items against those allocated resources, the resources are not deallocated and do not become available for new requests. Resources that are allocated but unused because of a scale failure are known as dangling resources. You can try to repair partially successful scale operations by attempting to scale the deployment again. However, you cannot scale a deployment to its current size, and fixing a partially successful scale this way does not deallocate the dangling resources. You can view the request execution details screen and find out which tasks failed on which nodes to help you decide whether to fix the partially successful scale with another scale operation. Failed and partially successful scale operations do not impact the functionality of your original deployment, and you can continue to use your catalog items while you troubleshoot any failures.

For a clustered deployment, in which the deployment created from a blueprint contains more than one VM, scaling fails if the blueprint uses a hostname custom property but does not contain a machine prefix value. To avoid this issue, you can use the machine prefix option in the blueprint definition. Otherwise, the scaling function attempts to use the same hostname setting for each VM in the cluster. For more information, see VMware Knowledge Base article 2148213 at <http://kb.vmware.com/kb/2148213>.

Scale Up or Scale Down By Using Reconfigure

After you provision a vSphere, vCloud Air, or vCloud Director virtual or cloud machine you can adjust to changing workload demands by requesting a machine reconfigure to increase (scale up) or decrease (scale down) machine resource specifications for CPU, memory, storage, or networks. You can also add, edit, or remove custom properties and change descriptions. You can request to reconfigure machines for scale up or scale down that are in the On or Off state.

When you reconfigure a virtual or cloud machine for scale up, vRealize Automation allocates the requested resources on the current reservation before proceeding. If the resources are not available, the machine reconfigure fails. If a machine reconfigure request fails, any resources allocated for scale up are deallocated and available for new requests. When you reconfigure a virtual or cloud machine for scale down, resources are not made available to new requests unless the reconfigure finishes successfully.

Table 10. Required Entitlements for Machine Reconfigure for Scaling Scenarios (vSphere , vCloud Air , and vCloud Director only

Virtual or Cloud Machine Owner wants to...	Required Entitlements
Run the reconfigure for scaling immediately after any required approvals are given.	Reconfigure
Specify a date and time to run the reconfiguration for scaling.	Reconfigure
Reschedule a reconfigure for scaling because the request was not approved until after the scheduled time.	Reconfigure
Retry a failed reconfigure request.	Execute reconfigure
Cancel a failed reconfigure request.	Cancel reconfigure
Cancel a scheduled reconfigure request.	Cancel reconfigure

XaaS Blueprints and Resource Actions

XaaS architects can use the XaaS options to create blueprints and publish them to the service catalog. They can also create and publish post-provisioning operations that the consumers can perform on provisioned items.

Creating XaaS Blueprints and Actions

By using the XaaS blueprints and resource actions, you define new provisioning, request, or action offerings and publish them to the common catalog as catalog items.

You can create XaaS blueprints and actions for either requesting or provisioning. The XaaS blueprints for requesting do not provision items and provide no options for post-provisioning operations. Examples of XaaS blueprints for requesting include blueprints for sending emails, generating reports, performing complex calculations, and so on. For an XaaS blueprint, the result is a provisioned item. You can create a custom resource so that you can access and manage the items on the **Items** tab.

To define the XaaS specification, you create a blueprint and publish it as a catalog item. After you publish a catalog item, you must include it in a service category. You can use an existing service or create one. A tenant administrator or business group manager can entitle the whole service or only the catalog item to specific users.

If you created a custom resource for a provisioned item, you can create resource actions to define the post-provisioning operations that the consumers can perform. You can also create resource actions for an item that is provisioned by a source different from the XaaS blueprints, for example by IaaS. For this purpose, first you must create a resource mapping to define the type of the catalog item.

Custom Resources

You must create a custom resource so that you can create an XaaS blueprint for provisioning with the option to access and manage the provisioned items. Custom resources define the items for provisioning, and you can use them to define post-provisioning operations that the consumers can perform.

You create a custom resource to define a new type of provisioned item and map it to an existing vRealize Orchestrator object type. vRealize Orchestrator object types are the objects exposed through the APIs of the vRealize Orchestrator plug-ins. The custom resource is the output type of a blueprint workflow for provisioning and can be the input type for a resource action workflow.

For example, if you have a running vCenter Server instance, and you also have the vCenter Server plug-in that is configured to work with vRealize Orchestrator, all of the object types from the vCenter Server API are exposed in vRealize Orchestrator. The vCenter Server plug-in exposes the vSphere inventory objects in the vRealize Orchestrator inventory. The vSphere inventory objects include data centers, folders, ESXi hosts, virtual machines and appliances, resource pools, and so on. You can perform operations on these objects. For example, you can create, clone, or destroy virtual machines.

For more information about the vRealize Orchestrator object types exposed through the vCenter Server API, see the *vCenter Server Plug-In API Reference for vCenter Orchestrator*.

Resource Mappings

You create resource mappings between the vRealize Automation catalog resource type and the vRealize Orchestrator inventory type to manage resources provisioned outside of XaaS.

For example, you might want to create an action so that users can take a snapshot of their Amazon machines. For this action to work on an Amazon machine provisioned, the three components involved, XaaS, vRealize Orchestrator, and IaaS, need a common language. You create that common language by adding a resource mapping in XaaS that runs a vRealize Orchestrator scripting action or workflow to map the IaaS Cloud Machine resource type to the vRealize Orchestrator AWS:EC2Instance inventory type.

vRealize Automation provides resource mappings, and the underlying vRealize Orchestrator script actions and workflows, for vSphere, vCloud Director, and vCloud Air machines.

XaaS Blueprints

An XaaS blueprint is a complete specification of a resource.

With XaaS blueprints, you publish predefined and custom vRealize Orchestrator workflows as catalog items for either requesting or provisioning. Blueprints for requesting run workflows with no provisioning and provide no options for managing a provisioned item. Before you create a blueprint for provisioning, you must map the workflow output parameter as a custom resource. Then you can assign resource actions that define post-provisioning operations.

Resource Actions

You can create custom resource actions to configure the post-provisioning operations that the consumers can perform.

To create post-provisioning operations, you must publish vRealize Orchestrator workflows as resource actions. To create a resource action for an item provisioned by using XaaS, you use a custom resource as an input parameter for the workflow. To create a resource action for an item that is provisioned by a source different from XaaS, you use a resource mapping as an input parameter for the workflow. When you entitle the resource actions, they appear in the **Actions** drop-down menu of the provisioned items on the **Items** tab.

Common Components

vRealize Automation includes several common components in addition to the service catalog and catalog item sources such as Infrastructure as a Service and XaaS.

Notifications

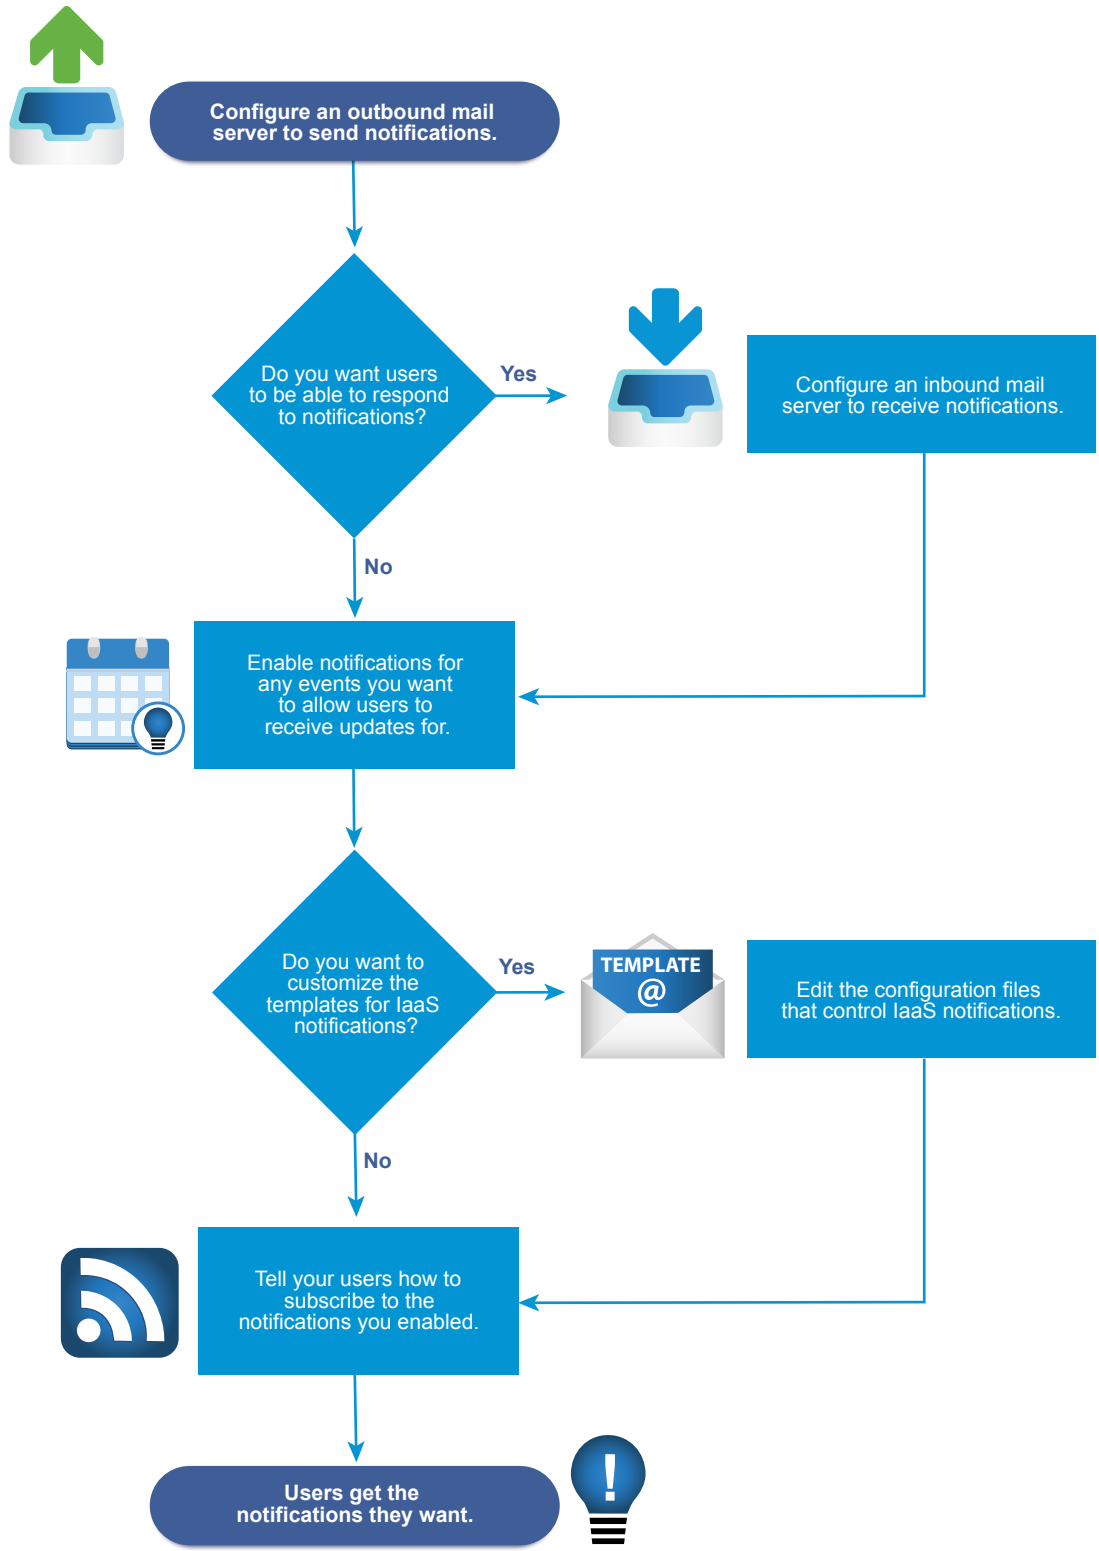
You can send automatic notifications for several types of events, such as the successful completion of a catalog request or a required approval.

System administrators can configure global email servers that process email notifications. Tenant administrators can override the system default servers, or add their own servers if no global servers are specified.

Tenant administrators select which events cause notifications to be sent to users in their tenants. Each component, such as the service catalog or IaaS, can define events that can trigger notifications, but none of them are selected by default.

Each user can choose whether to receive notifications. Users either receive all notifications configured by the tenant administrator or no notifications, they do not have fine-grained control over which notifications to receive.

Some emails have links that users can use to respond to the notification. For example, a notification about a request that requires approval can have one link for approving the request and one for rejecting it. When a user clicks one of the links, a new email opens with content that is automatically generated. The user can send the email to complete the approval.



Branding

Each tenant can change the appearance of the vRealize Automation console and login pages.

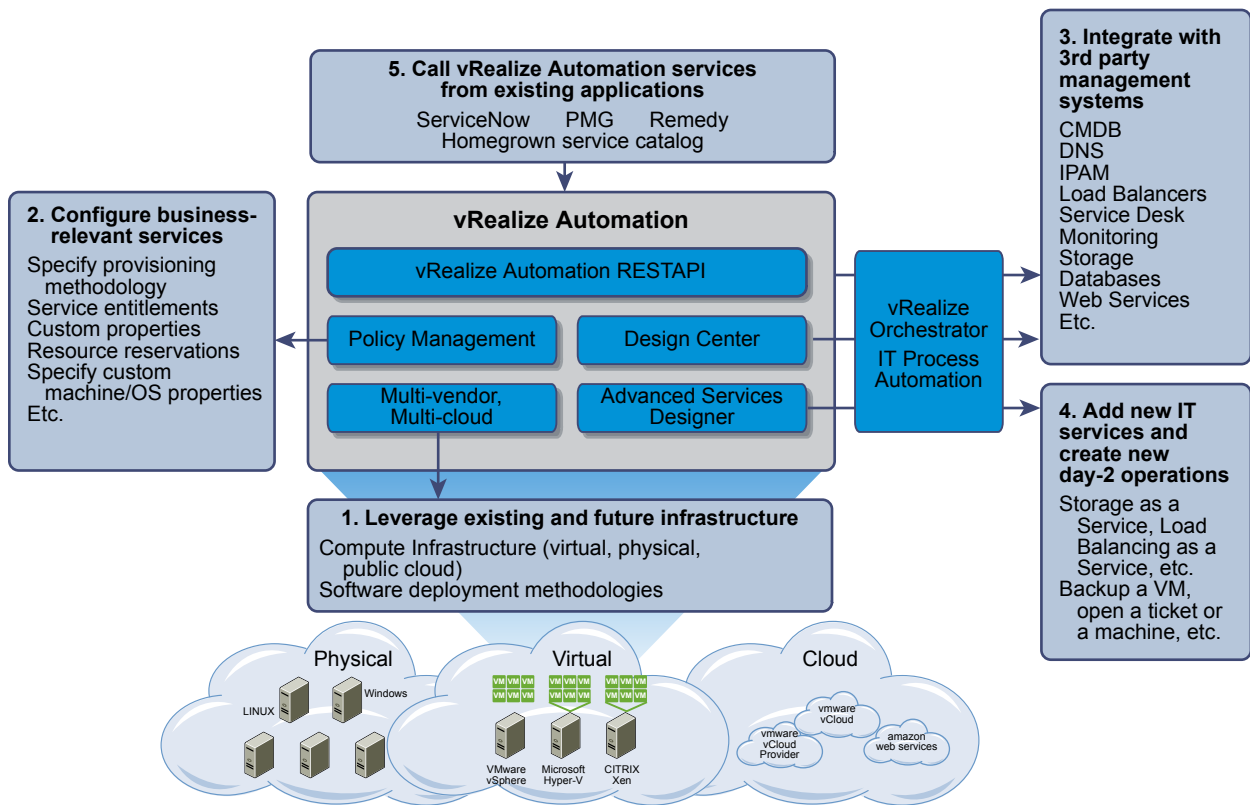
System administrators control the default branding for all tenants. A tenant administrator can change the branding of the portal including the login pages, logo, the background color, and the information in the header and footer. If the branding for a tenant is changed, a tenant administrator can always revert back to the system defaults.

Life Cycle Extensibility

The architecture of vRealize Automation is designed with extensibility in mind. To satisfy different extensibility use cases, vRealize Automation offers a variety of configuration options and tools.

vRealize Automation Extensibility Options

vRealize Automation is a flexible cloud management platform that enables customization and extensibility at multiple levels.



Leveraging Existing and Future Infrastructure

vRealize Automation provides support for many types of infrastructure and provisioning methods.

IaaS administrators can integrate with several infrastructure sources including virtual hypervisors, such as vSphere, Hyper-V, KVM (RHEV), and so on, public clouds including VMware vCloud[®] Air[™] and Amazon AWS, and physical infrastructure.

Blueprint authors can control many machine options, including provisioning methods, by configuring blueprints for various types of infrastructure.

For a full list of supported infrastructure types and provisioning methods, see *vRealize Automation Support Matrix*. For information about configuring infrastructure blueprints, see *Configuring vRealize Automation*.

Configuring Business-Relevant Services

The vRealize Automation console enables administrators to configure business- and user-specific policies through a web-based user interface without writing any code.

These business policies include entitlements and approvals for the service catalog, resource reservation policies for infrastructure, and many others.

For information about customization tasks that you can perform through the vRealize Automation console, see *Configuring vRealize Automation*.

Using custom properties, machine blueprint authors can define additional machine properties or override their standard attributes for a variety of purposes.

For details about the use and configuration of custom properties, see *Configuring vRealize Automation*.

Extending vRealize Automation with Event-Based Workflows

You can use workflow subscriptions to run vRealize Orchestrator workflows based on events.

vRealize Automation provides event topics to which you can subscribe, triggering your custom vRealize Orchestrator workflows when an IaaS resource is provisioned or modified.

For more information, see *Life Cycle Extensibility*.

Integrating with Third-Party Management Systems

Provisioning or decommissioning a new machine, especially for mission-critical systems, typically requires interacting with a number of different management systems, including DNS servers, load balancers, CMDBs, IP Address Management and other systems.

Administrators can inject custom logic (known as workflows) at various predetermined IaaS life cycle stages. These IaaS workflows can call out to vRealize Orchestrator for bi-directional integration with external management systems.

For details about machine life cycle extensibility, see *Life Cycle Extensibility*.

Adding New IT Services and Creating New Actions

The XaaS enables XaaS architects to define new services and new management operations on provisioned resources.

vRealize Automation provides a range of management operations that you can perform on machines. Your organization may find it valuable to extend the default IaaS machine menus with new options, such as creating a machine backup or running a security check.

It can also be beneficial to expose entirely new services in the service catalog, so that users can automate other initiatives directly via the portal. Service architects can create XaaS blueprints for storage-as-a-service, networking services or virtually any kind of IT service by using XaaS.

For details about how to create new catalog items, see *Configuring vRealize Automation*.

Calling vRealize Automation Services from External Applications

In some cases, organizations may want to interact with vRealize Automation programmatically rather than via the vRealize Automation console.

For such scenarios, the vRealize Automation API provides a standardized, secured RESTful interface for cloud access and interaction, controlled through business-aware policy for consumers such as users, infrastructure, devices, and applications.

All blueprints, including the ones created via the XaaS, are automatically exposed through the vRealize Automation API.

Distributed Execution

All core vRealize Automation workflows are executed in a distributed execution environment.

The vRealize Automation runtime environment consists of one or more DEM Worker instances that can execute any workflow installed in the core engine. Additional Worker instances can be added as needed for scalability, availability and distribution.

Skills can be used to associate DEMs and workflows, restricting execution of a given workflow to a particular DEM or set of DEMs with matching skills. Any number and combination of skills can be associated with a given workflow or DEM. For example, workflow execution can be restricted to a specific datacenter, or to environments that support a specific API the workflow requires. The vRealize Automation Designer and the CloudUtil command-line tool provide facilities for mapping skills to DEMs and workflows.

For more information about distributed execution and working with skills, see *Life Cycle Extensibility*.