

Foundations and Concepts

vRealize Automation 6.2



vmware®

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

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

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

docfeedback@vmware.com

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

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

Contents

Foundations and Concepts	5
Updated Information	5
Using the Goal Navigator	6
1 Introducing vRealize Automation	7
Service Catalog Overview	7
Infrastructure as a Service Overview	8
Advanced Service Designer Overview	8
vRealize Business Standard Edition Overview	8
Application Services Overview	9
2 Tenancy and User Roles	10
Tenancy Overview	10
User and Group Management	11
Comparison of Single-Tenant and Multitenant Deployments	11
User Roles Overview	15
System-Wide Role Overview	15
System-Wide Roles and Responsibilities	16
Tenant Role Overview	17
Tenant Roles and Responsibilities in vRealize Automation	18
3 Service Catalog	21
Requesting and Managing Items in the Catalog	21
Creating and Publishing Catalog Items	22
Service Categories	22
Catalog Items	22
Actions	23
Entitlements	23
4 Infrastructure as a Service	25
Configuring Infrastructure Fabric	25
Infrastructure Source Endpoints	26
Compute Resources	26
Data Collection	27
Fabric Groups	28
Business Groups	28
Machine Prefixes	29
Resource Reservations	29

- Configuring Reservation Policies 30
- Cost Profiles 30
- Machine Blueprints 31
- Machine Leases and Reclamation 32
 - Machine Leases 32
 - Reclamation Overview 32
- Machine Reconfiguration 33

5 Advanced Service Designer 34

- Creating Advanced Services 34
- Custom Resources 35
- Resource Mappings 35
- Service Blueprints 35
- Resource Actions 36
- Form Designer 36

6 Common Components 37

- Approvals Overview 37
- Notifications 38
- Branding 38

7 Extensibility 39

- vRealize Automation Extensibility Options 39
- Leveraging Existing and Future Infrastructure 40
- Configuring Business-Relevant Services 40
- Integrating with Third-Party Management Systems 41
- Adding New IT Services and Creating New Actions 41
- Calling vRealize Automation Services from External Applications 41
- Distributed Execution 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
- Advanced Service Designer

For information about VMware vRealize™ Business™ Standard Edition or vRealize Automation Application Services, see the documentation for the respective product or feature.

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>.

Updated Information

This *Foundations and Concepts* is updated with each release of the product or when necessary.

This table provides the update history of the *Foundations and Concepts*.

Revision	Description
001650-04	Updated Data Collection to add default data collection intervals.
001650-03	Made minor editorial updates.
001650-02	Updated Actions .
001650-01	Updated information about machine lease expiration and archive behavior. See Machine Leases .
001650-00	Initial 6.2 release.

Using the Goal Navigator

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

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

The goal navigator can answer the following questions:

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

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

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

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

Introducing vRealize Automation

1

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 service, including infrastructure, applications, desktops, and many others, are processed through a common service catalog to provide a consistent user experience.

vRealize Automation improves cost control by exposing the cost of cloud resources and by providing simplified resource management including reporting on capacity use.

This chapter includes the following topics:

- [Service Catalog Overview](#)
- [Infrastructure as a Service Overview](#)
- [Advanced Service Designer Overview](#)
- [vRealize Business Standard Edition Overview](#)
- [Application Services Overview](#)

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.

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 virtual, cloud, or physical machine. Blueprints are published as catalog items in the common service catalog. When a 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.

Advanced Service Designer Overview

With the Advanced Service Designer, service architects can create advanced services and publish them as catalog items.

With advanced services, you can provide anything as a service using the capabilities of VMware vRealize™ Orchestrator™. For example, you can create a service 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.

With the Advanced Service Designer, a service architect can create custom resource types mapped to vRealize Orchestrator object types and define them as items to be provisioned. A service 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 Advanced Service Designer to design additional actions that the consumer of the service 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 Advanced Service Designer, 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.

vRealize Business Standard Edition Overview

With vRealize Business Standard Edition, directors of cloud operations can monitor their expenditures and design more cost-efficient cloud services.

vRealize Business Standard Edition provides the following benefits:

- Drives accountability by providing visibility into the cost of virtual infrastructure and public cloud providers.
- Promotes efficiencies in the virtual infrastructure by making it possible to compare the costs, 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 Standard Edition, see the vRealize Business Standard Edition documentation set.

Application Services Overview

Formerly known as Application Director, Application Services automates and manages the life cycle of deployments for multitier enterprise applications in hybrid cloud environments.

Enterprise users can standardize, deploy, configure, update, and scale complex applications in dynamic cloud environments. These applications can range from simple Web applications to complex custom applications and packaged applications. Artifact management supports the use of logical names for build files and other types of software artifacts, allowing users to deploy applications without regard for the physical location or identifier of such files.

Application architects can use the drag-and-drop interface to create visual application blueprints. Application architects can use the prepopulated and extensible library of standard logical templates, application infrastructure service, components, and scripts to model an application blueprint. These blueprints standardize the structure of the application, including software components, dependencies, and configurations, for repeated deployments.

Teams can deploy application blueprints across different environments. Teams can repeatedly deploy a standard blueprint, customize configurations as allowed, and deploy in IT-approved deployment environments.

From an application blueprint, you can create different deployments by using deployment profiles to test prototypes or deploy mission-critical multitier applications in production environments. You can use composite deployment files to bundle deployment profiles, bind their properties, and deploy them in a specified order. You can also publish deployment profiles as catalog items in vRealize Automation.

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.

This chapter includes the following topics:

- [Tenancy Overview](#)
- [User Roles Overview](#)

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 2-1. 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 2-1. 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 single sign-on during the installation of vRealize Automation, a default tenant is created with the built-in system administrator account to log in to the vRealize Automation console. The system administrator can then configure the default tenant and create additional tenants.

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

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

User and Group Management

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

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

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

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

Comparison of Single-Tenant and Multitenant Deployments

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

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

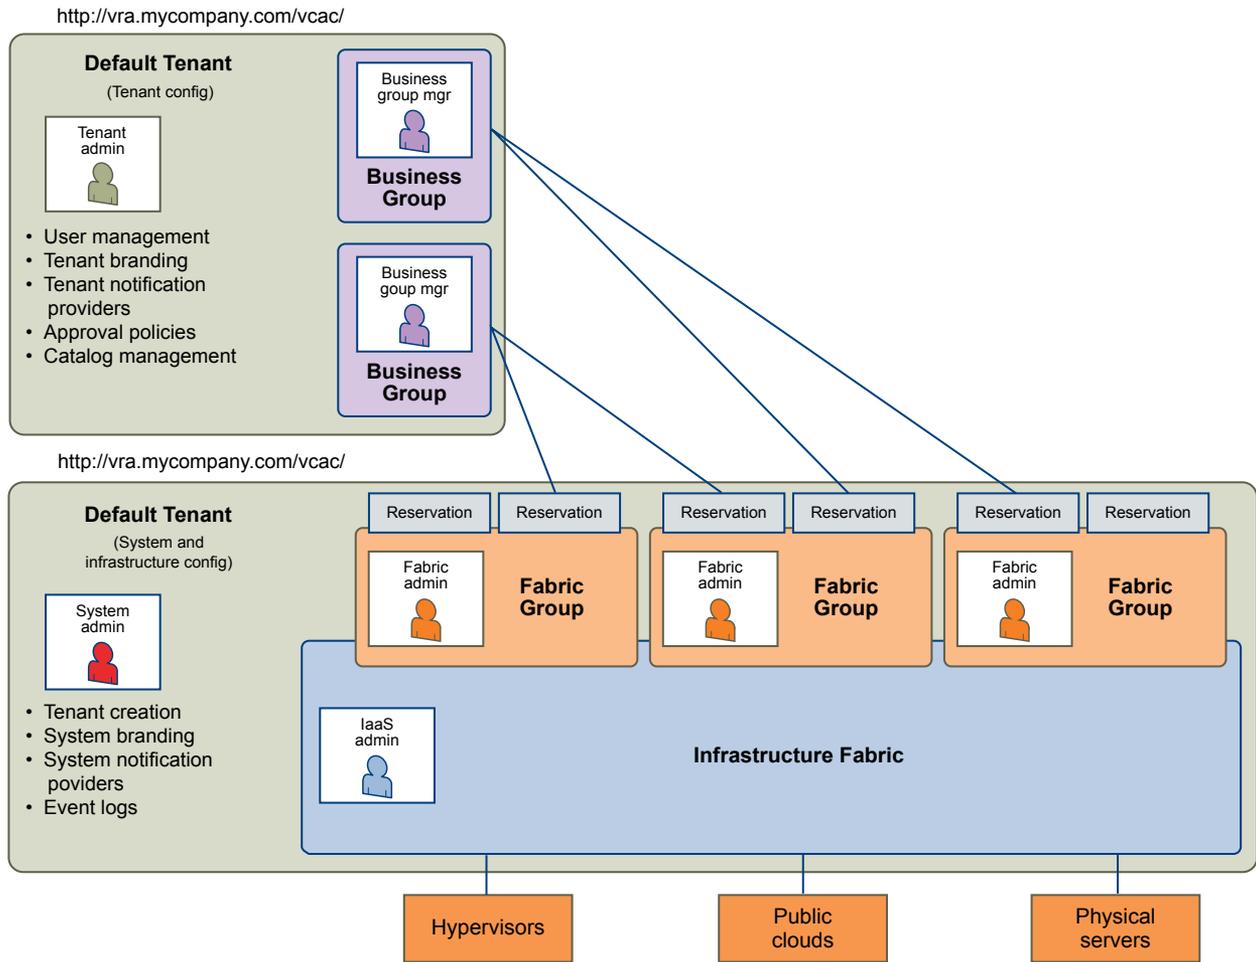
Infrastructure configuration, including the infrastructure sources that are available for provisioning, can be configured in any tenant and is shared among all tenants. The infrastructure resources, such as cloud or virtual compute resources or physical machines, can be divided into fabric groups managed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

Single-Tenant Deployment

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

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

Figure 2-1. Single-Tenant Example



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

Multitenant Deployment

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

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

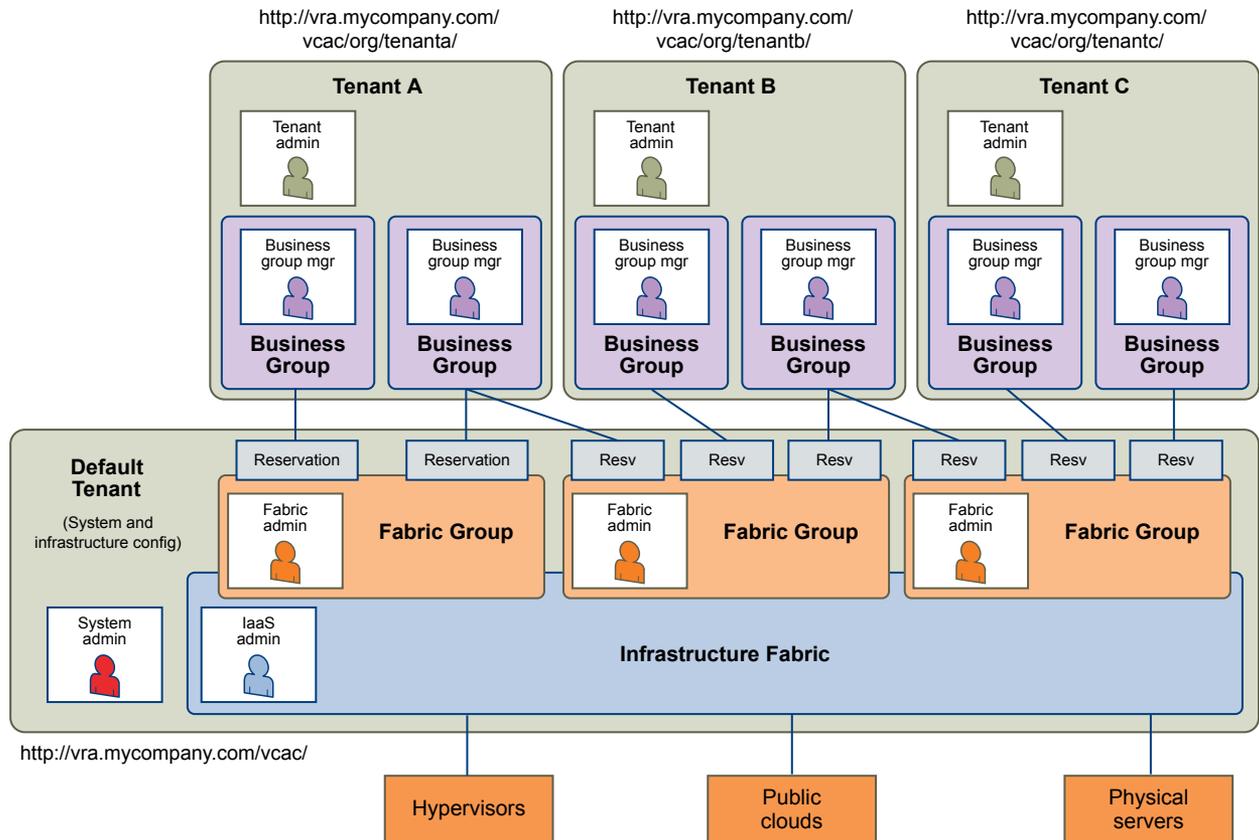
Table 2-2. Multitenant Deployment Examples

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

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

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

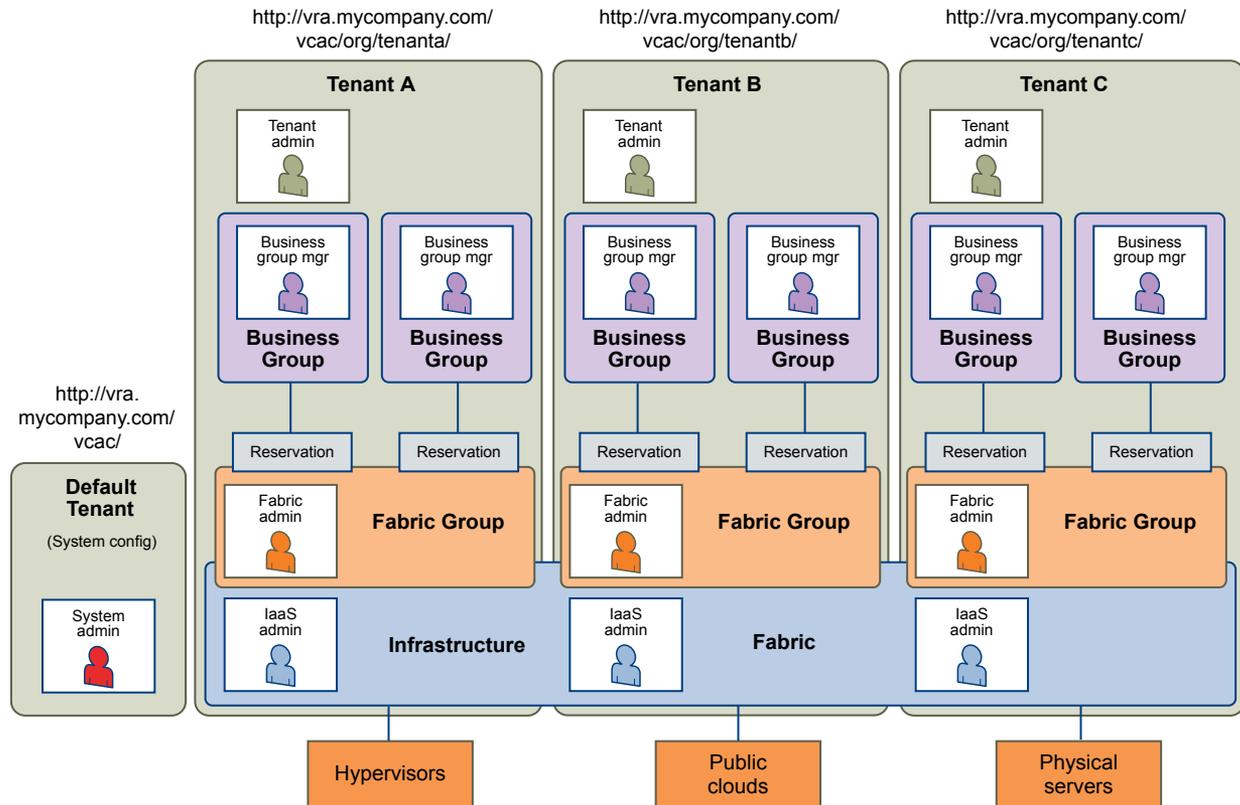
Figure 2-2. Multitenant Example with Infrastructure Configuration Only in Default Tenant



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

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

Figure 2-3. Multitenant 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 and fabric administrator roles 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 endpoints and endpoint credentials, create fabric groups, and configure virtualization proxy agents. They also manage cloud service accounts and physical machines and storage devices. They also monitor logs that are specific to IaaS.

Fabric Administrator

A fabric administrator is the administrator of one or more fabric groups. Fabric administrators manage physical machines and compute resources in their groups and the reservations and reservation policies associated with those resources. They also manage build profiles, machine prefixes, and the property dictionary that are used across all tenants and business groups.

System-Wide Roles and Responsibilities

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

Table 2-3. 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 in the Advanced Services Designer 	Built-in administrator credentials are specified when configuring single sign-on.
IaaS Administrator	<ul style="list-style-type: none"> ■ Configure IaaS features, global properties ■ Manage IaaS licenses ■ Create and manage fabric groups ■ Create and manage endpoints ■ Manage endpoint credentials ■ Configure proxy agents ■ Manage Amazon AWS instance types ■ Monitor IaaS-specific logs 	The system administrator designates the IaaS administrator when configuring a tenant.
Fabric Administrator	<ul style="list-style-type: none"> ■ Manage build profiles ■ Manage compute resources ■ Manage cost profiles ■ Manage network profiles ■ Manage Amazon EBS volumes and key pairs ■ Manage machine prefixes ■ Manage property dictionary ■ Manage reservations and reservation policies 	The IaaS administrator designates the fabric administrator when creating or editing fabric groups.

Tenant Role Overview

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

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.

Service Architect

Umbrella term for the individuals who are responsible for creating the blueprints for the catalog items that consumers can request from the service catalog. This role is typically someone in the IT department, such as an architect or analyst.

In Infrastructure as a Service, tenant administrators and business group managers can create machine blueprints and multi-machine blueprints. In Application Services, application architects can create application blueprints and deployment profiles. In the Advanced Service Designer, service architects can create service blueprints.

Business Group Manager

Manages one or more business groups. Typically a line manager or project manager. Business group managers manage catalog items and entitlements for their groups in the service catalog. They can request and manage items on behalf of users in their groups. They are also service architects in Infrastructure as a Service.

Support User

A role in a business group. Support users can request and manage catalog items on behalf of other members of their groups. This role is typically an executive administrator or department administrator.

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.

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.

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 2-4. Tenant Roles and Responsibilities

Role	Responsibilities	How Assigned
Tenant administrator	<ul style="list-style-type: none"> ■ Manage tenant identity stores. ■ Manage user and group roles. ■ Create custom groups. ■ Customize tenant branding. ■ Manage notification providers. ■ Enable notification scenarios for tenant users. ■ Create and manage approval policies. ■ Manage catalog services. ■ Manage catalog items. ■ Manage actions. ■ Manage entitlements. ■ Monitor tenant machines and send reclamation requests. ■ Configure vRealize Orchestrator servers, plug-ins and workflows for use in the Advanced Services Designer. ■ Create and publish shared machine blueprints from IaaS. 	The system administrator designates a tenant administrator when creating a tenant. Tenant administrators can assign the role to other users in the tenant.
Service architect	<ul style="list-style-type: none"> ■ Define custom resource types. ■ Create and publish service blueprints in the Advanced Services Designer. ■ Create and publish custom actions. 	The tenant administrator can assign this role to give a user or group the rights to create custom services in the Advanced Service Designer.
Application architect	<ul style="list-style-type: none"> ■ Create, modify, and delete applications in Application Services. 	The tenant administrator can assign this role to a user or group. The user or group must be in the tenant registered with Application Services.
Application catalog administrator	<ul style="list-style-type: none"> ■ Define services, templates, operating systems, tasks, and tags in the Application Services library. 	The tenant administrator can assign this role to a user or group. The user or group must be in the tenant registered with Application Services.
Application cloud administrator	<ul style="list-style-type: none"> ■ Define resources and deployment environments. 	The tenant administrator can assign this role to a user or group. The user or group must be in the tenant registered with Application Services.
Application publisher and developer	<ul style="list-style-type: none"> ■ Deploy applications into the vRealize Automation catalog. ■ Create, update, and publish services, library items, and actions in Application Services. 	The tenant administrator can assign this role to a user or group. The user or group must be in the tenant registered with Application Services.

Table 2-4. Tenant Roles and Responsibilities (Continued)

Role	Responsibilities	How Assigned
Business group manager	<ul style="list-style-type: none"> ■ Create and publish business group–specific machine blueprints from IaaS. ■ Manage business group–specific catalog items and entitlements. ■ Monitor resource usage in a business group. 	The tenant administrator designates the business group manager when creating or editing business groups.
Support user	<ul style="list-style-type: none"> ■ Request and manage items on behalf of other users within their business groups. 	The tenant administrator designates the support user when creating or editing business groups.
Business user	<ul style="list-style-type: none"> ■ Request and manage services. 	The tenant administrator designates the business users who can consume IT services when creating or editing business groups.
Approval administrator	<ul style="list-style-type: none"> ■ Create and manage approval policies. 	The tenant administrator can assign this role to give a user or group rights to manage approval policies.
Approver	<ul style="list-style-type: none"> ■ Approve 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.

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.

This chapter includes the following topics:

- [Requesting and Managing Items in the Catalog](#)
- [Creating and Publishing Catalog Items](#)
- [Service Categories](#)
- [Catalog Items](#)
- [Actions](#)
- [Entitlements](#)

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

Service architects, tenant administrators, and business group managers can define new catalog items and publish them to the service catalog. Tenant administrators or business group managers can then manage the presentation of catalog items to the consumer and entitle the new item to consumers.

Typically, a catalog item is defined in a blueprint, which 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 specifications or lease duration, or any additional information that the requester is prompted to provide when submitting the request.

For example, Sean the service architect creates a blueprint in the Advanced Services Designer. While he is working on the blueprint, he can save it in the Advanced Services Designer in a draft state. When he is satisfied with his work he can publish it to the catalog as a catalog item. A tenant administrator or business group manager can perform similar steps in IaaS, and an application architect can perform similar steps in Application Services.

After the blueprint is published, Sean, 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 catalog 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.

Service Categories

Service categories organize catalog items into related offerings to make it easier for 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 service architect can specify information about the service category 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.

4

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.

This chapter includes the following topics:

- [Configuring Infrastructure Fabric](#)
- [Infrastructure Source Endpoints](#)
- [Compute Resources](#)
- [Data Collection](#)
- [Fabric Groups](#)
- [Business Groups](#)
- [Machine Prefixes](#)
- [Resource Reservations](#)
- [Configuring Reservation Policies](#)
- [Cost Profiles](#)
- [Machine Blueprints](#)
- [Machine Leases and Reclamation](#)
- [Machine Reconfiguration](#)

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.

In addition to managing physical machines and compute resources in a fabric group, fabric administrators also manage cross-tenant configuration such as machine prefixes.

A tenant administrator creates business groups and designates managers for each group.

After business groups are created, a fabric administrator can create reservations. A reservation allocates provisioning resources in the fabric group to a specific business group. Optionally, the fabric administrator can configure policies such as reservation policies that control placement of provisioned machines, or cost profiles that associate a cost with the use of specific compute resources.

When the fabric administrator has created reservations, the tenant administrator or business group managers can create machine blueprints and publish them to the catalog.

Infrastructure Source Endpoints

Infrastructure sources can include a group of virtualization compute resources, a cloud service account, or a physical management interface.

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. This information includes virtualization hosts, templates, and ISO images for virtualization environments; virtual datacenters for vCloud Director; regions and machines provisioned on them for Amazon; and installed memory and CPU count for physical management interfaces.

Table 4-1. Examples of Infrastructure Source Endpoints

Infrastructure Source	Endpoints
vSphere (vCenter)	vCenter server
Hyper-V (SCVMM)	Microsoft System Center Virtual Machine Manager server
KVM (RHEV)	Red Hat Enterprise Virtualization server
vCloud Director	vCloud Director server
Amazon AWS	Cloud service account
Dell iDrac or HP iLO	Single standalone machine
Cisco UCS Manager	Pool of physical machines

For a complete list of supported infrastructure source endpoints, see the *vRealize Automation Support Matrix*.

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 4-2. 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 4-3. Data Collection Types

Data Collection Type	Description
Infrastructure Source Endpoint Data Collection	<p>Updates information about virtualization hosts, templates, and ISO images for virtualization environments. Updates virtual datacenters and templates for vCloud Director. Updates regions and machines provisioned on them for Amazon. Updates installed memory and CPU count for physical management interfaces.</p> <p>Endpoint data collection runs every 4 hours.</p>
Inventory Data Collection	<p>Updates the record of the virtual machines whose resource use is tied to a specific compute resource, including detailed information about the networks, storage, and virtual machines. This record also includes information about unmanaged virtual machines, which are machines provisioned outside of vRealize Automation.</p> <p>Inventory data collection runs every 24 hours.</p> <p>The default timeout interval for inventory data collection is 2 hours.</p>
State Data Collection	<p>Updates the record of the power state of each machine discovered through inventory data collection. State data collection also records missing machines that vRealize Automation manages but cannot be detected on the virtualization compute resource or cloud endpoint.</p> <p>State data collection runs every 15 minutes.</p> <p>The default timeout interval for state data collection is 1 hour.</p>

Table 4-3. Data Collection Types (Continued)

Data Collection Type	Description
Performance Data Collection (vSphere compute resources only)	<p>Updates the record of the average CPU, storage, memory, and network usage for each virtual machine discovered through inventory data collection.</p> <p>Performance data collection runs every 24 hours.</p> <p>The default timeout interval for performance data collection is 2 hours.</p>
vCNS inventory data collection (vSphere compute resources only)	<p>Updates the record of network and security data related to vCloud Networking and Security and NSX, particularly information about security groups and load balancing, for each machine following inventory data collection.</p>
WMI data collection (Windows compute resources only)	<p>Updates the record of the management data for each Windows machine. A WMI agent must be installed, typically on the Manager Service host, and enabled to collect data from Windows machines.</p>
Cost data collection (compute resources managed by vRealize Business Standard Edition only)	<p>Updates the CPU, memory, and storage costs for each compute resource managed by vRealize Business Standard Edition. The costs of catalog items that can be provisioned by using the compute resources are updated.</p>

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 on the **Infrastructure** tab but are used throughout the service catalog. Entitlements in the catalog are based on business groups.

To request catalog items, a user must belong to at least one business group. 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. In IaaS, group managers also create and manage machine blueprints for the groups they manage. Business groups can also contain support users, who 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. Machine prefixes are shared across all tenants.

Every business group has a default machine prefix. Every blueprint must have a machine prefix or use the group default prefix.

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.

When tenant administrators create a business group, they must assign one of the existing machine prefixes as its default. This assignment does not restrict business group managers 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

A fabric administrator creates 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 physical reservation is a set of physical machines reserved for a business group to use. Unprovisioned physical machines must be added to a physical reservation before being provisioned or imported, and cannot be removed until they are decommissioned and become unprovisioned.

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 physical reservations containing any number of physical 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 100GB of storage available, a fabric administrator can create one reservation for 50GB of storage and another reservation using the same path for 60GB of storage. You can provision machines by using either reservation as long as sufficient resources are available on the storage host.

You can reserve physical machines only for a single business group. Because physical machines do not belong to fabric groups, all fabric administrators can manage all physical machines and reserve them for a particular 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.

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

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

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

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

Cost Profiles

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

A cost profile includes the following values for daily cost:

- Cost per GB of memory capacity specified in the virtual blueprint or installed in the physical machine

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

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

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

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

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

Machine Blueprints

A machine blueprint is the complete specification for a virtual, cloud, or physical machine and is used to determine a machine's attributes and how it is provisioned.

When a business group member requests a machine, the machine is provisioned according to the specifications in the blueprint, such as CPU, memory, and storage. 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. Finally, blueprints specify policies such as the lease period and which operations are supported on machines provisioned from the blueprint.

An example of a virtual blueprint might be one that specifies a Windows 7 developer workstation with one CPU, 2GB of memory, and a 30GB hard disk. A cloud blueprint might specify a Red Hat Linux web server image in a small instance type with one CPU, 2GB of memory, and 160GB of storage. A physical blueprint might specify Windows Server 2008 R2 installed on a server with exactly two CPUs and at least 4GB of memory.

Machine blueprints can be either specific to a business group or shared among groups in a tenant.

- Business group managers can create group blueprints that can only be entitled to users in a specific business group. A business group manager cannot modify or delete shared blueprints.
- Tenant administrators can create shared blueprints that can be entitled to users in any business group in the tenant. Tenant administrators cannot view or modify group blueprints unless they also have the business group manager role for the appropriate group.

- If a tenant administrator sets the shared blueprint as copyable, the business group manager can also copy it to use as a starting point to create a new group blueprint.

You can add custom properties to a blueprint to specify attributes of a machine or to override default specifications. You can also add build profiles to a blueprint as a convenience for specifying multiple custom properties.

Machine Leases and Reclamation

IaaS provides two mechanisms for controlling resource use and controlling costs. Leases provide access to a machine for a limited period. By using the reclamation process, a tenant administrator can identify underused resources and reclaim them from their owners.

Machine Leases

A machine 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

Tenant administrators can use metrics such as low CPU use, low memory use, or low hard disk use to reclaim virtual machines in their tenant and help control resources.

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.

You can use the metrics to identify underused machines that might be candidates for reclamation. Select the candidate machines 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 the machine is still in use, the 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.

Machine Reconfiguration

vSphere and vCloud Director platforms support reconfiguration of existing machines to modify specifications for CPU, memory, storage, or networks. A machine owner can request a reconfiguration when the machine is in either the On or Off state.

Reconfiguration requests are subject to approval based on entitlements, policies, and the actions enabled for the blueprint.

Table 4-4. vSphere and vCloud Director vApp component machine reconfigure scenarios

Scenario	Required Entitlements
The machine owner can begin a reconfiguration to run immediately after any required approvals are given.	Reconfigure
The machine owner can specify a date and time to run the reconfiguration. Reconfiguration can be rescheduled if the request is not approved until after the scheduled time.	Reconfigure
If the reconfiguration fails, the machine owner can retry the operation.	Execute reconfigure
The machine owner can cancel a scheduled reconfiguration, or if a reconfiguration attempt fails, the user can cancel the reconfiguration operation instead of retrying.	Cancel reconfigure

Before proceeding with the reconfiguration, IaaS validates that any additional resources that were requested are available on the machine's current reservation. If the reconfiguration request decreases the amount of resources that the machine uses, these resources are not deallocated and made available to other requests until after the reconfiguration finishes successfully.

If the reconfiguration fails, any additional resources, such as storage, that were allocated in anticipation of the reconfiguration are deallocated and made available to other machine requests.

Advanced Service Designer

Service architects can use the Advanced Service Designer to create advanced services and publish them to the service catalog. They can also create and publish post-provisioning operations that the consumers can perform on provisioned items.

This chapter includes the following topics:

- [Creating Advanced Services](#)
- [Custom Resources](#)
- [Resource Mappings](#)
- [Service Blueprints](#)
- [Resource Actions](#)
- [Form Designer](#)

Creating Advanced Services

By using the Advanced Service Designer, you can define new service offerings and publish them to the common catalog as catalog items.

You can create advanced services for either requesting or provisioning. The services for requesting do not provision items and provide no options for post-provisioning operations. Examples of services for requesting include services for sending emails, generating reports, performing complex calculations, and so on. For a service that results in 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 service specification, you create a service 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 Advanced Service Designer, 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 advanced service 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 5.5 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 Advanced Service Designer.

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 by using IaaS, the three components involved, Advanced Service Designer, vRealize Orchestrator, and IaaS, need a common language for the Amazon machine. You create that common language by adding a resource mapping in Advanced Service Designer 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 virtual machines, vCloud Director virtual machines, and vCloud Director vApps.

Service Blueprints

A blueprint is a complete specification of a service.

With service blueprints, you can 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 the Advanced Service Designer, 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 the Advanced Service Designer, 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.

Form Designer

The Advanced Service Designer includes a form designer that you can use to design submission and details forms for blueprints, resources, and actions. Based on the presentation of the workflows, the form designer dynamically generates default forms and fields you can use to modify the default forms.

You can create interactive forms that the users can complete for submission of catalog items and resource actions. You can also create read-only forms that define what information the users can see on the details view for a catalog item or a provisioned resource.

As you create objects in the Advanced Service Designer, forms are generated for common use cases.

Table 5-1. Advanced Service Designer Object Types and Associated Forms

Object Type	Default Form	Additional Forms
Service blueprint	Request submission form based on the presentation of the selected workflow.	<ul style="list-style-type: none"> ■ Catalog item details (read-only) ■ Submitted request details (read-only)
Custom resource	Resource details form based on the attributes of the vRealize Orchestrator plug-in inventory type (read-only).	<ul style="list-style-type: none"> ■ None
Resource action	Action submission form based on the presentation of the selected workflow.	<ul style="list-style-type: none"> ■ Submitted action details (read-only)

You can modify the default forms and design new forms. You can drag fields to add and reorder them on the form. You can place constraints on the values of certain fields, specify default values, or provide instructional text for the end user who is completing the form.

Because of their different purposes, the operations you can perform to design read-only forms are limited compared to the operations for designing submission forms.

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 the Advanced Service Designer.

This chapter includes the following topics:

- [Approvals Overview](#)
- [Notifications](#)
- [Branding](#)

Approvals Overview

Any catalog request, whether it is for a new catalog item or to perform an action on a provisioned item, can be subject to approval.

A tenant administrator or business group manager can associate an approval policy with any service, catalog item, or action as part of entitling the items or actions to users.

A tenant administrator or an approval administrator can create approval policies. Each approval policy has an approval policy type that determines to what requests that policy can apply. For example, you can apply some policies only to new catalog item requests, or requests for a specific type of item. You can apply other policies only to post-provisioning actions or to a specific action on provisioned items.

Each approval policy has at least one phase and each phase can have one or more levels.

An approval level represents a single step in a business process. For example, an approval policy can have one level for manager approval, followed by a level for finance approval. An approval level can be designated as always required or required based on certain conditions. The available conditions can vary depending on the approval policy type. For example, an approval policy for virtual machine requests can be conditional based on the number of CPUs in the machine request.

Each level specifies one or more approvers. If a level has multiple approvers, the policy can specify whether all the approvers must approve the request to complete the level, or any one of the approvers can complete the level. If any approver rejects the request, then the entire request is rejected.

For each approval level, you can specify attributes that an approver can edit when completing the approval. For example, in an approval level for IT to review a machine request, they might be able to update the requested CPU, memory, or storage specifications. The approvers must determine that the requested specifications are not appropriate for the proposed use of the machine.

An approval phase is a sequence of approval levels associated with a particular stage in the request. The approval policy type defines how many phases a policy has and which stage in a request triggers a particular approval phase. For example, an approval policy can define the following phases:

- A preprovisioning approval phase must be complete before provisioning a requested item
- A postprovisioning phase occurs after the item is already provisioned but before it is released to the owner

An approval policy in the Draft status can be edited but after it is activated, it becomes read-only. A policy with an Inactive status is also read-only.

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 will 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.

System administrators control the default branding for all tenants. A tenant administrator can change the branding of the portal including the 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.

Extensibility

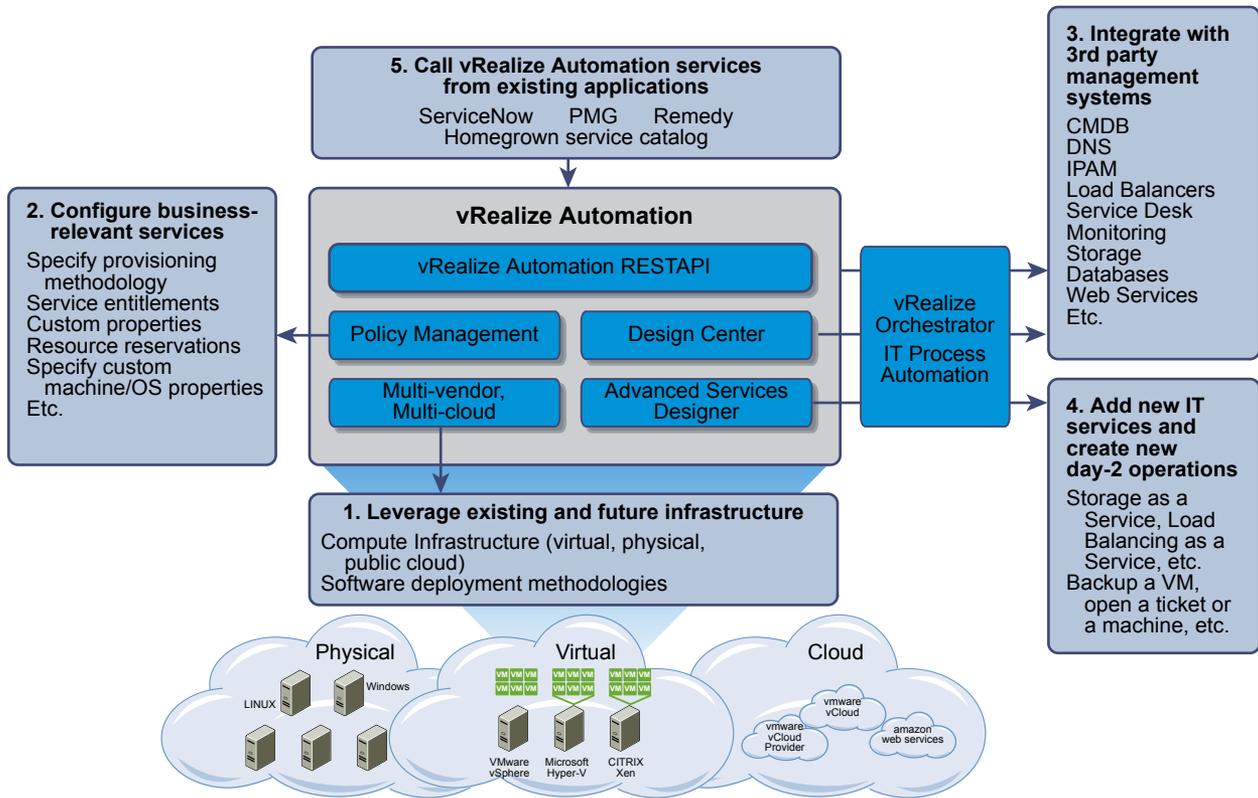
The architecture of vRealize Automation is designed with extensibility in mind. To satisfy different extensibility use cases, vRealize Automation offers a variety of configuration options and tools.

This chapter includes the following topics:

- [vRealize Automation Extensibility Options](#)
- [Leveraging Existing and Future Infrastructure](#)
- [Configuring Business-Relevant Services](#)
- [Integrating with Third-Party Management Systems](#)
- [Adding New IT Services and Creating New Actions](#)
- [Calling vRealize Automation Services from External Applications](#)
- [Distributed Execution](#)

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.

Infrastructure 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 the *vRealize Automation Support Matrix*. For information about configuring infrastructure blueprints, see the *IaaS Configuration* documentation for the relevant machine type.

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 *Tenant Administration* or the *IaaS Configuration* documentation for the relevant machine type.

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 the *IaaS Configuration* documentation for the relevant machine type or the *Custom Properties Reference*.

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 lifecycle extensibility, see *Machine Extensibility*.

Adding New IT Services and Creating New Actions

The Advanced Service Designer enables service architects to define new services and new management operations on provisioned services.

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 service blueprints for storage-as-a-service, networking services or virtually any kind of IT service by using the Advanced Service Designer.

For details about how to create new catalog items or menu options, see *Advanced Service Design*.

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 Advanced Service Designer, are automatically exposed through the vRealize Automation API. For more details, see the *REST API Reference*.

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 *Machine Extensibility* .