

What's New Guide

vCloud Automation Center 5.1

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

EN-001066-00

vmware[®]

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

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

The VMware Web site also provides the latest product updates.

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

docfeedback@vmware.com

Copyright © 2012 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

VMware, Inc.

3401 Hillview Ave.

Palo Alto, CA 94304

www.vmware.com

Contents

About the What's New Guide	6
Intended Audience	6
Chapter 1 What's New for Administrators	7
vCAC Installation Changes	7
Changes to HTTP/HTTPS Configuration	7
Other Installation Changes	7
Interoperability Changes	8
Changes to Email Configuration for vCAC	9
Updating Links in Email Templates	9
Distributed Execution Enhancements	10
The DEM Orchestrator	10
Multiple DEM Instances on the Same Machine	10
Changes to Skills	11
RunOneOnly Workflows	11
DEM Status Updates	11
Workflow Execution History	11
DEM Installation Changes	12
Viewing Distributed Execution Status and Workflow History	12
Skill Management Overview	14
Data Collection	16
Using Data Collection	16
vCAC Audit Log	17
Using the Audit Log Viewer	17
Enhanced Bulk Operations Client (BOC)	18
Chapter 2 What's New for Provisioning and Managing Machines	19
vCloud Director Integration	19
vCloud Director Overview	19
Preparing for vCloud Director Provisioning	20
Managing vApps in vCAC	24
Multi-Machine Services	25
About Multi-Machine Services	25
Configuring Multi-Machine Services	27
Requesting a Multi-Machine Service	31
Managing Multi-Machine Services	34
About Notification Emails and Multi-Machine Services	37
Differences in Multi-Machine Services Between the Self-Service Portal and the vCAC Console	38
Reconfigure Machine	38
About Reconfiguring Machines	38
Configuring the Reconfigure Machine Feature	40
Reconfiguring a Machine	43
Troubleshooting Reconfigure Machine	47

Changes to Reconfigure since 4.1 SP1 FP1	47
Using Multiple Datastores for Machine Volumes	48
Creating a Storage Reservation Policy	49
Assigning Storage Reservation Policies to a Compute Resource	49
Using Disk Properties with Volumes	50
Using Storage Reservation Policies in a Blueprint	52
Requesting a Machine with Storage Reservation Policies	54
Editing a Machine with Storage Reservation Policies	54
Viewing the Machine's Storage	55
Changes to Multiple Datastores since 4.5	55
Changes to vSphere and SCVMM Blueprints	56
SCVMM Integration	56
SCVMM Requirements	56
SCVMM Endpoints	57
Creating an SCVMM Endpoint	57
Using the Compute Resources List	58
Configuring SCVMM	58
Virtual Provisioning on SCVMM	59
Amazon Elastic Compute Cloud (EC2) Enhancements	60
Provisioning Workflow	60
Endpoints	61
Regions	62
Key Pairs	63
Virtual Private Cloud	66
Security Groups	67
Elastic Load Balancers	69
Elastic IP Addresses	72
Elastic Block Storage	74
Amazon Image (AMI) Selection	77
Property Dictionary	77
Creating Property Definitions	78
Configuring Common Property Attributes	81
Defining Relationships between Properties	82
Displaying a Property on the Confirm Machine Request Page	84
About Control Layouts	85
Examples of Using the Property Dictionary	86
Reservation Alerts	92
Snapshot Manager	93
Enabling and Disabling Snapshot Manager	94
Prerequisites	94
Working with Snapshots	94
Linked Clones	96
Working with Linked Clones	97
About Global Machine Limits for a Blueprint	98
Custom Properties	98
Performing Operations on Multiple Machines	99
Gugent	100

Chapter 3	What's New for Extensibility	101
	vCenter Orchestrator Integration	101
	Activity Library	101
	Productivity Tools	101
	CloudUtil Command-Line Tool	101
	Global changes	101
	Agent (DEM) Commands	101
	Client Commands	102
	Events Commands	102
	File Commands	103
	Model Commands	103
	Operation Commands	103
	Package Commands	104
	Schedule Commands	104
	Workflow Commands	104
	Workflow Generator Plugin for Visual Studio	106
	Documentation	106

About the What's New Guide

The *vCloud Automation Center What's New Guide* describes new features and enhancements to existing functionality for vCloud Automation Center™ (vCAC). It covers functionality that was introduced in release 5.1 and the previous release, 4.5.

Intended Audience

This guide is intended for several audiences:

- **Chapter 1, What's New for Administrators** – Focuses on new and enhanced features for installation, deployment, tracking, and data collection. The audience is system administrators.
- **Chapter 2, What's New for Provisioning and Managing Machines** – Describes system enhancements and new options for the overall provisioning process and machine requirements. The audience is users who manage, support, and request machines.
- **Chapter 3, What's New for Extensibility** – Explains updates to the VMware extensibility toolkits. The audience is system administrators and developers who use the extensibility toolkits to customize vCAC.

This guide complements the *vCloud Automation Center Operating Guide* which contains information about vCAC features and functionality that existed prior to the 4.5 release. The information in the *vCloud Automation Center Operating Guide* is current and provides details on the fundamentals of using vCAC.

For example, a reservation is and always has been required to provision a machine, so you would use the *vCloud Automation Center Operating Guide* to learn about a reservation. The storage reservation policy was introduced in release 4.5 and enhanced in release 5.1, so you would use the *vCloud Automation Center What's New Guide* to learn about how to use storage reservation policies.

Chapter 1 What's New for Administrators

vCAC Installation Changes

This section summarizes changes to the installation of vCAC. For details about installing vCAC, see the *vCloud Automation Center Installation Guide*.

Changes to HTTP/HTTPS Configuration

HTTPS is now the default transport protocol for installing all vCAC components.

As a result of this change, fields in the secondary vCAC installers (for example, the DEM, Self-Service Portal, and agent installations) to specify the host name and port of various vCAC components that were previously populated with a default value for an HTTP installation must now be specified manually for all configurations.

In an HTTPS installation, all host names must be specified as fully qualified domain names; for an HTTP installation, the short name or "localhost" is acceptable. The default port numbers for the main vCAC components are as follows:

Component	Installers	HTTPS Port	HTTP Port
Manager Service	DEM, Self-Service Portal	443	9003
Model Manager Web Service	DEM, Self-Service Portal, Agent, Bulk Operations Client, vCAC Designer and vCAC Development Kit	443	80
vCloud Automation Center	Agent only	443	80

Other Installation Changes

The following changes were made to installation in release 5.1:

- SMTP server configuration now supports connections using basic authentication over SSL or without SSL. Anonymous connections are still supported. For additional information about email configuration in release 5.1, see [Changes to Email Configuration for vCAC](#).
- The Prerequisite Checker can now verify prerequisites for websites other than the Default Web Site, including sites that do not run on port 80.

The following changes were made to the installation and deployment in release 4.5:

- The installer executables are now digitally signed.
- The installer now offers a new option for deploying the vCAC database in which the database schema is installed within an existing empty database.
- The sqlcmd.exe tool is no longer required when creating the vCAC database using the vCAC installer.
- The web components can now be deployed within a website in IIS other than the default website.

- The Manager Service can now be installed as a “disaster recovery cold standby node,” which sets vCloud Automation Center service startup type to Manual instead of Automatic.

Interoperability Changes

This section summarizes changes in interoperability support. For full details about interoperability, including deprecated platform support, see the *vCloud Automation Center Support Matrix*.

Support for the following platforms was added in release 5.1:

- SQL Server 2012 and SQL Server 2012 Express
- Provisioning with VMware vCenter Server 5.1 and ESXi 5.1
- Provisioning with vCloud Director 5.1
- Provisioning with SCCM 2012
- The following guest operating systems:
 - Red Hat Enterprise Linux 5.7, 5.8, and 6.3
 - VMware ESXi 5.1
- Ability to use the vCAC Console and Self-Service Portal in the following browsers:
 - Firefox 15
 - Chrome 22
 - Safari 5.1.7
 - Safari 6.0 (vCAC Console only)

Support for the following platforms was added in release 4.5:

- Enhanced Hyper-V support through SCVMM 2012
- Enhanced Amazon EC2 support
- Support for Dell iDRAC 7
- Support for HP iLO4
- Provisioning on Citrix XenServer 6.0.2
- Provisioning integration with Citrix Provisioning Server 6.0
- OnTAP version 8.1 for Netapp FlexClone-enabled storage
- Deployment of Red Hat Enterprise Linux 6.2 and SuSE 11 SP2 on virtual and physical machines
- Ability to use the vCAC console in the following browsers:
 - Firefox 11.0
 - Chrome 18
 - Safari 5 & 5.1.5

Changes to Email Configuration for vCAC

In previous releases, configuration (such as SMTP server information) for sending emails from vCAC was managed in several locations, including the Manager Service configuration and Distributed Execution Manager configuration. Beginning in release 5.1, email configuration is initially specified during the Model Manager installation and managed in global properties.

The following properties can be edited on the **vCAC Administrator > Global Properties** page:

Property	Description
Enable SSL	If true, vCAC communicates with the SMTP server using SSL.
SMTP Server	The host name or IP address of the outgoing SMTP server.
SMTP Port	The SMTP port on the server specified in the SMTP Server property. The default is 25.
From Address	The address with which to populate the "From" field of the email.
User Name	The user name with which to authenticate with the SMTP server.
Password	The password of the user with which to authenticate with the SMTP server.

The corresponding attributes have been removed from the Manager Service and Distributed Execution Manager configuration. Email configuration is no longer part of the DEM Worker installation.

Previously, the root URL used for links in notification emails was specified in the Manager Service configuration and was based on the Manager Service host name by default. The Model Manager installation now prompts for the web server host name and uses this to build the base URL for email links. This value can be updated post-installation by editing the **Website URI** global property.

The product name included in the subject line of all emails is now specified in the **Product Name** global property instead of the Manager Service configuration. The default product name is "vCAC."

Email profiles, introduced in the feature pack 4.1 SP1 FP1, are no longer supported for email configuration. Notification emails for the reconfigure machine feature use the same configuration as all other emails.

For more information about customizing vCAC emails, see the *vCloud Automation Center Operating Guide*.

Updating Links in Email Templates

If you have customized any email templates that use the **ServiceUriBase** parameter, you should update your templates to use the new **WebsiteURI** parameter, which is based on the **Website URI** global property.

The **ServiceUriBase** parameter was based on the legacy **websiteUri** attribute specified in the **ManagerService.exe.config** file. (If no value was specified for **websiteUri**, the Manager Service host name was used.) This parameter still exists in release 5.1 but is deprecated and will be removed in a future release.

The value of the **ServiceUriBase** parameter was in the format **http://example.com** while the **WebsiteURI** parameter includes the base path: **http://example.com/DCAC/**. When updating email templates, you should remove the **/DCAC/** from URLs to ensure that the links are properly formed.

Distributed Execution Enhancements

The Distributed Execution Manager has undergone an architectural overhaul in release 4.5 for improved scalability and reliability. These changes also enable more robust handling of both short- and long-term loss of connectivity between the DEM and the Model Manager. The architectural updates and the new features of the DEM are described in the following sections:

- The DEM Orchestrator
- Multiple DEM Instances on the Same Machine
- Changes to Skills
- RunOneOnly Workflows
- DEM Status Updates
- Workflow Execution History
- Installation changes

The DEM Orchestrator

Each DEM instance now performs one of two roles, Worker or Orchestrator. The Worker role is responsible for executing workflows as in previous releases of vCAC. The new Orchestrator role is responsible for monitoring DEM Worker instances, pre-processing workflows for execution, and scheduling workflows.

The DEM Orchestrator performs the following tasks:

- Monitors the status of DEM Workers and ensures that if a Worker instance stops or loses its connection to the Model Manager, its workflows are put back in the queue for another DEM Worker to pick up.
- Manages scheduled workflows by creating new workflow instances at the scheduled time. This means that scheduled workflows are now handled in the same way as non-scheduled workflows, giving the administrator more visibility into the execution of scheduled workflows.
- Ensures that only one instance of a particular scheduled workflow is running at a given time.
- Pre-processes workflows before execution, including checking preconditions for workflows (used in the implementation of the “RunOneOnly” feature) and creating the workflow execution history.

Exactly one DEM Orchestrator instance is designated as the active Orchestrator that performs the tasks listed above. The DEM Orchestrator is essential for the execution of workflows and therefore VMware strongly recommends that you install at least one additional Orchestrator instance on a separate machine for failover in the event that there is an issue with the machine on which the DEM Orchestrator is running or the Orchestrator loses its connection to the Model Manager. If a DEM Orchestrator is not the currently active Orchestrator, it monitors the active Orchestrator’s status so that it can take over as the active Orchestrator should the currently active Orchestrator go offline.

Multiple DEM Instances on the Same Machine

Beginning in version 4.5, multiple DEM instances (whether they are Workers or Orchestrators) can be installed on the same machine.

Each DEM installation has an instance name that must be unique across the entire system. The installer prompts you for this name, which is used in the following places:

1. The DEM service is named **VMware DEM — <Role> — <InstanceName>**.
2. The installation directory where the DEM is installed is named after the instance name. In a typical installation, this is **%SYSTEM_DRIVE%\Program Files (x86)\DynamicOps\Distributed Execution Manager\<InstanceName>**.
3. In the vCAC Console, it appears as the DEM name on the Distributed Execution Status and Workflow History pages.
4. It also appears in the vCloud Automation Center Designer (vCAC Designer) interface for managing skills and in the CloudUtil DEM commands.

Changes to Skills

Previously, when a skill was added or removed from a DEM the change would not take effect until after the DEM was restarted. Skill associations are now updated immediately without the need for a DEM restart.

In addition, the vCAC Designer console now provides a graphical user interface for managing skills. For more information about the new tooling for skills, see [Skill Management Overview](#).

RunOneOnly Workflows

When installing a workflow, you can now specify that it is a RunOneOnly workflow.

During the pre-processing stage, the DEM Orchestrator checks whether a workflow is designated as RunOneOnly. If so, it ensures that only one instance of that workflow is running at any one time.

DEM Status Updates

The Distributed Execution Status page now reflects the new DEM roles and displays the current list of scheduled workflows.

In addition, the Distributed Execution Status page now includes a column for the last completed workflow date and time. This enables the vCAC Administrator to see not only that a DEM is online, but also when last it completed executing a workflow.

Workflow Execution History

The execution of all workflows is now recorded and a Workflow History page has been added to the vCAC Console. The Workflow History page enables the vCAC Administrator to view all the workflows that have been executed. The information that it displays for each executed workflow includes:

- The DEM that executed it
- The created time, claimed time (when the DEM Worker picked up the workflow), and completed time
- Whether it was scheduled and any applicable schedule
- Whether the execution was successful (with a link to view the related error information if it failed)

The Workflow History page enables the administrator to perform such tasks as:

- Checking the outcome of all executed workflows

- Checking when scheduled workflows were run
- Monitoring each DEM to see which workflows they have executed
- Calculating how long it took for a workflow to execute

DEM Installation Changes

Both the DEM Worker and DEM Orchestrator applications are installed using the same installer. When you run the DEM installation program, you are prompted to select the role for the instance you are installing.

At least one DEM instance must be present for each role. Additional DEM instances can be added for redundancy and scalability, including multiple instances on the same machine.

The DEM Orchestrator should be installed on a machine with strong network connectivity to the Model Manager host (typically the same machine as the Manager Service). VMware strongly recommends that you install a second DEM Orchestrator on a different machine for failover. DEM Workers can be installed on any machine in your deployment architecture. As with the DEM Orchestrator, it is important for DEM Workers to have network connectivity to the Model Manager host. Some DEM Worker instances may have additional requirements depending on the types of workflows that they are intended to execute.

For details about installing DEM Worker and DEM Orchestrator instances, see the *vCloud Automation Center Installation Guide*.

Viewing Distributed Execution Status and Workflow History

The Distributed Execution Status and Workflow History pages display information about the status of DEMs and the workflows that they execute. Together they provide a complete picture of workflow execution.

Viewing the Status of Distributed Execution Managers

The Distributed Execution Status page displays information about each Distributed Execution Manager, including its name, status, and the number of workflow instances it is currently executing.

To view the status of Distributed Execution Managers:

1. Log on to the vCAC Console as a vCAC Administrator.
2. On the **vCAC Administrator** menu, click **Distributed Execution Status**.

The Distributed Execution Status page displays a list of all registered Distributed Execution Managers and their status: **Online** or **Offline**. The status **Online (Active)** on a DEM Orchestrator indicates that it is the Orchestrator instance that is currently orchestrating workflow execution. It also displays the following information:

- For each DEM instance, the name of the host on which it is running, so that you can locate and restart a DEM that is offline.
- For each DEM Worker, the number of workflow instances it is currently executing and the skills that are associated with that instance. Skills restrict which DEM Workers can execute which workflows.
- A list of all schedules installed in the Model Manager, including their most recent and next scheduled runs.

Viewing Current Workflow Execution Status

You can view the number of workflows currently being executed on the Distributed Execution Status page. You can also select the name of a DEM Worker to view the list of workflows currently being executed by that DEM, or view the list of all workflows currently being executed by all DEMs.

A workflow may have a status Executing or Pending. An executing workflow is currently being executed by a DEM Worker. A pending workflow has not yet been picked up by a DEM Worker and may be in one of the following states:

- Waiting to be preprocessed by the DEM Orchestrator.
- Preprocessed and is ready to be picked up by a DEM Worker.
- Delayed to allow for a similar workflow to complete.

To view the status of currently running workflows:

1. Log on to the vCAC Console as a vCAC Administrator.
2. On the **vCAC Administrator** menu, click **Distributed Execution Status**.
3. Do one of the following:
 - To view the status of all currently running workflows, click **View Workflows**.
 - To view the status of workflows being executed by a specific DEM Worker instance, click the name of the DEM Worker.

To return to the list of all currently running workflows, point to **(unnamed) Filter** and click **Clear**.

After a workflow has finished executing, its details, status, and the result of its execution appear on the Workflow History page.

Viewing Workflow History

The Workflow History page displays a list of every workflow that has been run by a Distributed Execution Manager.

The Workflow History list includes only workflows that have completed. You can view information about currently running workflows on the Workflows page, accessible via the Distributed Execution Status page.

Actions performed by vCAC Agents are not included on the Workflow History page.

To view information about completed workflows:

1. Log on to the vCAC Console as a vCAC Administrator.
2. On the **vCAC Administrator** menu, click **Workflow History**.

The Workflow History page displays the following information:

- A list of all completed workflows and their status: **Succeeded**, **Failed**, or **Stopped**.
- A workflow may be stopped if the DEM Worker that was executing it stops or loses its connection to the Model Manager. Stopped workflows are re-executed when a DEM Worker becomes available to pick up the workflow.
- In the case of failed workflows, the **Result Details** field displays a brief exception message.
- The **Multiple Attempts?** field indicates when a workflow was stopped and required

multiple attempts to complete successfully.

3. To view detailed information about a workflow instance, click the name of the workflow.
 - In the case of a workflow that has been stopped and restarted, the Workflow Execution Details page displays a row for each attempt to run this workflow.
 - In the case of a failed workflow, the Workflow Execution Details page provides a link to the associated log entry.

Skill Management Overview

Skills enable you to create a mapping between specific workflows and the DEM Workers that can execute them.

A skill is similar to a tag that you can apply to both workflows and DEM Worker instances. If a workflow is not associated with any skills, any DEM Worker can execute it. If a workflow is associated with one or more skills, then only DEM Workers that are associated with all of the same skills can execute it.

Skills are useful when a particular workflow requires a DEM Worker installed on a host with specific prerequisites. For example, you may want cloud provisioning workflows to execute only on a DEM Worker running on a host with the required network access to Amazon URLs.

Skills can also associate workflows with a particular data center location. For example, you might install one DEM Worker in your Boston data center and another in your London data center, and use skills to direct certain vCAC operations to one data center or the other.

This section describes how to manage skills using the vCAC Designer console. You can also use the CloudUtil command-line tool to add skills to the Model Manager and associate them with workflows and DEM Workers. For information about the CloudUtil syntax, see the *vCloud Automation Center Extensibility Guide*.

Associating Workflows and DEM Workers Using Skills

You associate workflows with a specific DEM Worker or set of Worker instances by adding a skill to the Model Manager and then associating the skill with one or more workflows and DEM Workers.

To associate workflows with DEM Workers:

1. Launch the vCAC Designer console.
2. On the ribbon, click **Manage Skills**.
3. In the text field at the upper left of the **Manage Skills** dialog, type the name of a new skill and click the **Add** button .

The skill name must be unique. If the name of the new skill matches the name of an existing skill, the **Add** button is unavailable.
4. Select the name of the skill in the list on the left.
5. Associate the skill with one or more DEM Workers:
 - a. Click the **Add** button next to **Distributed Execution Managers**.
 - b. In the **Select DEMs** dialog, select one or more DEM Worker instances.
 - c. Click **OK**.
6. Associate the skill with one or more Workflows:
 - a. Click the **Add** button next to **Workflows**.

- b. In the **Select Workflows** dialog, select one or more workflows.
- c. Click **OK**.

The workflows associated with this skill can only be executed by the DEM Workers that are associated with this skill.

7. When you are done adding skills and associating them with DEM workers and workflows, click **OK** to close the **Manage Skills** dialog and save your changes to the Model Manager.

Removing Associations between Skills and DEM Workers

When you remove the association between a skill and a DEM Worker, that Worker instance can no longer execute the workflows associated with the skill.

To remove an association between a skill and DEM Workers:

1. Launch the vCAC Designer console.
2. On the ribbon, click **Manage Skills**.
3. In the **Manage Skills** dialog, select a skill from the list on the left.
4. Select the name of one or more DEM Worker instances from the **Distributed Execution Managers** list and click the **Remove** button .
5. Click **OK** to close the **Manage Skills** dialog and save your changes to the Model Manager.

Removing Associations between Skills and Workflows

When you remove the association between a skill and a workflow, that workflow is no longer restricted to the DEM Workers that are associated with the same skill.

To remove an association between a skill and workflows:

1. Launch the vCAC Designer console.
2. On the ribbon, click **Manage Skills**.
3. In the **Manage Skills** dialog, select a skill from the list on the left.
4. Select the name of one or more workflows from the **Workflows** list and click the **Remove** button .
5. Click **OK** to close the **Manage Skills** dialog and save your changes to the Model Manager.

Removing a Skill

Removing a skill using vCAC Designer also removes its associations to any DEM Workers and workflows.

To remove a skill:

1. Launch the vCAC Designer console.
2. On the ribbon, click **Manage Skills**.
3. In the **Manage Skills** dialog, select a skill from the list on the left.
4. Click the **Remove** button  at the top of the list of skills.

After you confirm that you want to delete the skill, its name appears dimmed to indicate that it is marked for deletion.

5. Click **OK** to close the **Manage Skills** dialog and save your changes to the Model Manager or **Cancel** if you do not want to delete the skill and its associations with DEMs and workflows.

Data Collection

vCAC regularly collects data from all compute resources that belong to enterprise groups. Data collection updates information about a compute resource or endpoint and provides details about the machines provisioned on it. In the case of compute resources, this includes unmanaged virtual machines (provisioned outside of vCAC).

There are two types of data collection: inventory and state data collection.

- **Inventory** data collection on a compute resource updates the vCAC record of the virtual or cloud machines whose resource usage is tied to that compute resource.
- **State** data collection on a compute resource updates the vCAC record of the power state of each machine discovered through inventory data collection.

vCAC 4.5 introduced new functionality to manage and use data collection.

Using Data Collection

An EA can enable or disable Inventory and State data collection and set the frequency of Inventory data collection in hours and the frequency of State data collection in minutes.

To manage data collection:

- Point to the compute resources on the Compute Resources page (Enterprise Administrators) and click **Data Collection**.

The Data Collection page displays:

- The time of the last data collection from the compute resource
- Whether the last data collection succeeded or failed
- Whether data collection is currently running

To initiate data collection:

- Click **Request Now** to initiate data collection.

Although data collection can be set to run at regular intervals, you may want to work with the most recent data by initiating a data collection.

The time required for data collection to complete depends on several factors, including the capacity of the compute resource and the number of machines on the compute resource, the number of compute resources in an enterprise group, current system and network load, among other variables.

When data collection is complete for a compute resource, the **Status** label is updated with the result which can be success or failure and the **Last Completed** label is updated with the date and time.

When data collection is complete for an endpoint, the **Status** label is updated with the result which can be success or failure and the date and time.

By default, vCAC limits concurrent data collection activities to two per agent and queues requests that are over this limit. This allows data collection activities to complete quickly while not affecting overall performance. It is possible to raise the limit to take advantage of concurrent data collection but this should be weighed against any degradation in overall performance.

If you do increase the configured vCAC per-agent limit, you may want to increase one or more of these execution timeout intervals. See the *vCloud Automation Center Operating Guide* for more information about establishing and increasing timeout intervals.

vCAC Audit Log

vCAC Logs track important system information and interactions. vCAC 4.5 added the ability for a vCAC Administrator to audit and review all activities performed on virtual machines. The Audit Log provides details about activities that were performed successfully and activities that failed.

The Audit Log Viewer provides details about actions performed by instance types including Amazon EC2 cloud machines and snapshots of machines created with Snapshot Manager.

Using the Audit Log Viewer

The Audit Log provides an audit trail of all actions performed on virtual machines.

To review all messages from the Audit Log tracking the activities of virtual machines:

- Select **vCAC Administrator > Audit Log Viewer**.

By default the Audit Log Viewer provides details including Time, Action, Source, Entity, Entity Instance, and Message.

Other information can be displayed by adding Column headers to the table.

To add columns:

- Click the Columns drop-down list and click to display the following additional information:

Field	Description
Log ID	Display a unique ID for the log.
Time	Displays the date and time when the action occurred in the format MM/DD/YYYY 12-hour clock.
Action	Describes the action. For example, Amazon Instance Destroyed or Snapshot Created, Snapshot Applied.
Source	Identifies the component that performed the action. For example, an action taken by the DEM Worker displays DEM Worker.
Entity Instance ID	Displays a unique ID for the entity instance.
Entity Instance	Identifies the instance that the action was performed on.
Username	Display the username of the user who performed or requested the activity.
Message	Describes the details of the action taken. For example, Amazon Instance (name) Destroyed.
Properties	Displays properties associated with the activity.

To review all messages from all vCAC logs since the vCAC instance was installed:

- Click **vCAC > vCAC Administrator > Log Viewer**.

To view an audit of actions performed on managed machines:

- Click **vCAC > vCAC Administrator > Audit Log Viewer**.

Enhanced Bulk Operations Client (BOC)

Enhancements to the Bulk Operations Client (BOC) in release 4.5 included support for Physical machines through an easy to use interface and overall improvements to performance. The BOC client can be installed on any host with access to the vCAC server and allows Enterprise Administrators to use the options on multiple machines simultaneously. For more details about the BOC, see the *vCloud Automation Center Operating Guide*.

In release 5.1, the multi-machine service feature introduces another way to manage several machines as a group. See [Multi-Machine Services](#).

Chapter 2 What's New for Provisioning and Managing Machines

vCloud Director Integration

This section describes the integration between vCAC and VMware vCloud® Director™, which is new in release 5.1.

vCloud Director Overview

The vCloud Director integration in vCAC enables a user to provision vApps in a virtual datacenter (VDC) based on an existing vApp template, bringing the software-defined services of vCloud Director together with the user-centric governance of vCAC. It also enables vCAC to serve as a single self-service portal in a heterogeneous environment with multiple provisioning platforms.

The following table shows the mapping between entities in vCloud Director and vCAC:

vCloud Director	vCAC	Notes
vCloud Director server	endpoint	A vCAC endpoint can be mapped to a vCloud Director server or the combination of a vCloud Director server and an organization, which provides access only to the relevant organization vDCs.
organization	provisioning group	User roles and membership in vCloud Director organizations and vCAC provisioning groups are maintained independently in each tool. See vCloud Director Requirements .
organization virtual datacenter	compute resource	vCAC enforces a one-to-one mapping between a compute resource that represents a VDC and a reservation.
vApp template	virtual blueprint	A vCAC blueprint for vCloud Director must specify a vApp template from which to clone a vApp. A blueprint in vCAC that is specific to a provisioning group is analogous to a template that is part of an organization's catalog, while global blueprints are analogous to templates in a public catalog.
vApp	virtual machine	vApps are managed as machines in vCAC even though they are containers for virtual machines in vCloud Director.

vApps are one way to provision several machines with a single request. vCAC also has a multi-machine service feature that provides similar functionality. For a feature comparison between vCloud Director integration and multi-machine services, see [About Multi-Machine Services and vCloud Director Integration](#).

vCloud Director Requirements

Before you integrate vCloud Director with vCAC, you must install and configure your vCloud Director instance including setting up organizations, cloud resources, and vApp templates.

The user roles within a vCloud Director organization do not need to correspond with the roles in the vCAC provisioning group. For example, it is not necessary for the vCAC provisioning group manager to have vApp author permissions in order to configure vCloud Director blueprints in vCAC, unless the

user also needs the ability to create templates in vCloud Director. The only requirement for vCloud Director integration is that any vCAC user who requests a vApp must be a member of the domain associated with the organization VDC on which the vApp is provisioned and must have permissions to log in to vCloud Director Web console.

In addition, you must have configured the cloud resources (such as VDCs and networks) that you will use to provision vApps and also have created vApp templates for any vApps that you want to provision through vCAC.

For information about configuring vCloud Director and working with vApp templates, see the VMware vCloud Director documentation.

Known Limitations

The following features are not supported for vApps in the current release of vCAC:

- Creating new vApps or vApp templates
- Reconfiguring virtual machines in a vApp
- Snapshots of vApps
- Static IP
- Customizing Linux VMs that are part of vApps
- Moving vApps between VDCs

vCloud Director blueprints cannot be added as component blueprints to a multi-machine blueprint.

Preparing for vCloud Director Provisioning

This section describes the tasks that are required to enable provisioning of vApps through vCAC.

The high-level process for configuring vCloud Director provisioning in vCAC is as follows:

1. Create a credential to authenticate with the vCloud Director server.
2. Create a vCloud Director endpoint.
3. Add a VDC to an enterprise group.
4. Create a provisioning group that corresponds to a vCloud Director organization.
5. Create a reservation for each organization VDC.
6. Create a vCloud Director blueprint.

This section only describes procedures that are specific to provisioning with vCloud Director. For general information about configuring vCAC for provisioning, see the *vCloud Automation Center Operating Guide*.

Creating a vCloud Director Credential

You must create a vCAC credential for authenticating with the vCloud Director server. The procedure for creating a credential for vCloud Director is the same as for any credentials in vCAC.

The credentials you provide are typically those of a vCloud Director organization administrator. Alternatively, you can use the credentials for the vCloud Director system administrator.

Note: The credentials that you specify must include the username only and not the domain.

Creating a vCloud Director Endpoint

A vCloud Director endpoint is required to connect to a vCloud Director server.

You can specify the name of a vCloud Director organization when you create an endpoint. You must choose one of the following options:

- Use the credentials for a vCloud Director organization administrator to connect to the vCloud Director server and specify an organization for which the user has the administrator role. The endpoint can only access the associated organization VDCs. You can add endpoints for each additional organization in your vCloud Director instance that you want to integrate with vCAC.
- Use the credentials for a vCloud Director system administrator and leave the **Organization** field empty to enable access to all organization VDCs within the vCloud Director instance.

If you set up multiple organizations (either as separate endpoints or a single endpoint), VMware recommends that you create reservation policies for each organization to ensure that vApps are provisioned to the appropriate VDCs.

You cannot set up both organization-specific endpoints and global endpoints for the same vCloud Director instance.

Note: A vCloud Director instance is backed by vSphere resources that provide the CPU, memory, and storage for its virtual machines. When vCAC performs data collection on a vCloud Director endpoint, it collects information about the underlying vSphere resources. To avoid conflicts in resource accounting, if you create a vCloud Director endpoint, do not create an endpoint for its underlying vSphere resources.

Although vCAC uses a proxy agent to manage vSphere resources, no agent is needed for vCloud Director.

To create a vCloud Director endpoint:

1. Log in to the vCAC Console as a vCAC Administrator.
2. Select **vCAC Administrator > Endpoints**.
3. Click **New Endpoint > vCloud Director**.
4. Specify a **Name** for the endpoint.
5. Optionally, specify a description of the endpoint. This information displays on the Endpoints page.
6. Specify a URL with the fully qualified host name or IP address of the vCloud Director server, for example: **https://vcd.example.com**.

Note: You must specify **https** as the transport protocol.

7. Specify the credentials to use to connect to this endpoint.
8. Specify the name of an **Organization** to filter the organization VDCs that are available via this endpoint to those that are associated with that organization.

If the credentials you specified are those of a vCloud Director system administrator, you can leave this field empty to enable access to all VDCs on the vCloud Director server.

9. Optionally, specify custom properties to apply to all machines provisioned using this endpoint.
10. Click **OK**.

Adding a VDC to an Enterprise Group

You must add each organization VDC to an enterprise group to make it available for provisioning. The procedure is the same for configuring any compute resource with an enterprise group.

Before you add a VDC to an enterprise group, ensure that you have run data collection on a vCloud Director endpoint. Data collection runs automatically when you create a new endpoint. You can also request data collection manually for a vCloud Director endpoint from the Endpoints page.

Mapping a Provisioning Group to a vCloud Director Organization

A vCAC provisioning group corresponds to an organization in vCloud Director. The procedure for setting up a provisioning group for vCloud Director is the same as for any provisioning group.

The user roles for vCAC and vCloud Director are maintained independently in each tool. Ensure that any user who requests a vApp also has permissions to log in to vCloud Director Web console. Console Access Only permissions in the appropriate organization are sufficient to provision a machine through vCAC.

Creating a vCloud Director Reservation

Only one vCAC reservation can be created for each vCloud Director compute resource or organization VDC.

To create a vCloud Director reservation:

1. Log in to the vCAC Console as an enterprise administrator.
2. Select **Enterprise Administrator > Reservations**.
3. Click **New Reservation > Virtual**.
4. On the **Reservation Information** tab:
 - a. From the **Compute resource** list, select the name of the organization VDC.
 - b. Specify a **Name** of the reservation, or accept the auto-generated name.
 - c. Select the **Provisioning group** that corresponds to an organization that has access to this VDC in vCloud Director.
 - d. Optionally, specify a **Reservation policy** to apply to this reservation.
 - e. Optionally, specify the **Machine quota** for this reservation. This quota limits the number of vApps that can be provisioned on this reservation, each of which may contain one or more virtual machines. VMware recommends leaving this field empty for a vCloud Director reservation.
 - f. Specify a priority for this reservation.
5. On the **Resources** tab, select the available network paths for this reservation.

The values for Memory and Storage are automatically populated at the maximum available for each resource on the VDC. All storage profiles are automatically selected for use in the reservation.

To specify the storage cost for this reservation, set up a cost profile and apply it to the compute resource representing this VDC. For information about working with storage cost profiles, see the *vCloud Automation Center Operating Guide*.

Note: You can configure the maximum available Memory and Storage for the VDC in the vCloud Director Web Console. After you have made the change in vCloud Director, run data collection on the corresponding endpoint in vCAC to update the reservation.

6. Optionally, on the **Alerts** tab, configure capacity alerts for this reservation.
7. Click **OK**.

Creating a vCloud Director Blueprint

To enable users to provision vApps in vCAC, you must create a blueprint for cloning vApps from an existing vApp template.

The blueprints that are available to members of a provisioning group are analogous to catalog items that are available to members of a vCloud Director organization. Public catalog items in vCloud Director are similar to global blueprints in vCAC.

To create a vCloud Director blueprint:

1. Log in to the vCAC Console as an enterprise administrator or provisioning group manager.
2. Select **Enterprise Administrator > Global Blueprints** or **Provisioning Group Manager > Blueprints**.
3. Click **New Blueprint > Virtual**.
4. On the **Blueprint Information** tab, specify blueprint details as you would for any blueprint.

Note: The machine prefix that you specify for the blueprint is used for all vApps provisioned from this blueprint as well as all the virtual machines that are contained in the vApps.

5. On the **Build Information** tab:
 - a. From the **Platform type** list, select **vCloud Director**.
 - b. Select the **Blueprint type**.
 - c. The **Provisioning workflow** must be **CloneWorkflow**.
 - d. In the **Clone from** field, click the ellipsis to select a vApp template.

Note: If you have multiple endpoints for different organizations in the same vCloud Director instance and there are vApp templates that are shared between organizations as part of a published catalog, the shared templates appear once in the list for each endpoint or organization. You can filter the list of blueprints by endpoint to ensure that you are choosing from templates that are available to a particular organization. Optionally, you can save the filter for later use.

The **Machine resources** section is populated with values representing the total resources consumed by all virtual machines in the vApp template. The only setting that can be edited in vCAC is the duration and approval threshold for the lease.

6. Optionally, specify custom properties and security options.

7. Click **OK**.

Enabling a vApp to join an Active Directory Domain

If you want a vApp to join an Active Directory domain, you must specify certain custom properties on the blueprint. Joining an Active Directory domain is only supported for Windows machines.

To enable a vApp and its virtual machines to join an Active Directory domain:

1. Log in to the vCAC Console as an enterprise administrator or provisioning group manager.
2. Select **Enterprise Administrator > Global Blueprints** or **Provisioning Group Manager > Blueprints**.
3. Create a new blueprint or edit an existing vCloud Director blueprint.
4. On the **Properties** tab, add the following custom properties (case sensitive):

Property	Value
Sysprep.Identification.DomainAdmin	The username of an administrator of the target domain in Active Directory. Do not include the user domain in the credentials that you send to vCloud Director.
Sysprep.Identification.DomainAdminPassword	The password for the user specified in Sysprep.Identification.DomainAdmin .
Sysprep.Identification.JoinDomain	The name of the domain in Active Directory to join.

5. Click **OK**.

About Leases and vApps

Both vCAC and vCloud Director provide the ability to specify a lease on vApps. For vApps provisioned through vCAC, the lease specified in vCAC overrides any lease specified in vCloud Director.

To avoid conflicts in resource accounting, specify a lease period of **Never Expires** for vApps within an organization whose vApps can be managed by vCAC.

Managing vApps in vCAC

A vApp is treated as a single entity in vCAC even though it may contain multiple virtual machines. The status displayed for the vApp in vCAC corresponds to its status in vCloud Director.

When a user selects an operation from the machine menu for a vApp, such as Destroy, Reprovision, Expire, or any power operation, it is performed on the entire vApp and all its virtual machines.

Viewing the vApp in the vCloud Director Web Console

The vCAC Console does not display details of the virtual machines within a vApp. To view virtual machine details, you open the vApp in the vCloud Director Web console.

Note: If you are requesting a vApp in vCAC for the first time, wait for the vApp to be fully provisioned and powered on before connecting to the vCloud Director Web console.

To view a vApp in vCloud Director:

1. Log in to vCAC.
2. Do one of the following:
 - a. As a machine owner, click **Self-Service > My Machines**.
 - b. As a provisioning group manager, click **Provisioning Group Manager > Group Machines**.
 - c. As an enterprise administrator, click **Enterprise Administrator > Enterprise Machines**.
3. Point to the name of a vApp and click **Connect to Virtual Desktop**.
4. The vCloud Director Web console opens in a new browser window.
5. Log in to the vCloud Director Web console if you do not have an active session.

The vCloud Director page for the corresponding vApp opens.

Multi-Machine Services

This section describes the multi-machine service feature, which is new in release 5.1.

Note: Multi-machine services are sometimes referred to as application services or “AppServices” in some logs and error messages.

About Multi-Machine Services

A multi-machine service is a compound service composed of multiple machines that can be provisioned and managed through vCAC as a single entity.

For example, in a typical web application deployment, you may have one or more database servers, one or more application servers, and any number of web servers. In vCAC, in addition to creating blueprints for each of the server types, you can also create a multi-machine blueprint that includes all the machines that are needed for the entire application deployment. Users can then provision the multi-machine service including all its component machines in a single request and perform operations such as rebooting all the machines in a multi-machine service with a single action.

The component machines in a multi-machine service may be virtual, physical, or cloud, or any combination of the three.

Multi-machine services in vCAC are similar to machines in several ways:

- They are defined by a blueprint.
- Multi-machine service requests can be made subject to approval.
- They may have a lease duration associated with them.
- They are listed on the My Machines, Group Machines, and Enterprise Machines pages, and many machine operations can be performed on multi-machine services as a whole.

However, a multi-machine service is in essence a container for its component machines and is not itself a machine. Component machines are treated just like any machine managed by vCAC for reporting and license purposes, but the multi-machine service container is not counted as a machine in reports or licensing.

Hypervisor and vCloud Director (vApp) blueprints cannot be added as components of a multi-machine blueprint.

About Multi-Machine Services and vCloud Director Integration

vCAC provides two ways to provision one service composed of multiple machines: multi-machine services and support for vCloud Director vApps. Each method has different benefits and capabilities.

The following table provides a comparison of the two features:

vCAC Multi-Machine Service	vCloud Director vApp
Create a multi-machine blueprint that references individual machine blueprints in vCAC.	Leverage existing vApp templates created in vCloud Director.
Provision machines of any type (virtual, physical, or cloud) as part of a service.	Provision virtual machines on vSphere resources as part of a vApp.
Manage component machines of a multi-machine service through vCAC.	Access vApp virtual machines through the vCloud Director Web console.

Terminology

The following table provides definitions of several terms specific to multi-machine services.

Term	Definition
multi-machine service	A compound service composed of multiple machines that can be provisioned and managed through vCAC as a single entity. A multi-machine service is defined by a <i>multi-machine blueprint</i> , which references blueprints for component machines and specifies the number of machines to provision from each component blueprint as part of the multi-machine service.
multi-machine service menu	A context menu, similar to the machine menu that appears when a user points to the name of the multi-machine service on the My Machines, Group Machines, or Enterprise Machines page, enabling the user to perform operations on the multi-machine service (and its component machines) as a whole.
component blueprint	A blueprint that defines a type of component machine, or <i>component type</i> , that is part of a multi-machine service. A component blueprint is referenced by a multi-machine blueprint. The same blueprint may also be available for users to request standalone machines that are not part of a multi-machine service.
component machine	An individual machine instance that is managed as part of a multi-machine service. A multi-machine service may include several machines of the same component type (provisioned using the same component blueprint).

Multi-Machine Service Life Cycle

Multi-machine services follow the same general life cycle as machines, from requesting (with optional approval) and provisioning through managing and decommissioning.

A multi-machine blueprint contains references to blueprints for the component machines. For each component blueprint, it specifies the minimum and maximum number of machines for the multi-machine service. An enterprise administrator or provisioning group manager can optionally specify scripts or workflows to execute during the provisioning process or when powering the multi-machine service on or off.

Note: The scripts execute on the vCAC DEM Worker machine, not on the guest operating system of the component machine.

When users request a multi-machine service, they have the ability to configure:

- How many machines based on each component blueprint to provision, within the limits specified in the multi-machine blueprint
- The specifications of the component machines (such as CPU, Memory, and Storage) based on the individual component blueprints
- General settings such as lease duration and custom properties to be applied to machines in the multi-machine service

Before provisioning the multi-machine service, vCAC allocates resources for all of the component machines that are being requested. If the multi-machine service would cause a reservation to become overallocated, its provisioning fails. Once the resources have been allocated, the component machines are provisioned and powered on.

Once the multi-machine service is provisioned, the machine owner can perform various machine menu tasks on the multi-machine service as a whole, such as powering the multi-machine service off and on, or destroying the multi-machine service and its component machines.

If the multi-machine blueprint allows for a varying number of machines for any component type, the machine owner can add or remove machines from the multi-machine service after it is provisioned.

The machine owner can also view the components that make up a multi-machine service and manage them individually. Most machine operations are available for component machines, except for changing the owner, changing the lease, or expiring component machines.

Configuring Multi-Machine Services

You enable users to provision multi-machine services by creating one or more multi-machine blueprints as an enterprise administrator or provisioning group manager.

Preparing Reservations for Multi-Machine Services

You should create reservations for all the resources that may be used in multi-machine services in the same enterprise group.

To ensure that component machines are provisioned on the appropriate resources, use reservation policies to associate specific reservations with each blueprint that is a component of a multi-machine service. For information about reservation policies, see the *vCloud Automation Center Operating Guide*.

Creating a Multi-Machine Blueprint

A multi-machine blueprint is a specific type of blueprint that contains references to component blueprints.

Before you create a multi-machine blueprint, you should first create blueprints for each of the component machines you want to include in the multi-machine service.

Note: The component blueprints must be available to the same provisioning group or groups as the multi-machine blueprint.

There are some limitations on what you can edit in multi-machine blueprints and their component blueprints after you create them. For details, see [About Managing Multi-Machine Blueprints](#).

To create a multi-machine blueprint:

1. Log in to the vCAC Console as an enterprise administrator or provisioning group manager.

2. Select **Enterprise Administrator > Global Blueprints** or **Provisioning Group Manager > Blueprints**.
3. Click **New Blueprint > Multi-Machine**.
4. On the **Blueprint Information** tab, specify the multi-machine service details as you would for a regular machine.

The multi-machine blueprint only includes information that applies to a multi-machine service as a whole. Configuration options such as reservation policy apply only to component machines and are specified in the component blueprints. Multi-machine services do not have an archive period; the service and its component machines are destroyed when they expire.

5. Add component blueprints to the multi-machine service:
 - a. On the **Build Information** tab, click **Add Blueprints**.
 - b. In the **Add Blueprints** dialog box, select the component blueprints to include in the multi-machine service.

Note: If you are editing a local blueprint, only blueprints in the same provisioning group are available in this list. If you are editing a global blueprint, only global blueprints are available. Hypervisor blueprints and vCloud Director blueprints are not available to add as component blueprints.

- c. Click **OK**.

The **Component machines** grid is populated on the **Build Information** tab.

6. For each component blueprint:
 - a. In **Name**, specify a display name for the component type. Default is the name of the component blueprint.
 - b. Specify a **Minimum** number of machines provisioned from that blueprint to be included in the multi-machine service.

A machine owner cannot request a multi-machine service with fewer than the minimum number of machines for each component type. Additionally, this is the threshold that determines whether a multi-machine service provisioned from this blueprint is healthy. When the number of machines that are provisioned or powered on for a component type is lower than the minimum, the multi-machine service is considered unhealthy and its state is set to **Off**.

- c. Optionally, specify a **Maximum** number of machines to be provisioned from the blueprint in the multi-machine service. If no maximum is specified, then the multi-machine service cannot have more than the minimum number of machines for this component type.
 - d. Optionally, adjust the **Startup Order** and **Shutdown Order** for the machines.

Note: The startup and shutdown order do not apply at provisioning time; they are only used when powering the multi-machine service on or off after the initial provisioning.

- e. **Description** displays the description of the component blueprint.
7. Optionally, specify a lease duration. The multi-machine service lease (whether there is a specific expiration date or no expiration date) overrides the individual component machine leases.

8. Optionally, specify scripts, custom properties, and security options for this blueprint.
9. Click **OK** to save the blueprint.

About Specifying Scripts for Multi-Machine Service Provisioning, Startup, and Shutdown

You can designate scripts or workflows to execute at specific points during the multi-machine service life cycle.

Note: The scripts execute on the vCAC DEM Worker machine, and not on the guest operating system of the component machine.

The **Scripting** tab of the multi-machine blueprint enables you to specify custom logic to execute at each stage:

Stage	Description
Pre-provisioning	Executes after all necessary approvals are complete but before provisioning any machines.
Post-provisioning	Executes after all component machines are provisioned and powered on.
Pre-startup	Executes before powering on the multi-machine service and all its component machines.
Post-startup	Executes immediately after powering on the multi-machine service and all its component machines. The multi-machine service state is set to On after the post-startup scripts execute.
Pre-shutdown	Executes immediately before shutting down or powering off the multi-machine service and its component machines.
Post-shutdown	Executes after shutting down or powering off the multi-machine service and its component machines. The multi-machine service state is set to Off after the post-shutdown scripts execute.

The provisioning scripts execute only during the initial provisioning of a multi-machine service. The startup and shutdown scripts execute every time the multi-machine service is powered on or shutdown (see [Power Options for Multi-Machine Services](#)), with the exception of the first time the multi-machine service is powered on during provisioning. The pre-provisioning and post-provisioning scripts should include everything that you want to run before and after the initial startup.

Scripts or workflows must be installed in the Model Manager before you can use them in a multi-machine blueprint. For details about installing files (including scripts) and workflows using the CloudUtil command-line utility, see the *vCloud Automation Center Extensibility Guide*.

To specify logic to execute during the multi-machine service life cycle:

1. Create a new multi-machine blueprint or edit an existing blueprint.
2. Click the **Scripting** tab.
3. For each stage of the life cycle where you want to execute custom logic:
 - a. Select the name of a script in the Model Manager to execute that script, or select **Workflow** to execute a workflow.
 - b. If you selected **Workflow**, enter the name of the workflow in the Model Manager in the **Workflow** box.

4. Click **OK**.

About Specifying Custom Properties for Multi-Machine Services

You can specify custom properties that apply to all component machines within a multi-machine service in the multi-machine blueprint.

Custom properties specified in the multi-machine blueprint override properties specified in the component blueprint. However, runtime properties on the component machine (specified either at request time or by editing the machine after it has been provisioned) override runtime properties at the multi-machine service level. This enables the multi-machine blueprint to apply a consistent behavior across all of its component machines by default, while allowing a user to override the multi-machine service properties for each component type.

The full order of precedence for custom properties on component machines, including those defined at the multi-machine service level, is as follows (any property value specified in a source later in the list overrides values for the same property specified in sources earlier in the list):

1. Build Profile (specified on component blueprint)
2. Component Blueprint
3. Build Profile (specified on multi-machine blueprint)
4. Multi-Machine Blueprint
5. Provisioning Group
6. Compute Resource
7. Reservations
8. Endpoint
9. Runtime (specified on multi-machine service)
10. Runtime (specified on component machine)

Note: Custom properties specified in the multi-machine service are applied to all component machines within the service. Additionally, if a property is designated as **Prompt User** on a component blueprint, the value specified at request time is applied to all machines of the same component type that are provisioned as part of that request. Certain properties, such as Hostname, must be unique to each machine and therefore should not be specified at the multi-machine service level or at the component type level.

About Specifying Security Options for Multi-Machine Services

The **Security** tab for multi-machine blueprints functions in the same way as the same tab for all blueprints. It enables you to control access to the blueprint and specify which menu operations are available to machine owners for multi-machine services provisioned from this blueprint.

The security settings for the multi-machine blueprint override the access controls on the component blueprints. For example, if you have a component blueprint that is restricted to UserA only and you add it as a component of a multi-machine blueprint that is available to Everyone, then all members of the provisioning group can provision the restricted machine as part of a multi-machine service, but only UserA can provision it as a standalone machine.

The multi-machine service menu is similar to the machine menu, although a different set of operations are available by default. Some operations, such as reprovision, are not supported for multi-machine services. Additional operations are available only for multi-machine services, such as the ability to add component machines.

You can add custom operations to the multi-machine service menu as you can for the machine menu. For information about implementing custom menu operations, see the *vCloud Automation Center Extensibility Guide*.

About Managing Multi-Machine Blueprints

You can edit multi-machine blueprints after they are created; however, there are some special considerations and limitations on the changes you can make to a blueprint.

Once a blueprint has been added as a component blueprint of a multi-machine blueprint, you cannot change the provisioning groups associated with the multi-machine blueprint or the component blueprint.

There are additional restrictions to the configuration of the component machines on the **Build Information** tab if there are existing multi-machine services that were provisioned using a blueprint:

- You cannot add new component types to the multi-machine service. (The **Add Blueprint** link is not available.)
- You cannot delete component blueprints from the multi-machine service.
- When editing component types, the **Minimum** field is read-only.

These restrictions prevent an existing, healthy multi-machine service from changing to an unhealthy state due to changes in the blueprint. If you want to make any of these changes to the component machine configuration in a multi-machine blueprint, either decommission all existing services that use the blueprint or create a copy of the blueprint and update the configuration in the copy.

If you change the **Maximum** value for a component type, existing services that have more component machines of that type than the new maximum do not become invalid, although users cannot add any more machines of that component type until the number of machines falls below the maximum.

When you change custom properties on a multi-machine blueprint, the new properties are not only applied to new multi-machine services that are provisioned from this blueprint, but are also inherited by component machines that are added to existing multi-machine services provisioned from this blueprint.

When you change the available menu operations for a multi-machine blueprint and select **Apply machine operation changes to existing machines**, the changes are applied to all multi-machine services provisioned from the blueprint as well as their component machines.

When you change the available menu operations for a machine blueprint that is a component of a multi-machine blueprint and select **Apply machine operation changes to existing machines**, the changes are applied to all machines provisioned from the blueprint including any machines that may be components of a multi-machine service.

Requesting a Multi-Machine Service

When you request a new multi-machine service, you can specify configuration options for the component machines as well as the multi-machine service as a whole.

Requests for multi-machine services (and their component machines) may be subject to approval based on the approval policies specified in their respective blueprints.

To request a multi-machine service:

1. Log in to vCAC and click **Self-Service > Request Machine**.
2. Click the name of a blueprint whose **Type** is **Multi-Machine**.
3. Specify the settings that apply to the multi-machine service as a whole:
 - a. Ensure that the **Multi-Machine Service** is selected from the grid at the top.

- b. In the **Request Information** tab, specify a value for the **Number of multi-machine services** to provision as part of this request and the **Lease duration** in days. Depending on the blueprint, these values may not be editable.
 - c. Optionally, specify a **Description** to use for the multi-machine service.
 - d. Optionally, enter a **Reason for request**. This information is passed along to approvers for the multi-machine service and the component machines.
4. Configure the component machines as needed. For each component type that you want to configure:
 - a. Select the component from the grid at the top.

The **Request Information** tab updates with machine details based on the component blueprint.
 - b. Confirm or update the values for **Number of machines**, **Number of CPUs**, **Memory**, and **Storage**. Depending on the blueprint, these values may not be editable.

Note: For component blueprints that use vSphere, SCVMM, or Amazon EC2, storage volumes are configured on a separate **Storage** tab.

- c. Optionally, specify a **Description** to use for all machines of the same component type in this multi-machine service.
 - d. For cloud component machines only, specify the instance type, whether to use Virtual Private Cloud (VPC), and the location for the Amazon instance. Depending on the blueprint, these options may not be available. Selecting VPC may require additional configuration. For details, see [Virtual Private Cloud](#).
5. Click **OK**.

Note: If you are an enterprise administrator or provisioning group manager, you can also specify custom properties at request time. You can add custom properties to the **Multi-Machine Service** that apply to all component machines or add properties that apply only to machines of a particular component type.

About Approval and Multi-Machine Services

You can apply an approval policy to a multi-machine blueprint just as you would for any blueprint.

Multi-machine service requests use the same approval policies as machine requests. Advanced approval policies are supported for multi-machine service requests. For information about creating approval groups and approval policies, see the *vCloud Automation Center Operating Guide*.

The individual component machines in a multi-machine service are also subject to approval based on the approval policies and thresholds specified in the component blueprints just like any standalone machine, and may be approved or rejected independent of the multi-machine service.

After the multi-machine service or any component machine has been approved (or if it does not require approval), its state changes to **WaitingToBuild**. When the multi-machine service and all its component machines are in this state (that is, all required approvals are complete), the multi-machine service is provisioned.

One or more component machine requests may be rejected during the approval process. If any machines are rejected, one of the following occurs:

- If enough machines are approved to meet the minimum machine requirements for each

component type, the multi-machine service is provisioned with the component machines that were approved.

- In a standard approval workflow, if the number of machines that are approved for a particular component type is less than the minimum for that type, then the multi-machine service is destroyed.

A message appears under Recent Events similar to the following: "Machine Service-01: Failed to Launch Multi-Machine Service: Not enough components of blueprint "Database Server" are found; Minimum number of machines: 1, number of machines that are found: 0."

- If advanced approvals are enabled and the approval policy for a multi-machine service includes approval groups after the **[MachineProvisioning]** step, it is possible for enough components to be rejected *after* the multi-machine service has been provisioned to cause the multi-machine service to become unhealthy. In this case, the multi-machine service remains provisioned but its state is set to **Off**.

The following messages appear under Recent Events: "Machine Service-01: Not all requested components of blueprint "Web Server" are provisioned; Requested number of machines: 4; Provisioned number of machines: 1." followed by "Machine Service-01: Not enough components of blueprint "Web Server" are On; Minimum number of machines: 2, number of machines that are On: 1."

Emails are sent to approvers to notify them of multi-machine service or component machine requests awaiting their approval. For more information about notification emails that are sent for multi-machine services, see [About Notification Emails and Multi-Machine Services](#).

Troubleshooting Multi-Machine Service Requests

Requesting a multi-machine service is a way to provision multiple machines with a single request. Status messages for each individual machine request as well as the overall multi-machine service request are displayed in the Recent Events log.

If an error prevents one or more of the component machines from provisioning, one of the following occurs:

- If enough machines are successfully provisioned to meet the minimum machine requirements for each component type, the multi-machine service is provisioned and turned on. A warning message appears under Recent Events similar to the following: "Machine Service-01: Not all requested components of blueprint "Web Server" are provisioned; Requested number of machines: 4; Provisioned number of machines: 3." After you have resolved the issue that prevented the machine or machines from provisioning, you can add machines to the multi-machine service until it has the number of machines that you originally requested.
- If the minimum number of machines failed to provision successfully for one or more component types, the multi-machine service is provisioned but its state is set to **Off**. A message appears under Recent Events similar to the following: "Machine Service-01: Failed to Launch Multi-Machine Service: Not enough components of blueprint "Database Server" are found; Minimum number of machines: 1, number of machines that are found: 0." After you have resolved the issue that prevented the machine or machines from provisioning, you can add machines to the multi-machine service until it is healthy.
- If no component machines provision successfully, then the multi-machine service is destroyed. The following message appears under Recent Events: "Multi-Machine Service with no components is invalid." You must resolve any issues that prevented the machines from provisioning and then re-request the entire multi-machine service.

Managing Multi-Machine Services

Once the multi-machine service is provisioned, the machine owner can view its status and perform various machine menu tasks on the multi-machine service as a whole from the My Machines page.

Multi-machine service menu operations include:

- Editing the multi-machine service
- Viewing component machines
- Adding machines to a multi-machine service
- Rebooting or powering the multi-machine service (and its component machines) on and off
- Editing the multi-machine service lease
- Destroying a multi-machine service and its component machines

Reprovisioning is not available at the multi-machine service level, although individual component machines can be reprovisioned.

Note: The Bulk Operations Client is not supported for multi-machine services. If you use the Bulk Operations Client on a server that has multi-machine services, do not select machines of type “Multi-Machine” to perform bulk operations.

Additionally, the Bulk Operations Client does not indicate whether a machine is a component of a multi-machine service. Performing operations on component machines, such as powering off a group of component machines, can lead to the multi-machine service being in an unhealthy state.

About Editing Multi-Machine Services

Machine owners can edit the description of a multi-machine service. Support users can change the owner of a multi-machine service. Enterprise administrators and provisioning group managers can also change the properties of a multi-machine service.

Changing the owner of a multi-machine service updates the owner of all its component machines.

Editing the description or properties of a multi-machine service does not update the description or properties of existing component machines. If you add component machines to the multi-machine service after editing its properties, the newly added machines inherit the current properties on the multi-machine service.

Viewing the Status of Multi-Machine Services

The My Machines page displays multi-machine services in a separate grid from standalone machines.

By default, the multi-machine services grid displays the name, status, and description of each multi-machine service along with its expiration and destruction dates. Additional columns can be added via the **Columns** menu.

A multi-machine service is considered to be healthy and has a status of **On** if at least the minimum number of machines for each component type are powered on. When the number of machines that are powered on for a component type is lower than the minimum, the multi-machine service is considered unhealthy and its state is set to **Off**. During state data collection, vCAC detects the power state of each component machine, and the status of the multi-machine service is updated as necessary based on the power state of the component machines.

Managing Component Machines

Most machine menu operations are available for component machines in a multi-machine service.

To access the component machines in a multi-machine service, select **View Components** from the multi-machine service menu. The Component Machines page displays a list of machines similar to the My Machines page, except that it only shows the machines that are components of the multi-machine service.

You can edit component machines just like standalone machines, except that you cannot change the owner of a component machine. To change the owner of a component machine, edit the multi-machine service, which updates the owner for all component machines.

Lease operations (**Change Lease**, **Extend Lease**, or **Expire**) are not available on individual component machines. Leases can only be controlled at the multi-machine service level for all component machines. For more information, see [About Managing Leases for Multi-Machine Services](#).

Adding Component Machines to a Multi-Machine Service

If the multi-machine service has not reached its maximum number of allowed machines for a component type, you can add component machines to the multi-machine service.

To add component machines to an existing multi-machine service:

1. Log in to vCAC.
2. Do one of the following:
 - a. As a machine owner, click **Self-Service > My Machines**.
 - b. As a provisioning group manager, click **Provisioning Group Manager > Group Machines**.
 - c. As an enterprise administrator, click **Enterprise Administrator > Enterprise Machines**.
3. Point to the name of a multi-machine service and click **Add Components**.
4. Select the component type for which you want to add a machine. Component types that are already at their maximum number of machines are disabled.

Note: Adding components to a multi-machine service in the vCAC Console is treated as a new request from the component blueprint. You must specify the configuration for the new component machine as with any request.

5. In the **Request Information** tab, specify the following:
 - a. In the **# Machines** box, specify the number of machines of this type that you want to add to the multi-machine service. For example if you have 2 machines of this component type and specify 2 additional machines on the Add Components page, the updated multi-machine service has 4 machines of that type.
 - b. Configure the resources, such as **CPU**, **Memory**, and **Storage** for the additional machines. Depending on the blueprint, some or all of these options may not be editable.

Note: For component blueprints that use vSphere, SCVMM, or Amazon EC2, storage volumes are configured on a separate **Storage** tab.

- c. For cloud component machines only, specify the instance type, whether to use Virtual Private Cloud (VPC), and the location for the Amazon instance. Depending on the

blueprint, these options may not be available. Selecting VPC may require additional configuration. For details, see [Virtual Private Cloud](#).

6. Optionally, if you are an enterprise administrator or provisioning group manager, specify custom properties to apply to the new machines of this component type on the **Properties** tab.
7. Repeat steps 4-6 to add component machines of a different component type.
8. If any of the machines you added require approval, specify a reason for your request.
 - a. Select the name of the multi-machine service.
 - b. In **Reason for request**, enter a brief summary of why you need the additional machines.

This information is passed along to approvers for the component machines.

9. Review the information in the summary grid at the top and click **OK** to add the machines.

Power Options for Multi-Machine Services

The multi-machine service menu provides several options for managing the power status of a multi-machine service and all its component machines.

The following table describes the behavior of power options on multi-machine services as opposed to individual machines.

Menu option	Description
Power On	Power on all component machines in the startup order specified on the multi-machine blueprint.
Power Off	Power off all component machines without shutting down the guest OS in the shutdown order specified on the multi-machine blueprint.
Power Cycle	Power off all component machines in shutdown order (see Power Off), then power on all machines in the startup order.
Reboot	Shut down all component machines in shutdown order (see Shut Down), then power on all machines in the startup order.
Shut Down	Shut down the guest OS and power off all component machines in the shutdown order specified in the multi-machine blueprint.

If a multi-machine service has multiple machines of the same component type, all machines of that type are powered on before machines of the next component type in the startup order, and likewise for shutdown order. There is no set order in which individual machines of the same type are powered on or off.

Destroying Component Machines

If a component machine is no longer needed, you can destroy the unwanted machine as long as it would not cause the multi-machine service to become unhealthy.

To remove a component machine from a multi-machine service:

1. Log in to vCAC.
2. Do one of the following:
 - As a machine owner, click **Self-Service > My Machines**.

- As a provisioning group manager, click **Provisioning Group Manager > Group Machines**.
 - As an enterprise administrator, click **Enterprise Administrator > Enterprise Machines**.
3. Point to the name of a multi-machine service and click **View Components**.
 4. Point to the name of a machine and click **Destroy**.

To remove several machines at once, select the machines using the checkbox next to the machine name, then point to one of the selected machines and click **Destroy**.
 5. On the Confirm Machine Action page, review the list of machines that you selected to destroy.

If destroying one or more of the selected machines would cause the multi-machine service to become unhealthy, you can clear the selection on this page, or click **Cancel** to return to the Component Machines page.
 6. Click **OK** to destroy the machines.

About Managing Leases for Multi-Machine Services

The lease period for a multi-machine service always overrides the lease on its component machines.

- A component machine always has the same lease period as the multi-machine service, regardless of any minimum or maximum lease period specified in the component blueprint. If a component machine's lease changes for any reason, its expiration date is automatically synchronized with the multi-machine service at regular intervals.
- Operations such as **Change Lease**, **Extend Lease**, and **Expire** are available on the multi-machine service menu and not on the component machine menus. When the expiration date for a multi-machine service changes, that change is applied to all its component machines.
- Component machines cannot expire earlier than the multi-machine service to which they belong.
- When a multi-machine service expires, all its component machines expire at the same time. The multi-machine service and component machines are destroyed immediately when they expire.

About Notification Emails and Multi-Machine Services

vCAC sends notification emails for events such as provisioning and lease notifications for the multi-machine service. By default, vCAC does not send these emails for individual component machines.

You can override the email behavior for component machines with the custom property **MultiMachine.ComponentsEmail.Enabled**. If this property exists on a component machine (regardless of the value), vCAC sends event notification emails for that machine.

Approval emails are always sent for both multi-machine services and component machines.

Notification emails for multi-machine service events have a different subject heading and body text than emails for machine events, but they are configured using the same email template files. Multi-machine service emails share the same configuration to enable or disable machine emails for specific events in the Manager Service configuration.

For information about configuring notification emails, see the *vCloud Automation Center Operating Guide*.

Differences in Multi-Machine Services Between the Self-Service Portal and the vCAC Console

There are some behavioral differences when requesting and managing multi-machine services using the vCAC Self-Service Portal.

- When you use the Self-Service Portal to request a multi-machine service that includes a cloud machine as a component, some configuration options are not available at request time for the cloud component. You can click **Customize** to select the instance type of a cloud component machine, but you cannot specify network options (including VPC or non-VPC location) or add EBS storage volumes.
- When you add components to a multi-machine service from the Self-Service Portal, vCAC provisions the new machines with the same specifications as existing machines of the same component type. (Cloud component machines have the same instance type but are not guaranteed to have the same network and storage configuration.) If the component machine was provisioned using vSphere, then it may be possible to reconfigure the machine after provisioning to have different specifications.

Reconfigure Machine

The following sections describe the reconfigure machine feature, which is new in release 5.1. Reconfiguring machines is only available for virtual machines provisioned using VMware vSphere.

About Reconfiguring Machines

The reconfigure machine feature enables machine owners to modify the compute resources of a virtual machine.

A machine owner can make any of the following changes to a provisioned machine:

- Increase or decrease **Memory**
- Increase or decrease the number of **CPUs**
- Modify **Storage** by adding, removing, or increasing the size of volumes (SCSI disks only)
- Modify **Networks** by adding, removing, or updating network adapters

Changes to each of these parameters are subject to limits defined in the blueprint that was used to provision the machine. Reconfigure requests can also be made subject to approval, as with provisioning requests.

By default, the reconfiguration executes immediately after any necessary approvals, but a machine owner can optionally schedule a reconfiguration for a specific time or opt to trigger the reconfiguration manually after it is approved.

The virtual machine must be powered off when the reconfiguration occurs. To enable vCAC to shut down virtual machines cleanly before reconfiguration, VMware Tools must be installed on the guest operating system of the machine being reconfigured.

Reconfiguration Life Cycle

To reconfigure a machine, it must have been provisioned from a blueprint that allows reconfiguration. For details about enabling the reconfigure machine feature for a blueprint, see [Enabling Reconfigure Machine on Blueprints](#).

A machine owner can request a reconfiguration when the machine is in either the **On** or **Off** state. vCAC checks whether approval is required based on the approval policy and approval thresholds specified in the blueprint.

Once the request is approved, the reconfiguration is ready to be executed. The three options for when to execute the reconfiguration are:

- **Immediately** — The reconfiguration is automatically queued for execution after all required approvals have been given.
- **Scheduled** — The machine owner specifies a date and time to execute the reconfiguration. Execution can be rescheduled if the request is not approved until after the scheduled time.
- **Queue for Owner** — The reconfiguration must be triggered manually by the machine owner after all approvals have been given.

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

Note: Machines can be reconfigured using a reservation that has been disabled. Therefore, it is important to monitor resources that are part of disabled reservations.

If the virtual machine is powered on when the reconfiguration begins, it is shut down or powered off before reconfiguring and then powered back on at the end of the reconfiguration. If the virtual machine is powered off when reconfiguration begins, it remains powered off at the end of the reconfiguration process.

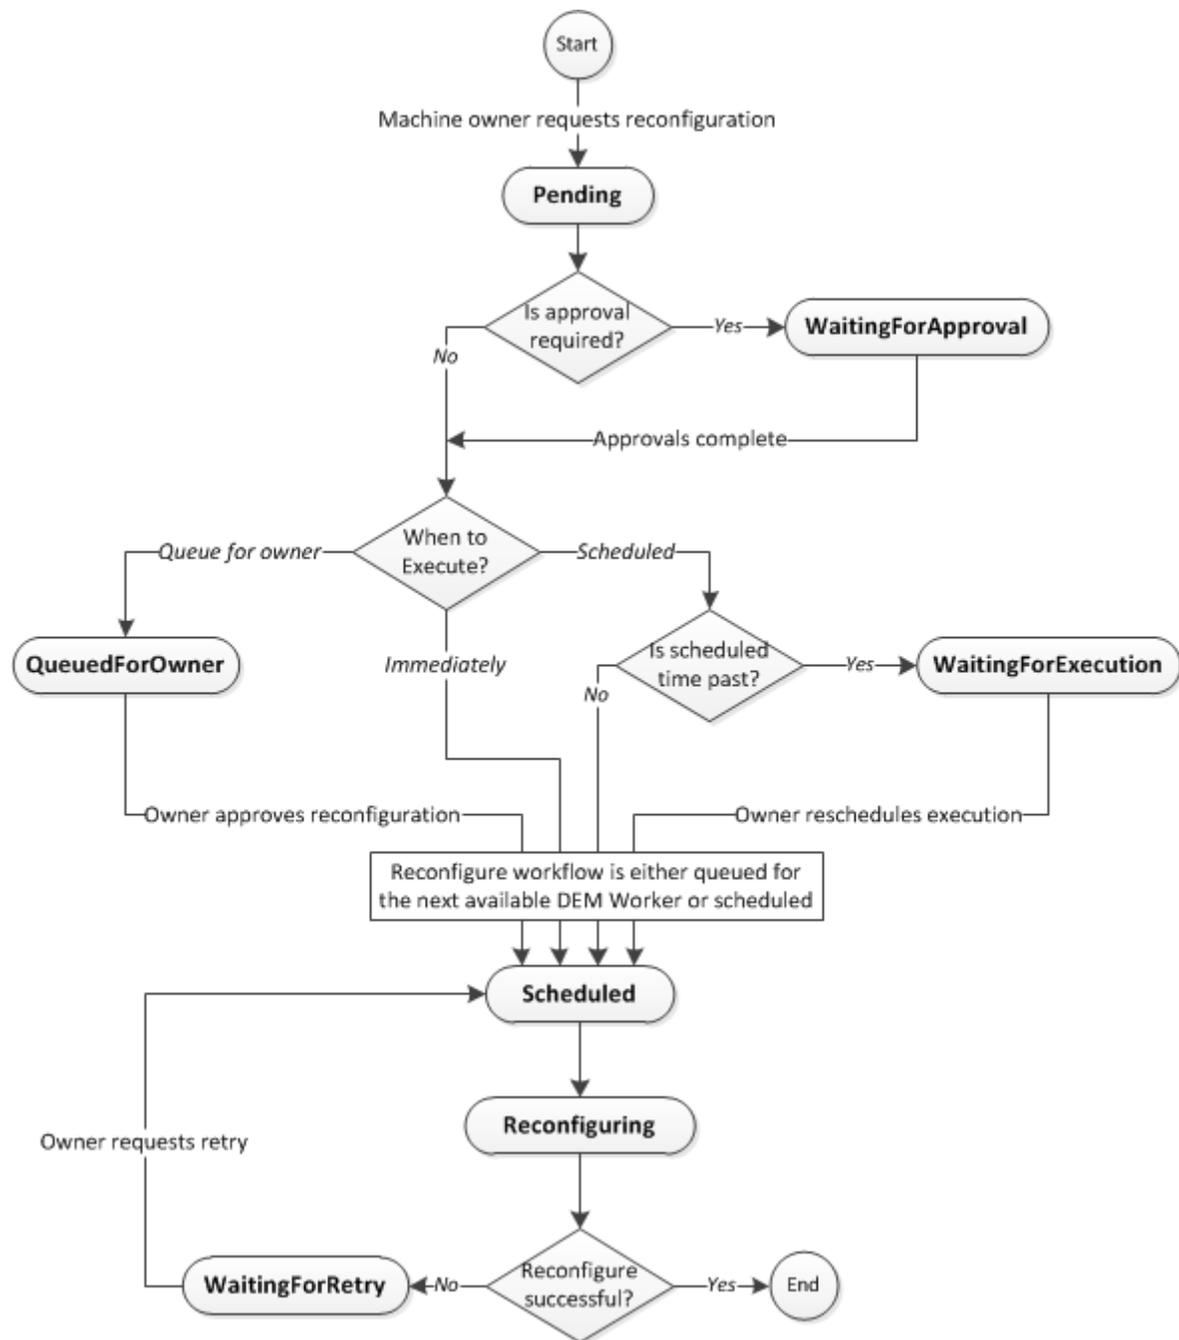
If the reconfiguration fails, any additional resources (such as storage) that were allocated in anticipation of the reconfiguration are deallocated and made available to other machine requests. The machine owner has the option to retry the reconfiguration (without any additional approval process) or to cancel the reconfiguration.

Email notifications are sent, as with a new provisioning request, at each stage of the approval process and upon completion of the reconfiguration, whether it was successful or not.

About Reconfiguration Status

The status of the machine is updated with a reconfiguration status that reflects the current stage of the reconfiguration life cycle as it goes through approval and execution. The reconfiguration status appears in parentheses after the machine status. For example, when a reconfigure request has been submitted and requires approval, the machine may have the status **On (Reconfigure.WaitingForApproval)**.

The following diagram shows the reconfiguration statuses and their place within the overall reconfiguration life cycle.



Configuring the Reconfigure Machine Feature

This section assumes that you are familiar with creating blueprints and the basic provisioning and approval process. They describe how the process differs when using the reconfigure machine feature.

Enabling Reconfigure Machine on Blueprints

You can enable the reconfigure machine feature on virtual blueprints that use vSphere for provisioning.

To enable reconfigure for a specific blueprint:

1. Create a new blueprint or edit an existing blueprint that uses vSphere.
2. On the **Blueprint Information** tab, specify blueprint details as you would for any blueprint.
3. On the **Build Information** tab, ensure that the **Platform type** is **vSphere (vCenter)**.
4. In the **Machine Resources** section:
 - a. Specify the **Minimum** value for **# CPUs** and **Memory (MB)**.
 - b. Optionally, specify a minimum and maximum number of days and an approval threshold for the machine lease. Leave the fields blank for no expiration date.
 - c. Add at least one volume to the blueprint.
 - i. Click **Add Volume**.
 - ii. Specify a minimum value for **Capacity (GB)**.
 - iii. Optionally, specify other values such as **Drive Letter, Label, or Storage Reservation Policy**.
 - iv. Click the green check mark to save the new volume.
 - v. Optionally, specify custom properties for this volume by clicking **Edit** under **Custom Properties**.

For more information about configuring storage volumes, see [Using Multiple Datastores for Machine Volumes](#).

- d. To offer the machine owner a choice within a range of values for CPUs, memory, or storage, specify a **Maximum** value for the appropriate resource. This maximum applies to both the initial machine request and any reconfigure request.

Note: Specifying a maximum value is optional in the blueprint, however, if you do not specify a range of values for a particular resource, that aspect of the machine cannot be reconfigured.

If you have specified an approval policy on the **Blueprint Information** tab and a maximum value for a machine resource, you can optionally specify an approval threshold in the **Approval At** column. When the machine owner selects a value at or above the threshold, or if there is a maximum value but no threshold for a particular resource, approval is required before the machine is provisioned or reconfigured.

- e. Specify the maximum number of volumes. Leave the field blank to accept the default maximum of 15 volumes. Specify a value of **0** to disable the addition of volumes.
 - f. Specify the maximum number of network adapters. Leave the field blank for no limit. Specify a value of **0** to disable the addition of network adapters.
5. On the **Properties** tab, specify custom properties as you would for any blueprint.
 6. On the **Security** tab, specify the blueprint access and enable machine operations as you would for any blueprint.
 7. In the **Reconfigure** section:
 - a. Specify whether to **Allow reconfigure** for machines provisioned from this blueprint.
 - b. To enable users requesting a reconfigure to specify when to execute the reconfiguration (immediately, at a scheduled time, or on owner approval), select

Enable execution selector. If this option is disabled, the reconfiguration is always executed immediately after any necessary approvals.

- c. To enable users requesting a reconfigure to specify whether vCAC should attempt to shut down a machine before reconfiguring it, select **Enable force shutdown**. If this option is disabled on the blueprint, force shutdown is always selected when a user requests a reconfigure.

Recall that the machine must be powered off for reconfiguration. If the machine is powered on when the reconfiguration begins, the force shutdown option enables vCAC to perform a graceful shutdown of the guest operating system before powering off the machine so that the reconfiguration can proceed. Shutting down a virtual machine requires VMware Tools to be installed on the guest OS.

About Approving Reconfigure Requests

Machine reconfiguration requests use the same approval policy as for machine provisioning requests.

Reconfigure requests are listed in a separate section on the Machine Requests page from provisioning requests. You can view a summary of the changes that are being requested for the machine by clicking the magnifying glass next to the machine name. You must either approve or reject the entire reconfigure request; you cannot approve a change in one resource (such as CPU) and reject another aspect (such as storage) if both resources were modified in the same request.

Advanced approval, including the option to delegate approvals, is not supported for reconfiguration. If a machine is provisioned from a blueprint that uses an advanced approval policy, then the reconfigure approval ignores any <<Owner>> steps and any steps after the [MachineProvisioning] step.

By default, vCAC determines the whether a reconfigure request requires approval according to the same rules as for provisioning requests. That is, if a blueprint has an approval policy, and any machine resource has both a minimum and maximum value, then approval is required if the user requesting the reconfigure selects a value for that resource at or above the approval threshold or if no approval threshold is specified in the blueprint.

You can require approval for all reconfigure requests (regardless of whether the approval thresholds are met) by setting the **Require approval for all requests** property to **true** on the **vCAC Administrator > Global Properties** page. This applies only to blueprints that have an approval policy (including Default Approval).

About Automatic Approval of Reconfigure Requests

Reconfigure requests may be automatically approved under certain circumstances, even if the approval thresholds are reached.

There are two cases where a reconfigure request is automatically approved:

- In the case of default (or group manager) approval, if the user requesting the reconfigure is a manager of the provisioning group
- Within a custom approval policy, if the requester is the **only** approver in the approval group

The above rules for automatic approval apply to a single step within an approval workflow. For example, suppose that a blueprint has a custom approval policy with several steps beginning with provisioning group manager approval followed by approval from Finance and IT. Amy is a provisioning group manager as well as one of the approvers in the Finance approval group. When she requests a reconfiguration for one of her machines (or any machine in her provisioning group), the following occurs:

1. The first step is automatically approved, so Amy does not have to approve her own request as provisioning group manager.
2. At the second step, because Amy is not the only approver in the Finance group, the request is sent to all the approvers in the group. Amy may approve her own request or another approver in the group may approve it.
3. At the final step, the request must be approved by someone in the IT group.

Automatic approval is enabled by default and applies to reconfigure requests only. You can disable automatic approval by setting the **Allow automatic approvals** property to **false** on the **vCAC Administrator > Global Properties** page.

About Notification Emails for Reconfigure

Notification emails are sent for the following events related to reconfiguring machines:

- When a reconfigure request requires approval
- When a reconfigure request is approved or rejected
- When a reconfiguration is queued for the owner to execute it
- When the reconfigure completes, with information about whether it was successful or not

The configuration for reconfiguration emails (including SMTP server information and the email address from which the emails are sent) are specified through the same global properties as other vCAC emails. For more information, see the *vCloud Automation Center Operating Guide*. The text of reconfiguration emails cannot be customized.

Reconfiguring a Machine

You can request reconfiguration of a machine by editing the machine when its status is either **On** or **Off**.

The aspects of the machine that you can reconfigure (CPU, Memory, Storage, Network) are defined in the blueprint that you used to provision the machine. Some of the options described below may not be available for a particular machine depending on its blueprint.

To request a reconfiguration:

1. Log in to vCAC.
2. Do one of the following:
 - a. As a machine owner, click **Self-Service > My Machines**.
 - b. As a provisioning group manager, click **Provisioning Group Manager > Group Machines**.
 - c. As an enterprise administrator, click **Enterprise Administrator > Enterprise Machines**.

The machine you want to reconfigure should have the status **On** or **Off** with no active reconfigure status.

3. Point to the name of the machine and click **Edit**.
4. On the **Machine Info** tab, you can adjust the number of CPUs and amount of Memory on the machine. The ranges are the same as when you requested the new machine from this blueprint, and are subject to the same approval thresholds.

5. On the **Storage** tab, you can add volumes, remove volumes, or increase the size of existing volumes.

Note: Reconfigure is not supported for IDE disks.

- To add a new volume, click **Add Volume**, specify the **Capacity** in GB, and click the green check mark to save your changes. Optionally, you can also specify a **Storage Reservation Policy** for the new volume or specify custom properties for the volume.
- To remove a volume, click the trash can icon next to the disk. You cannot remove volume 0.
- To increase the size of an existing volume, click the pencil icon next to the volume, specify a new value for **Capacity**, and click the green check mark to save your changes. You cannot reduce the size of existing volumes.

Volume size is limited by the total amount of storage specified in the blueprint. For example, if the blueprint has a maximum storage capacity of 30GB, and you provision a machine with 10GB of initial storage, you have 20GB remaining to allocate to existing volumes or additional volumes.

6. On the **Network** tab, you can add network adapters, remove network adapters, or edit existing adapters.

- To add a new network adapter, click **Add Network Adapter**, select the network under **Name**, and click the green check mark to save your changes. All networks enabled on the machine's reservation are available. Each adapter on the machine must be connected to a unique network.
- To remove a network adapter, click the trash can icon next to the network adapter. You cannot remove the first network adapter.
- To edit an existing network adapter, click the pencil icon next to the network adapter, select a new network under **Name**, and click the green check mark to save your changes.

7. On the **Execution** tab, specify additional information about the reconfigure request.

- a. From the **Execute request** dropdown list, select an option for when to execute the reconfiguration:

- If you select **Immediately**, the reconfiguration begins as soon as possible after it is approved.
- If you select **Scheduled**, you can specify a date and time to execute the reconfiguration.

Note: The scheduled time is the local time where the vCAC web server is located. If you are not in the same time zone as the server, you should manually calculate the server time when scheduling a reconfiguration.

- If you select **Queue for Owner**, the reconfiguration is put on hold after any necessary approvals until you execute it. For more information, see [Executing an Approved Reconfiguration](#).

If the **Execute request** option is not available, the reconfiguration is executed immediately.

- b. If you want vCAC to attempt to shut down the machine safely before reconfiguring, select **Shut down machine before reconfigure**. If the shutdown operation fails for any reason, the reconfiguration is canceled. If the machine is powered off before requesting the reconfiguration, you can clear the selection for this option. If you do not select this option and the machine is powered on when the reconfiguration begins, vCAC performs a hard power-off (the equivalent of pulling the plug) before reconfiguring the machine.

If the virtual machine is powered on when the reconfiguration begins, it is shut down or powered off before reconfiguring and then powered back on at the end of the reconfiguration. If the virtual machine is powered off when reconfiguration begins, it remains powered off at the end of the reconfiguration process.

- c. Optionally, specify a reason for the reconfiguration. This information is passed along with the request to approvers.

8. Click **OK** to submit the reconfigure request.

An entry appears in the Recent Events pane for the machine owner and the requester (if different), including a brief summary of the reconfigure request, and the request is sent for approval if necessary.

Executing an Approved Reconfiguration

If you select **Queue for Owner** when specifying when a reconfigure request should be executed, the request is held after it goes through any required approvals until you trigger it manually.

When the request is approved, you receive an email notifying you that the reconfiguration is waiting for your final approval to proceed.

To execute an approved reconfiguration:

1. Click the link in the notification email or log in to vCAC and click **My Machines**.
The machine that is being reconfigured should have the status **On** or **Off** with a reconfigure status of **Reconfigure.QueuedForOwner**.
2. Expand the **My Pending Requests** section.
The **My Pending Requests** grid appears below **My Machines** and **My Multi-Machine Services**. Expand it by clicking on the arrows at the right end of the header.
3. Locate the machine that you want to reconfigure.
4. Click **Approve**.

The reconfiguration proceeds. Refresh the My Machines page to view the status of the reconfiguration.

Retrying a Reconfiguration

If a reconfiguration fails, you receive a notification email and you have the option to retry the reconfiguration without seeking approval.

To retry a reconfiguration:

1. Click the link in the notification email or log in to vCAC and click **My Machines**.
2. The machine should have the status **On** or **Off** with one of the reconfigure status **Reconfigure.WaitingForRetry**.
3. In the **Recent Events** section, locate the message that describes the error that prevented the reconfiguration from completing successfully.

4. Resolve the error or work with an administrator to resolve it.
5. Point to the name of the machine and click **Execute Reconfigure**.
6. Specify a date and time to execute the reconfiguration or select **Now** to execute it immediately.
7. Click **OK**.

Rescheduling a Reconfiguration

You can reschedule a reconfiguration (which includes the option to execute it immediately) if the reconfiguration has not already occurred.

Additionally, if a reconfigure request is not approved until after the time when it was scheduled to execute, you can specify a new time or execute it immediately. You receive a notification email when a reconfiguration request is approved. The approval email does not indicate whether the reconfiguration is scheduled or whether the scheduled time has already passed, so you should check the status of your request after it is approved to see if you need to take any action.

To reschedule a reconfiguration:

1. Click the link in the notification email or log in to vCAC and click **My Machines**.
2. The machine should have the status **On** or **Off** with one of the following reconfigure statuses:
 - **Reconfigure.Pending**
 - **Reconfigure.WaitingForApproval**
 - **Reconfigure.Scheduled**
 - **Reconfigure.WaitingForExecution** (indicates that approval occurred after the scheduled time)
3. Point to the name of the machine and click **Execute Reconfigure**.
4. Specify a date and time to execute the reconfiguration or select **Now** to execute it immediately.
5. Click **OK**.

Canceling a Reconfiguration

You can cancel a reconfiguration if it is waiting for approval or if it has been approved and is waiting for execution.

To cancel a reconfiguration:

1. Log in to vCAC and click **My Machines**.
2. The machine should have the status **On** or **Off** with one of the following reconfigure statuses:
 - **Reconfigure.WaitingForApproval**
 - **Reconfigure.QueuedForOwner**
 - **Reconfigure.Scheduled**
 - **Reconfigure.WaitingForExecution**
 - **Reconfigure.WaitingForRetry**
3. Point to the name of the machine and click **Cancel Reconfigure**.

4. Click **OK**.

Troubleshooting Reconfigure Machine

This section describes issues you may encounter when using the reconfigure machine feature.

Data Collection

vCAC may display temporarily incorrect information about virtual machines immediately after reconfiguring.

This can occur when reconfiguration and data collection are running concurrently. For example:

- A virtual machine is powered on and is automatically shut down in preparation for reconfiguration. While the machine is powered off, state data collection runs and records the power state of the machine as **Off**. The reconfiguration process completes and the machine is powered on, but it still appears as **Off** in the vCAC Console.
- A virtual machine is reconfigured while inventory data collection is being performed. Depending on the point within the data collection process at which the machine is reconfigured, the machine may display the old value for the reconfigured resource (detected by the agent before the reconfiguration occurred) after the reconfiguration is complete.

Another case is when a virtual machine that is managed by vCAC is reconfigured directly in vSphere rather than through vCAC and the information in vCAC no longer reflects the current specifications of the machine.

In these and similar situations, the machine is updated with the correct information the next time data collection runs. Both state and inventory data collection are scheduled to run at regular intervals. Alternatively, as an enterprise administrator, you can run data collection on the affected compute resource from the **Enterprise Administrator > Compute Resources** page to refresh the information in vCAC.

Destroying Machines

If a machine has an active reconfigure status, the Destroy operation is not available in the machine menu.

If you want to destroy a machine that has an active reconfigure status, you must first cancel the reconfiguration. This includes the case where a machine goes missing or it expires while a reconfigure is not yet complete. See [Canceling a Reconfiguration](#).

Changes to Reconfigure since 4.1 SP1 FP1

If you have previously used the reconfigure machine feature available in the feature pack release 4.1 SP1 FP1, please note the following differences between the behavior of reconfigure machine in release 5.1 and the earlier implementation of the feature:

- The feature pack included support for hot swap and hot plug. In release 5.1, the virtual machine must be powered off before reconfiguring. A Force Shutdown option is available to enable vCAC to shut down the machine before beginning the reconfiguration. This helps prevent possible damage to the machine and minimizes downtime compared to manually powering off the machine before reconfiguring and powering it back on after the reconfigure is complete. Force Shutdown execution requires VMware Tools to be installed on the guest operating system.
- Advanced approvals were supported in the feature pack but not in release 5.1.

- In the feature pack, configuration related to reconfigure was specified in a reconfigure profile that was associated with a blueprint. In release 5.1, many of these settings are configured in the blueprint directly. Some of the options in the reconfiguration profile are no longer supported, such as **Require approval for resource decrease** and **Retry on failure**.
- In the feature pack, the ability to retry a failed reconfiguration was controlled at the blueprint level. In release 5.1, retry is always available but is limited to one retry attempt per request.
- In the feature pack, notification emails related to reconfigure were configured using email profiles. In release 5.1, reconfigure emails use the same global properties for email configuration as all emails sent by vCAC.
- Previously, it was possible to attach multiple network adapters on the same network to a single machine. In release 5.1, duplicate networks are not supported.
- Global custom properties (configured on the Customization page for vCAC Administrators) that were used in the feature pack are not applicable for this release.
- In the feature pack, menu options related to reconfigure were controlled on the **Security** tab of the blueprint under the **Machine owners can** section. In release 5.1, menu options related to reconfigure are automatically enabled when reconfigure is enabled for the blueprint.

Using Multiple Datastores for Machine Volumes

The multi-storage functionality introduced in the 4.5 release lets you assign the volumes of a virtual machine to different datastores for the vSphere and SCVMM platform types. In the 5.1 release, the custom properties used in multi-storage are now managed by the user interface (except for Linked Clones). See [Changes to Multiple Datastores since 4.5](#) for a brief description of the multi-storage changes implemented since version 4.5.

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

You can assign a single datastore or a storage reservation policy representing multiple datastores to a volume. When you assign a single datastore to a volume, vCAC uses that datastore at provisioning time, if possible. When you assign a storage reservation policy to a volume, vCAC uses one of its datastores at provisioning time, if possible. A storage reservation policy is essentially a tag applied to one or more datastores. It allows you to group datastores that have similar characteristics, such as speed or price. A datastore can be assigned to only one storage reservation policy at a time, but a storage reservation policy can have many different datastores.

The enterprise administrator creates a storage reservation policy and assigns it to one or more datastores. The enterprise administrator or provisioning group manager then assigns the storage reservation policy to a volume in a virtual blueprint. When a user requests a virtual machine that uses the blueprint, vCAC uses the storage reservation policy specified in the blueprint to select a datastore for the machine's volume.

You can do the following with datastores when adding or editing a volume in a virtual blueprint:

- Assign a single datastore to the volume. This one-to-one relationship allows vCAC to assign only the specified datastore at provisioning time. If the datastore is not available or lacks sufficient space, provisioning fails. See [Assigning Datastores to Volumes](#).
- Assign a storage reservation policy to the volume. This allows vCAC to select one of the datastores associated with the storage reservation policy at provisioning time. If the policy's datastores are unavailable or lack sufficient space, provisioning fails. You can override this restriction by [Allowing Alternative Datastores at Provisioning Time](#).

- *Not* assign a datastore or a storage reservation policy to a volume. This allows vCAC to select from all available datastores at provisioning time. Provisioning fails if a datastore with sufficient space is not available.

vCAC also uses a datastore's priority (as assigned by its reservation) to select a datastore from among eligible datastores at provisioning time.

You should generally avoid specifying a storage path and a storage reservation policy for the same volume, though there may be times when this is desirable. If this situation does arise, the storage path takes precedence over the storage reservation policy.

Creating a Storage Reservation Policy

The enterprise administrator can create, update, delete, and assign a storage reservation policy.

To create a storage reservation policy:

1. Log on as **Enterprise Administrator**.
2. Select **Reservation Policies**.
3. Click the **New Storage Reservation Policy** link in the **Storage Reservation Policies** title bar.

In the **Name** field, enter a name for the storage reservation policy. For example, a tiered strategy might have policies named Tier 1, Tier 2, and Tier 3.
4. In the **Description** field, you can type additional information.
5. Click  to save the storage reservation policy.

Your current storage reservation policies appear at the bottom of the page. After creating a storage reservation policy, you can assign it to a compute resource.

Assigning Storage Reservation Policies to a Compute Resource

A storage reservation policy must be assigned to one or more datastores before you can use it with a volume. Assigning a storage reservation policy to multiple datastores creates an exclusive group of datastores for use with a blueprint.

The enterprise administrator assigns storage reservation policies to datastores on the Compute Resources page.

To assign a storage reservation policy to a datastore:

1. Log on as **Enterprise Administrator**.
2. Select **Compute Resources**.
3. Point to an existing compute resource and click **Edit**.
4. Go to the **Configuration** tab.
5. In the **Storage Path** column, click  next to the datastore you want to assign to a storage reservation policy.
6. In the **Storage Reservation Policy** column, select the storage reservation policy to assign to the datastore.
7. To assign storage reservation policies to additional datastores, repeat steps 5 and 6.
8. Click  to save the assignment.

9. Click **OK** to save the configuration.

You can see which datastores are associated with a storage reservation policy by clicking **Compute Resources** and then going to the **Configuration** tab.

Using Disk Properties with Volumes

Disk properties allow you to specify a volume's storage capacity, labeling, and other properties. For create, clone, and Flexclone blueprints, you must use the **Volumes** grid to specify these properties. For a linked clone blueprint, you must use the **Properties** tab. Disk properties are for use only with virtual blueprints for the vSphere and SCVMM platforms.

The disk properties described in this section allow you to assign datastores to volumes, assign storage reservation policies to volumes, and to let vCAC use alternative datastores for volumes at provisioning time.

Using Disk Properties on the Volumes Grid

To create a disk property for a volume in the **Volumes** grid:

1. Log on as **Enterprise Administrator** or **Provisioning Group Manager**.
2. Select **Blueprints**.
 1. Select a blueprint:
 - New – Choose **New Blueprint > Virtual**.
 - Existing – Point to a create, clone, or Flexclone blueprint for a virtual machine and click **Edit**.
 4. Go to the **Build Information** tab. If creating a blueprint, select either **Hyper-V (SCVMM)** or **vSphere (vCenter)** as the **Platform type** and then, in the **Volumes** grid, click **Add Volume**.
 5. In the **Custom Properties** column of the **Volumes** grid, click **Edit** for the targeted volume.

The **Custom Properties** popup appears.

6. Select **New Property**.
7. In the **Name** and **Value** fields, enter the disk property name (use the last part only; omit **VirtualMachine.DiskN.**) and value.
8. Click  to save the property.
9. Repeat steps 6 through 8 for to create more properties.
10. Click the x at the top of the popup to close it.
11. Click **OK**.

Using Disk Properties on the Properties Tab

To create a disk property for a volume on the **Properties** tab:

1. Log on as **Enterprise Administrator** or **Provisioning Group Manager**.
2. Select **Blueprints**.
3. Point to a blueprint for a virtual machine and click **Edit**.
4. Go to the **Properties** tab.

5. Select **New Property**.
6. In the **Name** and **Value** fields, enter the full disk property name and value.
7. (Optional) Select the **Prompt User** check box. When you select this option, machine owners can see the property and change the value during a machine request. When you unselect the option, they cannot see or change the property during a machine request.
8. Click  to save the property.
9. Click **OK**.

Allowing Alternative Datastores at Provisioning Time

VirtualMachine.DiskN.StorageReservationPolicyMode (with values **Exact** and **NonExact**) determines whether vCAC can use datastores not associated with a volume's storage reservation policy. This disk property can be used with create, clone, Flexclone, and linked clone blueprints. **NonExact** allows vCAC to use other datastores when no space is found in the datastores associated with the storage reservation policy. **Exact** requires vCAC to use only the datastores associated with the storage reservation policy, and if none of the datastores has sufficient space for the volume, then provisioning fails. **Exact** is the default value.

With create, clone, and Flexclone blueprints, you must use the **Volumes** grid to enter the disk property. With linked clone blueprints, you must use the **Properties** tab. When using the property in the **Volumes** grid, you specify only the last part of the property name and not the full property name. When using the disk property on the **Properties** tab, you must specify the full property name.

Where Used	Format	Value
Volumes grid	StorageReservationPolicyMode	Exact (default) NonExact
Properties tab	VirtualMachine.DiskN.StorageReservationPolicyMode	Exact (default) NonExact

Assigning Datastores to Volumes

VirtualMachine.DiskN.Storage assigns the specified datastore to volume *N* and can be used with create, clone, Flexclone, and linked clone blueprints. This property is equivalent to the **Storage Path** field in the **Volumes** grid. Enter the datastore name as it appears in the **Storage Path** column on the **Configuration** tab of **Compute Resources**.

Note: When you assign a datastore to a volume, do not specify a storage reservation policy for the same volume. If a volume does have a storage path and a storage reservation policy, the storage path takes precedence during provisioning.

With create, clone, and Flexclone blueprints, you must use the **Volumes** grid to enter the disk property. With linked clone blueprints, you must use the **Properties** tab.

Where Used	Format	Value
Volumes grid	Storage	<i>datastore_name</i>
Properties tab	VirtualMachine.DiskN.Storage	<i>datastore_name</i>

Assigning Storage Reservation Policies to Volumes (Linked Clone Blueprint Only)

VirtualMachine.DiskN.StorageReservationPolicy assigns the specified storage reservation policy to volume *N* and can be used only with linked clone blueprints. This property is equivalent to the **Storage Reservation Policy** field in the **Volumes** grid.

Where Used	Format	Value
Volumes grid	n/a	
Properties tab	VirtualMachine.DiskN.StorageReservationPolicy	storage_reservation_policy_name

Note: If you use `VirtualMachine.Disk0.StorageReservationPolicy` or `VirtualMachine.Disk0.Storage`, the virtual machine is provisioned on the same storage as Volume 0. To change this, use the property `VirtualMachine.Storage.Name`. For more information on this property, see the *vCloud Automation Center Operating Guide*.

Using Storage Reservation Policies in a Blueprint

Storage reservation policies can be assigned to virtual machine volumes at several points during the provisioning process. An enterprise administrator or provisioning group manager can assign storage reservation policies when editing or creating virtual blueprints. A machine owner can assign a storage reservation policy when requesting a virtual machine, if the current blueprint allows it. This section focuses on using storage reservation policies in blueprints.

The new 5.1 multi-storage user interface assists you in assigning storage reservation policies in create, clone, and Flexclone blueprints. With linked clone blueprints, you use custom properties to assign storage reservation policies. This section explains how to use both methods.

Note: Create, Clone, Flexclone, and Linked Clone are the **Action** field values found on the **Build Information** tab of the **Edit Blueprint – Virtual** page.

When using clones, the volumes of the cloned template appear in the **Machine Resources** section on the **Build Information** tab. You cannot remove them, change their capacity, or specify a drive letter/mount path or label for them. However, you can specify a storage path or assign a storage reservation policy to the volumes.

You can use a build profile to specify the same disk property values for multiple blueprints. Simply specify the full names and values of the disk properties in the build profile and apply the build profile to the blueprints. If a blueprint and the build profile have the same disk property but different values, the blueprint takes precedence over the build profile.

Using Create, Clone, or Flexclone Blueprints

To assign a storage reservation policy to a volume using a create, clone, or Flexclone blueprint:

1. Log on as **Enterprise Administrator** or **Provisioning Group Manager**.
2. Select **Blueprints**.
2. Select a blueprint:
 - New – Choose **New Blueprint > Virtual**.

- Existing – Point to a create, clone, or Flexclone blueprint for a virtual machine and click **Edit**.
3. Go to the **Build Information** tab. If creating a blueprint, select either **Hyper-V (SCVMM)** or **vSphere (vCenter)** as the **Platform type**.
 4. In the **Volumes** grid, click **Add Volume** or click  next to an existing volume.
 5. In the volume's **Storage Reservation Policy** list, select a storage reservation policy. If adding a volume, you must specify a value for **Capacity (GB)**.
 6. Click  to save the volume assignment.
 7. To assign storage reservation policies to other volumes, repeat steps 5, 6, and 7.
 8. (Optional) Select the **Allow user to see and change storage reservation policies** check box. Selecting this box allows the user to see and change the storage reservation policy in the volume information on the **Storage** tab of the **Confirm Machine Request** page, and when adding a volume on the **Storage** tab of the **Edit Machine – Virtual** page.
 9. (Optional) Enter a maximum value for **Storage (GB)** in the **Machine Resources** section. Doing so allows the user to increase the size of the volumes in the blueprint and/or add volumes. If you selected the check box in the previous step, the user can also select a storage reservation policy when adding volumes.
 10. Click **OK** to save the configuration.

The **Custom Properties** column in the **Volumes** grid does not accept values for the **Size**, **Letter**, **Label**, and **Storage Reservation Policy** disk properties (disk properties have **Disk** in their full names; for example, **VirtualMachine.DiskN.Size**). To enter a value for one of these disk properties, use its corresponding field in the **Volumes** grid. For example, to specify the **Label** disk property value for a volume, enter the value in the **Label** field when you add or edit the volume. Note that the **Size** disk property corresponds to the **Capacity (GB)** field, the **Letter** disk property to the **Drive Letter/Mount Path** field, and the **Storage Reservation Policy** property to the field of the same name.

Using Linked Clone Blueprints

To assign a storage reservation policy to a volume using a linked clone blueprint:

1. Log on as **Enterprise Administrator** or **Provisioning Group Manager**.
2. Select **Blueprints**.
3. Select a blueprint:
 - **New** – Choose **New Blueprint > Virtual**. On the **Build Information** tab, select **vSphere (vCenter)** as the **Platform type** and **Linked Clone** as the **Action**; then select a vSphere machine and snapshot in **Clone from**.
 - **Existing** – Point to a linked clone blueprint for a virtual machine and click **Edit**.
3. Go to the **Properties** tab.
4. Select **New Property**.
5. In the **Name** field, type **VirtualMachine.DiskN.StorageReservationPolicy** where *N* is the index (starting at 0) of the volume you want to assign to a storage reservation policy.
6. In the **Value** field, type the name of the storage reservation policy to assign to the volume.

7. (Optional) Select the **Prompt User** check box. When you select this option, users can see the property and change the storage reservation policy value during a machine request. When you unselect the option, users cannot see or change the property during a machine request.
8. Click  to save the property.
9. Select **New Property** again.
10. In the **Name** field, type **VirtualMachine.Disk/N.Size** where *N* is the index you used in step 5.
11. In the **Value** field, type the amount of storage capacity (in gigabytes) to assign to the volume.
12. Click  to save the property.
13. Click **OK**.

Requesting a Machine with Storage Reservation Policies

When requesting a machine, you can add, update, and remove storage reservation policies, if the blueprint allows it. If you are using a create, clone, or Flexclone blueprint to provision the machine and you have access to the storage reservation policies, then the **Confirm Machine Request** page displays the **Storage** tab and on it the **Volumes** grid where you can work with the policies. If you are using a linked clone blueprint and you have access to the storage reservation policies, then the **Confirm Machine Request** page does not display the **Storage** tab. Instead, it displays the storage reservation policies as editable property values.

To add, update, or remove storage reservation policies:

1. Click **Self-Service > Request Machine**.
2. Open a provisioning group.
3. Click on a virtual blueprint.
4. On the **Confirm Machine Request** page:
 - (Storage tab with **Volumes** grid is present)
 - a. Go to the **Storage** tab.
 - b. In the **Volumes** grid, click **Add Volume** to create a volume or  to edit an existing volume.
 - c. In the **Storage Reservation Policy** field, select a storage reservation policy, or select blank to remove the current one.
 - d. Click  to save your changes.
 - (Storage tab is *not* present, but editable storage reservation policy names are present)
 - a. Click on the storage reservation policy name you want to change.

Note: Leaving the field blank without a name deletes the property. This is not recommended.

 - b. Edit the name and click **OK**.

Editing a Machine with Storage Reservation Policies

When you edit a provisioned machine (see [Reconfiguring a Machine](#)), you can add a storage reservation policy when you add a volume.

To add a storage reservation policy to a provisioned machine:

1. Click **Self-Service > My Machines**.
2. Point to an existing virtual machine and click **Edit**. The machine's status must be On or Off.
3. Go to the **Storage** tab.
4. In the **Volumes** grid, Click **Add Volume**.

Note: If **Add Volume** is not present on the tab, then the blueprint used to provision the machine does not allow editing of storage on a provisioned machine and you are unable to add a storage reservation policy.

5. In the **Storage Reservation Policy** field, select a storage reservation policy.
6. Click  to save your changes.

Viewing the Machine's Storage

By using multiple datastores, you can have the configuration file on a different datastore than volume 0 and other volumes. Two vCAC Console columns allow you to keep track of all datastores.

To view the locations of the configuration file, the machine's datastores, and the volume/datastore assignments:

1. Based on your privileges, navigate to your machines:
 - **Enterprise Administrator** – Choose **Enterprise Machines**.
 - **Provisioning Group Manager** – Choose **Group Machines**.
 - **Machine Owner** – Choose **My Machines**.
2. Expand the drop-down list of **Columns**.
3. Select the following columns:
 - **Storage Path** – Displays the location of the configuration file.
 - **Disk Storage Path Summary** – Displays the datastores assigned to the machine.
4. Point to a virtual machine and click **Edit**.
5. Click the **Storage** tab.

The **Volumes** grid displays the volumes and their assigned datastores.

Changes to Multiple Datastores since 4.5

If you have used multiple datastores with virtual volumes since the feature became available in release 4.5, please note the following differences introduced in the 5.1 release for the vSphere and SCVMM platform types:

- A user interface now manages storage reservation policies for virtual machines provisioned with Create, Clone, or Flexclone blueprints. For Linked Clone blueprints, you still use custom properties to add storage reservation policies to volumes.
- The **Build Information** tab in Create, Clone, and Flexclone blueprints has a new **Volumes** grid where you add and modify volumes for virtual machines.
- Also on the **Build Information** tab, the minimum value for the **Storage (GB)** field is now read-only. It is the sum of the capacity of the existing volumes.

- The **Properties** tab no longer shows volume-related properties in Create, Clone, and Flexclone blueprints. In release 5.1, you specify these properties with the **Volumes** grid.
- The default value for the **VirtualMachine.DiskN.StorageReservationPolicyMode** custom property has changed to an **Exact** match.

Changes to vSphere and SCVMM Blueprints

The user interface for specifying storage on vSphere and SCVMM blueprints has changed in release 5.1 to enable support for multiple volumes. These changes apply to blueprints that use Create or Clone to provision machines and do not apply to blueprints that use Linked Clone.

Instead of specifying the minimum value for storage capacity as you would for other machine resources, you must add at least one volume to the blueprint on the **Build Information** tab. The **Minimum** field for storage is now read-only and displays the total capacity of the volumes specified on the blueprint. You can still specify a **Maximum** value and an **Approval at** threshold for overall storage capacity.

Additionally, you now enter properties that are specific to a storage volume (that is, properties in the format `VirtualMachine.DiskN.PropertyName`) in the **Volumes** grid on the **Build Information** tab rather than on the **Properties** tab. For more information, see [Using Disk Properties with Volumes](#).

SCVMM Integration

Beginning in release 4.5, vCAC supports provisioning of virtual machines on compute resources within a Microsoft System Center Virtual Machine Manager (SCVMM) 2012 instance. (Only SCVMM 2012 is supported in this release.) SCVMM monitors and enforces policies on the VMs for which it is responsible keeping track of VM activities. When necessary, it provides resources, redirects the VM to resources, or denies access to resources.

SCVMM Requirements

- Each SCVMM instance must be joined to the domain containing the vCAC server.
- The credentials used to manage the endpoint representing an SCVMM instance must have Administrator privileges on the SCVMM server. These credentials must also have Administrator privileges on the Hyper-V servers within the instance.
- Hyper-V servers within an SCVMM instance to be managed by vCAC are Windows 2008 Servers with Hyper-V installed. The processor must be equipped with the necessary virtualization extensions .NET Framework 4 must be installed and Windows Management Instrumentation (WMI) must be enabled.
- At least one Distributed Execution Manager (DEM) Worker within the vCAC instance must be installed on a compute resource on which the SCVMM Console is installed. The DEM Worker must have access to the SCVMM PowerShell module installed with the console, and MS PowerShell Execution Policy must be set to "RemoteSigned" or "Unrestricted." See the *vCloud Automation Center Installation Guide* for more information.
- If all DEM Workers within the vCAC instance are not on compute resources meeting these requirements, Skills must be used to direct all SCVMM-related workflows to those that are. See [Associating Workflows and DEM Workers Using Skills](#).
- To provision machines on an SCVMM compute resource, a user must be added in at least one security role within the SCVMM instance, for example the Self-Service role.

SCVMM Endpoints

vCAC provisions virtual machines on compute resources. To make these compute resources available for provisioning involves two steps:

- Discovering the compute resource
- Adding the compute resource to an enterprise group

When these steps are complete, vCAC initiates data collection to enable creation of reservations on the compute resource. SCVMM Endpoint Data Collection also collects the SCVMM library objects including the following:

- Guest OS Profiles
- Hardware Profiles
- ISOs
- Virtual Hard Disks (VHDs)
- Virtual Machine Templates

If a user creates, updates or deletes one of these library objects, the user must then run SCVMM Endpoint data collection to collect that object, before it can be used in vCAC.

To discover compute resources for SCVMM, create an SCVMM endpoint representing the SCVMM instance.

Creating an SCVMM Endpoint

With SCVMM, there is no agent acting as a proxy between vCAC and the SCVMM endpoint. In this case, every requested action triggers a workflow executed by a DEM Worker. Follow these steps to create an SCVMM endpoint.

To create an endpoint:

1. Click **vCAC Administrator > Endpoints**.
2. Click the **New Endpoint** link in the upper-right corner of the title bar and click **Hyper-V (SCVMM)** to display the New Endpoint - Hyper-V (SCVMM) page.
3. Enter a name for the endpoint and a detailed description. Typically, the name indicates the SCVMM instance that the endpoint is for.
4. Specify the address of the instance using a compute resource FQDN or IP address.
5. Click the button to display the **Credentials** list.
6. Click the **New Credentials** link in the upper-right corner of the title bar.
7. Enter credentials for the SCVMM instance that the endpoint will represent.
8. Enter a name and detailed description.
9. Click the icon to save the credentials.
10. Select them and click **OK** to return to the New Endpoint page.
11. Click the green “go” icon to save the endpoint.

Repeat this procedure for all the needed endpoints.

To import an endpoint:

You can add multiple endpoints of multiple types all at once by loading a CSV file containing the required information.

- Select the **Import Endpoints** link at the right-hand end of the title bar and browse for the CSV file.

The CSV file must contain a header row that determines the order of the fields (which are case-sensitive). Fields include:

- InterfaceType – For SCVMM, enter SCVMM
- Address
- Credentials
- Name
- Description
- DataCenter
- Row
- Rack

Using the Compute Resources List

The Compute Resources page displays a list of compute resources in your enterprise group(s) including SCVMM.

To view and edit compute resources:

- Select **Enterprise Administrator > Compute Resources** in the activity pane.

The following are displayed by default:

- The name of the compute resource and an icon indicating the type.
- Whether regular data collection is enabled.
- The endpoint an SCVMM virtual compute resource is on
- How many reservations have been made on the compute resource and how many machines are provisioned on it.
- Information about resource usage, including all machines allocated from all reservations on the compute resource as a percentage of those reservations' total machine quota (virtual), total reserved memory and storage as a percentage of the compute resource's physical capacity (virtual only).

Note: Hyper-V (SCVMM) virtual compute resources are not managed by proxy agents. Every requested action triggers a workflow executed by a DEM Worker.

Configuring SCVMM

By default the DEM Worker configuration file (**DynamicOps.DEM.exe.config**) points to the standard installation path of the SCVMM Console.

{ProgramFiles}\Microsoft System Center 2012\Virtual Machine Manager\bin

If the SCVMM Console has been installed in another location, the configuration file of the DEM Worker must be updated to change the default path in the assemblyLoadConfiguration section to point to the new folder:

```
<assemblyLoadConfiguration>
  <assemblies>
    <!-- List of required assemblies for Scvmm -->
    <add name="Errors" path="{ProgramFiles}\Microsoft System Center 2012\Virtual
Machine Manager\bin" />
    [...]
  </assemblies>
</assemblyLoadConfiguration>
```

After the DEM .configuration file has been modified and saved, the DEM Worker must be restarted.

Virtual Provisioning on SCVMM

When setting up a virtual machine template in SCVMM, a Guest OS Profile can be added directly to a Windows template using SCVMM Console. See "Preparing for Virtual Provisioning by Cloning" in the *vCloud Automation Center Operating Guide* for more information on Windows and Linux template configuration.

Some restrictions apply to SCVMM template and hardware profile names. Specifically, these names cannot start with the following words:

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

Because of naming conventions that SCVMM and VMware use for temporary templates and hardware profiles, these words are ignored during data collection.

A compute resource running under SCVMM can have multiple paths in the placement section that are collected and exposed to be assigned in a reservation.

On a Hyper-V cluster under SCVMM management, data collection is for Shared Volumes only since a work load can only be provisioned on a shared resource of a cluster.

- View the SCVMM Console, then right-click your Hyper-V cluster to select properties. Browse to the Shared Volumes section to view the storage properties.

When running data collection on Standalone hosts for storage used in the reservation, vCAC collects the default virtual machine paths. This can be configured through SCVMM Console under the Placement section.

- View the SCVMM Console, then right-click your Hyper-V standalone host to select properties. Browse to the Placement section to view the storage properties.

Amazon Elastic Compute Cloud (EC2) Enhancements

This section describes enhancements and new features for Amazon EC2 made in release 4.5:

- Endpoints (enhancement)
- Regions
- AMI selection (enhancement)
- Security groups
- Virtual Private Cloud (VPC)
- Key pairs
- Elastic Load Balancer
- Elastic IP Addresses
- Elastic Block Storage

For information on EC2 functionality in the previous release, refer to the following sections in the *vCloud Automation Center Operations Guide*:

Feature	Section
Credentials	Access and Credentials
Endpoints	Creating an Amazon EC2 Endpoint
Reservations	Creating a Cloud Reservation
Blueprints	Creating a Blueprint for Amazon EC2 Provisioning
Amazon Machine Images	Amazon Machine Images
Instance Types	Instance Types

Provisioning Workflow

The following table provides a high-level summary of the workflow to provision a cloud machine. As you read about new EC2 features, refer to the table to see where the feature fits in the overall provisioning process.

Workflow Item	Description	Role
1. Credentials	Manage the credentials associated with endpoints.	vCAC Administrator
2. Endpoints	Manage the endpoints representing virtualization platforms, cloud service accounts, physical machines and storage devices.	vCAC Administrator
3. Enterprise Group	Create an enterprise group by assigning enterprise administrators and selecting the compute resources they can manage.	vCAC Administrator
4. Key Pairs	Manage the key pairs associated with Amazon access keys.	Enterprise Administrator

Workflow Item	Description	Role
5. Provisioning Group	Create and manage provisioning groups, by specifying the members and roles within it.	Enterprise Administrator
6. Compute Resources	Manage existing compute resources, view or add reservations, and force rediscovery.	Enterprise Administrator
7. Machine Prefix	Add, edit or delete the prefixes used on the blueprint to create machine names.	Enterprise Administrator
8. Reservation Policy (Optional)	Add a reservation policy to specify a set of reservations. In the machine blueprint, you can restrict provisioning to only those reservations in the policy.	Enterprise Administrator
9. Reservation	Create and manage cloud reservations that are accessible to a provisioning group.	Enterprise Administrator
10. Blueprint	Create and manage cloud blueprints that are accessible to provisioning groups.	Provisioning Group Manager
11. Request Machine	Select a cloud blueprint and request a cloud machine as specified in the blueprint. A provisioned cloud machine is also referred to as an instance.	Basic User, sometimes referred to as Machine Owner

Endpoints

There are two enhancements to Amazon EC2 endpoints.

One Endpoint per Access Key

Previously, a credential associated with an Amazon Access Key ID could be associated with multiple EC2 endpoints. During data collection, the same information was collected for multiple endpoints.

Now, only one EC2 endpoint can be associated with an Access Key ID.

Proxy Server

For additional security, you can make all connections to Amazon Web Services go through a proxy server.

To use a proxy server:

1. Log in as **vCAC Administrator > Endpoints**.
2. Click **New Endpoint > Amazon EC2** or point to an existing cloud endpoint and click **Edit**.
3. Select **Use proxy server** and enter the following information:
 - Hostname – Name of your proxy server
 - Port – Port number on proxy server for the connection
 - Credentials (Optional) – Username and password for the proxy server, if required by proxy configuration
4. Click **OK**.

Regions

Each Amazon Web Service account is represented by a cloud endpoint. When you create an endpoint in vCAC, regions are data collected. For example, us-east-1 and us-west-2 are regions.

For more information on regions, refer to Amazon documentation:

<http://docs.amazonwebservices.com/general/latest/gr/rande.html>

Terminology

vCAC refers to regions as compute resources.

Functionality by Role

- vCAC Administrator can select compute resources for an enterprise group.

After the vCAC Administrator selects compute resources for an enterprise group, inventory and state data collection occur automatically.

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

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

State data collection occurs automatically every 15 minutes by default. It gathers information about the state of managed instances, which are instances vCAC creates. Examples of state data include:

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

Note: Data is never collected for compute resources that do not belong to an enterprise group.

- Enterprise Administrator can manually start inventory and state data collection on compute resources, as well as disable or change the frequency of the inventory and state data collection.

Adding Compute Resources to an Enterprise Group

1. Log in as **vCAC Administrator > Enterprise Groups**.
2. Choose **New Enterprise Group** or point to an existing enterprise group and click **Edit**.
3. In the **Compute resources** list, select each resource for the enterprise group.
4. Click **OK**.

Manually Starting Data Collection on a Compute Resource

1. Log in as **Enterprise Administrator > Compute Resources**.
2. Point to an existing compute resource and click **Data Collection**.

A page opens that includes four sections:

- **Compute Resource**
- **Inventory**
- **State**

For the selected compute resource, you see when the last data collection was completed, if data collection is currently on or off and a frequency box to change the automatic schedule.

3. To manually start **Inventory** or **State** data collection, go to the section and click **Request Now**.
4. Click **OK**.

Then, return to check the status of the inventory and state data collections you started.

Key Pairs

A key pair is required to provision a cloud instance. It is used to decrypt Windows passwords and to log in to a Linux machine.

There are two components to a key pair. One component is public and the other is private and unique. Those components are:

- Name – This is the public component. vCAC discovers the name during data collection.
- Secret key – This is the private component.

For more information on key pairs, refer to Amazon documentation:

<http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/generating-a-keypair.html>

Functionality by Role

- Enterprise Administrator can:
 - Manage key pairs:
 - Create a new Amazon key pair
 - Upload .PEM file contents to an existing Amazon key pair
 - Export .PEM file contents for an existing key pair
 - Delete a key pair
 - Specify a method to auto-generate key pairs for machines provisioned on a reservation:
 - Per machine – Each machine has a unique key pair. This is the more secure method because no key pairs are shared among machines.
 - Per provisioning group (Default) – Every machine provisioned on a reservation has the same key pair. This setting applies to machines provisioned on other reservations when the machine has the same compute resource and provisioning group.
- Provisioning Group Manager can:
 - Specify a method to auto-generate key pairs for machines provisioned on a blueprint:
 - Per machine

- Per provisioning group
- Basic User can export .PEM file contents to a certificate.

Managing Key Pairs

This section describes tasks you perform with key pairs.

To create a new key pair:

1. Log in as **Enterprise Administrator > Key Pairs**.
2. Click **New Key Pair**.
3. Provide the following required information:
 - **Key pair** – Type a name.
 - **Compute resource** – Select a compute resource assigned to your enterprise group.
4. Click  to save the key pair.
5. Press **F5** to refresh the page. The key pair is ready to use when **Being Created** changes to
*****.

To upload .PEM file contents for a key pair:

1. Log in as **Enterprise Administrator > Key Pairs**.
 - Go to the key pair name that will use the .PEM file, Click  next it.
2. Use one of the following methods to upload .PEM file contents:
 - In the **Secret key** field, type or copy and paste the private key, including both of the following lines and all text between them: -----BEGIN RSA PRIVATE KEY----- and -----END RSA PRIVATE KEY-----.
 - Click **Browse**, select a .PEM file and click **Open**. Then click **Upload** for the .PEM file contents to appear in the **Secret key** field.
3. Click  to save the key pair.

To export .PEM file contents for a new key pair:

1. Follow the steps to create a new key pair or to upload the secret key for an existing key pair.
2. On the **Key Pairs** page, in the **Secret Key** column, the new key pair appears as a string of asterisks followed by an **Export** icon.
***** 
3. Click  and click **Save**.

To delete an unused key pair:

Caution: *This procedure permanently deletes the key pair.*

The following steps delete the key pair from the local database and AWS.

1. Log in as **Enterprise Administrator > Key Pairs**.
2. Click  next to the key pair to delete.
3. Click **OK** to confirm.

Managing Key Pairs on a Reservation

To select a key pair on a reservation:

1. Log in as **Enterprise Administrator > Reservations**.
2. Click **New Reservation > Cloud** or point to an existing reservation and click **Edit**.
3. Go to the **Resources** tab.
4. For **Key pair**, select **Specific Key Pair** and click .
5. In the **Key Pairs** page, select a key pair that has a secret key.

To specify how to auto-generate key pairs for machines provisioned on a reservation:

1. Log in as **Enterprise Administrator > Reservations**.
2. Click **New Reservation** or point to an existing reservation and click **Edit**.
3. Go to the **Resources** tab.
4. For **Key pair**, choose an option from the list:
 - Auto-Generated per Provisioning Group – Every machine provisioned on a reservation has the same key pair. Machines provisioned on other reservations also have the same key pair when the machine has the same compute resource and provisioning group.
Because key pairs generated this way are associated with a provisioning group, the key pairs are deleted when the provisioning group is deleted.
 - Auto-Generated per Machine – Each machine has a unique key pair.
5. Click **OK**.

Managing Key Pairs on a Blueprint

To specify how to auto-generate key pairs for machines provisioned on a blueprint:

Prerequisite: Key pairs must not be specified on the reservation, which overrules the blueprint.

1. Log in as **Provisioning Group Manager > Blueprints**.
2. Click **New Blueprint > Cloud** or point to an existing blueprint and click **Edit**.

3. Go to the **Build Information** tab.
4. Refer to the previous task about auto-generating key pairs. Follow steps 4-5.

To export .PEM file contents to a certificate:

1. Log in as **Basic User > My Machines**.
1. Point to a machine and click **Export Certificate**.
2. Click **Save**.

Virtual Private Cloud

Amazon Virtual Private Cloud (VPC) allows users to provision a private, isolated section of the Amazon Web Services (AWS) Cloud where instances can be launched in a virtual network you design. AWS users can create a VPC to design a virtual network topology according to your specifications.

You can assign a VPC in vCAC. However, vCAC does not track the cost of using it.

Prerequisites

You must use the AWS management console to create:

- An Amazon VPC, which includes Internet gateways, routing table, security groups and subnets. The Classless Inter-Domain Routing (CIDR) block determines the available IP addresses.
- A VPN if you want to login to machines outside of the AWS console.

Functionality by Role

vCAC users can perform the following tasks when working with a VPC:

- Enterprise Administrator can assign a VPC and subnets within it to a cloud reservation.
- Basic User can assign an instance to a VPC and subnets within it.

Caveats

When the basic user provisions a machine, assignments for VPC and non-VPC locations are random when both of the following conditions exist:

- The cloud blueprint has access to reservations with VPC and non-VPC locations.
- The cloud blueprint does not display the **Network** tab, which is the default.

To prevent this random assignment, the Provisioning Group Manager has two options:

- Display the **Network** tab by selecting the following option in the **Build Information** tab of the cloud blueprint:
Enable Amazon network options on the machine
- Tie the blueprint to the desired reservation by using a reservation policy

Using a Virtual Private Cloud

This section describes each task you perform with a VPC.

To assign a VPC and subnets within it to a cloud reservation:

1. Log in as **Enterprise Administrator > Reservations**.
2. Click **New Reservation > Cloud** or point to an existing cloud reservation and click **Edit**.
3. Go to the **Resources** tab.
4. Select the **Assign to a subnet in a VPC** checkbox.
5. In the **VPC** list, click  next to the VPC you want to assign.
6. In the **Edit VPC** window:
 - In the **Subnets** list, select each subnet in the VPC that will be available for provisioning.
 - In the **Security Groups** list, select each group that can be assigned to a machine during provisioning.
 - **Load balancers** that have the selected subnets display in a list. You can select load balancers to use in the VPC.
7. Click  to save the VPC settings.
8. If you want to assign additional VPCs, repeat steps 6-8.
9. Click **OK** to save the reservation.

To assign an instance to a VPC and subnets within it:

Prerequisite: The **Enable Amazon network options** option must be selected in the blueprint:

1. Log in **Basic User**, or any role that can request a machine.
2. Select **Self-Service > Request Machine**.
3. Click a cloud blueprint.
4. On the **Request Information** tab, click the **Provision into** list box and select **Subnet in a VPC**.
5. On the **Network** tab, click the **Subnet** list box to select a subnet within the VPC.
The VPC and Location of the selected subnet display in boxes below.
6. Click **OK**.

Security Groups

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

For more information on security groups, refer to Amazon documentation:

<http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/using-network-security.html>

Prerequisites

In the Amazon EC2 Management Console, you must create security groups, configure ports for RDP or SSH and set up a VPN for VPCs.

Functionality by Role

- Enterprise Administrator must select security groups on a reservation.
- Basic User can:
 - Select security groups when requesting a machine
 - View security groups on an existing machine

Selecting and Viewing Security Groups

This section describes tasks you perform with security groups.

To select security groups on a reservation:

1. Log in as **Enterprise Administrator > Reservations**.
2. Click **New Reservation > Cloud** or point to an existing cloud reservation and click **Edit**.
3. After the **Reservation Information** tab is complete, go to the **Resources** tab.
4. For a non-VPC reservation, in the **Security groups** list, select each group. Then click **OK**.
5. For a VPC reservation, select the **Assign to a subnet in a VPC** checkbox:
 - In the **VPC** list, click  next to the VPC you want to assign.
 - In the **Security Groups** list, select each group that can be assigned to a machine during provisioning.
 - Click  to save the VPC settings.

For information on other VPC settings, see the [Virtual Private Cloud](#) section.

To select security groups for a machine:

Prerequisites: You must specify that you are provisioning in a VPC or non-VPC. Otherwise, the machine would be associated to all security groups on the reservation and could be created in either a VPC or non-VPC location.

To specify VPC or non-VPC, the Provisioning Group Manager must select **Enable Amazon network options** in the blueprint.

1. Log in as **Basic User > Request Machine**.
2. For a non-VPC reservation, on the **Request Information** tab:
 - In the **Provision into** field, select **Non-VPC Location**.
 - Select a **Location**.
3. For a VPC reservation:
 - On the **Request Information** tab in the **Provision into** field, select **Subnet in a VPC**.
 - On the **Network** tab, select a **subnet**.
4. On the **Network** tab in the **Security groups** list, select each group the machine will belong to when it is provisioned

Note: If you do not select security groups, all security groups from the reservation will be applied to the machine.

5. Click **OK**.

To view security groups for a machine:

1. Log in as **Basic User > My Machines**.
2. Point to a cloud machine and click **Edit**.
3. Go to the **Network** tab.

Elastic Load Balancers

Elastic Load Balancers automatically distribute incoming application traffic across multiple Amazon EC2 instances. They enable you to achieve greater fault tolerance and performance in EC2 applications by providing the amount of load balancing capacity needed in response to incoming application traffic.

vCAC automatically discovers Elastic Load Balancers during routine inventory data collections or when an Enterprise Administrator runs a manual inventory data collection on selected compute resources. vCAC does not create, delete, manage or track the cost of using Elastic Load Balancers.

For more information, refer to AWS documentation on Elastic Load Balancing:

<http://aws.amazon.com/elasticloadbalancing/>

Prerequisites

You must create Elastic Load Balancers in the AWS Management Console.

The Elastic Load Balancer must be available in the EC2, VPC and provisioning location. For example, a load balancer is available in us-east1c and a machine location is us-east1b. The machine cannot use the available load balancer.

Functionality by Role

vCAC users can perform the following load balancing tasks:

- Enterprise Administrator can select load balancers on a reservation.
- Provisioning Group Manager can specify a custom property to assign a load balancer to machines provisioned on the blueprint.
- Basic user can:
 - Associate an instance to and disassociate an instance from a load balancer, provided that the following checkbox is selected in the blueprint:
Enable Amazon network options
 - View the following information about a load balancer:
 - Status
 - Port configuration
 - Locations
 - Ping target, timeout and interval

- Health
- Health thresholds
- Total number of instances assigned to a load balancer and how many of them are healthy.

Caveats

You can associate a vCAC-managed instance to a load balancer in two places:

- AWS Management Console – The association you make in AWS will show in vCAC after data collection.
- vCAC – You make the association on the reservation. If the load balancer is removed from the reservation, the association is lost.

Using Load Balancers

To select load balancers on a reservation:

1. Log in as **Enterprise Administrator > Reservations**.
2. Click **New Reservation > Cloud** or point to an existing cloud reservation and click **Edit**.
3. Go to the **Resources** tab.
4. The list of load balancers in the **Load balancers** list depends on whether **Assign to a subnet in a VPC** is selected:
 1. Non-VPC reservation – The list shows only load balancers that belong to selected locations.
 2. VPC reservation – In step 7, the list will show only load balancers that have selected subnets.
5. For a non-VPC reservation, select the **Location**. Then select each load balancer you want to be available for the reservation. Go to step 7.
6. For a VPC reservation, click  next to the VPC you want to assign. The **Edit VPC** window opens.
7. Select the following elements in the **Edit VPC** window:
 - In the **Subnets** list, select each subnet in the VPC that will be available for provisioning.
 - In the **Security Groups** list, select each group that can be assigned to a machine during provisioning.
 - In the **Load balancers** list, select each load balancer to use in the VPC.
 - Click  to save the VPC settings.
8. Click **OK**.

To associate and disassociate an instance from a load balancer:

Prerequisite: The **Enable Amazon network options** checkbox must be selected in the blueprint.

1. Log in as **Basic User**.

2. Select **Self-Service > My Machines**.
3. Point to an existing cloud machine and click **Edit**.
4. Go to the **Network** tab.
5. In the **Load balancers** list:
 3. Select each load balancer to assign to the machine.
 4. Deselect each load balance to disassociate it from the machine.
6. Click **OK**.

To view detailed information about a load balancer:

1. Log in as **Basic User**.
2. Select **Self-Service > My Machines**.
3. Point to an existing cloud machine and click **Edit**.
4. To view information about a load balancer, click  next to it.

Using Custom Properties

Custom Property Name:

Amazon.ElasticLoadBalancer.Names

Purpose:

To assign Elastic Load Balancers, that match values you specify in Amazon.ElasticLoadBalancer.Names, to provisioned machines

Best Practice:

When you use this custom property, we recommend using a reservation policy to restrict provisioning to the appropriate reservation. Without a policy, provisioning cannot be guaranteed.

Example:

There are two reservations:

- Reservation 1 – Region is US-East, with Subnet 1 of VPC 1 selected.
- Reservation 2 – Region is US-East, with Subnet 1 of VPC 2 selected.

On the blueprint, this custom property specifies a load balancer that is created for VPC1.

Because you cannot share a load balancer between VPCs, a reservation policy ensures that machines requested from the blueprint are provisioned on Reservation 1.

Procedure:

To specify a custom property that assigns a load balancer to a machine provisioned on a blueprint:

1. Log in as **Provisioning Group Manager > Blueprints**.
2. Click **New Blueprint > Cloud** or point to an existing cloud blueprint and click **Edit**.
3. Go to the **Properties** tab.
4. In the **Custom properties** list, click **New Property**.
5. In the **Name** field, type Amazon.ElasticLoadBalancer.Names.
6. In the **Value** field, type the elastic load balancer name. You can specify multiple load balancers by typing a comma-separated list of names.

7. Click  to save the property.
8. Click **OK** to save the blueprint.

Elastic IP Addresses

Elastic IP Addresses are static IP addresses designed for dynamic cloud computing. An Elastic IP Address:

- Is associated with an AWS account
- Is controlled by the account owner until the owner chooses to release it or a custom property takes it
- Is available to all provisioning groups that have rights to the region
- Can be assigned to an instance

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

Prerequisites

You must allocate Elastic IP Addresses to your account in the Amazon AWS Management Console. Also, you must allocate one range of addresses in a non-VPC region and another range in a VPC.

Functionality by Role

- Provisioning Group Manager can:
 - Assign an Elastic IP Addresses to a machines after it is provisioned
 - Specify a custom property to automate the process of assigning an Elastic IP Address to an instance during provisioning
- Basic User can view the Elastic IP Address assigned to an instance.

Caveats

- In Amazon, only one instance at a time can use the same Elastic IP Address.
- When you use a custom property to assign an Elastic IP Address, the following behavior occurs to an instance that was previously assigned to the address.
 - Non-VPC – Amazon takes the Elastic IP Address from an existing instance and assigns it to a new instance.
 - VPC – Amazon fails to assign the Elastic IP Address.
- You can use vCAC to assign an Elastic IP Address. However, vCAC does not track the cost of using the address.

Using Elastic IP Addresses

This section describes tasks you perform with Elastic IP Addresses.

To assign an Elastic IP Addresses to an existing instance:

1. Log in as **Provisioning Group Manager > Group Machines**.
2. Point to a cloud machine and click **Edit**.
3. Go to the **Network** tab.
4. In the **Elastic IP address** list, select an address.
5. Click **OK**.

To view the Elastic IP Address assigned to an instance.

Prerequisite: This capability must be allowed on the blueprint.

1. Log in as **Basic User > Edit Machine**.
2. Point to a cloud machine and click **Edit**.
3. Go to the **Network** tab.

Using Custom Properties

Custom Property Name:

Amazon.ElasticIpAddress.IpAddress

Purpose:

Machines provisioned via this blueprint will be assigned the Elastic IP Address with the corresponding address. If the Elastic IP Address is in use, the workflow will try to transfer it to the newly-created Amazon instance. If multiple instances are created in the machine request, the specified Elastic IP Address will be assigned to the last machine created.

Best Practice:

Before you specify a custom property, you should use a reservation policy as a bridge to link a blueprint to the correct reservation. Without a reservation policy, successful provisioning cannot be guaranteed. Provisioning could fail because the blueprint could use a reservation that does not have the desired resources.

The following example illustrates this best practice.

Example:

You have two reservations:

- Reservation 1 is for region US-East with location us-east-1b.
- Reservation 2 is also for region US-East, with subnet 1 created on location us-east-1b.

In the blueprint, you specify this custom property:

- Name: Amazon.ElasticIpAddress.IpAddress
- Value: Non-VPC elastic IP address

You need to make sure that when a user requests a machine from this blueprint, the machine ends up on Reservation 1. Because it is not possible to assign a non-VPC Elastic IP Addresses to a machine provisioned on a subnet in a VPC, Reservation 2 would not work.

To ensure this, you create a reservation policy and apply it to Reservation 1 and the blueprint.

Procedure:

To automatically assign an Elastic IP Address to an instance during provisioning:

4. Log in as **Provisioning Group Manager > Blueprints**.
5. Select a blueprint:
 - New – Choose **New Blueprint > Cloud**.
 - Existing – Point to a cloud machine and click **Edit**.
6. Go to the **Properties** tab.
7. In the **Custom properties** list, select **New Property**.
8. In the **Name** field, type Amazon.ElasticIpAddress.IpAddress.
9. In the **Value** field, type an IPv4 elastic IP address.
10. Click  to save the property.

The Provisioning Group Manager can also assign an Elastic IP Address after an instance is provisioned in **Group Machines**.

Elastic Block Storage

Amazon Elastic Block Store (EBS) provides block level storage volumes to use with an Amazon EC2 instance and VPC. Amazon EBS volumes are off-instance storage that can persist independently from the life of an instance.

Functionality by Role

- Basic User can:
 - Request and attach additional storage capacity, up to the maximum storage capacity specified in the blueprint, when provisioning an instance or when editing an existing instance
 - Attach and detach an EBS volume from an existing instance
 - Modify existing EBS volumes
 - Delete an EBS volume
- Provisioning Group Manager can limit, on the cloud blueprint, the maximum threshold of elastic block storage a user can request.
- Enterprise Administrator can destroy an unattached EBS volume that belongs to a compute resource in an enterprise group.
- Only a user with both Provisioning Group Manager and Enterprise Administrator privileges can change the owner of an EBS volume.

Caveats

When you use an EBS volume in vCAC, the following caveats apply:

- You can never attach an *existing* elastic block storage volume when you provision an instance.

However, if you create a new volume and request more than one machine at a time, the volume will be created and attached to each instance.

For example, if you create one volume named Volume 1 and request three machines, a volume will be created for each machine. Three volumes named Volume 1 will be created and attached to each machine. Each volume will be the same size and in the same location.

- The volume must be the same OS and in the same location as the machine to which you attach it.
- vCAC does not track the cost of using an EBS volume.
- vCAC does not manage the primary EBS volume of an EBS-backed instance.

Using Elastic Block Storage

This section describes each task you perform with EBS.

To request additional storage capacity for a new instance:

1. Log in **Basic User**.
2. In Self-Service, click **Request Machine**.
3. Point to a cloud blueprint and click **Edit**.
4. Go to the **Storage** tab. The number of GB attached and available is shown at the bottom.
5. Click  or **Add volume** and enter the following information for the EBS volume to add:
 5. **Name**
 6. **Storage (GB)** – Select the amount of storage without exceeding the amount available
 7. **Device** – Select the platform. Devices in the xvdk format are for Windows instances and devices in the /dev/sdk format are for Linux instances.
 8. **Description** – You can enter a description.
6. Click  to save the volume.
7. Click **OK** to submit the request.

To modify existing storage capacity for a provisioned instance:

1. Log in **Basic User**.
2. In Self-Service, click **My Machines**.
3. Point to a cloud machine and click **Edit**.
4. Go to the **Storage** tab. The number of GB attached and available is shown at the bottom.
5. Click  next to the volume to modify.
6. For volumes you own that are not attached to other machines, you can change the following fields:
 - **Attached** – Check the box to attach the volume or clear the checkbox to detach it.
 - **Name**

- **Description**

7. Click  to save the changes.
8. Click **OK** to submit the request.

To delete an EBS volume as a Basic User:

Best Practice: Log in to the instance and unmount the volume before you delete it. This prevents the system from crashing and the data from being corrupted.

1. Log in **Basic User**.
2. In Self-Service, click **My Machines**.
3. Point to a cloud machine and click **Edit**.
4. Go to the **Storage** tab.
5. Click  next to the volume to delete.
6. Click **OK** to confirm.
7. Click **OK** to submit the request.

To limit the maximum threshold of an EBS volume:

1. Log in as **Provisioning Group Manager > Blueprints**.
2. Click **New Blueprint > Cloud** or point to an existing cloud blueprint and click **Edit**.
3. Go to the **Build Information** tab.
4. To set a threshold, in **EBS storage (GB)** select a number in the **Maximum** column.
For unlimited storage capacity, leave the field blank. To disable the feature, set capacity to 0.
5. Click **OK**.

To change the owner of an EBS volume:

Prerequisite: You must have both **Enterprise Administrator** and **Provisioning Group Manager** privileges.

1. In the **Enterprise Administrator** menu, select **EBS Volume Management** to see a list of all volumes in the compute resources assigned to the enterprise group.
2. Click  next to the volume you want to modify. The **Edit EBS Volume** window opens and provides the following information:
 - Name
 - Description
 - Device Size (GB)
 - Owner
3. In the **Owner** field, type or select a name from the list.

4. Click  to save.

To delete an unattached EBS volume as an Enterprise Administrator:

1. Log in as **Enterprise Administrator > EBS Volume Management**.
2. Click  next to the unattached volume to delete.
3. Click **OK** to confirm.

To delete all EBS volumes when destroying a machine:

1. Log in as **Basic User**.
2. In Self-Service, click **My Machines**.
3. Point to a cloud machine and click **Destroy**.
4. Check the **Delete all attached EBS Volumes?** checkbox.
5. Click **OK** to confirm.

Amazon Image (AMI) Selection

When you select an Amazon image for a cloud blueprint, you can see the region to which the AMI belongs. The Select Amazon Machine Images list now has a sortable Regions column.

Property Dictionary

The property dictionary feature, introduced in release 4.5, enables an enterprise administrator to provide a more robust user interface for custom properties that a machine owner enters at request time.

Properties are used throughout the product to provide settings for many features. When users request new machines they are prompted for any required properties. Enterprise administrators or provisioning group managers designate which properties are required by selecting the **Prompt User** option on the blueprint or build profile. By default, the Confirm Machine Request page displays the literal name of the property as a required text box and does not provide any validation other than that a value has been entered.

The property dictionary allows you define characteristics of properties that are used to tailor the behavior of the request user interface, for example:

- Associate a property name with a particular user control, such as a check box, calendar control, or drop-down list
- Specify constraints such as minimum and maximum values or validation against a regular expression
- Provide descriptive display names for properties or specify text (for a tooltip or text label) with additional information
- Designate a property as optional rather than required
- Group sets of property controls together and specify the order in which they appear by using control layouts

You can create property definitions for reserved custom properties that are used by core vCAC features as well as any new custom properties that you define. For a list of reserved custom properties, see the *vCloud Automation Center Operating Guide*.

New custom properties are typically used by custom workflows that specify the logic to execute based on the value of that property. You can define these custom workflows using either vCAC Designer or the vCAC Development Kit. For more information about working with custom workflows, see the *vCloud Automation Center Extensibility Guide*.

Creating Property Definitions

A property definition associates a property by name to attributes which define its behavior.

The high-level process for using property definitions is as follows:

1. Create a new property definition and select the control type to use for editing the property. The name of the property in the property dictionary is the same name that you add to a blueprint or build profile.
2. Specify any relevant property attributes for the new property. Attributes contain additional information such as minimum and maximum values, or help text to display with the property. Attribute types can vary depending on the control type.
3. To enable the interface for editing the property, add the property to a blueprint or build profile and select **Prompt User**.

The following sections describe how to accomplish each of these steps. For a detailed example of creating and using a property definition, see [Property Definition Example: Network Drop-Down List](#).

User Controls for Property Definitions

Each property definition is associated with a particular control type that determines the user interface for editing that property.

The property dictionary provides the following types of user controls for custom properties:

Control Type	Description
CheckBox	Check box for specifying true or false values.
DateTimeEdit	A calendar and time control that enables users to specify a date and time.
DropDown	Drop-down-style combo box. A user can enter arbitrary text in addition to selecting from the list.
DropDownList	Drop-down list-style combo box. A user must select a value from the list.
Integer	Numeric spin box. A user can enter an integer between the minimum and maximum values or step up and down by a defined interval.
Label	Displays a read-only text label.
Link	Displays a link with the property display name as the link text and the property value as the URL.
Notes	Multi-line text box.
Password	Text box that masks the current value.
TextBox	Single-line text box.

Some control types may have additional configuration, for example, a drop-down property has a set of values from which a user can select and an integer property can have minimum and maximum values.

Defining a Basic Property

This procedure describes how to create a property definition with no attributes. Most control types do not require any attributes to be configured.

To define a basic property:

1. Click **Enterprise Administrator > Property Dictionary**.
2. Click **Add Property Definition**.
3. In the **Name** box, specify the property name. This name must be an exact match (case- and white space–sensitive) to the name that you specify in the blueprint or build profile.
4. In the **Display Name** box, specify a descriptive name for the property. This text appears as the field label on the Confirm Machine Request page.
5. Optionally, specify a **Description**. This text describes the behavior to the property dictionary administrator and is not displayed to end users.
6. From the **Control Type** list, select the user control that you want to associate with this property.
7. Select **Required** if the machine owner must specify a value for this property.
8. Click the green check mark to save the definition.

Defining a Date-Time Property

A Date-Time property displays as a calendar and time control that enables a user to specify a date and time value.

To define a date-time property:

1. Click **Enterprise Administrator > Property Dictionary**.
2. Click **Add Property Definition**.
3. In the **Name** box, specify the property name. This name must be an exact match (case- and white space–sensitive) to the name that you specify in the blueprint or build profile.
4. In the **Display Name** box, specify a descriptive name for the property. This text appears as the field label on the Confirm Machine Request page.
5. Optionally, specify a **Description**. This text describes the behavior to the property dictionary administrator and is not displayed to end users.
6. From the **Control Type** list, select **DateTimeEdit**.
7. Select **Required** if the machine owner must specify a value for this property.
8. Click the green check mark to save the definition.
9. Optionally, specify a minimum (earliest) value for the property.
 - a. Under **Property Attributes**, click **Edit**.
 - b. Click **Add Property Attribute**.
 - c. From the **Type** list, select **MinValue**.
 - d. Specify a **Name** (the attribute name is not displayed to the end user).

- e. In the **Value** box, specify a date in **YYYY-MM-DD** or **MM/DD/YYYY** format or a date followed by a time in **HH:MM** format (24-hour clock or followed by **AM** or **PM**).
- f. Click the green check mark to save the attribute.
- g. Click **OK**.

Defining an Integer Property

An integer property displays as a spin box that enables a user to specify an integer value.

To define an integer property:

1. Click **Enterprise Administrator > Property Dictionary**.
2. Click **Add Property Definition**.
3. In the **Name** box, specify the property name. This name must be an exact match (case- and white space–sensitive) to the name that you specify in the blueprint or build profile.
4. In the **Display Name** box, specify a descriptive name for the property. This text appears as the field label on the Confirm Machine Request page.
5. Optionally, specify a **Description**. This text describes the behavior to the property dictionary administrator and is not displayed to end users.
6. From the Control Type list, select **Integer**.
7. Select **Required** if the machine owner must specify a value for this property.
8. Click the green check mark to save the definition.
9. Optionally, specify the minimum, maximum, and interval values for the property.
 - a. Under **Property Attributes**, click **Edit**.
 - b. Click **Add Property Attribute**.
 - c. From the **Type** list, select **MinValue**.
 - d. Specify a **Name** (the attribute name is not displayed to the end user).
 - e. In the **Value** box, specify an integer for the minimum value of the property.
 - f. Click the green check mark to save the attribute.
 - g. Repeat steps b-f for the **MaxValue** and **Interval**.

The interval is the increment by which the property value is increased or decreased when a user clicks the up or down arrows.
 - h. Click **OK**.

Defining a Drop-Down Property

A drop-down property enables a user to select from a list of string values. You can optionally enable the user to enter arbitrary values.

To define a drop-down property:

1. Click **Enterprise Administrator > Property Dictionary**.
2. Click **Add Property Definition**.
3. In the **Name** box, specify the property name. This name must be an exact match to the name that you specify in the blueprint or build profile.

4. In the **Display Name** box, specify a descriptive name for the property. This text appears as the field label on the Confirm Machine Request page.
5. Optionally, specify a **Description**. This text describes the behavior to the property dictionary administrator and is not displayed to end users.
6. From the Control Type list, select **DropDown** for a control that enables the user to enter arbitrary values or **DropDownList** for a control that limits the user to selecting from the values in the list.
7. Select **Required** if the machine owner must specify a value for this property.
8. Click the green check mark to save the definition.
9. Specify the values for the drop-down list.
 - a. Under **Property Attributes**, click **Edit**.
 - b. Click **Add Property Attribute**.
 - c. From the **Type** list, select **ValueList**.
 - d. Specify a **Name** (the attribute name is not displayed to the end user).
 - e. In the **Value** box, specify a comma-separated list of values that the user can select from. Enclose any values that contain commas in quotation marks (").
 - f. Click the green check mark to save the attribute.
 - g. Click **OK**.

Configuring Common Property Attributes

This section describes aspects of a property definition can apply to multiple property types.

Specifying Tooltip Text for a Property

You can specify help text that displays as a tooltip for any property type except for a link property.

To specify help text for a property:

1. Under **Property Attributes**, click **Edit**.
2. Click **Add Property Attribute**.
3. From the **Type** list, select **HelpText**.
4. Specify a **Name** (the attribute name is not displayed to the end user).
5. In the **Value** box, specify the text that you want to display as a tooltip when the user pauses over the property field.
6. Click the green check mark to save the attribute.
7. Click **OK**.

Specifying Display Order for a Property

You can specify an order index for a property that determines the order in which it appears on the Confirm Machine Request page.

All properties that have a specified display order display before properties with no display order.

Another way to designate the order in which properties are displayed is to create a control layout. For more information, see [About Control Layouts](#).

To specify the display order:

1. Under **Property Attributes**, click **Edit**.
2. Click **Add Property Attribute**.
3. From the **Type** list, select **OrderIndex**.
4. Specify a **Name** (the attribute name is not displayed to the end user).
5. In the **Value** box, specify an alphanumeric value that represents the order in which a property should display. Index values do not have to be sequential.

Note: Index values are sorted alphabetically. For example, if you use numerical indexes from 1-12, they are sorted in the following order: 1, 10, 11, 12, 2, 3 ...

6. Click the green check mark to save the attribute.
7. Click **OK**.

Configuring Regular Expression Validation

You can specify a regular expression against which to validate text input in a text box, drop-down, notes, or password field.

Note: Regular expression validation only applies if a value is supplied; it does not make the property required. If you create a password field with a regular expression attribute that enforces a length of 6-8 characters and do not make the property required, vCAC does not attempt to validate a blank password. However, if the user enters a password that is only 4 characters long, then it would fail validation.

To specify a regular expression to use to validate a property:

1. Under **Property Attributes**, click **Edit**.
2. Click **Add Property Attribute**.
3. From the **Type** list, select **RegEx**.
4. Specify a **Name** (the attribute name is not displayed to the end user).
5. In the **Value** box, specify the regular expression to use to validate the property value.
6. Click the green check mark to save the attribute.
7. Click **OK**.

Defining Relationships between Properties

You can define a relationship between a drop-down property and another property such that the value of one property determines the possible values of the drop-down property.

In a property relationship, one property is the *parent* (the one whose value determines the behavior of the child property control) and the other is the *child* (the drop-down property whose values are populated based on the value of the parent). When the value of the parent property changes, the child property is refreshed with a list of possible values depending on the value of the parent.

Instead of specifying a ValueList for the child property, you specify the mapping between the values of the parent property and the possible values of the child property in an XML string called a *value expression*.

The high-level process for creating relationships is as follows:

1. Create two property definitions, the parent and child.
2. Add a **Relationship** attribute to the child property whose value is the name of the parent property.
3. Create a value expression that describes which values to display in the child drop-down list for each value of the parent property.
4. Add the value expression as an attribute of the child property.
5. Add both properties to a blueprint or build profile.

For a detailed example of defining relationships between properties, see [Property Relationship Example](#).

Creating a Value Expression

A value expression is an XML string that maps the values of the parent property and the child property in a property relationship.

To create a value expression:

1. Create a new XML document.

```
<?xml version="1.0" encoding="utf-8" standalone="yes"?>
```
2. Create the root element, **ArrayOfPropertyValue**.

```
<ArrayOfPropertyValue xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
</ArrayOfPropertyValue>
```
3. Insert a **PropertyValue** element.

```
<ArrayOfPropertyValue xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
  <PropertyValue>
  </PropertyValue>
</ArrayOfPropertyValue>
```
4. Insert the following child elements into the **PropertyValue** element:

Element	Description
FilterName	The name of the parent property
FilterValue	The value of the parent property that enables a specific value for the child property
Value	The value of the child property to enable when the parent property has a value of FilterValue

5. Repeat steps 3-4 for each property mapping you want to define for the parent property and the child property.
6. Format the XML as one continuous string with no line breaks.

The following is an example of a value expression (prior to formatting it as a single line):

```
<?xml version="1.0" encoding="utf-8" standalone="yes"?>
```

```

<ArrayOfPropertyVal ue xmlns:xsi="http://www.w3.org/2001/XMLSchema-i nstance">
  <PropertyVal ue>
    <Fi lterName>Custom. Property. Parent</Fi lterName>
    <Fi lterVal ue>Col ors</Fi lterVal ue>
    <Val ue>Bl ue</Val ue>
  </PropertyVal ue>
  <PropertyVal ue>
    <Fi lterName>Custom. Property. Parent</Fi lterName>
    <Fi lterVal ue>Col ors</Fi lterVal ue>
    <Val ue>Red</Val ue>
  </PropertyVal ue>
  <PropertyVal ue>
    <Fi lterName>Custom. Property. Parent</Fi lterName>
    <Fi lterVal ue>Shap es</Fi lterVal ue>
    <Val ue>Ci rcl e</Val ue>
  </PropertyVal ue>
  <PropertyVal ue>
    <Fi lterName>Custom. Property. Parent</Fi lterName>
    <Fi lterVal ue>Shap es</Fi lterVal ue>
    <Val ue>Square</Val ue>
  </PropertyVal ue>
</ArrayOfPropertyVal ue>

```

Displaying a Property on the Confirm Machine Request Page

After you have defined a property in the property dictionary, you must add the property to a blueprint or build profile to expose it to the user at request time.

Note: The following procedure assumes that you are adding the property directly to a blueprint or global blueprint, but the property can also be defined in a build profile that you reference in a blueprint.

To add a property from the property dictionary to a blueprint:

1. Create a new blueprint or edit an existing blueprint.
2. On the **Properties** tab, click **New Property**.
3. In the **Name** box, specify the name of the property definition. This must be an exact match (case- and white space–sensitive).
4. Optionally, specify a **Value**.
 - For most property types, the value you specify in the blueprint is a default value that can be overridden by the user at request time.
 - For a label property, the value is the text of the label that displays.
 - For a link property, the value is the target URL of the link. (The property display name serves as the link text.)
5. Select **Prompt User**.

Although the label and link properties are read only, you must still select **Prompt User** in order to for them to appear on the **Confirm Machine Request** page.

6. Click the green check mark to save the property.
7. Click **OK**.

About Control Layouts

Control layouts enable you to define a group of properties that go together and specify the order in which they are displayed on the Confirm Machine Request page.

You can specify exactly one control layout for a blueprint and it should include all the property definitions with custom controls that you want to use with that blueprint. You can use the same layout on multiple blueprints.

If there are additional custom properties that are specified as **Prompt User** on the same blueprint as a control layout, the properties in the layout appear first, followed by the individual properties. These individual properties display using the default behavior for custom properties (that is, with the property name as the display name, and a plain text field for entering the value) even if you have defined a property in the property dictionary with the same name.

The high-level process for creating a control layout is as follows:

1. Create the property definitions that you want to include in a layout.
2. Create a new control layout.
3. Add the properties to the layout and specify the order in which they should be displayed.
4. Add the control layout and the properties that it contains to a blueprint.

For a detailed example of creating and using control layouts, see [Control Layout Example](#).

Creating a Control Layout

A control layout acts as a named container to which you can add properties and specify the order in which their user controls should appear.

Before creating a control layout, you should first create all the property definitions that you want to include in the layout.

To create a control layout:

1. Click **Enterprise Administrator > Property Dictionary**.
2. Click **Add Property Layout**.
3. In the **Name** box, specify a name for the layout.
4. Optionally, specify a **Description**. This text describes the behavior to the property dictionary administrator and is not displayed to end users.
5. Click the green check mark to save the layout.
6. Under **Property Instances**, click **Edit**.
7. In the **Order** box, specify an alphanumeric value that represents the order in which a property should display. Order values do not have to be sequential.

Note: Order values are sorted alphabetically. For example, if you use numerical values from 1-12, they are sorted in the following order: 1, 10, 11, 12, 2, 3 ...

8. From the **Property Definition** drop-down, select one of the properties you have previously defined.

9. Click the green check mark to save the property as part of the layout.
10. Repeat steps 6-9 as needed to add additional properties to the layout.
11. Click **OK**.

Displaying a Control Layout on the Confirm Machine Request Page

To enable a control layout, you add it to a blueprint along with the properties that it contains.

You specify a control layout on a blueprint using a reserved custom property named **VirtualMachine.Request.Layout**.

If you use a control layout in a blueprint, any custom properties that you add to the blueprint that are not part of the control layout have the default behavior on the Confirm Machine Request page. For example, suppose you define a property named **CustomProperty.NotInLayout** with a display name and custom user control. If you add the property to a blueprint that has a control layout that does not include this property, the property appears on the Confirm Machine Request page with the label **CustomProperty.NotInLayout** and a basic text field for input rather than the display name and user control specified in the property dictionary.

Note: The following procedure assumes that you are adding the properties directly to a blueprint or global blueprint, but the properties can also be defined in a build profile that you reference in a blueprint.

To add a control layout to a blueprint:

1. Create a new blueprint or edit an existing blueprint.
2. On the **Properties** tab, click **New Property**.
3. In the **Name** box, type **VirtualMachine.Request.Layout**.
4. In the **Value** box, specify the name of the control layout.
5. Leave the **Prompt User** check box unselected.
6. Click the green check mark to save the property.
7. For each of the properties contained in the layout, add a custom property to the blueprint, ensuring that you select **Prompt User**.
8. Click **OK**.

Examples of Using the Property Dictionary

The following sections include detailed examples of using the property dictionary and control layouts.

Note: The following examples assume that you are adding the properties directly to a blueprint or global blueprint, but the properties can also be defined in a build profile that you reference in a blueprint.

Property Definition Example: Network Drop-Down List

Typically, when you request a machine on a reservation with more than one available network, the network is automatically assigned by vCAC during provisioning. The following example shows how you can add a drop-down list that enables users to select from a list of specific values for the network, giving the users more control over the resulting machine.

To create a blueprint with network selection:

1. Log on as an Enterprise Administrator.
2. Click **Property Dictionary**.
3. Create the property definition.
 - a. Click **Add Property Definition**.
 - b. In the **Name** box, specify **VirtualMachine.Network0.Name**. (**0** is the number zero.)
 - c. In the **Display Name** box, specify **Select Network**.
 - d. Optionally, specify a **Description**.
 - e. From the **Control Type** drop-down list, choose **DropDownList**.
 - f. Select the **Required** check box to make this property required at request time.
 - g. Click the green check mark to save the definition.
4. Specify the set of networks to display in the drop-down list by creating a property attribute.
 - a. Under **Property Attributes**, click **Edit**.
 - b. Click **Add Property Attribute**.
 - c. From the **Type** drop-down list, select **ValueList**.
 - d. In the **Name** box, specify **Values**.
 - e. In the **Value** box, enter a comma-separated list of networks a user can choose from, for example: **Development Network,QA Network 1,QA Network 2,Production Network,Failover Network**.

Note: You should use values that correspond to the networks that are available on your reservation.

- f. Click the green check mark to save the attribute.
 - g. Click **OK**.
5. Add the **VirtualMachine.Network0.Name** custom property to a blueprint (or global blueprint).
 - a. Create a new blueprint or edit an existing blueprint.
 - b. On the **Properties** tab, click **New Property**.
 - c. In the **Name** box, specify the name of the property definition you just created: **VirtualMachine.Network0.Name**. This must be an exact match (case- and white space-sensitive).
 - d. Leave the **Value** blank, or specify one of the possible values in the drop-down list to use as the default value.
 - e. Select **Prompt User**.
 - f. Click the green check mark to save the property.
 - g. Click **OK**.

When a user requests a machine using this blueprint, the "Select Network" drop-down list now appears on the Confirm Machine Request page, with the values specified in the property attributes. The user can select a network from this list for the machine that they are requesting.

Property Relationship Example

Building on the previous example, this example shows how to define a relationship so that the value of a new property (the parent) determines the values of the networks drop-down list (the child).

We introduce a property, **VirtualMachine.Network.Environment**, that is used to filter the available network options based on the environment for the machine. The user must first select the environment before selecting from the list of networks that are applicable to the selected environment.

Value selected for VirtualMachine.Network.Environment	Available values for VirtualMachine.Network0.Name
Development	Development Network
QA	QA Network 1 QA Network 2
Production	Production Network Failover Network

The **VirtualMachine.Network.Environment** has no effect in vCAC other than to act as a filter for the networks that a user can select on the Confirm Machine Request page.

To enable a user to select a network for a machine based on its environment:

1. Log on as an Enterprise Administrator.
2. Click **Property Dictionary**.
3. Create the parent property.
 - a. Click **Add Property Definition**.
 - b. In the **Name** box, specify **VirtualMachine.Network.Environment**.
 - c. In the **Display Name** box, specify **Environment**.
 - d. Optionally, specify a **Description**.
 - e. From the **Control Type** drop-down list, choose **DropDownList**.
 - f. Select the **Required** check box to make this property required at request time.
 - g. Click the green check mark to save the definition.
4. Define the values for the parent property.
 - a. Under **Property Attributes**, click **Edit**.
 - b. Click **Add Property Attribute**.
 - c. From the **Type** drop-down list, select **ValueList**.
 - d. In the **Name** box, specify **Values**.
 - e. In the **Value** box, specify **Development,QA,Production**.
 - f. Click the green check mark to save the attribute.
 - g. Click **OK**.
5. Create the child property if you have not already created it in the previous example.
 - a. Click **Add Property Definition**.
 - b. In the **Name** box, specify **VirtualMachine.Network0.Name**.

- c. In the **Display Name** box, specify **Select Network**.
 - d. Optionally, specify a **Description**.
 - e. From the **Control Type** drop-down list, choose **DropDownList**.
 - f. Click the green check mark to save the definition.
6. Define the relationship between the child and the parent properties.
 - a. Under **Property Attributes** for the **VirtualMachine.Network0.Name** property, click **Edit**.
 - b. Click **Add Property Attribute**.
 - c. From the **Type** drop-down list, select **Relationship**.
 - d. In the **Name** box, specify **Parent**.
 - e. In the **Value** box, specify **VirtualMachine.Network.Environment**.
 - f. Click the green check mark to save the attribute.
 - g. Click **OK**.
 7. Create the following gvalue expression that specifies the values of the child property based on the values of the parent property:

```
<?xml version="1.0" encoding="utf-8" standalone="yes"?>
<ArrayOfPropertyVal ue xmlns:xsi="http://www.w3.org/2001/XMLSchema-i nstance">
  <PropertyVal ue>
    <Fi lterName>Vi rtual Machi ne. Network. Envi ronment</Fi lterName>
    <Fi lterVal ue>Devel opment</Fi lterVal ue>
    <Val ue>Devel opment Network</Val ue>
  </PropertyVal ue>
  <PropertyVal ue>
    <Fi lterName>Vi rtual Machi ne. Network. Envi ronment</Fi lterName>
    <Fi lterVal ue>QA</Fi lterVal ue>
    <Val ue>QA Network 1</Val ue>
  </PropertyVal ue>
  <PropertyVal ue>
    <Fi lterName>Vi rtual Machi ne. Network. Envi ronment</Fi lterName>
    <Fi lterVal ue>QA</Fi lterVal ue>
    <Val ue>QA Network 2</Val ue>
  </PropertyVal ue>
  <PropertyVal ue>
    <Fi lterName>Vi rtual Machi ne. Network. Envi ronment</Fi lterName>
    <Fi lterVal ue>Producti on</Fi lterVal ue>
    <Val ue>Producti on Network</Val ue>
  </PropertyVal ue>
  <PropertyVal ue>
    <Fi lterName>Vi rtual Machi ne. Network. Envi ronment</Fi lterName>
    <Fi lterVal ue>Producti on</Fi lterVal ue>
    <Val ue>Fai l over Network</Val ue>
  </PropertyVal ue>
</ArrayOfPropertyVal ue>
```

8. Add the value expression to the child property.
 - a. In a text editor or similar tool, format the value expression so that it is a single string with no line breaks.
 - b. Under **Property Attributes** for the **VirtualMachine.Network0.Name** property, click **Edit**.
 - c. Click **Add Property Attribute**.
 - d. From the **Type** drop-down list, select **ValueExpression**.
 - e. In the **Name** box, type **expression**.
 - f. Copy the value expression from the text editor and paste it into the **Value** box.
 - g. Click the green check mark to save the attribute.
 - h. Click **OK**.
9. Add both properties to a blueprint.
 - a. Create a new blueprint or edit an existing blueprint.
 - b. Click the **Properties** tab.
 - c. Click **New Property**.
 - d. In the **Name** box, type **VirtualMachine.Network.Environment**. This must be an exact match (case- and white space-sensitive).
 - e. Leave the **Value** blank.
 - f. Select **Prompt User**.
 - g. Click the green check mark to save the property.
 - h. Repeat steps c-g for **VirtualMachine.Network0.Name**.
 - i. Click **OK**.

When a user requests a machine using this blueprint, the “Environment” and “Select Network” drop-down lists now appear on the Confirm Machine Request page and are initially empty. The user can select an environment, which then narrows the list of network they can select from for the machine that they are requesting.

Control Layout Example

In this example, we demonstrate how to create and use a control layout.

To use a control layout on a blueprint:

1. Click **Enterprise Administrator > Property Dictionary**.
2. Create the properties to use in the layout.
 - a. Click **Add Property Definition**.
 - b. In the **Name** box, type **ControlLayout.Network.Title**.
 - c. In the **Display Name** box, type **Title**.
 - d. From the **Control Type** list, select **Label**.
 - e. Click the green check mark to save the definition.

- f. If you have not created the **VirtualMachine.Network.Environment** and **VirtualMachine.Network0.Name** properties from a previous example, create two simple text box properties in their place:

Name	Display Name	Control Type
VirtualMachine.Network.Environment	Environment	TextBox
VirtualMachine.Network0.Name	Network	TextBox

3. Create a control layout.
 - a. Click **Enterprise Administrator > Property Dictionary**.
 - b. Click **Add Property Layout**.
 - c. In the **Name** box, type **NetworkLayout**.
 - d. Click the green check mark to save the layout.
4. Add the properties you just created to the layout.
 - a. Under **Property Instances**, click **Edit**.
 - b. In the **Order** box, type **1**.
 - c. From the **Property Definition** drop-down, select **ControlLayout.Network.Title**.
 - d. Click the green check mark to save the property as part of the layout.
 - e. Repeat steps b-d with the following values:

Order	Property Definition
2	VirtualMachine.Network.Environment
3	VirtualMachine.Network0.Name

- f. Click **OK**.
5. Add the layout to a blueprint.
 - a. Create a new blueprint or edit an existing blueprint.
 - b. On the **Properties** tab, click **New Property**.
 - c. In the **Name** box, type **VirtualMachine.Request.Layout**.
 - d. In the **Value** box, type **NetworkLayout**.
 - e. Leave the **Prompt User** check box unselected.
 - f. Click the green check mark to save the property.
6. Add the properties contained in the layout to the blueprint.
 - a. Click **New Property**.
 - b. Add the three properties as follows:

Name	Value	Prompt User
ControlLayout.Networks.Title	Select an environment and network:	Yes
VirtualMachine.Network.Environment	<i>[leave empty]</i>	Yes
VirtualMachine.Network0.Name	<i>[leave empty]</i>	Yes

- c. Click **OK**.

When a user requests a machine using this blueprint, the Confirm Machine Request page looks like the following:

The screenshot shows a dialog box titled 'Confirm Machine Request' with two tabs: 'Request Information' and 'Storage'. The 'Request Information' tab is active and contains the following fields:

- # Machines: 1
- # CPUs: 1
- Memory (MB): 512
- Storage (GB): 1
- Description: (empty text area)
- Select an environment and network: (label above dropdowns)
- * Environment: (dropdown menu)
- * Select Network: (dropdown menu)
- Reason for request: (empty text area)

At the bottom of the dialog are 'OK' and 'Cancel' buttons.

If there are any other custom properties that are specified as Prompt User on this blueprint, they would appear between **Select Network** and **Reason for request**.

Reservation Alerts

Reservation alerts, introduced in release 4.5, enables an enterprise administrator to alert users when thresholds on reservations are reached so the user can track capacity usage. Thresholds can be set as a percentage of the total compute resources available. When a threshold is exceeded, an email is sent to a designated person or to a set of people. The enterprise administrator can adjust the frequency at which alerts are sent as well as receive copies of all email alerts. Alerts can be created for the following:

- Physical resources of a reservation including memory, CPU, and storage
- Total number of machines allocated against the reservation

Note: If the reminder is blank, you will receive the alert only once. Otherwise, you will receive an email at the frequency you have established.

Reservations are checked daily by running a scheduled workflow. The workflow collects any alerts for all existing reservations. If there is at least one alert set up, the workflow collects data on

- current capacities (adding the capacities of each machines to get the vCPUs used)

- allocated capacities
- limits for each capacity

For each defined alert, the workflow compares the allocated amount against the prescribed limit. All capacities that are over the prescribed limit are summarized in an email that is sent (asynchronously) to the specified set of recipients.

To establish thresholds and send email alerts (virtual only):

1. Click **Enterprise Administrator > Reservations**.
2. Click the **Enable alerts** check box to turn on the thresholds alert option.
3. Use the sliders to set the thresholds for each resource, including:
 - Storage
 - Memory
 - CPU
 - Machine quota
4. Enter the email addresses of recipients of the alert emails.
Length of emails is limited to 256 characters.
5. (Optional) Click the check box Send email to Group Manager.
6. Enter the reminder frequency (in days) or leave the field blank to not send email reminders after sending the initial alert email.

To establish machine quota thresholds and send email alerts (cloud, physical):

1. Click **Enterprise Administrator > Reservations** and click the **Reservation Information** tab.
2. Click the Enable alerts check box to turn on the thresholds alert option.
3. Use the sliders to set the machine quota thresholds.
4. Enter the email addresses of recipients of the alert emails.
Length of emails is limited to 256 characters.
5. (Optional) Click the check box to Send email to Group Manager.
6. Enter the reminder frequency (in days) or leave the field blank to not send email reminders after sending the initial alert email.

Snapshot Manager

A snapshot is an image of a virtual machine at a certain point in time. It is a space-efficient, pointer-based copy of the original VM image. Snapshots can be an easy way to recover a system from damage, data loss and security threats.

Beginning in release 4.5, you can use Snapshot Manager to create a snapshot of the virtual machine state with or without the state of the memory. If you create snapshots in VirtualCenter, vCAC discovers them during inventory data collection. Snapshot Manager is supported for VMware vSphere only, excluding version 2.5.

To ensure that snapshots are used properly, you can specify:

- Who is allowed to create, apply and delete snapshots by enabling or disabling the feature for machine owners

- How many snapshots of a machine are allowed, which controls storage capacity
- When a snapshot becomes too old to apply, which prevents a machine owner from rolling back to obsolete or non-compliant system states

Enabling and Disabling Snapshot Manager

Snapshot Manager is enabled in vCAC by default on the Customization page in the Console Settings tab. The vCAC Administrator can disable Snapshot Manager by deselecting the **Enable vSphere vCenter Snapshot Management** checkbox. This removes all references to Snapshot Manager from the vCAC website, but does not destroy existing snapshots or change the database in any way.

The setting that allows users to take snapshots is on the Blueprint page in the Security tab. In the list of features that machine owners can use, **Snapshot Manager** is selected by default. If the Enterprise Administrator or Provisioning Group Manager deselects the checkbox, the machine owner does not see Snapshot Manager on the right-click menu for the machine.

Prerequisites

For the VirtualCenter permissions, see the *vCloud Automation Center Installation Guide*.

Working with Snapshots

Use Snapshot Manager to create, apply, delete and view snapshots.

Snapshot Manager displays a list of snapshots of the selected machine. For each snapshot, the following information displays:

- **Name** – Snapshot name. Duplicate names are not allowed.
- **Created** – Date and time when the snapshot was created
- **Is Memory Snapshot?** – Yes if the snapshot includes memory state, No if it does not. This column has data for only snapshots that were created from vCAC.
- **Description** – You can enter optional information.

As you work with snapshots, you can see a history of snapshot actions in the Recent Events log.

To create a snapshot:

Prerequisites:

- The maximum number of snapshots allowed on a machine cannot be exceeded. The default is 1.
- The datastore on the reservation on which the machine is provisioned must have sufficient space to store the snapshot. For vCAC, sufficient space is enough to store another copy of the virtual machine's hard drive.
- The machine must be in the on or off state.

1. Navigate to your machines:

- **Enterprise Administrator** – Choose **Enterprise Machines**.
- **Provisioning Group Manager** – Choose **Group Machines**.
- **Support** – Choose **Support Group > Group Machines** or **Support User > Machines for User**
- **Machine Owner** – Choose **My Machines**.

2. Point to a virtual machine and click **Snapshot Manager**.
3. Click **New Snapshot** in the upper-right corner.
4. (Optional) The **Name** field displays the machine name and present data and time, which you can change.
5. (Optional) Add a description.
6. (Optional) Click **Snapshot the machine's memory** to create a disk and memory snapshot.
7. Click **OK** to create the snapshot and add it to the list in Snapshot Manager.

To apply a snapshot:

1. Navigate to Snapshot Manager, as described in steps to create a snapshot.
2. Point to a snapshot and click **Apply**.
3. Confirm your request on the **Confirm Snapshot Action** dialog. While the snapshot is being applied, the **Apply** option is not available and you cannot create or delete other snapshots. When you refresh or reload the page after the snapshot is applied, all options are available again.

Caution: *When you apply a snapshot, all changes made or data added to the machine since the last snapshot was created are permanently lost.*

To delete a snapshot:

1. Navigate to Snapshot Manager, as described in steps to create a snapshot.
2. Point to a snapshot and click **Delete**.
3. Confirm your request on the **Confirm Snapshot Action** dialog.

Caution: *When you delete a snapshot, the machine state preserved in the snapshot is permanently lost. If you delete the current snapshot, indicated by **You are here**, the parent snapshot is applied.*

To customize a snapshot policy:

By default, vCAC allows one snapshot per machine and provides no age limit for snapshots. Both settings can be changed with custom properties:

Custom Property Name	Purpose	Default
Snapshot.Policy.Limit	To increase the number of snapshots allowed per machine Snapshots are shown in a hierarchical tree: <ul style="list-style-type: none"> • Depth – Maximum is 31, as limited by VMware. • Width – There is no limit. 	1
Snapshot.Policy.AgeLimit	To specify an age limit, in days, for a snapshot that can be applied. When a snapshot exceeds the age, the Apply option is no longer available for it.	None

To customize the limit or age of snapshots allowed:

1. Based on your privileges, navigate to a blueprint or virtual machine:
 - Enterprise Administrator – Open a new or existing global virtual blueprint.
 - Provisioning Group Manager – Open a new or existing virtual blueprint.
 - User who is both EA and PGM – In Self-Service, point to a virtual machine and click Edit.
2. Go to the **Properties** tab.
3. In the **Custom properties** list, select **New Property**.
4. In the **Name** field, type:
 - **Snapshot.Policy.Limit** to increase the number of snapshots
 - **Snapshot.Policy.AgeLimit** to specify the number of days a snapshot can exist before it becomes too old to apply
5. In the **Value** field:
 - For Snapshot.Policy.Limit, type a number that does not exceed the VMware limit of 31.
 - For Snapshot.Policy.AgeLimit, type the number of days a snapshot can exist before it becomes too old to apply.
6. Click  to save the property.
7. Click **OK**.

Note: Any administrative user, such as the Enterprise Administrator or Provisioning Group manager, can add these two snapshot policy properties or modify their values on Machine Edit page.

Linked Clones

A linked clone is a space-efficient copy of a virtual machine. It is based on a snapshot of a VM and uses a chain of delta disks to track differences from the parent machine. This feature was introduced in release 4.5.

By pointing to a snapshot, the linked clone uses less disk space than a full clone. This reduces storage costs. Linked clones can be provisioned quickly and are ideal to use when performance is not a high priority.

The following principles apply:

- Linked clones must have access to the parent disk. If the parent disk is deleted or moved to a different datastore, the linked clone is disabled.
- A linked clone of a machine with multiple disks has a delta file for each disk.
- If the machine or snapshot on which the linked clone is based is deleted, the linked clone still works because shared disks remain on the server. Only unshared disks get deleted with a machine.
- To ensure that linked clones are used properly, you can specify:
 - Who is allowed to provision a linked clone

- How many linked clones of one machine are allowed
- What pool of reservation resources to use

Working with Linked Clones

The Provisioning Group Manager and Enterprise Administrator can:

- Create a linked clone blueprint
- Customize the number of linked clones allowed

The Self-Service user can provision a virtual machine using a linked clone blueprint.

To create a linked clone blueprint:

Prerequisite: At least one virtual machine must be provisioned and be in the ON or OFF state.

1. Navigate to a blueprint:
 - a. Enterprise Administrator > Global Blueprints.
 - b. Provisioning Group Manager > Blueprints.
2. Click **New Blueprint > Virtual** or point to an existing virtual blueprint and click **Edit**.
3. Go to the **Build Information** tab.
4. In the **Action** field, select **Linked Clone** which is available for the vSphere (vCenter) platform only. This changes the **Provisioning workflow** to **CloneWorkflow**.
5. In the **Clone from** field, click .
6. In the wizard:
 - a. Select a **Machine to Clone** and click **Next**.
 - b. Select a **Snapshot** and click **Finish**.
7. Click **OK** to confirm.

By default, vCAC allows you to make 20 linked clones of one machine snapshot. You can change this setting with a custom property.

To set the maximum number of linked clones:

1. Navigate to a blueprint:
 - a. Enterprise Administrator > Global Blueprints.
 - b. Provisioning Group Manager > Blueprints.
2. Click **New Blueprint > Virtual** or point to an existing virtual linked clone blueprint and click **Edit**.
3. Go to the **Properties** tab.
4. In the **Custom properties** list, select **New Property**.
5. In the **Name** field, type MaximumProvisionedMachines.
6. In the **Value** field, type the maximum number of machines you want to provision on the linked clone blueprint.

7. Click  to save the property.
8. Click **OK**.

To provision a linked clone:

1. Select **Self-Service > Request Machine**.
2. Click a **virtual blueprint** that specifies Linked Clone, as described in the previous task.
3. If runtime properties are required, provide the information.
4. Click **OK**.

About Global Machine Limits for a Blueprint

Each blueprint has an option to limit the number of machines that any one user can provision based on that blueprint. vCAC also provides an option to limit the number of machines that can be provisioned across all users.

You can add the custom property `MaximumProvisionedMachines` to specify a global limit on machines that can be provisioned from that blueprint. This property applies to all machine types including multi-machine services and component machines of multi-machine services.

Most blueprint types are unlimited by default. Linked clone blueprints have a default maximum of 20 machines that can be provisioned from a given blueprint.

Custom Properties

This section describes new custom properties added to vCAC and provides information about how to add and use custom properties.

Custom Property Name	Description
Amazon.ElasticLoadBalancer.Names	Assigns machines provisioned on a blueprint to the elastic load balancers that match the values you specify.
Image.ISO.UserName	Specify a user name to access the boot ISO in the following form: <code>username@domain</code> .
Image.ISO.Password	Specify the password for the user specified in <code>Image.ISO.UserName</code> .
MaximumProvisionedMachines	To set the number of machines that can be provisioned across all users based on a given blueprint. Most blueprint types are unlimited by default. Linked clone blueprints have a default maximum of 20 machines.
Snapshot.Policy.Limit	Sets the limit of Snapshots (depth) that can be created for machines.

Custom Property Name	Description
Snapshot.Policy.AgeLimit	Sets the age limit (in days) for Snapshots that can be applied to machines.
VirtualMachine.DiskN.Storage	To specify the datastore on which to place machine disk N.
VirtualMachine.DiskN.StorageReservationPolicy	To specify a storage reservation policy to find storage for disk N.
VirtualMachine.DiskN.StorageReservationPolicyMode	To allocate disk N on the best available storage reservation policy
VirtualMachine.NetworkN.PortID	When using a dvPort group with a vSphere distributed switch, this property is used to specify the Port ID to use for the provisioned virtual machine. In this example, the property, VirtualMachine.NetworkN.PortID uses $N = 0, 1, 2, \dots$ to represent the numerical index of the network card. Replace N with the required number.

This version of vCAC clarifies the order in which custom properties take precedence and allows custom properties to be applied to and used on compute resources. See the *vCloud Automation Center Operating Guide* for more information about custom properties.

Performing Operations on Multiple Machines

Beginning in release 4.5, provisioning group managers, enterprise administrators, support users and machine owners can select more than one machine when performing an operation. Machine operations that can be performed on more than one machine include:

- Change Lease
- Send Reclamation Request
- Destroy
- Expire
- Install Tools
- Power Cycle
- Power Off
- Reboot
- Shut Down
- Suspend

Note: All workflow tasks except Reprovision can be performed on multiple machines. Reprovision requires user input for each machine so the operation cannot be performed on more than one machine at a time.

To select multiple machines:

1. In the machine list, click the check box to the left of a machine name.

Or

2. Click the Select All on Page link.

To deselect multiple machines:

- Click the Deselect All on Page link.

To perform operations on multiple machines:

1. Point to the selected machines to display the menu and select an operation.
 - If more than one machine is selected, only tasks that can be performed on all selected machines are enabled.
 - If a user is hovering over a machine that is on, the operations available for a machine that is on are displayed with some operations grayed out.
 - If machines that are turn off are selected in the batch, machines that were selected and are not turned on already, will be unchecked and grayed out at the confirmation screen.
 - Only machines that support the operation are enabled when you confirm the action. Other machines are displayed but disabled.
2. Confirm the operation.

If one of the machines in the list of selected machines is a hypervisor, a warning message is displayed with an additional check box to confirm the action. Machines that are hypervisors display an icon that distinguishes them as hypervisors.

When selecting Change Lease, if only one machine is selected, the existing lease status is displayed. If more than one machine is selected and available, the current expiration date is not displayed.

Machines that are not accessible are disabled and cannot be selected. For example, physical and cloud machines do not support all of the operations that are available for virtual machines. You cannot use the multiple machine selection across all machine types.

Note: Custom machine operations can be performed on more than one machine if the operation includes the confirmVMTasks page.

Gugent

The legacy guest agent for Linux has been deprecated as of release 4.5. The functionality is now in a consolidated Gugent agent. For information on the updates, refer to the *vCloud Automation Center Operating Guide*.

Chapter 3 What's New for Extensibility

The following sections describe updates to the VMware extensibility toolkits.

vCenter Orchestrator Integration

vCloud Automation Center 5.1 provides the ability to call vCenter Orchestrator workflows from a vCAC workflow using new activities available in vCloud Automation Center Designer.

For information about vCenter Orchestrator integration, see the *vCloud Automation Center Extensibility Guide*.

Activity Library

The activity library has been expanded in release 4.5 with the addition of a SendEmail activity and a RunProcess activity to launch external processes on the DEM host.

Productivity Tools

The vCAC Designer user interface has been updated in release 4.5 for greater ease-of-use including a new tool for managing skills and their associations with workflows and DEMs. For more information about the new tooling for skills, see the [Skill Management Overview](#) section.

The vCloud Automation Center Development Kit is now generally available, including the Workflow Generator plugin for Visual Studio and an expanded set of CloudUtil commands. The Workflow Generator plugin accelerates development and deployment of custom workflows for vCAC. The vCloud Automation Center Development Kit requires an additional license on top of the VMware Cloud Suite.

New utilities have been also added to the CloudUtil command-line tool to package models and their supporting constructs, making it easier to export and import custom components between VMware instances.

CloudUtil Command-Line Tool

The CloudUtil command-line tool has been updated in release 4.5 for new functionality and consistency in naming. Specific changes since version 4.1 SP1 are described below.

Global changes

CloudUtil commands are no longer case-sensitive in version 4.5.

All commands except for `Client-Generate` have a new optional argument, `--repository`, which enables a user to specify the root URI of the Model Manager. The default URI is specified in the application configuration file.

Agent (DEM) Commands

All Agent commands in 4.1 SP1 have been renamed as follows:

4.1 SP1	4.5	Description
Agent-Add-Skills	DEM-Add-Skills	Associate skills with a registered Distributed Execution Manager.
Agent-List	DEM-List	List all Distributed Execution Managers registered with the Model Manager.
Agent-Remove-Skills	DEM-Remove-Skills	Remove association between skills and a Distributed Execution Manager.

Client Commands

The Client-Generate command has new arguments as follows:

4.1 SP1	4.5	Description
-f -File <Value>	-f --filename <Value>	Path to a Microsoft Entity Framework Data Model file (EDMX).
-s -SecurityConfig <Value>	-c --securityConfig <Value>	Path to a model security configuration file (XML).
-o -OutputPath <Value>	-o --output <Value>	Output path for the generated C# code file.
<i>new in 4.5</i>	-s --server <Value>	SQL Server host name or IP address.
<i>new in 4.5</i>	-d --database <Value>	SQL catalog within a given SQL Server instance.
<i>new in 4.5</i>	-u --user <Value>	The SQL user for the database where the data described by the model resides.
<i>new in 4.5</i>	-p --password <Value>	The SQL password for the database where the data described by the model resides.

Events Commands

Events-Install

4.1 SP1	4.5	Description
-c -EventConfig <Value>	-c --eventConfig <Value>	Path to a Model Manager event configuration file (XML).

Events-Uninstall

4.1 SP1	4.5	Description
-n --WorkflowName <Value>	-n --name <Value>	Name of the workflow for which to uninstall events.

File Commands

File-Import

4.1 SP1	4.5	Description
-f --file <Value>	-f --filename <Value>	Path to a file to import into the Model Manager.

File-Update

4.1 SP1	4.5	Description
-f --file <Value>	-f --filename <Value>	Path to a file to update in the Model Manager.

Model Commands

Model-Config-Export

4.1 SP1	4.5	Description
-m --model <Value>	-n --name <Value>	Name of the model in the Model Manager for which to export security configuration.
-f --out-filename <Value>	-o --output <Value>	Path for file output.

Model-Config-Import

4.1 SP1	4.5	Description
-a --authorization-info <Value>	-c --securityConfig <Value>	Path to a model security configuration file (XML).

Model-Install

4.1 SP1	4.5	Description
-a --authorization-info <Value>	-c --securityConfig <Value>	Path to a model security configuration file (XML).
-c --catalog <Value>	-d --database <Value>	SQL catalog within a given SQL Server instance.

Operation Commands

Operation-Create

4.1 SP1	4.5	Description
-i --InputFile <Value>	-c --operationConfig <Value>	Path to an operation definition file (XML).

Operation-Delete

4.1 SP1	4.5	Description
-n --OperationName	-n --name <Value>	Name of the custom operation in the Model

<Value>		Manager.
-y --yes	--force	Force deletion of the operation.

Package Commands

Two new commands, Package-Export and Package-Import, have been added in 4.5 to facilitate migration of custom models between VMware instances. For details about working with packages, see the *vCloud Automation Center Extensibility Guide*.

Schedule Commands

All Schedules commands in 4.1 SP1 have been renamed as follows:

4.1 SP1	4.5	Description
Schedules-Install	Schedule-Install	Install a workflow schedule or set of schedules in the Model Manager based on a schedule configuration file.
Schedules-List	Schedule-List	List all workflow schedules installed in the Model Manager.
Schedules-Uninstall	Schedule-Uninstall	Uninstall a workflow schedule by ID or uninstall all workflow schedules for a given workflow.

Schedule-Install

4.1 SP1	4.5	Description
-c --ScheduleConfig <Value>	-c --scheduleConfig <Value>	Path to a workflow schedule configuration file (XML).

Schedule-Uninstall

4.1 SP1	4.5	Description
-s --Schedule <Value>	-s --schedule <Value>	The ID of the workflow schedule or the name of the workflow for which to uninstall schedules.

Workflow Commands

Workflow-Add-Skills

4.1 SP1	4.5	Description
-n --WorkflowName <Value>	-n --name <Value>	Name of a workflow in the Model Manager.

Workflow-Install

4.1 SP1	4.5	Description
-d --WorkflowDefinition <Value>	-f --filename <Value>	Path to a file (XAML) containing the workflow to install in the Model Manager.
-n --WorkflowName	-n --name <Value>	Name for the workflow in the Model Manager.

<Value>		
-p --Priority <Value>	<i>removed in 4.5</i>	
-e --Description <Value>	-d --description <Value>	Description of workflow.
<i>new in 4.5</i>	-o --runOneOnly	Limit this workflow to a single running instance.
<i>new in 4.5</i>	--force	Force installation of the workflow.

Workflow-Install-Three

This command is no longer supported in 4.5.

Workflow-Remove-Skills

4.1 SP1	4.5	Description
-n --WorkflowName <Value>	-n --name <Value>	Name of a workflow in the Model Manager.

Workflow-Rollback

4.1 SP1	4.5	Description
-n --WorkflowName <Value>	-n --name <Value>	Name of the workflow in the Model Manager.
-r --Version <Value>	-r --revision <Value>	Revision of the workflow to revert to.

Workflow-Uninstall

4.1 SP1	4.5	Description
-n --WorkflowName <Value>	-n --name <Value>	Name of the workflow to uninstall from the Model Manager.

Workflow-Update

4.1 SP1	4.5	Description
-d -- WorkflowDefinition <Value>	-f --filename <Value>	Path to a file (XAML) containing the updated workflow.
-n --WorkflowName <Value>	-n --name <Value>	Name of installed workflow to update.
-e --Description <Value>	-d --description <Value>	Description of workflow.

Note: The above is not intended to be a comprehensive list of CloudUtil commands or a complete description of their usage. For more information about CloudUtil, see the *vCloud Automation Center Extensibility Guide*.

Workflow Generator Plugin for Visual Studio

The Workflow Generator plugin for Visual Studio guides you in creating several types of vCAC workflows, including state change workflows. The workflow states available in the Workflow Generator wizard have changed in release 4.5 for consistency with the customizable workflows available in vCAC Designer.

The following workflow states have been added to the Workflow Generator plugin:

- Expired
- UnprovisionMachine
- Disposing

The following workflow state has been removed from the Workflow Generator plugin:

- DeactivateMachine

For a complete list of workflow states available in the vCAC Designer and Workflow Generator tools, see the *vCloud Automation Center Extensibility Guide*.

Documentation

The *vCloud Automation Center Extensibility Guide* has been expanded with information about working with custom models, including the new model packaging capabilities. There is also more detailed information about the supporting configuration for workflows, such as custom events and schedules.