Cisco UCS Director Fundamentals Guide, Release 6.5

First Published: 2017-07-21

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

CONTENTS

Preface

Audience

This guide is intended primarily for data center administrators who use Cisco UCS Director and who have responsibilities and expertise in one or more of the following:

- Server administration
- Storage administration
- Network administration
- Network security
- Virtualization and virtual machines

Conventions

Text Type	Indication
GUI elements	GUI elements such as tab titles, area names, and field labels appear in **this font**. Main titles such as window, dialog box, and wizard titles appear in **this font**.
Document titles	Document titles appear in *this font*.
TUI elements	In a Text-based User Interface, text the system displays appears in `this font`.

Text Type	Indication
System output	Terminal sessions and information that the system displays appear in `this font`.
CLI commands	CLI command keywords appear in **this font**. Variables in a CLI command appear in *this font*.
[ ]	Elements in square brackets are optional.
{x | y | z}	Required alternative keywords are grouped in braces and separated by vertical bars.
[x | y | z]	Optional alternative keywords are grouped in brackets and separated by vertical bars.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.
< >	Nonprinting characters such as passwords are in angle brackets.
[ ]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

Note Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the document.

Caution Means *reader be careful*. In this situation, you might perform an action that could result in equipment damage or loss of data.

Tip Means *the following information will help you solve a problem*. The tips information might not be troubleshooting or even an action, but could be useful information, similar to a Timesaver.

Timesaver Means *the described action saves time*. You can save time by performing the action described in the paragraph.

Warning

IMPORTANT SAFETY INSTRUCTIONS

This warning symbol means danger. You are in a situation that could cause bodily injury. Before you work on any equipment, be aware of the hazards involved with electrical circuitry and be familiar with standard practices for preventing accidents. Use the statement number provided at the end of each warning to locate its translation in the translated safety warnings that accompanied this device.

SAVE THESE INSTRUCTIONS

Related Documentation

Cisco UCS Director Documentation Roadmap

For a complete list of Cisco UCS Director documentation, see the *Cisco UCS Director Documentation Roadmap* available at the following URL: http://www.cisco.com/en/US/docs/unified_computing/ucs/ucs-director/doc-roadmap/b_UCSDirectorDocRoadmap.html.

Cisco UCS Documentation Roadmaps

For a complete list of all B-Series documentation, see the *Cisco UCS B-Series Servers Documentation Roadmap* available at the following URL: http://www.cisco.com/go/unifiedcomputing/b-series-doc.

For a complete list of all C-Series documentation, see the *Cisco UCS C-Series Servers Documentation Roadmap* available at the following URL: http://www.cisco.com/go/unifiedcomputing/c-series-doc.

Note The *Cisco UCS B-Series Servers Documentation Roadmap* includes links to documentation for Cisco UCS Manager and Cisco UCS Central. The *Cisco UCS C-Series Servers Documentation Roadmap* includes links to documentation for Cisco Integrated Management Controller.

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to ucs-director-docfeedback@cisco.com. We appreciate your feedback.

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, using the Cisco Bug Search Tool (BST), submitting a service request, and gathering additional information, see What's New in Cisco Product Documentation.

To receive new and revised Cisco technical content directly to your desktop, you can subscribe to the What's New in Cisco Product Documentation RSS feed. RSS feeds are a free service.

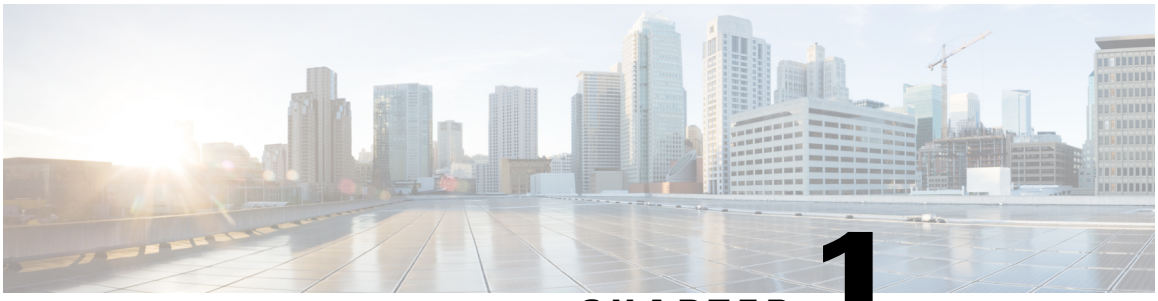

CHAPTER 1

New and Changed Information for this Release

This chapter contains the following sections:

New and Changed Information for this Release

No significant changes were made to this guide for the current release.

New and Changed Information for this Release

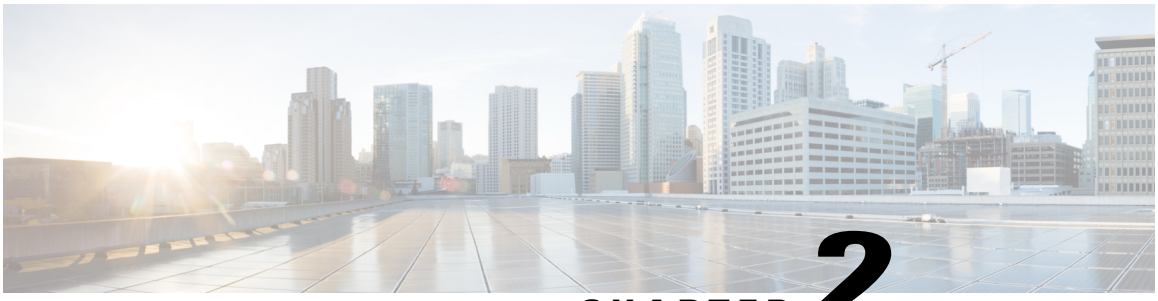

CHAPTER 2

Overview

Cisco UCS Director Overview

Cisco UCS Director focuses on delivering Infrastructure as a Service (IaaS) through a highly secure, end-to-end management, orchestration, and automation solution for a wide array of Cisco and non-Cisco data center infrastructure components. Cisco UCS Director can deliver IaaS for individual components and for the industry's leading converged infrastructure solutions based on the Cisco Unified Computing System (Cisco UCS) and Cisco Nexus platforms.

For a complete list of supported infrastructure components and solutions, see the Cisco UCS Director Compatibility Matrix. In addition to the supported components, you can develop modules that extend support to unsupported third party devices that are not listed in the Compatibility Matrix with the Cisco UCS Director Open Automation Framework.

With Cisco UCS Director, you can manage, automate, and orchestrate your physical and virtual compute, network, and storage resources. In addition, through the End User Portal, you can use those infrastructure components to deploy the desired virtual machines to support applications in cloud environments. To support this functionality, Cisco UCS Director enables you to do the following:

- Manage and support heterogeneous data centers that include compute, network, storage, and virtualization resources from multiple vendors.
- Provision physical and virtual compute, Layer 4-7 network services, and storage resources.
- Create and implement single and multi-tier application profiles.
- Define application containers that describe a set of tiers that include physical and/or virtual compute resources, their connectivity policy and communication policy. You can further define those application containers with network services, such as load balancing and firewalls, across these tiers.
- Establish secure multi-tenant environments, so that users, whether internal to your company or external, can work only within the secure constraints of their own resource pool. With the policies and user roles that you establish, your users can view, manage, and use only the infrastructure components appropriate for their roles.

- Automate the IT processes necessary to accomplish infrastructure provisioning and decommissioning using a role- and policy-based model that limits administrator and user capabilities.
- Implement a process-oriented approach to infrastructure orchestration that automates the processes you define using built-in workflows or customized workflows created from Cisco UCS Director task library or from tasks you create yourself.
- Implement metering, chargeback, and showback features so your organization can be properly compensated for the IT services you provide.

System Overview

Cisco UCS Director connects all the elements of your data center infrastructure, including the users, and the physical and virtual infrastructure. You can not only provision, configure, monitor, and automate your data center management, you can also use the Cisco UCS Director REST API or Open Automation Framework to extend the out-of-the-box functionality.

Figure 1: Cisco UCS Director System Overview

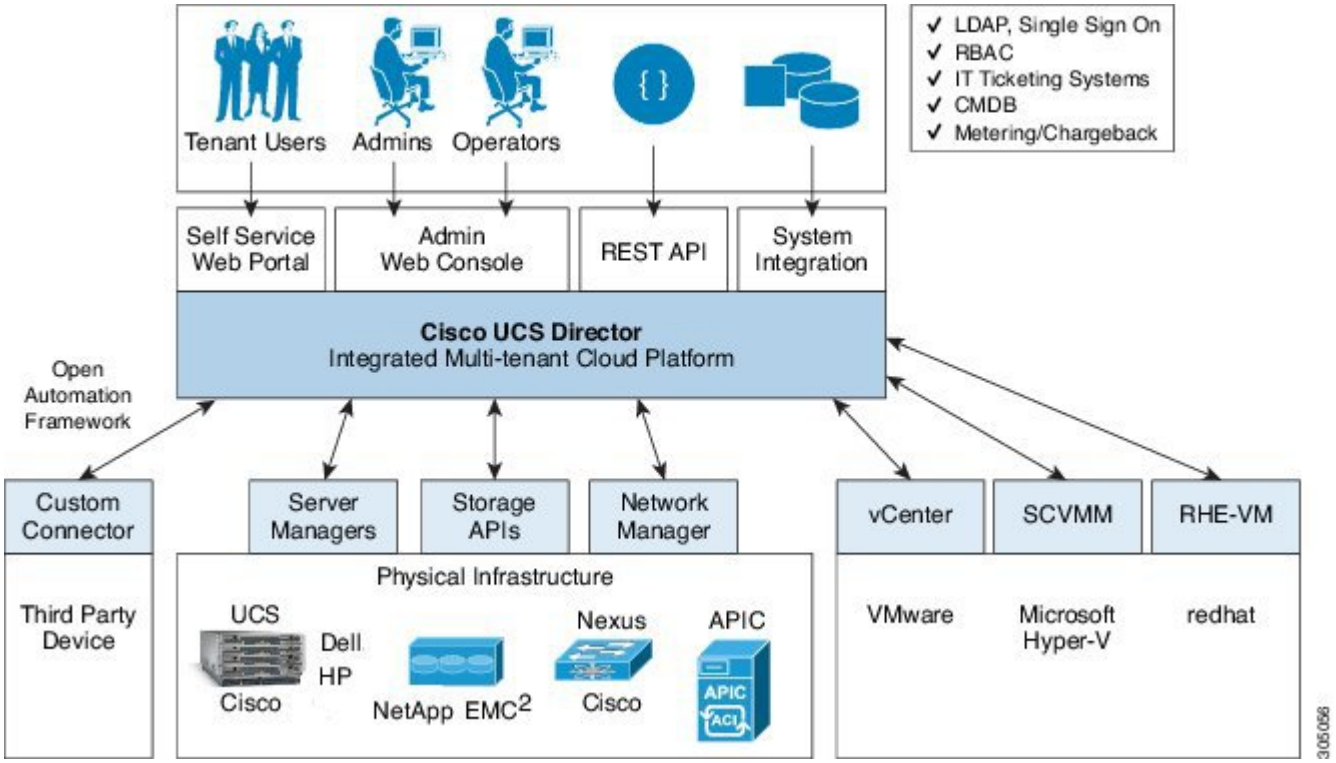

Infrastructure Configuration and Management

Cisco UCS Director extends the unification of compute, network, virtualization, and storage layers and provides you with comprehensive visibility into your data center infrastructure. By communicating with the appropriate domain managers or domain controllers, Cisco UCS Director can act as a single appliance to manage all your

infrastructure. This central management capability enables operations teams to configure, administer, manage, and monitor supported Cisco and non-Cisco physical and virtual compute, network, and storage components.

Cisco UCS Director provides out-of-the-box integration with virtual and physical components, including the following:

- Hypervisors, such as VMware vSphere, Microsoft Hyper-V, and RedHat KVM
- Compute servers and devices, such as Cisco UCS, HP, and Dell servers
- Network devices, such as Cisco Nexus and Brocade
- Storage components, such as NetApp, EMC, and IBM Storwize
- Hyperconverged storage solutions, such as VMware Virtual SAN (VSAN)

For a complete list of supported infrastructure components and solutions, see the Cisco UCS Director Compatibility Matrix.

Orchestration and Automation

Cisco UCS Director provides model-based orchestration through workflows. These workflows can include complex logic, can be imported into or exported from Cisco UCS Director, and can configured to resume from the point of last failure. You can also include advanced orchestration features, such as rollback of workflows, that enable you to automate the provisioning and de-provisioning of resources, which provide much-needed agility in a cloud environment. Unlike script based execution systems, if you include rollback in a workflow, Cisco UCS Director intelligently deprovisions resources in reverse order of execution. This functionality is possible because Cisco UCS Director is model-aware and state-aware.

Cisco UCS Director enables you to build workflows that provide automation services, and to publish those workflows and extend their services, on demand, to your users through Cisco UCS Director and the End User Portal. You develop your workflows in the Workflow Designer that is built in to Cisco UCS Director. The Workflow Designer is a drag and drop orchestration editor that includes a large library of out-of-the-box workflow tasks and workflows.

Depending upon your business needs, you can use or modify the out-of-the-box workflows and workflow tasks or you can develop your own custom workflows or workflow tasks. Custom workflow tasks can use CloupiaScript, a Javascript-like programming language, REST API or PowerShell cmdlets. In the workflows you can combine your custom tasks with out-of-the-box generic tasks that perform common tasks, such as executing scripts through SSH and PowerShell.

If required, you can embed approvals inside a workflow to ensure that resources are not provisioned until they have been approved.

Once built and validated, these workflows perform the same way every time, no matter who runs the workflows or where they run the workflows. An experienced data center administrator can run them from the Workflow Designer. You can also publish workflows to the End User Portal to implement Infrastructure as a Service (IaaS) and allow your users and customers to view and run the workflows that are applicable to their needs on a self-service, on demand basis.

For more information about orchestration and automation in Cisco UCS Director, see the Cisco UCS Director Orchestration Guide and Orchestration and Automation, on page 21.

Infrastructure as a Service

Cisco UCS Director delivers Infrastructure as a Service (IaaS) for both virtual and physical infrastructure. With Cisco UCS Director, you can create an application container template that defines the infrastructure required for a specific application or how a customer or business unit is expected to use that application.

Cisco UCS Director helps you and your IT team define the rules for your business's infrastructure services. You first onboard tenants and then define the boundaries of the physical and virtual infrastructure that they can use, or you can allow your onboarded tenants to define the infrastructure boundaries. You can then create policies, orchestration workflows, and application container templates in Cisco UCS Director that define the requirements for a specific type of application that can be used by a tenant, such as a web server, database server, or generic virtual machine (VM).

You can then publish these templates as a catalog in the End User Portal. Your users can go to the End User Portal, select the catalog that meets their needs, and make a service request for that particular application or VM. Their service request triggers the appropriate orchestration workflow to allocate the required infrastructure and provision the application or VM. If you specified that this type of service request requires approvals, Cisco UCS Director sends emails to the specified approver(s). Once the service request is approved, Cisco UCS Director assigns the infrastructure to those users, creating a virtual machine if necessary, and doing the base configuration such as provisioning the operating system.

You can also configure an orchestration workflow to ask questions before allowing a user to choose a catalog item. For example, you could configure the workflow to ask the user what type of application they plan to run and automatically select a catalog for them based on the answers to those questions. The end user in this case does not have to worry about whether to request a physical server or a VM, what kind of storage they require, or which operating system to install. Everything is predefined and prepackaged in the catalog.

For example, you can create policies, orchestration workflows and an application container template for an SAP application that uses a minimum level of infrastructure, requires approvals from a director in the company, and has a chargeback to the department. When an end user makes a service request in the End User Portal for that catalog item, Cisco UCS Director does the following:

1. Sends an email to the director, who is the required approver.
2. When the approval is received, creates a VM in the appropriate pod with 4 CPUs, 10GB of memory, and 1TB of storage.
3. Installs an operating system (OS) on the VM.
4. Notifies the end user that the VM is available for them to use.
5. Sets up the chargeback account for the cost of the VM.

With the available APIs from Cisco UCS Director, you can also script custom workflows to pre-install the SAP application in the VM after the OS is installed.

Extensibility of Cisco UCS Director

In addition to the built-in functionality, you can extend Cisco UCS Director to provide customized functionality and better meet the needs of your data center. The Cisco UCS Director extensibility model supports several extension and customization techniques at different levels of complexity.

For more information about how to extend Cisco UCS Director, see the Cisco UCS Director Orchestration Guide and the Cisco UCS Director Programming Guides.

Custom Workflows with Predefined Tasks

At the most basic level, you can extend the orchestration and automation functionality of Cisco UCS Director with custom workflows that you create with the tasks provided with Cisco UCS Director. You can use the workflows included with Cisco UCS Director as a base or you can develop a brand-new workflow.

This customization requires that you understand the operation that you want to automate with the workflow. You must then break down that operation into steps, or smaller operations, that are represented by a workflow task. You then place the tasks into the sequence required to perform the more complex operation. That sequence is the workflow.

Generic Tasks for Custom Workflows

If Cisco UCS Director does not provide a predefined task for a step that requires SSH or a PowerShell script, you can modify a generic task with a custom script to perform that step. For example, you can use a generic task to connect with and execute commands on a Linux server or an unsupported device.

Cisco UCS Director includes the following generic tasks:

SSH Command Task

You can use the SSH Command task to connect to and have Cisco UCS Director execute CLI commands on a remote system as part of a workflow. This task is useful in a Linux environment. It requires that the Cisco UCS Director administrator has CLI access, including password and credentials, to the remote system. It also requires an investment in creating the shell scripts required for the commands that you want Cisco UCS Director to execute.

For example, you can use an SSH Command task to push firewall rules for iptables or patch Linux applications with yum.

Execute PowerShell Command Task

You can use the Execute PowerShell Command task to launch PowerShell scripts from Cisco UCS Director. This task is useful in a Microsoft environment. It requires an investment in creating the PowerShell scripts required for the commands you want Cisco UCS Director to execute.

For example, you can use the Execute PowerShell Command task to add a host to additional domains or DNS systems.

Custom Tasks for Custom Workflows

Custom tasks enable you to create functionality for a workflow that is not available in the predefined tasks and workflows that are included with Cisco UCS Director. A custom task includes inputs and outputs and works like any other task in the workflow. You can call other tasks from within a custom task.

You create custom tasks with CloupiaScript, a version of JavaScript with Cisco UCS Director Java libraries that enable orchestration operations. CloupiaScript supports all JavaScript syntax. The Cisco UCS Director Java libraries allow you to access Cisco UCS Director components from a custom task. However, because CloupiaScript runs only on the server, client-side objects are not supported

Custom tasks can perform a wide variety of functions. For example, you can create a custom task to get static IP pool policy attributes, get mail settings, invoke a service request, or move an ESXi host to a different cluster.

Northbound APIs and Integrations

Northbound integrations with the Cisco UCS Director APIs enable you to perform the following tasks:

- Invoke operations and workflows
- Access reports and data for the following:
  - Physical infrastructure
  - Virtual devices
  - Network and storage devices
- User, groups, policies and other administrative functions

You can use the Cisco UCS Director REST API and PowerShell API to implement northbound integrations.

Cisco UCS Director REST API

Cisco UCS Director offers a REST API that enables applications to consume or manipulate the data stored in Cisco UCS Director. Cisco UCS Director REST API (hereafter, simply "the REST API") is a lightweight framework that requires little overhead for an application to use.

Applications use HTTP and HTTPS requests from the REST API to perform Create/Read/Update/Delete (CRUD) operations on Cisco UCS Director resources.

The REST API supports the following protocols and formats:

- JSON (JavaScript Object Notation)
- XML
- Java

When to Use Cisco UCS Director REST API

Cisco UCS Director REST API is a language-independent interface that can be used by any program or script capable of making HTTP or HTTPS requests. Use the REST API when you want to invoke operations on Cisco UCS Director from a separate program or process.

Applications can use the REST API to do the following:

- Retrieve Cisco UCS Director reports on physical and virtual devices, networks, appliances, groups and users, policies, resource accounting, funds, and other monitored entities within your Cisco UCS domains.
- Invoke Cisco UCS Director Orchestrator workflow and task operations.
- Invoke additional operations specific to Cisco UCS Director.

Cisco UCS Director PowerShell API

Cisco UCS Director offers JSON-based REST APIs that enable you to submit workflow requests, examine workflow inputs and output schemas, and fetch reports. You can integrate Cisco UCS Director APIs with the Cisco UCS Director PowerShell Console for improved automation of datacenter management.

Cisco UCS Director PowerShell Console provides cmdlet wrappers for the JSON-based APIs. Each cmdlet performs a single operation. The cmdlets are executed in a Microsoft Windows server. Depending on the data returned by the JSON-based APIs, the cmdlets automatically interpret the data and convert them to Windows PowerShell objects. You can chain multiple cmdlets together. To view a list of available cmdlets, see the Cmdlet List in the Cisco UCS Director PowerShell API Getting Started Guide. For more information about REST APIs, see the Cisco UCS Director REST API Getting Started Guide.

When to Use Cisco UCS Director PowerShell API

Use the Cisco UCS Director PowerShell API to execute the JSON API's of Cisco UCS Director that are provided as cmdlets.

Cisco UCS Director Open Automation

Cisco UCS Director provides the Cisco UCS Director Open Automation module (hereafter, simply "Open Automation") to enable you to enhance the functionality of the Cisco UCS Director appliance.

Using Open Automation, you can add a module to Cisco UCS Director. The module adds a new capability to Cisco UCS Director, such as the ability to control a new device type or to generate a new type of report.

For more information about existing functionality in Cisco UCS Director, see the Cisco UCS Director Administration Guide, the Cisco UCS Director Application Container Guide, the Cisco UCS Director Orchestration Guide, and the numerous other guides that document functionality available using the Cisco UCS Director application.

For system requirements and information about how to set up a development environment, install libraries, and begin using the Open Automation SDK, see the Cisco UCS Director Open Automation Getting Started Guide.

For annotated examples of using the Open Automation SDK, see the Cisco UCS Director Open Automation Cookbook.

For a reference to the Cisco UCS Director Open Automation API, see the Cisco UCS Director API Javadoc.

When to Use Cisco UCS Director Open Automation

Open Automation is a Java SDK and framework for extending the functionality of Cisco UCS Director. Use Open Automation to enhance Cisco UCS Director in the following ways:

- Develop your own or third-party components with Cisco UCS Director.
- Design a custom menu for displaying your device or component.
- Inventory your devices.
- Provde the ability to test connections between your device and Cisco UCS Director.
- Develop tasks that can be used in workflows.

- Expose your tasks in the form of a REST API.
- Develop and schedule repeatable tasks.
- Develop your own Cisco UCS Director reports and report actions.
- Develop new Cisco UCS Director Cloudsense reports.
- Track changes made to the system through your module.
- Customize your dashboard display by providing your own dashboard (stack) builder.
- Develop your own items that can be displayed in your dropdown boxes.
- Provide support for new account types.

Cisco UCS Director Components

Cisco UCS Director includes several components that are designed to assist you in the tasks involved in the management and automation of your infrastructure and cloud.

Cisco UCS Director Bare Metal Agent

Cisco UCS Director Bare Metal Agent (BMA) automates the process of using a Preboot Execution Environment (PXE) to install operating systems on bare metal servers or virtual machines. Bare Metal Agent provides the following services that are required for a functional PXE install environment:

- Dynamic Host Control Protocol (DHCP)
- Hypertext Transfer Protocol (HTTP)
- Trivial File Transfer Protocol (TFTP)

When this environment is operational and Bare Metal Agent and Cisco UCS Director are correctly configured, you can build PXE installation tasks into any Cisco UCS Director infrastructure workflow.

You can access Bare Metal Agent through Secure Shell (SSH). You can also perform services on Bare Metal Agent, such as DHCP configuration and starting and stopping services, through Cisco UCS Director. A single Cisco UCS Director node can support multiple Bare Metal Agent applications.

For more information about Bare Metal Agent, see the Installation and Configuration Guides.

Cisco UCS Director Shell

The Cisco UCS Director Shell is a text-based menu that you access through a secure shell (SSH) application and Cisco UCS Director administrator credentials. With the Shell, you can execute commands to perform various system administration tasks, including:

- Patch updates
- Database backup and restore
- Certificate imports
- Services management

For more information about the Shell and its commands, see the Cisco UCS Director Shell Guide.

End User Portal

The End User Portal is a self-service portal that includes a catalog of services provided by your administrator. After you request one of the services available to you, the End User Portal completes the service request workflow configured by your administrator. This workflow may include approval of your self-service provisioning request, assignment of the necessary compute, storage and network resources, and configuration of security and performance settings. After your service is provisioned, you can track the status of your services through the summary dashlets and summary reports on your landing page, and through the service request report available within the End User Portal.

As an end user, based on your administration setup, you can perform one or more of the following operations:

- Provision virtual machines (VMs), application specific infrastructure, and bare metal servers
- Review and manage your service requests
- Upload and deploy OVFs and other images
- Monitor and create reports for your provisioned virtual and physical resources
- Approve service requests to provision infrastructure

Additional functionality may be available to you if your administrator provides you with the necessary permissions.

For more information about the End User Portal, see the Cisco UCS Director End User Portal Guide.

Cisco UCS Director Express for Big Data

Cisco UCS Director Express for Big Data is a single-touch solution within Cisco UCS Director that automates deployment of Big Data infrastructure. Cisco UCS Director Express for Big Data provides a single management pane across physical infrastructure and across Hadoop and Splunk Enterprise software. It supports key Hadoop distributions, including Cloudera, MapR, and Hortonworks.

Cisco UCS Director Express for Big Data delivers end-to-end automation of Hadoop cluster deployment, allowing you to spin up and expand clusters on-demand. The physical infrastructure configuration is handled automatically, with minimal user input. The configuration includes compute, internal storage, network, and installation of operating system, Java packages, and Hadoop, along with the provisioning of Hadoop services. This is achieved through Cisco UCS service profiles wherein both the physical infrastructure and Hadoop configuration are incorporated into a Hadoop cluster deployment profile.

Cisco UCS Director Express for Big Data also delivers end-to-end automation of Splunk cluster deployment, with minimal user input. This is achieved through Cisco UCS service profiles wherein both the physical infrastructure and Splunk configuration are incorporated into a Splunk cluster deployment profile.

For more information about Cisco UCS Director Express for Big Data, see the Cisco UCS Director Express for Big Data Documentation Roadmap.

Cisco UCS Director SDK

The Cisco UCS Director SDK is a collection of technologies that enable you to extend the capabilities of Cisco UCS Director, access Cisco UCS Director data, and invoke Cisco UCS Director's automation and orchestration operations from any application. The Cisco UCS Director SDK includes the REST APIs and Open Automation. Scripting technologies include the Cisco UCS Director PowerShell API, custom tasks bundled in Cisco UCS Director script modules, and the ability to write your own custom tasks using CloupiaScript, a server-side JavaScript implementation.

With Cisco UCS Director SDK technologies, you can:

- Access Cisco UCS Director programmatically—Use the Cisco UCS Director REST API to invoke workflows and obtain reports.
- Customize Cisco UCS Director—Create custom workflow tasks. Customize Cisco UCS Director by deploying your own jar files and script libraries in script modules. Use custom tasks from script bundles.
- Extend Cisco UCS Director—Use Cisco UCS Director Open Automation to build connectors that support additional devices and systems. Use the Cisco UCS Director PowerShell API to connect to Microsoft System Center Virtual Machine Manager (SCVMM) and other PowerShell enabled devices.

For more information about Cisco UCS Director SDK, see the Cisco UCS Director API Integration and Customization Guide.

Guided Setup Wizards in Cisco UCS Director

Cisco UCS Director includes a set of wizards that guide you through configuring important features. The following are the available guided setup wizards:

- Device Discovery—This wizard enables you to discover devices and assign them to a pod.
- Initial System Configuration—This wizard helps you complete initial tasks to set up Cisco UCS Director, such as uploading licenses, and setting up SMTP, NTP, and DNS servers.
- vDC Creation—This wizard enables you to configure the policies required to provision VMs in private clouds.
- FlexPod Configuration—This wizard helps you set up a FlexPod account.
- Vblock Pod Configuration—This wizard enables you to discover and assign accounts to Vblock pods.
- VSPEX Pod Configuration—This wizard enables you to discover and assign accounts to VSPEX pods.
- Virtual SAN Pod Configuration—This wizard enables you to set up a Virtual SAN Pod and add devices.

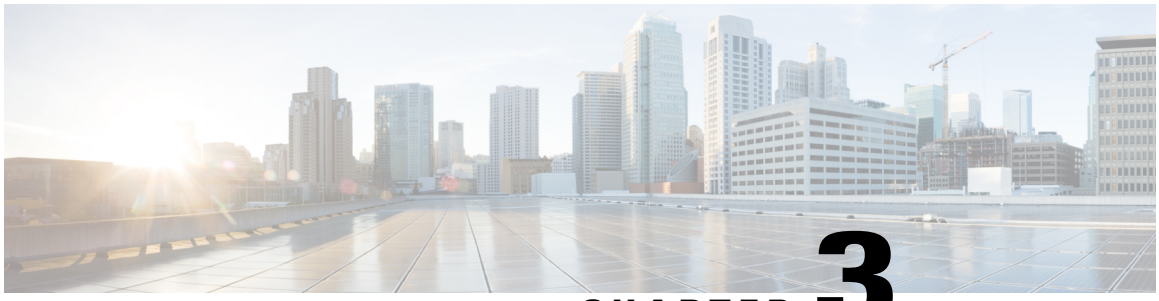

CHAPTER 3

Administration

Administration of Core Platform Features

Cisco UCS Director administration includes core platform features, such as configuration for users and groups, and the ways in which you can add your own brand to Cisco UCS Director and End User Portal. These core platform features enable you to securely control how your users and customers access Cisco UCS Director and what they see after they log in.

For more information how to administer the core platform features, see the Cisco UCS Director Administration Guide.

Users

You can create locally authenticated users for Cisco UCS Director, or you can integrate Cisco UCS Director with your LDAP server and synchronize the LDAP users and groups with Cisco UCS Director. You can do the following with your users, whether they are local to Cisco UCS Director or they access Cisco UCS Director through an LDAP integration:

- Assign user roles and permissions to limit their access and the operations that they can perform
- Enable the Managed Service Provider feature
- Group internal and external customers to control which applications, resources, and tenants they can view
- Set resource limits on user groups or customer organizations
- Create a password policy to define the minimum password requirements for all users

User Groups

User groups provide a way for you to define logical associations between your users. For example, you could create user groups for the Development, Sales, and Marketing users within your business. Through the user group, you can then associate the users in that group with the following:

- One or more virtual data centers (vDCs) that define the infrastructure resources which the users can access when making service requests
- Resource limits that further control the maximum resources that users in a group can use.
- A cost center for chargeback on the infrastructure resources

Managed Service Provider and Customer Organizations

Managed service provider and customer organizations are available only when you enable the Service Provider feature in Cisco UCS Director.

A Managed Service Provider (MSP) organization is a categorization of a customer group in Cisco UCS Director. It is a high-level categorization of a customer organization, and can also be considered as a parent organization. Within this MSP organization, you can group multiple sub-categories or child organizations. For example, a company name such as Cisco Systems is an MSP organization. Within this MSP organization, you can create multiple customer groups such as HR, Finance and Operations. These groups are considered to be customer organizations within the MSP organization

Like user groups, you can associate users in customer organizations with the following:

- One or more virtual data centers (vDCs) that define the infrastructure resources which the users can access when making service requests
- Resource limits that further control the maximum resources that users in a group can use
- A cost center for chargeback on the infrastructure resources

Group Budget Policy

Resources are accounted for by using the Chargeback feature. For resource usage by a group or customer organization, you associate the entity with a budget policy.

You can configure a group or customer organization with a budget watch, and configure a group or customer organization to stay within or exceed the provisioned budget.

Resource Limits

You can configure resource limits for a group, a customer organization, or a tenant, to manage resource utilization. You can specify limits for the following:

- Virtual resources
- Operating system resources
- Physical resources

Note Configuration of operating system resource and physical resource limits are not supported for public clouds.

Guidelines for Resource Limits with Service Provider Enabled

If you have enabled the Service Provider feature in Cisco UCS Director, keep in mind the following considerations while configuring resource limits:

- The limit set for the parent organization determines the limits that you can set for customer groups and containers within the parent organization.
- If you have not added resource limits to a specific customer group but have specified limits for containers within that group, you must consider the total resource limits before setting a limit for another customer group within the parent organization.
- The total number of resource limits configured for all customer groups and containers within those groups cannot exceed the resource limit set for the parent organization.
- If you do not add resource limits to a specific customer group but do specify resource limits for containers within that group, you must consider the total of all container resource limits before you set a limit for the parent organization.

For example, the parent organization, Tenant 1, includes three customer groups: Group A, Group B, and Group C. If you configure a resource limit of ten for Tenant 1, the cumulative resource limits specified for all customer groups within Tenant 1 must not exceed ten. The cumulative resource limits include resource limits applied to customer groups and containers.

Group A includes containers C1 and C2, and has a resource limit of 4 assigned to the customer group. Group B includes containers C3 and C4, and each of these containers has a resource limit of two. This configuration means that two is the maximum available resource limit you can configure for Group C and all its containers (the resource limit for Tenant 1, minus the total resource limits for Group A, Group, B, and containers C1, C2, C3, and C4).

Guidelines for Resource Limits with Service Provider Disabled

If you have disabled the Service Provider feature, there is only one parent organization. Therefore, if you set a resource limit for the parent organization, the total of the limits specified for all customer groups must not exceed the parent resource limit.

For example, the parent organization includes two customer groups: Group A and Group B. If you set a limit of ten for the parent organization and five for Group A. The resource limit that you set for Group B, must not exceed five (the resource limit for the parent organization minus the resource limit for Group A).

User Roles

Cisco UCS Director supports the following user roles:

- All Policy Admin
- Billing Admin
- Computing Admin

- Group Admin—An end user with the privilege of adding users. This user can use the End User Portal.
- IS Admin
- MSP Admin
- Network Admin
- Operator
- Service End User—This user can only view and use the End User Portal.
- Storage Admin
- System Admin

These user roles are system-defined and available in Cisco UCS Director by default. You can determine if a role is available in the system by default, if the **Default Role** column in the **Users** page is marked with **Yes**.

- Create a custom user role in the system, and create user accounts with this role or assign the role to existing users.

  When you create a new user role, you can specify if the role is that of an administrator or an end user.
- Modify existing user roles, including default roles, to change menu settings and read/write permissions for users associated with that role.

  The procedure to modify menu settings and permissions for a role is the same as the procedure to add a user role.

Multi-Role Access Profiles

A user can be assigned to more than one role, which is reflected in the system as a user access profile. For example, a user might log into Cisco UCS Director as a group administrator and as an all-policy administrator, if both types of access are appropriate.

Access profiles also define the resources that can be viewed by a user. With Cisco UCS Director Release 5.4, support for multiple profiles for a single user was introduced. So when you install version 5.4, and if a user account is associated with multiple groups, the system creates multiple profiles for the user account. But if you upgrade the system from a prior version to version 5.4, and if the **LDAPSyncTask** system task is not yet run, then, by default, only one profile is listed for a user account in the system.

When LDAP users are integrated with Cisco UCS Director, if a user belongs to more than one group, then the system creates a profile for each group. But by default, the domain users profile is added for LDAP users.

Note The **Manage Profiles** feature enables you to add, log into, edit, or delete a user access profile.

Default User Permissions

Each admin user has a set of permissions to access Cisco UCS Director. The types of user permissions are as follows:

- Read—An admin user with Read permission has the ability to only read a file.

- Write—An admin user with Write permission has the ability to read, write, and modify a file. This permission includes the ability to delete or rename files.
- Read/Write—An admin user with Read/Write permission has the ability to read and write a file.

LDAP Integration

You can use LDAP integration to synchronize the LDAP server's groups and users with Cisco UCS Director. LDAP authentication enables synchronized users to authenticate with the LDAP server. You can synchronize LDAP users and groups automatically or manually. While adding an LDAP account, you can specify a frequency at which the LDAP account is synchronized automatically with Cisco UCS Director. Optionally, you can manually trigger the LDAP synchronization by using the **LDAPSyncTask** system task.

After you configure an LDAP account and after the synchronization process is run, either manually or automatically, the recently added LDAP information is displayed in Cisco UCS Director. This information is displayed in a tree view that depicts the hierarchical structure of all organizational units (OUs), groups, and users that have been added to the system. You can view this information by choosing **Administration** > **LDAP Integration**. You can select and expand an OU in the left pane to view all the groups within it. If you select a group in this left pane, you can view a list of users that are associated with that group. If an OU has several sub OUs within it, then you can click the **Organization** tab in the right pane to view detailed information. In addition, the **Groups** and **Users** tabs in the right pane display groups and users respectively that are synchronized within the selected OU.

In addition to running a system task, Cisco UCS Director also provides an additional option for you to synchronize the LDAP directory with the system:

- **Cleanup LDAP Users** system task—This system task determines if the synchronized users in the system are deleted from the LDAP directories or not. If there are user records that have been deleted from the LDAP directories, then after this system task is run, these user accounts are marked as disabled in the system. As an administrator, you can unassign resources of these disabled user accounts. By default, this task is in the enabled state. After the second system restart, this system task is changed to the disabled state. This applies to both, a standalone and multi-node setup.

  In a multi-node setup, this system task runs only on the primary node, even if there is a service node configured.

Important Users that do not belong to a group or to a domain user's group display in LDAP as **Users with No Group**. These users are added under the domain user's group in Cisco UCS Director.

You can add LDAP users with the same name in different LDAP server accounts. The domain name is appended to the login user name to differentiate multiple user records. For example: abc@vxedomain.com. This rule is applicable to user groups as well.

Appending the domain name to the user name to login to the system is only applicable to LDAP users. It does not apply to local users. All local users can login to the system with the user names.

When a single LDAP account is added, and a user logs in by specifying only the user name, Cisco UCS Director first determines if the user is a local user or is an LDAP user. If the user is identified as a local user and as an external LDAP user, then at the login stage, if the user name matches the local user name, then the local user is authenticated into Cisco UCS Director. Alternatively, if the user name matches that of the external user, then the LDAP user is authenticated into Cisco UCS Director.

LDAP Integration Rules and Limitations

Group Synchronization Rules

- If a chosen LDAP group already exists in Cisco UCS Director and the source is type **Local**, the group is ignored during synchronization.
- If a chosen LDAP group already exists in Cisco UCS Director and the group source is type **External**, the group's description and email attributes are updated in Cisco UCS Director.
- While adding an LDAP server, you can now specify user filters and group filters. When you specify a group filter, all users that belong to the specified group are added to the system. In addition, the following actions are also performed:
  - If the specified group includes sub-groups, then the group, the sub-groups and the users in those sub-groups are added to the system (only applicable when you manually synchronize the LDAP directory).
  - If the user is part of multiple groups, and the other groups do not match the group specified as the group filter, then those additional groups are not added to the system.
- A user can be part of multiple user groups. However, the group that is mentioned first in the list of groups that the user is part of is set as the default primary group for the user. If the user is not part of any group, then the default primary group is set as **Domain Users**.

Note You can view information on all the groups that a user is part of only after the **LDAPSyncTask** system task is run.

- When an LDAP group is synchronized, all users that are in the group are first added to the system. Also, if users in the specified LDAP group are associated with other groups that are in the same OU or in a different OU, then those groups are also retrieved and added to the system.
- The LDAP synchronization process will retrieve the specified LDAP groups for the system, along with nested groups, if any.
- Prior to this release, a user was part of only one group. After an upgrade to the current release, and only after the **LDAPSyncTask** system task is run, the **Manage Profiles** dialog box displays the other groups that the user is part of. This is applicable only when the other groups match the group filters that you specified while configuring the LDAP server.

User Synchronization Rules

- LDAP users that have special characters in their names are now added to Cisco UCS Director.
- While adding an LDAP server, you can now specify user filters and group filters. When you specify a user filter, all the users that match the filter you specified, and the groups that they belong to, are retrieved for the system.
- An LDAP user can have multiple group memberships. When the LDAP user is synchronized with the system, this multiple group membership information is retained. Cisco UCS Director provides you with an option to view this information for every user. For more information on viewing group membership

information, see the Cisco UCS Director Administration Guide. In addition, multiple access profiles are also automatically created for the user.

Note You can view this information only when the groups match the filter you specified while configuring the LDAP server, and when the groups have been assimilated into the system.

- Cisco UCS Director now displays the User Principal Name (UPN) for each user that is added into the system. This is applicable for users that have been added into the system in prior releases. Users can log in to the system using their login name or their user principal name. Logging in using the user principal name along with the profile name is not supported.
- If a chosen LDAP user already exists in Cisco UCS Director and the source type is **External**, then that user is ignored at synchronization. The user's name, description, email, and other attributes are not updated again.
- If a user account is created in two different LDAP directories, then the user details of the LDAP directory that was synchronized first are displayed. The user details from the other LDAP directory are not displayed.
- After multiple LDAP directories are synchronized, the LDAP external users must log in to Cisco UCS Director by specifying the complete domain name along with their user name. For example: vxedomain.com\username. However, this rule does not apply if there is only one LDAP server directory added to Cisco UCS Director.

Note After an LDAP synchronization process, verify that the user is assigned to the correct group.

Branding Cisco UCS Director

Cisco UCS Director supports branding and customizing Cisco UCS Director and End User Portal at the following levels:

- Global level—This system level branding can be modified by the global administrator.
- MSP organization level or tenant level—The branding at this level can be modified by the administrator or the MSP administrator.
- Customer organization level—The branding at this level can be modified by an MSP administrator or a global administrator. Customer organizations are usually grouped with an MSP organization.

Branding for Customer Organizations

With the introduction of branding support at the MSP organization level, certain rules apply to what branding changes users may view. The settings that are applied depend on the following:

- User role—Is the user an end user, a group administrator, or an MSP administrator?

- User's customer organization and the branding set for it.
- MSP Organization branding settings.

The following table elaborates the branding behavior in Cisco UCS Director.

Table 1: Branding Behavior in Cisco Cisco UCS Director

Branding set at MSP Organization Level	Branding set at Customer Organization Level	MSP Administrator	Group Administrator	End User
Yes	Yes	Branding details set at the MSP organization level are displayed.	Branding details set at the customer organization level is displayed.	Branding details set at the customer organization level to which this user belongs to is displayed.
No	Yes	Global branding details are displayed.	Branding details set at the customer organization level is displayed.	Branding details set at the customer organization level to which this user belongs to is displayed.
Yes	No	Branding details set at the MSP organization level are displayed.	Branding details set at the MSP organization level to which this customer organization belongs to is displayed.	Branding details set at the MSP organization level to which the customer organization belongs to is displayed.
No	No	Global branding details are displayed.	Global branding details are displayed.	Global branding details are displayed.

Login Page Branding

A login page can be configured to display a logo that is associated with a domain name. When the end user logs in from that domain, the user sees the custom logo on the login page. The optimal image size for a logo is 890 pixels wide and 470 pixels high, with 255 pixels allowed for white space. Cisco recommends that you keep the image size small to enable faster downloads.

Note The group or customer organization login page must first be configured (enabled) for branding.

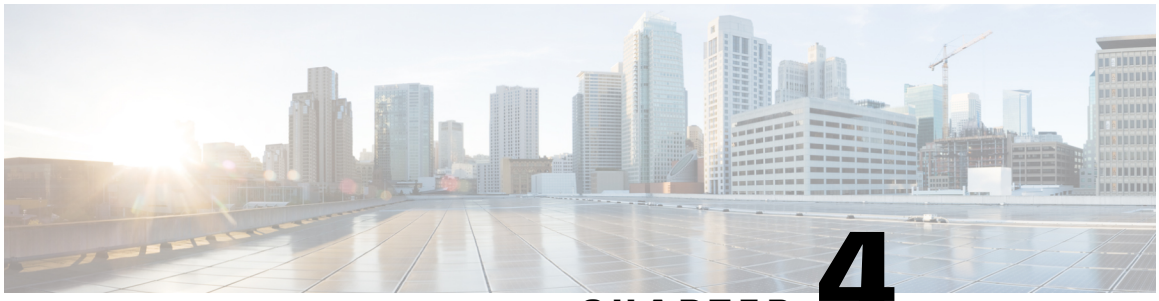

CHAPTER 4

Orchestration and Automation

Orchestration and Automation Overview

Cisco UCS Director enables you to automate a wide array of tasks and use cases across a wide variety of supported Cisco and non-Cisco hardware and software data center components, including physical infrastructure automation at the compute, network, and storage layers. A few examples of the use cases that you can automate include, but are not limited to:

- VM provisioning and lifecycle management
- Network resource configuration and lifecycle management
- Storage resource configuration and lifecycle management
- Tenant onboarding and infrastructure configuration
- Application infrastructure provisioning
- Self-service catalogs and VM provisioning
- Bare metal server provisioning, including installation of an operating system

For each of the processes that you decide to automate with orchestration workflows, you can choose to implement the processes in any of the following ways:

- Use the out-of-the-box workflows provided with Cisco UCS Director
- Modify the out-of-the-box workflows with one or more of the tasks provided with Cisco UCS Director

- Create your own custom tasks and use them to customize the out-of-the-box workflows
- Create your own custom workflows with custom tasks and the out-of-the-box tasks

For more information about automation and orchestration in Cisco UCS Director, see the Cisco UCS Director Orchestration Guide.

Cisco UCS Director Orchestrator

The Cisco UCS Director Orchestrator (also called the Cisco UCS Director Orchestration Engine, or just Orchestration) enables IT administrators to automate cloud deployment and provisioning and to standardize IT services.

At the scale of cloud computing, manually performing actions such as creating VMs and provisioning networks is prohibitively time-consuming. In Cisco UCS Director, these actions, or *tasks*, are executable elements that can be run from the GUI. Cisco UCS Director Orchestrator further automates these complex tasks by organizing them into *workflows*.

Orchestrator works by executing a series of scripted actions called *tasks*. Each task performs one action. By connecting tasks so that the input of one task is the output of a previous task, you build up a *workflow* to automate administrative processes such as creating VMs, provisioning bare metal servers, setting up storage, compute, and network resources, and many others.

Tasks

A task is the atomic unit of work in Cisco UCS Director Orchestrator. A task is a single action or operation with inputs and outputs (except in rare cases where the task operation does not require inputs or outputs). A task cannot be decomposed into smaller operations (exception: a compound task is made up of atomic tasks. But for now, think of a task as an indivisible unit of orchestration work).

Cisco UCS Director has a task library containing hundreds of predefined tasks that cover most of the actions an administrator must perform using Orchestration. In cases where there is no suitable predefined task, you can create custom tasks; see the Cisco UCS Director Custom Task Getting Started Guide.

Some examples of predefined tasks are:

- SSH Command task - Execute a command in a secure shell (SSH) session.
- Collect Inventory Task - Gathers information about available devices.
- Create New User - Add a user to the system.
- New VM Provision - Create a new VM (hypervisor-specific).

Workflows

A workflow is a series of tasks arranged to automate a complex operation. The simplest possible workflow contains a single task, but workflows can contain any number of tasks.

Workflows are the heart of the Cisco UCS Director Orchestrator; they enable you to automate processes of any level of complexity on your physical and virtual infrastructure.

You build workflows using the **Workflow Designer**, a drag-and-drop interface. In the **Workflow Designer**, you arrange tasks in sequence and define inputs and outputs to those tasks. Outputs from earlier tasks are available to use as inputs to any subsequent task.

Looping and conditional branching can be implemented using special flow-of-control tasks.

Service Requests

Service Requests are closely related to workflows. You create service requests by running workflows; a service request is generated every time you execute a workflow in Cisco UCS Director. A service request is a process under the control of Cisco UCS Director.

You can schedule a workflow for later execution, and Cisco UCS Director stores details of completed service requests. Thus a service request can have one of several states depending on its execution status: it can be scheduled, running, blocked (for example, awaiting an approval), completed, or failed (a service request can fail when one of its component tasks fails to execute properly).

Approvals

An approval is a "gate" task that requires the intervention of a Cisco UCS Director user to allow a workflow to run to completion. This user is typically an administrator who has go-or-no-go authority over the workflow process.

You can create *custom approvals* that allow users to enter input values when they approve a workflow. For example, you can create a custom approval to enable an IT administrator to approve provisioning of a VM, then specify the memory size of the VM before the workflow creates the VM.

Rollback

Workflows can be "rolled back" to a state identical or similar to the state before the workflow was executed. You can do this, for example, to remove virtual components that were created in error.

The term "rollback" often implies that a process is transactional, especially in systems using relational technology. However, it is important know that workflows are *not* transactional in a relational database sense. Instead, rollbacks work like this:

- Each task consists of two scripts. One script performs the work that the task was designed for. The other is a rollback script designed to "undo" the task. For example, if a task creates a VM, the rollback script deletes the VM.
- When you run a workflow, the scripts of the tasks in that workflow are executed in the order indicated by the workflow.
- When you roll back a workflow, the task rollback scripts are run in the reverse of the order they are run during normal workflow execution.
- A request to roll back a workflow creates a new service request, unconnected to the original service request.
- State information can be saved and used to roll back a task. For example, if you run a task to resize a VM, the pretask size can be saved to enable a rollback. This state persistence is a feature of the API that is used to write tasks; see the *Cisco UCS Director Custom Task* documentation for details.

- Tasks are atomic; workflows are not. You can roll back a workflow starting with any task in the workflow, but you cannot partially roll back a task.
- It is possible for a workflow to partially succeed; that is, for some but not all of its tasks to execute. Similarly, it is possible for a rollback to completely or partially fail; that is, that some or all of the rollback scripts could fail to execute.
- It is possible for a task to be written with a defective rollback script, or without a rollback script altogether. (Omitting the rollback script is not recommended.)

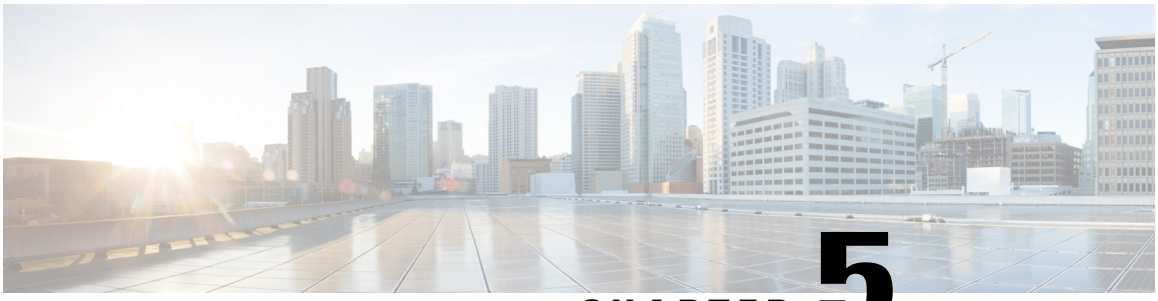

CHAPTER 5

Infrastructure Configuration and Management

Infrastructure Configuration and Management Overview

You can use Cisco UCS Director to configure, administer, manage, and monitor your infrastructure components. For example, out of the box, Cisco UCS Director can provision a VLAN on both a Cisco switch and the Cisco UCS servers with which the switch communicates, and it can also provision a LUN on a Netapp or EMC storage array and provision the vHBA on the Cisco UCS server to communicate with that LUN.

Some additional configuration and management tasks you can perform with Cisco UCS Director including the following:

- Create, clone, and deploy service profiles and templates for all Cisco UCS servers and compute applications.
- Monitor organizational usage, trends, and capacity across a converged infrastructure on a continuous basis. For example, you can view heat maps that show virtual machine (VM) utilization across all your data centers.
- Deploy and add capacity to converged infrastructures in a consistent, repeatable manner.
- Manage, monitor, and report on data center components, such as Cisco UCS domains or Cisco Nexus network devices.
- Configure, manage, and monitor all components of a converged infrastructure stack, such as FlexPod, Vblock, or VSPEX.

For more information about how to configure and manage the supported infrastructure components through Cisco UCS Director, see the Cisco UCS Director Management Guides for your physical and virtual compute, network, and storage resources.

Infrastructure Organization

Before you add any supported infrastructure to Cisco UCS Director, we recommend that you carefully consider how you want to organize your physical and virtual compute, network, and storage devices and resources. You can organize your infrastructure components into combinations of the following logical groupings:

- Accounts that represent the physical and virtual infrastructure components
- Pods that include one or more accounts
- Sites that include one or more pods

This hierarchy can reflect the physical organization and locations of your data center. However, you can also use the logical groupings in Cisco UCS Director to create an organization that maps to how the resources of your data center are used, the types of tenants that you support, or your customers.

Sites

A site is a logical grouping of one or more pods that represents a division of the resources in your data center. Sites are optional in Cisco UCS Director. However, if you choose to implement them, you can create sites that represent one or more of the following:

- A location within your data center. For example, if your data center is spread across multiple physical locations, you could create a site for each location, such as the U.S., India, and China.
- An internal customer who uses one or more tenants in a multi-tenant data center. For example, you could create a site for each of Finance, Sales, and Support.
- An external customer whose resources you want to segregate and keep separate from other customers.

You can view details of the sites on the **Converged** tab in Cisco UCS Director. For more information about how to create a site, see the Cisco UCS Director Administration Guide.

Pods

A pod is a logical grouping of physical and virtual components, including one or more physical or virtual accounts, such as a Cisco UCS Manager account for computing, a network account, or a cloud account. Each pod is a module of network, compute, storage, and application components that work together to deliver services for your data center and users. The pod is a repeatable pattern, and its components maximize the modularity, scalability, and manageability of data centers.

When you create a pod, consider what you want it to represent. For example, you can create a pod to represent the following:

- A single converged infrastructure stack, such as FlexPod, Vblock, or VSPEX
- A grouping of resources assigned to a specific customer or tenant
- The resources within a specific range of IP addresses

For systems that include Cisco UCS Central, we recommend that you create a pod for each Cisco UCS domain or domain group.

If needed, you can group pods into sites. However, a pod can only belong to one site.

You can view details of the pods and their components on the **Converged** tab in Cisco UCS Director. For information about how to create pods, see the Cisco UCS Director Administration Guide.

Accounts

Each account in Cisco UCS Director represents a physical or virtual infrastructure component. These accounts contain the information about the component that Cisco UCS Director requires to discover and manage the component, such as IP address, login name and password, port, and protocol.

For more information about the specific type of account you need to create, see the Cisco UCS Director Management Guides for your physical and virtual compute, network, and storage resources.

Note The following is a representative sample of the types of accounts that you can create in Cisco UCS Director. It is not an exhaustive list.

Virtual Accounts

A virtual account represents a supported hypervisor or cloud service. For example, you can create a virtual account for one or more of the following:

- VMware vSphere
- Microsoft Hyper-V
- RedHat KVM
- Rackspace Cloud
- Amazon Web Services EC2

Physical Accounts

A physical account represents a supported compute or storage device or domain. Physical accounts are identified by the element manager of the component, such as the following:

- UCSM account for Cisco UCS servers managed by Cisco UCS Manager
- HP OA account or HP iLO account for HP servers
- NetApp OnCommand account or NetApp ONTAP account for NetApp storage arrays
- EMC VNX account, EMC VNXe account, or EMC VMAX accounts for EMC storage arrays

Multi-Domain Manager Accounts

A multi-domain manager is a physical account that represents a software application that can manage more than one domain. Multi-domain manager accounts are identified by the software name, such as the following:

- Cisco UCS Central
- Cisco Prime Data Center Network Manager (DCNM)
- Cisco Prime Network Services Controller (PNSC)

- Cisco Application Policy Infrastructure Controller (APIC)
- EMC RecoverPoint
- EMC VPLEX

Rack Accounts

A rack account represents a supported rack server that is not managed by a physical account or a multi-domain manager account, such as a Cisco C-Series server or a Cisco E-Series server. There is only one type of rack account.

Managed Network Elements

A managed network element is a physical account that represents a supported network device, switch, firewall, or load balancer. Managed network elements are identified by the operating system or name of the network device, such as the following:

- Cisco Nexus OS
- Cisco IOS
- Cisco ASA
- F5 Load Balancer
- Brocade Fabric OS
- Brocade Network OS

Device Discovery

With Cisco UCS Director, you can perform device discovery manually or use the Device Discovery guided setup wizard. Some discovery steps are the same whether you do them manually or use the guided setup wizard. However, if you choose to use the wizard, you must use credential policies.

The protocol that Cisco UCS Director uses to communicate with the device and retrieve the inventory during discovery depends upon the type of device. For example, Cisco UCS Director uses XML to communicate with Cisco UCS Manager and CLI to communicate with Cisco Nexus devices.

During discovery and the retrieval of inventory, Cisco UCS Director creates at least one system task that polls the device at regular intervals and retrieves the inventory. The polling time for these system tasks is different for each device type. However, you can customize that interval to meet your data center's needs. For some devices, Cisco UCS Director creates additional system tasks to retrieve other information, such as faults and statistics.

Manual Device Discovery

The manual process discovers one infrastructure component at a time. You use the same basic procedure for all devices, whether the device is physical or virtual, or whether it's a compute, network, or storage device.

1. (Optional) You create a credential policy for the device.
2. You add a Cisco UCS Director account for the device to a pod, including the IP address, user name, password, and other relevant information.

3 Cisco UCS Director tests the connection to the device, using the credentials and other information that you provided in the account to authenticate.

4 If the connection test and authentication are successful, Cisco UCS Director does the following:

- Adds the device to its inventory and to the selected pod.
- Discovers the basic configuration and infrastructure elements associated with that device. For example, with a Cisco UCS Manager account, Cisco UCS Director discovers all infrastructure elements related to that account, including the chassis, servers, fabric interconnects, service profiles, and pools.
- Creates system tasks with default values for the device to poll for inventory and, if appropriate, other information at regular intervals.

This discovery process and inventory collection cycle takes several minutes to complete.

You need to repeat this manual process for all devices that you want Cisco UCS Director to manage.

Device Discovery through Guided Setup

With the Device Discovery guided setup wizard, you can perform discovery for all devices that you want to add to the same pod. You use the same basic procedure for all pods to which you want to add devices. You must create credential policies if you want to use the guided setup wizard.

1 You enter the following information in the **Policy** screen of the Device Discovery guided setup wizard:

- IP addresses for the devices that you want to discover (can be a range or a comma-separated list)
- Credential policies for the devices that you can choose or create

2 You choose a pod and start discovery on the **Discover and Assign** screen.

3 Cisco UCS Director tests the connection to all devices that can be accessed through the information added on the **Policy** screen.

4 If the connection test is successful, Cisco UCS Director does the following:

- Adds all the devices in the IP address list and range to its inventory.
- Discovers the basic configuration and infrastructure elements associated with all the devices. For example, with a Cisco UCS Manager account, Cisco UCS Director discovers all infrastructure elements related to that account, including the chassis, servers, fabric interconnects, service profiles, and pools.
- Creates system tasks with default values for the devices to poll for inventory and, if appropriate, other information at regular intervals.

This discovery process and inventory collection cycle takes several minutes to complete.

5 You choose the pod and the discovered devices that you want to add to it.

6 Cisco UCS Director adds all of the selected devices to the pod.

Credential Policies

Credential policies contain the login, port, and protocol information to log into a supported physical or virtual compute, network, or storage device or management application. Each credential policy represents a specific type of account. With a credential policy, when the password for a device login ID changes, you only need to change the password in the credential policies for those devices. You do not need to update the accounts for every individual device that Cisco UCS Director manages.

When to Use a Credential Policy

You can use the same credential policy for one or more of the same type of resources that share the same login user ID, password, port, and protocol. For example, you can use the same credential policy for one or more Cisco Nexus OS devices, NetApp ONTAP resources, Cisco UCS Manager domains, or VMware hypervisors.

When Not to Use a Credential Policy

You cannot use the same credential policy for different types of accounts. For example, you cannot use the same credential policy for a Cisco Nexus OS device and a Cisco UCS Manager domain, even if they share the same login user ID, password, port, and protocol.

Infrastructure Reporting

Cisco UCS Director provides a wide variety of reporting information out of the box. In addition, you can create custom report templates to build your own reports for the parameters that are relevant to your needs.

Additional information about reporting is available in the Cisco UCS Director Administration Guide and in the appropriate Cisco UCS Director Management Guides.

Reports

Cisco UCS Director can help you monitor virtual infrastructure and system resources. It displays a wide variety of reports that provide insight into how the system is performing

Following are the types of reports:

- Tabular reports for system information, including overview, host nodes, new VMs, and deleted VMs.
- Bar and pie graph comparisons, including VMs active versus inactive, and CPU provisioned versus capacity.
- Trend graphs about system resources, including CPU trends, memory trends, and VM additions and deletions.
- Other reports include Top 5 reports at the group, VDC, host node, and VM levels. The Top 5 reports focus on groups with the highest number of VMs, groups with the greatest CPU usage, VDCs with the highest number of VMs, and host nodes with the greatest CPU usage.
- Map reports, displaying the system resource information in the form of heat maps or color-coded maps.

Additional trend reports are available for certain accounts (for example: KVM accounts). Trend reports display data over a selected time frame.

CloudSense Analytics

CloudSense Analytics in Cisco UCS Director provide visibility into the infrastructure resources utilization, critical performance metrics across the IT infrastructure stack, and capacity in real time. CloudSense significantly improves capacity trending, forecasting, reporting, and planning of virtual and cloud infrastructures.

You can generate the following reports with CloudSense:

- Billing Report for a Customer
- EMC Storage Inventory Report
- NetApp Storage Inventory Report
- NetApp Storage Savings Per Group
- NetApp Storage Savings Report
- Network Impact Assessment Report
- Organizational Usage of Virtual Computing Infrastructure
- PNSC Account Summary Report
- Physical Infrastructure Inventory Report for a Group
- Storage Dedupe Status Report
- Storage Inventory Report For A Group
- Thin Provisioned Space Report
- UCS Data Center Inventory Report
- VM Activity Report by Group
- VMware Host Performance Summary
- Virtual Infrastructure and Assets Report

Note This is a complete list of reports available in the system. However, the number of reports available in the system for a user depends on the user role. By default, the **CloudSense** option is not visible to MSP administrators. The system administrator needs to enable this option for MSP administrators. Once this is done, then when an MSP administrator logs in, only reports relevant to customer organizations are displayed.

Report Builder for Custom Report Templates

Using the **Report Builder** option in Cisco UCS Director, you can create custom report templates to run reports on specific parameters. You can specify the context, the type of report to run, and the duration of the data samples for the report. You can also create multiple templates.

After you have created a report template, you can use it to generate a report in either PDF or HTML formats. You can view custom reports in Cisco UCS Director or you can email reports, either to yourself and to other users in your organization. You can review and archive these reports outside Cisco UCS Director.

In addition to creating a template, you can edit, clone, and delete custom report templates.

Note You cannot generate daily and hourly trend cost reports using the report builder. You can generate trend reports only for weekly and monthly duration. While generating trend reports for a month, the data is caculated from the first day of the month till the current date. For example, if you are generating a trend report on 5th March, this report includes data from March 1st, to March 5th.

Dashboard

In Cisco UCS Director, you can enable the **Dashboard** option in the user interface. On the **Dashboard** screen, you can add important, or frequently accessed report widgets. If you have enabled the **Dashboard** option, then this is the first window that you see when you log in to the user interface. After enabling the **Dashboard**, you can create additional dashboards, and delete them when you no longer need them.

Summary

The **Summary** screen allows you to manage system inventory. It gives you access to a wide array of tabular, graphical, and map reports, and also helps in managing inventory lifecycle actions.

Each report is displayed as a widget. Reports can be hidden through customization.

Inventory Management

You can monitor the system inventory using the **Dashboard**. The **Dashboard** displays the entire system level infrastructure information for administrative management.

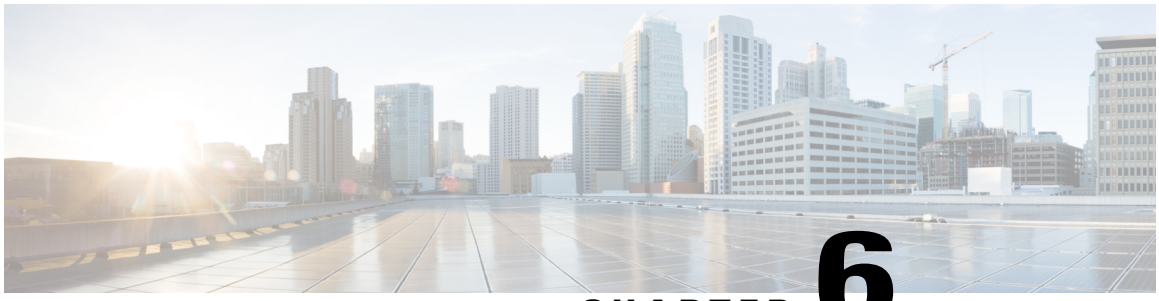

CHAPTER 6

Tenant Onboarding

Tenant Onboarding Overview

In Cisco UCS Director, tenants enable you to securely control and allocate the virtual and physical infrastructure of your data center to the different user groups, organizations, and customers. Your IT teams no longer need to manually provision infrastructure for users to deploy virtual machines (VMs) to run end user applications. Instead, you can configure tenant onboarding through Cisco UCS Director to define the infrastructure boundaries and resource limits that are applied automatically when a user makes a service request to provision a VM or an application.

Depending upon the infrastructure in your data center, you can use one of the following configurations for tenant onboarding:

- Resource groups for systems that include an integration with Cisco Application Centric Infrastructure (Cisco ACI) and Cisco Application Policy Infrastructure Controller (Cisco APIC)
- Virtual data centers for systems that does not include Cisco ACI or Cisco APIC

Cisco UCS Director tenants are essentially customers that share the compute, network, and storage resources that is configured for ACI in Cisco UCS Director.

Tenant Onboarding with Virtual Data Centers

Virtual data centers (vDCs) in Cisco UCS Director enable you to securely create and onboard tenants that represent user groups or customer organizations. For each tenant, the vDCs define the boundaries of the infrastructure that your user groups or customer organizations can access and use.

You can also use vDCs to control the VMs and infrastructure resources that user groups can view in the End User Portal. When users log in to request or manage VMs, those users can only see the vDCs assigned to their user group and the VMs that have been provisioned from those vDCs. The users cannot see the VMs or vDCs assigned to other user groups or customer organizations.

You can assign more than one vDC to a user group. This configuration enables you to set different approvers and/or different resources for VMs, based upon the type of application that the end user requests. If you assign multiple vDCs to a user group or customer organization, you can allow your users to choose the vDC where the VM should be created.

Note You cannot configure tenant onboarding with vDCs if your system includes an integration between Cisco UCS Director and Cisco Application Policy Infrastructure Controller (APIC).

For more information about tenant onboarding with vDCs, see the Cisco UCS Director Administration Guide.

Virtual Data Centers

A Virtual Data Center (VDC) is a logical grouping that combines virtual resources, operational details, rules, and policies to manage specific group requirements.

A group or organization can manage multiple VDCs, images, templates, and policies. Organizations can allocate quotas and assign resource limits for individual groups at the VDC level.

You can also define approvers specific to a VDC. The approvers assigned to a particular VDC must approve all service requests from users for VM provisioning.

Note There is a default VDC in Cisco UCS Director, and all discovered VMs are part of this default VDC. Discovered VMs are VMs that are created outside of Cisco UCS Director or were already created on VMware vCenter before Cisco UCS Director was installed. Cisco UCS Director automatically discovers such VMs and adds them to the default VDC.

A VM that is provisioned using a service request can be associated with a specific VDC. When you create a service request, you can choose the VDC on which this VM is provisioned. You can view a list of the VDCs that are available for a particular group and choose the required VDC when provisioning VMs.

Policies for Virtual Data Centers

You specify the infrastructure allocations in a vDC through a set of policies and configuration options. Each policy defines a subset of resources that the associated user group or customer organization can use and view. These policies include the following:

- System policy—System specific information, including which template to use, time zone, and OS specific information.
- Computing policy—Available compute resources, including the host or cluster, number of vCPUs, CPU configuration, and memory configuration.
- Network policy—Available network resources, including IP address configuration (DHCP or static IP address) and the option for the end user to choose vNICs for any VM.
- Storage policy—Available storage resources, including the datastore scope, type of storage to use, minimum conditions for capacity, and latency.
- Approvers—Names and contact email addresses of the required approvers for VMs provisioned from the vDC resources. By default, you can configure a maximum of two approvers per vDC

- Cost model—Cost and chargeback policy for all VMs provisioned from the vDC resources.
- ISO image mapping policy—ISO images available for VMs provisioned from the vDC resources.
- End user self-service policy—Actions or tasks that a user can perform on VMs provisioned from the vDC resources.

Computing Policies

Computing policies determine the compute resources that can be used during provisioning to satisfy group or workload requirements.

As an administrator, you can define advanced policies by mixing and matching various conditions in the computing policy.

Note We recommend that you thoroughly understand all the fields in the computing policy. Some combinations of conditions can result in no host machines being available during self-service provisioning.

Network Policies

The network policy includes resources such as network settings, DHCP or static IP, and the option to add multiple vNICs for VMs provisioned using this policy.

Storage Policies

A storage policy defines resources such as the datastore scope, type of storage to use, minimum conditions for capacity, latency, and so on.

The storage policy also provides options to configure additional disk policies for multiple disks, and to provide datastore choices for use during a service request creation.

Note Cisco UCS Director supports datastore choice during a service request creation for VM provisioning. You can enable or disable datastore choices for the end user during service request creation. The datastores listed depend upon the scope conditions specified in the storage policy that is associated with the VDC during the service request creation.

To use the datastore selection feature while creating a service request, the template for VM provisioning must have the disk type assigned as **System**. This is applicable for templates with single or multiple disks.

End User Self-Service Policy

An End User Self-Service Policy controls the actions or tasks that a user can perform on a VDC. The starting point for creating this policy is to specify an account type (for example, VMware). After you specify an account type, you can continue with creating the policy. After you create the policy, you must assign the policy to a vDC that is created with the same account type. For example, if you have created an end user

policy for VMware, then you can specify this policy when you create a VMware vDC. You cannot view or assign policies that have been created for other account types.

In addition to creating an end user self-service policy, Cisco UCS Director allows you to perform the following tasks:

- View—Displays a summary of the policy.
- Edit—Opens the **End User Policy** screen from which you can modify the description or the end user self-service options.
- Clone—Opens the **End User Policy** screen through which you can create an additional policy using the parameters defined in an existing end user self-service policy.
- Delete—Deletes the policy from the system. You cannot delete a policy that has a VDC assigned to it.

Important The tasks that a user can perform on a VDC are defined by the role that the user is mapped to and by the end user self-service policy assigned to the VDC. If you have upgraded to the current release, then the permissions to perform VM management tasks are retained in any pre-existing end user self-service policy. However, the permissions defined in the user role to which the user belongs takes precedence.

Resource Limits

If you want to ensure that a user group or customer organization does not exceed specific compute, network, or storage resource thresholds, you can add resource limits to the user group that apply across all vDCs assigned to that group. If you use resource limits, even if one or more vDCs assigned to the user group have resources available, Cisco UCS Director will reject service requests that exceed the resource limits.

For example, if you have a data center that includes the following resources:

- ESXi cluster with 10 hosts
- Compute resources that total 1,000 GHz
- Network resources with all the required network portgroups, IP subnets, and VLANs
- Storage resources of 100 datastores
- User groups for Finance, Sales, and Development
- Resource limits for the Finance user group of 300GHz compute and 20 datastores.
- Three vDCs assigned to the Finance user group, each with 100 GHz compute and 10 datastores

With this configuration, users in the Finance group can log into the End User Portal and view their available resource limits, the vDCs that are associated with their user group, and other information about their available resources, and deployed applications. If they have deployed applications across all 20 datastores in their resource limits, they cannot make any further service requests without first requesting an additional resource allocation.

As you can see from the following illustration, the combination of user groups, vDCs, policies, and resource limits determine the boundaries of the compute, network, and storage infrastructure that is available to each of the Finance and Sales organizations for application provisioning.

Figure 2: Example: Resource Limits for Finance and Sales Organizations

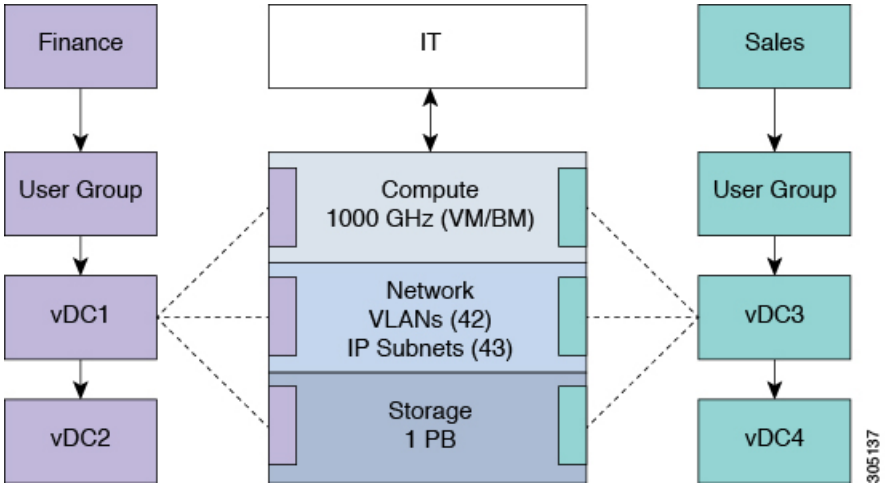

Tenant Onboarding with Resource Groups

Resource groups enable you to securely onboard resources for tenants that represent user groups and customer organizations. Each resource group is a pool of physical and/or virtual resources with specific capacities and capabilities. Based on the specific resource requirements, resources are assigned to a tenant (user groups or organizations) from this pool. You can share resources in a resource group across tenants or you can dedicate them to a specific tenant.

When you onboard a tenant through a resource group, you need to do the following:

1. Create one or more resource tags, if desired.
2. If you created resource tags, tag the physical and virtual resources that you want to make available to the tenant. If you did not create resource tags, you can allocate resources to a tenant based on the required capacity and capability.
3. Create a resource group.
4. Add infrastructure resources (devices) to the resource group, either manually or through associating resource tags with the group.

   You can add physical infrastructure resources, virtual infrastructure resources, or a combination of physical and virtual infrastructure resources to the group.
5. Define the environment specific inputs for each type of infrastructure resource associated with the resource group. These inputs can include domain and other information required to identify and assign the resources.
6. Create one or more service offerings to describe the requirements of a particular application that a user can provision.

7 Create one or more service classes within that service offering to define the requirements for a particular type of infrastructure resource required for the service offering. If you use resource tags, you can associate a service class with a particular resource tag.

8 Define a tenant profile to associate one or more service offerings with the resource group.

9 Develop the orchestration workflow required to deliver the service offering to the user as part of a service request.

Figure 3: Tenant Onboarding with Resource Groups Process Flow

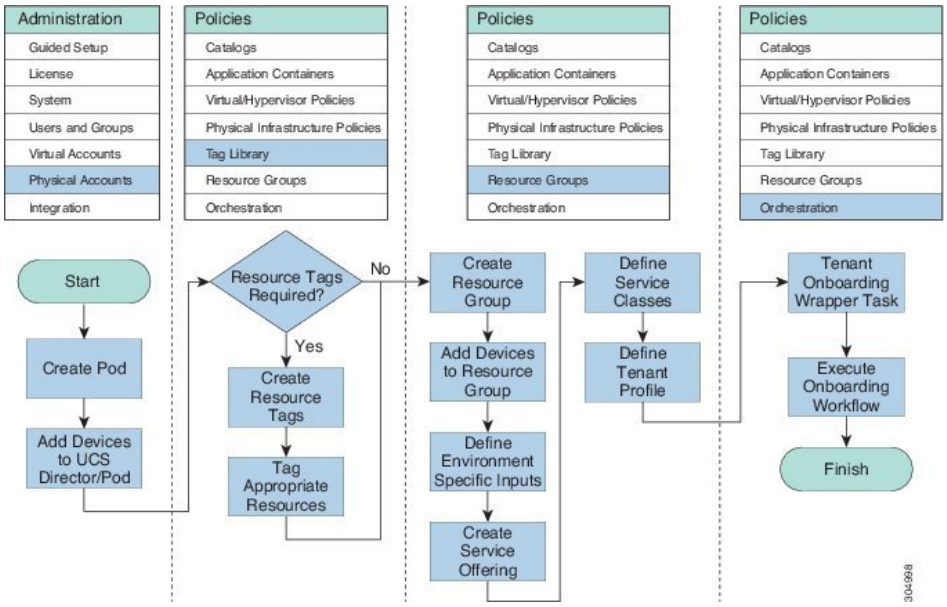

Note Resource groups support tenant onboarding only for systems that include Cisco Application Centric Infrastructure (ACI) and Cisco Application Policy Infrastructure Controller (APIC).

For more information about tenant onboarding with resource groups, including specific examples, see the Cisco UCS Director APIC Management Guide.

Resource Tags

Resource tags are an optional component that allow you to assign a label to a physical or virtual infrastructure resource or to a category of infrastructure resources. Each resource tag is a name-value pair that identifies the resource or category of resources in a way that is meaningful to your data center.

When you tag an object, you can also define applicability rules that determine how the tags are filtered and displayed. You can also organize your resource tags into tag libraries, which allow you to determine how you want to categorize the resources that have that tag applied. For example, you can create one or more tag libraries based on the following criteria:

- Resource type

- Capacity
- Quality
- Capability

You can use resource tags to identify specific resources that should be used by a tenant or to identify resources that are appropriate for a specific tier of service. For example, you can create a tag with a name of tier and a value of gold. The resources that you associate this tag are then identified as appropriate for the gold tier of service, whether they are compute, storage, or network infrastructure devices.

Once you have tagged a set of resources, you can associate those resources with resource groups. The resource tags enable you to easily identify how the infrastructure resources should be grouped. They also define the boundaries of the infrastructure resources that are available to a particular resource group.

Resource Groups

You can use a resource group to select the appropriate resources for a tenant based on the requirements of an application. Additional concepts, such as a service offering, tenant profile, application profile, and resource group, are all required. Using these resource group concepts, you can onboard tenants and deploy applications based on a dynamic selection of resources. You can share resources in a resource group across tenants or you can dedicate them to a specific tenant.

A resource group is a pool of resources. Each group can contain physical infrastructure resources, virtual infrastructure resources, or a combination of physical and virtual infrastructure resources. Resource groups enable you to onboard tenants into Cisco UCS Director with minimum intervention.

As an infrastructure administrator or system administrator, you can add physical or virtual accounts to a resource group one at a time. Also, you can assign a pod to a resource group where all the accounts in the pod are added to the resource group.

When an account is added to a resource group, the resource group by default announces all the capabilities and capacities for objects for that account as resource group entity capacities and capabilities. With Cisco UCS Director, you can selectively disable certain capacities or capabilities from the resource group.

Service Offerings

A service offering defines the resources required to provision an application. Each service offering must include one or more service classes that represents the capacity and capability needed for the following resource layers:

- Virtual Compute
- Virtual Storage
- Virtual Network
- Physical Compute
- Physical Storage
- Physical Network
- Layer 4 to Layer 7 Services

When you define a service offering, you can specify the usage of resource groups as one of the following:

- Shared—The resources are shared among the applications or tenants.
- Dedicated —The resources are dedicated to a single application or tenant.

Based on the capacity, capability, and resource tags defined in the service offering, the resource groups are filtered and the matching resource groups are selected for further processing in the tenant onboarding and application deployment.

Tenant Profiles

Tenant profiles represent the pairing of one or more service offerings with one or more resource groups. Each tenant profile defines the characteristic of infrastructure requirements and application requirements.

You can create a tenant profile to meet each possible combination of customer and application. You can associate a tenant profile with multiple service offerings and choose a resource group for each service offering. A tenant profile can be shared by more than one tenant.

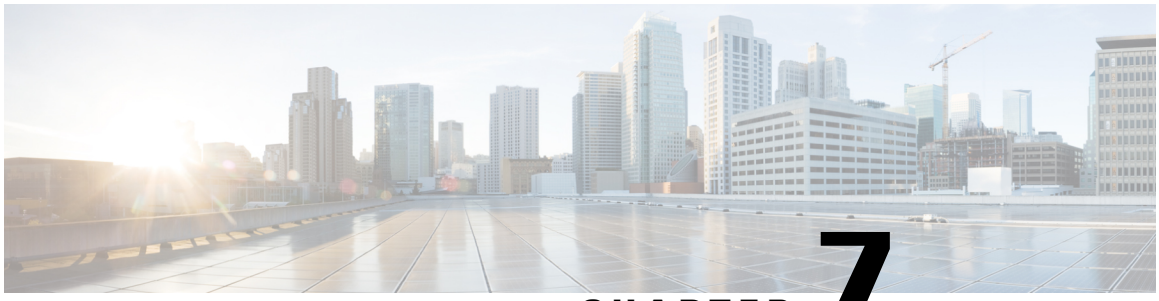

CHAPTER 7

Application Provisioning

Application Provisioning Overview

After you have allocated your resources among your user groups, you can set up application provisioning for your end users. Application provisioning includes the following:

- (Optional) Application categories for each cloud account that define the workload for a particular type of application. For example, you can create an application category for a CPU intensive web server application and another for a storage intensive database application.
- Application container templates that include policies, workflows, and VM tempates, and that determine how a specific application is to be provisioned for the end user. How you create the application container template depends upon whether your data center includes an integration with Cisco Application Centric Infrastructure (Cisco ACI).
- Catalog items that represent a VM template or application container template. End users can only see the catalog items that are available within the vDCs for their user group.

Application Categories

Application categories are an optional configuration that enable you to define the type of workload for a VM. If you do not use application categories, Cisco UCS Director assumes that all VMs provisioned for your users are generic VMs and configures them to handle CPU-intensive workloads. Whether you choose to use the default application categories or to create your own, you can provide your users with a pre-defined set of workloads that match their application needs.

The workload options for application categories include the following:

- CPU intensive
- Network I/O intensive
- Disk I/O intensive
- Memory intensive
- Any combination of the above

After you create your application categories, you can go to the desired cloud account and assign the vDC policies to the application categories. This assignment determines the boundaries of the infrastructure where the application can be provisioned. You can also use application categories to allocate clusters based on the type of application. For example, Cluster 1 is allocated for Web applications and Cluster 2 is allocated for database applications.

When an application category is chosen by a user, Cisco UCS Director uses the vDC assignment to determine which location, within the boundary of the vDC, best meets the application's workload needs. For example, if the user chooses a CPU-intensive application category, Cisco UCS Director provisions the application in the available infrastructure with the least CPU utilization.

For more information about application categories, see the Cisco UCS Director Administration Guide.

Application Containers

Application containers are a templatized approach to provisioning applications for end users. When you want to configure Cisco UCS Director to allow users to provision applications, you must create one or more application containers with the appropriate policies, workflows, and templates. The application container determines how the application is provisioned for the end user. It can define some or all of the following:

- Physical server or virtual machine
- Amount of reserved storage
- Maximum available CPU
- Maximum available memory
- Version of operating system
- Range of VLANs
- Gateway or firewall, if desired
- Load balancer, if desired
- Required approvals, if desired
- Costs associated with the application container, if desired

Cisco UCS Director supports multiple types of application containers. The type of application container that you need to implement depends upon your deployment configuration. The steps required to configure the application container depend upon which type you plan to implement.

For more information about application containers, see the Cisco UCS Director Application Container Guide.

Types of Application Containers

There are several application container types for use in various deployment scenarios:

- Fenced Virtual—The most common type of application container for use with VMs.
- Virtual Secure Gateway (VSG)—This container type is used to provide for enhanced security in virtual environments.
- Application Centric Infrastructure Controller (APIC) container—This container type is used in APIC deployments. For more information, see the Cisco UCS Director APIC Management Guide and the Cisco UCS Director Application Container Guide.
- Fabric—This container type is used in Dynamic Fabric Automation (DFA) network deployments. For more information on application containers in a DFA network, see the Cisco UCS Director Unified Fabric Automation Guide.
- Virtual Application Container Services (VACS) container—This container type is used in deployments of Cisco Virtual Application Cloud Segmentation Services.

Note This container type is available in Cisco UCS Director only if the VACS modules are installed. For more information, see the Cisco Virtual Application Cloud Segmentation Services Configuration Guide.

Options for Application Containers

In addition to the different types of application containers, Cisco UCS Director supports a variety of different options for those application containers. You can configure an application container to provision either a single tier or a multi-tier application, and you can choose to add application services, such as a load balancer or a firewall, that will be provisioned with the application. Before you create your application container, you need to decide which of these options you want to implement.

Single Tier Application

A single tier application is typically a simple application that is contained within a single VM. This VM includes the application interface and database and is usually in the application server tier.

Multi-Tier Application

A multi-tier application typically requires multiple VMs. The number of VMs depends upon several factors, such as the application complexity and the users. For example, you could have a three tier application with a VM for each of the web server, application server, and database server.

If required by your application architecture, a tier can include multiple networks.

Firewall

You can add a firewall as an internal or external gateway to the application. The gateway redirects and creates a tunnel through the firewall to the application. If you use a firewall, the user does not need to know the IP address of any of the application VMs. Instead, the user logs in through the IP address of the gateway and is

redirected to the application or database VM. You can also create rules that define the type of traffic that is permitted through the gateway.

You define the firewall that you want to use in a tiered application gateway policy. This policy is then included in the application container template. You can add one of the following types of firewalls:

- Linux VM—This default option provisions the appropriate firewalls and NAT rules on the VM.
- Cisco ASA—This physical gateway allows one-way (inside to outside) connections without an explicit configuration for each internal system and application.
- Cisco ASAv—This virtual gateway is typically cloned from a VM template during the application provisioning. It allows one-way (inside to outside) connections without an explicit configuration for each internal system and application.

Load Balancer

You can include a load balancer for a multi-tier application. The load balancer manages communication and workloads between the servers. For example, a load balancer can identify and redirect a user to one of several application servers to ensure that none of the application servers are overloaded.

For information about supported load balancers, see the Cisco UCS Director Compatibility Matrix.

General Application Container Creation Process

The following process explains the creation of an application container in Cisco UCS Director.

Note The process for creating an Application Centric Infrastructure Controller (APIC) managed container is different.

1. If a gateway is required, create a tiered application gateway policy.
2. If a load balancer is required, create a load balancer policy.
3. Create a virtual infrastructure policy to define the cloud account, the type of container and, if appropriate, the tiered application gateway and load balancer policies.
4. Create an application container template.
   a. Add networks (one network per application tier).
   b. Add virtual machines and bare metal servers.
   c. Add a compute policy, storage policy, network policy, and systems policy. If desired, you can also add a cost model.
   d. Add an end user self-service policy and configure the self-service options.
   e. Add the container setup workflow required to deliver the service offering to the user as part of a service request. The workflow must consider the type of container and the application to be provisioned.
5. Create a container based on the container template.

The figure illustrates the creation of the general application container template within Cisco UCS Director.

Figure 4: Process for Creating a General Application Container Template

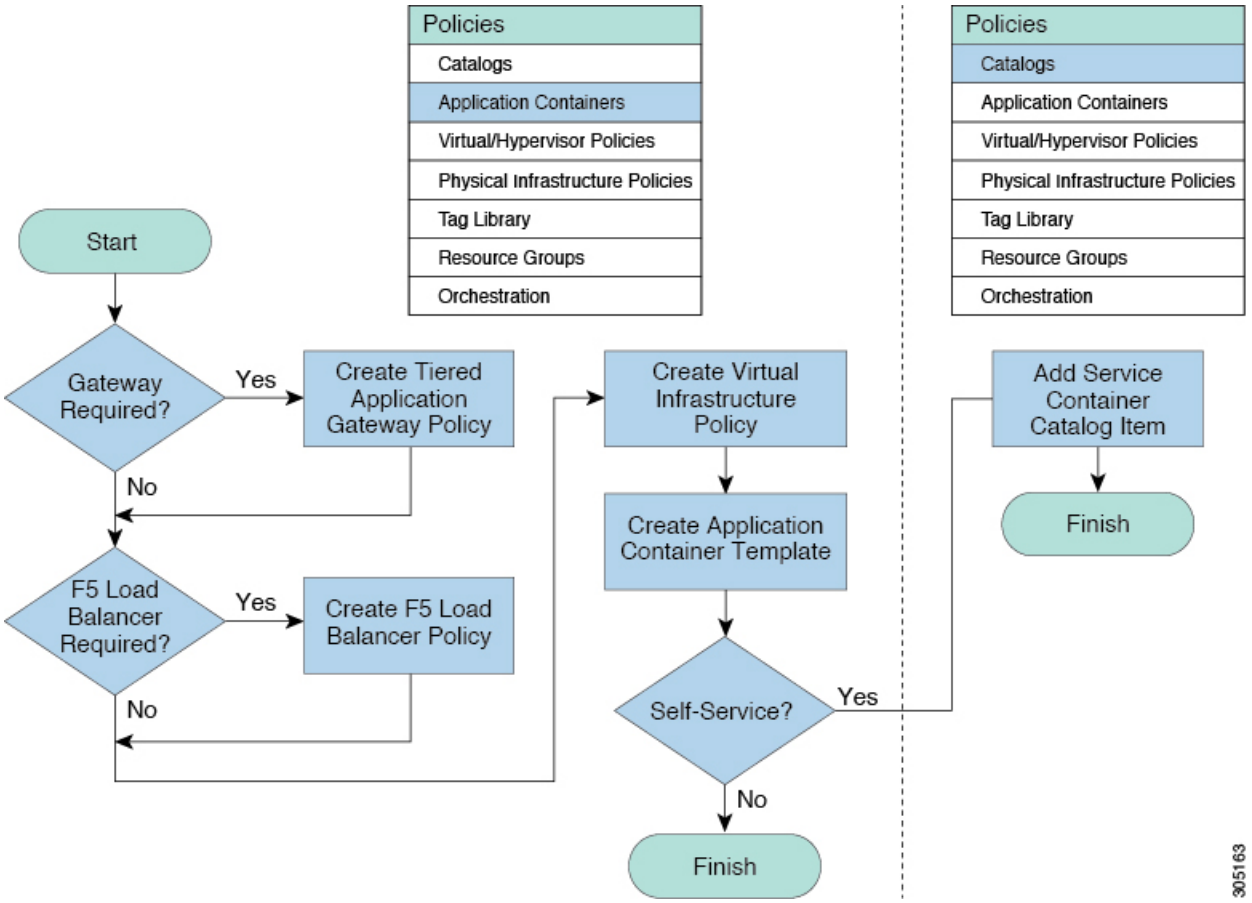

APIC Application Containers

Cisco UCS Director lets you create application containers that support a Cisco Application Policy Infrastructure Controller (APIC). For additional information, see the Cisco UCS Director APIC Management Guide for this release. APIC application containers let you do the following:

- Establish networks in a VMware environment.
- Provision multiple VMs from a network.
- Provide a way to isolate those networks using gateways (for example, ASAv).
- Allow load balancing the container network using VPX or SDX load balancers.
- Use a Cisco Application Centric Infrastructure (ACI).
- Provision a bare metal server and/or VMs.

APIC Application Container Limitations

Cisco UCS Director APIC application containers have the following limitations:

- Tenant onboarding must be done before container creation and usage.
- Resource groups must contain the accounts necessary to manage a container's resources. This can be any combination of storage, compute, network, and virtual resources.
- For application container configuration that requires physical servers, only UCS managed servers are currently supported.

APIC Application Container Creation Process

The figure below illustrates the flow of the APIC Application Container creation process within Cisco UCS Director.

Figure 5: Process for Creating an APIC Application Container

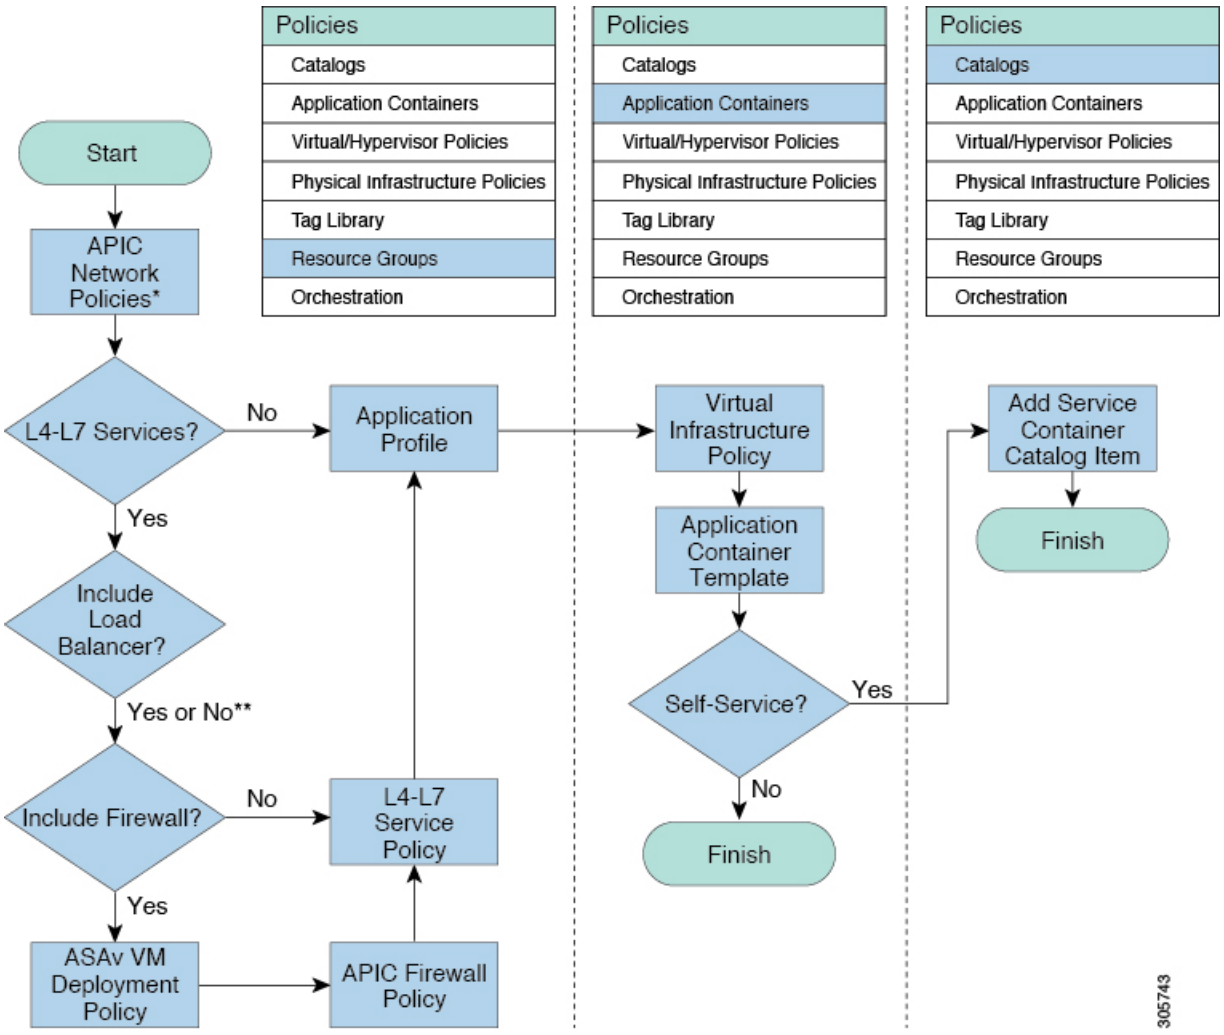

* Optional. APIC Network Policies are only needed to override default APIC network entity properties.

** Load Balancer is an L4-L7 service but does not require a separate policy.

Catalogs

Catalogs are collections of predefined workflows and application containers that enable a user to request and provision applications, VMs, or bare metal servers. You can choose to make a catalog available to end-users

for self-service provisioning through the End User Portal, only to administrators for provisioning through Cisco UCS Director, or both.

When you publish a catalog, you can choose one of the following types.

- Standard—For provisioning VMs with an operating system image from a specified cloud account
- Advanced—For publishing orchestration workflows as catalogs
- Service Container—For publishing application containers as catalogs

Catalog Organization

The End User Portal organizes the catalogs in folders by name. A set of default folders, named for the catalog types, are automatically created and sorted alphabetically. However, you do not have to use these default folders. When you create a catalog, you can place it in a default folder or into a custom folder with a name that you choose. For example, if you create a Standard Catalog to provision a Windows 2012 operating system for a VM, you can allow Cisco UCS Director to place it in the Standard folder, or you can create a custom folder with a name that identifies the operating system, such as `VM-Windows2012OS`.

Catalog folders created by administrators in Cisco UCS Director display in the End User Portal.

In the **Catalog** pane, you can use the **Move Up** and **Move Down** icons to change the organization of the catalog folders.

Note You can also use the **Manage Folder** icon to choose a folder and move it in relation to existing folders.

While creating a catalog in Cisco UCS Director, an administrator can specify a user group or a specific users that can view the catalog. If you are a user in the selected group, then your view is populated with the appropriate catalogs.

Self-Service Provisioning

You can provision virtual machines (VMs) or applications through self-service provisioning. To provision a VM or an application using self-service provisioning, you must first create a service request. This action initiates a VM-creation workflow that includes the following:

- Budget validation
- Dynamic resource allocation
- Approval
- Provisioning
- Lifecycle setup
- Notification about the status of service requests

Service Requests

You can use the self-service provisioning feature to create a service request to provision virtual machines (VMs), services, or applications. The service request process produces a provisioning workflow for VM creation that includes the following actions:

- Budget validation
- Dynamic resource allocation
- Approvals
- Provisioning
- Lifecycle setup and notification

Note If you change the number of CPU Cores or memory allocation while in the **Deployment Configuration** screen, the total cost is automatically updated and displayed.

To provision a VM or execute an orchestration workflow, you must first create a service request. If desired, you can require approval from one or two administrators or designated users before the VM is provisioned or the workflow executed. VMs can be immediately approved or scheduled to be approved within a maximum of 90 days from the original request.

Service Request Workflow and Details

After you create a service request, you can check its status and workflow, cancel the request, resubmit the request, and so on. These actions are controlled by the toolbar buttons at the top of the service request lists.

Service Request Workflow

The **Workflow Status** box displays details about the service request and the workflow steps. A typical service request workflow to provision a VM includes the following steps:

1. Initiation—Service request is initiated by the user.
2. Resource Allocation—Required resources, such as virtual compute, are allocated to the VM.
3. Approval—VM provisioning is approved, if required. During this step, an email is sent to the approvers defined in the catalog chosen for VM provisioning.
4. Provision—VM is created and provisioned.
5. Set Up Lifecycle Schedule—Lifecycle scheduling is configured with the setup, scheduled times, and termination times.
6. Notify—User is notified by email that the VM has been created and provisioned.

Following is a graphical representation of the workflow.

Figure 6: Catalog Service Request Workflow

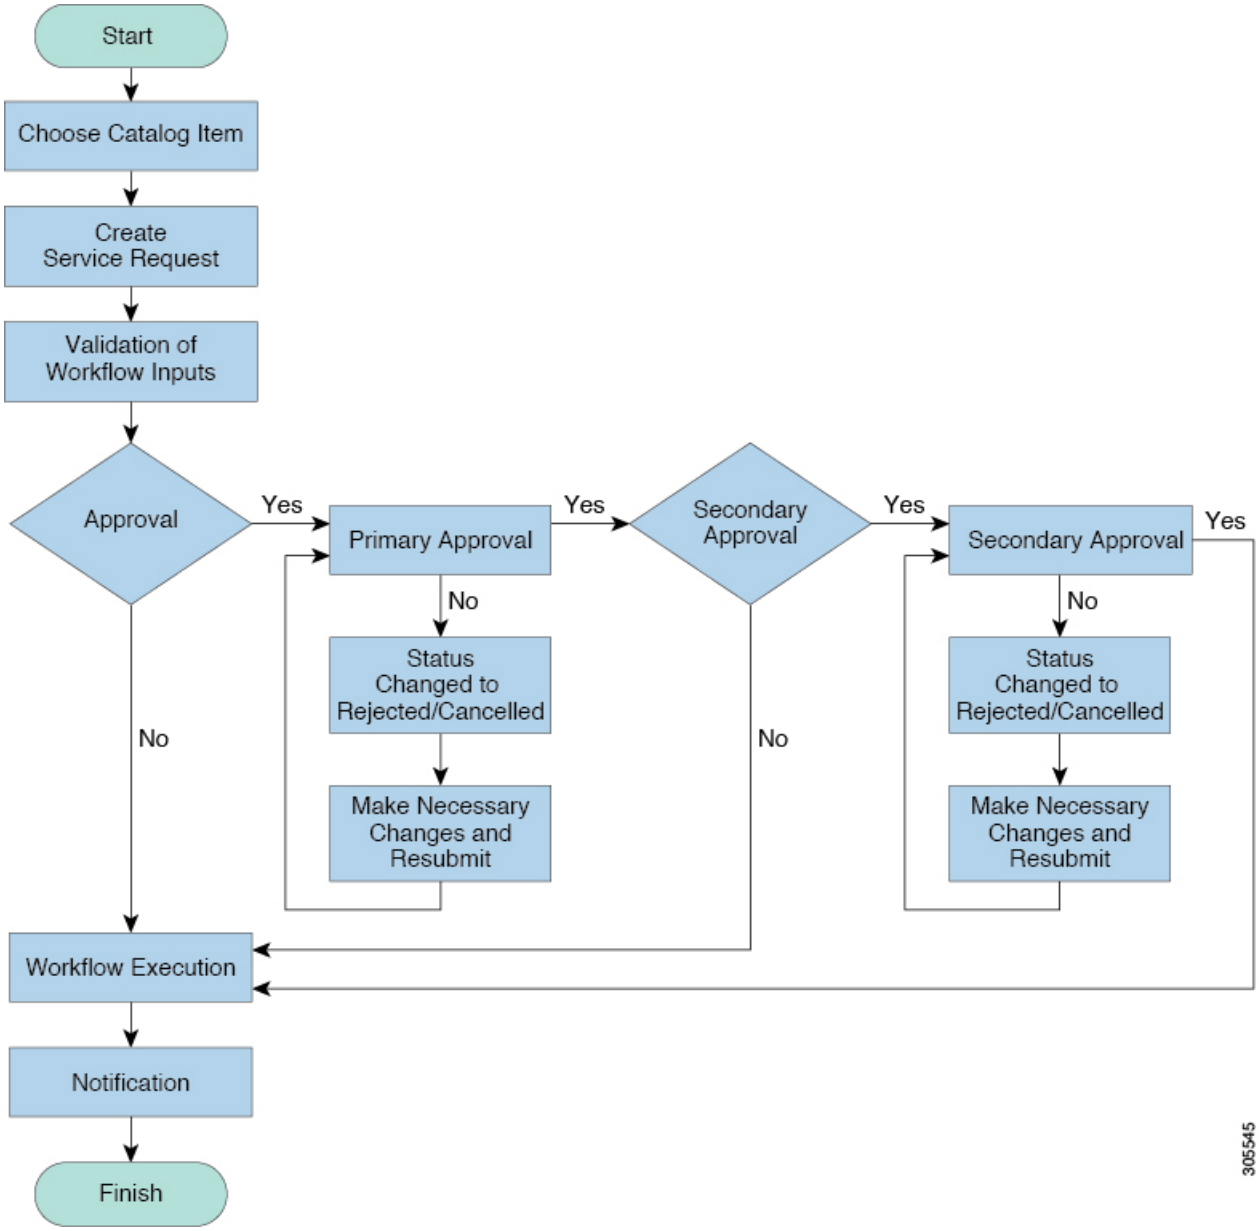

Optional service request workflow steps include Budget Watch and Check Resource Limits:

- Budget Watch—An administrator has to enable budgeting for a group. This step determines if a sufficient budget is available for provisioning a new VM in that group.
- Check Resource Limits—Resource limits for a group must be enabled by an administrator. This step determines if sufficient resources are available for provisioning a new VM in that group.

Any user who has been assigned the **Read-Group Service Request** permission can view the progress of a service request.

Service Request Details

Service Request details include items under Overview, Ownership, Catalog Information, and the Current Status of the service request, as follows:

Name	Description
Overview	
Request ID	The service request ID number.
Request Type	The type of request (in this case, creating a VM).
VDC	The VDC where the VM is provisioned.
Image	The image from which the VM is provisioned.
Request Time	The time of the service request creation.
Request Status	The status of the service request as Complete, Canceled, Failed, and so on.
Comments	Any comments.
Ownership	
Group	The group to which the service request initiating user belongs.
Initiating User	The user who has initiated the service request.
Duration Hours	The amount of time that the VM is active. If this time is defined, the VM is deleted after the specified time.
Scheduled Time	The time at which the VM is provisioned. If defined, the VM is provisioned at 6 a.m. on the scheduled date. If not defined, the VM is provisioned when the workflow steps for the service request are complete.
Catalog Information	
VDC Owner Email	The email ID provided by the administrator when creating a VDC.
Approving Users	The user (if defined) who must approve the service request for VM provisioning.

Name	Description
Catalog Name	The catalog item name from which the VM is provisioned.
Catalog Description	The catalog item description.
Service Request Cost	The cost (projected) of provisioning the VM. This cost is determined based on the Cost Model that is defined for the catalog item.

You can view the status of each workflow step. Details such as warning or error messages and the time of the request are also displayed. The workflow steps are color-coded to indicate their status:

Color Code	Description
Gray	The step is incomplete.
Green	The step completed successfully.
Red	The step failed. The reason for failure is also described.
Blue	More input is required for the step to complete. For example, an approver was defined for a service request, and until the request is approved, this step is incomplete.

Note Approvers may look under the **Approvals** tab to see their assigned service requests.