Cisco CloudCenter Cost Optimizer 5.5Documentation Documentation

First Published: July 19, 2021

Last Modified: July 19, 2021

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

© 2021 Cisco Systems, Inc. All rights reserved.

1. Cost Optimizer 5.5 Home . . . 3
1.1 Release Notes . . . 4
1.1.1 Cost Optimizer 5.5.0 . . . 5
1.2 What is Supported? . . . 7
1.2.1 Supported Public Clouds . . . 8
1.2.2 Supported Container Clouds . . . 12
1.2.3 Supported Datacenters and Private Clouds . . . 13
1.3 Getting Started . . . 14
1.3.1 Cost Optimizer Overview . . . 15
1.3.2 Cost Optimizer Architecture . . . 17
1.3.3 Access and Roles . . . 21
1.3.4 Email Settings in Cost Optimizer . . . 23
1.3.5 UI Behavior . . . 24
1.4 Configure Clouds . . . 26
1.4.1 Cloud Overview . . . 27
1.4.2 Configure Cloud End-to-End . . . 32
1.4.2.1 Configure a Google Cloud . . . 33
1.4.2.2 Configure a Kubernetes Cloud . . . 50
1.4.2.3 Configure a vCD Cloud . . . 70
1.4.2.4 Configure a vCenter Cloud . . . 84
1.4.2.5 Configure an AWS Cloud . . . 102
1.4.2.6 Configure an AzureRM Cloud . . . 116
1.4.2.7 Configure an IBM Cloud . . . 137
1.4.2.8 Configure an OpenStack Cloud . . . 150
1.4.2.9 Configure an Outscale Cloud . . . 170
1.4.3 Cloud Remote . . . 183
1.4.4 Cloud Maintenance . . . 184
1.5 Cost Groups Configuration . . . 187
1.5.1 Cost Groups UI . . . 188
1.5.2 How Do I... . . . 191
1.6 Allocate Budgets . . . 193
1.7 Cost Optimizer Dashboard . . . 197
1.8 Cost Reports . . . 203
1.8.1 Cost Reports Overview . . . 204
1.8.2 Cost by Cloud Group . . . 210
1.8.2.1 Cost by Billing Units in a Cloud Group . . . 211
1.8.2.2 Cost by Tags in a Cloud Group . . . 213
1.8.3 Cost by Category . . . 214
1.8.4 Invoice Report . . . 216
1.8.4.1 Invoice by Category . . . 217
1.8.4.2 Invoice by Region . . . 219
1.8.5 Cost Over Time . . . 221
1.8.6 Cost by Tags . . . 223
1.8.7 Cost by Cost Group Type . . . 225
1.8.8 Cost by Organization Hierarchy . . . 227
1.9 Budget Reports . . . 229
1.9.1 Budget Reports Overview . . . 230
1.9.2 Budget Overspenders . . . 231
1.9.3 Budget Underspenders . . . 232
1.9.4 Budget By Cloud . . . 233
1.9.5 Budget By Cost Group Type . . . 235
1.10 Inventory . . . 236
1.10.1 Inventory Overview . . . 237
1.10.2 Virtual Machines . . . 242
1.10.3 Kubernetes Workloads . . . 248
1.10.4 Storage Volumes . . . 252
1.10.5 Services . . . 255
1.10.6 Inventory States . . . 259
1.11 Rightsizing . . . 260
1.12 Suspension Candidates . . . 270
1.13 Unused Volumes . . . 277
1.14 Reserved Instances . . . 283
1.14.1 Reserved Instances Overview . . . 284
1.14.2 RI Subscription Report . . . 286
1.14.3 RI Opportunities Report . . . 295
1.15 Administration . . . 298
1.15.1 Admin Tasks in Cost Optimizer . . . 299
1.15.2 Settings Page . . . 300
1.15.3 Data Collection . . . 303
1.15.4 Alerts Page . . . 304
1.15.5 Tag-Based Cost Reporting . . . 306
1.16 Troubleshooting . . . 308
1.16.1 Cost Optimizer Troubleshooting . . . 309
1.16.2 Scheduling MongoDB . . . 311
1.17 Cost Optimizer API . . . 312
1.17.1 API Overview . . . 313
1.17.2 API Authentication . . . 320
1.17.3 API Key . . . 321
1.17.4 Base URI Format . . . 323

1.17.5 HTTP Status Codes 325
1.17.6 CSRF Token Protection 326
1.17.7 API Permissions 328
1.17.8 Synchronous and Asynchronous Calls 330
1.17.9 Cost and Inventory Calls 5.5.0 332
1.17.10 Recommendation Calls 5.5.0 333
1.17.11 Cost Groups Calls 5.5.0 334
1.17.12 Tags Collector Calls 5.5.0 335
1.17.13 Cloud Setup Calls 5.5.0 336
1.17.14 Remedial Actions on Recommendations for Cost Optimizer 5.5.0 337

Cost Optimizer 5.5 Home

CloudCenter Cost Optimizer 5.5 Documentation

System Announcements

Cisco released the following Cost Optimizer releases:

- Cost Optimizer 5.5.0released on July 19, 2021

Search Cost Optimizer 5.5 Documentation

Recent Updates

Supported Datacenters and Private Clouds
updated 4 minutes ago
view change
API Overview
updated 5 minutes ago
view change
Configure a Kubernetes Cloud
updated 5 minutes ago
view change

Release Notes

Cost Optimizer Release Notes

- Cost Optimizer 5.5.0

Cost Optimizer 5.5.0

Cost Optimizer 5.5.0 Release Notes

- Release Date
- Installation and Upgrade
- Upgrade Path
- Clouds
- Cost Optimizer UI
- API
- Resolved Issues

Release Date

First Published: July 19, 2020

Installation and Upgrade

- **CloudCenter Suite:**
  - Cost Optimizer cannot be installed separately and must be installed as a part of the CloudCenter Suite UI. SeeSuite Admin 5.2.0release notesfor additional details.

    > If you upgrade Cost Optimizer, you must also upgrade Workload Manager, and vice versa, as both modules use the same shared APIs.

- **Cost Optimizer:**
  - The Optimizer Admin can upgrade Cost Optimizer at the suite level to the latest version of the software. See Update Modulefor additional context.
  - When upgrading to Cost Optimizer 5.5.0, it is recommended that you upgrade from Cost Optimizer 5.4.3 and meet all upgrade requirements specified in the Workload Manager 5.5.0 release notes.

    > If you upgrade theWorkload Manager, you must also upgrade the Cost Optimizer and vice versa as both modules use shared APIs.

- **Cloud Remote:**
  - When updating to Cost Optimizer 5.5.0, you must also update all instances of Cloud Remote to Cloud Remote 5.5.0.
  - SeeCloud Remote (Conditional)for additional details.

Upgrade Path

You can upgrade from Cost Optimizer 5.4.3 and later versions to Cost Optimizer 5.5.0.

Clouds

- **VMware**:
  - Workload Manager 5.5.0 introduces support forVMware vCenter 7.0.
  - SeeSupported Datacenters and Private Cloudsfor additional details.
- **Kubernetes**:
  - The followingKubernetesversions are supported effective Workload Manager 5.5.0:
    - Kubernetes 1.18
    - Kubernetes 1.19
  - SeeSupported Container Clouds for additional details.

Cost Optimizer UI

- **Browser Compatibility**: SeeBrowser Compatibilityfor a list of compatible browsers.
- **Localization**: Cost Optimizeris only available in the English language.
- Refer to the Suite Admin for additional context onSuite ArchitectureandAdministration and Governance.

API

See theWorkload Manager 5.5.0Release Notes > *API* section for changes and updates to the Cloud Setup Calls 5.5.0.

Resolved Issues

- **CSCvw72237**: While the deployment was up and running and thedeployment was in the process of being resized from theWorkload Manager UI, the instance type was updated in cloud, the action history displays the resize operation as having completed successfully, but instance type in Workload ManagerUI was not updated. The**Sync VM** actionoperation too did not succeed at this point.
  Resolution:Cost Optimizer 5.5.0 includes a fix to address this issue and the instance type is updated in theWorkload Manager UI as designed. The **Sync VM** action operation succeeds as designed.

What is Supported?

What is Supported?

- Supported Public Clouds
- Supported Container Clouds
- Supported Datacenters and Private Clouds

Supported Public Clouds

Supported Public Clouds

Cisco supports the following public clouds and managed private clouds for the Workload Manager and Cost Optimizer modules.

The following table identifies the cloud regions that are currently available out-of-the-boxWorkload Manager and Cost Optimizer modules.

Cloud Family	Available Regions
Amazon Web Services (AWS)	Asia Pacific (Mumbai)
	Asia Pacific (Osaka-Local)
	Asia Pacific (Seoul)
	Asia Pacific (Singapore)
	Asia Pacific (Sydney)
	Asia Pacific (Tokyo)
	AWS GovCloud (US-East)
	AWS GovCloud (US-West)
	Canada (Central)
	CN North (Beijing)
	China (Ningxia) ⚠ Invoice reports in Cost Optimizer are not supported for China regions.
	EU (Frankfurt)
	EU (Ireland)
	EU (London)
	EU (Paris)
	EU (Stockholm)
	South America (Sao Paulo)
	US East (N. Virginia)
	US East (Ohio)
	US West (N. California)
	US West (Oregon)
Google Cloud Platform	Central US (Iowa)
	Eastern Asia-Pacific (Hong Kong)
	Eastern Asia-Pacific (Taiwan)
	Eastern US (Northern Virginia)
	Eastern US (South Carolina)
	European West (Frankfurt)
	European West (London)
	European West (Netherlands)
	Northeastern Asia-Pacific (Japan)
	Northern America (Canada)

	Northern Europe (Finland)
	South Eastern Asia-Pacific (Singapore)
	South Eastern Australia (Sydney)
	Southern America (Sao Paulo)
	Southern Asia-Pacific (Mumbai)
	Western Europe (Belgium)
	Western US (California)
	Western US (Oregon)
IBM	Amsterdam 01 (ams01)
	Amsterdam 03 (ams03)
	Chennai 01 (che01)
	Dallas 05 (dal05)
	Dallas 06 (dal06)
	Dallas 09 (dal09)
	Dallas 10 (dal10)
	Dallas 12 (dal12)
	Dallas 13 (dal13)
	Frankfurt 02 (fra02)
	Frankfurt 02 (fra02)
	Frankfurt 05 (fra05)
	Hong Kong 02 (hkg02)
	Houston 02 (hou02)
	London 02 (lon02)
	London 04 (lon04)
	London 05 (lon05)
	London 06 (lon06)
	Melbourne 01 (mel01)
	Milan 01 (mil01)
	Montreal 01 (mon01)
	Oslo 01 (osl01)
	Paris 01 (par01)
	Queretaro 01 (mex01)
	San Jose 01 (sjc01)
	San Jose 04 (sjc04)
	San Jose 04 (sjc04)
	Sao Paulo 01 (sao01)
	Seattle 01 (sea01)
	Seattle 01 (sea01)
	Seoul 01 (seo01)
	Singapore 01 (sng01)
	Sydney 01 (syd01)

	Sydney 04 (syd04)
	Sydney 05 (syd05)
	Tokyo 02 (tok02)
	Tokyo 04 (tok04)
	Tokyo 05 (tok05)
	Toronto 01 (tor01)
	Washington, DC 01 (wdc01)
	Washington, DC 04 (wdc04)
	Washington, DC 06 (wdc06)
	Washington, DC 07 (wdc07)
Microsoft Azure	Australia Central (Canberra)
	Australia Central 2 (Canberra)
	Australia East (New South Wales)
	Australia Southeast (Victoria)
	Brazil South (sao Paulo State)
	Canada Central (Toronto)
	Canada East
	Central India (Pune)
	China East (Shanghai)
	China North (Beijing)
	East Asia (Hong Kong)
	Europe North (Ireland)
	Europe West (Netherlands)
	France Central (Paris)
	France South (Marseille)
	Germany Central (Frankfurt)
	Germany North
	Germany Northeast (Magdeburg)
	Germany West Central
	Japan East (Saitama)
	Japan West (Osaka)
	Korea South (Busan)
	South Africa North (Johannesburg)
	South Africa West (Cape Town)
	South India (Chennai)
	Southeast Asia (Singapore)
	Switzerland North (Zurich)
	Switzerland West (Geneva)
	UAE Central (Abu Dhabi)
	UAE North (Dubai)
	UK South (London)

	UK West (Cardiff)
	US Central (Iowa)
	US East (Virginia)
	US East 2 (Virginia)
	US Gov Arizona
	US Gov Texas
	US Gov Virginia
	US North Central (Illinois)
	US South Central (Texas)
	US West (California)
	US West 2 (West US 2)
	US West Central (West Central US)
	West India (Mumbai)
Outscale	US East 2 (N. Virginia)
	US West 1 (N. California)
	EU (France)
	Asia Pacific (Hong-Kong)

Supported Container Clouds

Supported Container Clouds

- Overview
- Requirements
- Upstream Support and Capability

Overview

A container cloud relies on a*container*infrastructure that is configured by an administrator outside of Workload Manager. Currently, Workload Manager supports one container cloud: Kubernetes cloud.

Requirements

Kubernetes cloud configurations require:

- Kubernetesversion support
    - Kubernetes 1.8
    - Kubernetes 1.9
    - Kubernetes 1.10
    - Kubernetes 1.11
    - Kubernetes 1.12
    - Kubernetes 1.13
- A single Kubernetes cluster with an implicit default region
- One or more cloud accounts
- Cloud settings API endpoint
- Instance types (fractional CPU and memory)

Upstream Support and Capability

Workload Manager supports*upstream*Kubernetes setups.*Upstream*refers to any bare Kubernetes setup likeGoogle Kubernetes Engine(GKE), Amazon Elastic Container Service for Kubernetes (EKS),Cisco Container Platform, and so forth as these environments expose the Kubernetes APIs to users. This term does not include platforms that only use Kubernetes and then add on their own APIs.

Workload Manager's API layer handles configuration tasks such as application deployment for Kubernetes pods at the time of application deployment, Workload Manager dynamically creates the application pod information, which can be in Kubernetes as YAML or JSON files. Workload Managerdynamicall y deploys applications based on the Workload Managerapplication profile.While you cannot directly modify the application pod information that is dynamically created, you can edit the Workload Manager application profile in JSON format.

When creating an application profile, users define the network service. Workload Manager uses these user-configured network settings to automatically deploy load balancers through Kubernetes. SeeContainer Service>*Deploying a Container Service*>*Network Services*for details.

The Firewall Rules in the application profile correspond to a Network Policy Ingress rules in Kubernetes.SeeContainer Service>*Deploying a Container Service*>*Firewall Rules*for details.

Supported Datacenters and Private Clouds

Supported Datacenters and Private Clouds

The Workload Manager and Cost Optimizer modules support the datacenters or private clouds built using the following technology stacks.

Cloud Family	Version
VMware vCloud Director	VMware vCloud Director 1.1
VMware vCenter	VMware vCenter 6.5
	VMware vCenter 6.7
	VMware vCenter 7.0
OpenStack	OpenStackNewton
	OpenStack Mitaka
	OpenStack Pike
	OpenStack Queens

To compute costs in Cost Optimizer, you must specify thecompute and storage costs for an instance family that is auto-discovered.

Cisco does not provide out-of-box image mapping for datacenters or managed private clouds. You must manually import the physical images you need to deploy and map the appropriate logical images to those physical images. SeeImagesfor more context.

Getting Started

Getting Started

- Cost Optimizer Overview
- Cost Optimizer Architecture
- Access and Roles
- Email Settings in Cost Optimizer
- UI Behavior

Cost Optimizer Overview

Cost Optimizer Overview

- Overview
- Terminology
- Features
- Infrastructure
- Module Update Considerations
- Logging In to Cost Optimizer
- Related Information

Overview

Cost Optimizer is a comprehensive cloud cost management and optimization solution thatanalyzes cloud-deployed workloads and consumption patterns and identifies cost-optimization strategies. The Cost Optimizer solution helps you to rightsize your cloud workload instances, minimize overprovisioning, and avoid paying for resources that do not deliver business value.

Terminology

Throughout this document, you will refer to the following terms:

Term	Description
Cost Group Type	Maps to the various functions in an organization, for example, Development, HR, IT, and so on.
Cost Groups	Hierarchical structure to define your organization and distribute billing units.
Cloud Account	Credentials for logging in to a cloud provider.
Billing Units	Different entities depending on the cloud. These entities are account IDs in Amazon cloud, Project IDs in Google cloud, Subscription ID in AzureRMcloud, Datacenter name (prefixed with the cloud group) in vCenter clouds, Project ID in OpenStack cloud, and Namespace UID in Kubernetes cloud.
Budgets	Ability to allocate or reserve amounts per cloud or cost group type.
Tags	Key-value pairs associated with resources in a cloud.

Features

The new features inCost Optimizer 5.2.0 are:

- Perform recommended actions, such as, terminate, stop, or suspend multiple instancesin a go.
- Implement your cloud in CloudCenter Suite.
- Create reports between your chosen dates.
- Review and update historical prices for instances and storage types in private clouds.
- Enable resizing recommendations for VMware vCenter.
- Enhanced dashboard that displayscost and savings that can be achieved through recommendations.

Infrastructure

For setting up theCost Optimizer infrastructure, see*Suite Install 5.2.0 Home* > *Installation Approach* >Prepare Infrastructure.

Module Update Considerations

When updating the Cost Optimizer module, be aware that the update occurs for several minutes. During that time there may be a loss of connectivity between the CloudCenter Suite and individual cloud regions even after the Suite Admin UI indicates that the update has completed. Therefore, it is encouraged to keep this potential loss of connectivityin mind beforeapplying updates.

Logging In to Cost Optimizer

InSuite Admin Dashboard, click the**Cost Optimizer** card to open Cost Optimizer.

1. Enter the following:
    - **Email**
    - **Password**
    - **Tenant ID**of your organization
2. Click**Login**.

Cost Optimizer opens in the **Cost Optimizer Dashboard** page.

To log out, clickthe *Welcome <username>* text in the top-right corner and choose**Log Out**.

Related Information

See the following sections for detailed information about the Cost Optimizer features:

- Access and Roles
- Allocate Budgets
- Cost Groups Configuration
- Inventory
- Rightsizing
- Reserved Instances

Cost Optimizer Architecture

Cost Optimizer Architecture

- Deployment Architecture
- Basic Install Architecture
- Full Install Architecture
- Port Requirements
  - Without Cloud Remote
  - With Cloud Remote

Deployment Architecture

Cost optimizer is a module of CloudCenter Suite, that installs on a Kubernetes cluster througha Suite Chart. A suite chart is a common framework that allows the creating of tenants and users. %coiis deployed using the CloudCenter Shared and Cost Optimizer helm charts. The following diagram shows the Kubernetes deployment architecture inCost Optimizer.

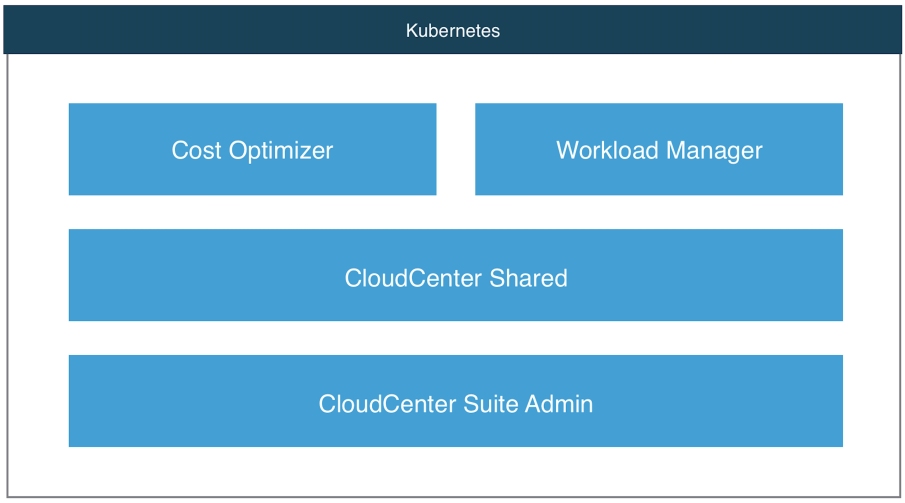

Cost Optimizer is a read-only module that connects to different cloud providers to collect information and use the collected information to generate recommendations. To act on a recommendation, to save costs, Workload Manager, another module in CloudCenter Suite is required. Workload Manager acts as an execution engine.

Basic Install Architecture

After installing Cost Optimizer from the Suite Admin, if your CloudCenter Suite Kubernetes cluster can receive connections from public internet addresses, you have everything that is required to use of Cost Optimizer's core features with VM-based public clouds. This includes collecting inventory, cost, metrics and generating recommendations. As mentioned above, it is recommended that you install the Workload Manager module to act on recommendations suggested by Cost Optimizer. The following diagram illustrates the basic install architecture for Cost Optimizer. Note that the icons indicatecompute, storage, database, load balancer, metrics, and invoice.

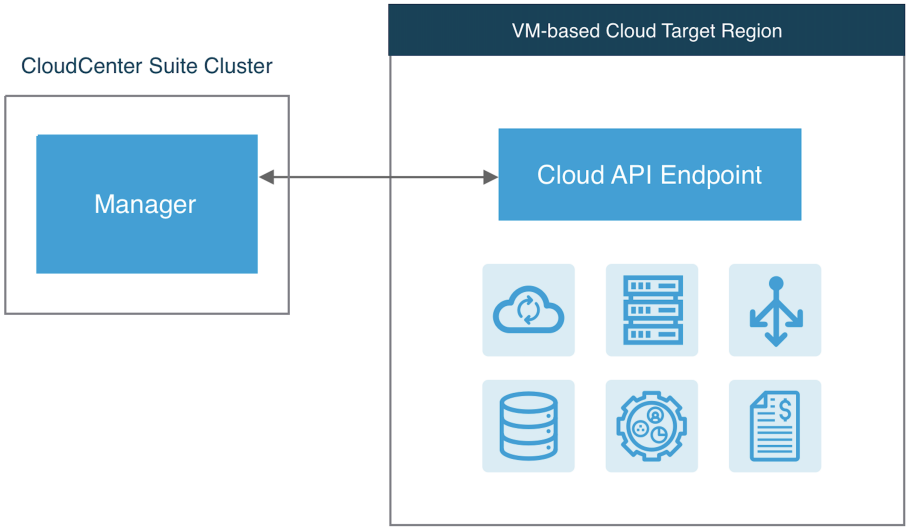

The**manager**component is the main component of Cost Optimizer.The basic install architecture installs the manager component, which is broken down into multiple microservices,running within pods in the CloudCenter Suite cluster. Some of these services are common framework services used by all CloudCenter Suite module. While some services are some are specific to Cost Optimizer, some services are shared between Workload Manager and Cost Optimizer. The manager communicates with the API endpoint of the target cloud region where your workloads will be launched. This communication is used to launch and control the VMs or pods running your workloads, and to extract data regarding cloud resource consumption. For Kubernetes target clouds, there are no worker VMs and the container-based workloads are controlled through the Kubernetes API.The basic install architecture relative to Kubernetes target clouds is summarized in the figure below.

Full Install Architecture

The basic install architecture has a limitation. The basic install architecture assumes that the manager and the target cloud regions can initiate connections to or receive connections from public internet addresses. If either of these cases is not true, or you want to restrict internet access for security reasons, you will need to install additional components to ensure full functionality of Cost Optimizer. For VM-based clouds, you will need to install**Cloud Remote**as an additional component. The full install architecture for VM-based cloud regions is as shown in the following diagram.

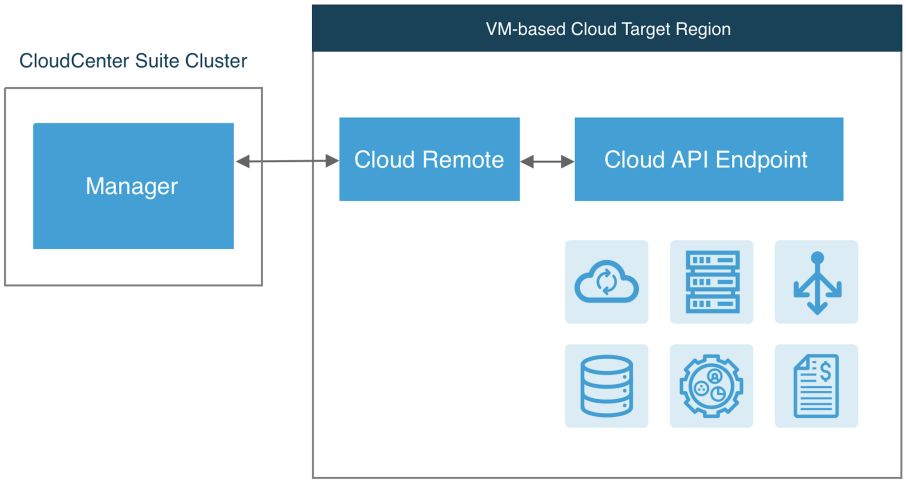

If you use Cloud Remote, you only access in one direction either from or to the CloudCenter Suite.The Cloud Remote handles communication in the other direction.

The Cloud Remote component is delivered as a virtual appliance that you import to your target VM-based cloud region. It is a CentOS 7 image that manages a collection of containerized services. Cloud Remote can be deployed as a single VM and later scaled to a cluster of VMs.

For VM-based cloud regions, Cloud Remote acts as a communication proxy between the manager and the cloud API endpoint (also used by Workload Manager).

If the manager cannot accept inbound connections from public addresses, you must install Cloud Remote in all VM-based target regions that are not within the same network as the manager.

The following is a full install architecture for Kubernetes target clouds, for which you mustinstall the Cloud Remote in an environment that is in the same network as the target Kubernetes cloud.

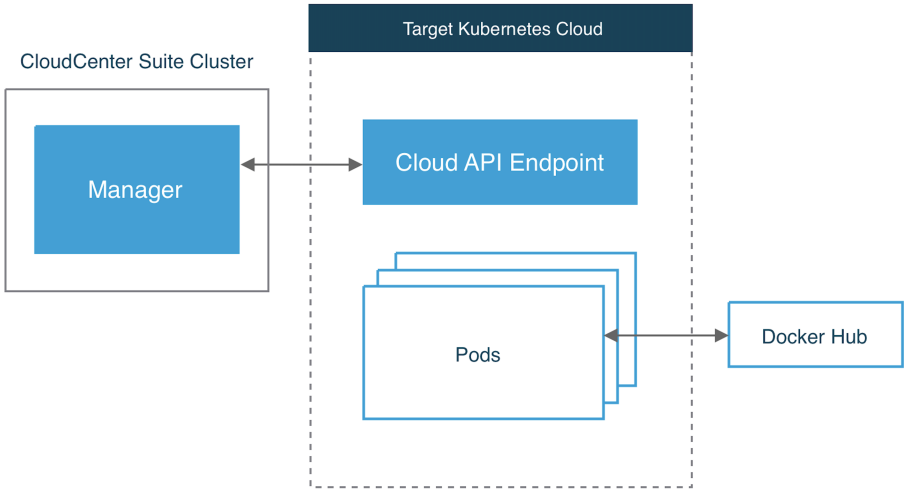

Port Requirements

Without Cloud Remote

The following image identifies the ports that must be open for Cost Optimizer.

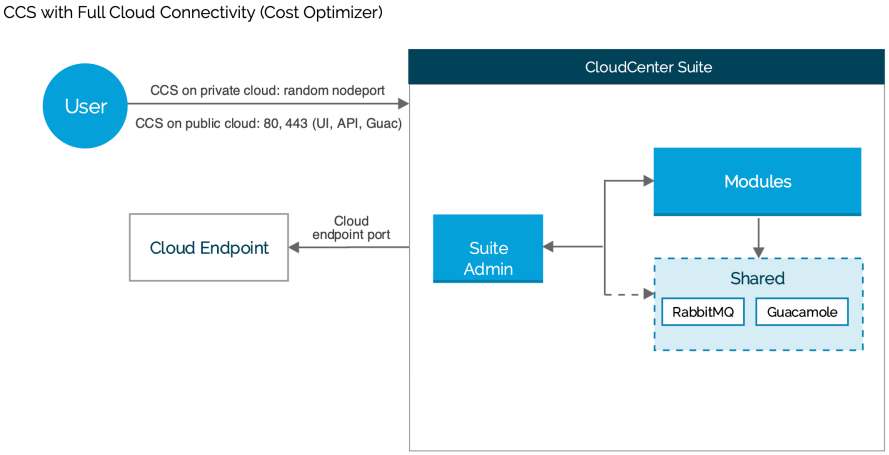

With Cloud Remote

The following image identifies the ports that must be open for Cost Optimizer when using the Cloud Remote component.

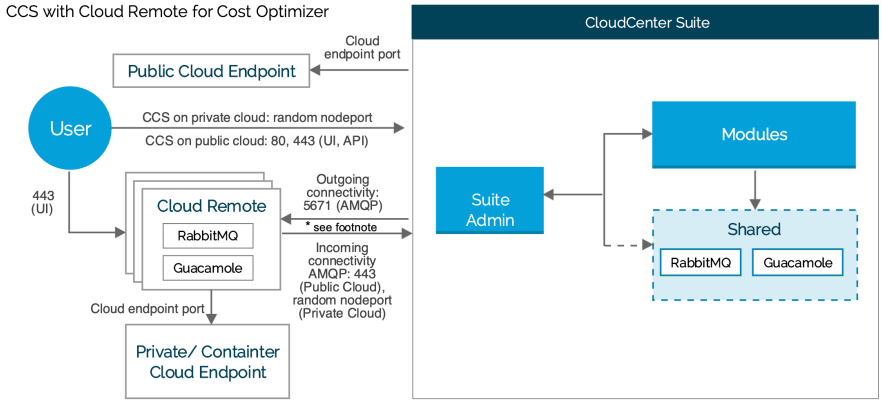

*Footnote

- Is CloudCenter Suite directly accessible from your Cloud Remote? =**YES**, the arrow from Cloud Remote to CloudCenter Suite is applicable
- Is CloudCenter Suite directly accessible from your Cloud Remote? =**NO**, the arrow from CloudCenter Suite to Cloud Remote applicable

Type NodePort: If you set the type field to NodePort, the Kubernetes control plane allocates a port from a range specified by service-node-port-range flag (default: 30000-32767). Refer tohttps://kubernetes.io/docs/concepts/services-networking/service/for additional context.

Access and Roles

Access and Roles

- Overview
- User Groups
- Roles
- Access Control Lists (ACLs)
- Personas

Overview

When you accessCost Optimizer you can see the cost, inventory, and recommendations reports and dashlets based on your group and role settings.

User Groups

A user must belong to at least one group to view resources authorized for that group.Cost Optimizer ships with the following user groups.

User Group	Description
Optimizer Admin	Root or module admin. Users belonging to this group have the ability to add budgets, view costs, inventory, recommendations for all billing units. Users do not need to be explicitly assigned to cost groups. Users are also permitted to perform administrative tasks like managing cloud accounts and settings in Cost Optimizer.
Optimizer User	Cost Groups must explicitly be shared with users belonging to this group, else users cannot see costs, inventory, or recommendations. Users assigned to this group can view data only pertaining to billing units associated with the cost groups. Users assigned to this group can only reallocate the budgets.
Financial Expert	Read-only users, who have view-only access to all data, regardless of cost group or billing unit association.

See:Create and Assign Groupsfor additional details.

Roles

Roles are a collection of privileges provided to users in a group. The users within each group can perform*permitted functions*on*permitted resources*by being part of the group.Roles are*only*associated withuser groups. Coupled with Access Control Lists (ACLs), roles offer the ability to perform specific tasks and view corresponding data.

Cost Optimizer ships with the following roles, which shares the samename as user groups.

- Optimizer Admin
- Optimizer User
- Financial Expert

See:Understand Rolesfor additional details.

Access Control Lists (ACLs)

While a role gives you visibility into a resource type, ACLs determine the users with who you share that resource. Using ACLs, a resource owner can share a specific resource directly with a user thereby allowinggranular privileges to individual resources. In Cost Optimizer, ACLs allow permitted users to share a resource with other users or groups by providing the following access levels to the usersthrough the **Share** dialog in Cost Groups Configuration.

Access Level	Description
View	User or group has read-only permissions but cannot modify or share this resource with others.
Manage	User or group can make changes as well as share this resource with others.

Personas

Based on the combination of user groups, roles, and ACLs, the following personas can be deduced for Cost Optimizer.

Persona	Maps to a Role or User Group in Cost Optimizer...	Function
Optimizer Administrator	Optimizer Admin	Access to every function in the module. An *Optimizer Administrator* can view data in *all* cost groups and types in a tenant. An *Optimizer Administrator* builds the organization hierarchy by creating cost groups types, cost groups, and assign billing units to one or more cost groups in the hierarchy. The Optimizer Administrator shares Cost Groups with **User A** by providing *Manage* access through ACLs. The Optimizer Administrator also manages tenant-level configuration parameters.
Cost Group Owner	Optimizer User	*Owner* of a Cost Group (for definition, see Cost Groups Configuration). A *Cost Group Owner* (**User A**) can redistribute billing units among the cost groups that the cost group owner can view and also share the cost group with others. However, **User A** *cannot* update or modify cost group hierarchies that an Optimizer Administrator has established.
Limited Viewer	Optimizer User	*View access* to one or more Cost Groups through an ACL. A *Limited Viewer* cannot share cost groups with other users nor reassign Billing Units. For example, **User B** may be granted the privilege to *view* cost, inventory reports, and recommendations within *Cost Group A*. **User B's** view is restricted based on Billing Unit associations to Cost Groups that **User B** can *view*.
Financial Expert	Financial Expert	Cannot make any changes to the system. Tenant-wide cost, inventory and recommendation views are displayed.

Email Settings in Cost Optimizer

EmailSettings in Cost Optimizer

Email settings are required to communicate withspecified users or user groups when the threshold limits are crossed as mentioned in the Alerts Pageand to send scheduled reports (see Cost Reports Overview > *Advanced Options* > *Scheduling Reports*).Cost Optimizer uses the SMTP settings specified in Suite Admin (see Email Settings).

The emails are sent from the default email addressnoreply.cloudcenter@cisco.com. You can specify an email of your choice by changing the value for the **email.from.address** in the email-config configmap, a sample of which is shown below.

```
# Please edit the object below. Lines beginning with a '#' will be ignored,
# and an empty file will abort the edit. If an error occurs while saving this file will be
# reopened with the relevant failures.
#
apiVersion: v1
data:
  email.from.address: noreply.cloudcenter@cisco.com
kind: ConfigMap
metadata:
  creationTimestamp: 2020-04-24T05:31:17Z
  labels:
    app: cloudcenter-task-scheduler-5.3.0
    chart: cloudcenter-task-scheduler-5.3.0
    heritage: Tiller
    purpose: configuration
    release: cloudcenter-shared
  name: email-config
  namespace: cisco
  resourceVersion: "118686337"
  selfLink: /api/v1/namespaces/cisco/configmaps/email-config
  uid: d21aef4c-85ec-11ea-bec6-42010af0019c
```

You must restart the *cloudcenter-shared-cloudcenter-task-scheduler* and *cost-optimizer-cloudcenter-alerts-manager* podsfor the changes to take effect.

UI Behavior

UI Behavior

- Icons
- Canceling without Saving

Icons

The following table identifies the Cost Optimizer icons.

Icon	Description
Actions	Perform action-oriented tasks on Cost Groups.
Add Cost Group	Add a Cost Group to a Cost Group Type.
Cost Group Type	Lists Cost Group Types (departments) set up in Cost Optimizer and adds Cost Group Type.
Date Range 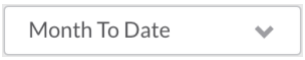	Choose a range to display the report.
Download	Downloads the report.
Dropdown list	Select resources from a list.
Filter	Filters information based on the selected category.
Module Navigation	Navigate back and forth between the module dashboards.
Move Billing Units or **Move Tags**	Move multiple billing units or tags to a resource.
Order	Sorts the listed items based on the latest or longest time period for the selected resource.
Schedule	Generates and sends a report at the specified date and time to the user or user group.

Search	Search resources based on the specified text for the allowed resources. ⚠ Not all fields and resources are searchable.
Select All	Select all items displayed on the page by clicking the checkbox in the table header or by clicking the checkbox against each item.
Switch	• The feature is disabled and configuration is unnecessary. • Enable the feature by turning it on and then inputting configuration values.
Toggle Chart	Toggles graphical report display between a line chart and a pie chart.
Unassigned Billing Units 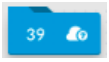	Unassigned cloud accounts with cloud resources.
Unassigned Tags	Unassigned tags associated with cloud resources.
Visibility Control	Visibility of default values can be toggled using this control.

Canceling without Saving

During configuration, you can cancel any changes or additions to a screen by clicking the**X**at the top right corner of the screen. This action takes you back to the original page that launched the screen.

The following screenshot shows how to cancel when assigning share access.

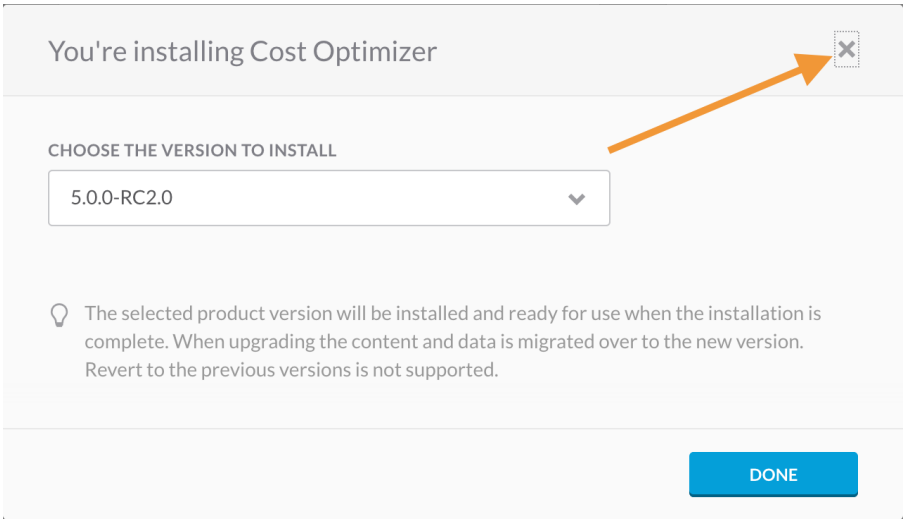

Configure Clouds

Configure Clouds

- Cloud Overview
- Configure Cloud End-to-End
- Cloud Remote
- Cloud Maintenance

Cloud Overview

Cloud Overview

- Overview
- Scope of a Cloud Region
- Minimum Permissions for Public Clouds

Overview

In CloudCenter Suite, the features to specify clouds are shared by Workload Manager and Cost Optimizer.

A cloud is an instance of one of the supported cloud types. A cloud has at least one region, but certain cloud types have multiple cloud regions.

Workload Manager and Cost Optimizer manage clouds on a per-region basis. The main point of control for a cloud region is the cloud region API endpoint. In the case of public VM-based clouds, such as AWS, GCP, and AzureRM, each cloud can have multiple regions that correspond to different geographic regions. OpenStack clouds also support multiple regions, but they are logical regions that do not have to be in different geographical areas.Kubernetes clouds and VMware vCenter clouds have only one region each.

A cloud must also have at least one cloud account associated with it. The cloud account information is needed to launch workloads, collect billing information, and in the case of VM-based clouds, list VMs associated with a particular cloud account that was launched outside of Workload Manager.

The workflow for specifying a cloud is as follows:

- Create the cloud: specify cloud name and cloud type
- For single-region cloud types (vCenter and Kubernetes): configure region details
- For multi-region cloud types: add a region,configure region details, repeat as necessary
- Add cloud accounts

If you are using Workload Manager, you will make your clouds available to users for deploying workloads using deployment environments.

Scope of a Cloud Region

For public clouds, a cloud region is associated with a geographic region defined by the cloud provider. For OpenStack clouds, a cloud region is a logical region defined within OpenStack. For VMware vCenter and vCDclouds, each instance of vCenter or vCD is considered a region. For Kubernetes clouds, each Kubernetes cluster is considered a region unto itself. The following table summarizes the scope of a region for each of the supported cloud types.

Cloud Family	Cloud Region Mapping	Supports any number of these per region
AWS	Geographical Region	• Accounts • Sub-Accounts • Identity and Access Management (IAM)
VMware vCenter	vCenter instance	• Datacenter • Clusters • Resource pools • Accounts • Datastores • Datastore clusters
VMware vCloud Director	vCD instance	• Datacenter • Clusters • Resource pools • Accounts • Datastores • Datastore clusters
Azure RM	Geographical Region	• Networks • Cloud services • Accounts

Google Cloud	Geographical Region	• Projects • Accounts
IBM Cloud	Geographical Region	• Accounts
OpenStack	Logical Region	• Tenants • Networks • Accounts
Kubernetes	Kubernetes cluster	• Accounts • Namespaces • VPCs • IAM policies
Outscale	Geographical Region	• Accounts • Sub-Accounts • Identity and Access Management (IAM)

Minimum Permissions for Public Clouds

The following table lists the minimum permissions for public cloud accounts supported in Cost Optimizer and Workload Manager modules of CloudCenter Suite Release 5.1.

> ⚠ You must enable AWS Cost Explorer to view AWS-specific costs on the Cost Optimizer dashboard. For additional details on enabling AWS Cost Explorer, seehttps://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/ce-enable.html.

Product	Function	AWS (IAM user)	Azure RM (Application)	Google (Service Account)
Cost Optimizer and Workload Manager	Discover billing units	iam:Get* iam:List*	*Cost management reader*	resourcemanager.projects.get,list
Cost Optimizer	Discover organization hierarchy	organizations:Describe* organizations:List*	N/A	billing.accounts.get,list orgpolicy.policy.get resourcemanager.folders.get,list resourcemanager.organizations.get
Cost Optimizer	Collect invoices	ce:* cur:Describe* ⚠ AWS Cost Explorer must be enabled to view AWS-specific costs on Cost Optimizer.	*Billing reader*	storage.objects.get,list storage.buckets.get,list

Cost Optimizer and Workload Manager	Collect VMs and volumes	ec2:DescribeAvailabilityZones ec2:DescribeAddresses ec2:DescribeInstances ec2:DescribeVolumes ec2:DescribeTags tag:getTagKeys tag:getTagValues ⚠ • The ec2:DescribeAvailabilityZones permission is mandatory and used for validating accounts. • The ec2:DescribeAddresses permission is optional and is used for Used to populated IP allocation type of NIC during inventory collection. • The ec2:DescribeTags permission is mandatory and used for discovering tags of PassService (ELB). • The tags permissions arerequired for tag-based reporting andonly applicable to Cost Optimizer.	VM: *VM contributor* Volume: *Reader* ⚠ The *Reader* role must be offered because no built-in role is provided.	compute.instances.get, list compute.disks.get,list
Cost Optimizer	Collect PAAS services	rds:Describe* elasticloadbalancing:Describe*	SQL Server and SQL database: *SQL Server contributor* MySQL and PostgreSQL Server: *Reader* ⚠ The *Reader* role must be offered because no built-in role is provided.	cloudsql.databases.get,list cloudsql.instances.get, list compute.forwardingRules.get, list compute.targetPools.get,list
Cost Optimizer and Workload Manager	Collect VM metrics	cloudwatch:Describe* cloudwatch:Get* cloudwatch:List*	*Monitoring reader* or *virtual machine contributor*	monitoring.metricsDescriptors.get, list monitoring.timeSeries.list
Cost Optimizer	Collect resource usage	s3:Get* s3:List*	N/A	N/A
Cost Optimizer	Collect RI subscriptions	ec2:DescribeReservedInstances*	N/A	N/A

<table>
<tr>
<td>Cost Optimizer and Workload Manager</td>
<td>Collect data for AWS member account</td>
<td>To allow a primary account to collect data on behalf of member accounts, the following is necessary:

• A primary account must be permitted to assume the role of a member account
• A member account must establish trust with the primary account

You must associate the following permission with the primary account's IAM user, as shown below:

<pre>{
 "Version": "2012-10-17",
 "Statement": [
  {
  "Effect": "Allow",
  "Action": [
  "sts:assumerole"
  ],
  "Resource": "*"
  }
 ]
}</pre>
Effective from Cost Optimizer 5.4, the AWS OrganizationAccountAccessRole is supported on member accounts.

Alternatively, you can create a role named Optimizer and do the following to the new role:

• Associate permissions listed above to collect invoices, inventory, metrics
• Add a trust relationship to the primary account

<pre>{
 "Version": "2012-10-17",
 "Statement": [
  {
  "Effect": "Allow",
  "Principal": {
  "AWS": "arn:aws:iam::
  <primary-account-number>:root"
  },
  "Action": "sts:AssumeRole",
  "Condition": {}
  }
 ]
}</pre></td>
<td>N/A</td>
<td>N/A</td>
</tr>
<tr>
<td>Workload Manager</td>
<td>Manage VMs and volumes</td>
<td>ec2:AssignPrivateIpAddresses

ec2:AttachNetworkInterface

ec2:AttachVolume

ec2:AuthorizeSecurityGroupEgress

ec2:AuthorizeSecurityGroupIngress

ec2:CreateImage

ec2:CreateKeyPair

ec2:CreateNetworkInterface

ec2:CreateSecurityGroup

ec2:CreateSnapshot

ec2:CreateTags

ec2:CreateVolume

ec2:DeleteKeyPair

ec2:DeleteNetworkInterface

ec2:DeleteSecurityGroup

ec2:DeleteSnapshot

ec2:DeleteTags

ec2:DeleteVolume

ec2:DescribeAccountAttributes</td>
<td>Offer the italicized roles to create, modify, or delete:

• NICs, Public IPs and security group: Network Contributor
• Diagnostics: Storage Account Contributor
• Unmanaged data disk: Storage Account Contributor
• Managed data disks: Owner
• VMs with managed data disks: Owner
• VMs with unmanaged data disks and diagnostic logs: Virtual Machine Contributor, Network Contributor, and Storage Account Contributor
• VMs with no data disks: Virtual Machine Contributor and Networ k Contributor

⚠ In some cases, the Owner role must be offered because no built-in role is provided.</td>
<td>Use the pre-defined Pr oject Editor role,

OR

compute.addresses. create,delete,get,list, use

compute.disks.create, delete,get,list,update, use

compute.firewalls. create,delete,get,list, update

compute.instances.*

compute. machineTypes.get

compute.neworks.get, list,use

compute.projects.get

compute.regions.get

compute.subnetworks. get,list,use, useExternalIp

compute.zones.get

iam.serviceaccounts. get,list</td>
</tr>
</table>

		ec2:DescribeAvailabilityZones ec2:DescribeDhcpOptions ec2:DescribeImageAttribute ec2:DescribeImages ec2:DescribeInstanceAttribute ec2:DescribeInstances ec2:DescribeInstanceStatus ec2:DescribeKeyPairs ec2:DescribeNetworkInterfaceAttribute ec2:DescribeNetworkInterfaces ec2:DescribeRegions ec2:DescribeSecurityGroups ec2:DescribeSnapshotAttribute ec2:DescribeSnapshots ec2:DescribeStaleSecurityGroups ec2:DescribeSubnets ec2:DescribeTags ec2:DescribeVolumeAttribute ec2:DescribeVolumes ec2:DescribeVolumesModifications ec2:DescribeVolumeStatus ec2:DescribeVpcAttribute ec2:DescribeVpcs ec2:DetachNetworkInterface ec2:DetachVolume ec2:EnableVolumeIO ec2:GetConsoleOutput ec2:GetConsoleScreenshot ec2:GetPasswordData ec2:ImportKeyPair ec2:ImportVolume ec2:ModifyImageAttribute ec2:ModifyInstanceAttribute ec2:ModifyNetworkInterfaceAttribute ec2:ModifyVolume ec2:ModifyVolumeAttribute ec2:RebootInstances ec2:RevokeSecurityGroupEgress ec2:RevokeSecurityGroupIngress ec2:RunInstances ec2:StartInstances ec2:StopInstances ec2:TerminateInstances ec2:UnassignPrivateIpAddresses		

Configure Cloud End-to-End

Configure Cloud End-to-End

- Configure a Google Cloud
- Configure a Kubernetes Cloud
- Configure a vCD Cloud
- Configure a vCenter Cloud
- Configure an AWS Cloud
- Configure an AzureRM Cloud
- Configure an IBM Cloud
- Configure an OpenStack Cloud
- Configure an Outscale Cloud

Configure a Google Cloud

Configure a Google Cloud

Configuring a Google cloud is a four-step process:

- Add a Google Cloud
- Add a GoogleRegion
- Configure a GoogleRegion
- Add a Google Cloud Account

Add a Google Cloud

To add a Google cloud follow these steps.

1. Navigate to**Admin > Clouds**. This brings you to the Clouds page. If you, or another tenant admin in your tenant, have already added clouds to your tenant, they will be listed here.
2. Click the**Add Cloud** linkin the upper right. The Add Cloud dialog box is displayed.
3. Enter the**Cloud Name**and select the**cloud provider**.

   When assigning the Cloud Name, use a unique name. Ensure that you do not assign a previously-used name (including deleted cloud names) as the Cloud Name as they may cause deployments to fail intermittently.

4. Click**Next**. The second page of the **Add Clouds** dialog box, **Connectivity Settings**, appears. Set the toggle to configure the Cloud Connectivity settings.
   - When adding a public VM cloud in the CloudCenter Suite UI, the Cloud Connectivity Settings page, the second page of the Add Cloud dialog box, appears with a single toggle displayed:**Worker VMs Directly Connect with CloudCenter Suite**.
   - Setting this toggle to No implies you will install Cloud Remote for each region of this cloud. This also causes a second toggle to appear:**CloudCenter Suite Directly Accessible from Cloud Remote**.
   - Follow the table below for guidance on setting these toggles.

Toggle settings	Use case	Diagram
Worker VMs Directly Connect with CloudCenter Suite = Yes	Unimpeded connectivity exists between the CloudCenter Suite cluster and the cloud region API endpoint AND Unimpeded connectivity exists between the CloudCenter Suite cluster and worker VMs Cloud Remote is not required	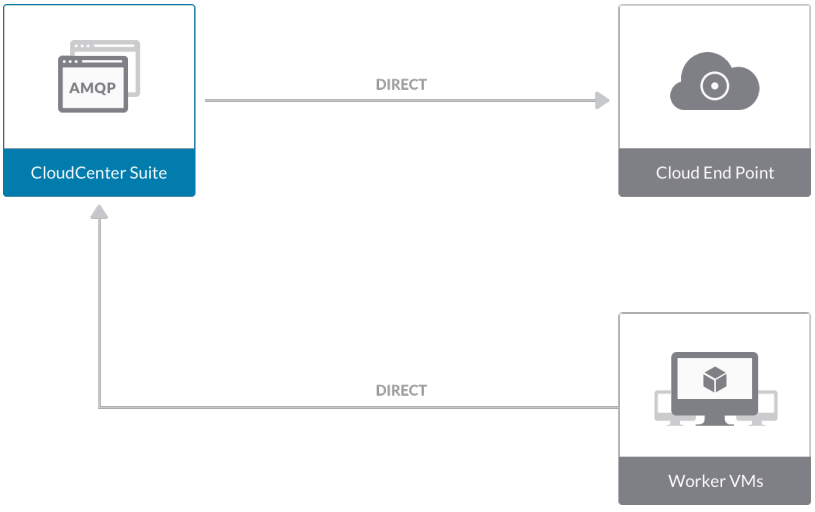

Worker VMs Directly Connect with CloudCenter Suite = No AND CloudCenter Suite Directly Accessible from Cloud Remote = Yes	Worker VMs need to communicate to the CloudCenter Suite cluster through Cloud Remote AND Cloud Remote can initiate the connection to the CloudCenter Suite cluster	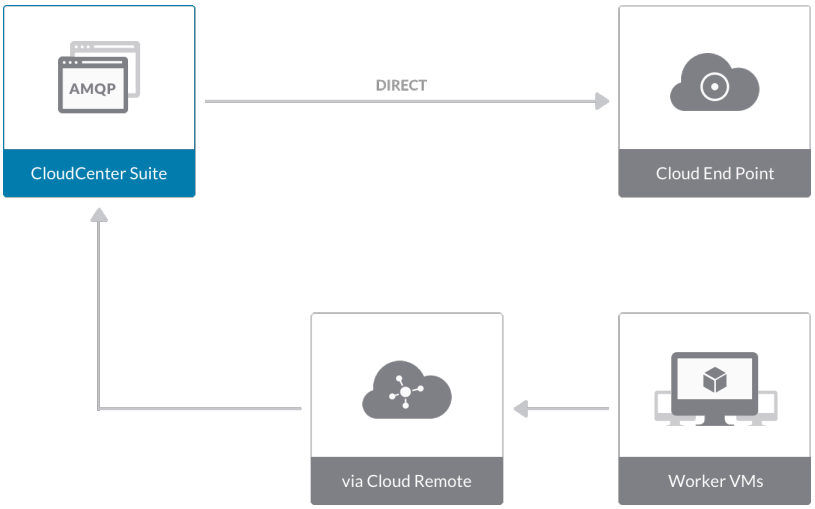
Worker VMs Directly Connect withCloudCenter Suite = No AND CloudCenter SuiteDirectly Accessible from Cloud Remote = No	Worker VMs need to communicate to the CloudCenter Suite cluster through Cloud Remote AND the CloudCenter Suite cluster cannot receive a connection initiated by Cloud Remote	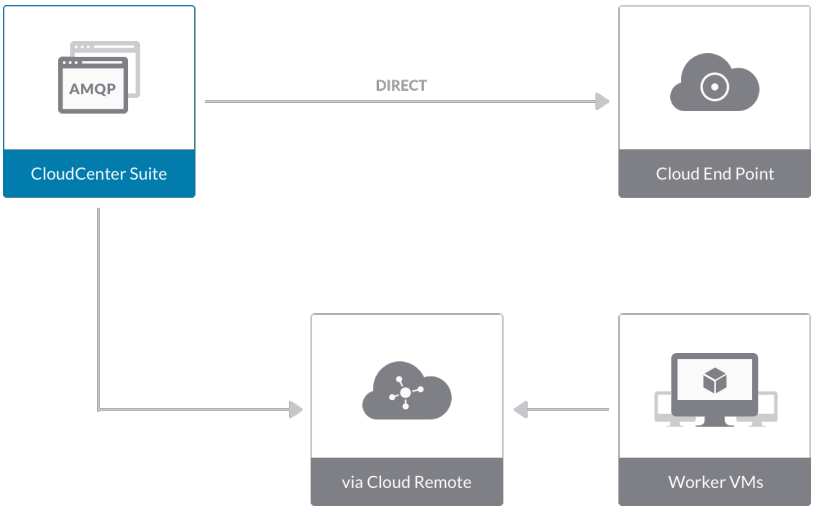

Note

The connectivity toggle settings set at the cloud level are inherited by each region you add to this cloud. However, it is possible to override these toggle settings on a per-region basis from the Regions tab for each region.

5. Click**Done**to save the configuration and close the dialog box. This brings you back to the Clouds page and the cloud you just created will be added to the bottom of the list on the left side of the page.

Add a GoogleRegion

After creatinga Google cloud, the next step is to create the first region for the cloud. Follow these steps.

1. Navigate to the Clouds page and select the cloud you created on the left side of the screen. Then click the**Add Region**button on the right side of the screen.
2. After clicking the Add Region button, the Add Region dialog box is displayed. Select a region from the list and click **Save.**
3. After clicking**Save**you are brought back to the Clouds page with theregion you added shown on the right side of the page.

Configure a GoogleRegion

To configure a region you added to your Google cloud, follow this procedure.

1. Navigate to Clouds page: **Admin > Clouds**. Find your Google cloud from the cloud list on the left half of the screen and click its **Configure Cloud** link. This displays the Regions tab for this cloud as shown in the figure below with theCloud Settings section displayed first.

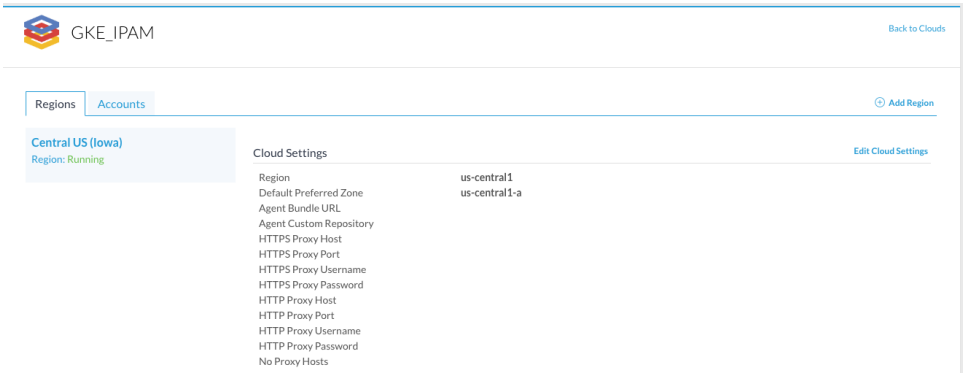

After you have added multiple regions to your Google cloud, the Regions tab will show multiple individual region tabs on the left side of the screen. Click the tab of the region you want to configure.

2. Click the **Edit Cloud Settings**link in the upper right of the Cloud Settings section. This displays the **Configure Cloud Settings** dialog box. The Cloud Settings section contains fields that are unique to Google and settings that are common to all cloud providers. Adjust these field values per the instructions in the following tables:

Google Specific Cloud Settings

Field	Usage
Region	This field is set by CloudCenter Suite based on the region location you selected from the Add Region dialog box.
Default Preferred Zone	This field is set by CloudCenter Suite based on the region location you selected from the Add Region dialog box.

Agnostic Cloud Settings

Field	Usage
Exclude these special characters for Windows password	When the Workload Manager agent is installed on a Windows worker VM, a special user account, called cliqruser, is created to support RDP sessions that may be initiated by the user through the Workload Manager UI. A Workload Manager process running on the CloudCenter Suite cluster creates a random password and passes it to the agent for creating the cliqruser account. Because some Windows deployments may restrict using certain characters for Windows passwords, this field is provided to tell the Workload Manager to exclude these special characters in the generation of the password for the cliqruser account.
Agent Bundle URL	If you plan to use a local repository to host the bundle store, you need to enter the URL of the local bundle store here. Otherwise, leave blank.
Agent Custom Repository	If you plan to use alocal repository to host the package store, you need to enter the URL of the local package store here. Otherwise, leave blank.
HTTP /HTTPS proxy fields (host, username, password)	If you require VMs in your region to access public addresses through a web proxy, enter the URL and credentials of the HTTP and HTTPS proxy servers in these fields.
No Proxy Hosts	If you have specified an HTTP or HTTP proxy using the above fields, you can specify that managed VMs in the region should bypass the proxy and connect directly to certain hosts. Use this field to create a comma-separated list of IP addresses or URLs that should be accessed directly. This field is ignored if an HTTP or HTTPS proxy is not specified.

Important information on proxy settings

In CloudCenter Suite, you can specify proxy settings at the region level, as described here, and at the suite level. To understand the expected behavior when proxy settings are specified at both levels, see the subsequent *Precedence of Proxy Settings* section.

Download Configuration and Encryption Key

After saving the Region Configuration settings in the Workload Manager or Cost Optimizer UI, you can download them to your local computer and then upload them to other conditional components such as Cloud Remote.

The Configuration and Encryption key is only visible when you have configured the Cloud Remote component. Click the**Download Configuration**l ink in the upper right of the Region Connectivity section, as shown in the following screenshot.

Clicking**DownloadConfiguration**causestwo things to happen:

- An encrypted zip file namedartifacts.zipis downloaded by your browser. Make a note of the location of this zip file as you will need if you are using Cloud Remote.
- The Region Connectivity section header updates to display a**Copy Encryption Key**link, as shown in the following screenshot.

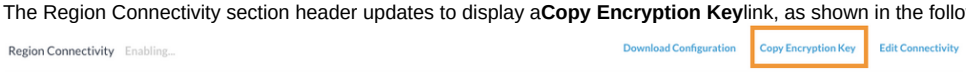

Click the**Copy Encryption Key**link to save the key to your clipboard. A success message will be displayed temporarily above the Region Connectivity section header. Make sure not to overwrite the clipboard with other data. You will need the key when you upload the configuration zip file to conditional components like Cloud Remote.

If you change the connectivity settings in the CloudCenter Suite UI and need to again download the zip file fromsoftware.cisco.com, use the automatically create a (new) encryption key, and copy the key to the clipboard by clicking the**Copy Encryption Key**link again.

Precedence of Proxy Settings

In CloudCenter Suite, you can specify HTTP and HTTPS proxy settings at the suite level as described inProxy Settings, and at the region level as described in the preceding*Agnostic Cloud Settings*section. The CloudCenter Suite cluster, Cloud Remote, and worker VMs will each use either the suite-level proxy settings or the region level proxy settings based on the conditions described below.

For Cloud Remote mode, whenever you change the region proxy settings, ensure you again download the region connectivity setting from the Workload Manager UI and reapply to Cloud Remote.

Mode	**Proxy Settings Used By**		
Mode (Is Cloud Remote used to communicate withcloud (or APIC) endpoint?)	CloudCenter Suite cluster to communicate with cloud (or APIC) endpoint	Cloud Remote to communicate with cloud endpoint	Worker VM to communicate with bundle store and package store
Cloud Remote Mode	NA	Region-level	Region-level
Non-Cloud Remote Mode	Suite-level	NA	Region-level

Not all clouds support all the proxy settings, and not all clouds support bothCloud Remote and non-Cloud Remote modes.Cloud Remote mode is not applicable to Google. In non-Cloud Remote mode, Google supports all proxy settings, except HTTP/HTTPS Username/Password, and proxy hosts.

When you are done editing the settings in the dialog box, click**Save**.

3. Determine if you need Cloud Remote for this region. Scroll down to the Region Connectivity section for the region and click on the**Configure Region**link in the upper right to open the Configure Region dialog box.The toggle settings should be the same as when you set them on the connectivity page of the Add Cloud dialog box.If all of the connectivity togglesinthe Region Connectivity dialog boxare set to Yes, then Cloud Remote is NOT needed for this cloud region. In this case,you would normally leave all region connectivity settings at their current values and continue to the next settings section.

   The exceptionto this guidance is when a NAT firewall or proxy server exists between theCloudCenter Suite management cluster andworker VMs, or between the CloudCenter Suite management cluster and users that would use Workload Manager to initiate a Guacamole remote connection to a worker VM. In either of these cases, override the address fields in the Region Connectivity dialog box as explained below.

Networking Constraint	**Field**	**Value**
Worker VMs must use a proxy server or NAT firewall to access the"local" AMQP server running in the CloudCenter Suite cluster.	Worker AMQP IP Address	IP address and port number that the firewall or proxy server presents to the worker VMs on behalf of the "local" AMQP server running in the CloudCenter Suite cluster.
Users must use a proxy server or NAT firewall to access theGuacamole server running in the CloudCenter Suite cluster.	Guacamole Public IP Address and Port	IP address and port number that the firewall or proxy server presents to users on behalf of theGuacamole server running in the CloudCenter Suite cluster.
Worker VMs must use a proxy server or NAT firewall to access theGuacamole server running in the CloudCenter Suite cluster.	Guacamole IP Address and Port for Application VMs	IP address and port number that the firewall or proxy server presents to the worker VMs on behalf of the Guacamole server running in the CloudCenter Suite cluster.

   Click**OK**to save the changes and dismiss the dialog box. You can now proceed to the next region settings section:VM Naming and IPAM Strategy.

4. If any of theconnectivity toggles inthe Region Connectivity dialog box are set to No, then **you must install and configure Cloud Remote for this region**.

Configure Cloud Remote in a Google Region

Configure Cloud Remote in a Google region as follows.

Obtain and Launch the Cloud Remote Appliance in Google

a. Request the Cloud Remote shared VMI form Cisco support by opening a CloudCenter Support case. In your request, specify the following details:

   i. Your GCP account number
   ii. Your GCP project ID number
   iii. Your CloudCenter Suite version
   iv. Your Customer ID (CID)
   v. Your customer name
   vi. Specify if your setup is in production or for a POC
   vii. Your Contact Email

b. After you open a case, your support case is updated with the shared VMI ID.**Proceed to the next step only after your supportcase is updated with the VMI ID.**

c. Navigate to the GCP dashboard and search for the VMI ID name provided in theCloudCenter Support case in the list of images for your project.

d. Launch an instance using the shared VMI.

   i. Click on the image name. This takes you to the page for the image

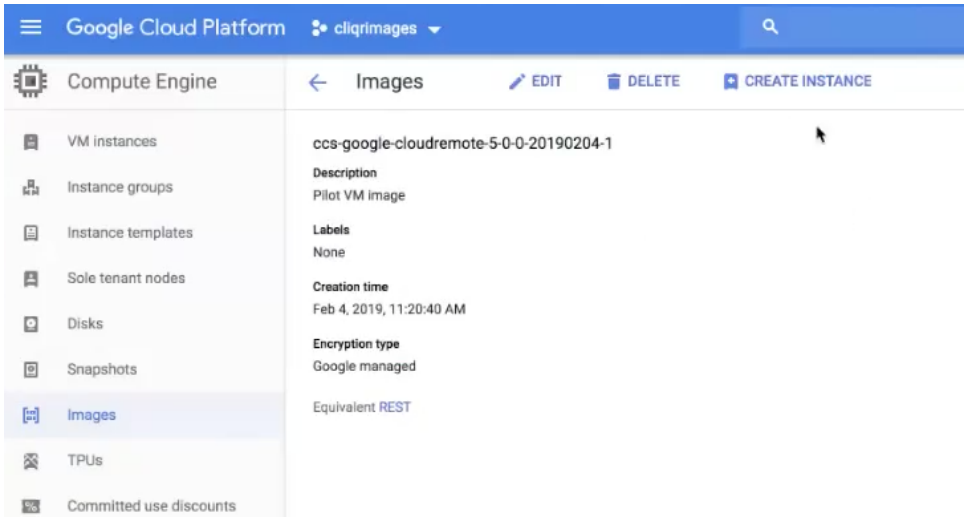

   ii. Click on Create Instance to display the Instance properties page

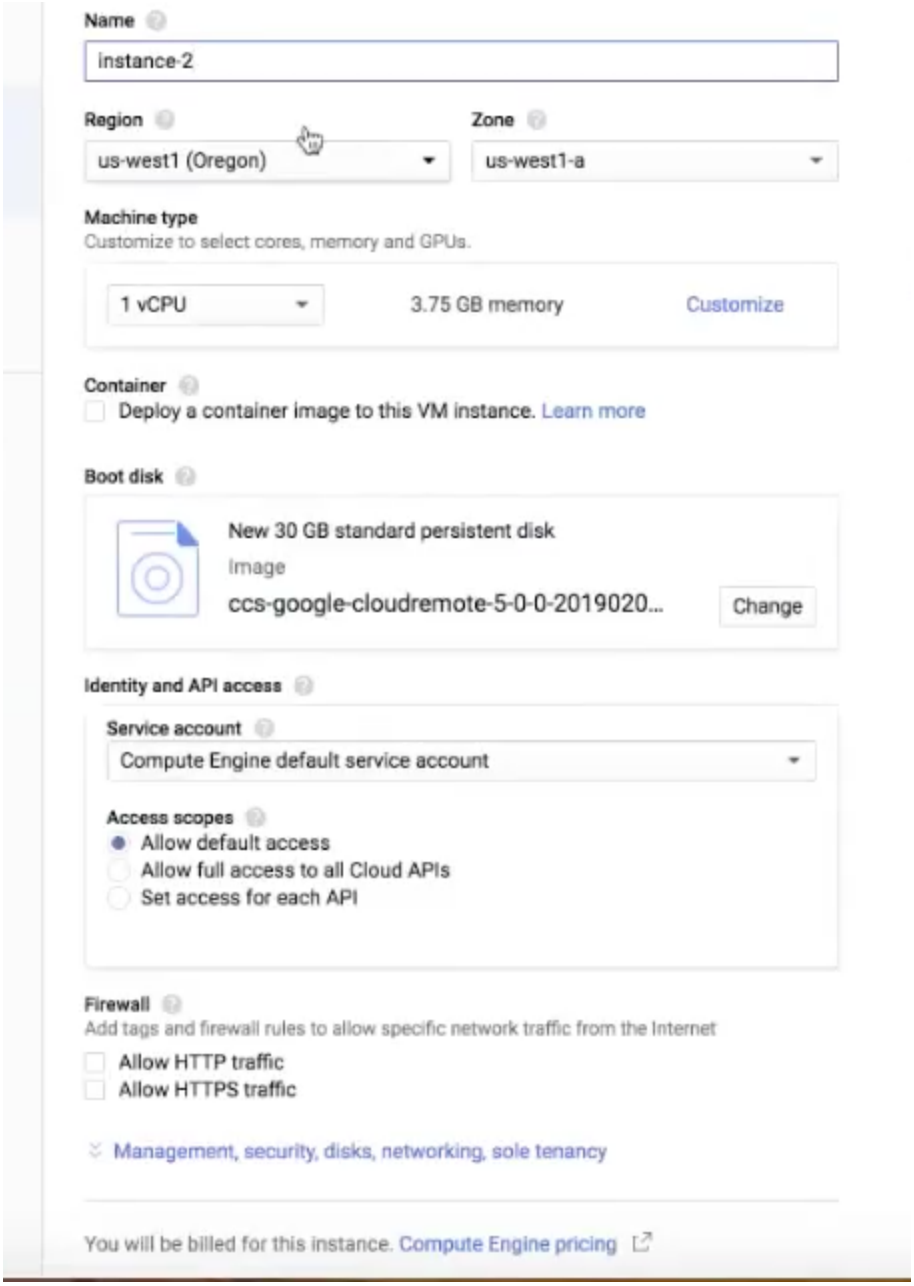

iii. Complete these fields:

1. Instance name
2. Region and zone
3. Machine type: select 2 vCPU, 7.5 GB RAM
4. Click the checkbox to allow HTTPS access
5. Click the Security tab (under the Allow HTTPS traffic checkbox). In the SSH key field, add your organization's public ssh key followed by a space and then the username you want to use to login to the Cloud Remote appliance. Click the Add Item button when done.

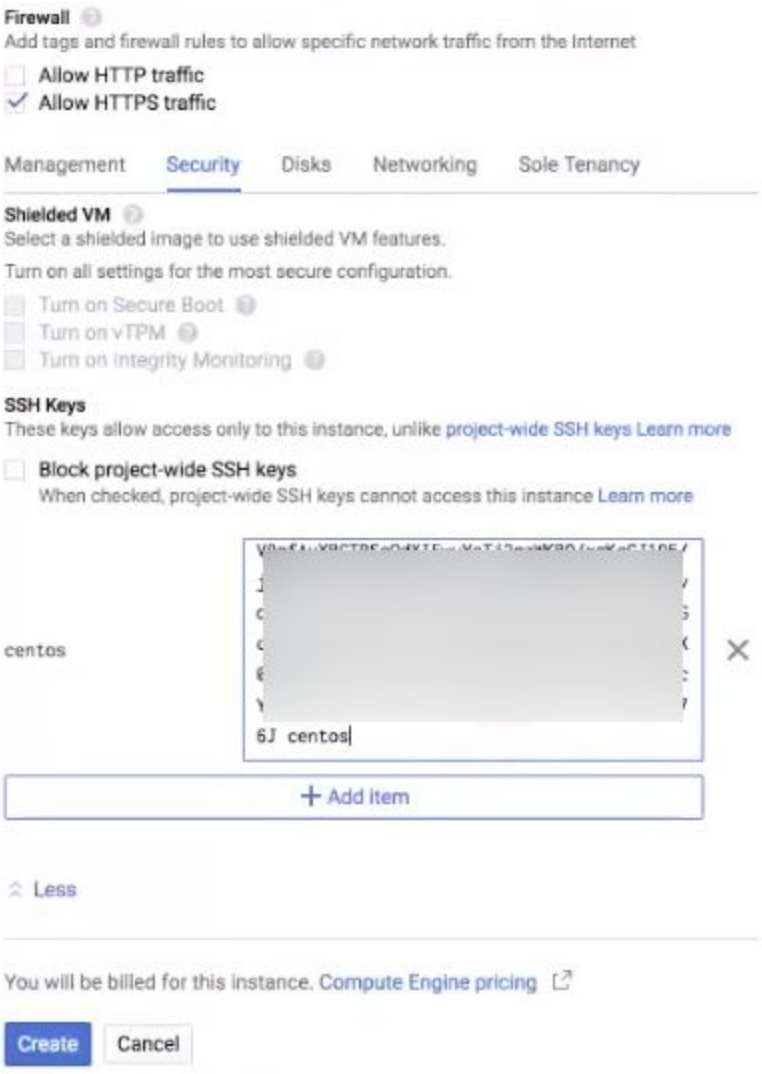

iv. Click Create to launch the instance.

e. Optional but recommended for production environments: Deploy two additional instances of the appliance to form a cluster for HA. Cloud Remote includes support for the clustering of multiple nodes. You will "add" these two additional instances to the first instance after the first instance is configured. SeeCloud Remote (Conditional)> *Scaling*for details.

f. Once the first instance of the appliance has been launched, use theGCP console to**note its IP public and private addresses**. You will need this information later on in order to login to the Cloud Remote web UI and to complete the Region Connectivity settings in the CloudCenter Suite Web UI. Also, note the IP addresses of any other appliances you launch.

Unable to render {include} The included page could not be found.

Specify AMQP and Guacamole Addresses for Supporting Cloud Remote

From the CloudCenter Suite UI, for the cloud region requiring Cloud Remote, navigate to the corresponding Regions or Details tab. Click the**C onfigure Region**link in the upper left of the Region Connectivity section to bring up the Configure Region dialog box. The toggle settings should be the same as when you set them on the connectivity page of the Add Cloud dialog box. You must update some of the address fields in the dialog box according to the scenarios summarized in the table below.

Toggle Settings	Field	Value
Worker VMs Directly Connect with CloudCenter = No AND CloudCenter Directly Accessible from Cloud Remote = Yes	Local AMQP IP Address	Pre-populated with the address and port number of the "local" AMQP server running in the CloudCenter Suite cluster. This address must be **accessible to Cloud Remote**. If Cloud Remote is accessing the CloudCenter Suite cluster through a user-supplied proxy server or NAT firewall, overwrite this field with the corresponding local AMQP IP address and port number provided by the user-supplied proxy server or NAT firewall and accessible to Cloud Remote.

Worker VMs Directly Connect with CloudCenter = No AND CloudCenter Directly Accessible from Cloud Remote = No	Remote AMQP IP Address	Enter <Cloud_Remote_IP>:<amqp_port>, where <Cloud_Remote_IP> = the IP address Cloud Remote which is **accessible to the CloudCenter Suite cluster**, and <amqp_port> = 5671 OR the custom AMQP port number you would later set with the **Change Ports shell script** on the Cloud Remote appliance (see Cloud Remote (Conditional) > *Custom Port Numbers (Conditional)*). If there is no user-supplied NAT firewall or proxy server between the CloudCenter Suite cluster and Cloud Remote, the IP address would be the public IP address of Cloud Remote. If there is a NAT firewall or proxy server between the CloudCenter Suite cluster and Cloud Remote, instead, enter the corresponding public IP address and port number that the firewall or proxy server presents to the internet on behalf of the "remote" AMQP server running in Cloud Remote.
Worker VMs Directly Connect with CloudCenter = No	Worker AMQP IP Address	Enter <Cloud_Remote_IP>:<amqp_port>, where <Cloud_Remote_IP> = the Cloud Remote IP address **accessible to the worker VMs**, and <amqp_port> = 5671 OR the custom AMQP port number you would later set with the **Change Ports shell script** on the Cloud Remote appliance (see Cloud Remote (Conditional) > *Custom Port Numbers (Conditional)*).
Worker VMs Directly Connect with CloudCenter = No	Guacamole Public IP and Port	Enter <Cloud_Remote_IP>:<guac_port>, where <Cloud_Remote_IP> = the Cloud Remote IP address **accessible to CloudCenter Suite users**, and <guac_port> = 8443 OR the custom Guacamole port number you would later set with the **Change Ports shell script** on the Cloud Remote appliance (see Cloud Remote (Conditional) > *Custom Port Numbers (Conditional)*).
Worker VMs Directly Connect with CloudCenter = No	Guacamole IP Address and Port for Application VMs	Enter <Cloud_Remote_IP>:<guac_port>, where <Cloud_Remote_IP> = the Cloud Remote IP address **accessible to worker VMs**, and <guac_port> = 7789

When done, click **OK** to save the setting and dismiss the dialog box.

Download Region Connectivity Settings and Upload to Cloud Remote

Cloud Remote uses the region connectivity settings set in the Workload Manager or Cost Optimizer UI. After saving the Region Configuration settings in the Workload Manager or Cost Optimizer UI, you must download them and to your local computer and then upload them to Cloud Remote as follows.

Click the**Download Configuration**link in the upper right of the Region Connectivity section, as shown in the figure below.

Clicking DownloadConfiguration causestwo things to happen:

- An encrypted zip file named**artifacts.zip**will be downloaded by your browser. Make note of the location of this zip file as you will need to upload it to Cloud Remote through the Cloud Remote web UI (see below).
- The Region Connectivity section header updates to display a **Copy Encryption Key**link, as shown in the figure below.

Click the**Copy Encryption Key**link to save the key to your clipboard. A success message will be displayed temporarily above the Region Connectivity section header. Make sure not to overwrite the clipboard with other data. You will need the key when you upload the configuration zip file to Cloud Remote.

> ⚠ If you change the connectivity settings in the CloudCenter Suite UI and need to again download the zip file, a new encryption key is automatically created and can be copied to the clipboard by clicking the **Copy Encryption Key** link again.

After you have downloaded the zip file and copied the encryption key to your clipboard, login toCloud Remote web UI.

a. Open another browser tab and login to https://<Cloud Remote_ip> with thedefault credentials: admin/cisco.
b. You will immediately be required to change your password. Do so now.
c. You are now brought to the Cloud Remote home page as shown in the figure below.

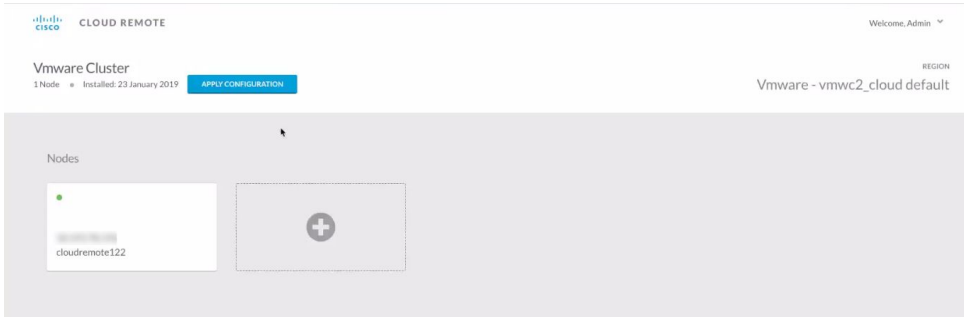

d. Click the **Apply Configuration** button in the page header. This prompts you to select a configuration file and enter the encryption key as shown in the figure below.

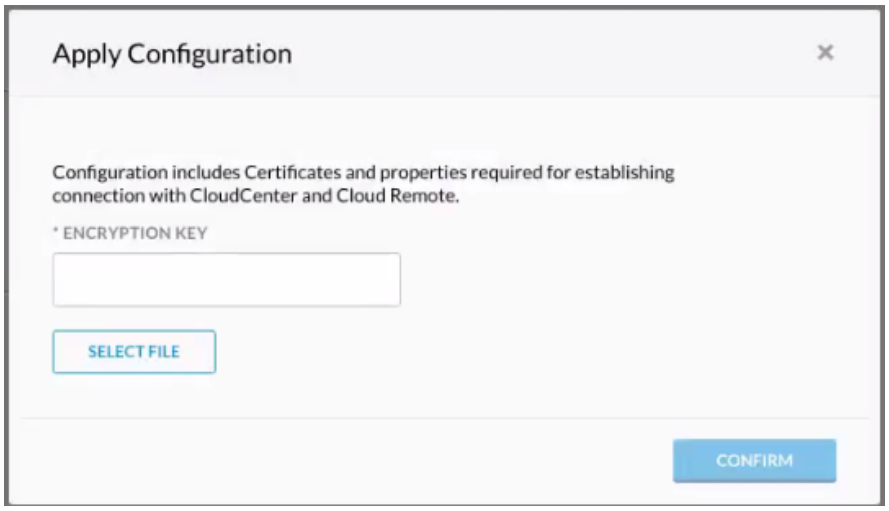

e. Paste the encryption key that was copied to the clipboard into the**Encryption Key**field in the dialog box.
f. Click**Select File**and browse tothe artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
g. Click**Confirm**.
h. Once the zip file is successfully transmitted and accepted, the Cloud Remote appliance attempts to establish communication with the CloudCenter Suite cluster and the Cloud Remote web UI home page is updated to show the name of the region it is connecting to in the upper right (see figure below).

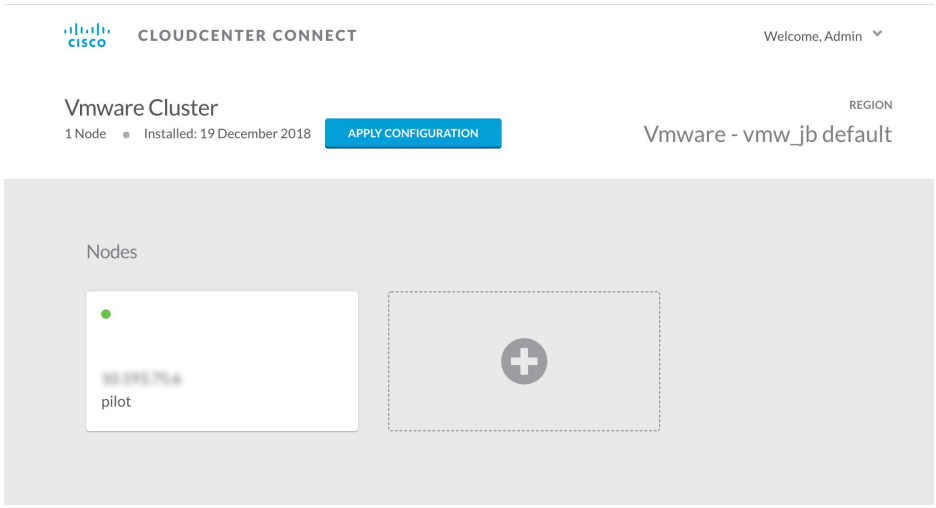

Switch your focus back to the Region Connectivity section of the target cloud region in the CloudCenter Suite web UI. The status indicator in the Region Connectivity section header will change from Not Configured to Running once connectivity between Cloud Remote and the CloudCenter Suite cluster is completely established (see figure below).

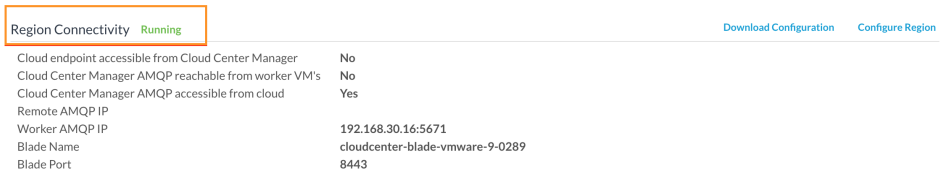

After completing these steps, Workload Manager and Cost Optimizer can use Cloud Remote for communicating with the target cloud region.

5. VM Naming and IPAM Strategy (conditional): Configure any VM naming or IPAM strategies in the Strategy section. If you leave the settings at the defaults, no IPAM strategy is applied and the default VM naming strategy is applied.

VM Naming and IPAM Strategies

Workload Manager lets you define a VM naming strategy and an IP address management (IPAM) strategy on a per cloud region basis. The VM naming strategy lets Workload Manager generate unique names for each of your VMs as they are launched. The IPAM strategy letsWorkl oad Manager generate IP addresses for each of the vNICs on each of your VMs. The VM naming strategy can reference a VM naming callout script, and the IPAM strategy requires that you specify an IP address allocation callout script andan IP address deallocation callout script. An IP address allocation callout scriptmay work in conjunction with a 3rd party IPAM server in your cloud environment.

Like External Lifecycle Actions, the scripts or commands associated with a VM naming strategy or IPAM strategy are executed in a containerized script execution engine running in the CloudCenter Suite cluster, if you are not using Cloud Remote in your region. If you are using Cloud Remote, these scripts or commands will execute in a containerized script execution engine running in Cloud Remote.

To set up the VM naming and IPAM strategies for your region, scroll down to the Strategy section of the Regions or Details tab and follow these steps.

Click the**Edit Strategy**link in the upper right to display the Edit Strategy dialog box as shown in the figure below.

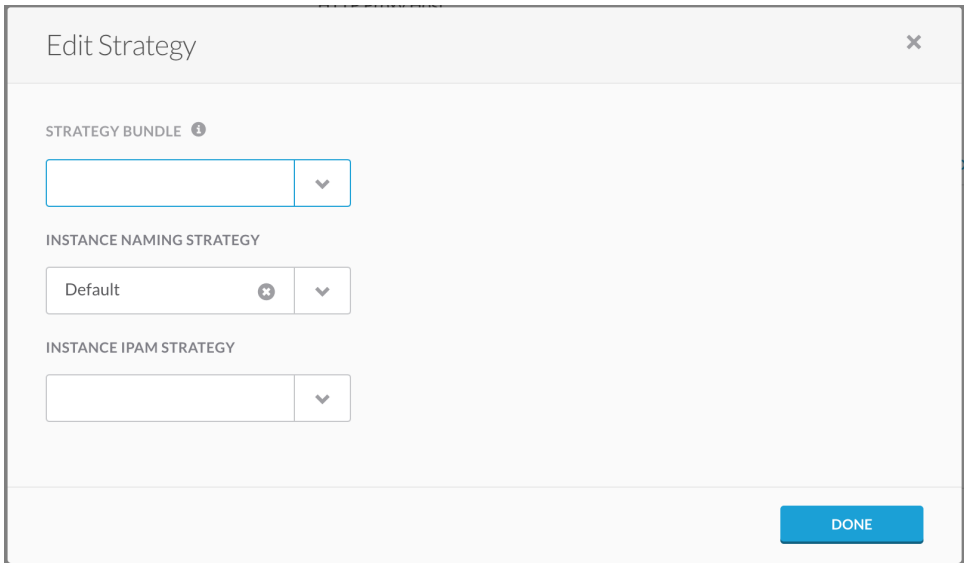

The dialog box can reveal a total of seven data entry fields as explained in the table below:

Field	Usage
Strategy Bundle	Zip file containing all of the scripts associated with your VM naming strategy or IPAM strategy. Select source from the dropdown an add the path in the new field that appears to the right. 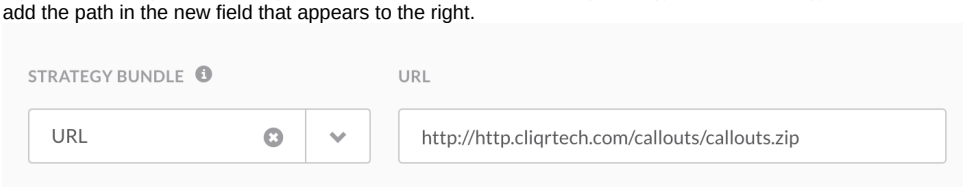
Instance Naming Strategy	Dropdown for selecting one of three naming strategies: Default, CliQr Macro Replacement, and Hostname Callout. If you select CliQ Macro Replacement, the **Node Name Config** field becomes visible and required. If you select Hostname Callout, the **Custom VM Name** field becomes visible and required. See the Instance Naming Strategy table below for details.

Node Name Config	Required if you selected "CliQr Macro Replacement" for theInstance Naming Strategyfield. Content: text string consisting of a combination of characters and numbers with embedded macro variables (see table below). The macros get translated to text string run time.
	
	The resulting VM name string mustmeet the following requirements for job submission to succeed:
	• Only contain the following characters:
	  • Lowercase a to z
	  • 0 to 9
	  • Hyphen
	• Start withan alpha character
	• Cannot end with a hyphen

Instance IPAM Strategy	Dropdown menu with two choices: No IPAM (default), and IPAM Callouts. If you select the latter, two additional fields are displayed and become required: **IPAM Alloc Rule**and**IPAM Dealloc Rule**. (See below).
Custom VM Name	Location of your custom VM naming script if you set the**Instance Naming Strategy**field to "Custom Callout". Location options inclu "Script from bundle" if you specify a zip file in the **Strategy Bundle** field. See Guidance for Callout Scripts for guidance on developi custom VM naming scripts.
IPAM Alloc Rule	Location of your custom IPAM address allocation script if you set the**Instance IPAM Strategy**field to "IPAM Callouts". Location opti include "Script from bundle" if you specify a zip file in the **Strategy Bundle** field. See Guidance for Callout Scripts for guidance on developing custom IPAM callout scripts.
IPAM Dealloc Rule	Similar to the**IPAM Alloc Rule**above except this is where you specify the location of yourcustom IPAM address deallocation script. Guidance for Callout Scripts for guidance on developing custom IPAM callout scripts.

Instance Naming Strategy Options:

Options	Description
Default	The default method provided by Cisco is the generation of a string consisting of A predefined token prefixed with **cq jw-**
CliQr Macro Replacement	Selecting this method causes the **Node Name Config** field to be displayed and to become a required field. See the fields table above for details.
Hostname Callout	Selecting this method causes the **Custom VM Name** field to be displayed and to become a required field. See the fields table above for details.

OOB macro variables forCliQr Macro Replacement:

Macro	Returned Value
%os1%	First character of OS type string in lower case
%os2%	Firsttwo characters of OS type string in lower case
%OS1%	Firstcharacter of OS type string in upper case
%OS2%	Firsttwo character of OS type string in upper case
%RND<number>%	Fixed length random string, e.g., %RND10%, 10-character long random string ⚠ This is a mandatory field, only values in the range: %RND6% to %RND11% or just %RND% which defaults to %RND6%"
%UID%	User ID
%VID%	Vendor ID
%VM_PREFIX%	VM name prefix must be added as a global parameter in the app profile. ⚠ Only macro whose value can be user-customized.

6. External Lifecycle Actions (conditional):Specify any external lifecycle actions to be performed on all VMs launched by Workload Manager in this region.

External Lifecycle Actions Settings

Whenever you use Workload Manager to deploy a VM-based application, Workload Manager can execute external actions on VMs in the deployment as the VMs progress through these five phases of their lifecycle:

- Pre VM start
- Pre VM initialization
- PostVM initialization
- Pre VM Stop
- Post VM Stop

These lifecycle actions will execute in a script execution engine running in a container in the CloudCenter Suite cluster if no Cloud Remote appliance is configured for the region. If you have configured Cloud Remote for this region, the external lifecycle actions will executein a script execution engine running in a container in the Cloud Remote appliance.

External Lifecycle Action scripts can be defined at the service level, the application level, and the cloud region level. Use the External Lifecycle Actions section of the Region tab or Details tab (see figure below) tospecify external lifecycle actions at the region level.

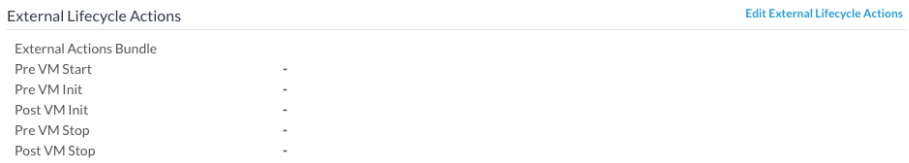

Click on the**Edit External Lifecycle Actions**link in the upper right. This displays theEdit External Lifecycle Actions dialog box as shown below.

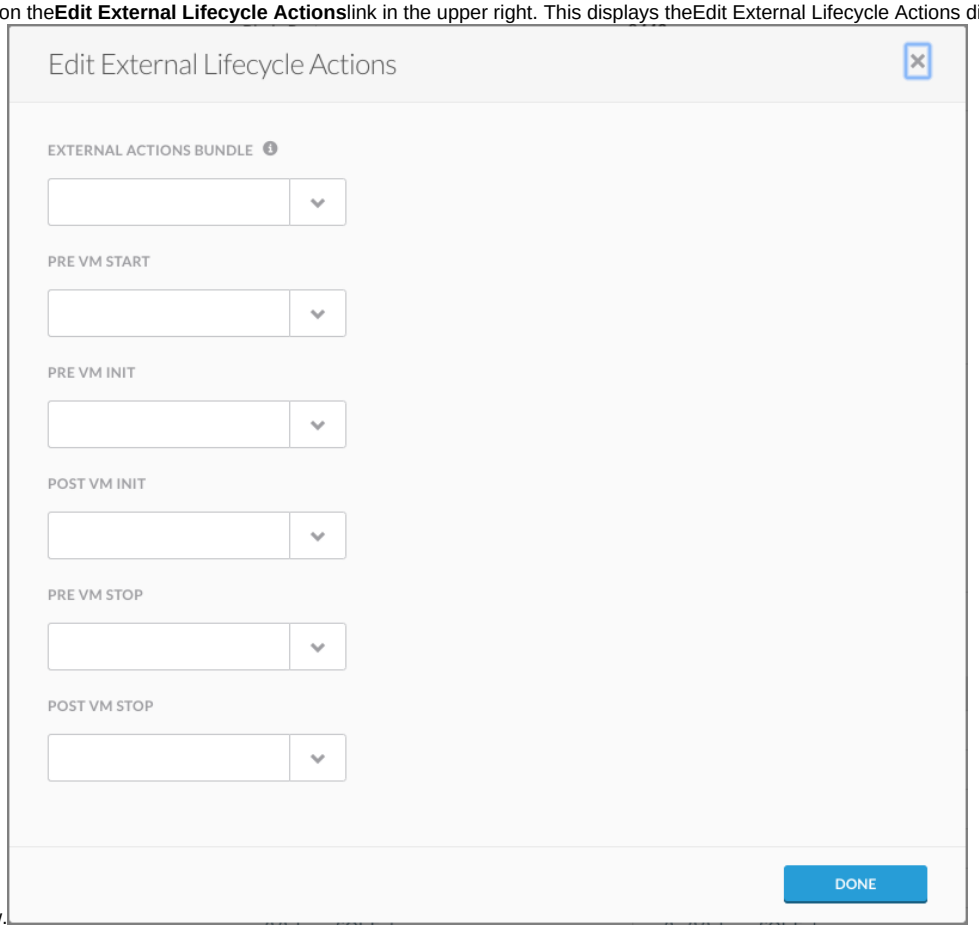

One way to specify an external lifecycle action is to use the first field,**External Actions Bundle**, to specify a zip file that contains all scripts for external lifecycle management. This zip file needs to contain a directory that contains all the scripts. Use the dropdown to select a repository location (URL, FTP, S3, etc.). This causes an additional field to be displayed to the right for entering the path to the zip file as shown in the figurebelow.

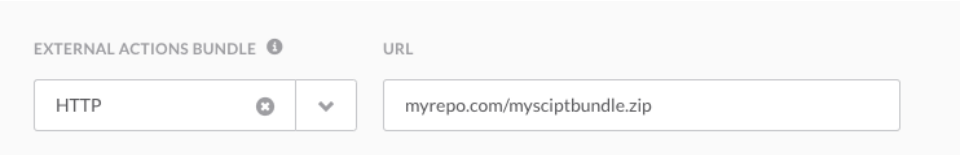

Once you specify actions bundle, for each of the individual lifecycle stage fields, select "Script from Bundle" from the dropdown. This displays an additional field to the right where you must specify the name of the script file in the bundle as shown in the figure below.

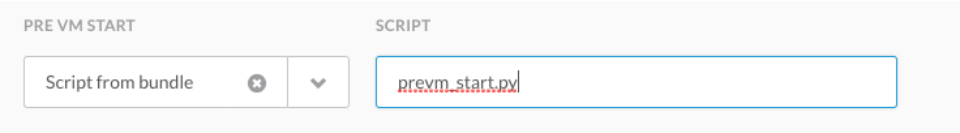

The scripts you specify can be bash scripts or Python scripts called from within a bash script.For more information on lifecycle scripts, seeDeployment Lifecycle Scripts.

7. Instance Types (informational): CloudCenter Suite automatically synchronizes instance types for public cloud regions on a daily basis. This data includes published pricing for each instance type. It is possible to edit Google region instance types, but only the changes in the cost are used by CloudCenter Suite.

Instance Types Settings

An instance type is a performance category of a VM for VM-based clouds or a performance category of a container for container-based clouds. The parameters for Google cloud instance type is as follows:

- Display Name
- Price per hour
- Instance Type ID
- CPU
- 32- or 64-bit architecture
- RAM (MB)
- NICs
- Instance storage (GB)
- SSD support

For Google cloud, the following apply:

- You can only sync, edit, and add **Auto-populate and Auto-sync** and **Edit** function.
- CloudCenter Suite auto-populates the instance types in the Instance Types section.
- CloudCenter Suite auto-syncs this information from the cloud provider once per 24 hours.

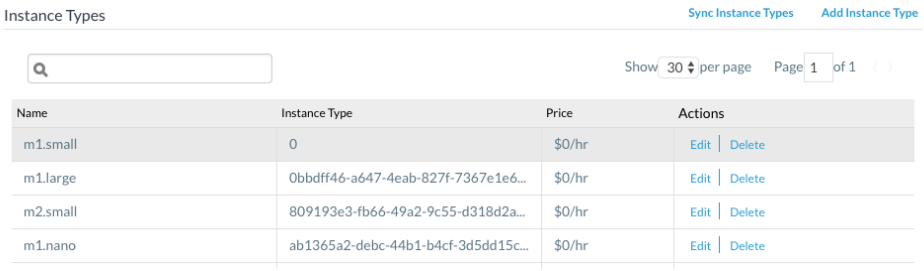

> ⚠ You cannot force sync on demand nor edit the field of an instance type (except instance type ID) by clicking the**Edit**link in the **Actions** column in the list of instance types.

8. Storage Types (conditional): CloudCenter Suite automatically synchronizesstorage types for public cloud regions on a daily basis. This data includes the cloud provider published pricing for each storage type. It is possible to edit Google region storage types, but only the changes in the cost are used by CloudCenter Suite.

Storage Types Settings

A storage type is a categorization of storage that can be provisioned for an application during deployment. It is a construct that applies to VM-based clouds only. Storage types are displayed in the Storage Types section of the Regions or Details tab.

The storage type construct contains the following fields for all VM-based clouds:

- Display Name
- Cost Per Month per GB
- Cloud Storage Type ID
- Minimum Volume Size (GB)
- Maximum Volume Size (GB)

- Maximum Throughput (MB/s)
- Maximum IOPS
- Provisioned IOPS
- IOPS Cost Per Month

Storage types are handled by CloudCenter Suite in a fashion similar to instance types. Hence, the following apply:

- You can only sync, edit, and add **Auto-populate and Auto-sync** and **Edit** functions.
- CloudCenter Suite auto-populates the instance types in the Instance Types section.
- CloudCenter Suite auto-syncs this information from the cloud provider once per 24 hours.

9. Image Mappings: Image mappings allow services based on Workload Manager logical images to be deployed using the appropriate physical image stored on the target cloud region. Workload Manager automatically maps the OOB logical images to public cloud region physical images when you add the region to your cloud. Cisco periodically updates these mappings when new versions of OS physical image are uploaded by the cloud provider. To apply these updates to your region after it is added to your cloud, click the **Sync Image Mappings** link in the upper right of this section. If you create any custom logical images, you mustmanually import the corresponding physical images into your region and then map the corresponding logical images to these physical images.SeeImagesfor more context.

Add a Google Cloud Account

> Be aware that these screenshots may change based on the Google Cloud platform changes. They are provided in this section as a point of reference.

Prerequisites

Before adding a Google cloud account, verify the followingGoogle requirements:

- A validGoogle Cloud Platform accountwith*Project Owner*permissions
- If using theShared VPC network feature, you also requiredShared VPC Admin permissions (seehttps://cloud.google.com/vpc/docs/provisioning-shared-vpcfor additional context).
- CloudCenter Suite appends the network name with a unique ID to form the firewall rule name; the network name can be a maximum of 24 (network name) + 39 (unique ID) = 63 total characters.For example:abcdefghijklmnopqrstuvwx-c3f-462828f37a06acd3ee194716bfe10de0
- Enable the following APIs for each Google cloud account you will be adding to CloudCenter Suite:
    - Google Compute Engine API
    - Google Cloud Resource Manager API
    - Google Cloud SQL Admin API (needed only for Cost Optimizer for PAAS services)

  The following image depicts the Google portal to enabled APIs:

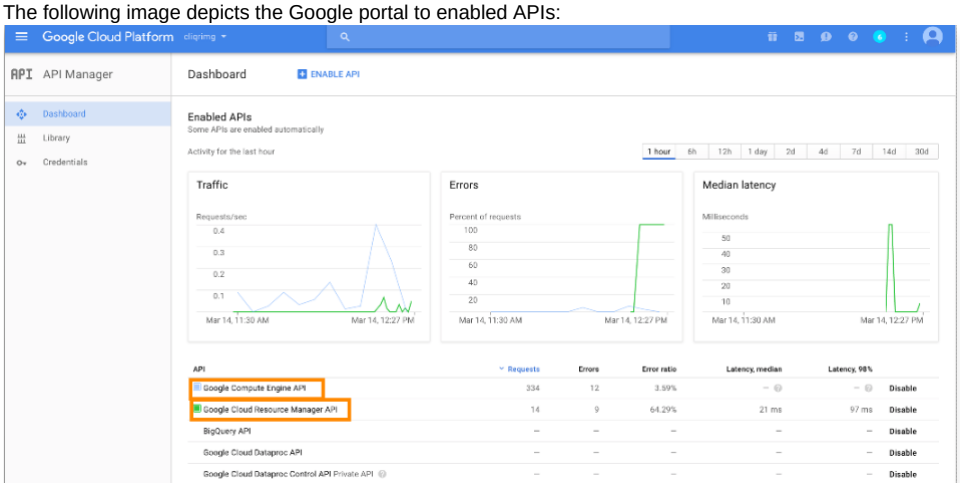

- Set the minimum permissions for your cloud account.SeeCloud Overview>*Minimum Permissions for Public Clouds*for additional details.

- Create a new **service account key in JSON format** per the GCP documentation:https://cloud.google.com/iam/docs/creating-managing-service-account-keys.
  Make sure you use the default JSON format as shown in the create key dialog box below.

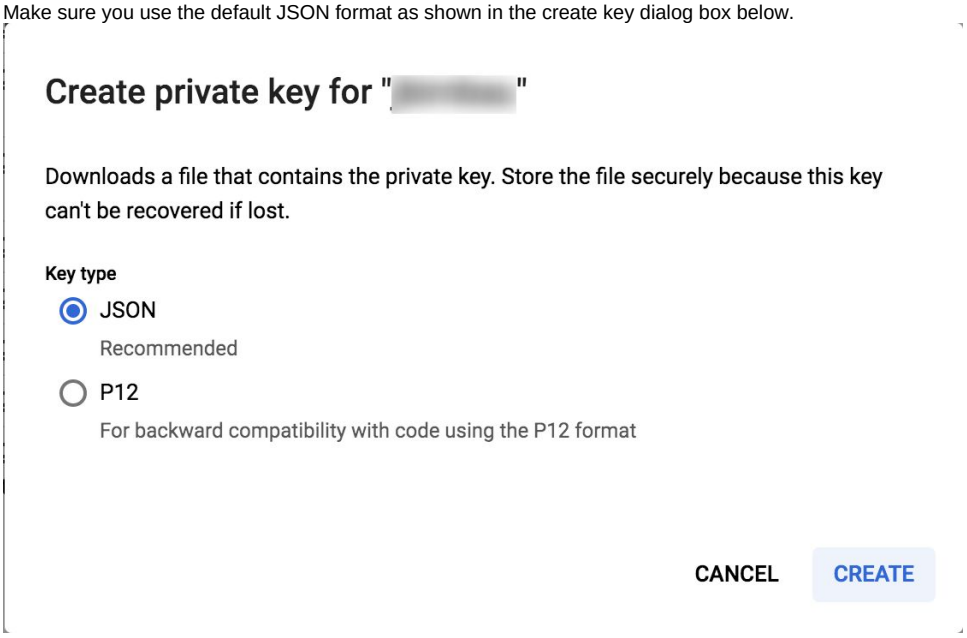

- Once you click **Create**, the file will be downloaded by your browser. Make note of its name and location as you will need to specify this inthe**Service Account JSON File**field in the CloudCenter Suite UI as explained below.

Configuration Process

To add a Google cloud account, follow this procedure.

1. Locate the newly-added cloud and click the**Add Cloud Account**link. The Add Cloud Account dialog box displays:

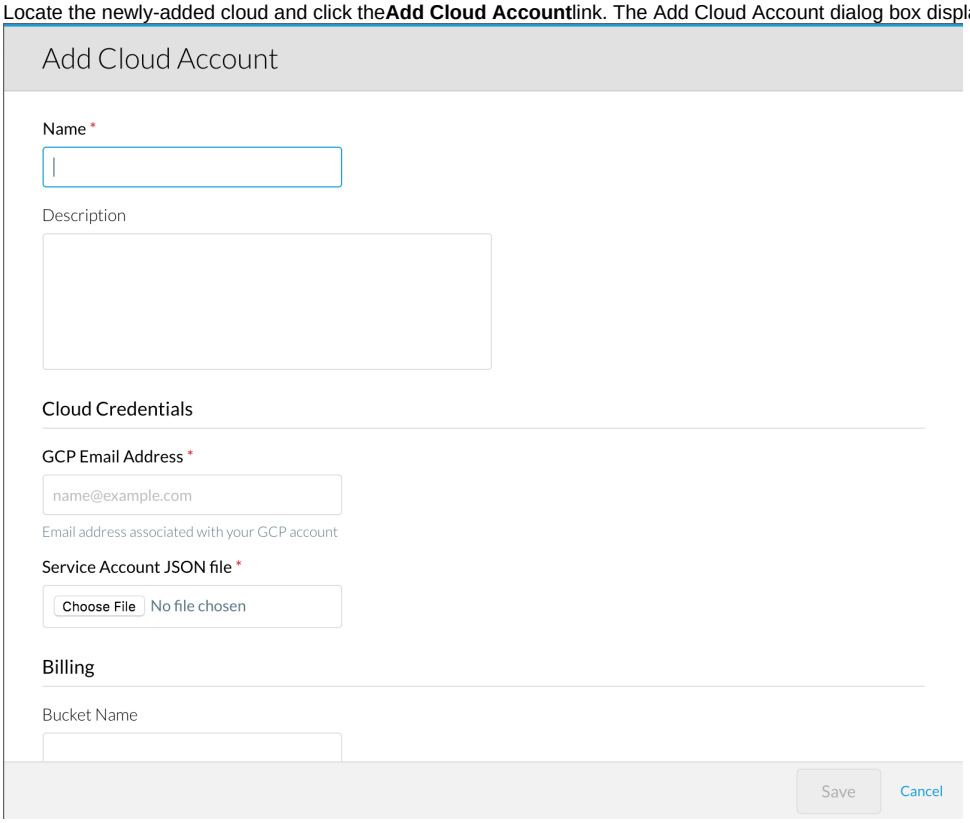

2. Assign a new cloud account name.

> **Tip**
>
> The name should not contain any space, dash, or special characters.

3. Add the following Cloud Credentials associated with your Google account.

   The location of these details in GCP is identified in the*Prerequisites*section.

Field	Description
GCP Email Address	The email address that you used to log into the GCP account.
GCP Service Account JSON File	The JSON private key associated with the Service Account. (See *Prerequisites*section)

4. Enter the **Bucket Name** and **Report Prefix** as shown in the figure below.For information on setting up billing information, seehttps://cloud.google.com/billing/docs/how-to/export-data-file.

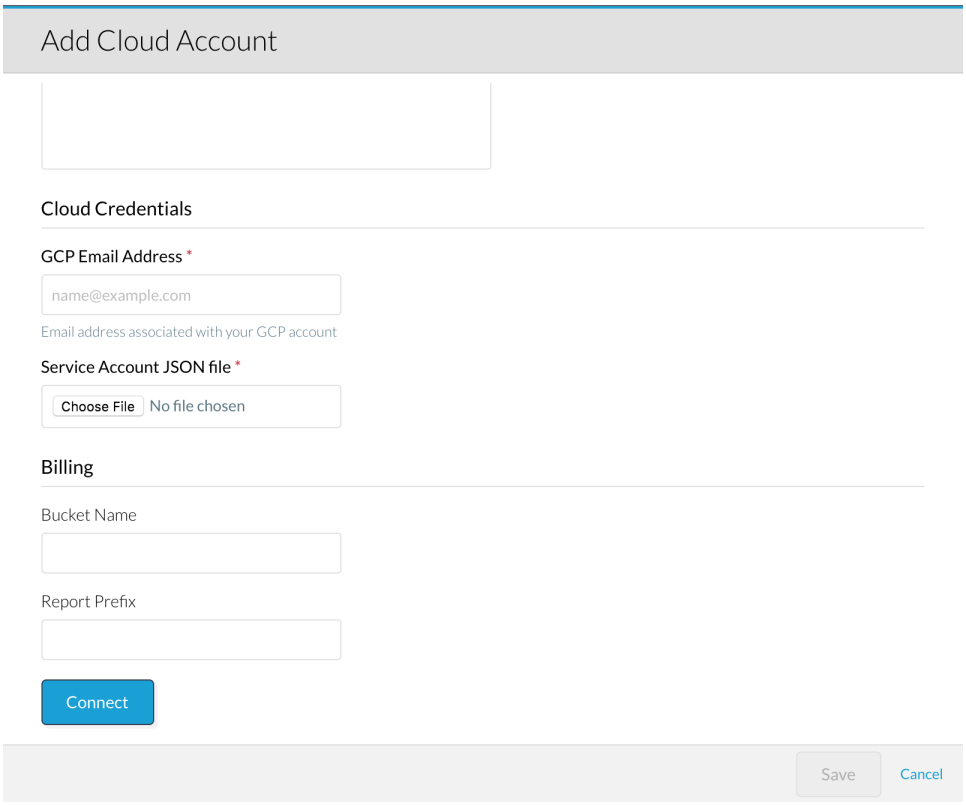

> In the cloud console, create a bucket, if it does not exist already, and navigate to **Billing** > **Billing Export** to view billing information.

5. Click the **Connect** button. CloudCenter Suite will now attempt to validate your account credentials.
6. After the credentials are verified, the**Connect**button changes to an**Edit**button and two new fields appear**Enable Account For**and**Enable Reporting By Org Structure**,

   a. Set the**Enable Account For**dropdown per the table below.

Value	Usage
Provisioning	Workload Manager can deploy jobs using this account.

Reporting	Cost Optimizer and Workload Manager will track cloud costs for this account. Typical usage: master cloud accounts that are used for billing aggregation. ⚠ It is recommended that you do not add a *Reporting* account to the same tenant through different cloud groups. ⚠ Enabling a public cloud account for *Reporting* may incur expenses to retrieve cost data. These expenses are proportional to the number of configured cloud accounts and regions.
Provisioning, Reporting	Default. Account is used for both provisioning and reporting.

b. **For AWS and Google clouds only**: Set the**Enable Reporting By Org Structure**toggle to On to cause Cost Optimizer to import the cost hierarchy createdin the cloud provider portal. This saves the time of manually creating a comparable cost hierarchy within Cost Optimizer. SeeCost Groups Configurationfor more information on cost hierarchies in Cost Optimizer.
c. Click the**Save**button when done.

Cloud Accounts Tab

After you add cloud accounts to a cloud, they will appear in the Accounts tab for the cloud as shown in the figure below.

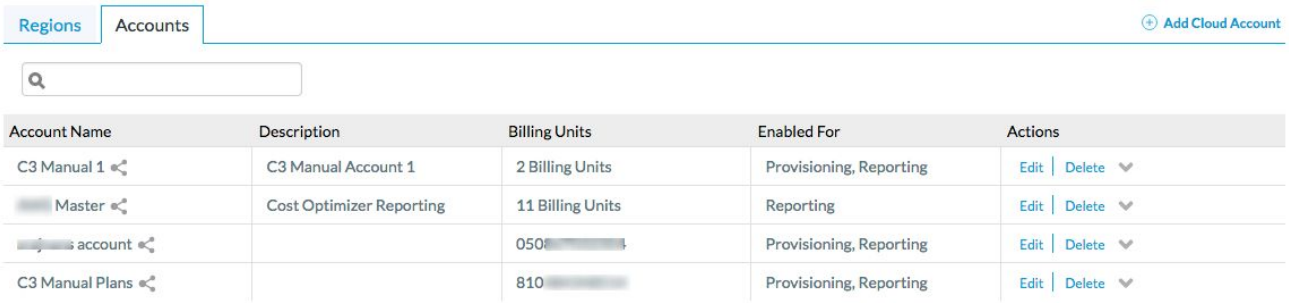

The Accounts tab contains columns for data entered when creating an account: Account Name, Description, Enabled For; and two additional columns:**Billing Units**and**Actions**.**Billing Units**is a dual function:

- If the cloud account contains only one billing unit, the ID for that billing unit is displayed.
- If the cloud account contains multiple billing units, such as an AWS master account, the number of billing units in that account is displayed followed by the text*Billing Units*.

A billing unit is the most granular level of cloud cost recording in CloudCenter Suite. The definition of a billing unit varies by a cloud provider as shown in the table below.

Cloud Provider	Billing Unit
AWS	Account ID
AzureRM	Subscription ID
Google	Project ID
IBM Cloud	Account ID
vCenter	Cloud Group Prefix - Datacenter Name
vCD	Organization Name
OpenStack	Project ID
Kubernetes	Namespace UID

The last column,**Actions**, contains links to let you edit or deleted the cloud account, ormanage instance typesfor the cloud account.

Configure a Kubernetes Cloud

Configure aKubernetes Cloud

Configuring aKubernetes cloud is a three-step process:

- Add a KubernetesCloud
- Configure aKubernetes Region
- Add a Kubernetes Cloud Account

Add a KubernetesCloud

To add a Kubernetescloud follow these steps.

1. Navigate to**Admin > Clouds**. This brings you to the Clouds page. If you, or another tenant admin in your tenant, have already added clouds to your tenant, they will be listed here.
2. Click the**Add Cloud link**in the upper right. The **Add Cloud** dialog box is displayed.
3. Enter the**Cloud Name**and select the**cloud provider**.

   When assigning the Cloud Name, use a unique name. Ensure that you do not assign a previously-used name (including deleted cloud names) as the Cloud Name as they may cause deployments to fail intermittently.

4. Since you are selecting select a Kubernetes cloud provider, a new data entry field appears at the bottom of the dialog box called **Kubernetes Cluster API Endpoint**. You must enter the URL of the Kubernetes API endpoint in this field before the**Next**button is enabled. When done click**Nex** .
5. After clicking**Next**, the second page of the Add Clouds dialog box, Connectivity Settings, appears. Set the toggle switches to indicate the Cloud Connectivity Settings for a Kubernetes Cloud

   - When adding a Kubernetes cloud in the Workload Manager or Cost Optimizer UI, the second page of the Add Clouds dialog box, Connectivity Settings,appears with a single toggle displayed:**Cloud Endpoint Directly Accessible.**
   - Setting this toggle to **No** implies you will install Cloud Remote in the VM cloud that is hosting this Kubernetes cloud. This also causes a second toggle to be displayed: **CloudCenter Suite Directly Accessible from Cloud Remote**
   - Follow the table below for guidance on setting these toggles.

Toggle settings	Use case	Network Diagram
Cloud Endpoint Directly Accessible = Yes	CloudCenter Suite cluster can initiate a connection to the Kubernetes API endpoint Cloud Remote is not required	
Cloud Endpoint Directly Accessible = No AND CloudCenter SuiteDirectly Accessible fromCloud Remote = Yes	CloudCenter Suite cluster cannot initiate a connection to the Kubernetes API endpoint AND Cloud Remote can initiate the connection to the CloudCenter Suite cluster	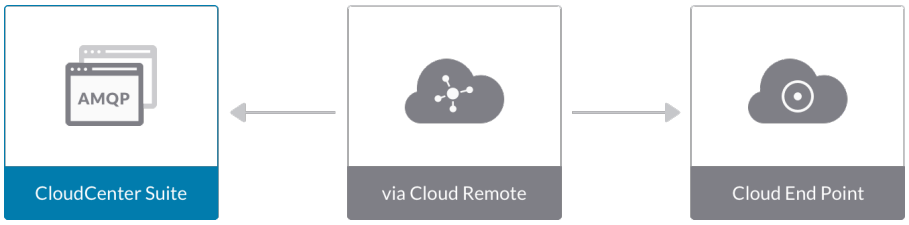

Cloud Endpoint Directly Accessible = No AND CloudCenter SuiteDirectly Accessible fromCloud Remote = No	CloudCenter Suite cluster cannot initiate a connection to the cloud region API endpoint AND Cloud Remote cannot initiate the connection to the Cloud Center Suite cluster	

6. Click**Done**to save the configuration and close the dialog box. This brings you back to the **Clouds** page and the cloud you just created will be added to the bottom of the list on the left side of the page.

Configure aKubernetes Region

AKubernetes cloud has one region that you configure from theKubernetes cloud Details tab. Follow this procedure:

1. Navigate to Clouds page: **Admin > Clouds**. Find your newly createdKubernetes cloud from the cloud list on the left half of the screen and click its **Configure Cloud** link. This displays the **Details** tab for this cloud.
2. Click the **Edit Kubernetes Settings** link in the upper right to open the **Configure Cloud Settings** dialog box. Adjust the field values in the dialog box per the instructions in the following table.

Field	Usage
Kubernetes cluster API Endpoint	This field is set to the value you set for the API endpoint when you created thisKubernetes cloud. You can edit it here but should only do so if the API endpoint address of yourKubernetes cloud has changed since you added it to CloudCenter Suite.
API version override	⚠ Kubernetes can add or deprecate API versions for Kubernetes resources used by the CloudCenter platform. In Kubernetes 1.16, for example, the following versions were deprecated fordeployments: extensions/v1beta1, apps /v1beta1, and apps/v1beta2 as listedhere. Based on the version of the Kubernetes cluster that you are managing, you can use the API Version Override feature to specify the supported version. This tellsCloudCenter Suite to use an API version other than the default version for certain Kubernetes resources.This field should normally be left blank. If errors occur in your deployments, contact support regarding using a different version for selected resources. This is a semicolon-separated list of key-value pairs in the format: <resource_name_1>:<api_version_1>; <resource_name_2>:<api_version_2>; etc. Possible examples are as follows: • Example 1: **Secret:***custom_api_version***;Service:***custom_api_version***;PersistentVolumeClaim:***custom_api_version***;NetworkPolicy:***custom_api_version***;Pod:c***ustom_api_version***;Deployment:***custom_api_version* • Example 2: **PersistentVolumeClaim:***custom_api_version***;NetworkPolicy:***custom_api_version***;Pod:c***ustom_api_version***;Deployment:***custom_api_version* • Example 3: **PersistentVolumeClaim:***custom_api_version***;NetworkPolicy:***custom_api_version*
Namespace(s)	If at least one of the cloud accounts that you add to this cloud has admin privileges for the cloud (recommended), CloudCenter Suite will automatically find all namespaces in the cloud. You can leave this field blank. If none of your cloud accounts for this cloud have sufficient privileges to retrieve the list of namespaces in the cluster, use this field to manually enter the comma-separated list of namespaces.

When you are done editing the settings in the dialog box, click **Save**.

3. Scroll down to the **Region Connectivity** section for the region and click on the **Configure Region** link in the upper right to open the **Configure Region** dialog box. The toggle settings should be the same as when you set them on the connectivity page of the **Add Cloud** dialog box. If all of the connectivity toggles in the **Region Connectivity** dialog box are set to **Yes**, then Cloud Remote is NOT needed for this cloud region. In this case, you would normally leave the region connectivity settings at their current values and continue to the next settings section.

4. If any of theconnectivity toggles inthe Region Connectivity dialog box are set to No, then you must install and configure Cloud Remote for this region. Since Cloud Remote is a VM-based appliance, when used to support a Kubernetes cloud it must be installed in a VM-based cloud region that is accessible from the Kubernetes cloud. Typically, this would be the same cloud region that hosts the nodes supporting the Kubernetes cloud. Choose the option that is appropriate for your Kubernetes target cloud:

Configure Cloud Remote in a Google Region for a Kubernetes Cloud

Configure Cloud Remote in a Google regionto support a Kubernetes target cloud as follows.

Obtain and Launch the Cloud Remote Appliance in Google

a. Request the Cloud Remote shared VMI form Cisco support by opening a CloudCenter Support case. In your request, specify the following details:

   i. Your GCP account number
   ii. Your GCP project ID number
   iii. Your CloudCenter Suite version
   iv. Your Customer ID (CID)
   v. Your customer name
   vi. Specify if your setup is in production or for a POC
   vii. Your Contact Email

b. After you open a case, your support case is updated with the shared VMI ID.**Proceed to the next step only after your supportcase is updated with the VMI ID.**
c. Navigate to the GCP dashboard and search for the VMI ID name provided in theCloudCenter Support case in the list of images for your project.
d. Launch an instance using the shared VMI.

   i. Click on the image name. This takes you to the page for the image

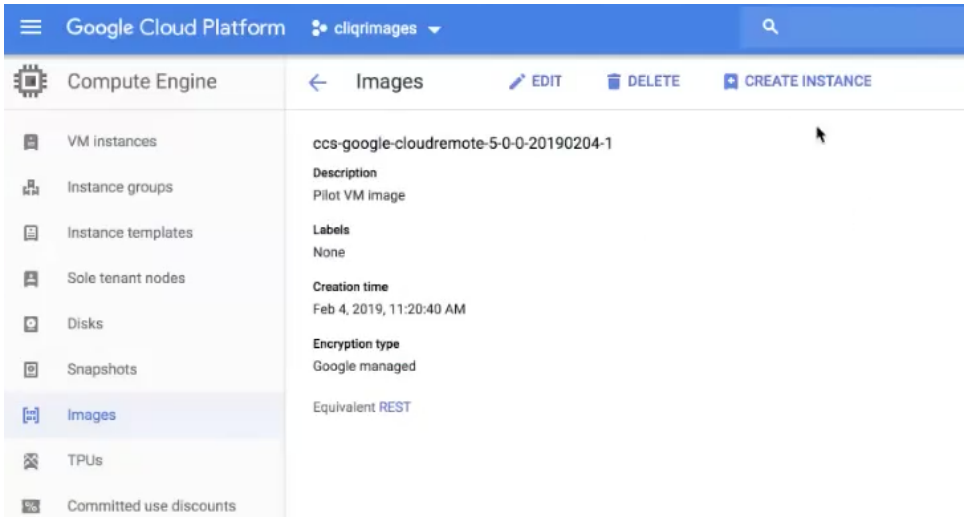

   ii. Click on Create Instance to display the Instance properties page

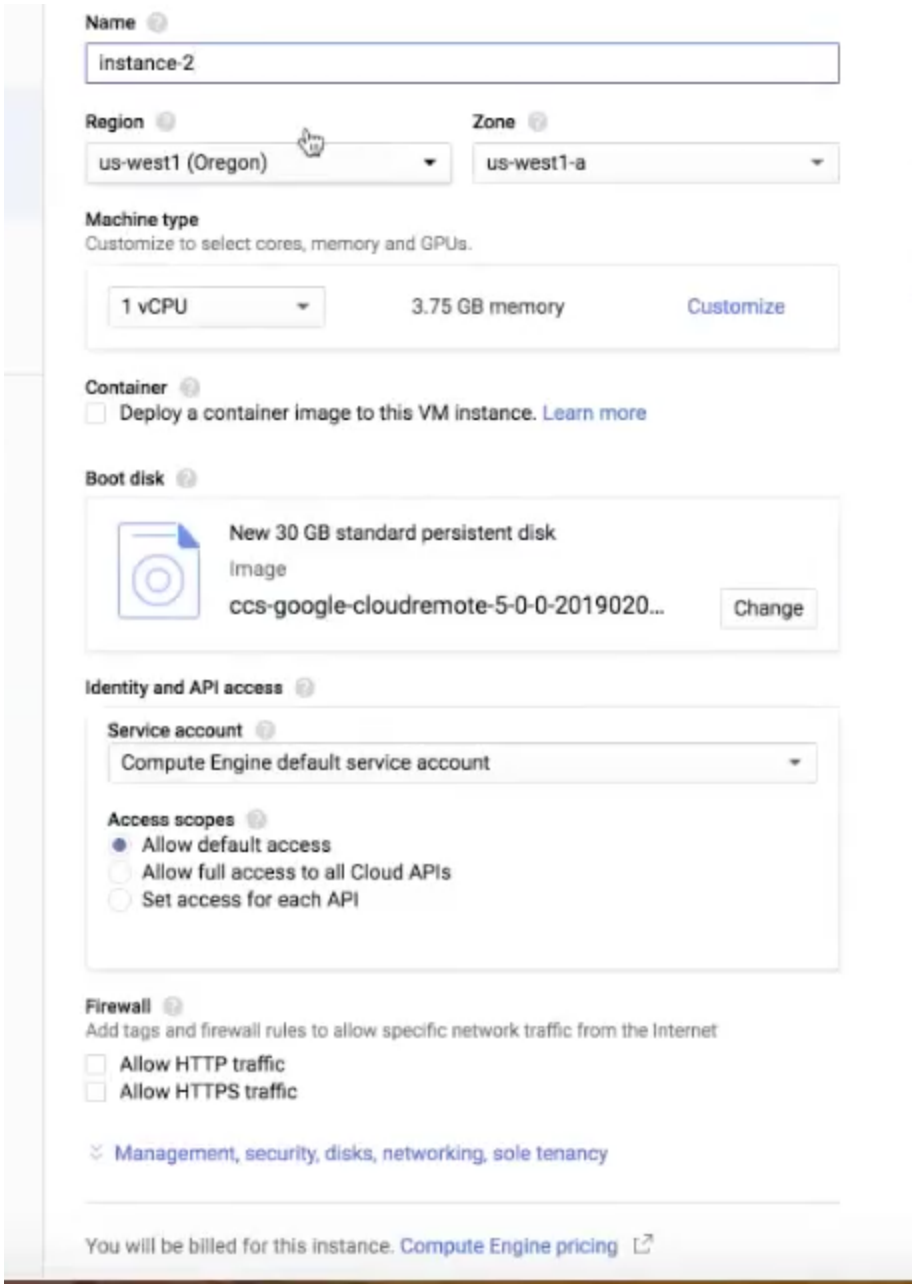

iii. Complete these fields:

1. Instance name
2. Region and zone
3. Machine type: select 2 vCPU, 7.5 GB RAM
4. Click the checkbox to allow HTTPS access
5. Click the Security tab (under the Allow HTTPS traffic checkbox). In the SSH key field, add your organization's public ssh key followed by a space and then the username you want to use to login to the Cloud Remote appliance. Click the Add Item button when done.

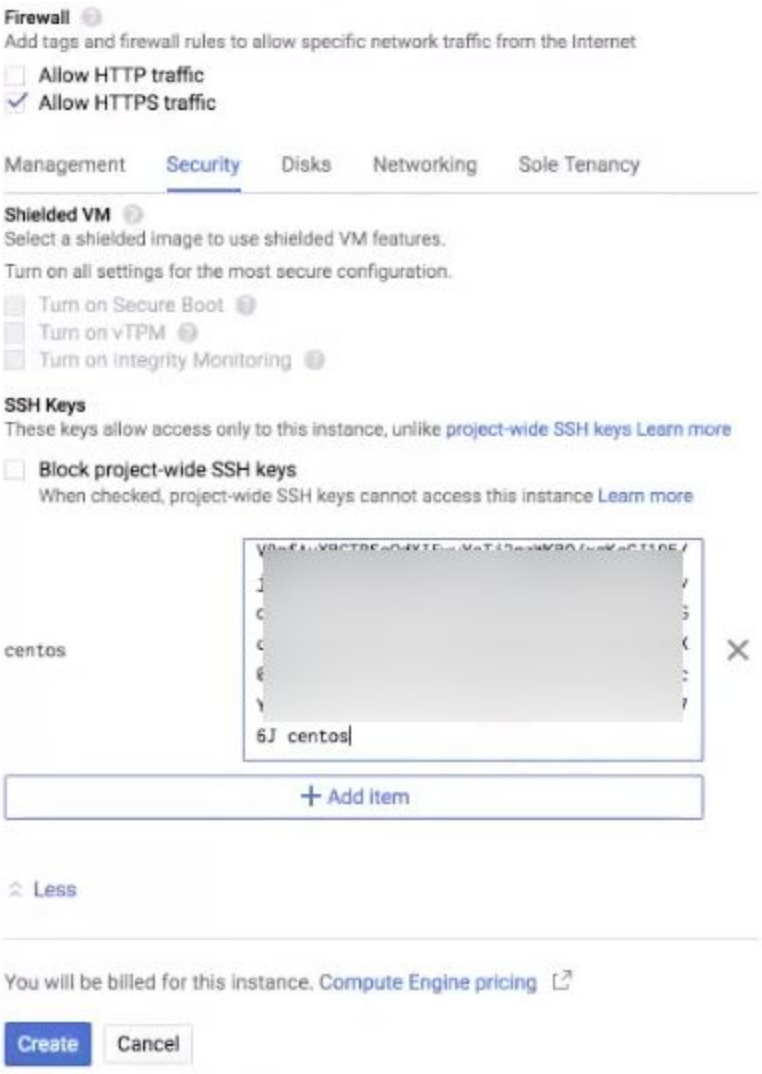

iv. Click Create to launch the instance.

e. Optional but recommended for production environments: Deploy two additional instances of the appliance to form a cluster for HA. Cloud Remote includes support for the clustering of multiple nodes. You will "add" these two additional instances to the first instance after the first instance is configured. SeeCloud Remote (Conditional)> *Scaling*for details.

f. Once the first instance of the appliance has been launched, use theGCP console to**note its IP public and private addresses**. You will need this information later on in order to login to the Cloud Remote web UI and to complete the Region Connectivity settings in the CloudCenter Suite Web UI. Also, note the IP addresses of any other appliances you launch.

Unable to render {include} The included page could not be found.

Specify AMQP Addresses for Supporting Cloud Remote for a Kubernetes Cloud

From the CloudCenter Suite UI, for the Kubernetes cloud requiring Cloud Remote, navigate to the corresponding Details tab. Click the**Configure Region**link in the upper left of the Region Connectivity section to bring up the Configure Region dialog box.

The toggle settings should be the same as when you set them on the connectivity page of the Add Cloud dialog box. You may need to update the **Local AMQP IP Address**or the**Remote AMQP IP Address**fields per the table below.

Toggle Settings	Field	Value
Cloud Endpoint Directly Accessible = No AND CloudCenter Directly Accessible from Cloud Remote = Yes	Local AMQP IP Address	Pre-populated with the address and port number of the "local" AMQP server running in the CloudCenter Suite cluster. If Cloud Remote is accessing the CloudCenter Suite cluster through a user-supplied proxy server or NAT firewall, overwrite this field with the corresponding local AMQP IP address and port number provided by the user-supplied proxy server or NAT firewall and accessible to Cloud Remote .

Cloud Endpoint Directly Accessible = No AND CloudCenter Directly Accessible from Cloud Remote = No	Remote AMQP IP Address	Enter <Cloud_Remote_IP>:<amqp_port>, where <Cloud_Remote_IP> = the IP address Cloud Remote which is accessible to the CloudCenter Suite cluster, and <amqp_port> = 5671 OR the custom AMQP port number you would later set with the**Change Ports shell script** on the Cloud Remote appliance (seeCloud Remote (Conditional) >*Custom Port Numbers (Conditional)*). If there is no user-supplied NAT firewall or proxy server between the CloudCenter Suite cluster and Cloud Remote , the IP address would be the public IP address of Cloud Remote . If there is a NAT firewall or proxy server between the CloudCenter Suite cluster and Cloud Remote , instead, enter the corresponding public IP address and port number that the firewall or proxy server presents to the internet on behalf of the "remote" AMQP server running in Cloud Remote .

When done, click **OK** to save the setting and dismiss the dialog box.

Download Region Connectivity Settings and Upload to Cloud Remote

Cloud Remote uses the region connectivity settings set in the Workload Manager or Cost Optimizer UI. After saving the Region Configuration settings in the Workload Manager or Cost Optimizer UI, you must download them and to your local computer and then upload them to Cloud Remote as follows.

Click the**Download Configuration**link in the upper right of the Region Connectivity section, as shown in the figure below.

Clicking DownloadConfiguration causestwo things to happen:

- An encrypted zip file named**artifacts.zip**will be downloaded by your browser. Make note of the location of this zip file as you will need to upload it to Cloud Remote through the Cloud Remote web UI (see below).
- The Region Connectivity section header updates to display a **Copy Encryption Key**link, as shown in the figure below.

Click the**Copy Encryption Key**link to save the key to your clipboard. A success message will be displayed temporarily above the Region Connectivity section header. Make sure not to overwrite the clipboard with other data. You will need the key when you upload the configuration zip file to Cloud Remote .

> If you change the connectivity settings in the CloudCenter Suite UI and need to again download the zip file, a new encryption key is automatically created and can be copied to the clipboard by clicking the **Copy Encryption Key** link again.

After you have downloaded the zip file and copied the encryption key to your clipboard, login toCloud Remote web UI.

a. Open another browser tab and login to https://<Cloud Remote _ip> with thedefault credentials: admin/cisco.
b. You will immediately be required to change your password. Do so now.
c. You are now brought to the Cloud Remote home page as shown in the figure below.

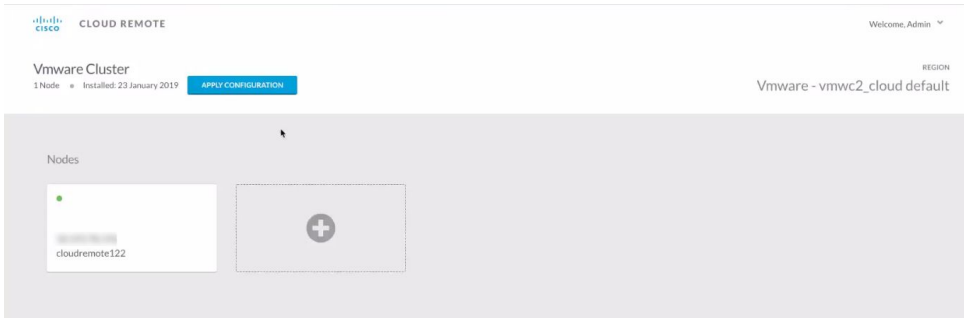

d. Click the **Apply Configuration** button in the page header. This prompts you to select a configuration file and enter the encryption key as shown in the figure below.

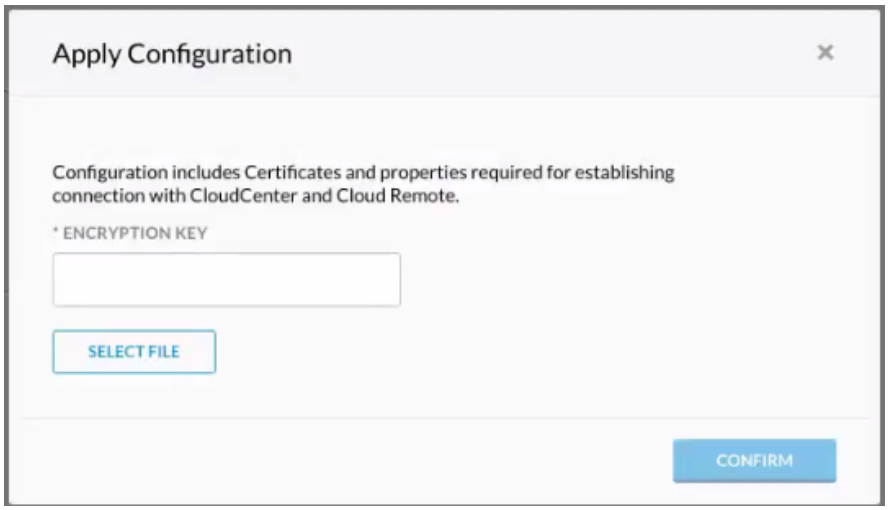

e. Paste the encryption key that was copied to the clipboard into the**Encryption Key**field in the dialog box.
f. Click**Select File**and browse tothe artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
g. Click**Confirm**.
h. Once the zip file is successfully transmitted and accepted, the Cloud Remote appliance attempts to establish communication with the CloudCenter Suite cluster and the Cloud Remote web UI home page is updated to show the name of the region it is connecting to in the upper right (see figure below).

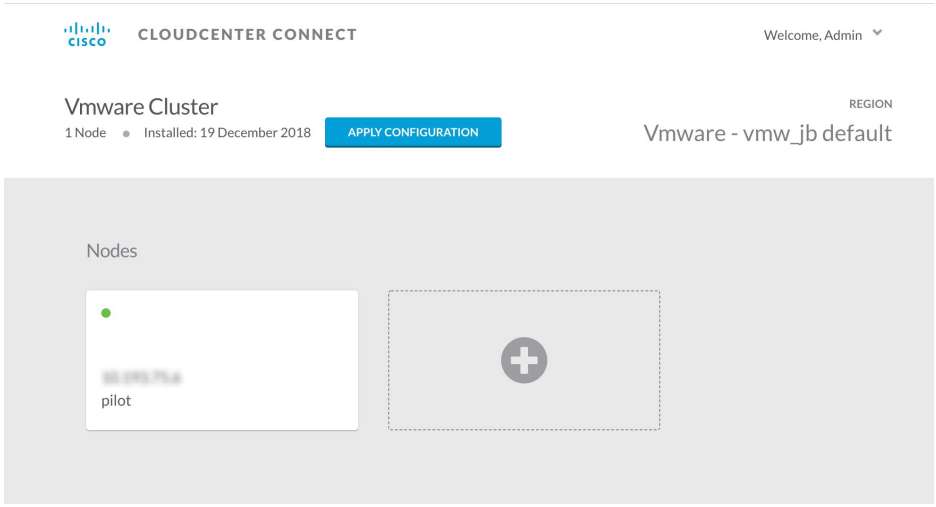

Switch your focus back to the Region Connectivity section of the target cloud region in the CloudCenter Suite web UI. The status indicator in the Region Connectivity section header will change from Not Configured to Running once connectivity between Cloud Remote and the CloudCenter Suite cluster is completely established (see figure below).

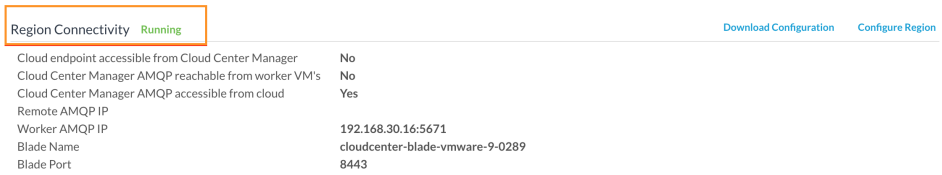

After completing these steps, Workload Manager and Cost Optimizer can use Cloud Remote for communicating with the target cloud region.

Configure Cloud Remote in a vCenter Region for a Kubernetes Cloud

Configure Cloud Remote in a vCenter region to support a Kubernetes target cloud as follows.

Download and Launch the Cloud Remote Appliance in vCenter

a. From your local computer, download the Cloud Remote appliance OVA fromsoftware.cisco.com.
b. Log in to the vCenter console using the vSphere web client with Flash, or with the vSphere Windows client. Do not use the HTML5 web client.

c. Navigate to the folder or resource pool where you want to deploy the OVA. Right-click on that resource pool or folder and select Deploy OVF Template.
d. From the Deploy OVF Template dialog box, for Source, select Local file and click Browse to find the OVA file you downloaded in step 1.
e. Complete the fields for Name and location, Host / Cluster, Resource Pool, Storage, and Disk Format appropriate for your environment.
f. For the Network Mapping section, make sure to properly map the Management network (public) and VM Network network (private) to the appropriate network names in your environment.
g. For the Properties section, make sure to check the box labeled Does the VM need a second interface? if the Cloud Remote appliance needs to be multi-homed on a public network and a private network.
h. Confirm your settings and click Finish to launch the VM.
i. Optional but recommended for production environments:Deploy two additional instances of the appliance to form a cluster for HA. Cloud Remote includes support for the clustering of multiple nodes. You will "add" these two additional instances to the first instance after the first instance is configured. SeeCloud Remote (Conditional)> *Scaling*for details.
j. Once the first instance of the appliance has been launched, use thevSphere clienttonote its IP public and private addresses. You will need this information later on in order to login to the Cloud Remote web UI and to complete the Region Connectivity settings in the CloudCenter Suite Web UI. Also, note the IP addresses of any other appliances you launch.

Unable to render {include} The included page could not be found.

Specify AMQP Addresses for Supporting Cloud Remote for a Kubernetes Cloud

From the CloudCenter Suite UI, for the Kubernetes cloud requiring Cloud Remote, navigate to the corresponding Details tab. Click the**Configure Region**link in the upper left of the Region Connectivity section to bring up the Configure Region dialog box.

The toggle settings should be the same as when you set them on the connectivity page of the Add Cloud dialog box. You may need to update the **Local AMQP IP Address**or the**Remote AMQP IP Address**fields per the table below.

Toggle Settings	Field	Value
Cloud Endpoint Directly Accessible = No AND CloudCenter Directly Accessible from Cloud Remote = Yes	Local AMQP IP Address	Pre-populated with the address and port number of the "local" AMQP server running in the CloudCenter Suite cluster. If Cloud Remote is accessing the CloudCenter Suite cluster through a user-supplied proxy server or NAT firewall, overwrite this field with the corresponding local AMQP IP address and port number provided by the user-supplied proxy server or NAT firewall and accessible to Cloud Remote .
Cloud Endpoint Directly Accessible = No AND CloudCenter Directly Accessible from Cloud Remote = No	Remote AMQP IP Address	Enter <Cloud_Remote_IP>:<amqp_port>, where <Cloud_Remote_IP> = the IP address Cloud Remote which is accessible to the CloudCenter Suite cluster, and <amqp_port> = 5671 OR the custom AMQP port number you would later set with the**Change Ports shell script** on the Cloud Remote appliance (seeCloud Remote (Conditional) >*Custom Port Numbers (Conditional)*). If there is no user-supplied NAT firewall or proxy server between the CloudCenter Suite cluster and Cloud Remote , the IP address would be the public IP address of Cloud Remote . If there is a NAT firewall or proxy server between the CloudCenter Suite cluster and Cloud Remote , instead, enter the corresponding public IP address and port number that the firewall or proxy server presents to the internet on behalf of the "remote" AMQP server running in Cloud Remote .

When done, click **OK** to save the setting and dismiss the dialog box.

Download Region Connectivity Settings and Upload to Cloud Remote

Cloud Remote uses the region connectivity settings set in the Workload Manager or Cost Optimizer UI. After saving the Region Configuration settings in the Workload Manager or Cost Optimizer UI, you must download them and to your local computer and then upload them to Cloud Remote as follows.

Click the**Download Configuration**link in the upper right of the Region Connectivity section, as shown in the figure below.

Clicking DownloadConfiguration causestwo things to happen:

- An encrypted zip file named**artifacts.zip**will be downloaded by your browser. Make note of the location of this zip file as you will need to upload it to Cloud Remote through the Cloud Remote web UI (see below).
- The Region Connectivity section header updates to display a **Copy Encryption Key**link, as shown in the figure below.

Click the**Copy Encryption Key**link to save the key to your clipboard. A success message will be displayed temporarily above the Region Connectivity section header. Make sure not to overwrite the clipboard with other data. You will need the key when you upload the configuration zip file to Cloud Remote .

If you change the connectivity settings in the CloudCenter Suite UI and need to again download the zip file, a new encryption key is automatically created and can be copied to the clipboard by clicking the **Copy Encryption Key** link again.

After you have downloaded the zip file and copied the encryption key to your clipboard, login toCloud Remote web UI.

a. Open another browser tab and login to https://<Cloud Remote _ip> with thedefault credentials: admin/cisco.
b. You will immediately be required to change your password. Do so now.
c. You are now brought to the Cloud Remote home page as shown in the figure below.

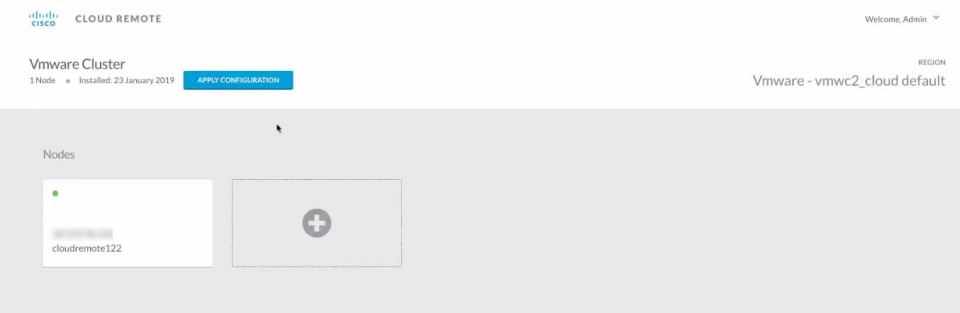

d. Click the **Apply Configuration** button in the page header. This prompts you to select a configuration file and enter the encryption key as shown in the figure below.

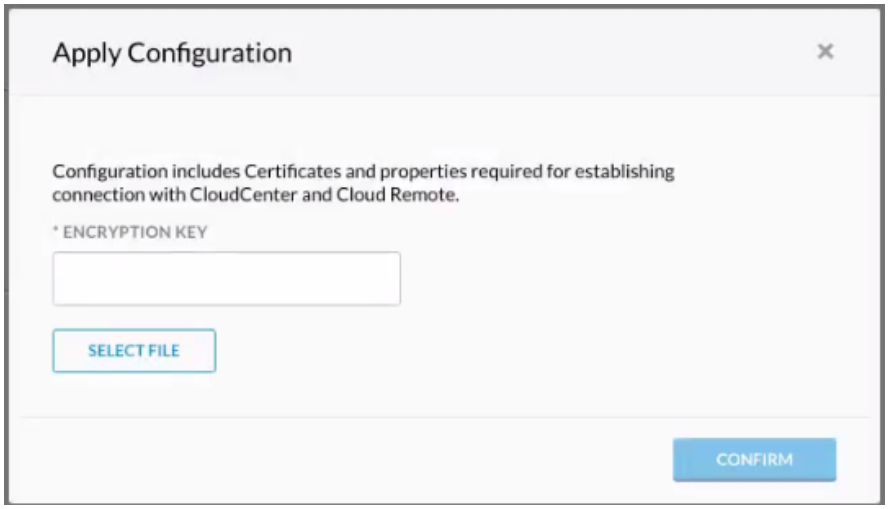

e. Paste the encryption key that was copied to the clipboard into the**Encryption Key**field in the dialog box.
f. Click**Select File**and browse tothe artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
g. Click**Confirm**.
h. Once the zip file is successfully transmitted and accepted, the Cloud Remote appliance attempts to establish communication with the CloudCenter Suite cluster and the Cloud Remote web UI home page is updated to show the name of the region it is connecting to in the upper right (see figure below).

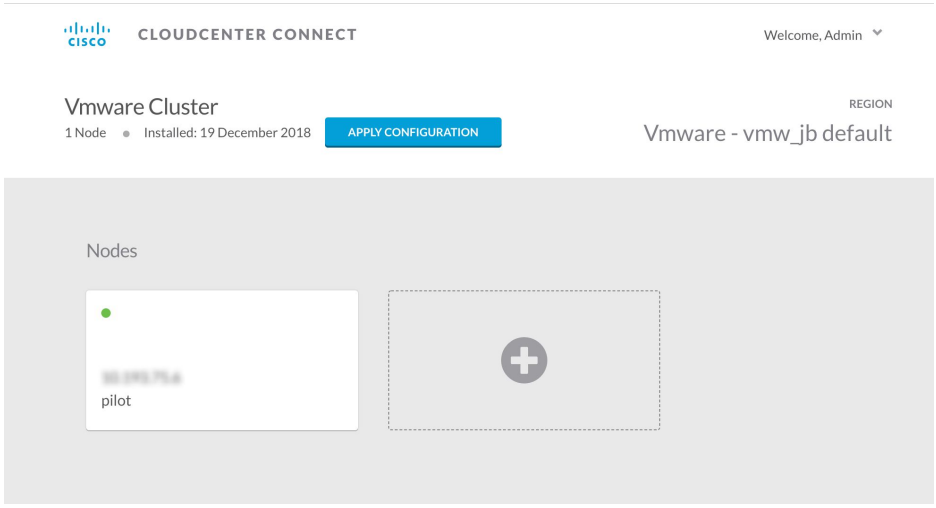

Switch your focus back to the Region Connectivity section of the target cloud region in the CloudCenter Suite web UI. The status indicator in the Region Connectivity section header will change from Not Configured to Running once connectivity between Cloud Remote and the CloudCenter Suite cluster is completely established (see figure below).

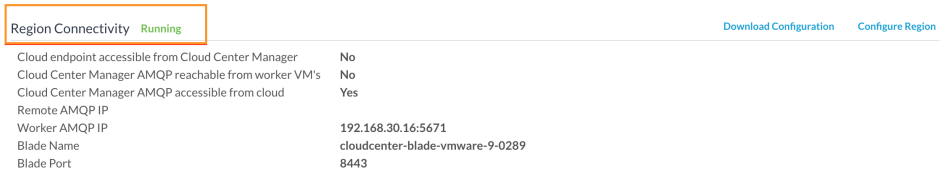

After completing these steps, Workload Manager and Cost Optimizer can use Cloud Remote for communicating with the target cloud region.

Configure Cloud Remote in an AWS Region for a Kubernetes Cloud

The SSH username used to be *ec2-user* for Cloud Remote images on AWS prior to Workload Manager 5.2.0. Effective Workload Manager 5.2.0, this username has been changed to **centos**.

Configure Cloud Remote in an AWS regionto support a Kubernetes target cloud as follows.

Obtain and Launch the Cloud Remote Appliance in AWS

a. Obtain the Cloud Remote shared AMI form Cisco support and launch it. Follow the same guidance for obtaining and launching theCloudCenter Suite installer appliance for AWS.
b. Optional but recommended for production environments: Deploy two additional instances of the appliance to form a cluster for HA.Cloud Remote includes support for the clustering of multiple nodes. You will "add" these two additional instances to the first instance after the first instance is configured. SeeCloud Remote (Conditional) > *Scaling*for details.
c. Once the first instance of the appliance has been launched, use your cloud console to**note its IP public and private addresses**. You will need this information later on in order to login to the Cloud Remote web UI and to complete the Region Connectivity settings in the CloudCenter Suite Web UI. Also, note the IP addresses of any other instances you launch.

Unable to render {include} The included page could not be found.

Specify AMQP Addresses for Supporting Cloud Remote for a Kubernetes Cloud

From the CloudCenter Suite UI, for the Kubernetes cloud requiring Cloud Remote, navigate to the corresponding Details tab. Click the**Configure Region**link in the upper left of the Region Connectivity section to bring up the Configure Region dialog box.

The toggle settings should be the same as when you set them on the connectivity page of the Add Cloud dialog box. You may need to update the **Local AMQP IP Address**or the**Remote AMQP IP Address**fields per the table below.

Toggle Settings	Field	Value
Cloud Endpoint Directly Accessible = No AND CloudCenter Directly Accessible from Cloud Remote = Yes	Local AMQP IP Address	Pre-populated with the address and port number of the "local" AMQP server running in the CloudCenter Suite cluster. If Cloud Remote is accessing the CloudCenter Suite cluster through a user-supplied proxy server or NAT firewall, overwrite this field with the corresponding local AMQP IP address and port number provided by the user-supplied proxy server or NAT firewall and accessible to Cloud Remote .
Cloud Endpoint Directly Accessible = No AND CloudCenter Directly Accessible from Cloud Remote = No	Remote AMQP IP Address	Enter <Cloud_Remote_IP>:<amqp_port>, where <Cloud_Remote_IP> = the IP address Cloud Remote which is accessible to the CloudCenter Suite cluster, and <amqp_port> = 5671 OR the custom AMQP port number you would later set with the**Change Ports shell script** on the Cloud Remote appliance (seeCloud Remote (Conditional) >*Custom Port Numbers (Conditional)*). If there is no user-supplied NAT firewall or proxy server between the CloudCenter Suite cluster and Cloud Remote , the IP address would be the public IP address of Cloud Remote . If there is a NAT firewall or proxy server between the CloudCenter Suite cluster and Cloud Remote , instead, enter the corresponding public IP address and port number that the firewall or proxy server presents to the internet on behalf of the "remote" AMQP server running in Cloud Remote .

When done, click **OK** to save the setting and dismiss the dialog box.

Download Region Connectivity Settings and Upload to Cloud Remote

Cloud Remote uses the region connectivity settings set in the Workload Manager or Cost Optimizer UI. After saving the Region Configuration settings in the Workload Manager or Cost Optimizer UI, you must download them and to your local computer and then upload them to Cloud Remote as follows.

Click the**Download Configuration**link in the upper right of the Region Connectivity section, as shown in the figure below.

Clicking DownloadConfiguration causestwo things to happen:

- An encrypted zip file named**artifacts.zip**will be downloaded by your browser. Make note of the location of this zip file as you will need to upload it to Cloud Remote through the Cloud Remote web UI (see below).
- The Region Connectivity section header updates to display a **Copy Encryption Key**link, as shown in the figure below.

Click the**Copy Encryption Key**link to save the key to your clipboard. A success message will be displayed temporarily above the Region Connectivity section header. Make sure not to overwrite the clipboard with other data. You will need the key when you upload the configuration zip file to Cloud Remote .

> If you change the connectivity settings in the CloudCenter Suite UI and need to again download the zip file, a new encryption key is automatically created and can be copied to the clipboard by clicking the **Copy Encryption Key** link again.

After you have downloaded the zip file and copied the encryption key to your clipboard, login toCloud Remote web UI.

a. Open another browser tab and login to https://<Cloud Remote _ip> with thedefault credentials: admin/cisco.
b. You will immediately be required to change your password. Do so now.
c. You are now brought to the Cloud Remote home page as shown in the figure below.

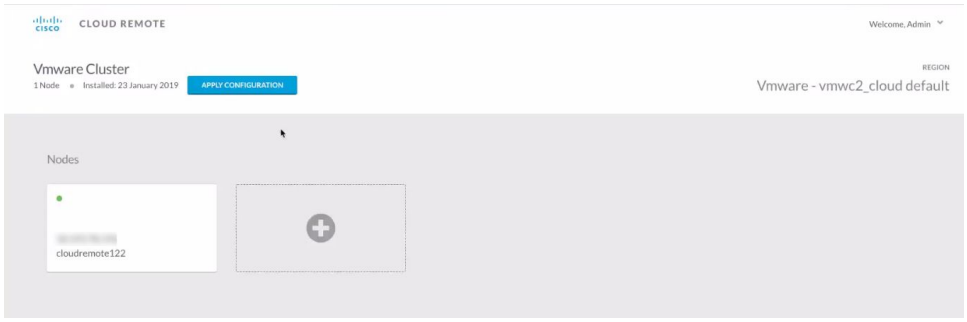

d. Click the **Apply Configuration** button in the page header. This prompts you to select a configuration file and enter the encryption key as shown in the figure below.

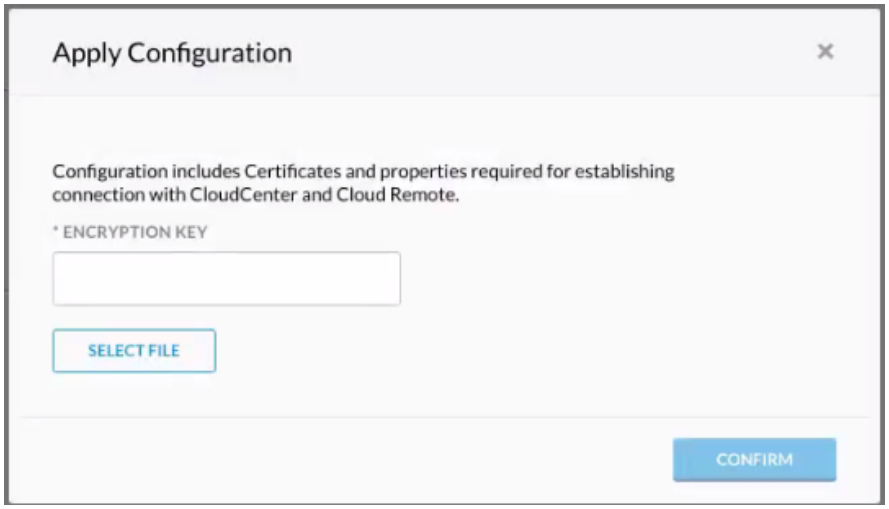

e. Paste the encryption key that was copied to the clipboard into the**Encryption Key**field in the dialog box.
f. Click**Select File**and browse tothe artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
g. Click**Confirm**.
h. Once the zip file is successfully transmitted and accepted, the Cloud Remote appliance attempts to establish communication with the CloudCenter Suite cluster and the Cloud Remote web UI home page is updated to show the name of the region it is connecting to in the upper right (see figure below).

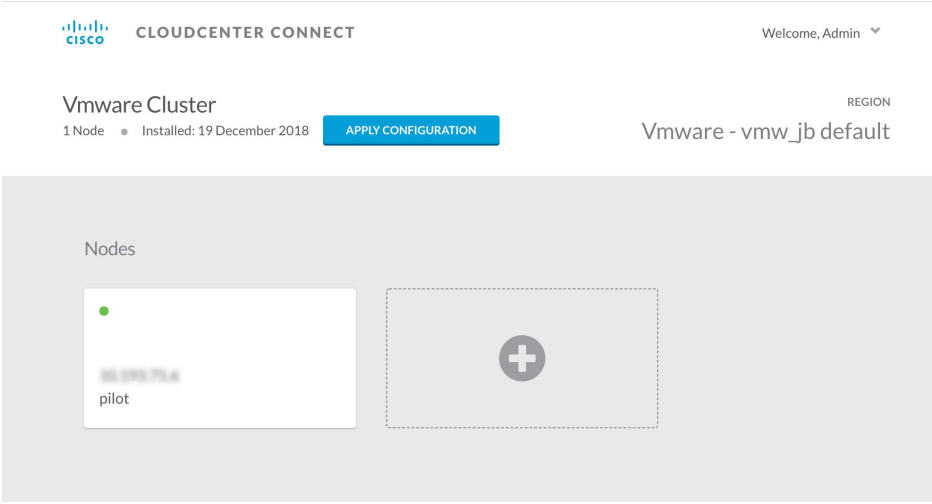

Switch your focus back to the Region Connectivity section of the target cloud region in the CloudCenter Suite web UI. The status indicator in the Region Connectivity section header will change from Not Configured to Running once connectivity between Cloud Remote and the CloudCenter Suite cluster is completely established (see figure below).

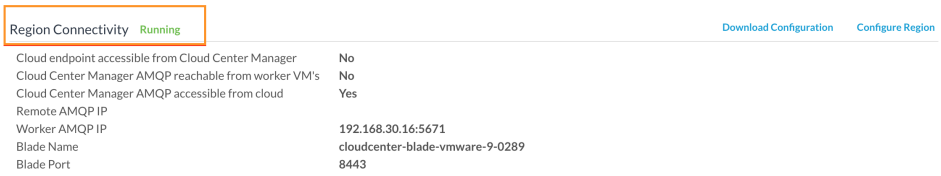

After completing these steps, Workload Manager and Cost Optimizer can use Cloud Remote for communicating with the target cloud region.

Configure Cloud Remote in an AzureRM Region for a Kubernetes Cloud

Configure Cloud Remote in an AzureRM regionto support a Kubernetes target cloudas follows.

Download and Launch the Cloud Remote Appliance in AzureRM

a. Download the Cloud Remote appliance for AzureRM as a zip file fromsoftware.cisco.comand then unzip it to reveal the VHD file.
b. Upload the Cloud Remote appliance VHD file to AzureRM using the AzureRM CLI, then launch the appliance from the AzureRM console web UI. This process is similar to uploading and launching theCloudCenter Suite installer appliance for AzureRM.

> You must use the AzureRM CLI to perform this upload.

c. Optional but recommended for production environments: Deploy two additional instances of the appliance to form a cluster for HA.Cloud Remote includes support for the clustering of multiple nodes. You will "add" these two additional instances to the first instance after the first instance is configured. SeeCloud Remote (Conditional)> *Scaling*for details.
d. Once the first instance of the appliance has been launched, use theAzureRM console to**note its IP public and private addresses**. You will need this information later on in order to login to the Cloud Remote web UI and to complete the Region Connectivity settings in the CloudCenter Suite Web UI. Also, note the IP addresses of any other appliances you launch.

Unable to render {include} The included page could not be found.

Specify AMQP Addresses for Supporting Cloud Remote for a Kubernetes Cloud

From the CloudCenter Suite UI, for the Kubernetes cloud requiring Cloud Remote, navigate to the corresponding Details tab. Click the**Configure Region**link in the upper left of the Region Connectivity section to bring up the Configure Region dialog box.

The toggle settings should be the same as when you set them on the connectivity page of the Add Cloud dialog box. You may need to update the **Local AMQP IP Address**or the**Remote AMQP IP Address**fields per the table below.

Toggle Settings	Field	Value

Cloud Endpoint Directly Accessible = No AND CloudCenter Directly Accessible from Cloud Remote = Yes	Local AMQP IP Address	Pre-populated with the address and port number of the "local" AMQP server running in the CloudCenter Suite cluster. If Cloud Remote is accessing the CloudCenter Suite cluster through a user-supplied proxy server or NAT firewall, overwrite this field with the corresponding local AMQP IP address and port number provided by the user-supplied proxy server or NAT firewall and accessible to Cloud Remote .
Cloud Endpoint Directly Accessible = No AND CloudCenter Directly Accessible from Cloud Remote = No	Remote AMQP IP Address	Enter <Cloud_Remote_IP>:<amqp_port>, where <Cloud_Remote_IP> = the IP address Cloud Remote which is accessible to the CloudCenter Suite cluster, and <amqp_port> = 5671 OR the custom AMQP port number you would later set with the**Change Ports shell script** on the Cloud Remote appliance (seeCloud Remote (Conditional) >*Custom Port Numbers (Conditional)*). If there is no user-supplied NAT firewall or proxy server between the CloudCenter Suite cluster and Cloud Remote , the IP address would be the public IP address of Cloud Remote . If there is a NAT firewall or proxy server between the CloudCenter Suite cluster and Cloud Remote , instead, enter the corresponding public IP address and port number that the firewall or proxy server presents to the internet on behalf of the "remote" AMQP server running in Cloud Remote .

When done, click **OK** to save the setting and dismiss the dialog box.

Download Region Connectivity Settings and Upload to Cloud Remote

Cloud Remote uses the region connectivity settings set in the Workload Manager or Cost Optimizer UI. After saving the Region Configuration settings in the Workload Manager or Cost Optimizer UI, you must download them and to your local computer and then upload them to Cloud Remote as follows.

Click the**Download Configuration**link in the upper right of the Region Connectivity section, as shown in the figure below.

Clicking DownloadConfiguration causestwo things to happen:

- An encrypted zip file named**artifacts.zip**will be downloaded by your browser. Make note of the location of this zip file as you will need to upload it to Cloud Remote through the Cloud Remote web UI (see below).
- The Region Connectivity section header updates to display a **Copy Encryption Key**link, as shown in the figure below.

Click the**Copy Encryption Key**link to save the key to your clipboard. A success message will be displayed temporarily above the Region Connectivity section header. Make sure not to overwrite the clipboard with other data. You will need the key when you upload the configuration zip file to Cloud Remote .

> If you change the connectivity settings in the CloudCenter Suite UI and need to again download the zip file, a new encryption key is automatically created and can be copied to the clipboard by clicking the **Copy Encryption Key** link again.

After you have downloaded the zip file and copied the encryption key to your clipboard, login toCloud Remote web UI.

a. Open another browser tab and login to https://<Cloud Remote _ip> with thedefault credentials: admin/cisco.
b. You will immediately be required to change your password. Do so now.
c. You are now brought to the Cloud Remote home page as shown in the figure below.

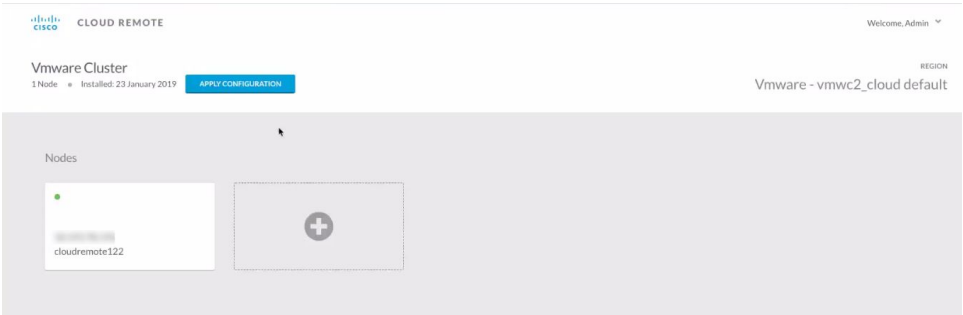

d. Click the **Apply Configuration** button in the page header. This prompts you to select a configuration file and enter the encryption key as shown in the figure below.

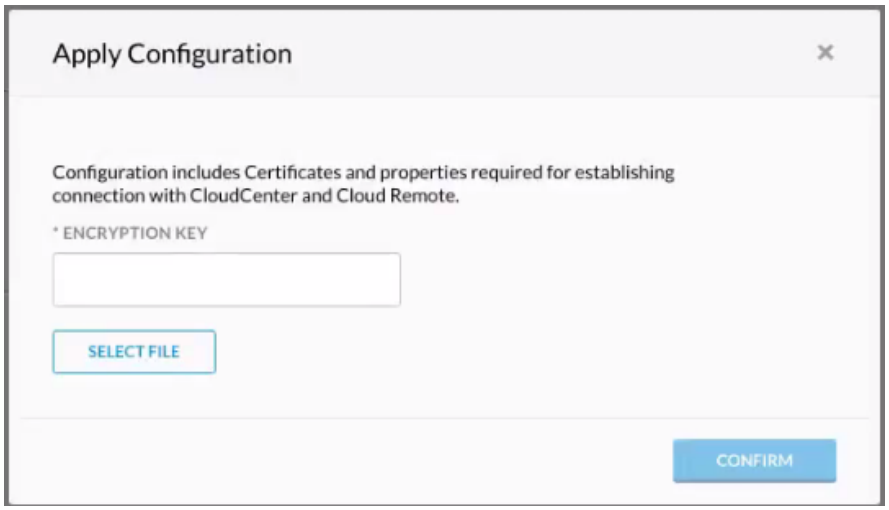

e. Paste the encryption key that was copied to the clipboard into the**Encryption Key**field in the dialog box.
f. Click**Select File**and browse tothe artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
g. Click**Confirm**.
h. Once the zip file is successfully transmitted and accepted, the Cloud Remote appliance attempts to establish communication with the CloudCenter Suite cluster and the Cloud Remote web UI home page is updated to show the name of the region it is connecting to in the upper right (see figure below).

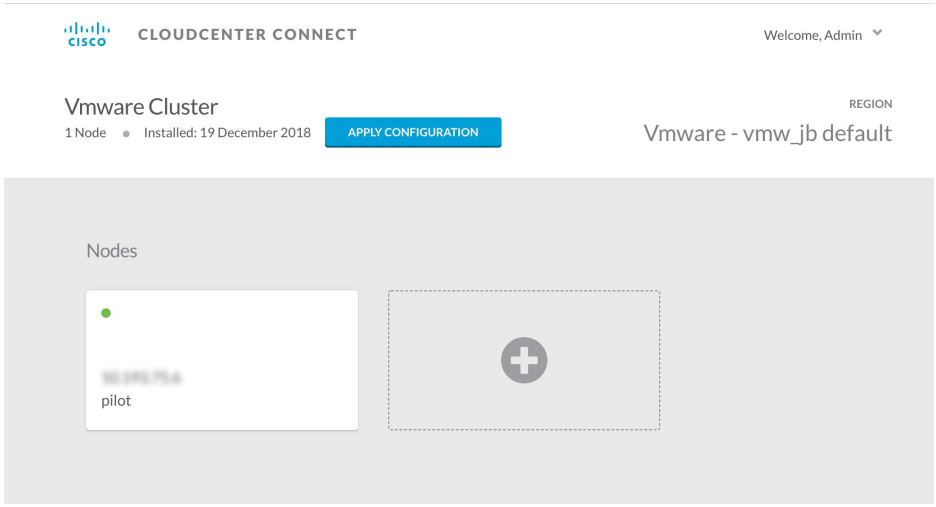

Switch your focus back to the Region Connectivity section of the target cloud region in the CloudCenter Suite web UI. The status indicator in the Region Connectivity section header will change from Not Configured to Running once connectivity between Cloud Remote and the CloudCenter Suite cluster is completely established (see figure below).

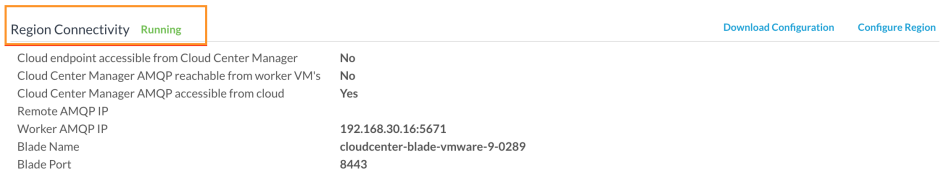

After completing these steps, Workload Manager and Cost Optimizer can use Cloud Remote for communicating with the target cloud region.

Configure Cloud Remote in an OpenStack Region for a Kubernetes Cloud

Configure Cloud Remote in an OpenStack region to support a Kubernetes target cloud as follows.

Download and Launch the Cloud Remote Appliance in OpenStack

a. Download the Cloud Remote appliance qcow2 file fromsoftware.cisco.com.
b. Through the OpenStack console, import and launch the Cloud Remote appliance. This process is similar to importing and launching theCloudCenter Suite installer appliance for OpenStack.

> ⚠ Do not add Network Ports while launching a Cloud Remote instance in OpenStack.

c. Optional but recommended for production environments: Deploy two additional instances of the appliance to form a cluster for HA.Cloud Remote includes support for the clustering of multiple nodes. You will "add" these two additional instances to the first instance after the first instance is configured. See Cloud Remote (Conditional) > *Scaling* for details.
d. Once the first instance of the appliance has been launched, use the OpenStack console to **note its IP public and private addresses**. You will need this information later on in order to login to the Cloud Remote web UI and to complete the Region Connectivity settings in the CloudCenter Suite Web UI. Also, note the IP addresses of any other appliances you launch.

Unable to render {include} The included page could not be found.

Specify AMQP Addresses for Supporting Cloud Remote for a Kubernetes Cloud

From the CloudCenter Suite UI, for the Kubernetes cloud requiring Cloud Remote, navigate to the corresponding Details tab. Click the **Configure Region** link in the upper left of the Region Connectivity section to bring up the Configure Region dialog box.

The toggle settings should be the same as when you set them on the connectivity page of the Add Cloud dialog box. You may need to update the **Local AMQP IP Address** or the **Remote AMQP IP Address** fields per the table below.

Toggle Settings	Field	Value
Cloud Endpoint Directly Accessible = No AND CloudCenter Directly Accessible from Cloud Remote = Yes	Local AMQP IP Address	Pre-populated with the address and port number of the "local" AMQP server running in the CloudCenter Suite cluster. If Cloud Remote is accessing the CloudCenter Suite cluster through a user-supplied proxy server or NAT firewall, overwrite this field with the corresponding local AMQP IP address and port number provided by the user-supplied proxy server or NAT firewall and accessible to Cloud Remote .
Cloud Endpoint Directly Accessible = No AND CloudCenter Directly Accessible from Cloud Remote = No	Remote AMQP IP Address	Enter <Cloud_Remote_IP>:<amqp_port>, where <Cloud_Remote_IP> = the IP address Cloud Remote which is accessible to the CloudCenter Suite cluster, and <amqp_port> = 5671 OR the custom AMQP port number you would later set with the **Change Ports shell script** on the Cloud Remote appliance (see Cloud Remote (Conditional) > *Custom Port Numbers (Conditional)*). If there is no user-supplied NAT firewall or proxy server between the CloudCenter Suite cluster and Cloud Remote , the IP address would be the public IP address of Cloud Remote . If there is a NAT firewall or proxy server between the CloudCenter Suite cluster and Cloud Remote , instead, enter the corresponding public IP address and port number that the firewall or proxy server presents to the internet on behalf of the "remote" AMQP server running in Cloud Remote .

When done, click **OK** to save the setting and dismiss the dialog box.

Download Region Connectivity Settings and Upload to Cloud Remote

Cloud Remote uses the region connectivity settings set in the Workload Manager or Cost Optimizer UI. After saving the Region Configuration settings in the Workload Manager or Cost Optimizer UI, you must download them and to your local computer and then upload them to Cloud Remote as follows.

Click the **Download Configuration** link in the upper right of the Region Connectivity section, as shown in the figure below.

Clicking DownloadConfiguration causes two things to happen:

- An encrypted zip file named **artifacts.zip** will be downloaded by your browser. Make note of the location of this zip file as you will need to upload it to Cloud Remote through the Cloud Remote web UI (see below).
- The Region Connectivity section header updates to display a **Copy Encryption Key** link, as shown in the figure below.

Click the **Copy Encryption Key** link to save the key to your clipboard. A success message will be displayed temporarily above the Region Connectivity section header. Make sure not to overwrite the clipboard with other data. You will need the key when you upload the configuration zip file to Cloud Remote .

If you change the connectivity settings in the CloudCenter Suite UI and need to again download the zip file, a new encryption key is automatically created and can be copied to the clipboard by clicking the **Copy Encryption Key** link again.

After you have downloaded the zip file and copied the encryption key to your clipboard, login toCloud Remote web UI.

a. Open another browser tab and login to https://<Cloud Remote _ip> with thedefault credentials: admin/cisco.
b. You will immediately be required to change your password. Do so now.
c. You are now brought to the Cloud Remote home page as shown in the figure below.

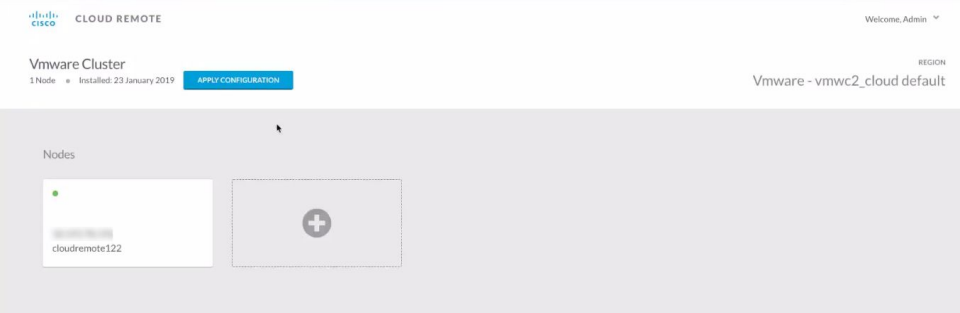

d. Click the **Apply Configuration** button in the page header. This prompts you to select a configuration file and enter the encryption key as shown in the figure below.

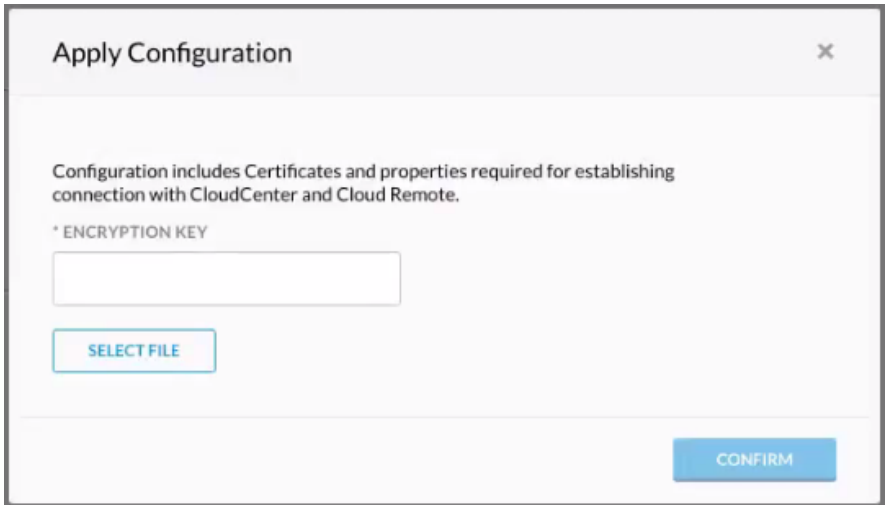

e. Paste the encryption key that was copied to the clipboard into the**Encryption Key**field in the dialog box.
f. Click**Select File**and browse tothe artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
g. Click**Confirm**.
h. Once the zip file is successfully transmitted and accepted, the Cloud Remote appliance attempts to establish communication with the CloudCenter Suite cluster and the Cloud Remote web UI home page is updated to show the name of the region it is connecting to in the upper right (see figure below).

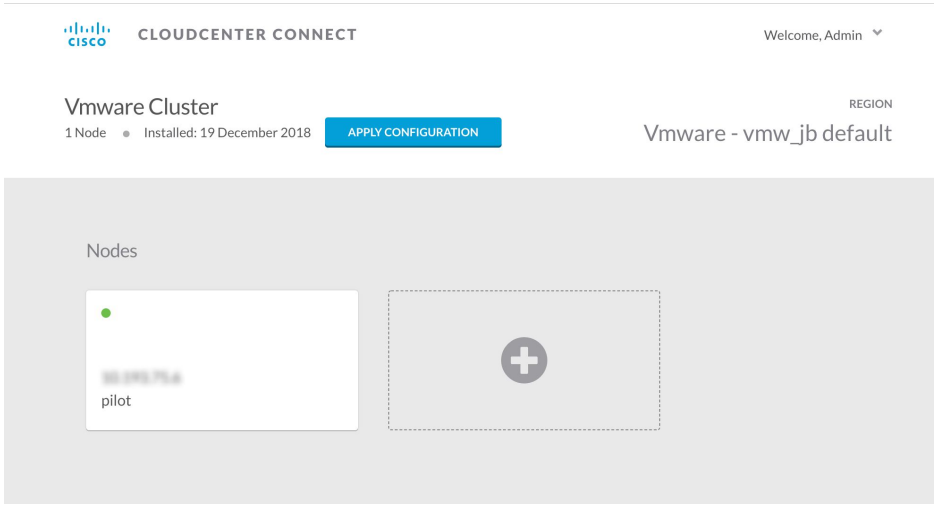

Switch your focus back to the Region Connectivity section of the target cloud region in the CloudCenter Suite web UI. The status indicator in the Region Connectivity section header will change from Not Configured to Running once connectivity between Cloud Remote and the CloudCenter Suite cluster is completely established (see figure below).

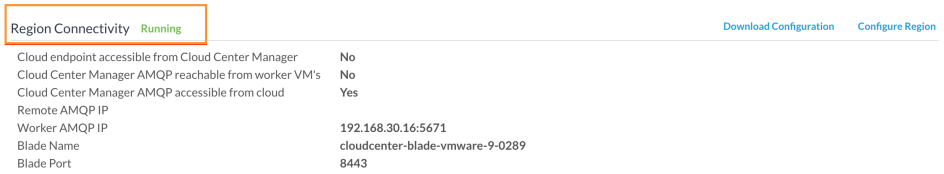

After completing these steps, Workload Manager and Cost Optimizer can use Cloud Remote for communicating with the target cloud region.

5. Instance Types: A Kubernetes cloud region does not include any instance type out-of-box. You must manually add instance types to your Kubernetes cloud if you want Workload Manager to deploy jobs to it.

Instance Types Settings

An instance type is a performance category of a VM for VM-based clouds or a performance category of a container for container-based clouds. The parameters for Kubernetes cloud instance type are as follows:

- Display Name
- Price per hour
- Instance Type ID
- CPU in milliCPUs
- RAM (MB)

You can sync, edit, and add instance types as shown below:

- Add
- Edit

You can add new instance types for Kubernetes clouds (by clicking the**Add Instance Types**link in the upper right of the section) and these instance types will control the parameters of the corresponding VMs or containers launched as part of your application. There is no concept of instance type native to Kubernetes, therefore the instance ID of the instance types you create in CloudCenter Suite can be arbitrary and the CPU and RAM parameters you specified in CloudCenter Suite will be used in the deployment.

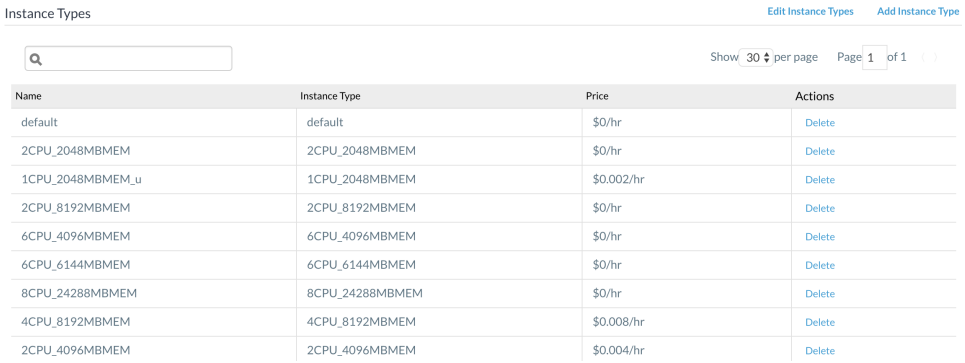

You must specify the price details in the**Price**field for Kubernetes cloudwhen adding instance types. The price is used to compute resource costs, which is displayed as invoice costs in theCost Optimizer Dashboard. The**Price**column (highlighted in red in the above screenshot) displays the cost of an instance.

You can edit the fields of an instance type (except instance type ID) by clicking the**Edit**link in the Actions column in the list of instance types. For Kubernetes cloud changing the values of parameters such as CPU and RAM will impact the properties of the VMs or containers deployed in those clouds.

Add a Kubernetes Cloud Account

Prerequisites

> Be aware that these screenshots may change based on the Kubernetes container changes. They are provided in this section as a point of reference.

Before adding a cloud account to a Kubernetes cloud in CloudCenter Suite, verify the following Kubernetes requirements:

- A valid Kubernetes service account.
- A**cluster-admin**cluster role binding exists on the API server (see theKubernetes Documentation).

- A valid**Service Account Token**. You can retrieve the Service Account Token from Kubernetes using one of two methods:
  - *Kubernetes Dashboard Method*:
    1. Access theKubernetes web UIand scroll the left menu bar down to Config and Storage and click**Secrets**. The list of secrets for the cluster is shown on the right panel:

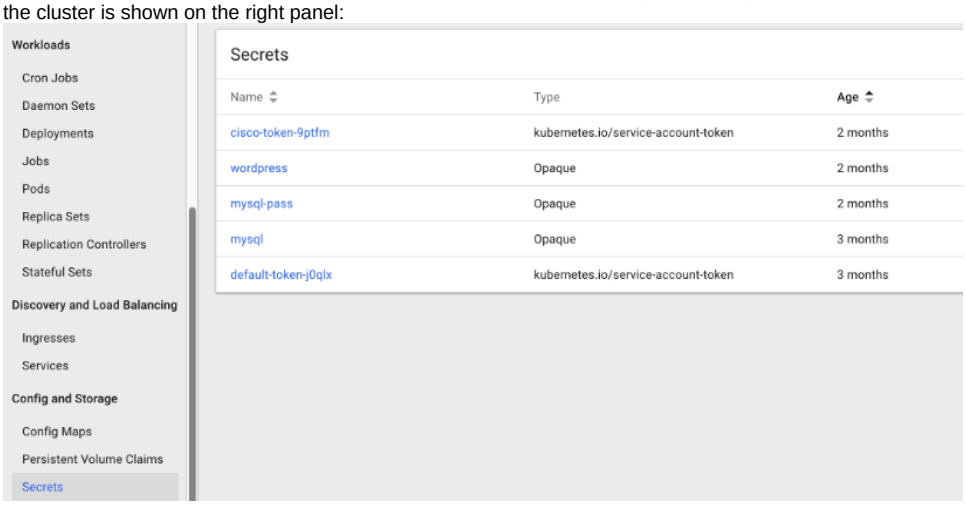

    2. Click the link corresponding to the**Service Account Token**to view the token details screen:

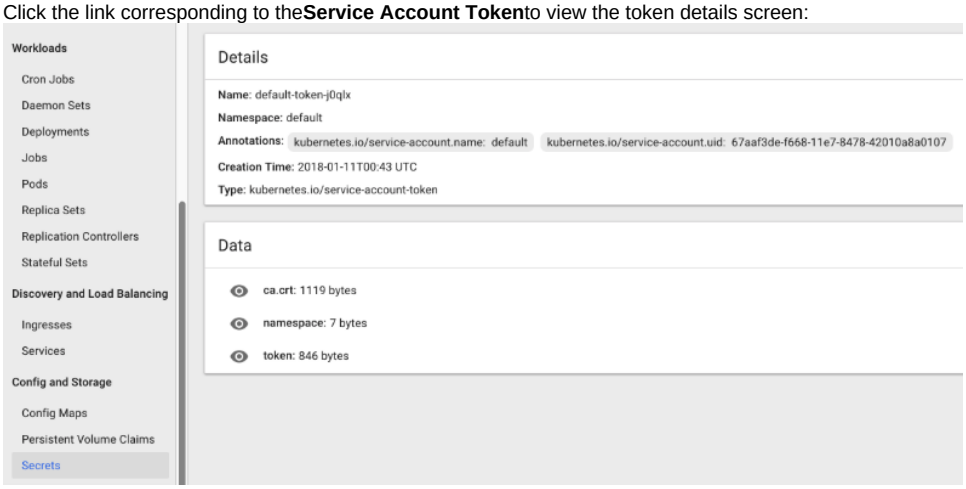

    3. Click the eyeball icon to the left of the token at the end of the Data section to reveal the token. Copy and paste to the**Service Account Token**field in the CloudCenter Suite's Add Cloud Account dialog box (see Configuration Process below).

> The service account token must be in base64 format before pasting into the Add Cloud Accounts page. Retrieving the token form the Kubernetes Web UI assures this to be true.

  - *The**kubectl**Command Method*:
    1. Issue the following commands in sequence the last command returns the token.

```
export NAMESPACE="default"

export SERVICE_ACCOUNT_NAME="bob-the-bot3"

kubectl create serviceaccount $SERVICE_ACCOUNT_NAME -n $NAMESPACE
serviceaccount "bob-the-bot3" created

kubectl create clusterrolebinding <name> --clusterrole=cluster-admin --
serviceaccount=$NAMESPACE:$SERVICE_ACCOUNT_NAME

export SECRET_NAME=$(kubectl get serviceaccount $SERVICE_ACCOUNT_NAME -n $NAMESPACE -o
'jsonpath={.secrets[0].name}' 2>/dev/null)

kubectl get secret $SECRET_NAME -n $NAMESPACE -o "jsonpath={.data.token}" | openssl enc -d -
base64 -
```

2. Copy and paste this token to the**Service Account Token**field in theCloudCenter Suite's Add Cloud Account dialog box (see Configuration Process below).

Configuration Process

To add a cloud account a Kubernetes cloud, follow this procedure.

1. Locate the Kubernetes cloud in the Clouds page and click the**Add Cloud Account**link. This displays the **Add Cloud Account** dialog box as shown in the figure below.

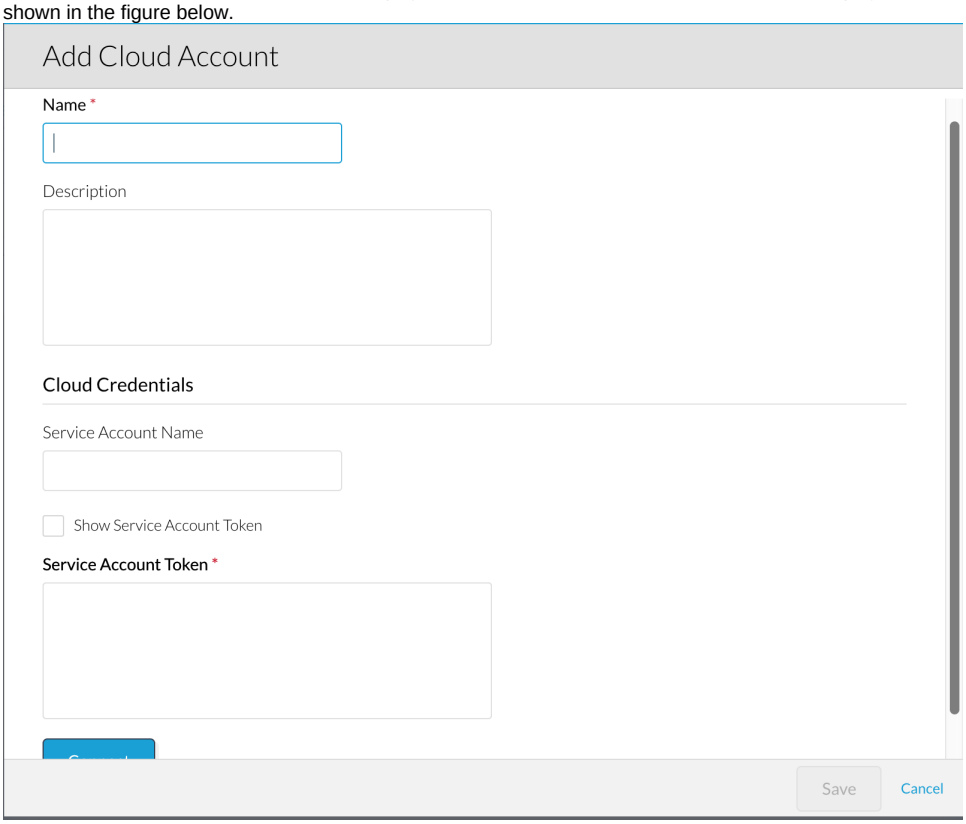

2. Assign a new cloud account name.

> **Tip**
>
> The name should not contain any space, dash, or special characters.

3. Add the following Cloud Credentials:

Field	Description
Service Account Name	The email address or username that you used to login to the Kubernetes cluster.
Service Account Token	The token used to access the Kubernetes service account as specified in the *Prerequisites* section above.

4. When done, click**Connect**.CloudCenter Suite will now attempt to validate your account credentials.
5. After the credentials are verified, the**Connect**button changes to an**Edit**button and two new fields appear**Enable Account For**and**Enable Reporting By Org Structure**,

   a. Set the**Enable Account For**dropdown per the table below.

Value	Usage
Provisioning	Workload Manager can deploy jobs using this account.
Reporting	Cost Optimizer and Workload Manager will track cloud costs for this account. Typical usage: master cloud accounts that are used for billing aggregation.
Provisioning, Reporting	Default. Account is used for both provisioning and reporting.

   b. **For AWS and Google clouds only**: Set the**Enable Reporting By Org Structure**toggle to On to cause Cost Optimizer to import the cost hierarchy createdin the cloud provider portal. This saves the time of manually creating a comparable cost hierarchy within Cost Optimizer. SeeCost Groups Configurationfor more information on cost hierarchies in Cost Optimizer.
   c. Click the**Save**button when done.

Cloud Accounts Tab

After you add cloud accounts to a cloud, they will appear in the Accounts tab for the cloud as shown in the figure below.

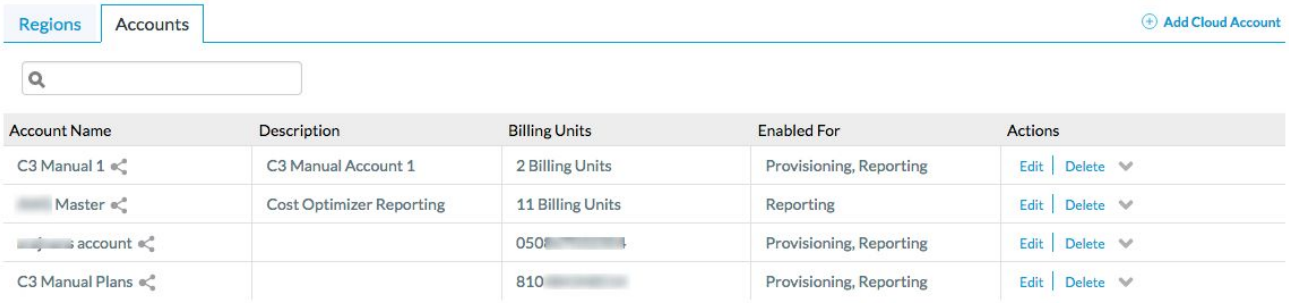

The Accounts tab contains columns for data entered when creating an account: Account Name, Description, Enabled For; and two additional columns:**Billing Units**and**Actions**.**Billing Units**is a dual function:

- If the cloud account contains only one billing unit, the ID for that billing unit is displayed.
- If the cloud account contains multiple billing units, such as an AWS master account, the number of billing units in that account is displayed followed by the text*Billing Units*.

A billing unit is the most granular level of cloud cost recording in CloudCenter Suite. The definition of a billing unit varies by a cloud provider as shown in the table below.

Cloud Provider	Billing Unit
AWS	Account ID
AzureRM	Subscription ID
Google	Project ID
IBM Cloud	Account ID
vCenter	Cloud Group Prefix - Datacenter Name
vCD	Organization Name
OpenStack	Project ID
Kubernetes	Namespace UID

The last column,**Actions**, contains links to let you edit or deleted the cloud account, ormanage instance typesfor the cloud account.

Configure a vCD Cloud

Configure a vCDCloud

Configuring a vCD cloud is a four-step process:

- Add a vCD Cloud
- Configure a vCD Region
- Add a vCD Cloud Account

Add a vCD Cloud

To add a vCD cloud follow these steps.

1. Navigate to**Admin > Clouds**. This brings you to the Clouds page. If you, or another tenant admin in your tenant, have already added clouds to your tenant, they will be listed here.
2. Click the**Add Cloud link**in the upper right. The **Add Cloud** dialog box is displayed.
3. Enter the**Cloud Name,**select the**cloud provider**, thenclick**Next**.

> When assigning the Cloud Name, use a unique name. Ensure that you do not assign a previously-used name (including deleted cloud names) as the Cloud Name as they may cause deployments to fail intermittently.

The second page of the **Add Clouds** dialog box, **Connectivity Settings**, appears. Set the toggle switches to configure the Cloud Connectivity settings.

- When adding a private VM cloud in the Workload Manager or Cost Optimizer UI, the second page of the Add Clouds dialog box, Connectivity Settings, appears with two toggles displayed:
  - **Worker VMs Directly Connect with CloudCenter Suite**
  - **VMs Directly Connect with CloudCenter Suite**
- Setting either of these toggles to No implies you will install Cloud Remote for each region of this cloud. This also causes a third toggle to appear:**CloudCenter Suite Directly Accessible from Cloud Remote**.
- Follow the table below for guidance on setting these toggles.

Toggle settings	Use case	Network Diagram
Cloud Endpoint Directly Accessible = Yes AND VMs Directly Connect with CloudCenter Suite = Yes	CloudCenter Suite cluster can initiate a connection to the cloud region API endpoint AND Worker VMs can initiate a connection to the CloudCenter Suite cluster Cloud Remote is not required	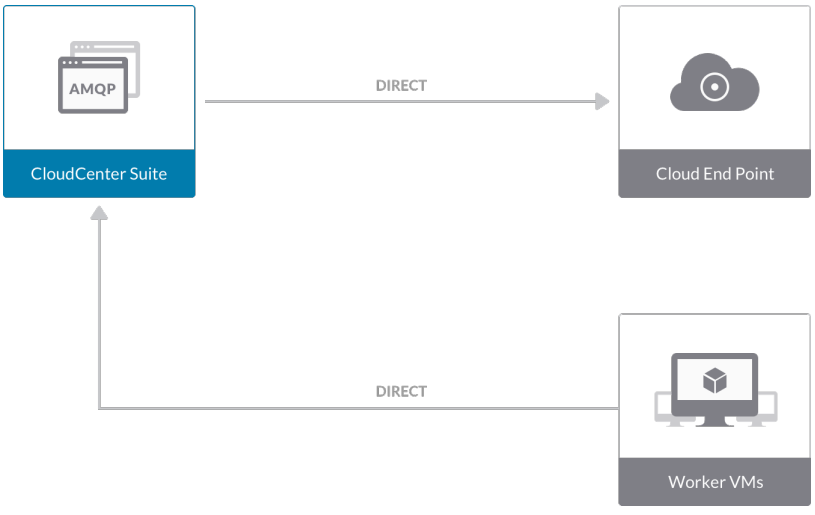

Cloud Endpoint Directly Accessible = No AND Worker VMs Directly Connect with CloudCenter Suite = No AND CloudCenter Suite Directly Accessible fromCloud Remote = Yes	CloudCenter Suite cluster cannot initiate a connection to the cloud region API endpoint AND Worker VMs cannot initiate a connection to the CloudCenter Suite cluster AND Cloud Remote can initiate the connection to the CloudCenter Suite cluster	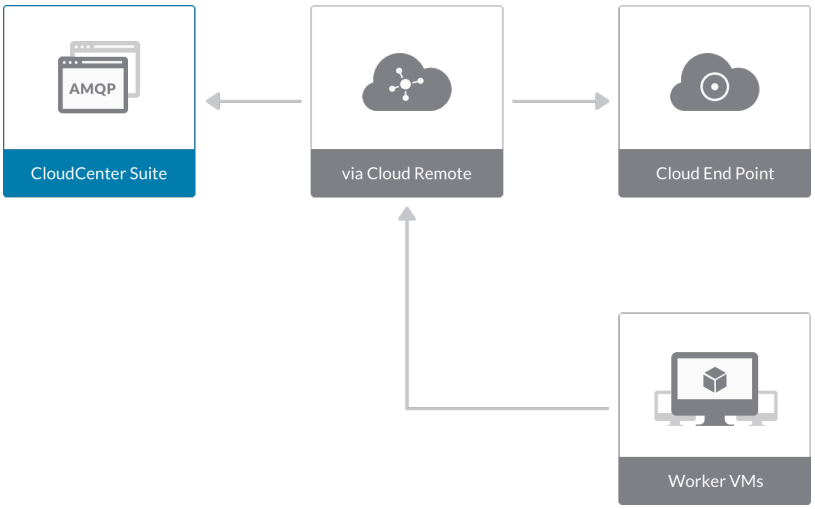
Cloud Endpoint Directly Accessible = No AND Worker VMs Directly Connect withCloudCenter Suite = No AND CloudCenter Suite Directly Accessible from Cloud Remote = No	CloudCenter Suite cluster cannot initiate a connection to the cloud region API endpoint AND Worker VMs cannot initiate a connection to the CloudCenter Suite cluster AND Cloud Remote cannot initiate the connection to the CloudCenter Suite cluster	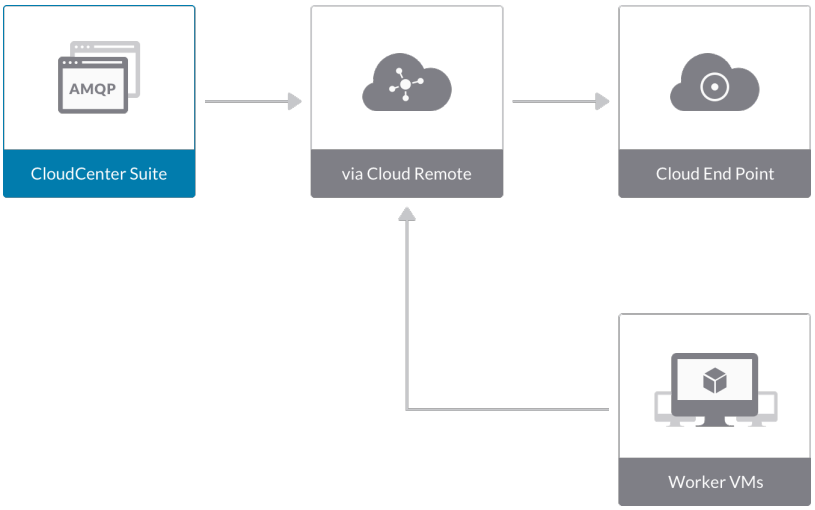
Cloud Endpoint Directly Accessible = Yes AND Worker VMs Directly Connect withCloudCenter Suite = No AND CloudCenter Suite Directly Accessible from Cloud Remote = No	CloudCenter Suite cluster can initiate a connection to the cloud region API endpoint AND Worker VMs cannot initiate a connection to the CloudCenter Suite cluster AND Cloud Remote cannot initiate the connection to the CloudCenter Suite cluster	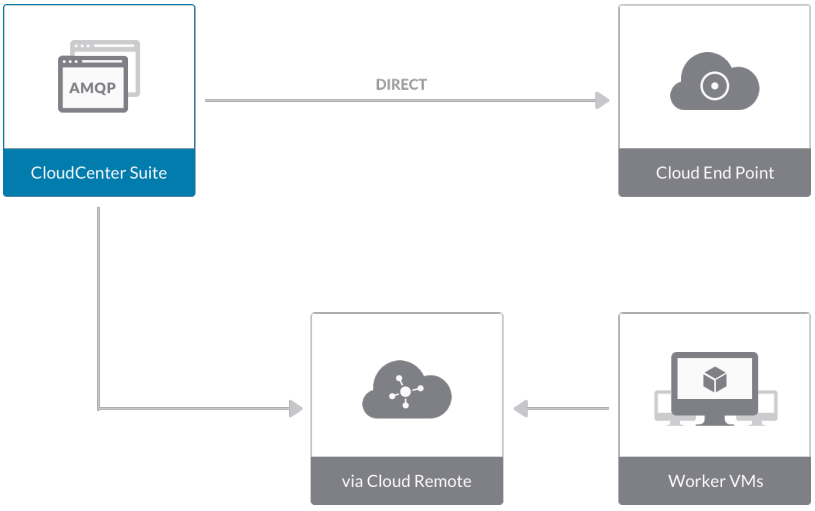

Cloud Endpoint Directly Accessible = Yes AND Worker VMs Directly Connect withCloudCenter Suite = No AND CloudCenter Suite Directly Accessible fromCloud Remote = Yes	CloudCenter Suite cluster can initiate a connection to the cloud region API endpoint AND Worker VMs need to communicate to the CloudCenter Suite cluster through Cloud Remote AND Cloud Remote can initiate the connection to the CloudCenter Suite cluster	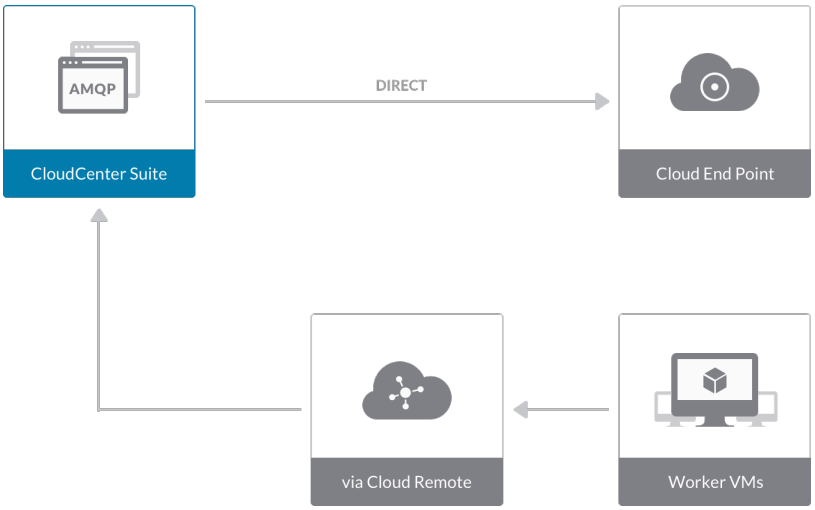

Note

The connectivity toggle settings set at the cloud level are inherited by each region you add to this cloud. However, it is possible to override these toggle settings on a per-region basis from the Regions tab for each region.

4. Click**Done**to save the configuration and close the dialog box. This brings you back to the Clouds page and the cloud you just created will be added to the bottom of the list on the left side of the page.

Configure a vCD Region

A vCD cloud has one region that you configure from the vCD cloud Details tab. Follow this procedure.

1. Navigate to Clouds page:**Admin > Clouds**. Find your newly created vCD cloud from the cloud list on the left half of the screen and click its**Configure Cloud**link. This displays the Details tab for this cloud as shown in the figure below.

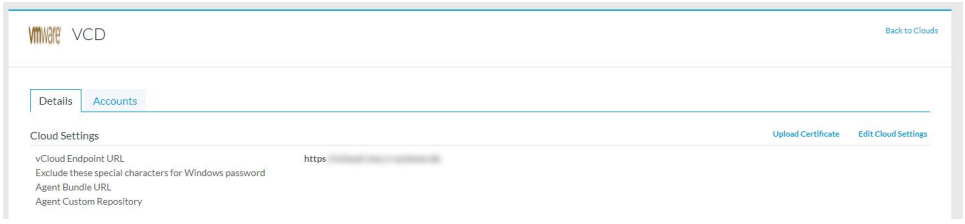

2. Upload a TLS certificate to the vCD system by clicking the**Upload Certificate**link and then using the dialog box to select a file from your PC.
3. Click**Edit Cloud Settings**to open the **Configure Cloud Settings** dialog box. The Cloud Settings section contains fields that are unique to the vCD cloud family and settings that are common to all cloud families. Adjust these field values per the instructions in the following tables.

vCD Specific Cloud Settings

Field	Usage
vCD API Endpoint	Address used by Workload Manager to deploy and manage deployment in the vCD cloud

Cloud Agnostic Cloud Settings

Field	Usage
Exclude these special characters for Windows password	When the Workload Manager agent is installed on a Windows worker VM, a special user account, called cliqruser, is created to support RDP sessions that may be initiated by the user through the Workload Manager UI. A Workload Manager process running on the CloudCenter Suite cluster creates a random password and passes it to the agent for creating the cliqruser account. Because some Windows deployments may restrict using certain characters for Windows passwords, this field is provided to tell the Workload Manager to exclude these special characters in the generation of the password for the cliqruser account.
Agent Bundle URL	If you plan to use a local repository to host the bundle store, you need to enter the URL of the local bundle store here. Otherwise, leave blank.

Agent Custom Repository	If you plan to use a local repository to host the package store, you need to enter the URL of the local package store here. Otherwise, leave blank.

When you are done editing the settings in the dialog box, click**Save**.

4. Determine if you need Cloud Remote for this region. Scroll down to the Region Connectivity section for the region and click on the**Configure Region**link in the upper right to open the Configure Region dialog box.The toggle settings should be the same as when you set them on the connectivity page of the Add Cloud dialog box.If all of the connectivity togglesinthe Region Connectivity dialog boxare set to Yes, then Cloud Remote is NOT needed for this cloud region. In this case,you would normally leave the region connectivity settings at their current values and continue to the next settings section.

   The exceptionto this guidance is when a NAT firewall or proxy server exists between theCloudCenter Suite management cluster andworker VMs, or between the CloudCenter Suite management cluster and users that would use Workload Manager to initiate a Guacamole remote connection to a worker VM. In either of these cases, override the address fields in the Region Connectivity dialog box as explained below.

Networking Constraint	Field	Value
Worker VMs must use a proxy server or NAT firewall to access the"local" AMQP server running in the CloudCenter Suite cluster.	Worker AMQP IP Address	IP address and port number that the firewall or proxy server presents to the worker VMs on behalf of the "local" AMQP server running in the CloudCenter Suite cluster.
Users must use a proxy server or NAT firewall to access theGuacamole server running in the CloudCenter Suite cluster.	Guacamole Public IP Address and Port	IP address and port number that the firewall or proxy server presents to users on behalf of theGuacamole server running in the CloudCenter Suite cluster.
Worker VMs must use a proxy server or NAT firewall to access theGuacamole server running in the CloudCenter Suite cluster.	Guacamole IP Address and Port for Application VMs	IP address and port number that the firewall or proxy server presents to the worker VMs on behalf of the Guacamole server running in the CloudCenter Suite cluster.

5. Click**OK**to save the changes and dismiss the dialog box. You can now proceed to the next region settings section:VM Naming and IPAM Strategy.
6. If any of theconnectivity toggles inthe Region Connectivity dialog box are set to No, then you must install and configure Cloud Remote for this region.

Configure Cloud Remote in a vCD Region

Configure Cloud Remote in a vCD region as follows.

> ⚠ Since CloudCenter Suite does not include a prebuilt appliance for Cloud Remote for vCD, the following procedure includes steps to build the Cloud Remote appliance from the Cisco-supplied Cloud Remote installer file.

Launch Cloud Remote Built from the Installer File

a. Launch a Centos 7 instance, ensure the prerequisites are installed, and run the Cloud Remote installer file:

   Unable to render {include} The included page could not be found.

b. Optional but recommended for production environments: Repeat the step above twice to create two additional instances of the appliance to be used to form a cluster for HA.Cloud Remote includes support for the clustering of multiple nodes. You will "add" these two additional instances to the first instance after the first instance is configured. SeeCloud Remote (Conditional)> *Scaling*for details.

Unable to render {include} The included page could not be found.

Specify AMQP and Guacamole Addresses for Supporting Cloud Remote

From the CloudCenter Suite UI, for the cloud region requiring Cloud Remote, navigate to the corresponding Regions or Details tab. Click the**Configure Region**link in the upper left of the Region Connectivity section to bring up the Configure Region dialog box. The toggle settings should be the same as when you set them on the connectivity page of the Add Cloud dialog box. You must update some of the address fields in the dialog box according to the scenarios summarized in the table below.

Toggle Settings	Field	Value

Worker VMs Directly Connect with CloudCenter = No AND CloudCenter Directly Accessible from Cloud Remote = Yes	Local AMQP IP Address	Pre-populated with the address and port number of the "local" AMQP server running in the CloudCenter Suite cluster. This address must be **accessible to Cloud Remote** . If Cloud Remote is accessing the CloudCenter Suite cluster through a user-supplied proxy server or NAT firewall, overwrite this field with the corresponding local AMQP IP address and port number provided by the user-supplied proxy server or NAT firewall and accessible to Cloud Remote .
Worker VMs Directly Connect with CloudCenter = No AND CloudCenter Directly Accessible from Cloud Remote = No	Remote AMQP IP Address	Enter <Cloud_Remote_IP>:<amqp_port>, where <Cloud_Remote_IP> = the IP address Cloud Remote which is **accessible to the CloudCenter Suite cluster**, and <amqp_port> = 5671 OR the custom AMQP port number you would later set with the **Change Ports shell script** on the Cloud Remote appliance (see Cloud Remote (Conditional) > *Custom Port Numbers (Conditional)*). If there is no user-supplied NAT firewall or proxy server between the CloudCenter Suite cluster and Cloud Remote , the IP address would be the public IP address of Cloud Remote . If there is a NAT firewall or proxy server between the CloudCenter Suite cluster and Cloud Remote , instead, enter the corresponding public IP address and port number that the firewall or proxy server presents to the internet on behalf of the "remote" AMQP server running in Cloud Remote .
Worker VMs Directly Connect with CloudCenter = No	Worker AMQP IP Address	Enter <Cloud_Remote_IP>:<amqp_port>, where <Cloud_Remote_IP> = the Cloud Remote IP address **accessible to the worker VMs**, and <amqp_port> = 5671 OR the custom AMQP port number you would later set with the **Change Ports shell script** on the Cloud Remote appliance (see Cloud Remote (Conditional) > *Custom Port Numbers (Conditional)*).
Worker VMs Directly Connect with CloudCenter = No	Guacamole Public IP and Port	Enter <Cloud_Remote_IP>:<guac_port>, where <Cloud_Remote_IP> = the Cloud Remote IP address **accessible to CloudCenter Suite users**, and <guac_port> = 8443 OR the custom Guacamole port number you would later set with the **Change Ports shell script** on the Cloud Remote appliance (see Cloud Remote (Conditional) > *Custom Port Numbers (Conditional)*).
Worker VMs Directly Connect with CloudCenter = No	Guacamole IP Address and Port for Application VMs	Enter <Cloud_Remote_IP>:<guac_port>, where <Cloud_Remote_IP> = the Cloud Remote IP address **accessible to worker VMs**, and <guac_port> = 7789

When done, click **OK** to save the setting and dismiss the dialog box.

Download Region Connectivity Settings and Upload to Cloud Remote

Cloud Remote uses the region connectivity settings set in the Workload Manager or Cost Optimizer UI. After saving the Region Configuration settings in the Workload Manager or Cost Optimizer UI, you must download them and to your local computer and then upload them to Cloud Remote as follows.

Click the**Download Configuration**link in the upper right of the Region Connectivity section, as shown in the figure below.

Clicking DownloadConfiguration causestwo things to happen:

- An encrypted zip file named**artifacts.zip**will be downloaded by your browser. Make note of the location of this zip file as you will need to upload it to Cloud Remote through the Cloud Remote web UI (see below).
- The Region Connectivity section header updates to display a **Copy Encryption Key**link, as shown in the figure below.

Click the**Copy Encryption Key**link to save the key to your clipboard. A success message will be displayed temporarily above the Region Connectivity section header. Make sure not to overwrite the clipboard with other data. You will need the key when you upload the configuration zip file to Cloud Remote .

> If you change the connectivity settings in the CloudCenter Suite UI and need to again download the zip file, a new encryption key is automatically created and can be copied to the clipboard by clicking the **Copy Encryption Key** link again.

After you have downloaded the zip file and copied the encryption key to your clipboard, login toCloud Remote web UI.

a. Open another browser tab and login to https://<Cloud Remote _ip> with thedefault credentials: admin/cisco.
b. You will immediately be required to change your password. Do so now.
c. You are now brought to the Cloud Remote home page as shown in the figure below.

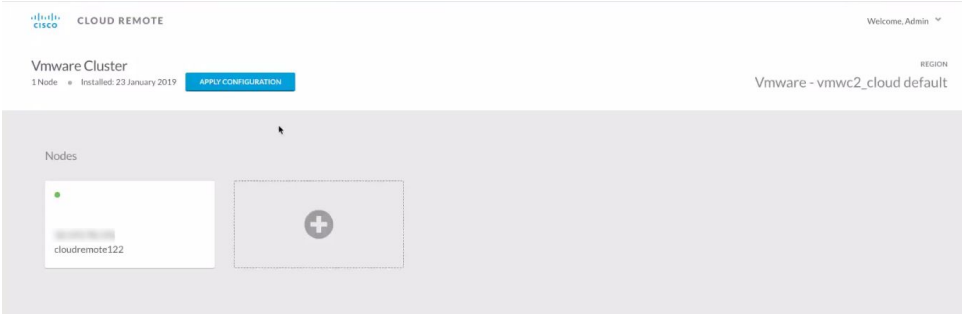

d. Click the **Apply Configuration** button in the page header. This prompts you to select a configuration file and enter the encryption key as shown in the figure below.

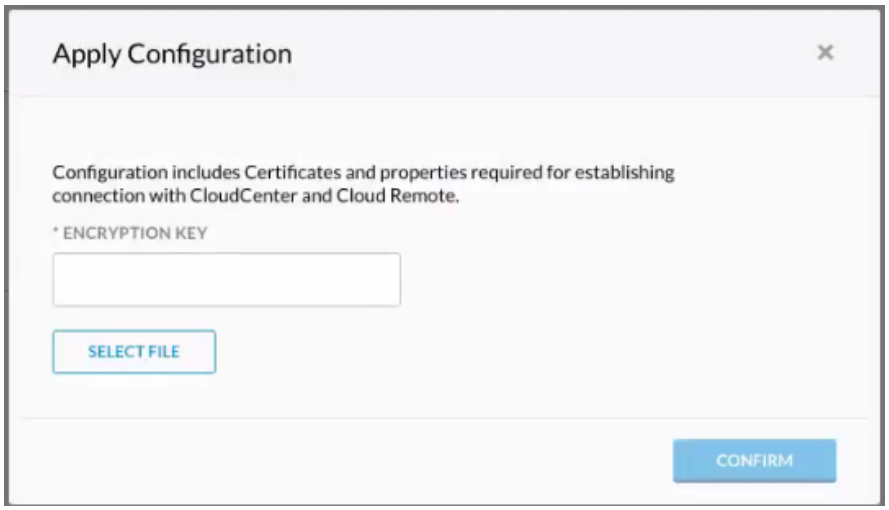

e. Paste the encryption key that was copied to the clipboard into the**Encryption Key**field in the dialog box.
f. Click**Select File**and browse tothe artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
g. Click**Confirm**.
h. Once the zip file is successfully transmitted and accepted, the Cloud Remote appliance attempts to establish communication with the CloudCenter Suite cluster and the Cloud Remote web UI home page is updated to show the name of the region it is connecting to in the upper right (see figure below).

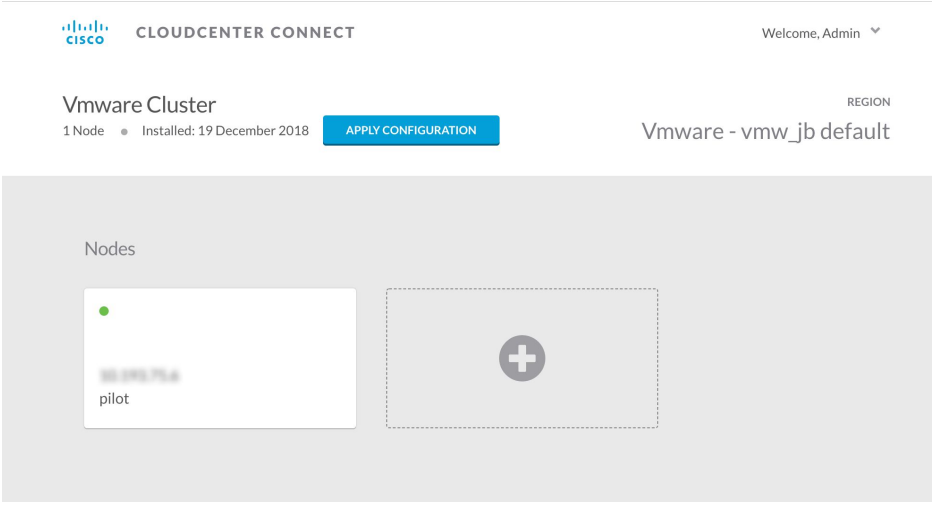

Switch your focus back to the Region Connectivity section of the target cloud region in the CloudCenter Suite web UI. The status indicator in the Region Connectivity section header will change from Not Configured to Running once connectivity between Cloud Remote and the CloudCenter Suite cluster is completely established (see figure below).

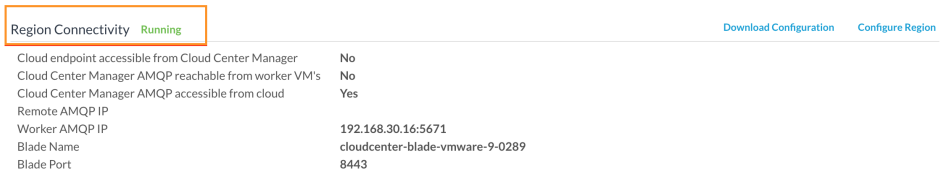

After completing these steps, Workload Manager and Cost Optimizer can use Cloud Remote for communicating with the target cloud region.

7. VM Naming and IPAM Strategy (conditional): Configure any VM naming or IPAM strategies in the Strategy section. If you leave the settings at the defaults, no IPAM strategy is applied and the default VM naming strategy is applied.

VM Naming and IPAM Strategies

Workload Manager lets you define a VM naming strategy and an IP address management (IPAM) strategy on a per cloud region basis. The VM naming strategy lets Workload Manager generate unique names for each of your VMs as they are launched. The IPAM strategy letsWorkl oad Manager generate IP addresses for each of the vNICs on each of your VMs. The VM naming strategy can reference a VM naming callout script, and the IPAM strategy requires that you specify an IP address allocation callout script andan IP address deallocation callout script. An IP address allocation callout scriptmay work in conjunction with a 3rd party IPAM server in your cloud environment.

Like External Lifecycle Actions, the scripts or commands associated with a VM naming strategy or IPAM strategy are executed in a containerized script execution engine running in the CloudCenter Suite cluster, if you are not using Cloud Remote in your region. If you are using Cloud Remote , these scripts or commands will execute in a containerized script execution engine running in Cloud Remote .

To set up the VM naming and IPAM strategies for your region, scroll down to the Strategy section of the Regions or Details tab and follow these steps.

Click the**Edit Strategy**link in the upper right to display the Edit Strategy dialog box as shown in the figure below.

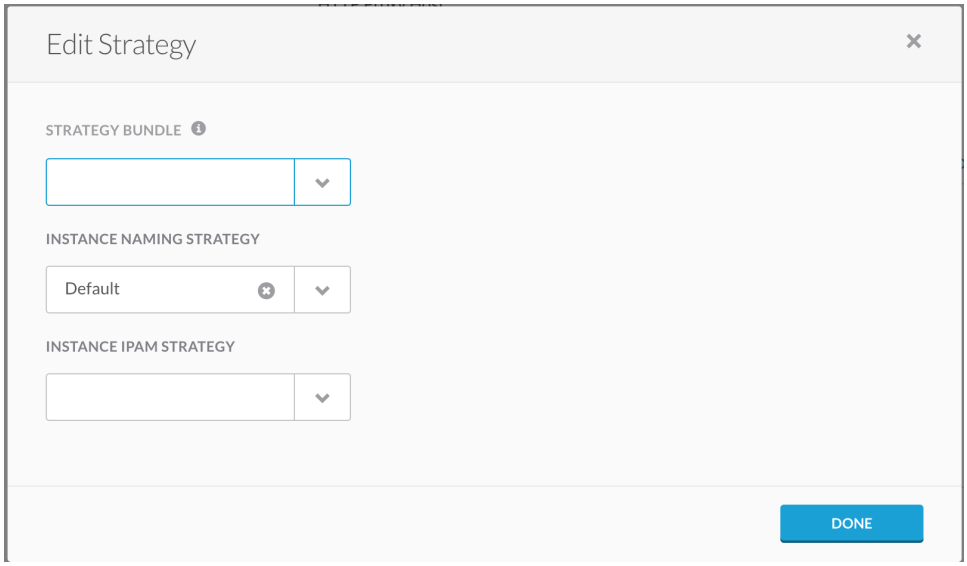

The dialog box can reveal a total of seven data entry fields as explained in the table below:

Field	Usage
Strategy Bundle	Zip file containing all of the scripts associated with your VM naming strategy or IPAM strategy. Select source from the dropdown an add the path in the new field that appears to the right. 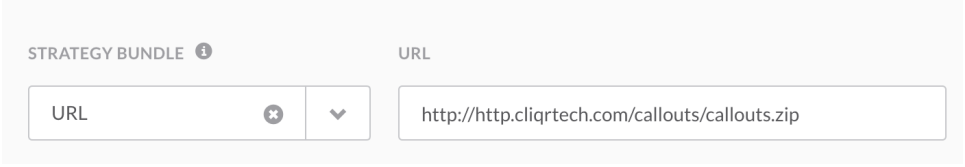

Instance Naming Strategy	Dropdown for selecting one of three naming strategies: Default, CliQr Macro Replacement, and Hostname Callout. If you select CliQ Macro Replacement, the **Node Name Config** field becomes visible and required. If you select Hostname Callout, the **Custom VM Name** field becomes visible and required. See the Instance Naming Strategy table below for details.

Node Name Config	Required if you selected "CliQr Macro Replacement" for theInstance Naming Strategyfield. Content: text string consisting of a combination of characters and numbers with embedded macro variables (see table below). The macros get translated to text strings run time.
	
	The resulting VM name string mustmeet the following requirements for job submission to succeed:
	
	• Only contain the following characters:
	  ◦ Lowercase a to z
	  ◦ 0 to 9
	  ◦ Hyphen
	• Start withan alpha character
	• Cannot end with a hyphen

Instance IPAM Strategy	Dropdown menu with two choices: No IPAM (default), and IPAM Callouts. If you select the latter, two additional fields are displayed and become required: **IPAM Alloc Rule**and**IPAM Dealloc Rule**. (See below).
Custom VM Name	Location of your custom VM naming script if you set the**Instance Naming Strategy**field to "Custom Callout". Location options inclu "Script from bundle" if you specify a zip file in the **Strategy Bundle** field. See Guidance for Callout Scripts for guidance on developi custom VM naming scripts.
IPAM Alloc Rule	Location of your custom IPAM address allocation script if you set the**Instance IPAM Strategy**field to "IPAM Callouts". Location opti include "Script from bundle" if you specify a zip file in the **Strategy Bundle** field. See Guidance for Callout Scripts for guidance on developing custom IPAM callout scripts.
IPAM Dealloc Rule	Similar to the**IPAM Alloc Rule**above except this is where you specify the location of yourcustom IPAM address deallocation script. Guidance for Callout Scripts for guidance on developing custom IPAM callout scripts.

Instance Naming Strategy Options:

Options	Description
Default	The default method provided by Cisco is the generation of a string consisting of A predefined token prefixed with **cqjw-**
CliQr Macro Replacement	Selecting this method causes the **Node Name Config** field to be displayed and to become a required field. See the fields table above for details.
Hostname Callout	Selecting this method causes the **Custom VM Name** field to be displayed and to become a required field. See the fields table above for details.

OOB macro variables forCliQr Macro Replacement:

Macro	Returned Value
%os1%	First character of OS type string in lower case
%os2%	Firsttwo characters of OS type string in lower case
%OS1%	Firstcharacter of OS type string in upper case
%OS2%	Firsttwo character of OS type string in upper case
%RND<number>%	Fixed length random string, e.g., %RND10%, 10-character long random string ⚠ This is a mandatory field, only values in the range: %RND6% to %RND11% or just %RND% which defaults to %RND6%"
%UID%	User ID
%VID%	Vendor ID

%VM_PREFIX%	VM name prefix must be added as a global parameter in the app profile. ⚠ Only macro whose value can be user-customized.

8. External Lifecycle Actions (conditional):Specify any external lifecycle actions to be performed on all VMs launched by Workload Manager in this region.

External Lifecycle Actions Settings

Whenever you use Workload Manager to deploy a VM-based application, Workload Manager can execute external actions on VMs in the deployment as the VMs progress through these five phases of their lifecycle:

- Pre VM start
- Pre VM initialization
- PostVM initialization
- Pre VM Stop
- Post VM Stop

These lifecycle actions will execute in a script execution engine running in a container in the CloudCenter Suite cluster if no Cloud Remote appliance is configured for the region. If you have configured Cloud Remote for this region, the external lifecycle actions will executein a script execution engine running in a container in the Cloud Remote appliance.

External Lifecycle Action scripts can be defined at the service level, the application level, and the cloud region level. Use the External Lifecycle Actions section of the Region tab or Details tab (see figure below) tospecify external lifecycle actions at the region level.

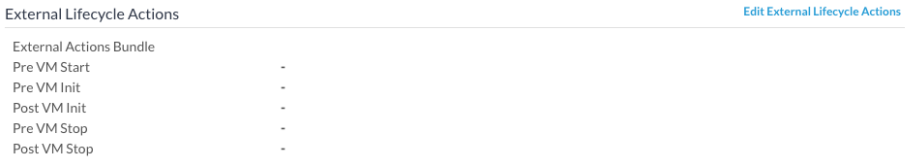

Click on the**Edit External Lifecycle Actions**link in the upper right. This displays theEdit External Lifecycle Actions dialog box as shown below.

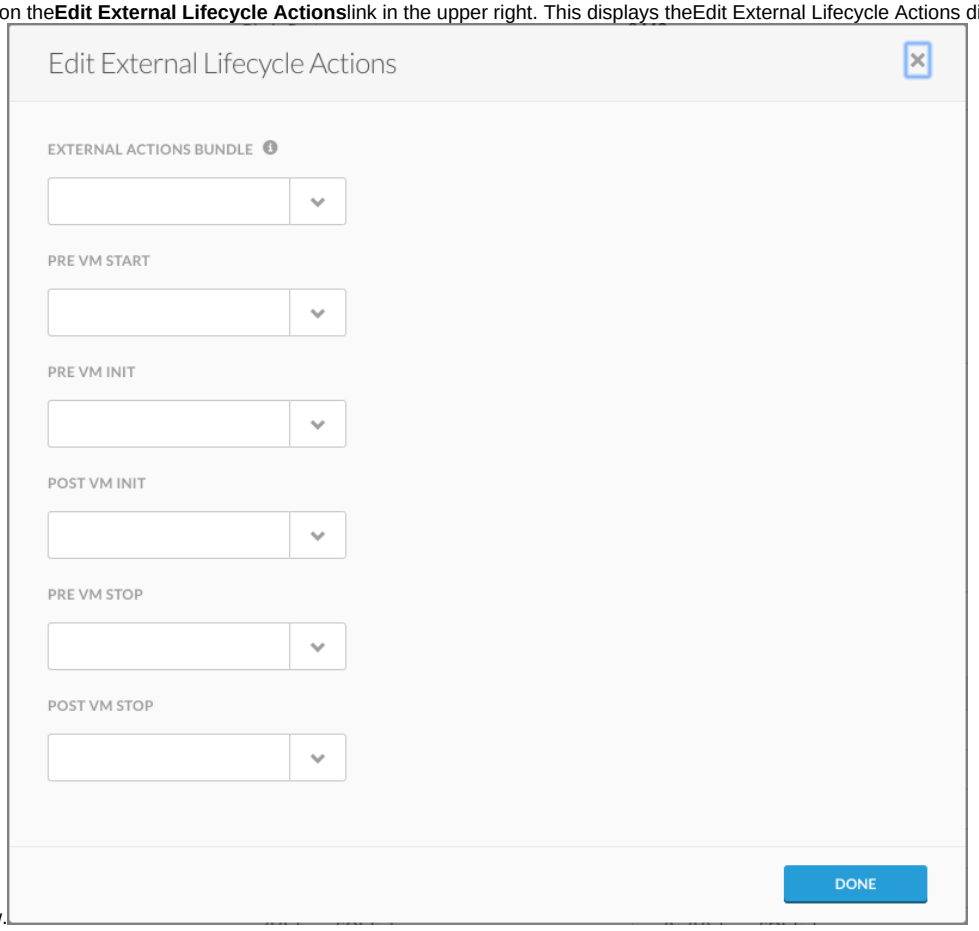

One way to specify an external lifecycle action is to use the first field,**External Actions Bundle**, to specify a zip file that contains all scripts for external lifecycle management. This zip file needs to contain a directory that contains all the scripts. Use the dropdown to select a repository location (URL, FTP, S3, etc.). This causes an additional field to be displayed to the right for entering the path to the zip file as shown in the figurebelow.

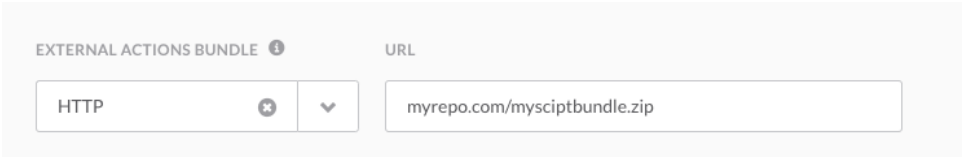

Once you specify actions bundle, for each of the individual lifecycle stage fields, select "Script from Bundle" from the dropdown. This displays an additional field to the right where you must specify the name of the script file in the bundle as shown in the figure below.

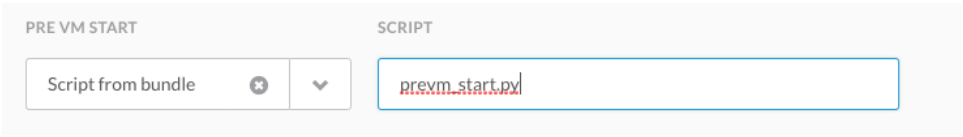

The scripts you specify can be bash scripts or Python scripts called from within a bash script.For more information on lifecycle scripts, seeDeployment Lifecycle Scripts.

9. Instance Types (conditional): A vCD cloud region includes one "default" instance type with 1 vCPU, 1 vNIC, 1024 MB RAM, and no additional disk storage. CloudCenter Suite will also automatically create instance types based on the parameters of VMs you deploy from within vCD. You would manually add more instance types to your vCD region if you want Workload Manager to deploy jobs to this region with differently sized instance types.

Instance Types Settings

An instance type is a performance category of a VM for VM-based clouds or a performance category of a container for container-based clouds. You can sync, edit, and add instance types depending on the cloud provider. For public clouds, CloudCenter Suite auto-populates the instance types in the Instance Types section from the cloud provider. CloudCenter Suite auto-syncs this information from the cloud provider once per 24 hours. For these clouds, you cannot force sync on demand.

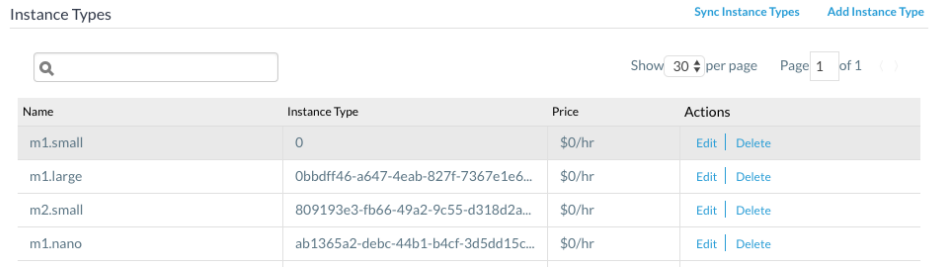

You can edit the fields of an instance type (except instance type ID) by clicking the**Edit**link in the Actions column in the list of instance types. The edits to these parameters are ignored during application deployment. The only parameter that will impact CloudCenter Suite behavior for public clouds is the price per hour, which is used by Workload Manager and Cost Optimizer for deployment cost calculations.

Historical Price Management

Sometimes, you want to update the historical cost to correct or update cost computation. In such cases, thecost of the affected resourcesis computed from the date and time the price is updated.You can update the prices of the instances types for the preceding three months only.

To update the price, follow these steps:

a. Click**Edit Instance Type**. The**Edit Instance Types**page appears.
b. In the**Price/Hour**column, click the**Pricing History**icon. The**EditPrice History**dialog appears.

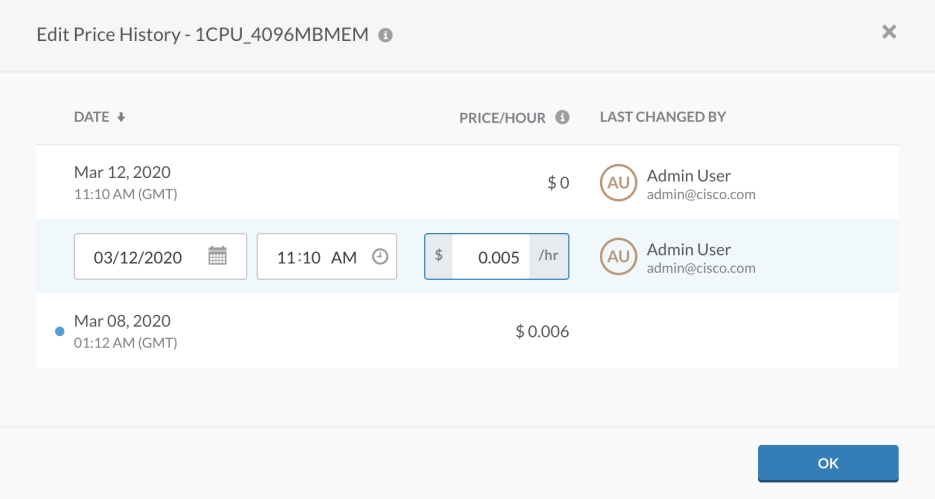

c. Specify a date and time in the appropriate fields and enter the desired price inthe**Price/Hour**field, enter the desired price.
d. Click**OK**, followed by**Done**(in the**Edit Instance Types**page) to save the changes.

After the price has been updated, the cost is recomputed for allresources(VMs and deployments, as applicable)affected by the price change. This could take a few minutes. When the recomputation is successful, you receive a notification that can be viewed in the Notification icon in the Cost Optimizer header (seeCost Optimizer Dashboard), as shown in the following sample screenshot.

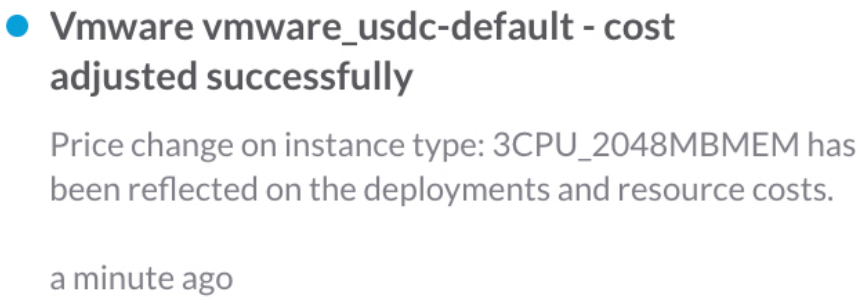

Simultaneously,the updated accruedcostof a VMis displayed on theVirtual Machinespage and the updated deployment cost is displayed on the Deployment Details page.

10. Storage Types (conditional): For private VM-based clouds like vCD, CloudCenter Suite uses storage types for cost tracking purposes. CloudCenter Suite creates a default storage type with zero cost.You would manually edit this storage type to enter your own cost factor. You can optionally add more storage types to your vCD region.

Storage Types Settings

A storage type is a categorization of storage that can be provisioned for an application during deployment. It is a construct that applies to VM-based clouds only. Storage types are displayed in the Storage Types section of the Regions or Details tab. The storage type construct contains the following fields for all VM-based clouds:

- Display Name
- Cost Per Month per GB
- Cloud Storage Type ID
- Minimum Volume Size (GB)
- Maximum Volume Size (GB)
- Maximum Throughput (MB/s)
- Maximum IOPS
- Provisioned IOPS
- IOPS Cost Per Month

Storage types are handled by CloudCenter Suite for various cloud providers in a fashion similar to instance types. CloudCenter Suite auto-populates the list of storage types in the Storage Types section based on information from the cloud provider. This data is stored in thepackage store and the package store auto-syncs this information from the cloud provider once per 24 hours. Storage types can be edited. From the list of storage types in the Storage Type section (see figure below), click the Edit link (in the Actions column) for that storage type.

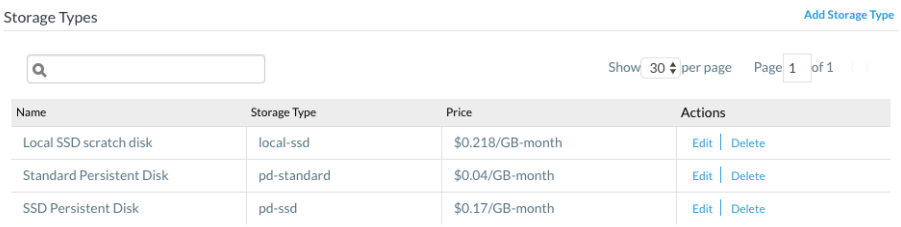

This brings up an Edit Storage Type dialog box (similar to the Add Storage Type dialog box) where you can edit all nine fields except the**Cloud Storage Type ID**field. For AzureRM and Google clouds, only changes to the pricing fields change the behavior of CloudCenter Suite: the cloud cost calculations are modified based on the new unit cost data. All other field changes are ignored.

For all cloud providers,**Minimum Volume Size**and**Maximum Volume Size**restrict the volume size you can specify at deploy time. For public cloud providers that support provisioned IOPS, the user can specify the guaranteed IOPS at deploy time.

Historical Price Management

Sometimes, you want to update the historical cost to correct or update cost computation. In such cases, the cost of the affected resources is computed from the date and time the price is updated.You can update the prices of the older storage types for the preceding three months only.

To update the price, follow these steps:

a. Click**Edit Instance Type**. The**Edit Instance Types**page appears.
b. In the**Price/Hour**column, click the**Pricing History**icon. The**EditPrice History**dialog appears.

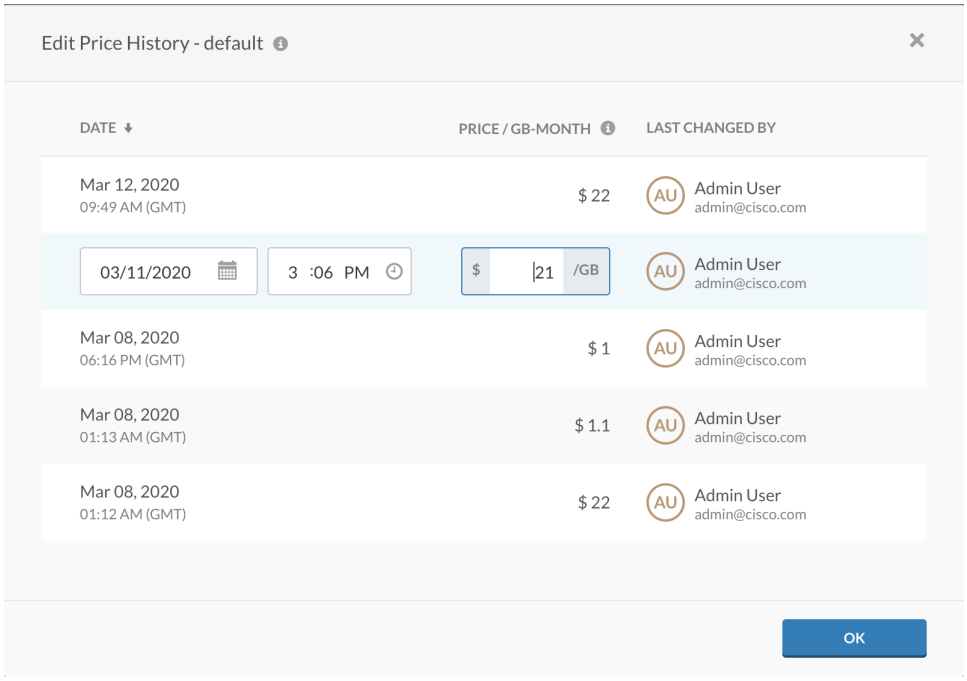

c. Specify a date and time in the appropriate fields and enter the desired price inthe**Price/Hour**field, enter the desired price.
d. Click**OK**, followed by**Done**(in the**Edit Instance Types**page) to save the changes.

After the price has been updated, the cost is recomputed for allresources(VMs and deployments, as applicable)affected by the price change. This could take a few minutes. When the recomputation is successful, you receive a notification that can be viewed in the Notification icon (seeCost Optimizer Dashboard) in the header.Simultaneously,the updated accruedcostof a VMis displayed on theStorage Volumespage and the updated deployment cost is displayed on the Deployment Details page.

11. Image Mappings: Image mappings allow services based on Workload Manager logical images to be deployed using the appropriate physical image stored on the target cloud region. You must manually import these physical images into your vCD region and then map the appropriate Workload Manager logical images to these physical images. SeeImagesfor more context.

Add a vCD Cloud Account

Prerequisites

For Workload Manager to deploy jobs in vCD using a particular user account, that account must have the permissions identified in the table below.

vCD Object	Required Permission	Reason
Network	Assign Network	If the default network in a template/snapshot must be changed
Datastore	Allocate space	For persistent disk operation
	Browse datastore	
	Low-level file operations	
	Remove file	
Folder	Create folder	For user folder creation
Resource	Apply recommendation	For datastore cluster support
	Assign VM to resource pool	For resource pool selection
Tasks	Create task	For VM operation
	Update task	
Virtual Machine	All permissions	
Global Role	Set Custom Attributes	To add custom attributes on virtual machines
	Manage Custom Attributes	

Configuration Process

To add a vCD cloud account, follow this process:

1. Locate the vCD cloud in the Clouds page and click**Add Cloud Account**button. This will display the **Add Cloud Account** dialog box as shown in the figure below.

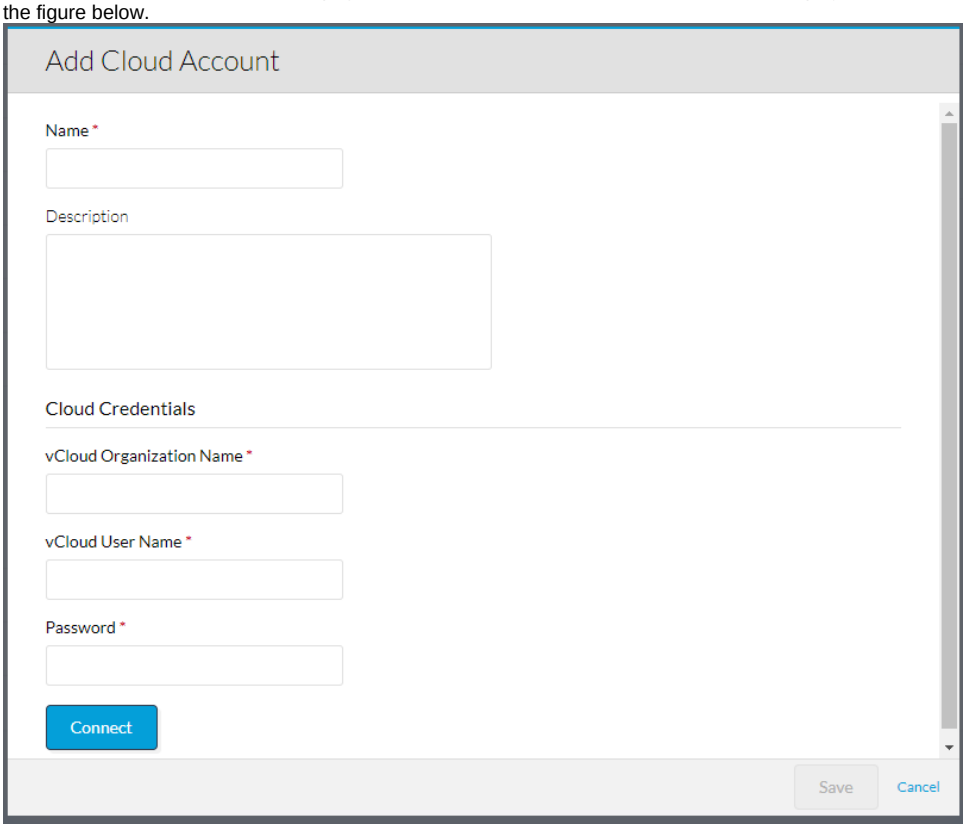

2. Assign a new cloud account**Name**.

> **Tip**
>
> The name should not contain any space, dash, or special characters.

3. Provide the vCD cloud account credentials:**vCloud Organization Name**,**vCloud User Name**, and**Password**.

4. Click the**Connect**button. CloudCenter Suitewill now attempt to validate your account credentials.
5. After the credentials are verified, the**Connect**button changes to an**Edit**button and two new fields appear**Enable Account For**and**Enable Reporting By Org Structure**.

    a. Set the**Enable Account For**dropdown per the table below.

Value	Usage
Provisioning	Workload Manager can deploy jobs using this account.
Reporting	Cost Optimizer and Workload Manager will track cloud costs for this account. Typical usage: master cloud accounts that are used for billing aggregation.
Provisioning, Reporting	Default. Account is used for both provisioning and reporting.

    b. **For AWS and Google clouds only**: Set the**Enable Reporting By Org Structure**toggle to On to cause Cost Optimizer to import the cost hierarchy createdin the cloud provider portal. This saves the time of manually creating a comparable cost hierarchy within Cost Optimizer. SeeCost Groups Configurationfor more information on cost hierarchies in Cost Optimizer.
    c. Click the**Save**button when done.

Cloud Accounts Tab

After you add cloud accounts to a cloud, they will appear in the Accounts tab for the cloud as shown in the figure below.

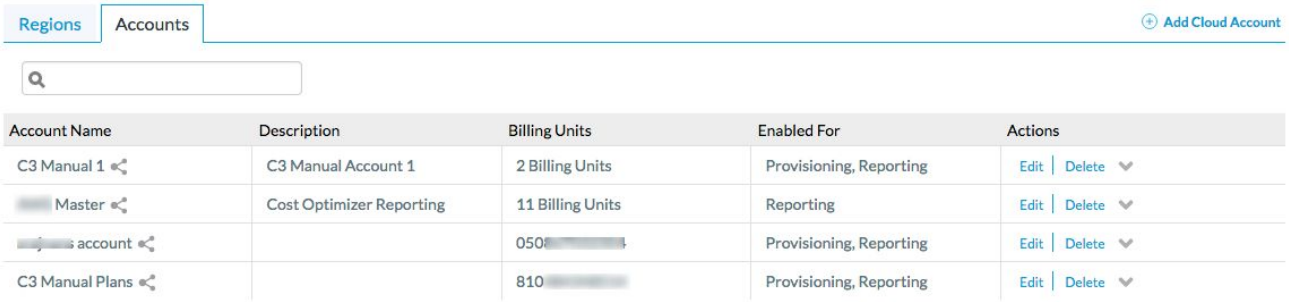

The Accounts tab contains columns for data entered when creating an account: Account Name, Description, Enabled For; and two additional columns:**Billing Units**and**Actions**.**Billing Units**is a dual function:

- If the cloud account contains only one billing unit, the ID for that billing unit is displayed.
- If the cloud account contains multiple billing units, such as an AWS master account, the number of billing units in that account is displayed followed by the text*Billing Units*.

A billing unit is the most granular level of cloud cost recording in CloudCenter Suite. The definition of a billing unit varies by a cloud provider as shown in the table below.

Cloud Provider	Billing Unit
AWS	Account ID
AzureRM	Subscription ID
Google	Project ID
IBM Cloud	Account ID
vCenter	Cloud Group Prefix - Datacenter Name
vCD	Organization Name
OpenStack	Project ID
Kubernetes	Namespace UID

The last column,**Actions**, contains links to let you edit or deleted the cloud account, ormanage instance typesfor the cloud account.

Configure a vCenter Cloud

Configure a vCenter Cloud

Configuring a vCenter cloud is a three-step process:

- Add a vCenter Cloud
- Configure a vCenter Region
- Add a vCenter Cloud Account

Add a vCenter Cloud

To add a vCenter cloud follow these steps.

1. Navigate to**Admin > Clouds**. This brings you to the Clouds page. If you, or another tenant admin in your tenant, have already added clouds to your tenant, they will be listed here.
2. Click the**Add Cloud link**in the upper right. The **Add Cloud** dialog box is displayed.
3. Enter the**Cloud Name**and select the**cloud provider**.

> When assigning the Cloud Name, use a unique name. Ensure that you do not assign a previously-used name (including deleted cloud names) as the Cloud Name as they may cause deployments to fail intermittently.

4. Since you are selecting select a vCenter cloud provider, a new data entry field appears at the bottom of the dialog box called**vCenter Region Endpoint**, as shown in the figure below. You must enter the URL of the vCenter API endpoint in this field before the**Next**button is enabled.
5. When done click**Next**. The second page of the **Add Clouds** dialog box, Connectivity Settings, appears.Set the toggle switches to configure the Cloud Connectivity settings.

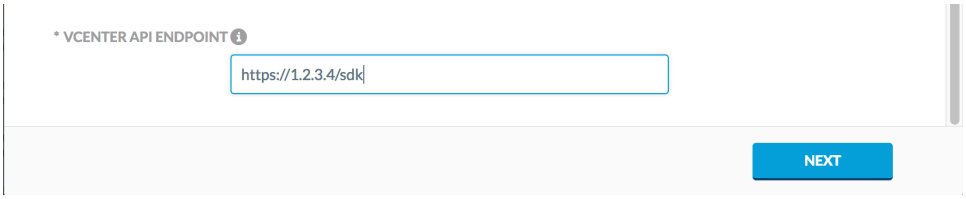

> **Note**
>
> For vCenter clouds, by default, the region endpoint URL is in the format:
> https://<vCenter_dns_name_or_IP>/sdk

- When adding a private VM cloud in the Workload Manager or Cost Optimizer UI, the second page of the Add Clouds dialog box, Connectivity Settings, appears with two toggles displayed:
  - **Worker VMs Directly Connect with CloudCenter Suite**
  - **VMs Directly Connect with CloudCenter Suite**
- Setting either of these toggles to No implies you will install Cloud Remote for each region of this cloud. This also causes a third toggle to appear:**CloudCenter Suite Directly Accessible from Cloud Remote**.
- Follow the table below for guidance on setting these toggles.

Toggle settings	Use case	Network Diagram

Cloud Endpoint Directly Accessible = Yes AND VMs Directly Connect with CloudCenter Suite = Yes	CloudCenter Suite cluster can initiate a connection to the cloud region API endpoint AND Worker VMs can initiate a connection to the CloudCenter Suite cluster Cloud Remote is not required	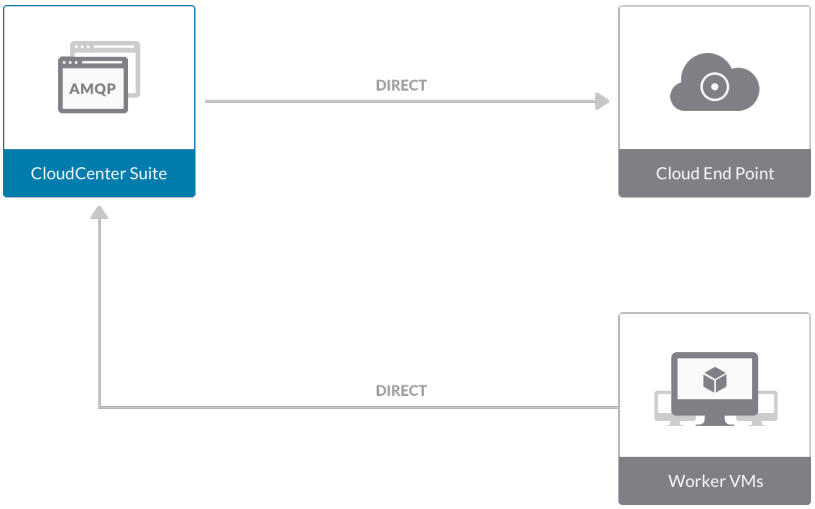
Cloud Endpoint Directly Accessible = No AND Worker VMs Directly Connect with CloudCenter Suite = No AND CloudCenter Suite Directly Accessible fromCloud Remote = Yes	CloudCenter Suite cluster cannot initiate a connection to the cloud region API endpoint AND Worker VMs cannot initiate a connection to the CloudCenter Suite cluster AND Cloud Remote can initiate the connection to the CloudCenter Suite cluster	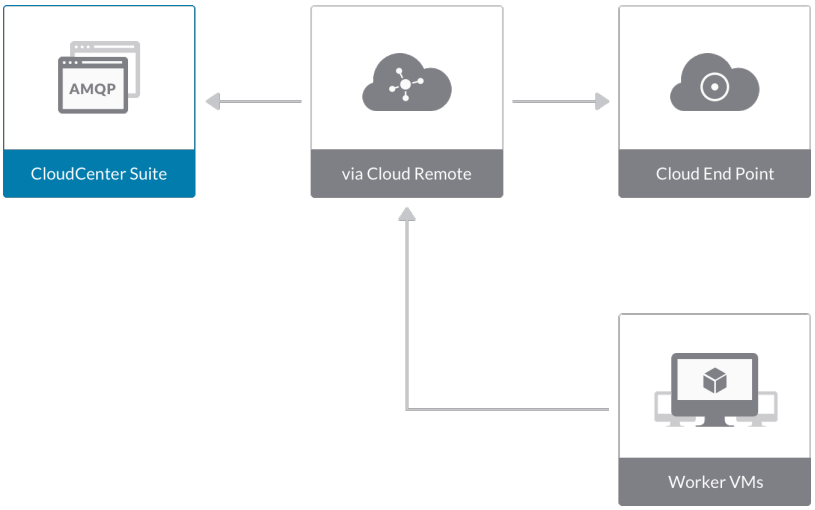
Cloud Endpoint Directly Accessible = No AND Worker VMs Directly Connect withCloudCenter Suite = No AND CloudCenter Suite Directly Accessible from Cloud Remote = No	CloudCenter Suite cluster cannot initiate a connection to the cloud region API endpoint AND Worker VMs cannot initiate a connection to the CloudCenter Suite cluster AND Cloud Remote cannot initiate the connection to the CloudCenter Suite cluster	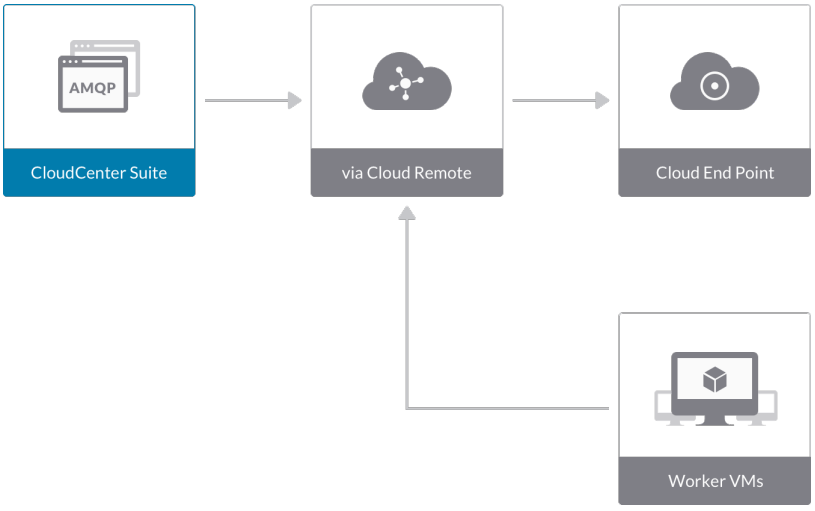

Cloud Endpoint Directly Accessible = Yes AND Worker VMs Directly Connect withCloudCenter Suite = No AND CloudCenter Suite Directly Accessible from Cloud Remote = No	CloudCenter Suite cluster can initiate a connection to the cloud region API endpoint AND Worker VMs cannot initiate a connection to the CloudCenter Suite cluster AND Cloud Remote cannot initiate the connection to the CloudCenter Suite cluster	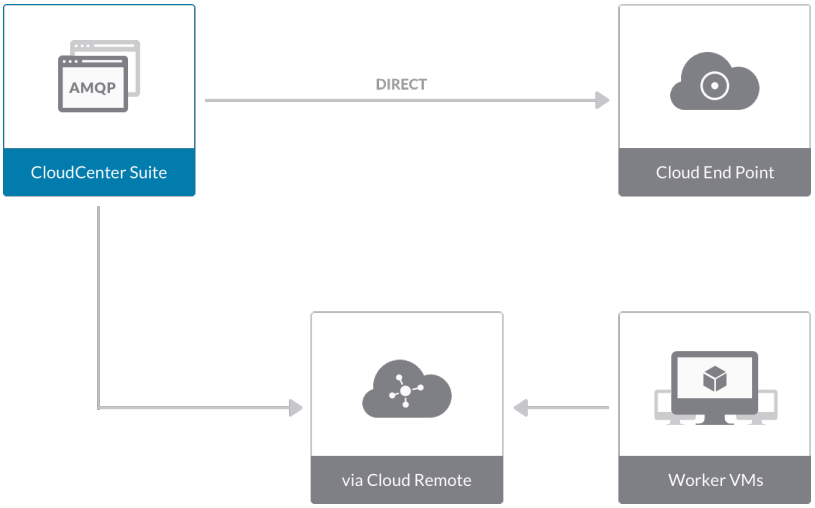
Cloud Endpoint Directly Accessible = Yes AND Worker VMs Directly Connect withCloudCenter Suite = No AND CloudCenter Suite Directly Accessible fromCloud Remote = Yes	CloudCenter Suite cluster can initiate a connection to the cloud region API endpoint AND Worker VMs need to communicate to the CloudCenter Suite cluster through Cloud Remote AND Cloud Remote can initiate the connection to the CloudCenter Suite cluster	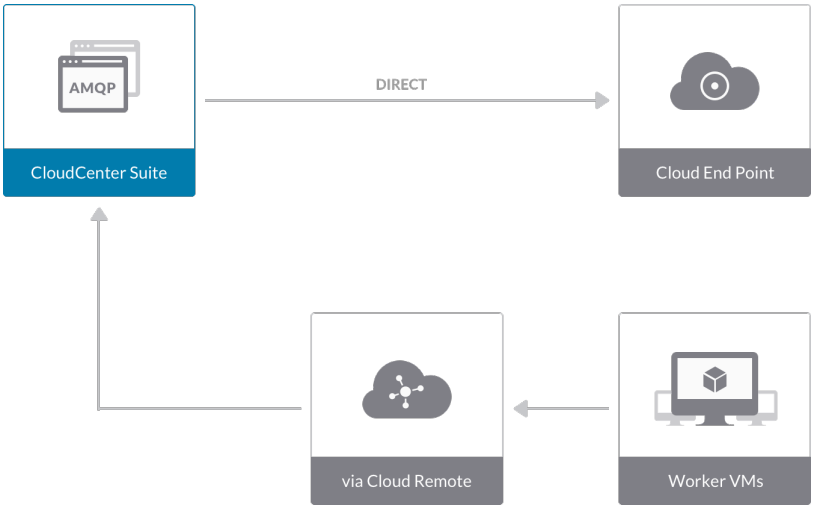

Note

The connectivity toggle settings set at the cloud level are inherited by each region you add to this cloud. However, it is possible to override these toggle settings on a per-region basis from the Regions tab for each region.

6. Click**Done**to save the configuration and close the dialog box. This brings you back to the Clouds page and the cloud you just created will be added to the bottom of the list on the left side of the page.

Configure a vCenter Region

A vCenter cloud has one region that you configure from the vCenter cloud Details tab. Follow this procedure.

1. Navigate to Clouds page:**Admin > Clouds**. Find your newly created vCenter cloud from the cloud list on the left half of the screen and click its**Configure Cloud**link. This displays the Details tab for this cloud as shown in the figure below.

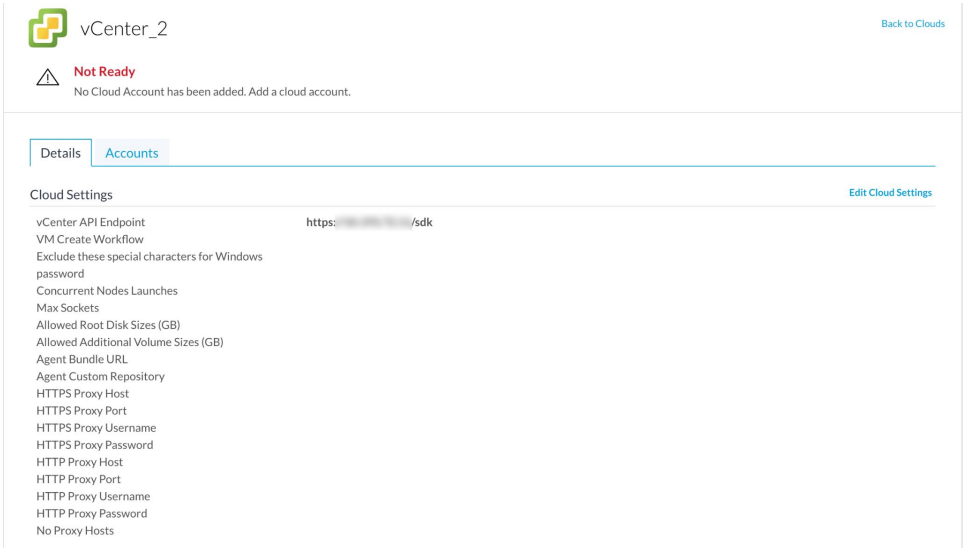

2. Click**Edit Cloud Settings**to open the Configure Cloud Settings dialog box.
   The Cloud Settings section contains fields that are unique to the vCenter cloud family and settings that are common to all cloud families. Adjust these field values per the instructions in the following tables.

vCenter Specific Cloud Settings

Field	Usage
vCenter API Endpoint	This field is set to the value you set for the API endpoint when you created this vCenter cloud. You can edit it here but should only do so if the API endpoint address of your vCenter cloud has changed since you added it to CloudCenter Suite.
VM Create Workflow	This field has two options that can be selected from a dropdown menu: • "Clone, Reconfig and Customize together" (default value) and • "Clone, Reconfig and Customize separately". Choose the second option only if the default value is resulting in failures to deploy VMs.
Concurrent Nodes Launches	This is the maximum number of VMs that can be launched simultaneously per application deployment. If left blank, the default value of 30 is applied. A value of 0 or 1 both means only one VM will be launched at a time.
Linux Max Sockets Windows Max Sockets	When the number of vCPUs assigned to a Linux VM is a prime number, Workload Manager will direct vCenter to configure the VM with that many cores on one socket. If the number of vCPUs assigned to a Linux VM is a not a prime number, Workload Manager will direct vCenter to configure the VM with X sockets of Y cores each, where X is the largest factor of the number of vCPUs which is no greater than Linux Max Sockets, and Y is vCPUs / X. The platform attempts to use the maximum number of sockets during deployment as well as when resizing instance types. During an application deployment: • If set, the Workload Manager ensures that the number of sockets set for the VM does not exceed the number specified in the setting. • If not set,the current behavior of setting the VMs vCPU as the number of sockets will continue. • Even if set,the Workload Manager doesnot use the Max Sockets settingwhen resizing the instance type.
Allowed Root Disk Sizes (GB)	Entering a comma-separated string of integers will result in corresponding options for root disk size being displayed in the Deploy form.
Allowed Additional Volume Sizes (GB)	Entering a comma-separated string of integers will result in corresponding options for secondary disk size being displayed in the Deploy form.
Disable Custom Attributes	Leaving this toggle at the default **Off** setting causes any tags specified for the VM, including tier level and deployment level tags, to be written to the attributes field in the VM. Setting this toggle to **On** prevents any tags from being written to the attributes field in the VM.

Snapshot Limit	Enter an integer for limiting the number of snapshots that can be created through Workload Manager based on the number of snapshots currently stored in vCenter. Once this limit is reached you will no longer be able to create new snapshots through Workload Manager until some of the snapshots are deleted through vCenter.

Agnostic Cloud Settings

Field	Usage
Exclude these special characters for Windows password	When the Workload Manager agent is installed on a Windows worker VM, a special user account, called cliqruser, is created to support RDP sessions that may be initiated by the user through the Workload Manager UI. A Workload Manager process running on the CloudCenter Suite cluster creates a random password and passes it to the agent for creating the cliqruser account. Because some Windows deployments may restrict using certain characters for Windows passwords, this field is provided to tell the Workload Manager to exclude these special characters in the generation of the password for the cliqruser account.
Agent Bundle URL	If you plan to use a local repository to host the bundle store, you need to enter the URL of the local bundle store here. Otherwise, leave blank.
Agent Custom Repository	If you plan to use alocal repository to host the package store, you need to enter the URL of the local package store here. Otherwise, leave blank.
HTTP /HTTPS proxy fields (host, username, password)	If you require VMs in your region to access public addresses through a web proxy, enter the URL and credentials of the HTTP and HTTPS proxy servers in these fields.
No Proxy Hosts	If you have specified an HTTP or HTTP proxy using the above fields, you can specify that managed VMs in the region should bypass the proxy and connect directly to certain hosts. Use this field to create a comma-separated list of IP addresses or URLs that should be accessed directly. This field is ignored if an HTTP or HTTPS proxy is not specified.

Important information on proxy settings

In CloudCenter Suite, you can specify proxy settings at the region level, as described here, and at the suite level. To understand the expected behavior when proxy settings are specified at both levels, see the subsequent *Precedence of Proxy Settings* section.

Download Configuration and Encryption Key

After saving the Region Configuration settings in the Workload Manager or Cost Optimizer UI, you can download them to your local computer and then upload them to other conditional components such as Cloud Remote .

The Configuration and Encryption key is only visible when you have configured the Cloud Remote component. Click the**Download Configuration** link in the upper right of the Region Connectivity section, as shown in the following screenshot.

Clicking**DownloadConfiguration**causestwo things to happen:

- An encrypted zip file namedartifacts.zipis downloaded by your browser. Make a note of the location of this zip file as you will need if you are using Cloud Remote .
- The Region Connectivity section header updates to display a**Copy Encryption Key**link, as shown in the following screenshot.

Click the**Copy Encryption Key**link to save the key to your clipboard. A success message will be displayed temporarily above the Region Connectivity section header. Make sure not to overwrite the clipboard with other data. You will need the key when you upload the configuration zip file to conditional components like Cloud Remote .

If you change the connectivity settings in the CloudCenter Suite UI and need to again download the zip file fromsoftware.cisco.com, use the automatically create a (new) encryption key, and copy the key to the clipboard by clicking the**Copy Encryption Key**link again.

Precedence of Proxy Settings

In CloudCenter Suite, you can specify HTTP and HTTPS proxy settings at the suite level as described inProxy Settings, and at the region level as described in the*Agnostic Cloud Settings*section. The CloudCenter Suite cluster, Cloud Remote , and worker VMs will each use either the suite-level proxy settings or the region level proxy settings based on the conditions described below.

For Cloud Remote mode, whenever you change the region proxy settings, ensure you again download the region connectivity setting from the Workload Manager UI and reapply to Cloud Remote .

Mode	Proxy Settings Used By		
Mode (Is Cloud Remote used to communicate withcloud (or APIC) endpoint?)	CloudCenter Suite cluster to communicate with cloud (or APIC) endpoint	Cloud Remote to communicate with cloud endpoint	Worker VM to communicate with bundle store and package store
Cloud Remote Mode	NA	Region-level	Region-level
Non-Cloud Remote Mode	Suite-level	NA	Region-level

Not all clouds support all the proxy settings, and not all clouds support bothCloud Remote and non-Cloud Remote modes.Cloud Remote mode is not applicable to public clouds. The following table lists the support for a vCenter cloud.

Cloud Category	Cloud Remote Mode	Non-Cloud Remote Mode
vCenter	All	N/A
vCenter with ACI extension	All for communication with APIC endpoint	All for communication with APIC endpoint

When you are done editing the settings in the dialog box, click**Save**.

3. Determine if you need Cloud Remote for this region. Scroll down to the Region Connectivity section for the region and click on the**Configure Region**link in the upper right to open the Configure Region dialog box.The toggle settings should be the same as when you set them on the connectivity page of the Add Cloud dialog box.If all of the connectivity togglesinthe Region Connectivity dialog boxare set to Yes, then Cloud Remote is NOT needed for this cloud region. In this case,you would normally leave the region connectivity settings at their current values and continue to the next settings section.

   The exceptionto this guidance is when a NAT firewall or proxy server exists between theCloudCenter Suite management cluster andworker VMs, or between the CloudCenter Suite management cluster and users that would use Workload Manager to initiate a Guacamole remote connection to a worker VM. In either of these cases, override the address fields in the Region Connectivity dialog box as explained below.

Networking Constraint	Field	Value
Worker VMs must use a proxy server or NAT firewall to access the"local" AMQP server running in the CloudCenter Suite cluster.	Worker AMQP IP Address	IP address and port number that the firewall or proxy server presents to the worker VMs on behalf of the "local" AMQP server running in the CloudCenter Suite cluster.
Users must use a proxy server or NAT firewall to access theGuacamole server running in the CloudCenter Suite cluster.	Guacamole Public IP Address and Port	IP address and port number that the firewall or proxy server presents to users on behalf of theGuacamole server running in the CloudCenter Suite cluster.
Worker VMs must use a proxy server or NAT firewall to access theGuacamole server running in the CloudCenter Suite cluster.	Guacamole IP Address and Port for Application VMs	IP address and port number that the firewall or proxy server presents to the worker VMs on behalf of the Guacamole server running in the CloudCenter Suite cluster.

4. Click**OK**to save the changes and dismiss the dialog box. You can now proceed to the next region settings section:VM Naming and IPAM Strategy.
5. If any of theconnectivity toggles inthe Region Connectivity dialog box are set to No, then you must install and configure Cloud Remote for this region.

Configure Cloud Remote in a vCenter Region

Configure Cloud Remote in a vCenter region as follows.

Download and Launch the Cloud Remote Appliance in vCenter

a. From your local computer, download the Cloud Remote appliance OVA fromsoftware.cisco.com.
b. Log in to the vCenter console using the vSphere web client with Flash, or with the vSphere Windows client. Do not use the HTML5 web client.
c. Navigate to the folder or resource pool where you want to deploy the OVA. Right-click on that resource pool or folder and select Deploy OVF Template.
d. From the Deploy OVF Template dialog box, for Source, select Local file and click Browse to find the OVA file you downloaded in step 1.
e. Complete the fields for Name and location, Host / Cluster, Resource Pool, Storage, and Disk Format appropriate for your environment.
f. For the Network Mapping section, make sure to properly map the Management network (public) and VM Network network (private) to the appropriate network names in your environment.
g. For the Properties section, make sure to check the box labeled Does the VM need a second interface? if the Cloud Remote appliance needs to be multi-homed on a public network and a private network.
h. Confirm your settings and click Finish to launch the VM.
i. Optional but recommended for production environments:Deploy two additional instances of the appliance to form a cluster for HA. Cloud Remote includes support for the clustering of multiple nodes. You will "add" these two additional instances to the first instance after the first instance is configured. SeeCloud Remote (Conditional)> *Scaling*for details.
j. Once the first instance of the appliance has been launched, use thevSphere clienttonote its IP public and private addresses. You will need this information later on in order to login to the Cloud Remote web UI and to complete the Region Connectivity settings in the CloudCenter Suite Web UI. Also, note the IP addresses of any other appliances you launch.

Unable to render {include} The included page could not be found.

Specify AMQP and Guacamole Addresses for Supporting Cloud Remote

From the CloudCenter Suite UI, for the cloud region requiring Cloud Remote, navigate to the corresponding Regions or Details tab. Click the**Configure Region**link in the upper left of the Region Connectivity section to bring up the Configure Region dialog box. The toggle settings should be the same as when you set them on the connectivity page of the Add Cloud dialog box. You must update some of the address fields in the dialog box according to the scenarios summarized in the table below.

Toggle Settings	Field	Value
Worker VMs Directly Connect with CloudCenter = No AND CloudCenter Directly Accessible from Cloud Remote = Yes	Local AMQP IP Address	Pre-populated with the address and port number of the "local" AMQP server running in the CloudCenter Suite cluster. This address must be **accessible to Cloud Remote** . If Cloud Remote is accessing the CloudCenter Suite cluster through a user-supplied proxy server or NAT firewall, overwrite this field with the corresponding local AMQP IP address and port number provided by the user-supplied proxy server or NAT firewall and accessible to Cloud Remote .
Worker VMs Directly Connect with CloudCenter = No AND CloudCenter Directly Accessible from Cloud Remote = No	Remote AMQP IP Address	Enter <Cloud_Remote_IP>:<amqp_port>, where <Cloud_Remote_IP> = the IP address Cloud Remote which is **accessible to the CloudCenter Suite cluster**, and <amqp_port> = 5671 OR the custom AMQP port number you would later set with the **Change Ports shell script** on the Cloud Remote appliance (see Cloud Remote (Conditional) > *Custom Port Numbers (Conditional)*). If there is no user-supplied NAT firewall or proxy server between the CloudCenter Suite cluster and Cloud Remote , the IP address would be the public IP address of Cloud Remote . If there is a NAT firewall or proxy server between the CloudCenter Suite cluster and Cloud Remote , instead, enter the corresponding public IP address and port number that the firewall or proxy server presents to the internet on behalf of the "remote" AMQP server running in Cloud Remote .
Worker VMs Directly Connect with CloudCenter = No	Worker AMQP IP Address	Enter <Cloud_Remote_IP>:<amqp_port>, where <Cloud_Remote_IP> = the Cloud Remote IP address **accessible to the worker VMs**, and <amqp_port> = 5671 OR the custom AMQP port number you would later set with the **Change Ports shell script** on the Cloud Remote appliance (see Cloud Remote (Conditional) > *Custom Port Numbers (Conditional)*).
Worker VMs Directly Connect with CloudCenter = No	Guacamole Public IP and Port	Enter <Cloud_Remote_IP>:<guac_port>, where <Cloud_Remote_IP> = the Cloud Remote IP address **accessible to CloudCenter Suite users**, and <guac_port> = 8443 OR the custom Guacamole port number you would later set with the **Change Ports shell script** on the Cloud Remote appliance (see Cloud Remote (Conditional) > *Custom Port Numbers (Conditional)*).
Worker VMs Directly Connect with CloudCenter = No	Guacamole IP Address and Port for Application VMs	Enter <Cloud_Remote_IP>:<guac_port>, where <Cloud_Remote_IP> = the Cloud Remote IP address **accessible to worker VMs**, and <guac_port> = 7789

When done, click **OK** to save the setting and dismiss the dialog box.

Download Region Connectivity Settings and Upload to Cloud Remote

Cloud Remote uses the region connectivity settings set in the Workload Manager or Cost Optimizer UI. After saving the Region Configuration settings in the Workload Manager or Cost Optimizer UI, you must download them and to your local computer and then upload them to Cloud Remote as follows.

Click the**Download Configuration**link in the upper right of the Region Connectivity section, as shown in the figure below.

Clicking DownloadConfiguration causestwo things to happen:

- An encrypted zip file named**artifacts.zip**will be downloaded by your browser. Make note of the location of this zip file as you will need to upload it to Cloud Remote through the Cloud Remote web UI (see below).
- The Region Connectivity section header updates to display a **Copy Encryption Key**link, as shown in the figure below.

Click the**Copy Encryption Key**link to save the key to your clipboard. A success message will be displayed temporarily above the Region Connectivity section header. Make sure not to overwrite the clipboard with other data. You will need the key when you upload the configuration zip file to Cloud Remote .

> ⚠ If you change the connectivity settings in the CloudCenter Suite UI and need to again download the zip file, a new encryption key is automatically created and can be copied to the clipboard by clicking the **Copy Encryption Key** link again.

After you have downloaded the zip file and copied the encryption key to your clipboard, login toCloud Remote web UI.

a. Open another browser tab and login to https://<Cloud Remote _ip> with thedefault credentials: admin/cisco.
b. You will immediately be required to change your password. Do so now.
c. You are now brought to the Cloud Remote home page as shown in the figure below.

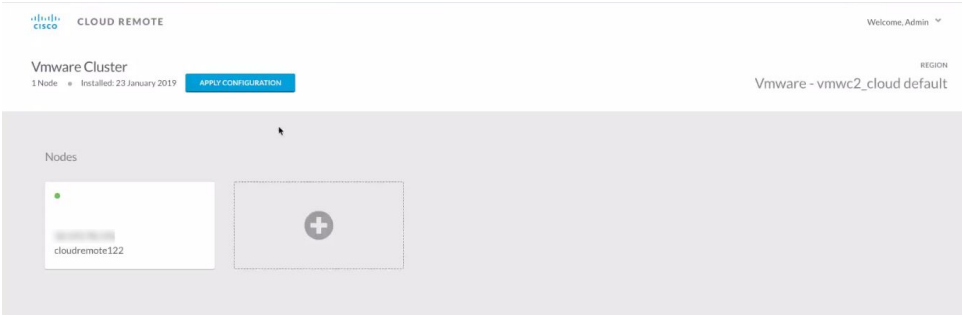

d. Click the **Apply Configuration** button in the page header. This prompts you to select a configuration file and enter the encryption key as shown in the figure below.

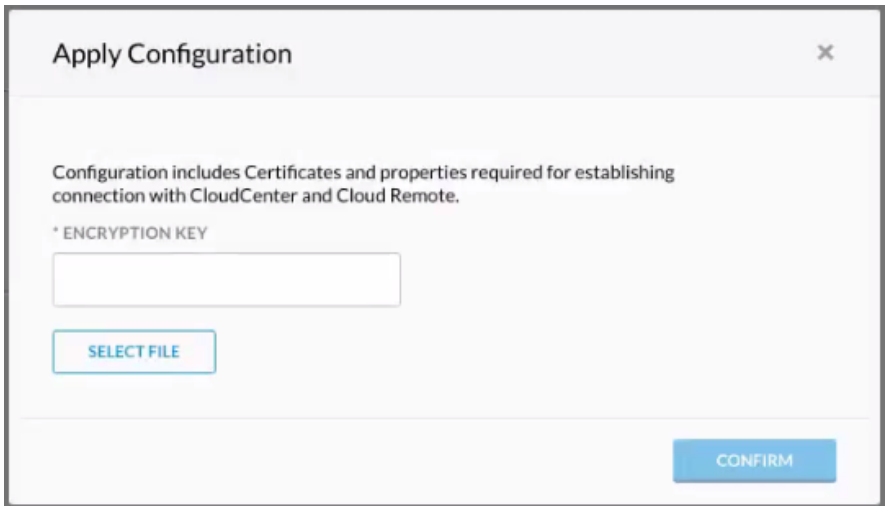

e. Paste the encryption key that was copied to the clipboard into the**Encryption Key**field in the dialog box.
f. Click**Select File**and browse tothe artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
g. Click**Confirm**.
h. Once the zip file is successfully transmitted and accepted, the Cloud Remote appliance attempts to establish communication with the CloudCenter Suite cluster and the Cloud Remote web UI home page is updated to show the name of the region it is connecting to in the upper right (see figure below).

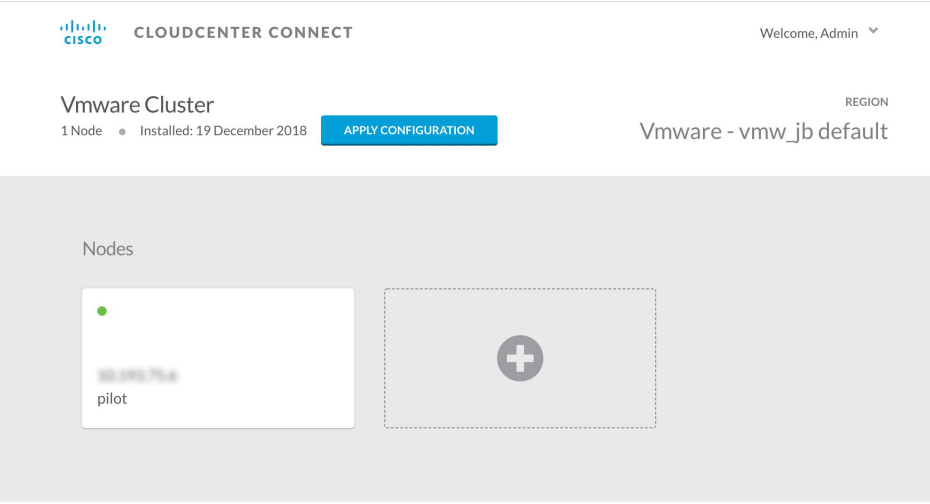

Switch your focus back to the Region Connectivity section of the target cloud region in the CloudCenter Suite web UI. The status indicator in the Region Connectivity section header will change from Not Configured to Running once connectivity between Cloud Remote and the CloudCenter Suite cluster is completely established (see figure below).

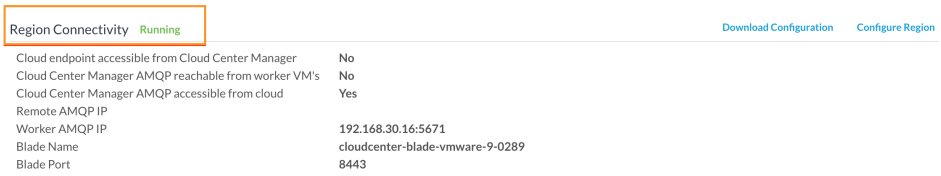

After completing these steps, Workload Manager and Cost Optimizer can use Cloud Remote for communicating with the target cloud region.

6. VM Naming and IPAM Strategy (conditional): Configure any VM naming or IPAM strategies in the Strategy section. If you leave the settings at the defaults, no IPAM strategy is applied and the default VM naming strategy is applied.

VM Naming and IPAM Strategies

Workload Manager lets you define a VM naming strategy and an IP address management (IPAM) strategy on a per cloud region basis. The VM naming strategy lets Workload Manager generate unique names for each of your VMs as they are launched. The IPAM strategy letsWorkl oad Manager generate IP addresses for each of the vNICs on each of your VMs. The VM naming strategy can reference a VM naming callout script, and the IPAM strategy requires that you specify an IP address allocation callout script andan IP address deallocation callout script. An IP address allocation callout scriptmay work in conjunction with a 3rd party IPAM server in your cloud environment.

Like External Lifecycle Actions, the scripts or commands associated with a VM naming strategy or IPAM strategy are executed in a containerized script execution engine running in the CloudCenter Suite cluster, if you are not using Cloud Remote in your region. If you are using Cloud Remote , these scripts or commands will execute in a containerized script execution engine running in Cloud Remote .

To set up the VM naming and IPAM strategies for your region, scroll down to the Strategy section of the Regions or Details tab and follow these steps.

Click the**Edit Strategy**link in the upper right to display the Edit Strategy dialog box as shown in the figure below.

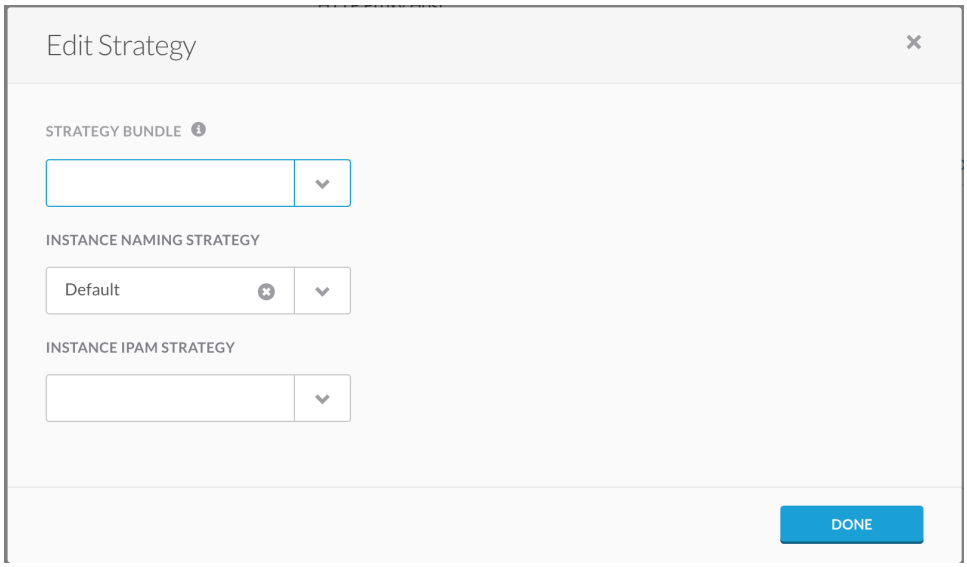

The dialog box can reveal a total of seven data entry fields as explained in the table below:

Field	Usage
Strategy Bundle	Zip file containing all of the scripts associated with your VM naming strategy or IPAM strategy. Select source from the dropdown an add the path in the new field that appears to the right. 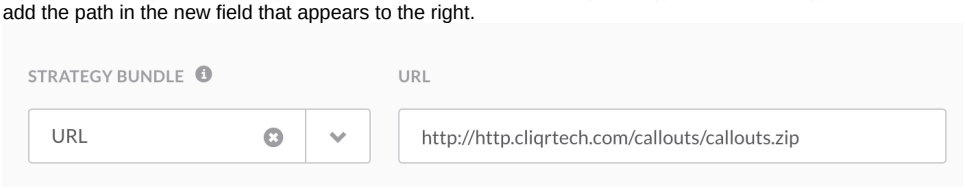
Instance Naming Strategy	Dropdown for selecting one of three naming strategies: Default, CliQr Macro Replacement, and Hostname Callout. If you select CliQ Macro Replacement, the **Node Name Config** field becomes visible and required. If you select Hostname Callout, the **Custom VM Name** field becomes visible and required. See the Instance Naming Strategy table below for details.

Node Name Config	Required if you selected "CliQr Macro Replacement" for theInstance Naming Strategyfield. Content: text string consisting of a combination of characters and numbers with embedded macro variables (see table below). The macros get translated to text strings run time.
	
	The resulting VM name string mustmeet the following requirements for job submission to succeed:
	• Only contain the following characters:
	  ◦ Lowercase a to z
	  ◦ 0 to 9
	  ◦ Hyphen
	• Start withan alpha character
	• Cannot end with a hyphen

Instance IPAM Strategy	Dropdown menu with two choices: No IPAM (default), and IPAM Callouts. If you select the latter, two additional fields are displayed and become required: **IPAM Alloc Rule**and**IPAM Dealloc Rule**. (See below).
Custom VM Name	Location of your custom VM naming script if you set the**Instance Naming Strategy**field to "Custom Callout". Location options inclu "Script from bundle" if you specify a zip file in the **Strategy Bundle** field. See Guidance for Callout Scripts for guidance on developi custom VM naming scripts.
IPAM Alloc Rule	Location of your custom IPAM address allocation script if you set the**Instance IPAM Strategy**field to "IPAM Callouts". Location opti include "Script from bundle" if you specify a zip file in the **Strategy Bundle** field. See Guidance for Callout Scripts for guidance on developing custom IPAM callout scripts.
IPAM Dealloc Rule	Similar to the**IPAM Alloc Rule**above except this is where you specify the location of yourcustom IPAM address deallocation script. Guidance for Callout Scripts for guidance on developing custom IPAM callout scripts.

Instance Naming Strategy Options:

Options	Description
Default	The default method provided by Cisco is the generation of a string consisting of A predefined token prefixed with **cq jw-**
CliQr Macro Replacement	Selecting this method causes the **Node Name Config** field to be displayed and to become a required field. See the fields table above for details.
Hostname Callout	Selecting this method causes the **Custom VM Name** field to be displayed and to become a required field. See the fields table above for details.

OOB macro variables forCliQr Macro Replacement:

Macro	Returned Value
%os1%	First character of OS type string in lower case
%os2%	Firsttwo characters of OS type string in lower case
%OS1%	Firstcharacter of OS type string in upper case
%OS2%	Firsttwo character of OS type string in upper case
% RND<number>%	Fixed length random string, e.g., %RND10%, 10-character long random string ⚠ This is a mandatory field, only values in the range: %RND6% to %RND11% or just %RND% which defaults to %RND6%"
%UID%	User ID
%VID%	Vendor ID
%VM_PREFIX%	VM name prefix must be added as a global parameter in the app profile. ⚠ Only macro whose value can be user-customized.

7. External Lifecycle Actions (conditional):Specify any external lifecycle actions to be performed on all VMs launched by Workload Manager in this region.

External Lifecycle Actions Settings

Whenever you use Workload Manager to deploy a VM-based application, Workload Manager can execute external actions on VMs in the deployment as the VMs progress through these five phases of their lifecycle:

- Pre VM start
- Pre VM initialization
- PostVM initialization
- Pre VM Stop
- Post VM Stop

These lifecycle actions will execute in a script execution engine running in a container in the CloudCenter Suite cluster if no Cloud Remote appliance is configured for the region. If you have configured Cloud Remote for this region, the external lifecycle actions will executein a script execution engine running in a container in the Cloud Remote appliance.

External Lifecycle Action scripts can be defined at the service level, the application level, and the cloud region level. Use the External Lifecycle Actions section of the Region tab or Details tab (see figure below) tospecify external lifecycle actions at the region level.

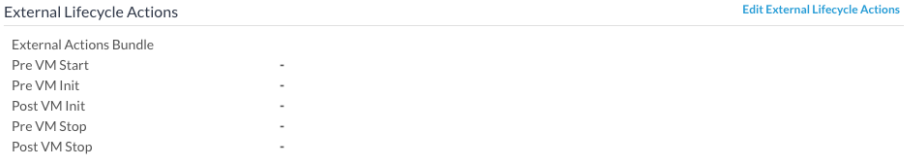

Click on the**Edit External Lifecycle Actions**link in the upper right. This displays theEdit External Lifecycle Actions dialog box as shown below.

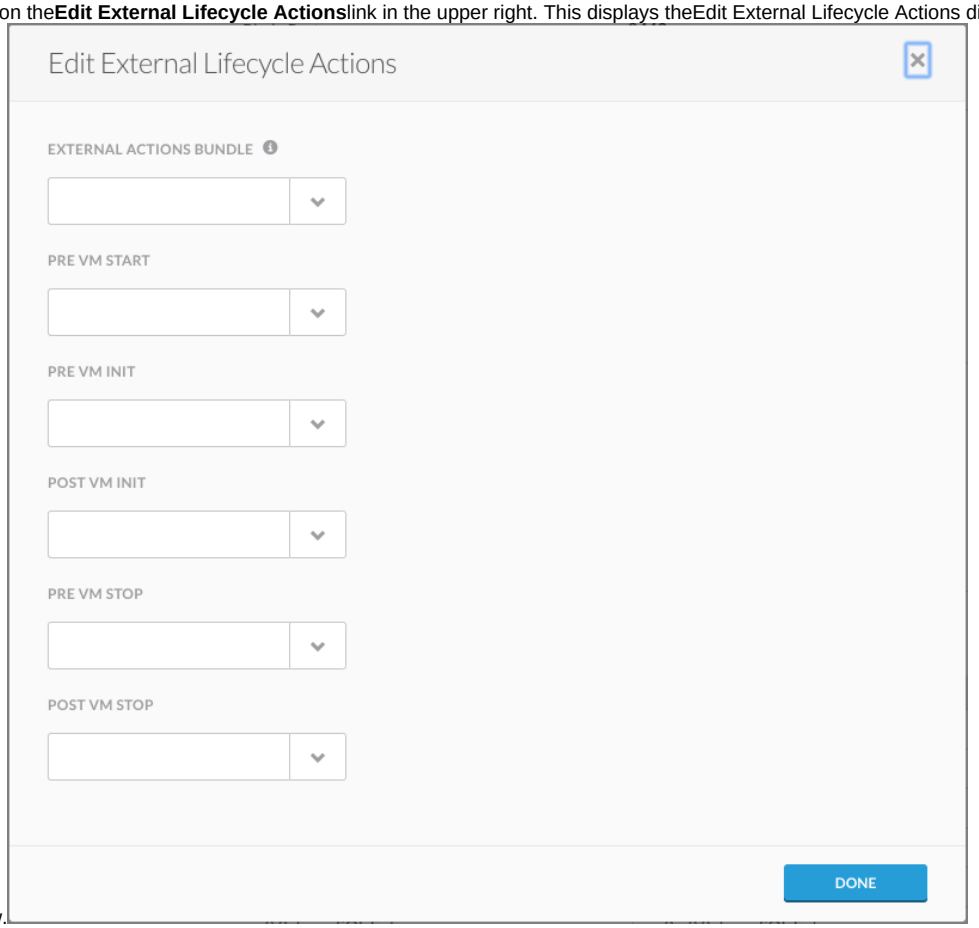

One way to specify an external lifecycle action is to use the first field,**External Actions Bundle**, to specify a zip file that contains all scripts for external lifecycle management. This zip file needs to contain a directory that contains all the scripts. Use the dropdown to select a repository location (URL, FTP, S3, etc.). This causes an additional field to be displayed to the right for entering the path to the zip file as shown in the figurebelow.

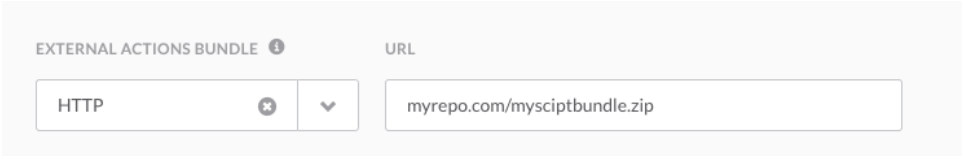

Once you specify actions bundle, for each of the individual lifecycle stage fields, select "Script from Bundle" from the dropdown. This displays an additional field to the right where you must specify the name of the script file in the bundle as shown in the figure below.

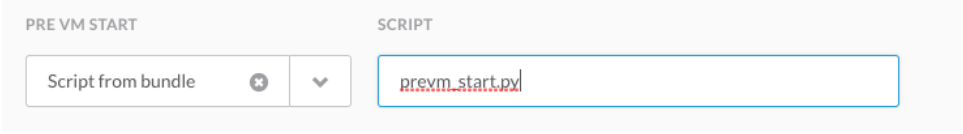

The scripts you specify can be bash scripts or Python scripts called from within a bash script.For more information on lifecycle scripts, seeDeployment Lifecycle Scripts.

8. Instance Types (conditional): A vCenter cloud region includes one "default" instance type with 1 vCPU, 1 vNIC, 1024 MB RAM, and no additional disk storage. CloudCenter Suite will also automatically create instance types based on the parameters of VMs you deploy from within vCenter. You would manually add more instance types to your vCenter region if you want Workload Manager to deploy jobs to this region with differently sized instance types.

Instance Types Settings

An instance type is a performance category of a VM for VM-based clouds or a performance category of a container for container-based clouds. The parameters for a vCenter instance type is as follows:

- Display name
- Price per hour

- Instance type ID
- CPU
- 32- or 64-architecture
- RAM (MB)
- NICs
- SSD support

You can sync, edit, and add instance types as shown below:

- Add
- Edit
- Auto-create

In addition, CloudCenter Suite will auto-create instance types for vCenter clouds based on observing the VM parameters of instances executing in vCenter. This is to facilitate cost calculations for vCenter clouds based on the instance type.

For public clouds, CloudCenter Suite auto-populates the instance types in the Instance Types section from the cloud provider. CloudCenter Suite auto-syncs this information from the cloud provider once per 24 hours. For these clouds, you cannot force sync on demand.

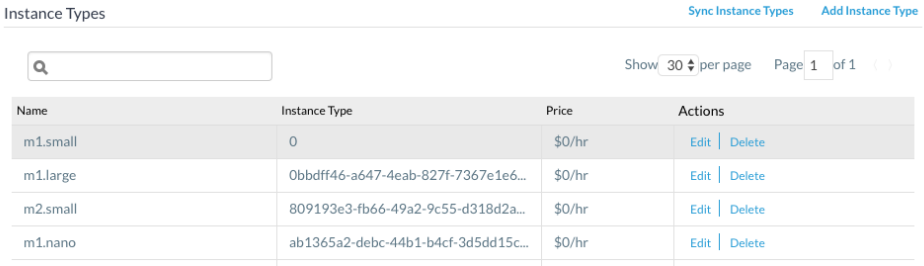

You can add new instance types for vCenter by clicking the**Add Instance Types**link in the upper right of the section. These instance types will control the parameters of the corresponding VMs or containers launched as part of your application. There is no concept of instance type native to vCenter, therefore the instance ID of the instance types you create in CloudCenter Suite can be arbitrary and the CPU and RAM parameters you specified in CloudCenter Suite will be used in the deployment.

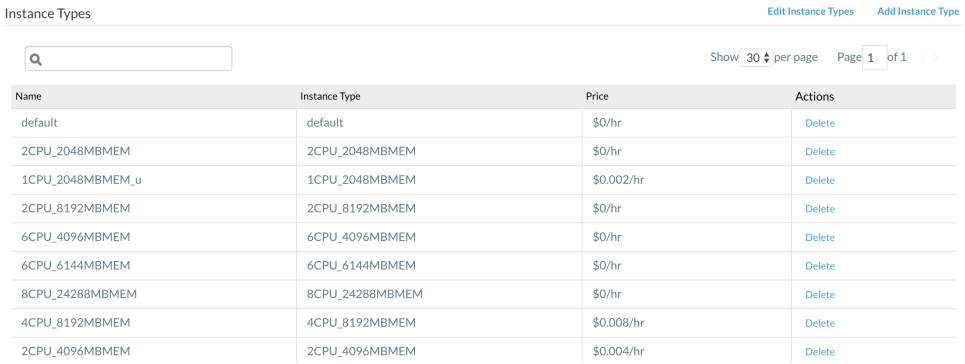

You must specify the price details in the**Price**field forvCenter cloudwhen adding instance types. The price is used to compute resource costs, which is displayed as invoice costs in theCost Optimizer Dashboard. The**Price**column (highlighted in red in the above screenshot) displays the cost of an instance.

For all cloud providers except AWS, you can edit the fields of an instance type (except instance type ID) by clicking the**Edit**link in the Actions column in the list of instance types. For vCenter cloud changing the values of parameters such as CPU and RAM will impact the properties of the VMs or containers deployed in those clouds.

Historical Price Management

Sometimes, you want to update the historical cost to correct or update cost computation. In such cases, thecost of the affected resourcesis computed from the date and time the price is updated.You can update the prices of the instances types for the preceding three months only.

To update the price, follow these steps:

a. Click**Edit Instance Type**. The**Edit Instance Types**page appears.
b. In the**Price/Hour**column, click the**Pricing History**icon. The**EditPrice History**dialog appears.

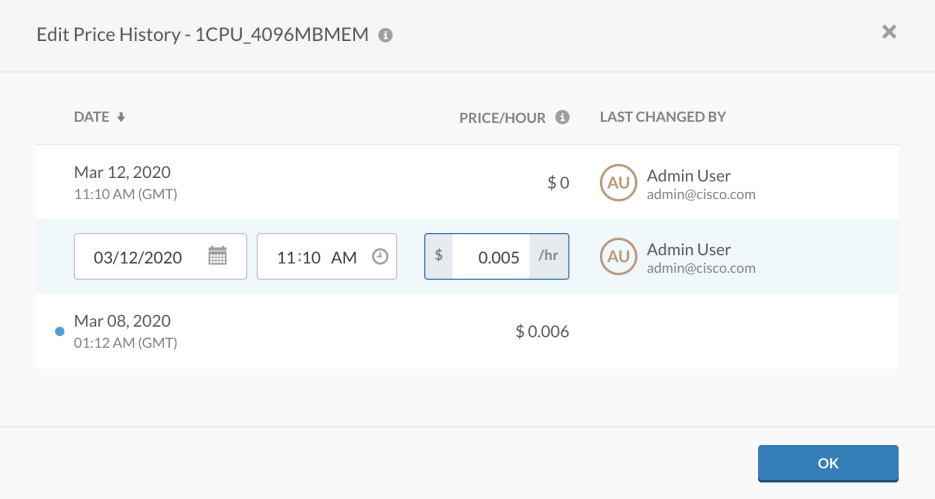

c. Specify a date and time in the appropriate fields and enter the desired price inthe**Price/Hour**field, enter the desired price.
d. Click**OK**, followed by**Done**(in the**Edit Instance Types**page) to save the changes.

After the price has been updated, the cost is recomputed for allresources(VMs and deployments, as applicable)affected by the price change. This could take a few minutes. When the recomputation is successful, you receive a notification that can be viewed in the Notification icon in the Cost Optimizer header (seeCost Optimizer Dashboard), as shown in the following sample screenshot.

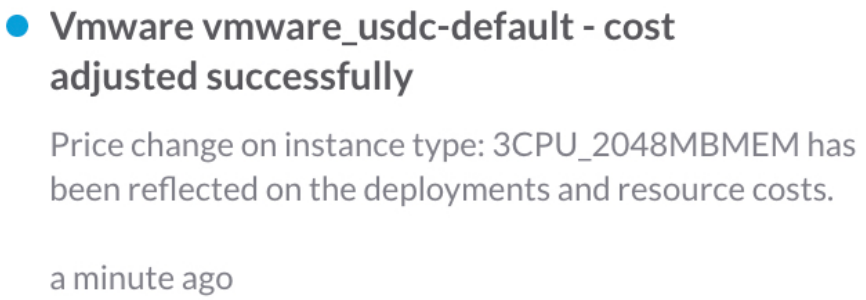

Simultaneously,the updated accruedcostof a VMis displayed on theVirtual Machinespage and the updated deployment cost is displayed on the Deployment Details page.

9. Storage Types (conditional): For private VM-based clouds like vCenter, CloudCenter Suite uses storage types for cost tracking purposes. CloudCenter Suite creates a default storage type with zero cost.You would manually edit this storage type to enter your own cost factor. You can optionally add more storage types to your vCenter region.

Storage Types Settings

A storage type is a categorization of storage that can be provisioned for an application during deployment. It is a construct that applies to VM-based clouds only. Storage types are displayed in the Storage Types section of the Regions or Details tab.

The storage type construct contains the following fields for all VM-based clouds:

- Display Name
- Cost Per Month per GB
- Cloud Storage Type ID
- Minimum Volume Size (GB)
- Maximum Volume Size (GB)
- Maximum Throughput (MB/s)
- Maximum IOPS
- Provisioned IOPS
- IOPS Cost Per Month

Storage types are handled by CloudCenter Suite for various cloud providers in a fashion similar to instance types as shown in the table below.

Function	AWS	AzureRM	Google	vCenter	OpenStack
Auto-populate and Auto-sync	x	x	x		
Add				x	x
Edit		x	x	x	x

For public clouds, CloudCenter Suite auto-populates the list of storage types in the Storage Types section based on information from the cloud provider. This data is stored in thepackage store and the package store auto-syncs this information from the cloud provider once per 24 hours.

For private VM-based clouds, such as vCenter and OpenStack, CloudCenter Suite auto-creates a single defaultstoragetype with zero cost. You may manually add additional storage types by clicking the**Add Storage Type**link in the upper right of the Storage Types settings section as shown in the figure below.

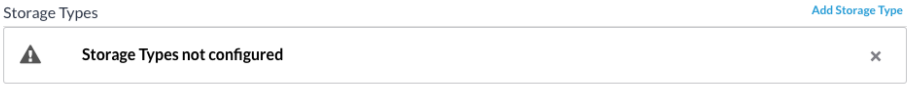

This brings up an**Add Storage Type**dialog box as shown in the figure below.

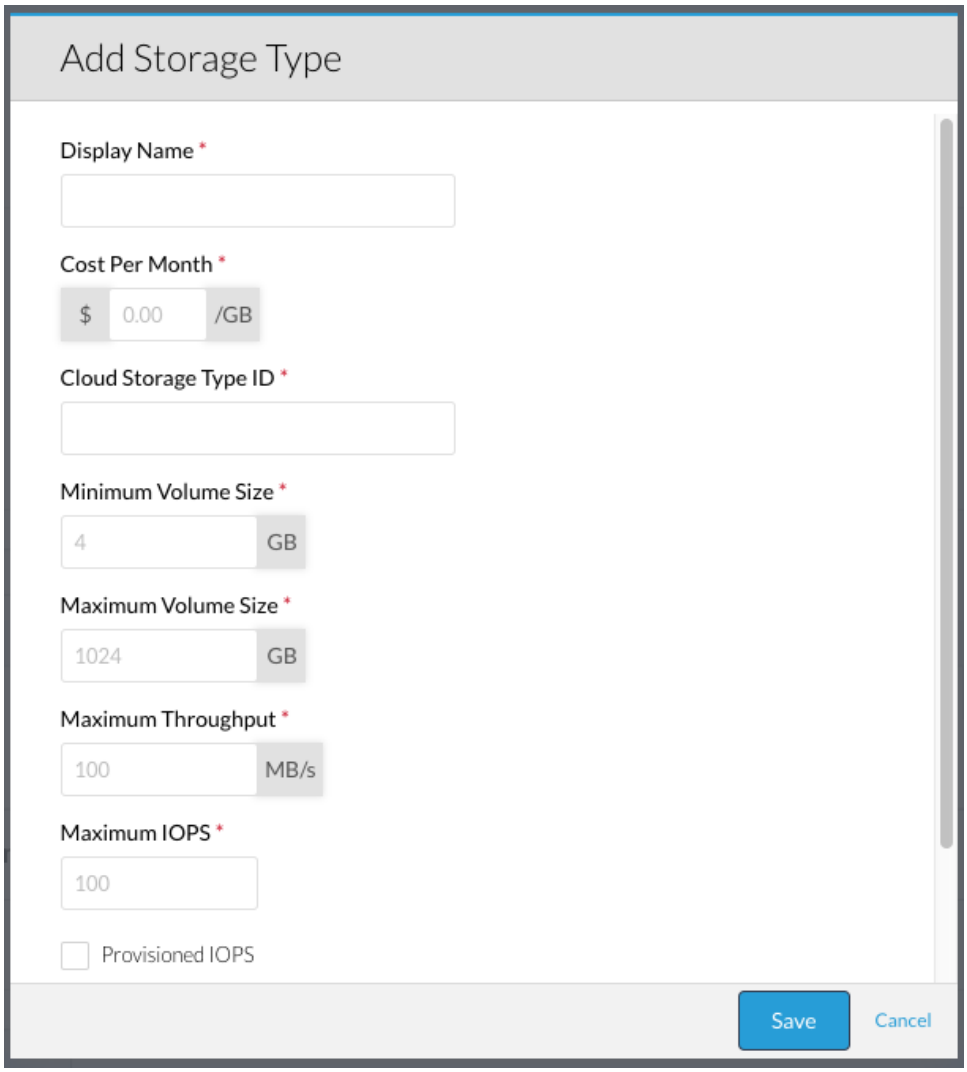

You must enter all required fields, however, the**Maximum Throughput**,**Maximum IOPS**,**Provisioned IOPS**, and**IOPS Cost Per Month**are not recognized by these cloud providers and will be ignored. You must specify the**Cost**to compute the service costs in Cost Optimizer. The total cost is displayed in theCost Optimizer Dashboard.

Storage types can be edited. From the list of storage types in the Storage Type section (see figure below), click the Edit link (in the Actions column) for that storage type. This brings up an Edit Storage Type dialog box (similar to the Add Storage Type dialog box) where you can edit all nine fields except the**Cloud Storage Type ID**field.

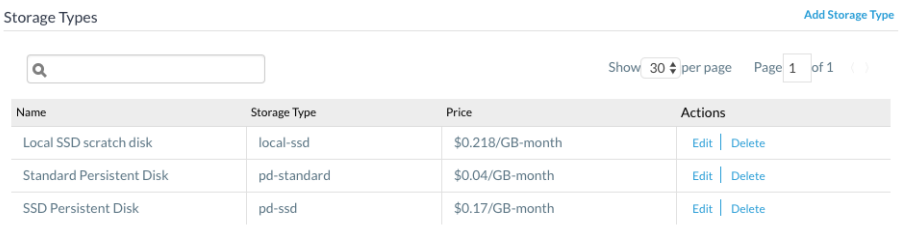

For all cloud providers,**Minimum Volume Size**and**Maximum Volume Size**restrict the volume size you can specify at deploy time. For public cloud providers that support provisioned IOPS, the user can specify the guaranteed IOPS at deploy time.

Historical Price Management

Sometimes, you want to update the historical cost to correct or update cost computation. In such cases, the cost of the affected resources is computed from the date and time the price is updated.You can update the prices of the older storage types for the preceding three months only.

To update the price, follow these steps:

a. Click**Edit Instance Type**. The**Edit Instance Types**page appears.
b. In the**Price/Hour**column, click the**Pricing History**icon. The**EditPrice History**dialog appears.

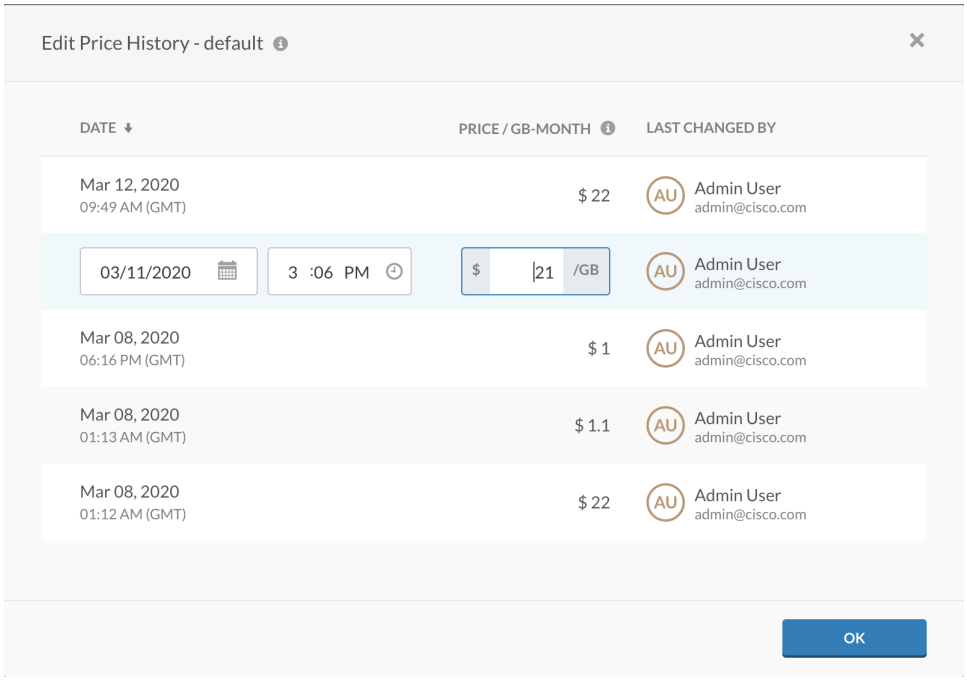

c. Specify a date and time in the appropriate fields and enter the desired price inthe**Price/Hour**field, enter the desired price.
d. Click**OK**, followed by**Done**(in the**Edit Instance Types**page) to save the changes.

After the price has been updated, the cost is recomputed for allresources(VMs and deployments, as applicable)affected by the price change. This could take a few minutes. When the recomputation is successful, you receive a notification that can be viewed in the Notification icon (seeCost Optimizer Dashboard) in the header.Simultaneously,the updated accruedcostof a VMis displayed on theStorage Volumespage and the updated deployment cost is displayed on the Deployment Details page.

10. Image Mappings: Image mappings allow services based on Workload Manager logical images to be deployed using the appropriate physical image stored on the target cloud region. You must manually import these physical images into your vCenter region and then map the appropriate Workload Manager logical images to these physical images. SeeImagesfor more context.

Add a vCenter Cloud Account

Prerequisites

For Workload Manager to deploy jobs in vCenter using a particular user account, that account must have the permissions identified in the table below.

vCenter Object	Required Permission	Reason

Network	Assign Network	If the default network in a template/snapshot must be changed
Datastore	Allocate space	For persistent disk operation
	Browse datastore	
	Low-level file operations	
	Remove file	
Folder	Create folder	For user folder creation
Resource	Apply recommendation	For datastore cluster support
	Assign VM to resource pool	For resource pool selection
Tasks	Create task	For VM operation
	Update task	
Virtual Machine	All permissions	
Global Role	Set Custom Attributes	To add custom attributes on virtual machines
	Manage Custom Attributes	

Configuration Process

To add a vCenter cloud account, follow this process:

1. Locate the vCenter cloud in the Clouds page and click**Add Cloud Account**button. This will display the Add Cloud Account dialog box as shown in the figure below.

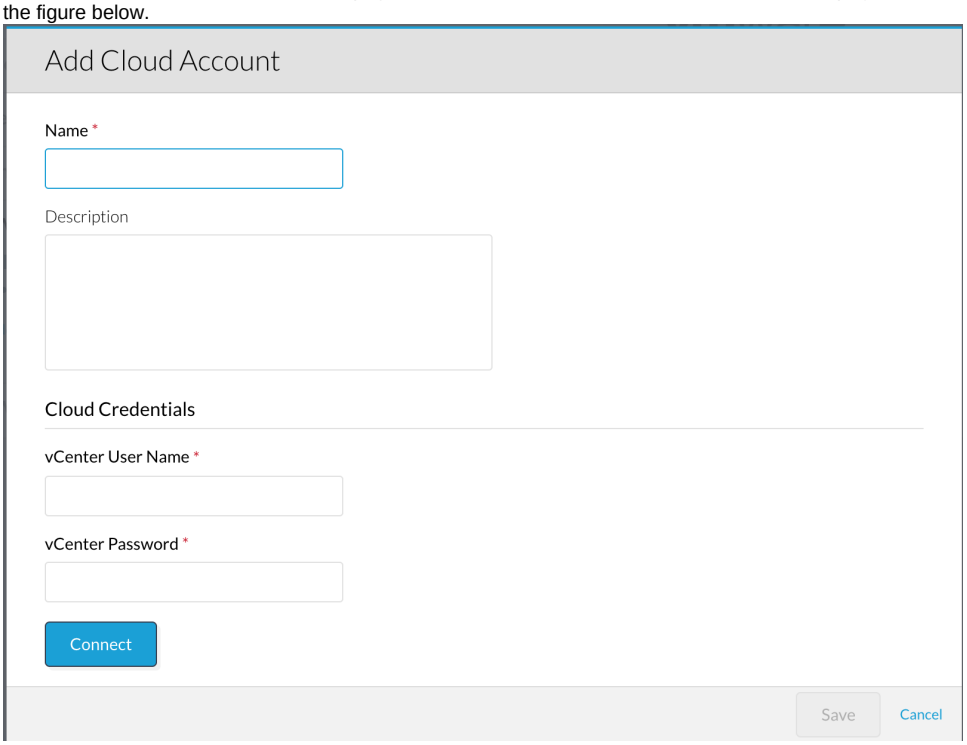

2. Assign a new cloud account**Name**.

> **Tip**
>
> The name should not contain any space, dash, or special characters.

3. Provide the vCenter cloud credentials:**vCenter User Name**and**vCenter Password**.
4. Click the**Connect**button. CloudCenter Suitewill now attempt to validate your account credentials.
5. After the credentials are verified, the**Connect**button changes to an**Edit**button and two new fields appear**Enable Account For**and**Enable Reporting By Org Structure**,

    a. Set the**Enable Account For**dropdown per the table below.

Value	Usage
Provisioning	Workload Manager can deploy jobs using this account.
Reporting	Cost Optimizer and Workload Manager will track cloud costs for this account. Typical usage: master cloud accounts that are used for billing aggregation.
Provisioning, Reporting	Default. Account is used for both provisioning and reporting.

b. **For AWS and Google clouds only**: Set the**Enable Reporting By Org Structure**toggle to On to cause Cost Optimizer to import the cost hierarchy createdin the cloud provider portal. This saves the time of manually creating a comparable cost hierarchy within Cost Optimizer. SeeCost Groups Configurationfor more information on cost hierarchies in Cost Optimizer.
c. Click the**Save**button when done.

Cloud Accounts Tab

After you add cloud accounts to a cloud, they will appear in the Accounts tab for the cloud as shown in the figure below.

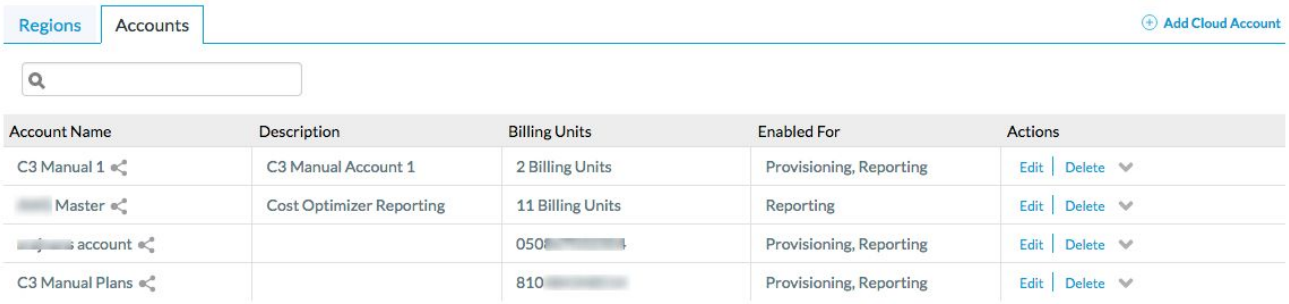

The Accounts tab contains columns for data entered when creating an account: Account Name, Description, Enabled For; and two additional columns:**Billing Units**and**Actions**.**Billing Units**is a dual function:

- If the cloud account contains only one billing unit, the ID for that billing unit is displayed.
- If the cloud account contains multiple billing units, such as an AWS master account, the number of billing units in that account is displayed followed by the text*Billing Units*.

A billing unit is the most granular level of cloud cost recording in CloudCenter Suite. The definition of a billing unit varies by a cloud provider as shown in the table below.

Cloud Provider	Billing Unit
AWS	Account ID
AzureRM	Subscription ID
Google	Project ID
IBM Cloud	Account ID
vCenter	Cloud Group Prefix - Datacenter Name
vCD	Organization Name
OpenStack	Project ID
Kubernetes	Namespace UID

The last column,**Actions**, contains links to let you edit or deleted the cloud account, ormanage instance typesfor the cloud account.

Configure an AWS Cloud

Configure an AWS Cloud

Configuring an AWS cloud is a four-step process:

- Add an AWS Cloud
- Add an AWS Region
- Configure an AWS Region
- Add an AWS Cloud Account

Add an AWS Cloud

To add an AWS cloud follow these steps.

1. Navigate to**Admin > Clouds**. This brings you to the Clouds page. If you, or another tenant admin in your tenant, have already added clouds to your tenant, they will be listed here. Click the**Add Cloud** linkin the upper right.
2. After clicking **Add Cloud**, the Add Cloud dialog box is displayed. Enter the**Cloud Name**and select the**cloud provider**.

   When assigning the Cloud Name, use a unique name. Ensure that you do not assign a previously-used name (including deleted cloud names) as the Cloud Name as they may cause deployments to fail intermittently.

3. After clicking**Next**, the second page of the **Add Clouds** dialog box, **Connectivity Settings**, appears. Set the toggle switches to configure the **Cloud Connectivity** settings.
   - When adding a public VM cloud in the CloudCenter Suite UI, the Cloud Connectivity Settings page, the second page of the Add Cloud dialog box, appears with a single toggle displayed:**Worker VMs Directly Connect with CloudCenter Suite**.
   - Setting this toggle to No implies you will install Cloud Remote for each region of this cloud. This also causes a second toggle to appear:**CloudCenter Suite Directly Accessible from Cloud Remote**.
   - Follow the table below for guidance on setting these toggles.

Toggle settings	Use case	Diagram
Worker VMs Directly Connect with CloudCenter Suite = Yes	Unimpeded connectivity exists between the CloudCenter Suite cluster and the cloud region API endpoint AND Unimpeded connectivity exists between the CloudCenter Suite cluster and worker VMs Cloud Remote is not required	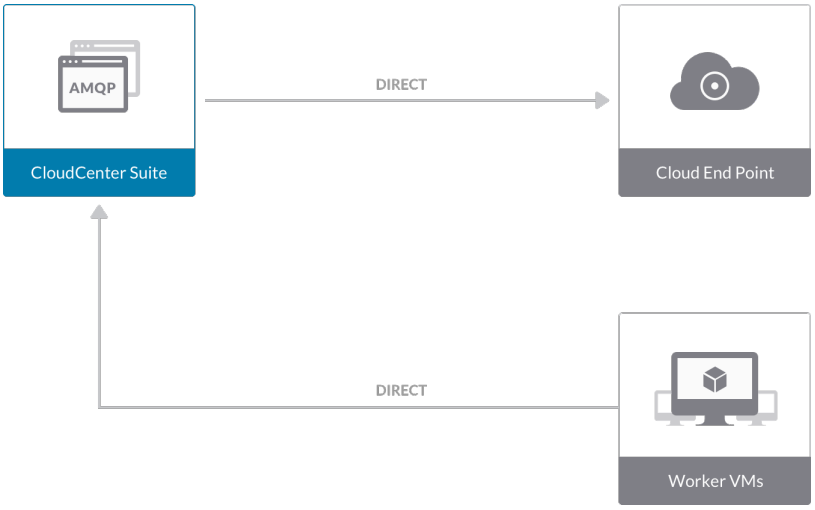

Worker VMs Directly Connect with CloudCenter Suite = No AND CloudCenter Suite Directly Accessible from Cloud Remote = Yes	Worker VMs need to communicate to the CloudCenter Suite cluster through Cloud Remote AND Cloud Remote can initiate the connection to the CloudCenter Suite cluster	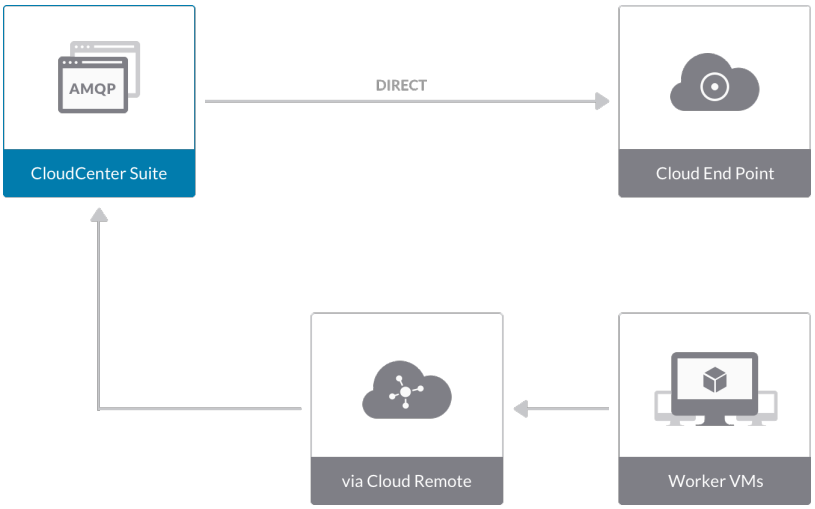
Worker VMs Directly Connect withCloudCenter Suite = No AND CloudCenter SuiteDirectly Accessible from Cloud Remote = No	Worker VMs need to communicate to the CloudCenter Suite cluster through Cloud Remote AND the CloudCenter Suite cluster cannot receive a connection initiated by Cloud Remote	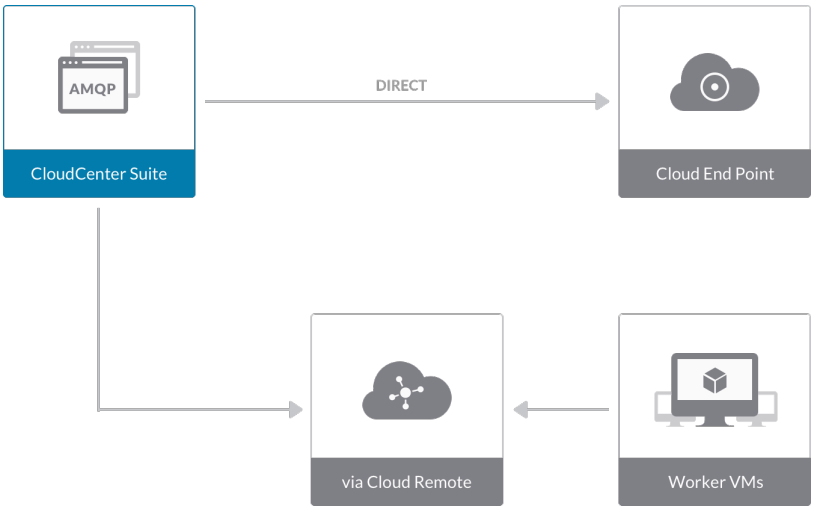

Note

The connectivity toggle settings set at the cloud level are inherited by each region you add to this cloud. However, it is possible to override these toggle settings on a per-region basis from the Regions tab for each region.

4. Click**Done**to save the configuration and close the dialog box. This brings you back to the **Clouds** page, and the cloud you just created will be added to the bottom of the list on the left side of the page.

Add an AWS Region

After creating an AWS cloud, the next step is to create the first region for the cloud. Follow these steps.

1. Navigate to the **Clouds** page and select the cloud you created on the left side of the screen. Then click the**Add Region**button on the right side of the screen.
2. After clicking the **Add Region** button, the Add Region dialog box is displayed. Select a region from the list and click **Save.**
3. After clicking**Save**you are brought back to the **Clouds** page with theregion you added shown on the right side of the page.

Configure an AWS Region

To configure a region you added to your AWS cloud, follow this procedure:

1. Navigate to Clouds page: **Admin > Clouds**. Find your AWS cloud from the cloud list on the left half of the screen and click its **Configure Cloud** link. This displays the Regions tab for this cloud as shown in the figure below with theCloud Settings section displayed first.

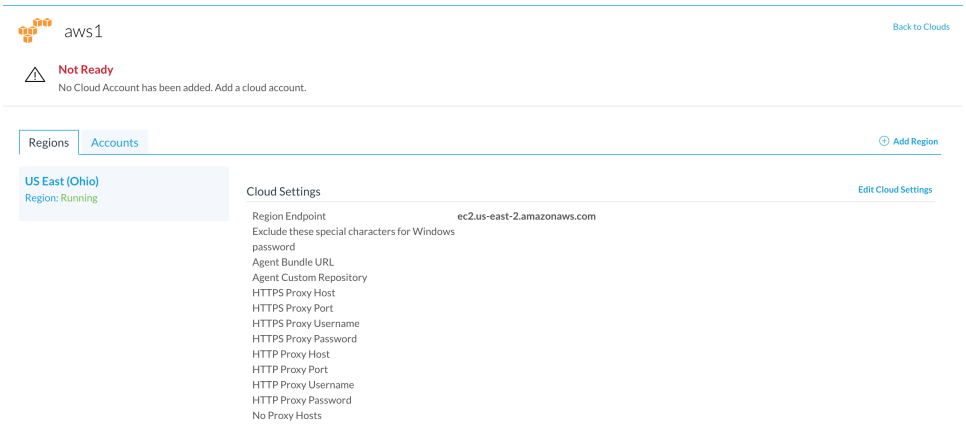

After you have added multiple regions to your AWS cloud, the Regions tab will show multiple individual region tabs on the left side of the screen. Click the tab of the region you want to configure.

2. Click the **Edit Cloud Settings**link in the upper right of the **Cloud Settings** section. This opens the **Configure Cloud Settings** dialog box.The **Cloud Settings** section contains fields that are unique to AWS and settings that are common to all cloud providers. Adjust these field values per the instructions in the following tables.

AWS Specific Cloud Settings

Field	Usage
Region Endpoint	This field is set by CloudCenter Suite based on the region location you selected from the Add Region dialog box.

Agnostic Cloud Settings

Field	Usage
Exclude these special characters for Windows password	When the Workload Manager agent is installed on a Windows worker VM, a special user account, called cliqruser, is created to support RDP sessions that may be initiated by the user through the Workload Manager UI. A Workload Manager process running on the CloudCenter Suite cluster creates a random password and passes it to the agent for creating the cliqruser account. Because some Windows deployments may restrict using certain characters for Windows passwords, this field is provided to tell the Workload Manager to exclude these special characters in the generation of the password for the cliqruser account.
Agent Bundle URL	If you plan to use a local repository to host the bundle store, you need to enter the URL of the local bundle store here. Otherwise, leave blank.
Agent Custom Repository	If you plan to use alocal repository to host the package store, you need to enter the URL of the local package store here. Otherwise, leave blank.
HTTP /HTTPS proxy fields (host, username, password)	If you require VMs in your region to access public addresses through a web proxy, enter the URL and credentials of the HTTP and HTTPS proxy servers in these fields.
No Proxy Hosts	If you have specified an HTTP or HTTP proxy using the above fields, you can specify that managed VMs in the region should bypass the proxy and connect directly to certain hosts. Use this field to create a comma-separated list of IP addresses or URLs that should be accessed directly. This field is ignored if an HTTP or HTTPS proxy is not specified.

Important information on proxy settings

In CloudCenter Suite, you can specify proxy settings at the region level, as described here, and at the suite level. To understand the expected behavior when proxy settings are specified at both levels, see the subsequent *Precedence of Proxy Settings* section.

Download Configuration and Encryption Key

After saving the Region Configuration settings in the Workload Manager or Cost Optimizer UI, you can download them to your local computer and then upload them to other conditional components such as Cloud Remote .

The Configuration and Encryption key is only visible when you have configured the Cloud Remote component. Click the **Download Configuration** link in the upper right of the Region Connectivity section, as shown in the following screenshot.

Clicking **DownloadConfiguration** causestwo things to happen:

- An encrypted zip file namedartifacts.zipis downloaded by your browser. Make a note of the location of this zip file as you will need if you are using Cloud Remote .
- The Region Connectivity section header updates to display a **Copy Encryption Key** link, as shown in the following screenshot.

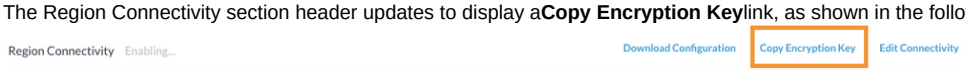

Click the **Copy Encryption Key** link to save the key to your clipboard. A success message will be displayed temporarily above the Region Connectivity section header. Make sure not to overwrite the clipboard with other data. You will need the key when you upload the configuration zip file to conditional components like Cloud Remote .

If you change the connectivity settings in the CloudCenter Suite UI and need to again download the zip file fromsoftware.cisco.com, use the automatically create a (new) encryption key, and copy the key to the clipboard by clicking the **Copy Encryption Key** link again.

Precedence of Proxy Settings

In CloudCenter Suite, you can specify HTTP and HTTPS proxy settings at the suite level as described inProxy Settings, and at the region level as described in *Agnostic Cloud Settings* section. The CloudCenter Suite cluster, Cloud Remote , and worker VMs will each use either the suite-level proxy settings or the region level proxy settings based on the conditions described below.

For Cloud Remote mode, whenever you change the region proxy settings, ensure you again download the region connectivity setting from the Workload Manager UI and reapply to Cloud Remote .

Mode	Proxy Settings Used By		
Mode (Is Cloud Remote used to communicate withcloud (or APIC) endpoint?)	CloudCenter Suite cluster to communicate with cloud (or APIC) endpoint	Cloud Remote to communicate with cloud endpoint	Worker VM to communicate with bundle store and package store
Cloud Remote Mode	NA	Region-level	Region-level
Non-Cloud Remote Mode	Suite-level	NA	Region-level

Not all clouds support all the proxy settings, and not all clouds support bothCloud Remote and non-Cloud Remote modes.Cloud Remote mode is not applicable to AWS but supports all proxy sections innon-Cloud Remote mode.
When you are done editing the settings in the dialog box, click **Save**.

3. Determine if you need Cloud Remote for this region.Scroll down to the **Region Connectivity** section for the region and click on the **Configure Region** link in the upper right to open the **Configure Region** dialog box.The toggle settings should be the same as when you set them on the connectivity page of the **Add Cloud** dialog box.If all of the connectivity togglesinthe **Region Connectivity** dialog boxare set to **Yes**, then Cloud Remote is NOT needed for this cloud region. In this case,you would normally leave the region connectivity settings at their current values and continue to the next settings section. The exceptionto this guidance is when a NAT firewall or proxy server exists between theCloudCenter Suite management cluster andworker VMs, or between the CloudCenter Suite management cluster and users that would use Workload Manager to initiate a Guacamole remote connection to a worker VM. In either of these cases, override the address fields in the **Region Connectivity** dialog box as explained below.

Networking Constraint	Field	Value
Worker VMs must use a proxy server or NAT firewall to access the"local" AMQP server running in the CloudCenter Suite cluster.	Worker AMQP IP Address	IP address and port number that the firewall or proxy server presents to the worker VMs on behalf of the "local" AMQP server running in the CloudCenter Suite cluster.
Users must use a proxy server or NAT firewall to access theGuacamole server running in the CloudCenter Suite cluster.	Guacamole Public IP Address and Port	IP address and port number that the firewall or proxy server presents to users on behalf of theGuacamole server running in the CloudCenter Suite cluster.
Worker VMs must use a proxy server or NAT firewall to access theGuacamole server running in the CloudCenter Suite cluster.	Guacamole IP Address and Port for Application VMs	IP address and port number that the firewall or proxy server presents to the worker VMs on behalf of the Guacamole server running in the CloudCenter Suite cluster.

Click **OK** to save the changes and dismiss the dialog box. You can now proceed to the next region settings section:VM Naming and IPAM Strategy.

4. If any of theconnectivity toggles inthe **Region Connectivity** dialog box are set to No, then you must install and configure Cloud Remote for this region.

Configure Cloud Remote in an AWS Region

Configure Cloud Remote in an AWS region as follows.

Obtain and Launch the Cloud Remote Appliance in AWS

a. Obtain the Cloud Remote shared AMI form Cisco support and launch it. Follow the same guidance for obtaining and launching theCloudCenter Suite installer appliance for AWS.
b. Optional but recommended for production environments: Deploy two additional instances of the appliance to form a cluster for HA.Cloud Remote includes support for the clustering of multiple nodes. You will "add" these two additional instances to the first instance after the first instance is configured. SeeCloud Remote (Conditional) > *Scaling*for details.
c. Once the first instance of the appliance has been launched, use your cloud console to**note its IP public and private addresses**. You will need this information later on in order to login to the Cloud Remote web UI and to complete the Region Connectivity settings in the CloudCenter Suite Web UI. Also, note the IP addresses of any other instances you launch.

Unable to render {include} The included page could not be found.

Specify AMQP and Guacamole Addresses for Supporting Cloud Remote

From the CloudCenter Suite UI, for the cloud region requiring Cloud Remote, navigate to the corresponding Regions or Details tab. Click the**Configure Region**link in the upper left of the Region Connectivity section to bring up the Configure Region dialog box. The toggle settings should be the same as when you set them on the connectivity page of the Add Cloud dialog box. You must update some of the address fields in the dialog box according to the scenarios summarized in the table below.

Toggle Settings	Field	Value
Worker VMs Directly Connect with CloudCenter = No AND CloudCenter Directly Accessible from Cloud Remote = Yes	Local AMQP IP Address	Pre-populated with the address and port number of the "local" AMQP server running in the CloudCenter Suite cluster. This address must be **accessible to Cloud Remote** . If Cloud Remote is accessing the CloudCenter Suite cluster through a user-supplied proxy server or NAT firewall, overwrite this field with the corresponding local AMQP IP address and port number provided by the user-supplied proxy server or NAT firewall and accessible to Cloud Remote .
Worker VMs Directly Connect with CloudCenter = No AND CloudCenter Directly Accessible from Cloud Remote = No	Remote AMQP IP Address	Enter <Cloud_Remote_IP>:<amqp_port>, where <Cloud_Remote_IP> = the IP address Cloud Remote which is **accessible to the CloudCenter Suite cluster**, and <amqp_port> = 5671 OR the custom AMQP port number you would later set with the **Change Ports shell script** on the Cloud Remote appliance (see Cloud Remote (Conditional) > *Custom Port Numbers (Conditional)*). If there is no user-supplied NAT firewall or proxy server between the CloudCenter Suite cluster and Cloud Remote , the IP address would be the public IP address of Cloud Remote . If there is a NAT firewall or proxy server between the CloudCenter Suite cluster and Cloud Remote , instead, enter the corresponding public IP address and port number that the firewall or proxy server presents to the internet on behalf of the "remote" AMQP server running in Cloud Remote .
Worker VMs Directly Connect with CloudCenter = No	Worker AMQP IP Address	Enter <Cloud_Remote_IP>:<amqp_port>, where <Cloud_Remote_IP> = the Cloud Remote IP address **accessible to the worker VMs**, and <amqp_port> = 5671 OR the custom AMQP port number you would later set with the **Change Ports shell script** on the Cloud Remote appliance (see Cloud Remote (Conditional) > *Custom Port Numbers (Conditional)*).
Worker VMs Directly Connect with CloudCenter = No	Guacamole Public IP and Port	Enter <Cloud_Remote_IP>:<guac_port>, where <Cloud_Remote_IP> = the Cloud Remote IP address **accessible to CloudCenter Suite users**, and <guac_port> = 8443 OR the custom Guacamole port number you would later set with the **Change Ports shell script** on the Cloud Remote appliance (see Cloud Remote (Conditional) > *Custom Port Numbers (Conditional)*).
Worker VMs Directly Connect with CloudCenter = No	Guacamole IP Address and Port for Application VMs	Enter <Cloud_Remote_IP>:<guac_port>, where <Cloud_Remote_IP> = the Cloud Remote IP address **accessible to worker VMs**, and <guac_port> = 7789

When done, click **OK** to save the setting and dismiss the dialog box.

Download Region Connectivity Settings and Upload to Cloud Remote

Cloud Remote uses the region connectivity settings set in the Workload Manager or Cost Optimizer UI. After saving the Region Configuration settings in the Workload Manager or Cost Optimizer UI, you must download them and to your local computer and then upload them to Cloud Remote as follows.

Click the**Download Configuration**link in the upper right of the Region Connectivity section, as shown in the figure below.

Clicking DownloadConfiguration causestwo things to happen:

- An encrypted zip file named**artifacts.zip**will be downloaded by your browser. Make note of the location of this zip file as you will need to upload it to Cloud Remote through the Cloud Remote web UI (see below).
- The Region Connectivity section header updates to display a **Copy Encryption Key**link, as shown in the figure below.

Click the**Copy Encryption Key**link to save the key to your clipboard. A success message will be displayed temporarily above the Region Connectivity section header. Make sure not to overwrite the clipboard with other data. You will need the key when you upload the configuration zip file to Cloud Remote .

> If you change the connectivity settings in the CloudCenter Suite UI and need to again download the zip file, a new encryption key is automatically created and can be copied to the clipboard by clicking the **Copy Encryption Key** link again.

After you have downloaded the zip file and copied the encryption key to your clipboard, login toCloud Remote web UI.

a. Open another browser tab and login to https://<Cloud Remote _ip> with thedefault credentials: admin/cisco.
b. You will immediately be required to change your password. Do so now.
c. You are now brought to the Cloud Remote home page as shown in the figure below.

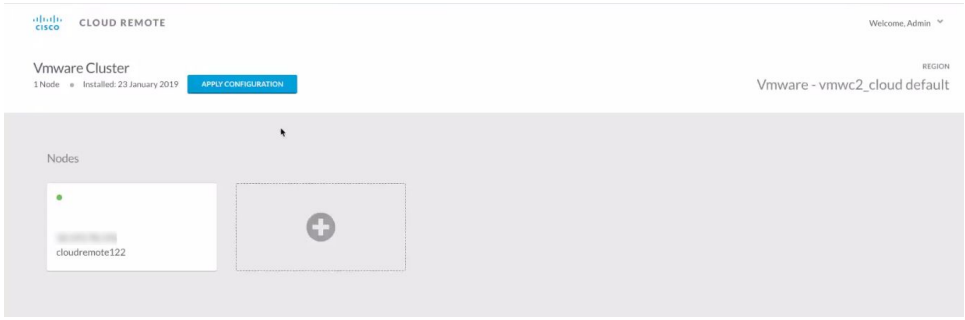

d. Click the **Apply Configuration** button in the page header. This prompts you to select a configuration file and enter the encryption key as shown in the figure below.

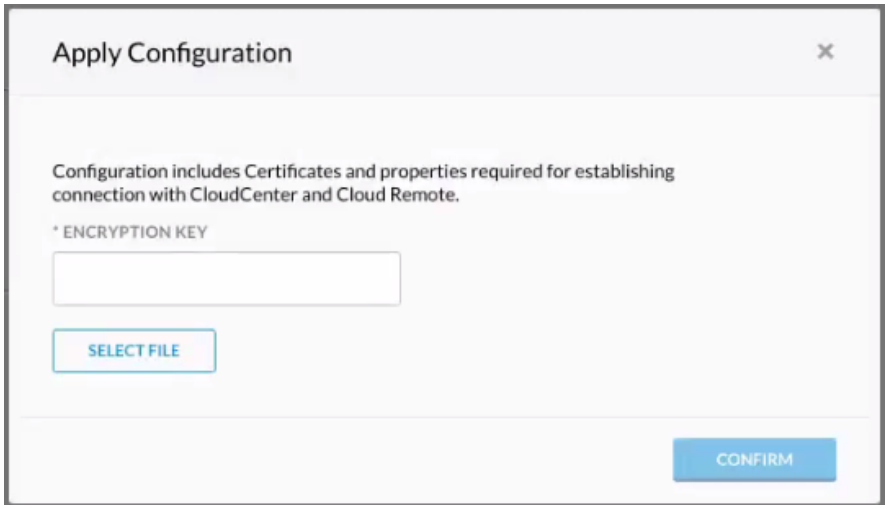

e. Paste the encryption key that was copied to the clipboard into the**Encryption Key**field in the dialog box.
f. Click**Select File**and browse tothe artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
g. Click**Confirm**.
h. Once the zip file is successfully transmitted and accepted, the Cloud Remote appliance attempts to establish communication with the CloudCenter Suite cluster and the Cloud Remote web UI home page is updated to show the name of the region it is connecting to in the upper right (see figure below).

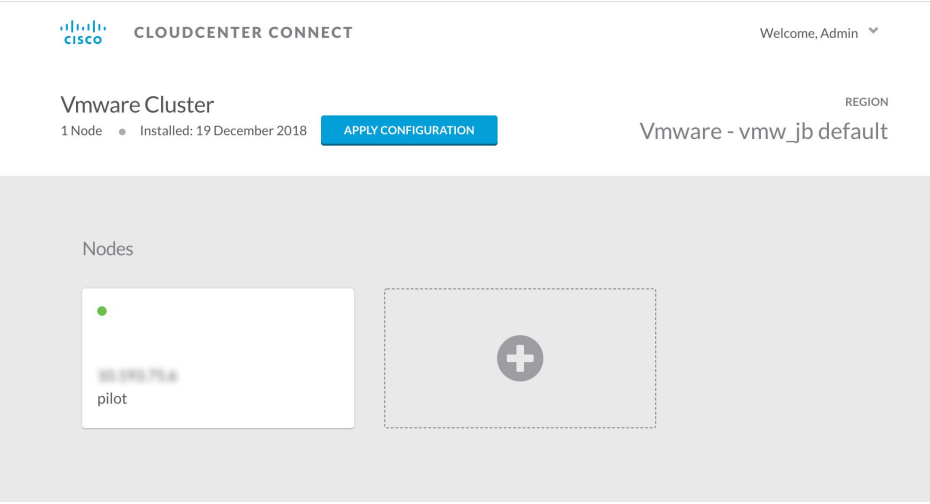

Switch your focus back to the Region Connectivity section of the target cloud region in the CloudCenter Suite web UI. The status indicator in the Region Connectivity section header will change from Not Configured to Running once connectivity between Cloud Remote and the CloudCenter Suite cluster is completely established (see figure below).

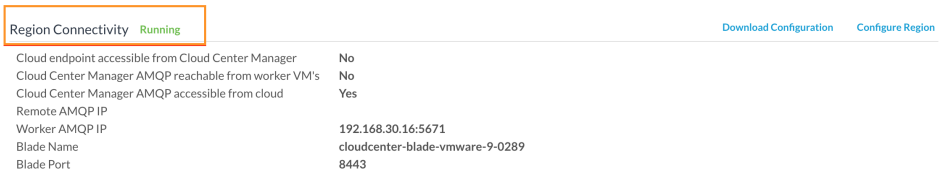

After completing these steps, Workload Manager and Cost Optimizer can use Cloud Remote for communicating with the target cloud region.

5. VM Naming and IPAM Strategy (conditional): Configure any VM naming or IPAM strategies in the Strategy section. If you leave the settings at the defaults, no IPAM strategy is applied and the default VM naming strategy is applied.

VM Naming and IPAM Strategies

Workload Manager lets you define a VM naming strategy and an IP address management (IPAM) strategy on a per cloud region basis. The VM naming strategy lets Workload Manager generate unique names for each of your VMs as they are launched. The IPAM strategy letsWorkl oad Manager generate IP addresses for each of the vNICs on each of your VMs. The VM naming strategy can reference a VM naming callout script, and the IPAM strategy requires that you specify an IP address allocation callout script andan IP address deallocation callout script. An IP address allocation callout scriptmay work in conjunction with a 3rd party IPAM server in your cloud environment.

Like External Lifecycle Actions, the scripts or commands associated with a VM naming strategy or IPAM strategy are executed in a containerized script execution engine running in the CloudCenter Suite cluster, if you are not using Cloud Remote in your region. If you are using Cloud Remote , these scripts or commands will execute in a containerized script execution engine running in Cloud Remote .

To set up the VM naming and IPAM strategies for your region, scroll down to the Strategy section of the Regions or Details tab and follow these steps.

Click the**Edit Strategy**link in the upper right to display the Edit Strategy dialog box as shown in the figure below.

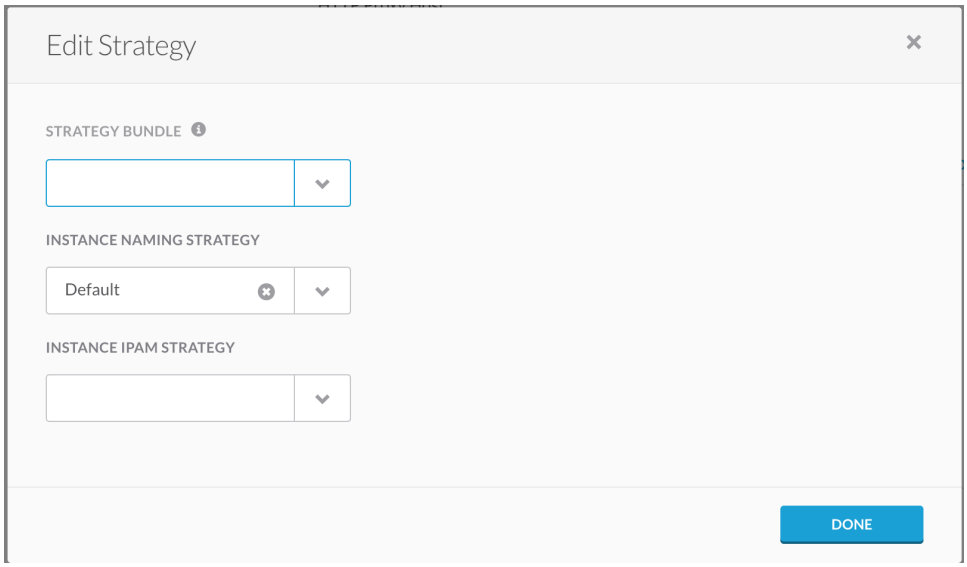

The dialog box can reveal a total of seven data entry fields as explained in the table below:

Field	Usage
Strategy Bundle	Zip file containing all of the scripts associated with your VM naming strategy or IPAM strategy. Select source from the dropdown an add the path in the new field that appears to the right. 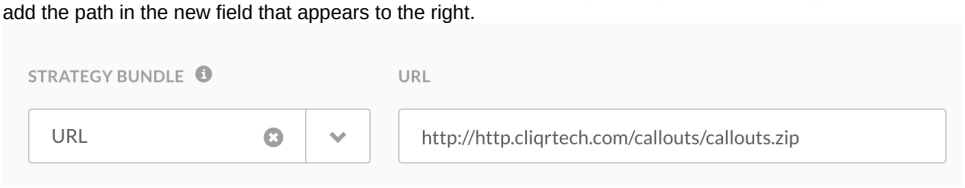
Instance Naming Strategy	Dropdown for selecting one of three naming strategies: Default, CliQr Macro Replacement, and Hostname Callout. If you select CliQ Macro Replacement, the **Node Name Config** field becomes visible and required. If you select Hostname Callout, the **Custom VM Name** field becomes visible and required. See the Instance Naming Strategy table below for details.

Node Name Config	Required if you selected "CliQr Macro Replacement" for theInstance Naming Strategyfield. Content: text string consisting of a combination of characters and numbers with embedded macro variables (see table below). The macros get translated to text strings run time.
	
	The resulting VM name string mustmeet the following requirements for job submission to succeed:
	• Only contain the following characters:
	  • Lowercase a to z
	  • 0 to 9
	  • Hyphen
	• Start withan alpha character
	• Cannot end with a hyphen

Instance IPAM Strategy	Dropdown menu with two choices: No IPAM (default), and IPAM Callouts. If you select the latter, two additional fields are displayed and become required: **IPAM Alloc Rule**and**IPAM Dealloc Rule**. (See below).
Custom VM Name	Location of your custom VM naming script if you set the**Instance Naming Strategy**field to "Custom Callout". Location options inclu "Script from bundle" if you specify a zip file in the **Strategy Bundle** field. See Guidance for Callout Scripts for guidance on developi custom VM naming scripts.
IPAM Alloc Rule	Location of your custom IPAM address allocation script if you set the**Instance IPAM Strategy**field to "IPAM Callouts". Location opti include "Script from bundle" if you specify a zip file in the **Strategy Bundle** field. See Guidance for Callout Scripts for guidance on developing custom IPAM callout scripts.
IPAM Dealloc Rule	Similar to the**IPAM Alloc Rule**above except this is where you specify the location of yourcustom IPAM address deallocation script. Guidance for Callout Scripts for guidance on developing custom IPAM callout scripts.

Instance Naming Strategy Options:

Options	Description
Default	The default method provided by Cisco is the generation of a string consisting of A predefined token prefixed with **cq jw-**
CliQr Macro Replacement	Selecting this method causes the **Node Name Config** field to be displayed and to become a required field. See the fields table above for details.
Hostname Callout	Selecting this method causes the **Custom VM Name** field to be displayed and to become a required field. See the fields table above for details.

OOB macro variables forCliQr Macro Replacement:

Macro	Returned Value
%os1%	First character of OS type string in lower case
%os2%	Firsttwo characters of OS type string in lower case
%OS1%	Firstcharacter of OS type string in upper case
%OS2%	Firsttwo character of OS type string in upper case
% RND<number>%	Fixed length random string, e.g., %RND10%, 10-character long random string ⚠ This is a mandatory field, only values in the range: %RND6% to %RND11% or just %RND% which defaults to %RND6%"
%UID%	User ID
%VID%	Vendor ID
%VM_PREFIX%	VM name prefix must be added as a global parameter in the app profile. ⚠ Only macro whose value can be user-customized.

6. External Lifecycle Actions (conditional):Specify any external lifecycle actions to be performed on all VMs launched by Workload Manager in this region.

External Lifecycle Actions Settings

Whenever you use Workload Manager to deploy a VM-based application, Workload Manager can execute external actions on VMs in the deployment as the VMs progress through these five phases of their lifecycle:

- Pre VM start
- Pre VM initialization
- PostVM initialization
- Pre VM Stop
- Post VM Stop

These lifecycle actions will execute in a script execution engine running in a container in the CloudCenter Suite cluster if no Cloud Remote appliance is configured for the region. If you have configured Cloud Remote for this region, the external lifecycle actions will executein a script execution engine running in a container in the Cloud Remote appliance.

External Lifecycle Action scripts can be defined at the service level, the application level, and the cloud region level. Use the External Lifecycle Actions section of the Region tab or Details tab (see figure below) tospecify external lifecycle actions at the region level.

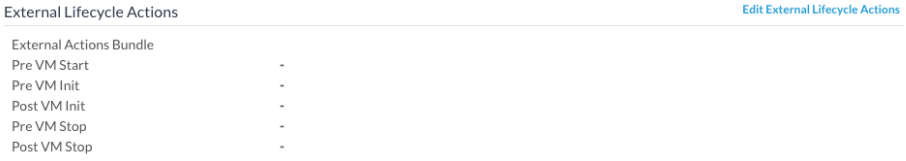

Click on the**Edit External Lifecycle Actions**link in the upper right. This displays theEdit External Lifecycle Actions dialog box as shown below.

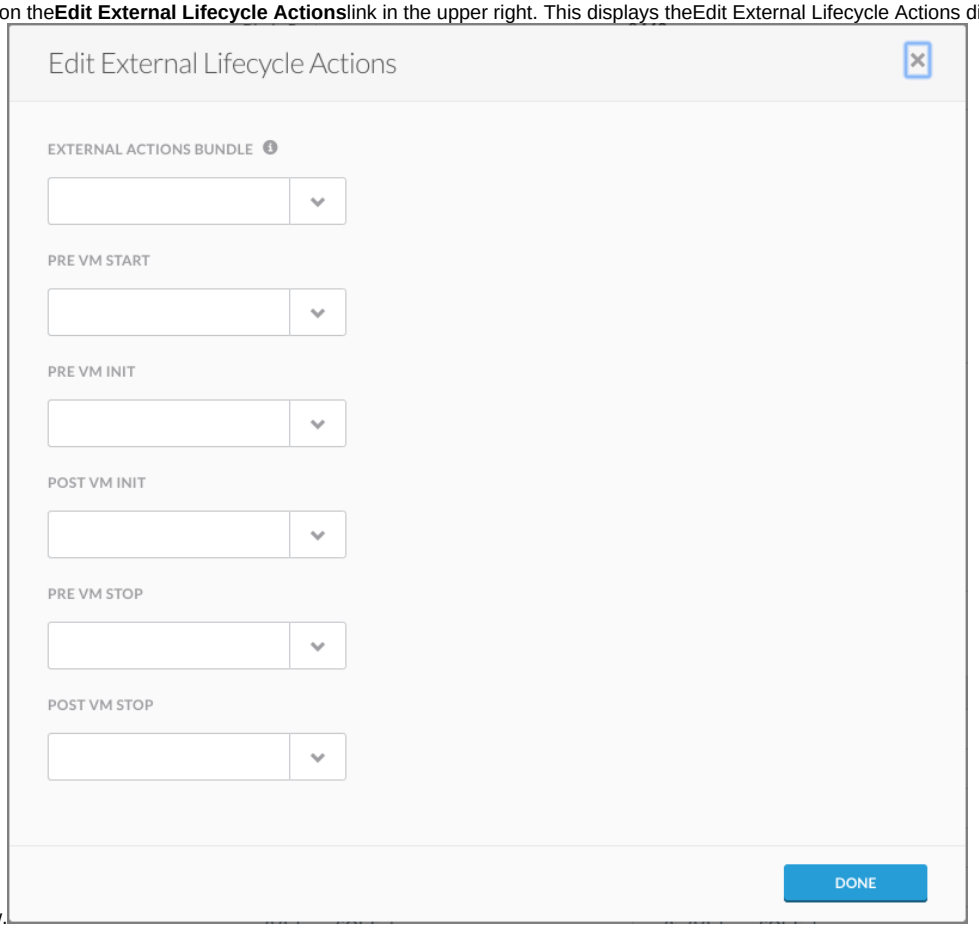

One way to specify an external lifecycle action is to use the first field,**External Actions Bundle**, to specify a zip file that contains all scripts for external lifecycle management. This zip file needs to contain a directory that contains all the scripts. Use the dropdown to select a repository location (URL, FTP, S3, etc.). This causes an additional field to be displayed to the right for entering the path to the zip file as shown in the figurebelow.

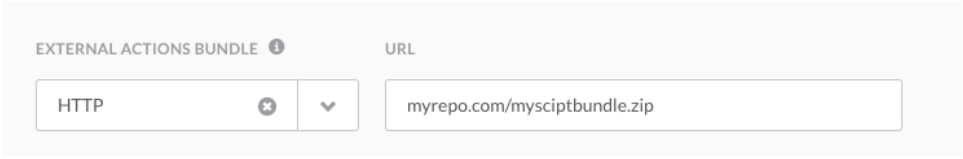

Once you specify actions bundle, for each of the individual lifecycle stage fields, select "Script from Bundle" from the dropdown. This displays an additional field to the right where you must specify the name of the script file in the bundle as shown in the figure below.

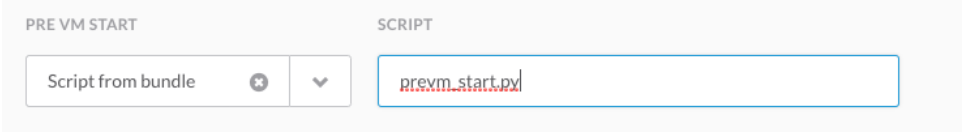

The scripts you specify can be bash scripts or Python scripts called from within a bash script.For more information on lifecycle scripts, seeDeployment Lifecycle Scripts.

7. Instance Types (informational): CloudCenter Suite automatically synchronizes instance types for public cloud regions on a daily basis. This data includes published pricing for each instance type. It is not possible to edit AWS region instance types.

Instance Types Settings

An instance type is a performance category of a VM for VM-based clouds or a performance category of a container for container-based clouds. The parameters for AWS cloud instance type is as follows:

- Display Name
- Price per hour
- Instance Type ID
- CPU

- 32- or 64-bit architecture
- RAM (MB)
- NICs
- Instance storage (GB)
- SSD support

For AWS cloud, the following apply:

- You can only sync, edit, and add **Auto-populate and Auto-sync** function.
- CloudCenter Suite auto-populates the instance types in the Instance Types section.
- CloudCenter Suite auto-syncs this information from the cloud provider once per 24 hours.

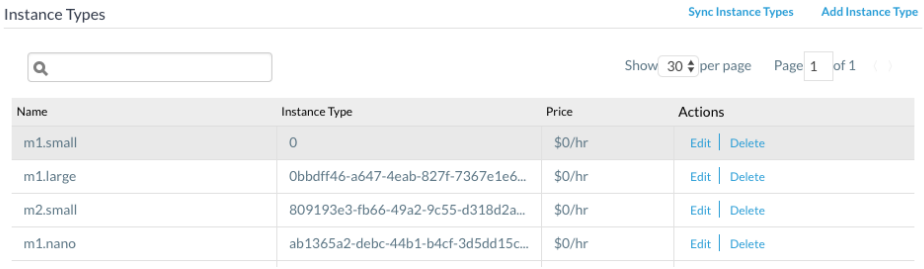

> You cannot force sync on demand nor edit the field of an instance type (except instance type ID) by clicking the**Edit**link in the **Actions** column in the list of instance types.

8. Storage Types (conditional): CloudCenter Suite automatically synchronizes storage types for public cloud regions on a daily basis. This data includes the cloud provider published pricing for each storage type. It is not possible to edit AWS region storage types.

   ### Storage Types Settings

   A storage type is a categorization of storage that can be provisioned for an application during deployment. It is a construct that applies to VM-based clouds only. Storage types are displayed in the Storage Types section of the Regions or Details tab.

   The storage type construct contains the following fields for all VM-based clouds:

   - Display Name
   - Cost Per Month per GB
   - Cloud Storage Type ID
   - Minimum Volume Size (GB)
   - Maximum Volume Size (GB)
   - Maximum Throughput (MB/s)
   - Maximum IOPS
   - Provisioned IOPS
   - IOPS Cost Per Month

   Storage types are handled by CloudCenter Suite in a fashion similar to instance types. Hence, the following apply:

   - You can only sync, edit, and add **Auto-populate and Auto-sync** function.
   - CloudCenter Suite auto-populates the instance types in the Instance Types section.
   - CloudCenter Suite auto-syncs this information from the cloud provider once per 24 hours.

   > Storage types cannot be edited.

9. Image Mappings: Image mappings allow services based on CloudCenter Suite logical images to be deployed using the appropriate physical image stored on the target cloud region. CloudCenter Suite automatically maps the OOB logical images to public cloud region physical images when you add the region to your cloud. Cisco periodically updates these mappings when new versions of OS physical images are uploaded by the cloud provider. To apply these updates to your region after it is added to your cloud, click the **Sync Image Mappings** link in the upper right of this section. If you create any custom logical images, you mustmanually import the corresponding physical images into your region and then map the corresponding logical images to these physical images. SeeImages for more context.

Add an AWS Cloud Account

Prerequisites

Before adding an AWS cloud account, do the following:

- Ensure the account has the minimum permissions.See Cloud Overview > *Minimum Permissions for Public Clouds*for additional details.

Configuration Process

To add an AWS cloud account, follow this procedure.

1. Locate your AWS cloud on the Clouds page and click theAdd Cloud Accountlink for this cloud.This displays the Add Cloud Account dialog box, as shown below.

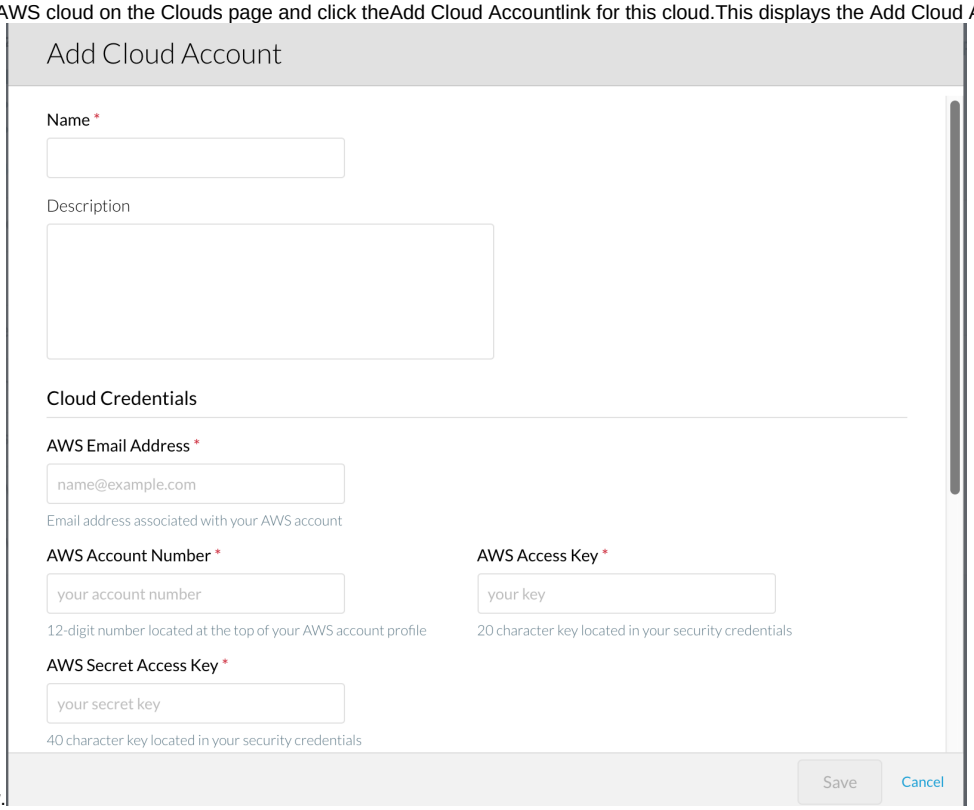

2. Assign a cloud account**Name**.

> **Tip**
>
> The name should not contain any space, dash, or special characters.

3. Provide the AWS cloud credentials:
    a. **AWS Email Address**: The email address associated with your AWS cloud account.
    b. **AWS Account Number**: The account number from your AWS account.
    c. **AWS Access Key and Secret Key**: The security credentials to access this AWS account.
4. Scroll the dialog box down and specify the location of your AWS account's billing reports: **S3 bucket region**, **S3 bucket name**, and **Report Path Prefix**, as shown in the figure below. For information on setting up billing information, seehttps://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/billing-reports-gettingstarted-s3.html.

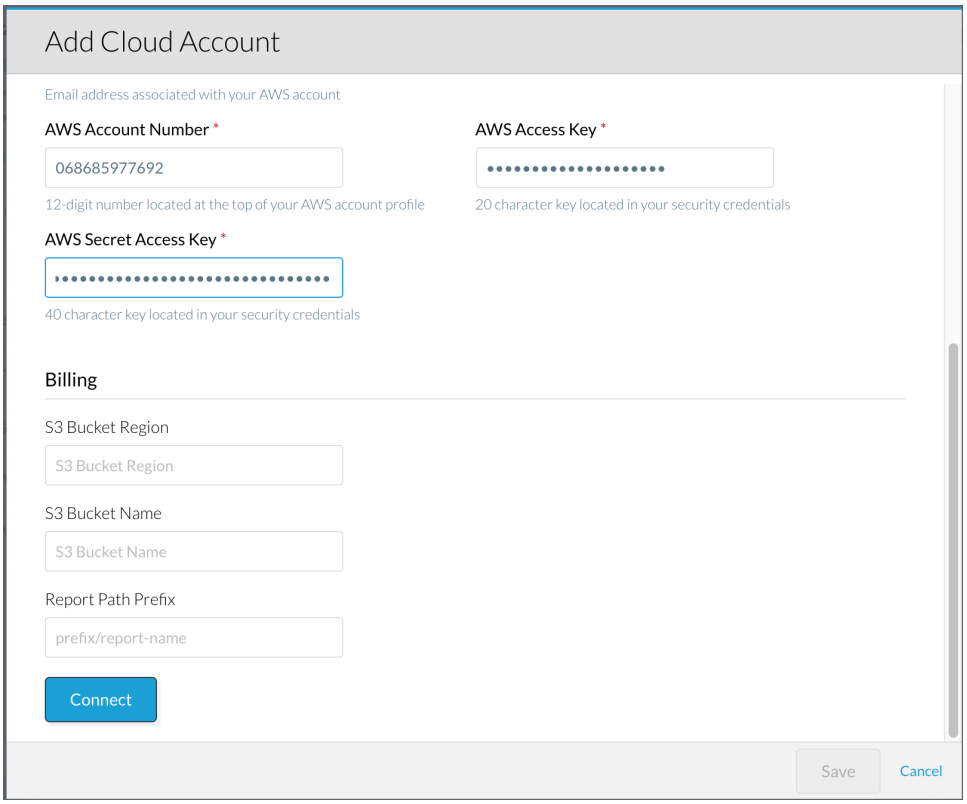

> ⚠ In the cloud console, create a bucket, if not already, and navigate to **Reports** to view billing information.

5. Click the **Connect** button. CloudCenter Suite will now attempt to validate your account credentials.
6. After the credentials are verified, the **Connect** button changes to an **Edit** button and two new fields appear, namely, **Enable Account For** and **Enable Reporting By Org Structure**.

    a. Set the **Enable Account For** dropdown per the table below.

Value	Usage
Provisioning	Workload Manager can deploy jobs using this account.
Reporting	Cost Optimizer and Workload Manager will track cloud costs for this account. Typical usage: master cloud accounts that are used for billing aggregation. ⚠ It is recommended that you do not add a *Reporting* account to the same tenant through different cloud groups. ⚠ Enabling a public cloud account for *Reporting* may incur expenses to retrieve cost data. These expenses are proportional to the number of configured cloud accounts and regions.
Provisioning, Reporting	Default. Account is used for both provisioning and reporting.

    b. **For AWS and Google clouds only**: Set the **Enable Reporting By Org Structure** toggle to On to cause Cost Optimizer to import the cost hierarchy created in the cloud provider portal. This saves the time of manually creating a comparable cost hierarchy within Cost Optimizer. See Cost Groups Configuration for more information on cost hierarchies in Cost Optimizer.
    c. Click the **Save** button when done.

> ⚠ You must enable **AWS Cost Explorer** to view AWS-specific costs on the Cost Optimizer dashboard. For additional details on enabling **AWS Cost Explorer**, see https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/ce-enable.html.

Cloud Accounts Tab

After you add cloud accounts to a cloud, they will appear in the Accounts tab for the cloud as shown in the figure below.

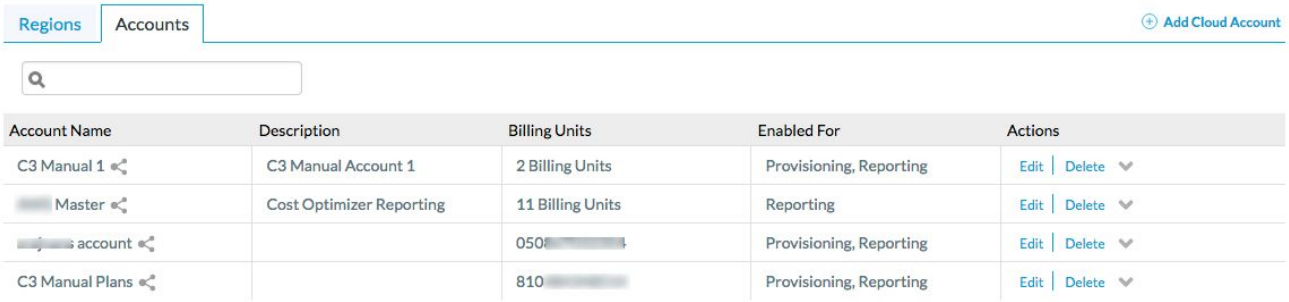

The Accounts tab contains columns for data entered when creating an account: Account Name, Description, Enabled For; and two additional columns:**Billing Units**and**Actions**.**Billing Units**is a dual function:

- If the cloud account contains only one billing unit, the ID for that billing unit is displayed.
- If the cloud account contains multiple billing units, such as an AWS master account, the number of billing units in that account is displayed followed by the text*Billing Units*.

A billing unit is the most granular level of cloud cost recording in CloudCenter Suite. The definition of a billing unit varies by a cloud provider as shown in the table below.

Cloud Provider	Billing Unit
AWS	Account ID
AzureRM	Subscription ID
Google	Project ID
IBM Cloud	Account ID
vCenter	Cloud Group Prefix - Datacenter Name
vCD	Organization Name
OpenStack	Project ID
Kubernetes	Namespace UID

The last column,**Actions**, contains links to let you edit or deleted the cloud account, ormanage instance typesfor the cloud account.

Configure an AzureRM Cloud

Configure an AzureRMCloud

Configuring an AzureRM cloud is a four-step process:

- Add an AzureRMCloud
- Add an AzureRMRegion
- Configure an AzureRMRegion
- Add an AzureRM Cloud Account

Add an AzureRMCloud

To add an AzureRMcloud, follow these steps.

1. Navigate to**Admin > Clouds**. This brings you to the Clouds page. If you, or another tenant admin in your tenant, have already added clouds to your tenant, they will be listed here. Click the**Add Cloud**linkin the upper right.
2. Click**Add Cloud.**The **Add Cloud** dialog box is displayed.
3. Enter the**Cloud Name**and select the**cloud provider**.

   When assigning the Cloud Name, use a unique name. Ensure that you do not assign a previously-used name (including deleted cloud names) as the Cloud Name as they may cause deployments to fail intermittently.

4. Click**Next.**The second page of the **Add Clouds** dialog box, **Connectivity Settings**, appears. Set the toggle to configure the Cloud Connectivity Settings.

   - When adding a public VM cloud in the CloudCenter Suite UI, the Cloud Connectivity Settings page, the second page of the Add Cloud dialog box, appears with a single toggle displayed:**Worker VMs Directly Connect with CloudCenter Suite**.
   - Setting this toggle to No implies you will install Cloud Remote for each region of this cloud. This also causes a second toggle to appear:**CloudCenter Suite Directly Accessible from Cloud Remote**.
   - Follow the table below for guidance on setting these toggles.

Toggle settings	Use case	Diagram
Worker VMs Directly Connect with CloudCenter Suite = Yes	Unimpeded connectivity exists between the CloudCenter Suite cluster and the cloud region API endpoint AND Unimpeded connectivity exists between the CloudCenter Suite cluster and worker VMs Cloud Remote is not required	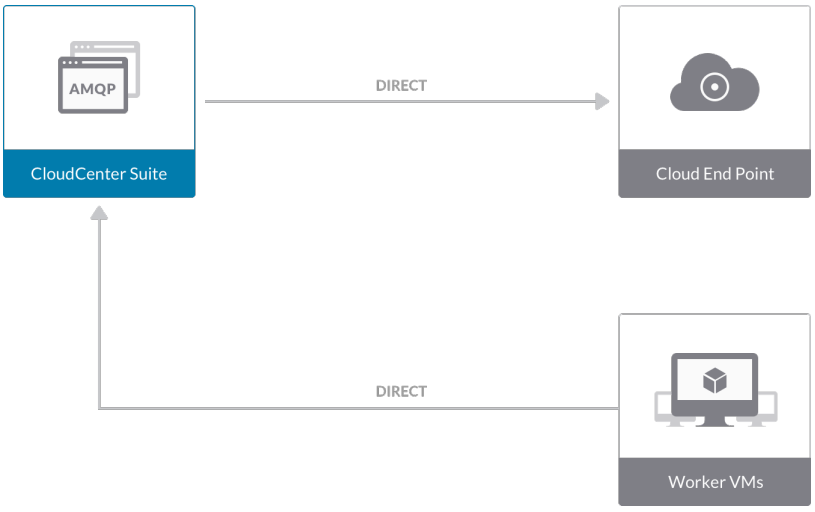

Worker VMs Directly Connect with CloudCenter Suite = No AND CloudCenter Suite Directly Accessible from Cloud Remote = Yes	Worker VMs need to communicate to the CloudCenter Suite cluster through Cloud Remote AND Cloud Remote can initiate the connection to the CloudCenter Suite cluster	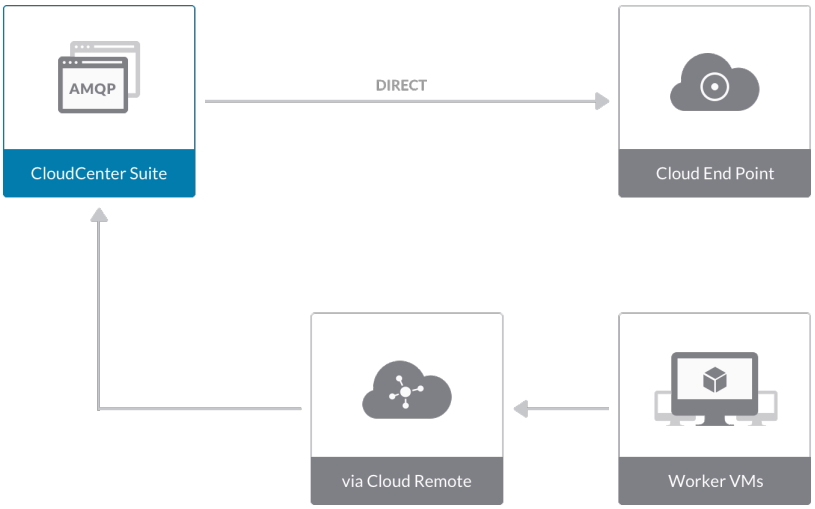
Worker VMs Directly Connect withCloudCenter Suite = No AND CloudCenter SuiteDirectly Accessible from Cloud Remote = No	Worker VMs need to communicate to the CloudCenter Suite cluster through Cloud Remote AND the CloudCenter Suite cluster cannot receive a connection initiated by Cloud Remote	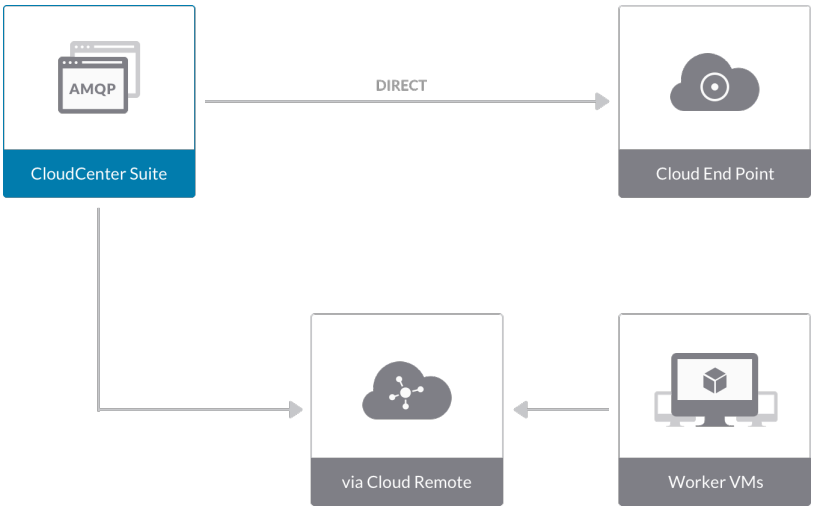

Note

The connectivity toggle settings set at the cloud level are inherited by each region you add to this cloud. However, it is possible to override these toggle settings on a per-region basis from the Regions tab for each region.

5. Click**Done**to save the configuration and close the dialog box. This brings you back to the **Clouds** page, and the cloud you just created will be added to the bottom of the list on the left side of the page.

Add an AzureRMRegion

After creating an AzureRMcloud, the next step is to create the first region for the cloud. Follow these steps.

1. Navigate to the **Clouds** page and select the cloud you created on the left side of the screen.
2. Click the**Add Region**button onthe right side of the screen. The Add Region dialog box is displayed.
3. Select a region from the list and click**Save**. You are back to the **Clouds** page with theregion you added shown on the right side of the page.

Configure an AzureRMRegion

To configure a region you added to your AzureRM cloud, follow this procedure.

1. Navigate to Clouds page:**Admin > Clouds**. Find your AzureRM cloud from the cloud list on the left half of the screen and click its**Configure Cloud** link. This displays the **Regions** tab for this cloud, as shown in the figure below, with theCloud Settings section displayed first.

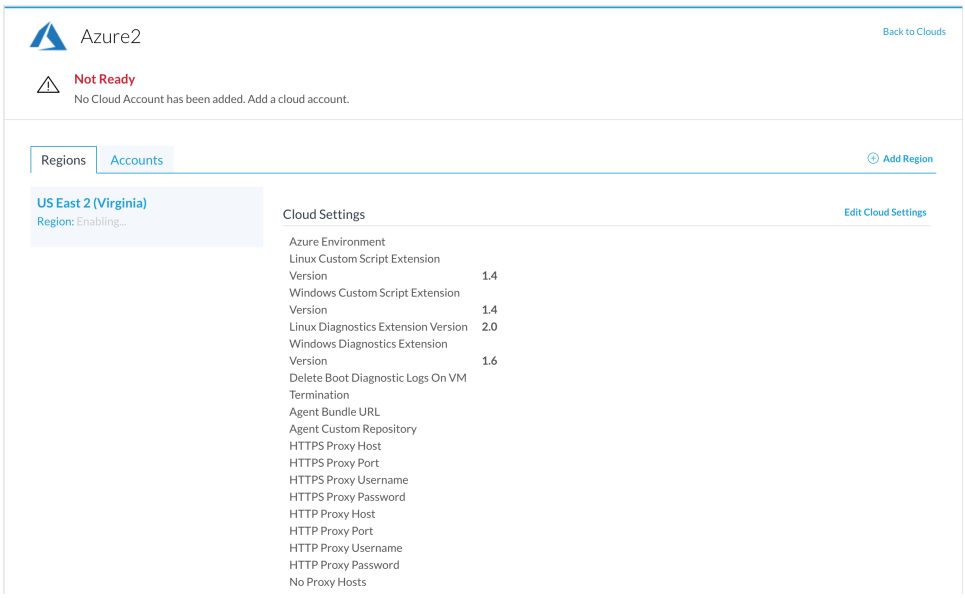

After you have added multiple regions to your AzureRM cloud, the **Regions** tab will show multiple individual region tabs on the left side of the screen. Click the tab of the region you want to configure.

2. Click the**Edit Cloud Settings**link in the upper right of the **Cloud Settings** section. This opens the **Configure Cloud Settings** dialog box. The **Cloud Settings** section contains fields that are unique to AzureRM and settings that are common to all cloud providers. Adjust these field values per the instructions in the following tables.

AzureRM Specific Cloud Settings

Field	Usage
Azure Environment	Automatically set by CloudCenter Suite based on the region you selected, but it can be overridden by using the dropdown list.
Linux and Windows extension versions	The custom script extensions are provided by Microsoft to support dynamic bootstrapping. The diagnostics extension is provided by Microsoft to support metrics monitoring.These four fields are set to recommended values by default by CloudCenter Suite,but you can override them.
Delete Boot Diagnostic Logs On VM Termination	AzureRM will store VM boot diagnostic logs after a VM terminates. CloudCenter Suite sets this value to false by default, but you can change the value to True using the dropdown.

Agnostic Cloud Settings

Field	Usage
Exclude these special characters for Windows password	When the Workload Manager agent is installed on a Windows worker VM, a special user account, called cliqruser, is created to support RDP sessions that may be initiated by the user through the Workload Manager UI. A Workload Manager process running on the CloudCenter Suite cluster creates a random password and passes it to the agent for creating the cliqruser account. Because some Windows deployments may restrict using certain characters for Windows passwords, this field is provided to tell the Workload Manager to exclude these special characters in the generation of the password for the cliqruser account.
Agent Bundle URL	If you plan to use a local repository to host the bundle store, you need to enter the URL of the local bundle store here. Otherwise, leave blank.
Agent Custom Repository	If you plan to use alocal repository to host the package store, you need to enter the URL of the local package store here. Otherwise, leave blank.
HTTP /HTTPS proxy fields (host, username, password)	If you require VMs in your region to access public addresses through a web proxy, enter the URL and credentials of the HTTP and HTTPS proxy servers in these fields.

No Proxy Hosts	If you have specified an HTTP or HTTP proxy using the above fields, you can specify that managed VMs in the region should bypass the proxy and connect directly to certain hosts. Use this field to create a comma-separated list of IP addresses or URLs that should be accessed directly. This field is ignored if an HTTP or HTTPS proxy is not specified.

Important information on proxy settings

In CloudCenter Suite, you can specify proxy settings at the region level, as described here, and at the suite level. To understand the expected behavior when proxy settings are specified at both levels, see the subsequent *Precedence of Proxy Settings* section.

Download Configuration and Encryption Key

After saving the Region Configuration settings in the Workload Manager or Cost Optimizer UI, you can download them to your local computer and then upload them to other conditional components such as Cloud Remote .

The Configuration and Encryption key is only visible when you have configured the Cloud Remote component. Click the**Download Configuration** link in the upper right of the Region Connectivity section, as shown in the following screenshot.

Clicking**DownloadConfiguration**causestwo things to happen:

- An encrypted zip file namedartifacts.zipis downloaded by your browser. Make a note of the location of this zip file as you will need if you are using Cloud Remote .
- The Region Connectivity section header updates to display a**Copy Encryption Key**link, as shown in the following screenshot.

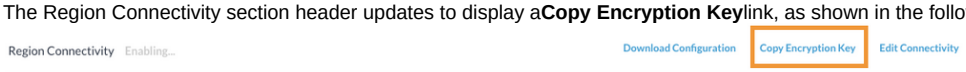

Click the**Copy Encryption Key**link to save the key to your clipboard. A success message will be displayed temporarily above the Region Connectivity section header. Make sure not to overwrite the clipboard with other data. You will need the key when you upload the configuration zip file to conditional components like Cloud Remote .

If you change the connectivity settings in the CloudCenter Suite UI and need to again download the zip file fromsoftware.cisco.com, use the automatically create a (new) encryption key, and copy the key to the clipboard by clicking the**Copy Encryption Key**link again.

Precedence of Proxy Settings

In CloudCenter Suite, you can specify proxy settings at the suite level as described inProxy Settings, and at the region level as described inthe*Agnostic Cloud Settings*section. The CloudCenter Suite cluster, Cloud Remote , and worker VMs will each use either the suite-level proxy settings or the region level proxy settings based on the conditions described below.

For Cloud Remote mode, whenever you change the region proxy settings, ensure you again download the region connectivity setting from the Workload Manager UI and reapply to Cloud Remote .

Mode	Proxy Settings Used By		
Mode (Is Cloud Remote used to communicate withcloud (or APIC) endpoint?)	CloudCenter Suite cluster to communicate with cloud (or APIC) endpoint	Cloud Remote to communicate with cloud endpoint	Worker VM to communicate with bundle store and package store
Cloud Remote Mode	NA	Region-level	Region-level
Non-Cloud Remote Mode	Suite-level	NA	Region-level

Not all clouds support all the proxy settings, and not all clouds support bothCloud Remote and non-Cloud Remote modes.Cloud Remote modeis not applicable to AzureRM.In non-Cloud Remote mode, AzureRM supports all proxy settings, except HTTP/HTTPS Username/Password, and proxy hosts.

When you are done editing the settings in the dialog box, click**Save**.

3. Determine if you need Cloud Remote for this region. Scroll down to the **Region Connectivity** section for the region and click on the**Configure Region**link in the upper right to open the **Configure Region** dialog box.The toggle settings should be the same as when you set them on the connectivity page of the **Add Cloud** dialog box.If all of the connectivity togglesinthe **Region Connectivit**y dialog boxare set to **Yes**, then Cloud Remote is NOT needed for this cloud region. In this case,you would normally leave the region connectivity settings at their current values and continue to the next settings section.

The exceptionto this guidance is when a NAT firewall or proxy server exists between theCloudCenter Suite management cluster andworker VMs, or between the CloudCenter Suite management cluster and users that would use Workload Manager to initiate a Guacamole remote connection to a worker VM. In either of these cases, override the address fields in the **Region Connectivity** dialog box, as explained below.

Networking Constraint	Field	Value
Worker VMs must use a proxy server or NAT firewall to access the"local" AMQP server running in the CloudCenter Suite cluster.	Worker AMQP IP Address	IP address and port number that the firewall or proxy server presents to the worker VMs on behalf of the "local" AMQP server running in the CloudCenter Suite cluster.

Users must use a proxy server or NAT firewall to access theGuacamole server running in the CloudCenter Suite cluster.	Guacamole Public IP Address and Port	IP address and port number that the firewall or proxy server presents to users on behalf of theGuacamole server running in the CloudCenter Suite cluster.
Worker VMs must use a proxy server or NAT firewall to access theGuacamole server running in the CloudCenter Suite cluster.	Guacamole IP Address and Port for Application VMs	IP address and port number that the firewall or proxy server presents to the worker VMs on behalf of the Guacamole server running in the CloudCenter Suite cluster.

4. Click**OK**to save the changes and dismiss the dialog box. You can now proceed to the next region settings section:VM Naming and IPAM Strategy.
If any of theconnectivity togglesinthe **Region Connectivity** dialog box are set to No, then**you must install and configure Cloud Remote for this region**.

Cloud Remote for AzureRM

Follow these steps to obtain, launch, and configure Cloud Remote for an AzureRM region.

Download and Launch the Cloud Remote Appliance in AzureRM

a. Download the Cloud Remote appliance for AzureRM as a zip file fromsoftware.cisco.comand then unzip it to reveal the VHD file.
b. Upload the Cloud Remote appliance VHD file to AzureRM using the AzureRM CLI, then launch the appliance from the AzureRM console web UI. This process is similar to uploading and launching theCloudCenter Suite installer appliance for AzureRM.

> You must use the AzureRM CLI to perform this upload.

c. Optional but recommended for production environments: Deploy two additional instances of the appliance to form a cluster for HA.Cloud Remote includes support for the clustering of multiple nodes. You will "add" these two additional instances to the first instance after the first instance is configured. SeeCloud Remote (Conditional)> *Scaling*for details.
d. Once the first instance of the appliance has been launched, use theAzureRM console to**note its IP public and private addresses**. You will need this information later on in order to login to the Cloud Remote web UI and to complete the Region Connectivity settings in the CloudCenter Suite Web UI. Also, note the IP addresses of any other appliances you launch.

Unable to render {include} The included page could not be found.

Specify AMQP and Guacamole Addresses for Supporting Cloud Remote

From the CloudCenter Suite UI, for the cloud region requiring Cloud Remote, navigate to the corresponding Regions or Details tab. Click the**Configure Region**link in the upper left of the Region Connectivity section to bring up the Configure Region dialog box. The toggle settings should be the same as when you set them on the connectivity page of the Add Cloud dialog box. You must update some of the address fields in the dialog box according to the scenarios summarized in the table below.

Toggle Settings	Field	Value
Worker VMs Directly Connect with CloudCenter = No AND CloudCenter Directly Accessible from Cloud Remote = Yes	Local AMQP IP Address	Pre-populated with the address and port number of the "local" AMQP server running in the CloudCenter Suite cluster. This address must be **accessible to Cloud Remote** . If Cloud Remote is accessing the CloudCenter Suite cluster through a user-supplied proxy server or NAT firewall, overwrite this field with the corresponding local AMQP IP address and port number provided by the user-supplied proxy server or NAT firewall and accessible to Cloud Remote .
Worker VMs Directly Connect with CloudCenter = No AND CloudCenter Directly Accessible from Cloud Remote = No	Remote AMQP IP Address	Enter <Cloud_Remote_IP>:<amqp_port>, where <Cloud_Remote_IP> = the IP address Cloud Remote which is **accessible to the CloudCenter Suite cluster**, and <amqp_port> = 5671 OR the custom AMQP port number you would later set with the **Change Ports shell script** on the Cloud Remote appliance (see Cloud Remote (Conditional) > *Custom Port Numbers (Conditional)*). If there is no user-supplied NAT firewall or proxy server between the CloudCenter Suite cluster and Cloud Remote , the IP address would be the public IP address of Cloud Remote . If there is a NAT firewall or proxy server between the CloudCenter Suite cluster and Cloud Remote , instead, enter the corresponding public IP address and port number that the firewall or proxy server presents to the internet on behalf of the "remote" AMQP server running in Cloud Remote .

Worker VMs Directly Connect with CloudCenter = No	Worker AMQP IP Address	Enter <Cloud_Remote_IP>:<amqp_port>, where <Cloud_Remote_IP> = the Cloud Remote IP address **accessible to the worker VMs**, and <amqp_port> = 5671 OR the custom AMQP port number you would later set with the **Change Ports shell script** on the Cloud Remote appliance (see Cloud Remote (Conditional) > *Custom Port Numbers (Conditional)*).
Worker VMs Directly Connect with CloudCenter = No	Guacamole Public IP and Port	Enter <Cloud_Remote_IP>:<guac_port>, where <Cloud_Remote_IP> = the Cloud Remote IP address **accessible to CloudCenter Suite users**, and <guac_port> = 8443 OR the custom Guacamole port number you would later set with the **Change Ports shell script** on the Cloud Remote appliance (see Cloud Remote (Conditional) > *Custom Port Numbers (Conditional)*).
Worker VMs Directly Connect with CloudCenter = No	Guacamole IP Address and Port for Application VMs	Enter <Cloud_Remote_IP>:<guac_port>, where <Cloud_Remote_IP> = the Cloud Remote IP address **accessible to worker VMs**, and <guac_port> = 7789

When done, click **OK** to save the setting and dismiss the dialog box.

Download Region Connectivity Settings and Upload to Cloud Remote

Cloud Remote uses the region connectivity settings set in the Workload Manager or Cost Optimizer UI. After saving the Region Configuration settings in the Workload Manager or Cost Optimizer UI, you must download them and to your local computer and then upload them to Cloud Remote as follows.

Click the**Download Configuration**link in the upper right of the Region Connectivity section, as shown in the figure below.

Clicking DownloadConfiguration causestwo things to happen:

- An encrypted zip file named**artifacts.zip**will be downloaded by your browser. Make note of the location of this zip file as you will need to upload it to Cloud Remote through the Cloud Remote web UI (see below).
- The Region Connectivity section header updates to display a **Copy Encryption Key**link, as shown in the figure below.

Click the**Copy Encryption Key**link to save the key to your clipboard. A success message will be displayed temporarily above the Region Connectivity section header. Make sure not to overwrite the clipboard with other data. You will need the key when you upload the configuration zip file to Cloud Remote .

> ⚠ If you change the connectivity settings in the CloudCenter Suite UI and need to again download the zip file, a new encryption key is automatically created and can be copied to the clipboard by clicking the **Copy Encryption Key** link again.

After you have downloaded the zip file and copied the encryption key to your clipboard, login toCloud Remote web UI.

a. Open another browser tab and login to https://<Cloud Remote _ip> with thedefault credentials: admin/cisco.
b. You will immediately be required to change your password. Do so now.
c. You are now brought to the Cloud Remote home page as shown in the figure below.

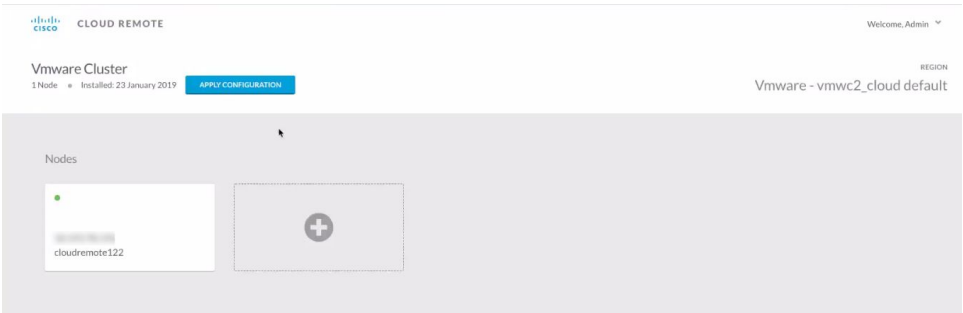

d. Click the **Apply Configuration** button in the page header. This prompts you to select a configuration file and enter the encryption key as shown in the figure below.

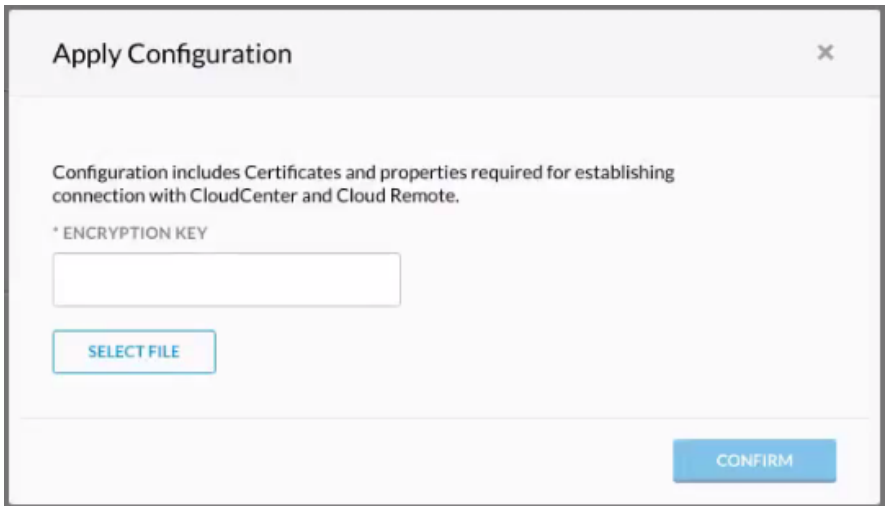

e. Paste the encryption key that was copied to the clipboard into the**Encryption Key**field in the dialog box.
f. Click**Select File**and browse tothe artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
g. Click**Confirm**.
h. Once the zip file is successfully transmitted and accepted, the Cloud Remote appliance attempts to establish communication with the CloudCenter Suite cluster and the Cloud Remote web UI home page is updated to show the name of the region it is connecting to in the upper right (see figure below).

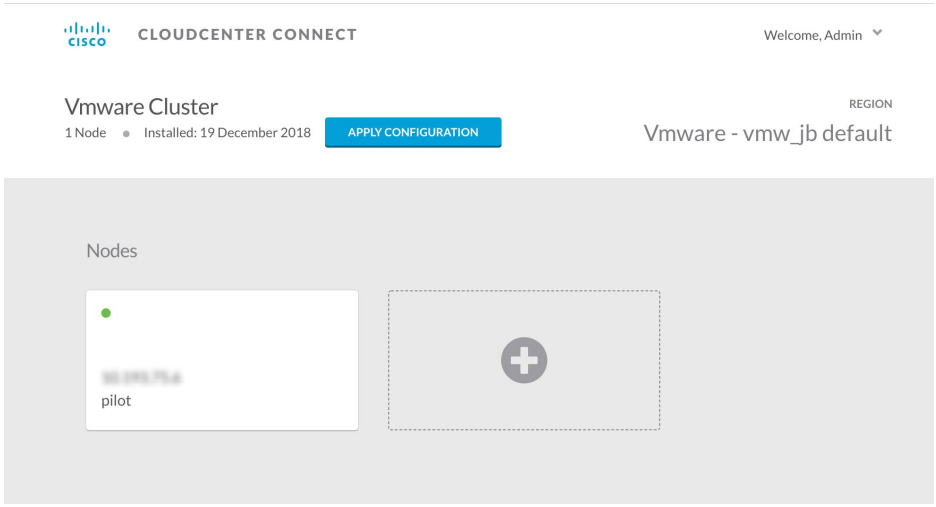

Switch your focus back to the Region Connectivity section of the target cloud region in the CloudCenter Suite web UI. The status indicator in the Region Connectivity section header will change from Not Configured to Running once connectivity between Cloud Remote and the CloudCenter Suite cluster is completely established (see figure below).

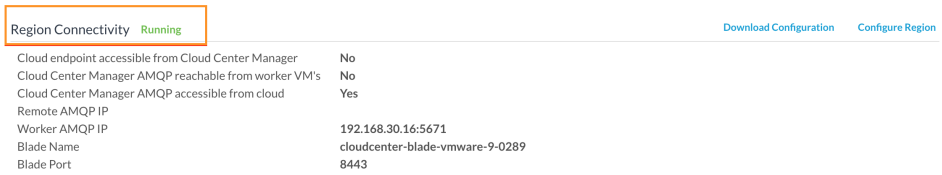

After completing these steps, Workload Manager and Cost Optimizer can use Cloud Remote for communicating with the target cloud region.

5. VM Naming and IPAM Strategy (conditional): Configure any VM naming strategy in the Strategy section. CloudCenter Suite currently does not support an IPAM strategy for AzureRM. If you leave the settings at the defaults, the default VM naming strategy is applied.

VM Naming and IPAM Strategies

Workload Manager lets you define a VM naming strategy and an IP address management (IPAM) strategy on a per cloud region basis. The VM naming strategy lets Workload Manager generate unique names for each of your VMs as they are launched. The IPAM strategy letsWorkl oad Manager generate IP addresses for each of the vNICs on each of your VMs. The VM naming strategy can reference a VM naming callout script, and the IPAM strategy requires that you specify an IP address allocation callout script andan IP address deallocation callout script. An IP address allocation callout scriptmay work in conjunction with a 3rd party IPAM server in your cloud environment.

Like External Lifecycle Actions, the scripts or commands associated with a VM naming strategy or IPAM strategy are executed in a containerized script execution engine running in the CloudCenter Suite cluster, if you are not using Cloud Remote in your region. If you are using Cloud Remote , these scripts or commands will execute in a containerized script execution engine running in Cloud Remote .

To set up the VM naming and IPAM strategies for your region, scroll down to the Strategy section of the Regions or Details tab and follow these steps.

Click the**Edit Strategy**link in the upper right to display the Edit Strategy dialog box as shown in the figure below.

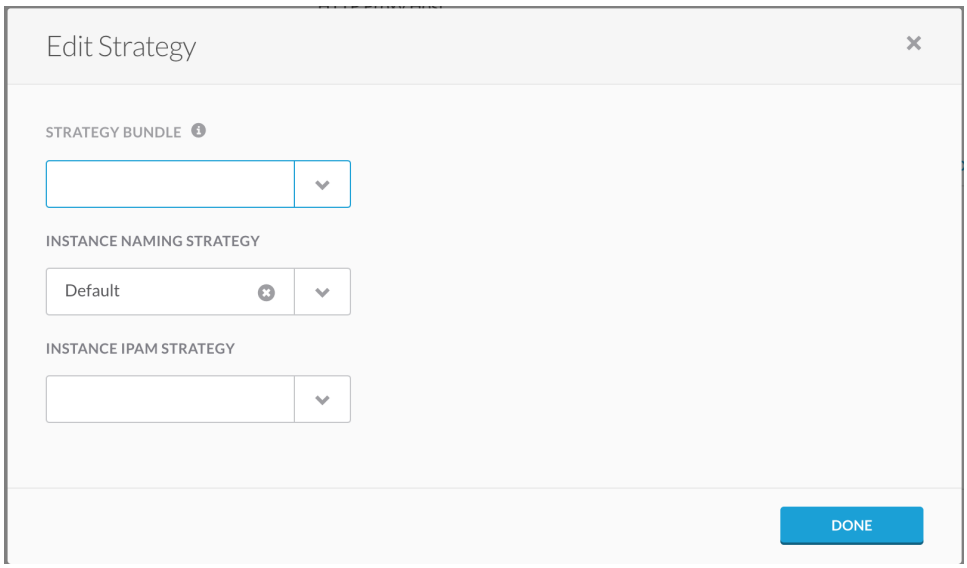

The dialog box can reveal a total of seven data entry fields as explained in the table below:

Field	Usage
Strategy Bundle	Zip file containing all of the scripts associated with your VM naming strategy or IPAM strategy. Select source from the dropdown an add the path in the new field that appears to the right. 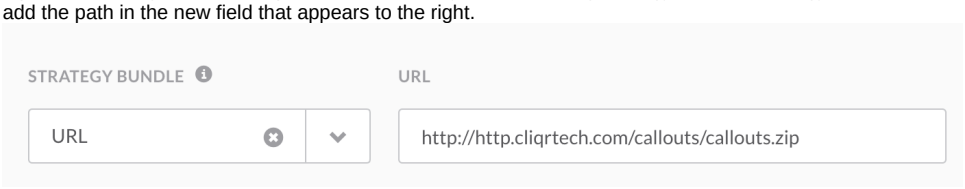
Instance Naming Strategy	Dropdown for selecting one of three naming strategies: Default, CliQr Macro Replacement, and Hostname Callout. If you select CliQ Macro Replacement, the **Node Name Config** field becomes visible and required. If you select Hostname Callout, the **Custom VM Name** field becomes visible and required. See the Instance Naming Strategy table below for details.

Node Name Config	Required if you selected "CliQr Macro Replacement" for theInstance Naming Strategyfield. Content: text string consisting of a combination of characters and numbers with embedded macro variables (see table below). The macros get translated to text string run time.
	
	The resulting VM name string mustmeet the following requirements for job submission to succeed:
	• Only contain the following characters:
	  • Lowercase a to z
	  • 0 to 9
	  • Hyphen
	• Start withan alpha character
	• Cannot end with a hyphen

Instance IPAM Strategy	Dropdown menu with two choices: No IPAM (default), and IPAM Callouts. If you select the latter, two additional fields are displayed and become required: **IPAM Alloc Rule**and**IPAM Dealloc Rule**. (See below).
Custom VM Name	Location of your custom VM naming script if you set the**Instance Naming Strategy**field to "Custom Callout". Location options inclu "Script from bundle" if you specify a zip file in the **Strategy Bundle** field. See Guidance for Callout Scripts for guidance on developi custom VM naming scripts.
IPAM Alloc Rule	Location of your custom IPAM address allocation script if you set the**Instance IPAM Strategy**field to "IPAM Callouts". Location opti include "Script from bundle" if you specify a zip file in the **Strategy Bundle** field. See Guidance for Callout Scripts for guidance on developing custom IPAM callout scripts.
IPAM Dealloc Rule	Similar to the**IPAM Alloc Rule**above except this is where you specify the location of yourcustom IPAM address deallocation script. Guidance for Callout Scripts for guidance on developing custom IPAM callout scripts.

Instance Naming Strategy Options:

Options	Description
Default	The default method provided by Cisco is the generation of a string consisting of A predefined token prefixed with **cq jw-**
CliQr Macro Replacement	Selecting this method causes the **Node Name Config** field to be displayed and to become a required field. See the fields table above for details.
Hostname Callout	Selecting this method causes the **Custom VM Name** field to be displayed and to become a required field. See the fields table above for details.

OOB macro variables forCliQr Macro Replacement:

Macro	Returned Value
%os1%	First character of OS type string in lower case
%os2%	Firsttwo characters of OS type string in lower case
%OS1%	Firstcharacter of OS type string in upper case
%OS2%	Firsttwo character of OS type string in upper case
%RND<number>%	Fixed length random string, e.g., %RND10%, 10-character long random string ⚠ This is a mandatory field, only values in the range: %RND6% to %RND11% or just %RND% which defaults to %RND6%"
%UID%	User ID
%VID%	Vendor ID
%VM_PREFIX%	VM name prefix must be added as a global parameter in the app profile. ⚠ Only macro whose value can be user-customized.

6. External Lifecycle Actions (conditional):Specify any external lifecycle actions to be performed on all VMs launched by Workload Manager in this region.

External Lifecycle Actions Settings

Whenever you use Workload Manager to deploy a VM-based application, Workload Manager can execute external actions on VMs in the deployment as the VMs progress through these five phases of their lifecycle:

- Pre VM start
- Pre VM initialization
- PostVM initialization
- Pre VM Stop
- Post VM Stop

These lifecycle actions will execute in a script execution engine running in a container in the CloudCenter Suite cluster if no Cloud Remote appliance is configured for the region. If you have configured Cloud Remote for this region, the external lifecycle actions will executein a script execution engine running in a container in the Cloud Remote appliance.

External Lifecycle Action scripts can be defined at the service level, the application level, and the cloud region level. Use the External Lifecycle Actions section of the Region tab or Details tab (see figure below) tospecify external lifecycle actions at the region level.

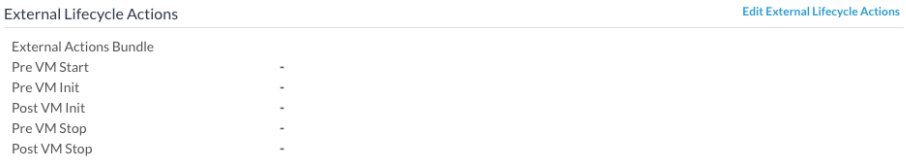

Click on the**Edit External Lifecycle Actions**link in the upper right. This displays theEdit External Lifecycle Actions dialog box as shown below.

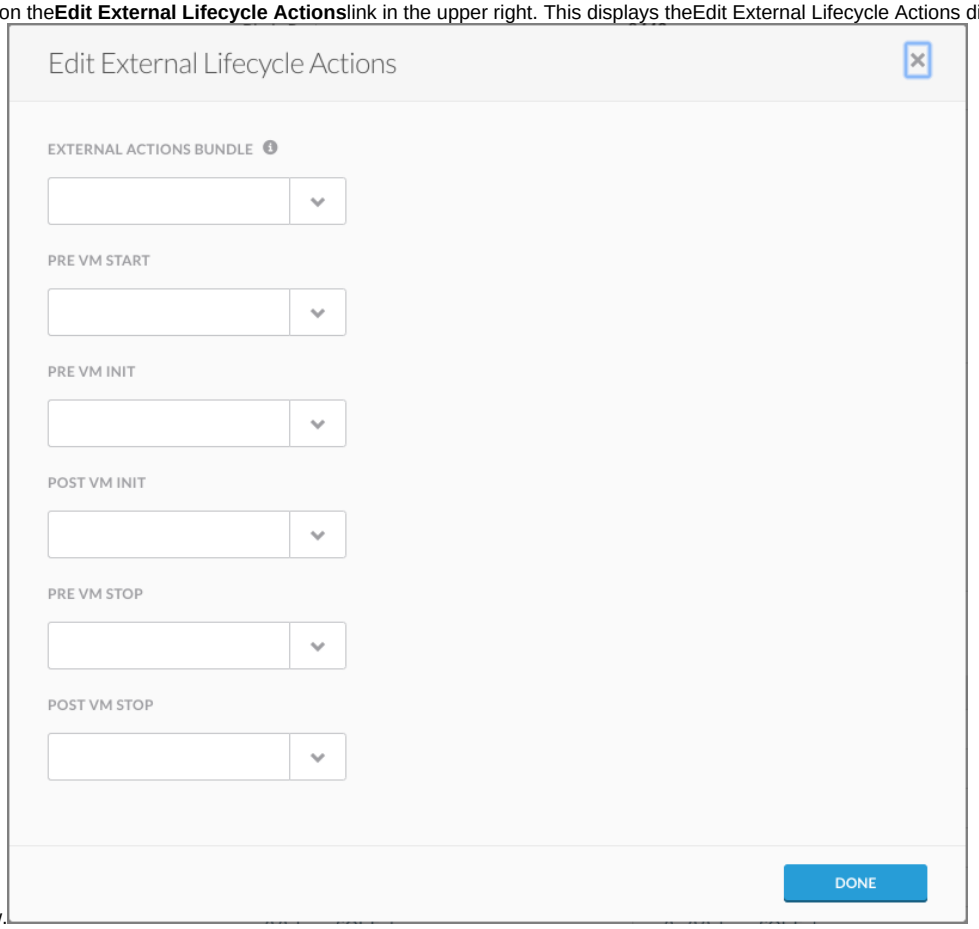

One way to specify an external lifecycle action is to use the first field,**External Actions Bundle**, to specify a zip file that contains all scripts for external lifecycle management. This zip file needs to contain a directory that contains all the scripts. Use the dropdown to select a repository location (URL, FTP, S3, etc.). This causes an additional field to be displayed to the right for entering the path to the zip file as shown in the figurebelow.

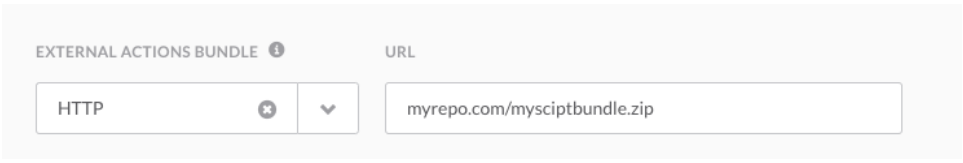

Once you specify actions bundle, for each of the individual lifecycle stage fields, select "Script from Bundle" from the dropdown. This displays an additional field to the right where you must specify the name of the script file in the bundle as shown in the figure below.

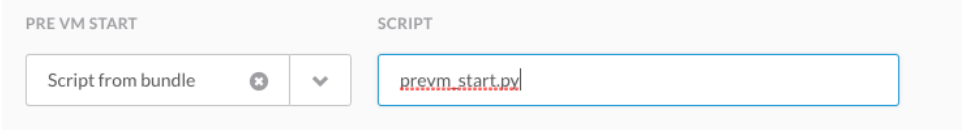

The scripts you specify can be bash scripts or Python scripts called from within a bash script.For more information on lifecycle scripts, seeDeployment Lifecycle Scripts.

7. Instance Types (informational): CloudCenter Suite automatically synchronizes instance types for public cloud regions daily. This data includes published pricing for each instance type. It is possible to edit AzureRM region instance types, but only the changes in the cost are used by CloudCenter Suite.

Instance Types Settings

An instance type is a performance category of a VM for VM-based clouds or a performance category of a container for container-based clouds. The parameters for an instance type for AzureRM is as follows:

- Display Name
- Price per hour
- Instance Type ID
- CPU
- 32- or 64-bit architecture
- RAM (MB)
- NICs
- Instance storage (GB)
- SSD support

You can sync, edit, and add instance types depending on the cloud provider as summarized in the table below. The following functions are synchronized for AzureRM clouds:

- Auto-populate and Auto-sync
- Edit

For public clouds, CloudCenter Suite auto-populates the instance types in the Instance Types section from the cloud provider. CloudCenter Suite auto-syncs this information from the cloud provider once per 24 hours. For these clouds, you cannot force sync on demand. For OpenStack clouds, you can sync all instance types (flavors) defined in OpenStack to CloudCenter Suite on demand. To manually sync OpenStack instance types, click the**Sync Instance Types**link in the upper left for the instances types section as shown in the figure below.

You can edit the fields of an instance type (except instance type ID) by clicking the**Edit**link in the Actions column in the list of instance types. Changing the values of parameters such as CPU and RAM are ignored during application deployment.

8. Storage Types (conditional): CloudCenter Suite automatically synchronizesstorage types for public cloud regions on a daily basis. This data includes the cloud provider published pricing for each storage type.It is possible to edit AzureRM region storage types, but only the changes in the cost are used by CloudCenter Suite.

Storage Types Settings

A storage type is a categorization of storage that can be provisioned for an application during deployment. It is a construct that applies to VM-based clouds only. Storage types are displayed in the Storage Types section of the Regions or Details tab.

The storage type construct contains the following fields for all VM-based clouds:

- Display Name
- Cost Per Month per GB
- Cloud Storage Type ID
- Minimum Volume Size (GB)
- Maximum Volume Size (GB)
- Maximum Throughput (MB/s)
- Maximum IOPS
- Provisioned IOPS
- IOPS Cost Per Month

Storage types are handled by CloudCenter Suite for various cloud providers in a fashion similar to instance types. The functions available for AzureRM cloud is as follows:

- Auto-populate and Auto-sync
- Edit

You may manually add additional storage types by clicking the**Add Storage Type**link in the upper right of the Storage Types settings section as shown in the figure below. CloudCenter Suite auto-populates the list of storage types in the Storage Types section based on information from the cloud provider. This data is stored in thepackage store and the package store auto-syncs this information from the cloud provider once per 24 hours.

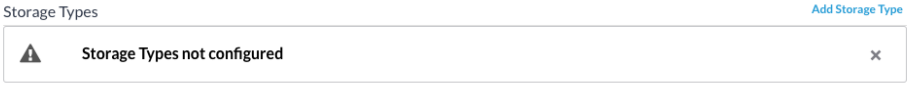

This brings up an**Add Storage Type**dialog box as shown in the figure below.

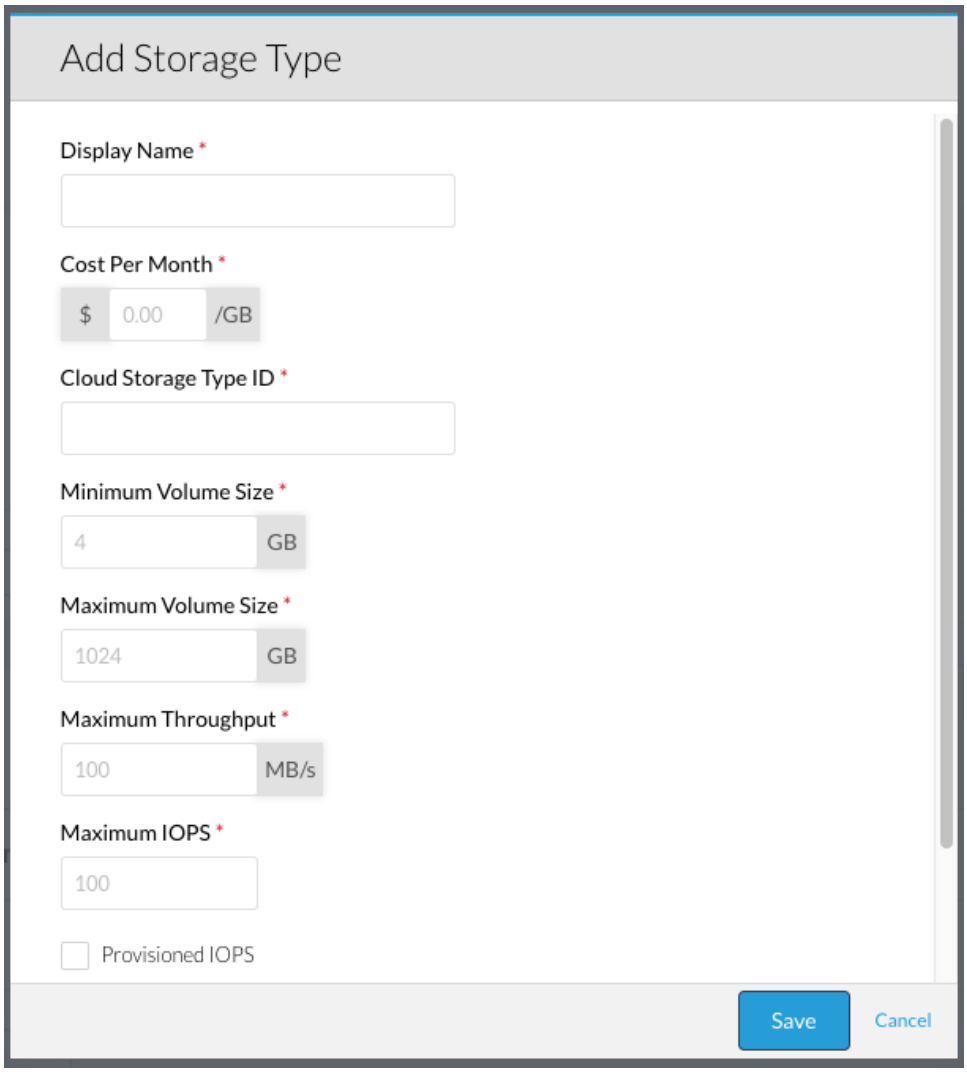

You must enter all required fields, however, the**Maximum Throughput**,**Maximum IOPS**,**Provisioned IOPS**, and**IOPS Cost Per Month**are not recognized by these cloud providers and will be ignored.

You must specify the**Cost**to compute the service costs in Cost Optimizer. The total cost is displayed in theCost Optimizer Dashboard.

Storage types for all cloud providers except AWS can be edited. From the list of storage types in the Storage Type section (see figure below), click the **Edit** link (in the **Actions** column) for that storage type.

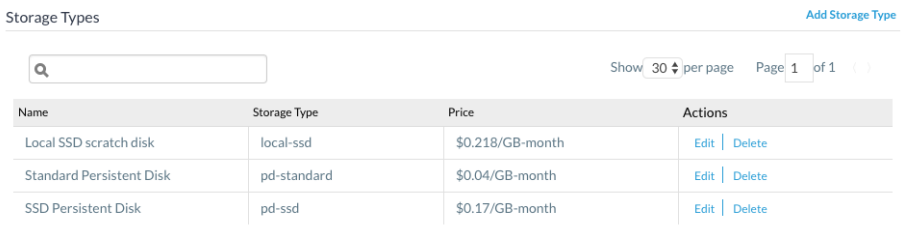

This brings up an **Edit Storage Type** dialog box (similar to the **Add Storage Type** dialog box) where you can edit all nine fields except the**Cloud Storage Type ID**field. For AzureRM and Google clouds, only changes to the pricing fields change the behavior of CloudCenter Suite: the cloud cost calculations are modified based on the new unit cost data. All other field changes are ignored.

For all cloud providers,**Minimum Volume Size**and**Maximum Volume Size**restrict the volume size you can specify at deploy time. For public cloud providers that support provisioned IOPS, the user can specify the guaranteed IOPS at deploy time.

9. Image Mappings: Image mappings allow services based on Workload Manager logical images to be deployed using the appropriate physical image stored on the target cloud region. Workload Manager automatically maps theOOB logical imagesto public cloud region physical images when you add the region to your cloud. Cisco periodically updates these mappings when new versions of OS physical images are uploaded by the cloud provider. To apply these updates to your region after it is added to your cloud, click the**Sync Image Mappings**link in the upper right of this section. If you create any custom logical images, you mustmanually import the corresponding physical images into your region and then map the corresponding logical images to these physical images.SeeImagesfor more context.

Add an AzureRM Cloud Account

Be aware that the screenshots may change based on the Azure portal changes. They are provided in this section as a point of reference.

Prerequisites

Before adding an AzureRM cloud, verify the following requirements:

- You have a validWindows Azure Resource Manager account.
- Register the required Azure providers from the Azure portal:

Previously, you could only perform this procedure using Azure commands.

Now, you can use the UI (**All Services** > **Subscriptions**) to register the following Azure providers:

- Microsoft.Compute (displayed in the following image)
- Microsoft.Storage (displayed in the following image)
- Microsoft.Network (displayed in the following image)
- Microsoft.Resources
- Microsoft.Authorization

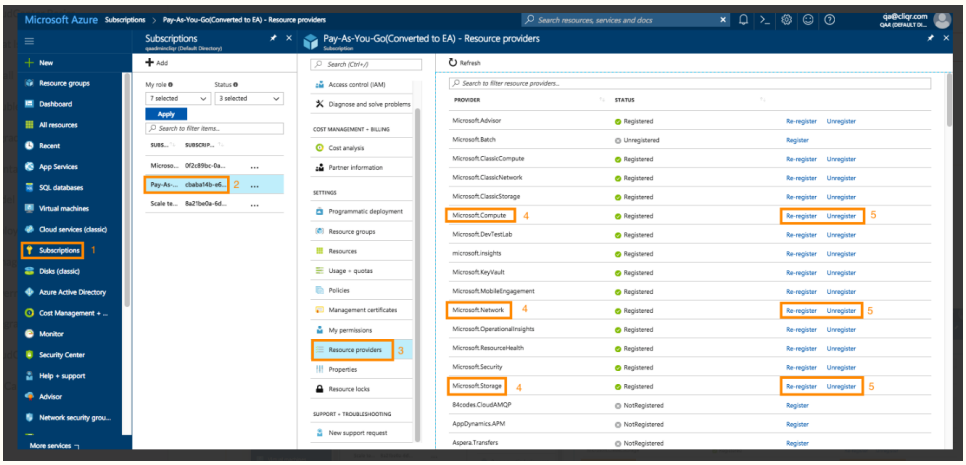

- In the**Azure Resource ManagerPortal**, navigate to**Azure Active Directory**page:

1. Select**App Registration**and click**Add**.
2. Provide the**Name**,**Sign-On URL**, and**Create**the application. This value must be a standard URL and is required by theAzureRM cloud configuration it is not used by the CloudCenter platform.

In the following screenshot, the Sign-On URL displays *http://www.cliqr.com*. This is just an example. Be sure to provide the base URL for your application using the required protocol (HTTP or HTTPS) for example:

```
http://<YourLocalHost or YourAppURL>
```

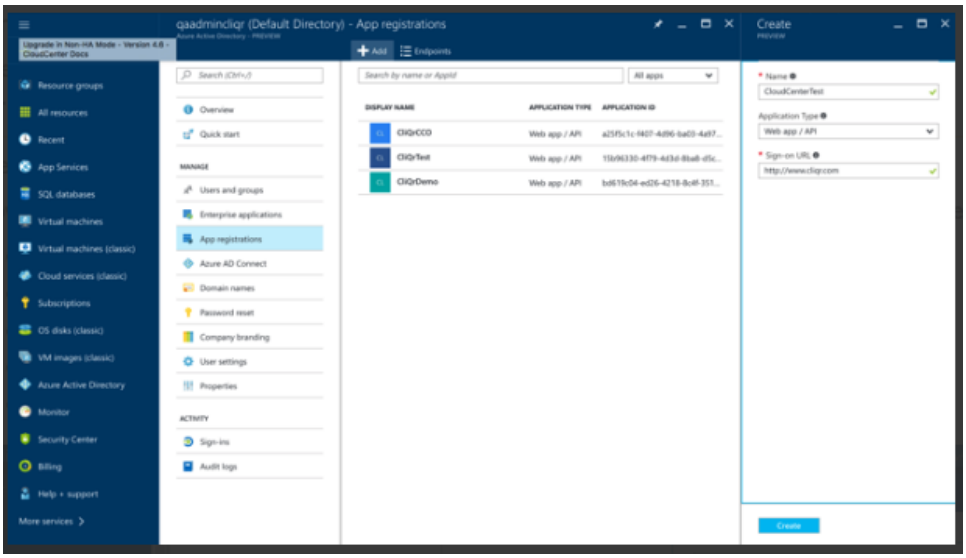

3. Select the newly created application.

> **Note down the Application ID; it is required to create a Cloud Account in CloudCenter this is the Client ID.**
>
> **If you prefer to use *Certificate-Based Authentication*, see the related bullet further in this section.**

4. Click**All Settings**.
5. Select**Required Permission**under API Access and click**Add**.SeeCloud Overview>*Minimum Permissions for Public Clouds*for additional details.

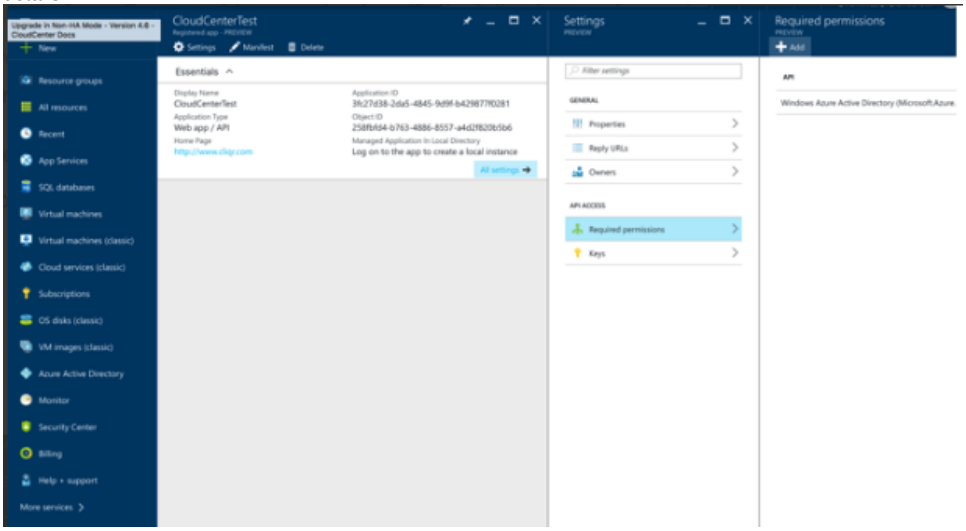

6. Select**Windows Azure Service Management API**.

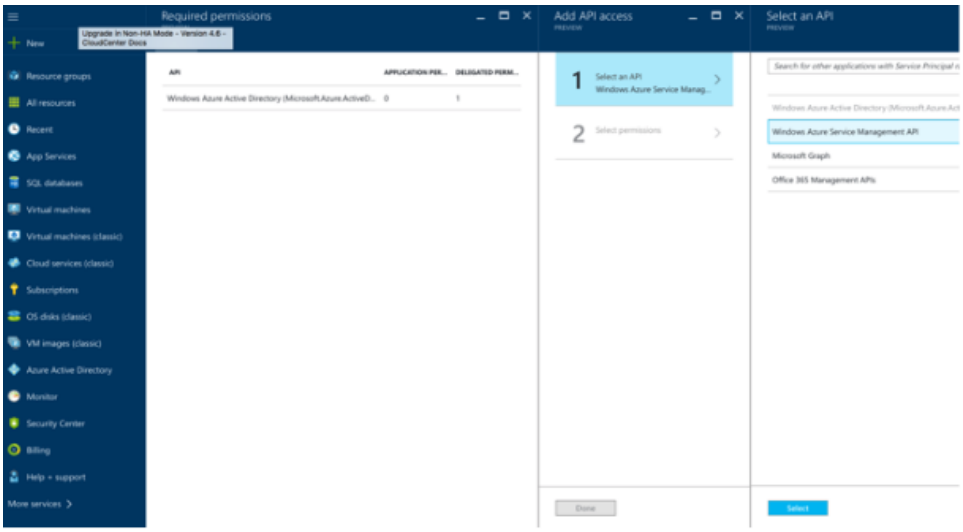

7. Select permissions as**Delegated Permission**and click**Done**.

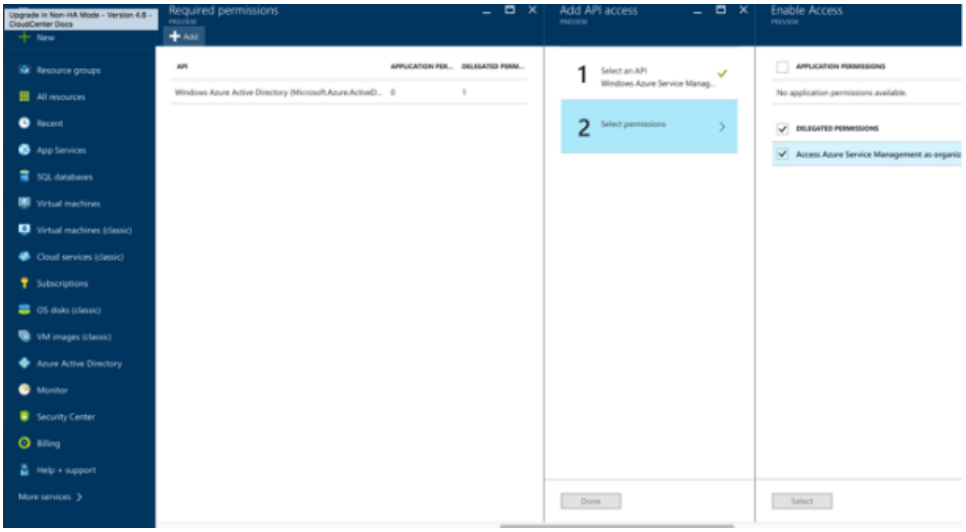

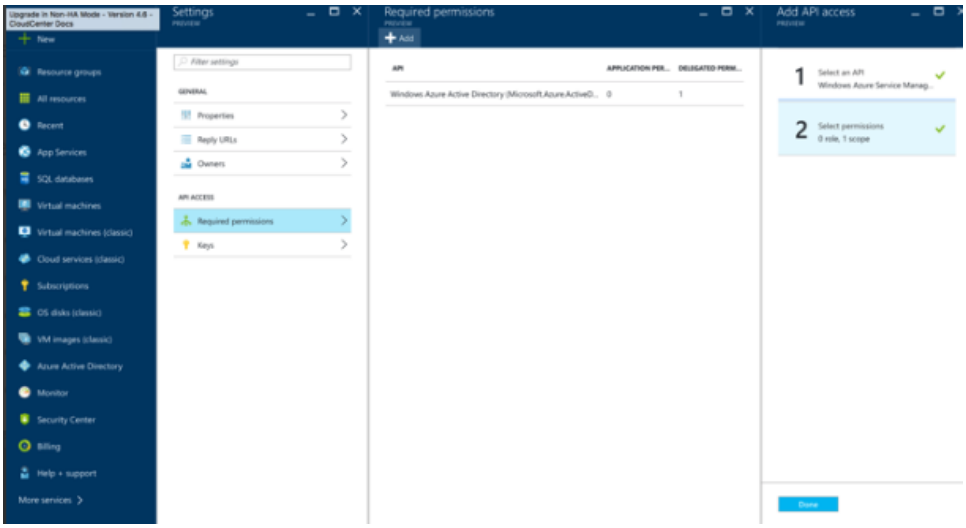

8. Select**Keys**under**API Access**.
9. Specify the**Description**, Expires,and click**Save**.

> Note down the key after you click save this key cannot be retrieved later from the portal, and it is used by the Workload Manager as the Client Key when creating the cloud account.

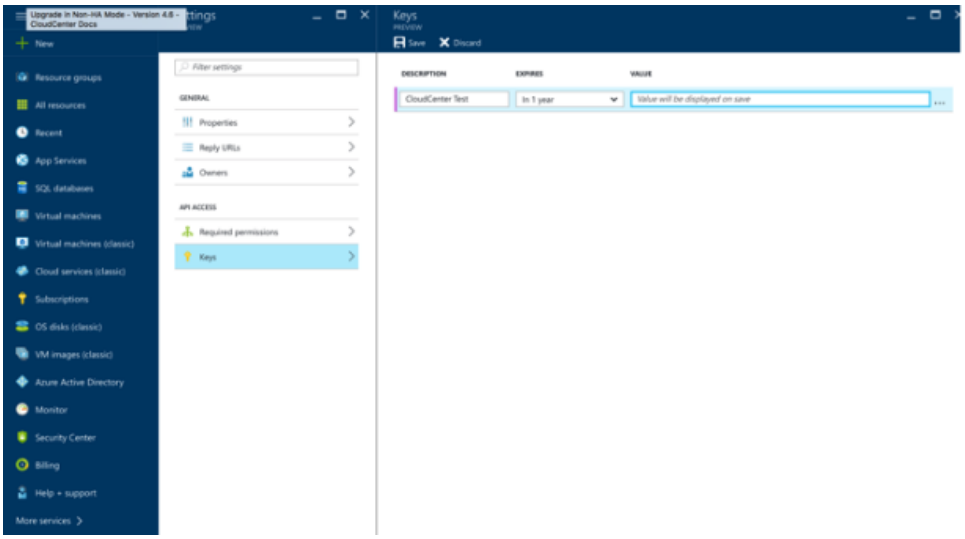

10. Select**App Registration**and click**Endpoints**.

> Note down the Tenant-ID from the OAuth 2.0 Authorization Endpoint this ID is used by the Workload Manager when creating a cloud account.

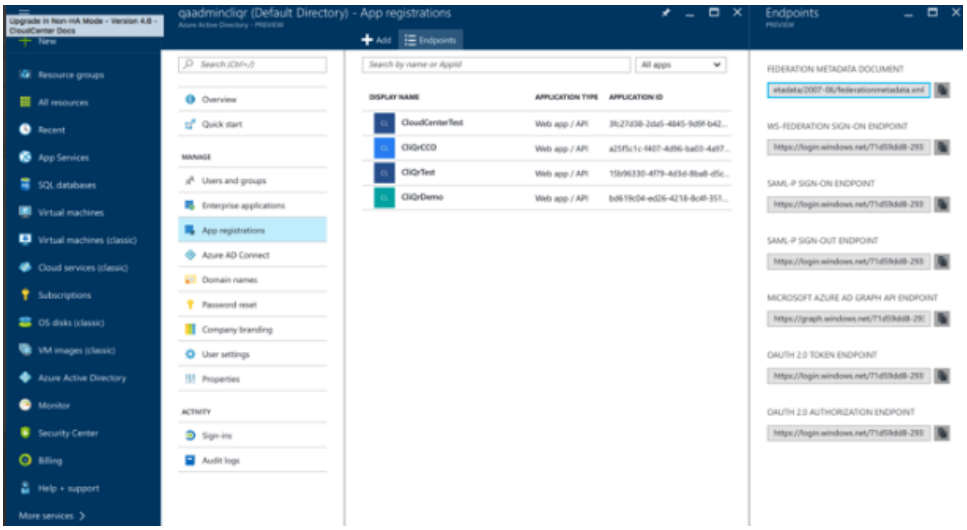

- *Certificate-Based Authentication*You can select either key-based authentication or the more secure certificate-based authentication.

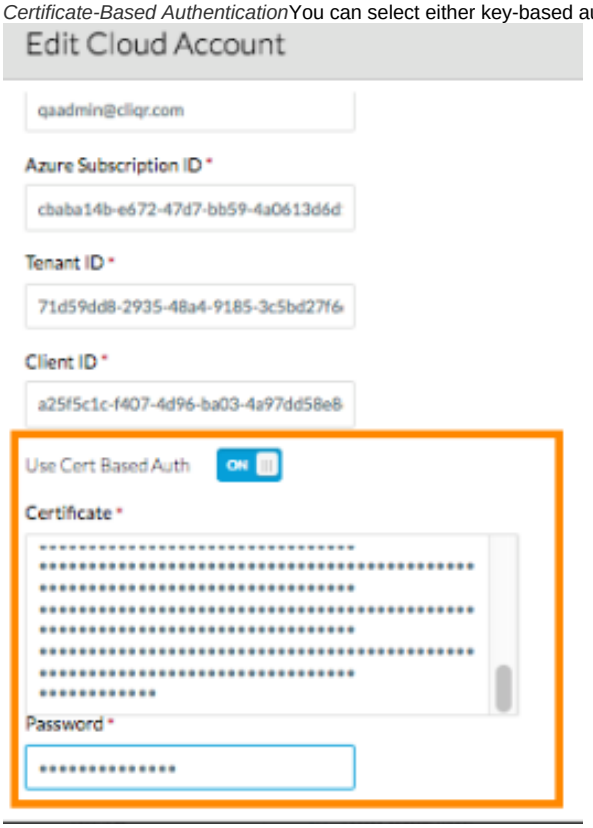

  - The certificate used can either be one of the following optionsYou can create either type using the*openssl*command from the command prompt of any Linux system:
    - A self-signed certificate: See the following example.

      > Remember this password as you will need to enter it in the CloudCenter Suite UI's Certificate and Password fields when you create or edit the Cloud Account.

      - Generate a key and certificate.

        ```
        openssl req -newkey rsa:2048 -nodes -keyout key.pem -x509 -days 365 -out certificate.
        pem
        ```

      - Convert the certificate.pem to PKCS 12 format.

        ```
        openssl pkcs12 -inkey key.pem -in certificate.pem -export -out certificate.p12
        ```

      - Provide a password to this command when prompted.
    - A Certificate Authority (CA) signed certificateGenerate a key and CSR,send/receive the certificate.csrfile(s) to the signature authority, convert the signed-certificate.pem to PKCS 12format, and provide a password to this command when prompted.

      > Remember this password as you will need to enter it in the Workload Manager UI's Certificate and Password fields when you create or edit the Cloud Account.

  - Convert the PKCS formatted certificate (certificate.p12 or signed-certificate.p12) to base64 format using the tool athttps://www.base64encode.org/.
  - Enter the base64 formatted certificate, and the export password used to create the PKCS formatted certificate, in the corresponding fields in theWorkload ManagerAdd or Edit Cloud Account dialog box.
  - Login to**Azure Resource Manager Portal**to upload the certificate PEM file (Azure Active Directory > AppRegistrations > Settings > keys > Upload public key) and save.

    > The corresponding public key for the certificate must be uploaded to the Azure RM portal for the Application Registration that the user must add to the CloudCenter Suite cloud account.

- In the**Azure Resource Manager Portal**, configure the user role settings for your web application:

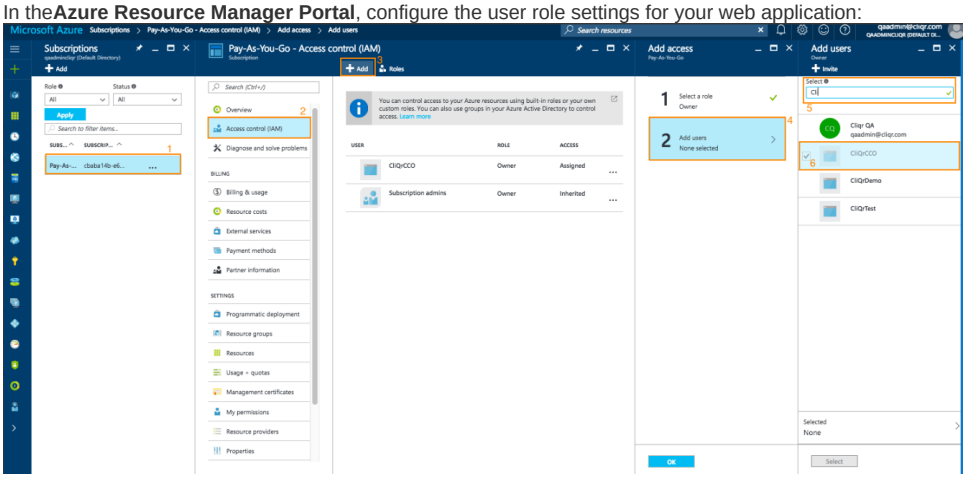

1. Select**Subscription**>**Valid subscription**(this isthe subscription you want to manage).
2. Click**Access control (IAM)**.
3. Click the**+Add**icon at the top right corner of the managed subscription pane.
4. Click**Add users**and select the**OWNER**role. You can also select other roles for more granular management.

> This role should be able to access and manage AzureRM resources like storage, compute, network, keyvault, and so forth to configure AzureRMfor the CloudCenter Suite.

5. In the User search box, enter the web application name you defined earlier. In this example, it is**CliQrCCO**.
6. Click**OK**to save your settings.

Configuration Process

To add an AzureRM cloud account, follow this procedure.

1. Locate the newly-added cloud and clickthe**Add Cloud Account**link. The Add Cloud Account dialog box displays, as shown in the figure below:

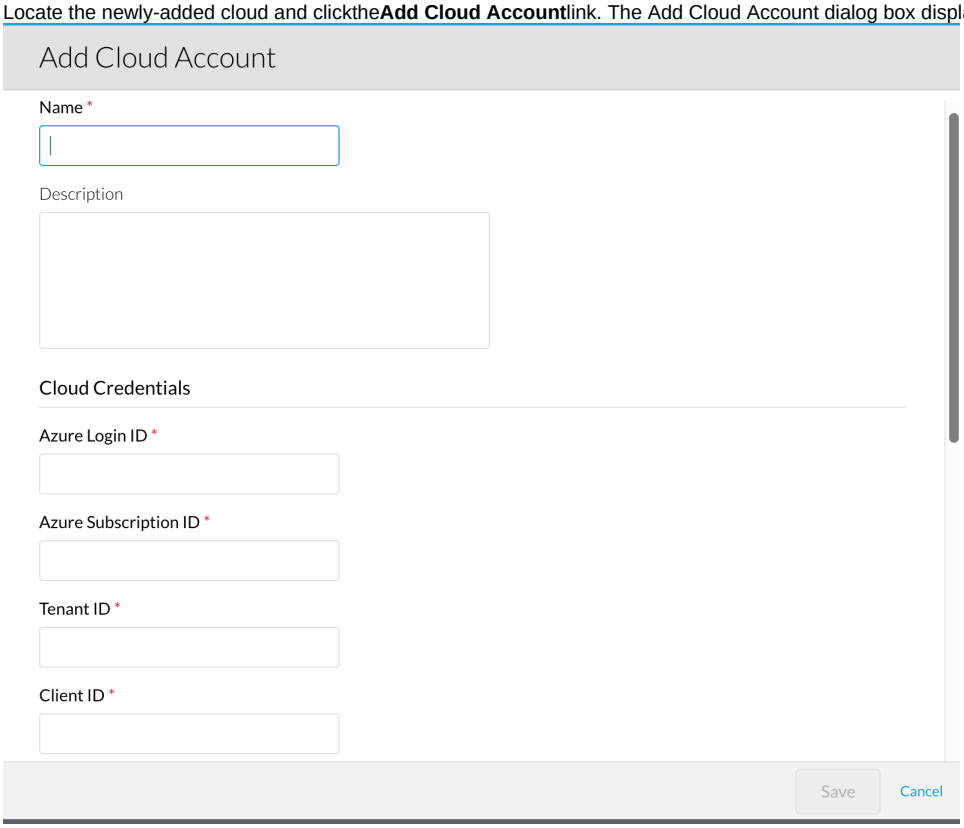

2. Assign a new cloud account name.

> **Tip**
>
> The name should not contain any space, dash, or special characters.

3. Add the following cloud credentials associated with your Azure account.
    a. **Azure Login ID**: The email address used to login to your Azure Resource Manager cloud account
    b. **Azure Subscription ID**: To retrieve the**Subscription ID**, toggle to the**Azure Portal**Interface as described in the*Prerequisites*section above and access Settings:

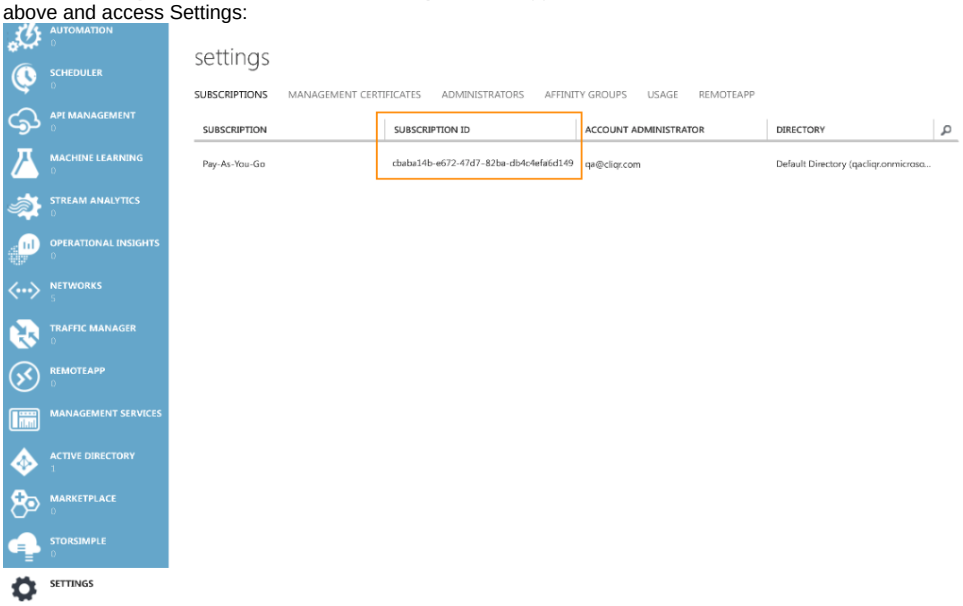

    c. **Tenant ID**: The UUID identified in the*VIEW ENDPOINTS*bullet in the*Prerequisites*section above.
    d. **Client ID**:The UUIDidentified in the blue icon bullet in the*Prerequisites*section above.
    e. **Use Cert Based Auth**:If you enable**Use Cert Based Auth**, the**Client Key**field is*hidden,*and the following fields are displayed:
        i. **Certificate** The certificate in PKCS 12 format as Base64 text asidentified in the*Certificate-Based Authentication*bullet in the*Prerequisites*section above.
        ii. **Password** Enter the password used to create the certificateasidentified in the*Certificate-Based Authentication*bullet in the*Prerequisites*section above.
    f. **Client Key**:If you do not enable**Use Cert Based Auth**,use the client key identified in the*keys*bullet in the*Prerequisites*section above.
4. Scroll the dialog box down to reveal the billing fields and enter the**Region Info**,**OfferId**,**EA Enrollment Number**, and**EA API Access Key,**as shown in the figure below.For information on setting up billing information, seehttps://docs.microsoft.com/en-us/rest/api/consumption/andhttps://docs.microsoft.com/en-us/azure/billing/billing-enterprise-api.

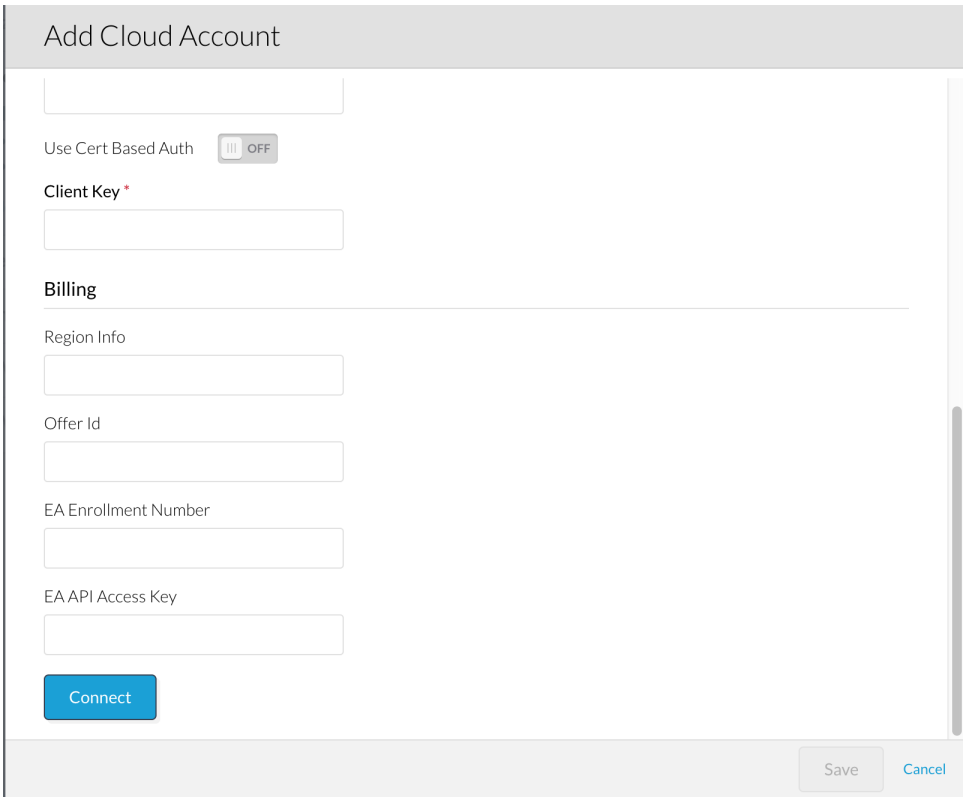

> The **Region Info** is the two-letter ISO code where the offer was purchased. For example, US.
>
> The **Offer Id** is tied to the account. To find the **Offer Id** for your account, navigate to **Azure Portal** > **Subscriptions page** and choose a subscription. The **Offer Id** is displayed in the **Overview** section.
>
> The **EA Enrollment Number** is displayed in the top left corner when you log in to https://ea.azure.com/.
>
> The **EA API Access Key** must be generated as follows:Log in to https://ea.azure.com/ as **EA Admin** andnavigate to **Reports** > **Download Usage** > **API Access Key** > **Generate**.

5. Click the **Connect** button. CloudCenter Suite will now attempt to validate your account credentials.
6. After the credentials are verified, the **Connect** button changes to an **Edit** button, two new fields appear **Enable Account For** and **Enable Reporting By Org Structure**.

    a. Set the **Enable Account For** dropdown per the table below.

Value	Usage
Provisioning	Workload Manager can deploy jobs using this account.
Reporting	Cost Optimizer and Workload Manager will track cloud costs for this account. Typical usage: master cloud accounts that are used for billing aggregation. ⚠ It is recommended that you do not add a *Reporting* account to the same tenant through different cloud groups. ⚠ Enabling a public cloud account for *Reporting* may incur expenses to retrieve cost data. These expenses are proportional to the number of configured cloud accounts and regions.
Provisioning, Reporting	Default. Account is used for both provisioning and reporting.

    b. **For AWS and Google clouds only**: Set the **Enable Reporting By Org Structure** toggle to On to cause Cost Optimizer to import the cost hierarchy createdin the cloud provider portal. This saves the time of manually creating a comparable cost hierarchy within Cost Optimizer. See Cost Groups Configuration for more information on cost hierarchies in Cost Optimizer.
    c. Click the **Save** button when done.

Cloud Accounts Tab

After you add cloud accounts to a cloud, they will appear in the Accounts tab for the cloud as shown in the figure below.

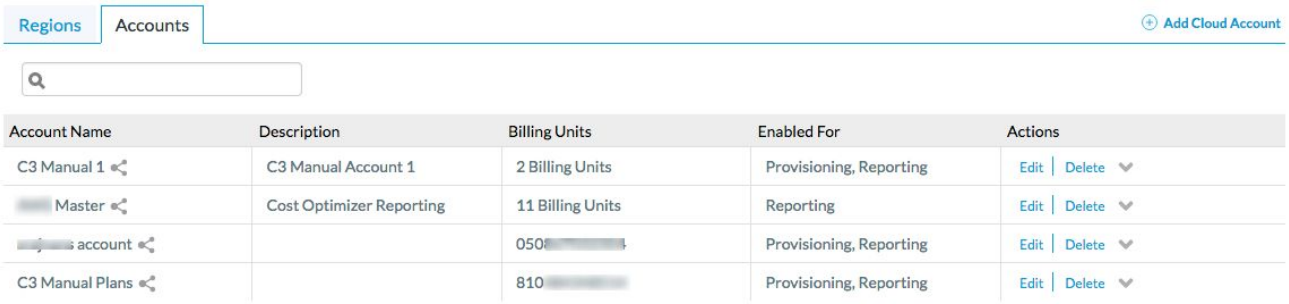

The Accounts tab contains columns for data entered when creating an account: Account Name, Description, Enabled For; and two additional columns:**Billing Units**and**Actions**.**Billing Units**is a dual function:

- If the cloud account contains only one billing unit, the ID for that billing unit is displayed.
- If the cloud account contains multiple billing units, such as an AWS master account, the number of billing units in that account is displayed followed by the text*Billing Units*.

A billing unit is the most granular level of cloud cost recording in CloudCenter Suite. The definition of a billing unit varies by a cloud provider as shown in the table below.

Cloud Provider	Billing Unit
AWS	Account ID
AzureRM	Subscription ID
Google	Project ID
IBM Cloud	Account ID
vCenter	Cloud Group Prefix - Datacenter Name
vCD	Organization Name
OpenStack	Project ID
Kubernetes	Namespace UID

The last column,**Actions**, contains links to let you edit or deleted the cloud account, ormanage instance typesfor the cloud account.

Configure an IBM Cloud

Configure an IBM Cloud

Configuring an IBM Cloud is a four-step process:

- Add an IBM Cloud
- Add an IBM CloudRegion
- Configure an IBM Cloud Region
- Add an IBM Cloud Cloud Account

Add an IBM Cloud

To add an IBM Cloud cloud follow these steps.

1. Navigate to**Admin > Clouds**. This brings you to the Clouds page. If you, or another tenant admin in your tenant, have already added clouds to your tenant, they will be listed here. Click the**Add Cloud**linkin the upper right.
2. Click**Add Cloud**, the Add Cloud dialog box is displayed.
3. Enter the**Cloud Name**,select the**cloud provider**, and click**Next**.

   When assigning the Cloud Name, use a unique name. Ensure that you do not assign a previously-used name (including deleted cloud names) as the Cloud Name as they may cause deployments to fail intermittently.

   The second page of the**Add Clouds**dialog box,**Connectivity Settings**, appears.Set the toggle switches to configure the Cloud Connectivity settings.

   - When adding a public VM cloud in the CloudCenter Suite UI, the Cloud Connectivity Settings page, the second page of the Add Cloud dialog box, appears with a single toggle displayed:**Worker VMs Directly Connect with CloudCenter Suite**.
   - Setting this toggle to No implies you will install Cloud Remote for each region of this cloud. This also causes a second toggle to appear:**CloudCenter Suite Directly Accessible from Cloud Remote**.
   - Follow the table below for guidance on setting these toggles.

Toggle settings	Use case	Diagram
Worker VMs Directly Connect with CloudCenter Suite = Yes	Unimpeded connectivity exists between the CloudCenter Suite cluster and the cloud region API endpoint AND Unimpeded connectivity exists between the CloudCenter Suite cluster and worker VMs Cloud Remote is not required	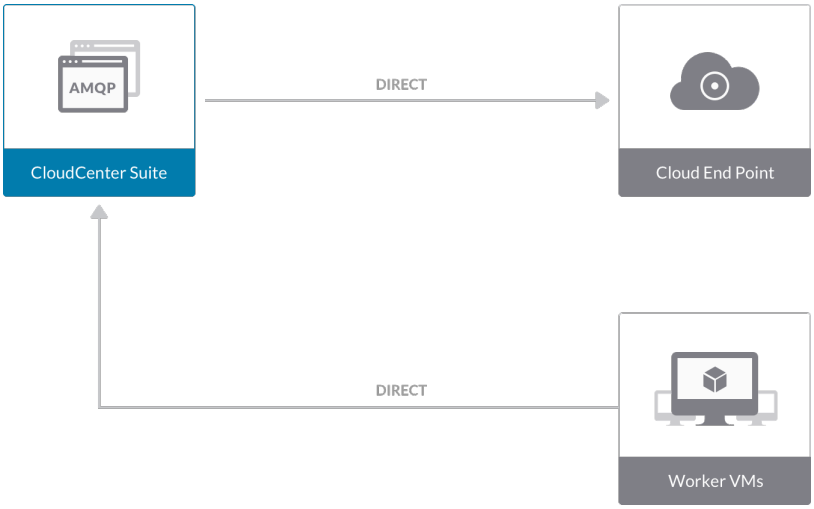

Worker VMs Directly Connect with CloudCenter Suite = No AND CloudCenter Suite Directly Accessible from Cloud Remote = Yes	Worker VMs need to communicate to the CloudCenter Suite cluster through Cloud Remote AND Cloud Remote can initiate the connection to the CloudCenter Suite cluster	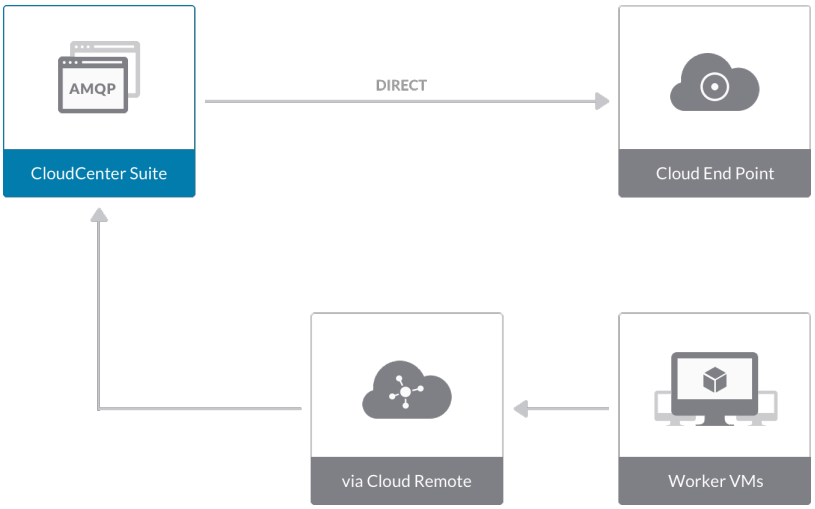
Worker VMs Directly Connect withCloudCenter Suite = No AND CloudCenter SuiteDirectly Accessible from Cloud Remote = No	Worker VMs need to communicate to the CloudCenter Suite cluster through Cloud Remote AND the CloudCenter Suite cluster cannot receive a connection initiated by Cloud Remote	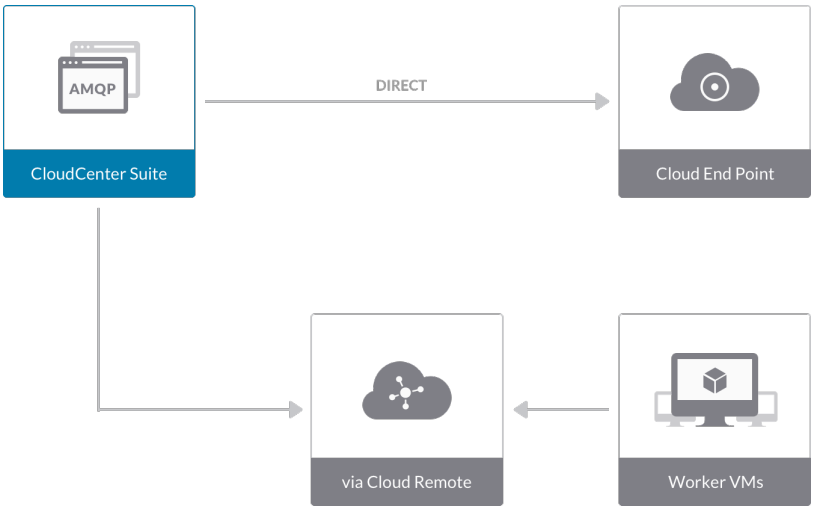

Note

The connectivity toggle settings set at the cloud level are inherited by each region you add to this cloud. However, it is possible to override these toggle settings on a per-region basis from the Regions tab for each region.

4. Click**Done**to save the configuration and close the dialog box. This brings you back to the **Clouds** page and the cloud you just created will be added to the bottom of the list on the left side of the page.

Add an IBM CloudRegion

After creating an IBM Cloud cloud, the next step is to create the first region for the cloud. Follow these steps.

1. Navigate to the **Clouds** page and select the cloud you created on the left side of the screen. Click the**Add Region**button onthe right side of the screen. The **Add Region** dialog box is displayed.
2. Select a region from the list and click**Save**. You are brought back to the **Clouds** page with theregion you added shown on the right side of the page.

Configure an IBM Cloud Region

To configure a region you added to your IBM Cloud cloud, follow this procedure:

1. Navigate to Clouds page:**Admin > Clouds**. Find your IBM Cloud cloud from the cloud list on the left half of the screen and click its**Configure Cloud**link. This displays the **Regions** tab for this cloud as shown in the figure below with the**Cloud Settings** section displayed first.If you have added multiple regions to your IBM Cloud cloud, the**Regions**tab will show multiple individual region tabs on the left side of the screen.

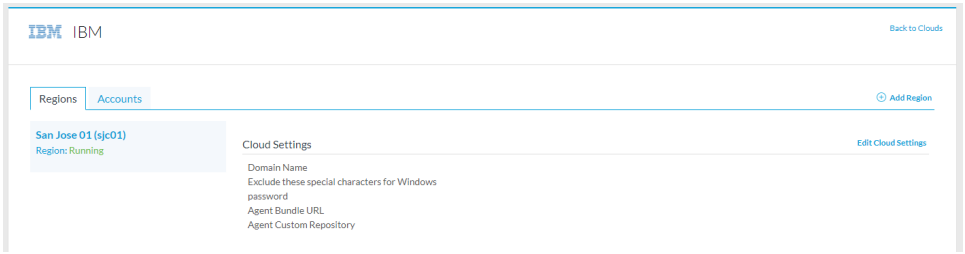

2. Click the tab of the region you want to configure.
3. Click the**Edit Cloud Settings**link in the upper right of the **Cloud Settings** section. This opens the **Configure Cloud Settings** dialog box.The **Cloud Settings** section contains fields that are unique to IBM Cloud and settings that are common to all cloud providers. Adjust these field values per the instructions in the following tables.

IBM Cloud Specific Cloud Settings

Field	Usage
Domain Name	The URL route allocated to your organization in IBM Cloud.

Agnostic Cloud Settings

Field	Usage
Exclude these special characters for Windows password	When the Workload Manager agent is installed on a Windows worker VM, a special user account, called cliqruser, is created to support RDP sessions that may be initiated by the user through the Workload Manager UI. A Workload Manager process running on the CloudCenter Suite cluster creates a random password and passes it to the agent for creating the cliqruser account. Because some Windows deployments may restrict using certain characters for Windows passwords, this field is provided to tell the Workload Manager to exclude these special characters in the generation of the password for the cliqruser account.
Agent Bundle URL	If you plan to use a local repository to host the bundle store, you need to enter the URL of the local bundle store here. Otherwise, leave blank.
Agent Custom Repository	If you plan to use a local repository to host the package store, you need to enter the URL of the local package store here. Otherwise, leave blank.

When you are done editing the settings in the dialog box, click**Save**.

4. Determine if you need Cloud Remote for this region.Scroll down to the **Region Connectivity** section for the region and click on the**Configure Region**link in the upper right to open the Configure Region dialog box.The toggle settings should be the same as when you set them on the connectivity page of the **Add Cloud** dialog box.If all of the connectivity togglesinthe **Region Connectivity** dialog boxare set to **Yes**, then Cloud Remote is NOT needed for this cloud region. In this case,you would normally leave the region connectivity settings at their current values and continue to the next settings section.

   The exceptionto this guidance is when a NAT firewall or proxy server exists between theCloudCenter Suite management cluster andworker VMs, or between the CloudCenter Suite management cluster and users that would use Workload Manager to initiate a Guacamole remote connection to a worker VM. In either of these cases, override the address fields in the **Region Connectivity** dialog box as explained below.

Networking Constraint	Field	Value
Worker VMs must use a proxy server or NAT firewall to access the"local" AMQP server running in the CloudCenter Suite cluster.	Worker AMQP IP Address	IP address and port number that the firewall or proxy server presents to the worker VMs on behalf of the "local" AMQP server running in the CloudCenter Suite cluster.
Users must use a proxy server or NAT firewall to access theGuacamole server running in the CloudCenter Suite cluster.	Guacamole Public IP Address and Port	IP address and port number that the firewall or proxy server presents to users on behalf of theGuacamole server running in the CloudCenter Suite cluster.
Worker VMs must use a proxy server or NAT firewall to access theGuacamole server running in the CloudCenter Suite cluster.	Guacamole IP Address and Port for Application VMs	IP address and port number that the firewall or proxy server presents to the worker VMs on behalf of the Guacamole server running in the CloudCenter Suite cluster.

   Click**OK**to save the changes and dismiss the dialog box. You can now proceed to the next region settings section:VM Naming and IPAM Strategy.

5. If any of theconnectivity toggles inthe Region Connectivity dialog box are set to No, then you must install and configure Cloud Remote for this region.

Configure Cloud Remote in an IBM Cloud Region

Configure Cloud Remote in an IBM Cloud region as follows.

⚠ Since CloudCenter Suite does not include a prebuilt appliance for Cloud Remote for IBM Cloud, the following procedure includes steps to build the Cloud Remote appliance from the Cisco-supplied Cloud Remote installer file.

Launch Cloud Remote Built from the Installer File

a. Launch a Centos 7 instance, ensure the prerequisites are installed, and run the Cloud Remote installer file:

Unable to render {include} The included page could not be found.

b. Optional but recommended for production environments: Repeat the step above twice to create two additional instances of the appliance to be used to form a cluster for HA.Cloud Remote includes support for the clustering of multiple nodes. You will "add" these two additional instances to the first instance after the first instance is configured. SeeCloud Remote (Conditional)> *Scaling*for details.

Unable to render {include} The included page could not be found.

Specify AMQP and Guacamole Addresses for Supporting Cloud Remote

From the CloudCenter Suite UI, for the cloud region requiring Cloud Remote, navigate to the corresponding Regions or Details tab. Click the**Configure Region**link in the upper left of the Region Connectivity section to bring up the Configure Region dialog box. The toggle settings should be the same as when you set them on the connectivity page of the Add Cloud dialog box. You must update some of the address fields in the dialog box according to the scenarios summarized in the table below.

Toggle Settings	Field	Value
Worker VMs Directly Connect with CloudCenter = No AND CloudCenter Directly Accessible from Cloud Remote = Yes	Local AMQP IP Address	Pre-populated with the address and port number of the "local" AMQP server running in the CloudCenter Suite cluster. This address must be **accessible to Cloud Remote** . If Cloud Remote is accessing the CloudCenter Suite cluster through a user-supplied proxy server or NAT firewall, overwrite this field with the corresponding local AMQP IP address and port number provided by the user-supplied proxy server or NAT firewall and accessible to Cloud Remote .
Worker VMs Directly Connect with CloudCenter = No AND CloudCenter Directly Accessible from Cloud Remote = No	Remote AMQP IP Address	Enter <Cloud_Remote_IP>:<amqp_port>, where <Cloud_Remote_IP> = the IP address Cloud Remote which is **accessible to the CloudCenter Suite cluster**, and <amqp_port> = 5671 OR the custom AMQP port number you would later set with the **Change Ports shell script** on the Cloud Remote appliance (see Cloud Remote (Conditional) > *Custom Port Numbers (Conditional)*). If there is no user-supplied NAT firewall or proxy server between the CloudCenter Suite cluster and Cloud Remote , the IP address would be the public IP address of Cloud Remote . If there is a NAT firewall or proxy server between the CloudCenter Suite cluster and Cloud Remote , instead, enter the corresponding public IP address and port number that the firewall or proxy server presents to the internet on behalf of the "remote" AMQP server running in Cloud Remote .
Worker VMs Directly Connect with CloudCenter = No	Worker AMQP IP Address	Enter <Cloud_Remote_IP>:<amqp_port>, where <Cloud_Remote_IP> = the Cloud Remote IP address **accessible to the worker VMs**, and <amqp_port> = 5671 OR the custom AMQP port number you would later set with the **Change Ports shell script** on the Cloud Remote appliance (see Cloud Remote (Conditional) > *Custom Port Numbers (Conditional)*).
Worker VMs Directly Connect with CloudCenter = No	Guacamole Public IP and Port	Enter <Cloud_Remote_IP>:<guac_port>, where <Cloud_Remote_IP> = the Cloud Remote IP address **accessible to CloudCenter Suite users**, and <guac_port> = 8443 OR the custom Guacamole port number you would later set with the **Change Ports shell script** on the Cloud Remote appliance (see Cloud Remote (Conditional) > *Custom Port Numbers (Conditional)*).
Worker VMs Directly Connect with CloudCenter = No	Guacamole IP Address and Port for Application VMs	Enter <Cloud_Remote_IP>:<guac_port>, where <Cloud_Remote_IP> = the Cloud Remote IP address **accessible to worker VMs**, and <guac_port> = 7789

When done, click **OK** to save the setting and dismiss the dialog box.

Download Region Connectivity Settings and Upload to Cloud Remote

Cloud Remote uses the region connectivity settings set in the Workload Manager or Cost Optimizer UI. After saving the Region Configuration settings in the Workload Manager or Cost Optimizer UI, you must download them and to your local computer and then upload them to Cloud Remote as follows.

Click the**Download Configuration**link in the upper right of the Region Connectivity section, as shown in the figure below.

Clicking DownloadConfiguration causestwo things to happen:

- An encrypted zip file named**artifacts.zip**will be downloaded by your browser. Make note of the location of this zip file as you will need to upload it to Cloud Remote through the Cloud Remote web UI (see below).
- The Region Connectivity section header updates to display a **Copy Encryption Key**link, as shown in the figure below.

Click the**Copy Encryption Key**link to save the key to your clipboard. A success message will be displayed temporarily above the Region Connectivity section header. Make sure not to overwrite the clipboard with other data. You will need the key when you upload the configuration zip file to Cloud Remote .

> If you change the connectivity settings in the CloudCenter Suite UI and need to again download the zip file, a new encryption key is automatically created and can be copied to the clipboard by clicking the **Copy Encryption Key** link again.

After you have downloaded the zip file and copied the encryption key to your clipboard, login toCloud Remote web UI.

a. Open another browser tab and login to https://<Cloud Remote _ip> with thedefault credentials: admin/cisco.
b. You will immediately be required to change your password. Do so now.
c. You are now brought to the Cloud Remote home page as shown in the figure below.

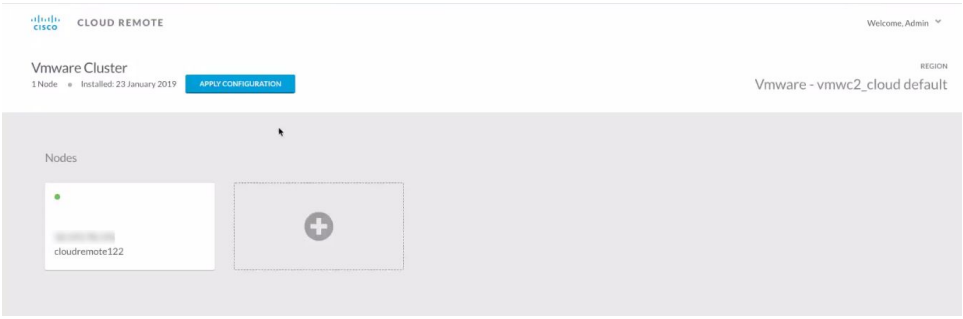

d. Click the **Apply Configuration** button in the page header. This prompts you to select a configuration file and enter the encryption key as shown in the figure below.

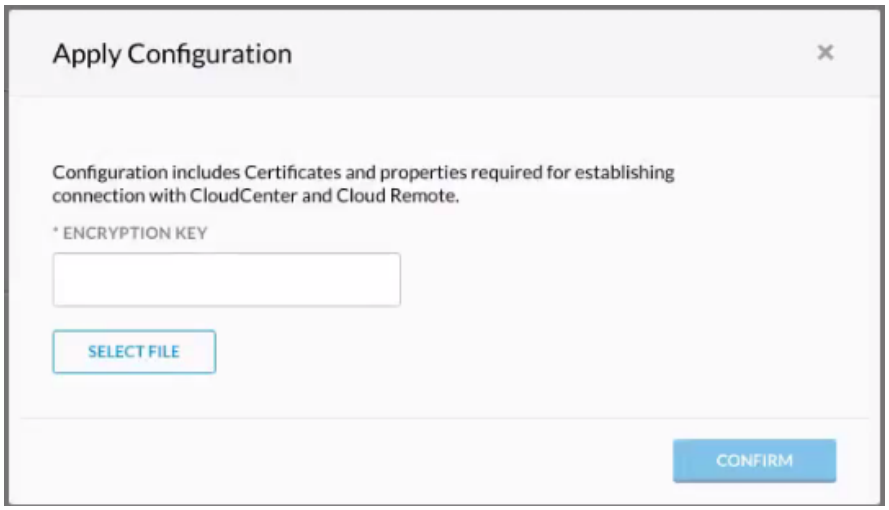

e. Paste the encryption key that was copied to the clipboard into the**Encryption Key**field in the dialog box.
f. Click**Select File**and browse tothe artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
g. Click**Confirm**.
h. Once the zip file is successfully transmitted and accepted, the Cloud Remote appliance attempts to establish communication with the CloudCenter Suite cluster and the Cloud Remote web UI home page is updated to show the name of the region it is connecting to in the upper right (see figure below).

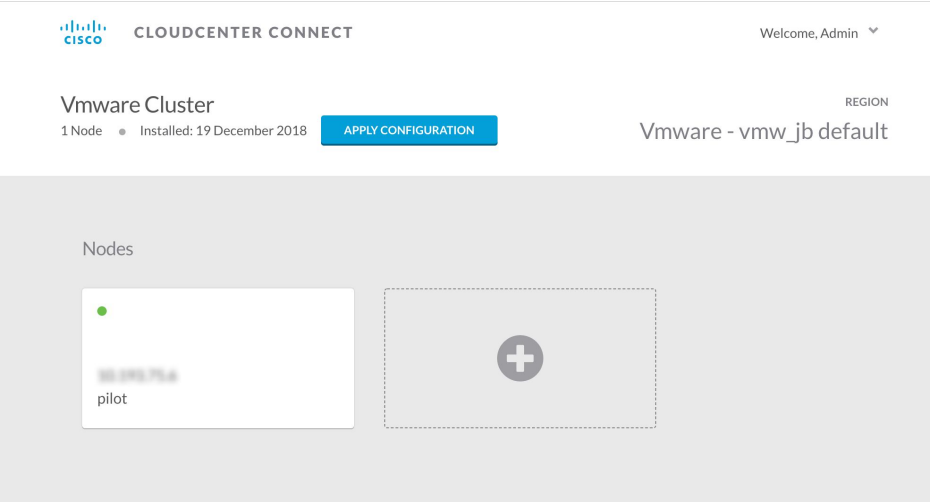

Switch your focus back to the Region Connectivity section of the target cloud region in the CloudCenter Suite web UI. The status indicator in the Region Connectivity section header will change from Not Configured to Running once connectivity between Cloud Remote and the CloudCenter Suite cluster is completely established (see figure below).

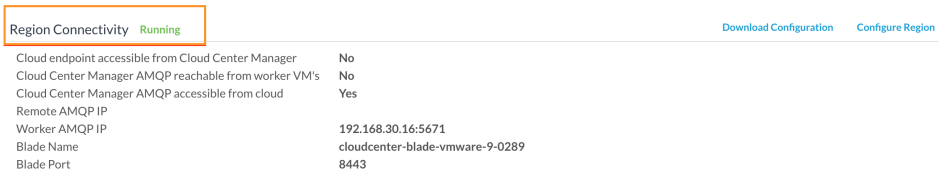

After completing these steps, Workload Manager and Cost Optimizer can use Cloud Remote for communicating with the target cloud region.

6. VM Naming and IPAM Strategy (conditional): Configure any VM naming or IPAM strategies in the Strategy section. If you leave the settings at the defaults, no IPAM strategy is applied and the default VM naming strategy is applied.

VM Naming and IPAM Strategies

Workload Manager lets you define a VM naming strategy and an IP address management (IPAM) strategy on a per cloud region basis. The VM naming strategy lets Workload Manager generate unique names for each of your VMs as they are launched. The IPAM strategy letsWorkl oad Manager generate IP addresses for each of the vNICs on each of your VMs. The VM naming strategy can reference a VM naming callout script, and the IPAM strategy requires that you specify an IP address allocation callout script andan IP address deallocation callout script. An IP address allocation callout scriptmay work in conjunction with a 3rd party IPAM server in your cloud environment.

Like External Lifecycle Actions, the scripts or commands associated with a VM naming strategy or IPAM strategy are executed in a containerized script execution engine running in the CloudCenter Suite cluster, if you are not using Cloud Remote in your region. If you are using Cloud Remote , these scripts or commands will execute in a containerized script execution engine running in Cloud Remote .

To set up the VM naming and IPAM strategies for your region, scroll down to the Strategy section of the Regions or Details tab and follow these steps.

Click the**Edit Strategy**link in the upper right to display the Edit Strategy dialog box as shown in the figure below.

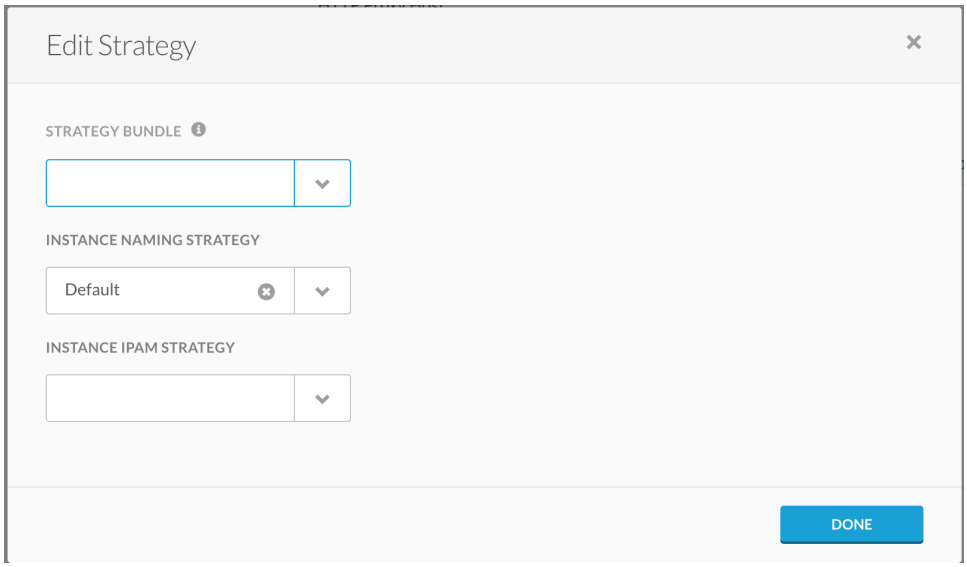

The dialog box can reveal a total of seven data entry fields as explained in the table below:

Field	Usage
Strategy Bundle	Zip file containing all of the scripts associated with your VM naming strategy or IPAM strategy. Select source from the dropdown an add the path in the new field that appears to the right. 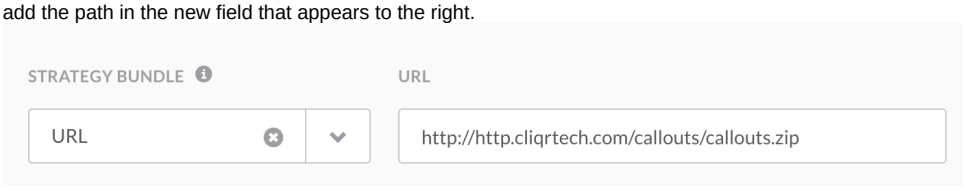
Instance Naming Strategy	Dropdown for selecting one of three naming strategies: Default, CliQr Macro Replacement, and Hostname Callout. If you select CliQ Macro Replacement, the **Node Name Config** field becomes visible and required. If you select Hostname Callout, the **Custom VM Name** field becomes visible and required. See the Instance Naming Strategy table below for details.

Node Name Config	Required if you selected "CliQr Macro Replacement" for theInstance Naming Strategyfield. Content: text string consisting of a combination of characters and numbers with embedded macro variables (see table below). The macros get translated to text strings run time.
	
	The resulting VM name string mustmeet the following requirements for job submission to succeed:
	• Only contain the following characters:
	  • Lowercase a to z
	  • 0 to 9
	  • Hyphen
	• Start withan alpha character
	• Cannot end with a hyphen

Instance IPAM Strategy	Dropdown menu with two choices: No IPAM (default), and IPAM Callouts. If you select the latter, two additional fields are displayed and become required: **IPAM Alloc Rule**and**IPAM Dealloc Rule**. (See below).
Custom VM Name	Location of your custom VM naming script if you set the**Instance Naming Strategy**field to "Custom Callout". Location options inclu "Script from bundle" if you specify a zip file in the **Strategy Bundle** field. See Guidance for Callout Scripts for guidance on developi custom VM naming scripts.
IPAM Alloc Rule	Location of your custom IPAM address allocation script if you set the**Instance IPAM Strategy**field to "IPAM Callouts". Location opti include "Script from bundle" if you specify a zip file in the **Strategy Bundle** field. See Guidance for Callout Scripts for guidance on developing custom IPAM callout scripts.
IPAM Dealloc Rule	Similar to the**IPAM Alloc Rule**above except this is where you specify the location of yourcustom IPAM address deallocation script. Guidance for Callout Scripts for guidance on developing custom IPAM callout scripts.

Instance Naming Strategy Options:

Options	Description
Default	The default method provided by Cisco is the generation of a string consisting of A predefined token prefixed with **cq jw-**
CliQr Macro Replacement	Selecting this method causes the **Node Name Config** field to be displayed and to become a required field. See the fields table above for details.
Hostname Callout	Selecting this method causes the **Custom VM Name** field to be displayed and to become a required field. See the fields table above for details.

OOB macro variables forCliQr Macro Replacement:

Macro	Returned Value
%os1%	First character of OS type string in lower case
%os2%	Firsttwo characters of OS type string in lower case
%OS1%	Firstcharacter of OS type string in upper case
%OS2%	Firsttwo character of OS type string in upper case
% RND<number>%	Fixed length random string, e.g., %RND10%, 10-character long random string ⚠ This is a mandatory field, only values in the range: %RND6% to %RND11% or just %RND% which defaults to %RND6%"
%UID%	User ID
%VID%	Vendor ID
%VM_PREFIX%	VM name prefix must be added as a global parameter in the app profile. ⚠ Only macro whose value can be user-customized.

7. External Lifecycle Actions (conditional):Specify any external lifecycle actions to be performed on all VMs launched by Workload Manager in this region.

External Lifecycle Actions Settings

Whenever you use Workload Manager to deploy a VM-based application, Workload Manager can execute external actions on VMs in the deployment as the VMs progress through these five phases of their lifecycle:

- Pre VM start
- Pre VM initialization
- PostVM initialization
- Pre VM Stop
- Post VM Stop

These lifecycle actions will execute in a script execution engine running in a container in the CloudCenter Suite cluster if no Cloud Remote appliance is configured for the region. If you have configured Cloud Remote for this region, the external lifecycle actions will executein a script execution engine running in a container in the Cloud Remote appliance.

External Lifecycle Action scripts can be defined at the service level, the application level, and the cloud region level. Use the External Lifecycle Actions section of the Region tab or Details tab (see figure below) tospecify external lifecycle actions at the region level.

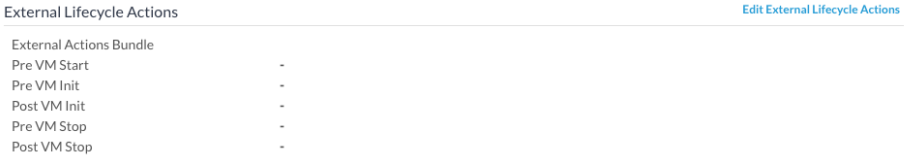

Click on the**Edit External Lifecycle Actions**link in the upper right. This displays theEdit External Lifecycle Actions dialog box as shown below.

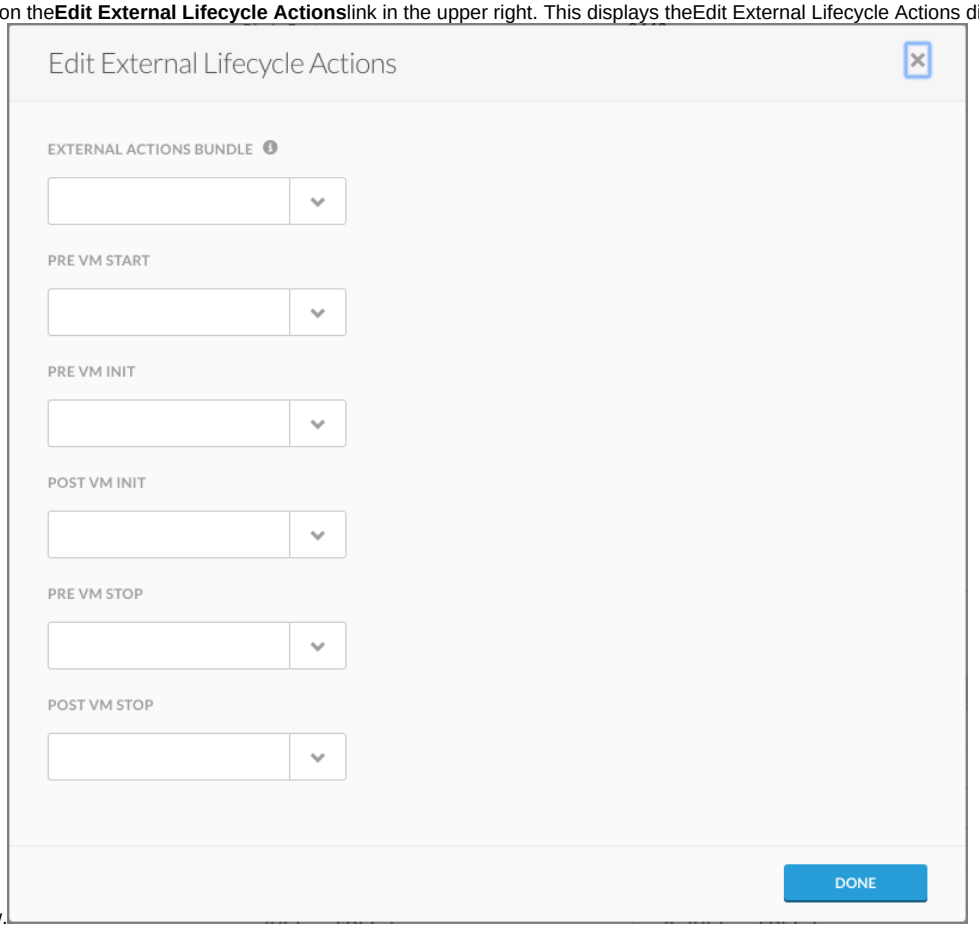

One way to specify an external lifecycle action is to use the first field,**External Actions Bundle**, to specify a zip file that contains all scripts for external lifecycle management. This zip file needs to contain a directory that contains all the scripts. Use the dropdown to select a repository location (URL, FTP, S3, etc.). This causes an additional field to be displayed to the right for entering the path to the zip file as shown in the figurebelow.

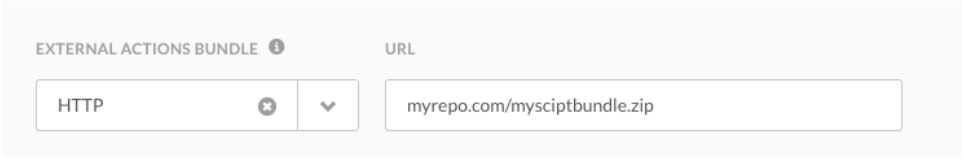

Once you specify actions bundle, for each of the individual lifecycle stage fields, select "Script from Bundle" from the dropdown. This displays an additional field to the right where you must specify the name of the script file in the bundle as shown in the figure below.

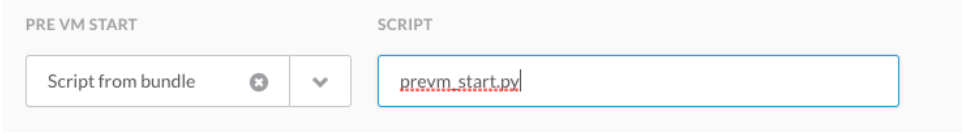

The scripts you specify can be bash scripts or Python scripts called from within a bash script.For more information on lifecycle scripts, seeDeployment Lifecycle Scripts.

8. Instance Types (informational): CloudCenter Suite automatically synchronizes instance types for public cloud regions on a daily basis. This data includes published pricing for each instance type. It is not possible to edit the IBM Cloud region instance types.

Instance Types Settings

An instance type is a performance category of a VM for VM-based clouds or a performance category of a container for container-based clouds. You can sync, edit, and add instance types depending on the cloud provider. For public clouds, CloudCenter Suite auto-populates the instance types in the Instance Types section from the cloud provider. CloudCenter Suite auto-syncs this information from the cloud provider once per 24 hours. For these clouds, you cannot force sync on demand.

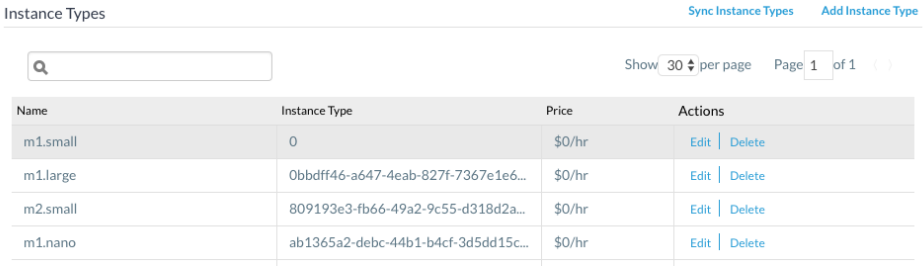

For all cloud providers except AWS, you can edit the fields of an instance type (except instance type ID) by clicking the**Edit**link in the Actions column in the list of instance types. The edits to these parameters are ignored during application deployment. The only parameter that will impact CloudCenter Suite behavior for public clouds is the price per hour, which is used by Workload Manager and Cost Optimizer for deployment cost calculations.

9. Storage Types (conditional): CloudCenter Suite automatically synchronizesstorage types for public cloud regions on a daily basis. This data includes the cloud provider published pricing for each storage type. It is not possible to edit the IBM Cloud region storage types.

Storage Types Settings

A storage type is a categorization of storage that can be provisioned for an application during deployment. It is a construct that applies to VM-based clouds only. Storage types are displayed in the Storage Types section of the Regions or Details tab. The storage type construct contains the following fields for all VM-based clouds:

- Display Name
- Cost Per Month per GB
- Cloud Storage Type ID
- Minimum Volume Size (GB)
- Maximum Volume Size (GB)
- Maximum Throughput (MB/s)
- Maximum IOPS
- Provisioned IOPS
- IOPS Cost Per Month

Storage types are handled by CloudCenter Suite for various cloud providers in a fashion similar to instance types. For public clouds, CloudCenter Suite auto-populates the list of storage types in the Storage Types section based on information from the cloud provider. This data is stored in thepackage store and the package store auto-syncs this information from the cloud provider once per 24 hours. You may manually add additional storage types by clicking the**Add Storage Type**link in the upper right of the Storage Types settings section as shown in the figure below.

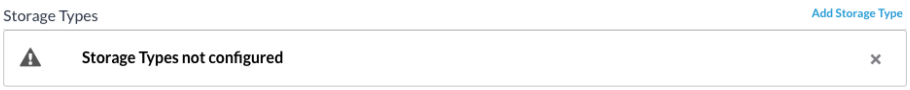

This brings up an**Add Storage Type**dialog box as shown in the figure below.

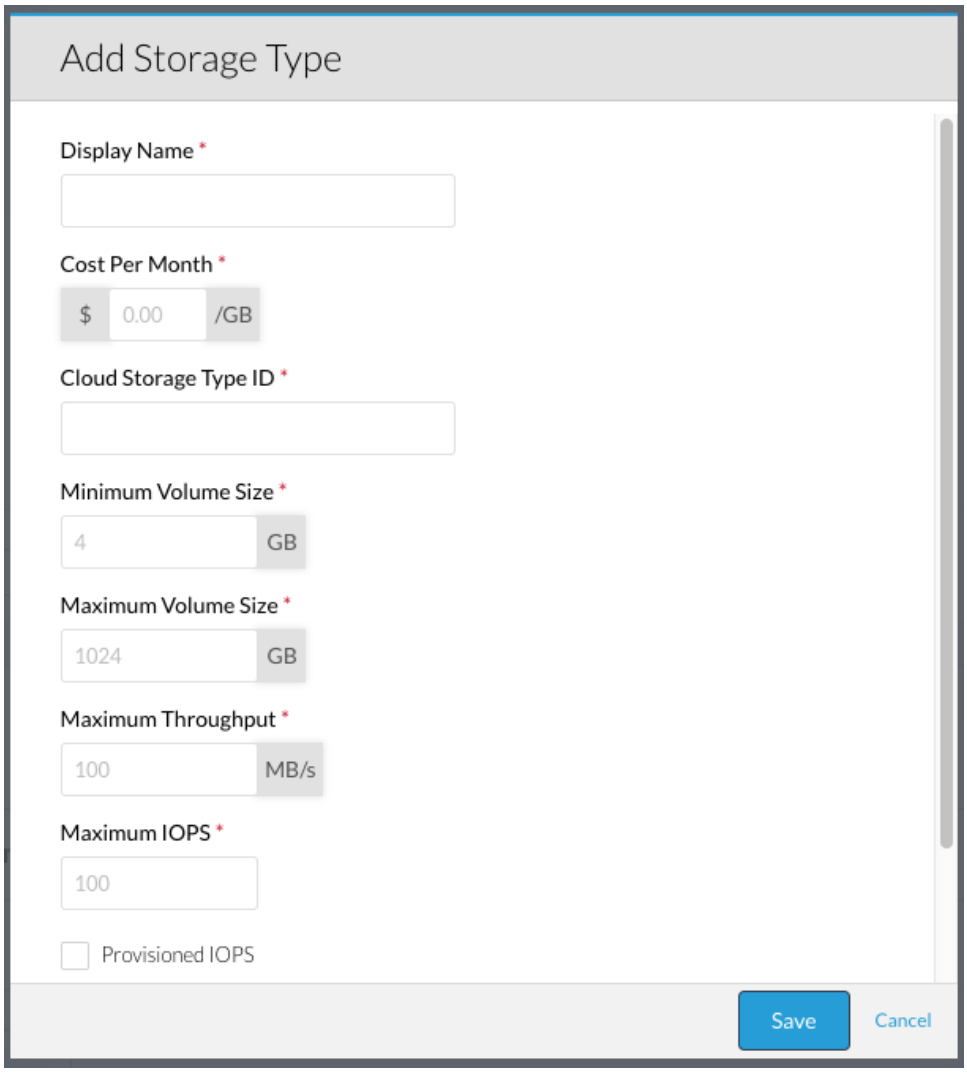

You must enter all required fields, however, the**Maximum Throughput**,**Maximum IOPS**,**Provisioned IOPS**, and**IOPS Cost Per Month**are not recognized by these cloud providers and will be ignored.

You must specify the**Cost**to compute the service costs in Cost Optimizer. The total cost is displayed in theCost Optimizer Dashboard.

Storage types for all cloud providers except AWS can be edited. From the list of storage types in the Storage Type section (see figure below), click the Edit link (in the Actions column) for that storage type.

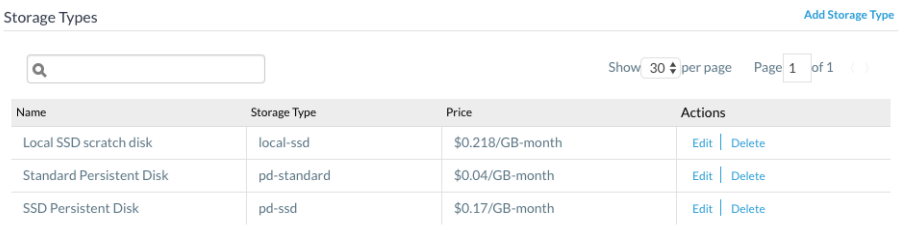

This brings up an Edit Storage Type dialog box (similar to the Add Storage Type dialog box) where you can edit all nine fields except the**Cloud Storage Type ID**field. For AzureRM and Google clouds, only changes to the pricing fields change the behavior of CloudCenter Suite: the cloud cost calculations are modified based on the new unit cost data. All other field changes are ignored.

For all cloud providers,**Minimum Volume Size**and**Maximum Volume Size**restrict the volume size you can specify at deploy time. For public cloud providers that support provisioned IOPS, the user can specify the guaranteed IOPS at deploy time.

10. Image Mappings: Image mappings allow services based on CloudCenter Suite logical images to be deployed using the appropriate physical image stored on the target cloud region. CloudCenter Suite automatically maps theOOB logical imagesto public cloud region physical images when you add the region to your cloud. Cisco periodically updates these mappings when new versions of OS physical images are uploaded by the cloud provider. To apply these updates to your region after it is added to your cloud, click the**Sync Image Mappings**link in the upper right of this section. If you create any custom logical images, you mustmanually import the corresponding physical images into your region and then map the corresponding logical images to these physical images. SeeImagesfor more context.

Add an IBM Cloud Cloud Account

Configuration Process

To add an IBM Cloud cloud account, follow this procedure.

1. Locate your IBM Cloud cloud on the **Clouds** page and click the**Add Cloud Account**link for this cloud.This displays the **Add Cloud Account** dialog box as shown below.

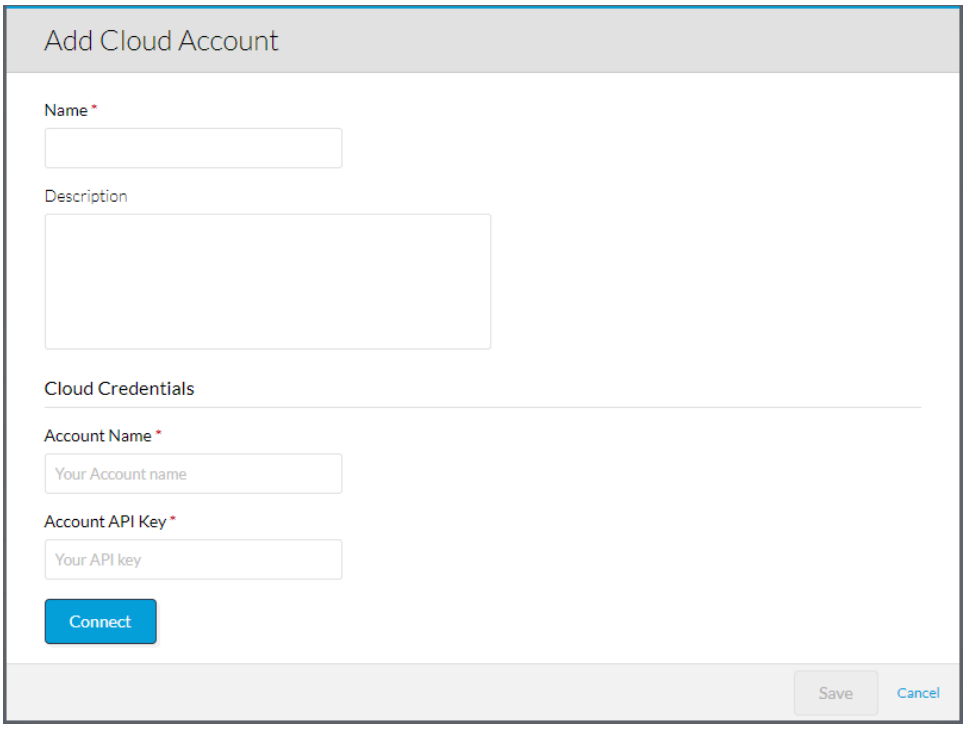

2. Assign a cloud account**Name**.

> **Tip**
>
> The name should not contain any space, dash, or special characters.

3. Provide the IBM Cloud cloud credentials:
    a. **IBM Cloud Account Name**
    b. **IBM Cloud Account API Key**
4. Click the**Connect**button. CloudCenter Suite will now attempt to validate your account credentials.
5. After the credentials are verified, the**Connect**button changes to an**Edit**button and two new fields appear**Enable Account For**and**Enable Reporting By Org Structure**,
    a. Set the**Enable Account For**dropdown per the table below.

Value	Usage
Provisioning	Workload Manager can deploy jobs using this account.

Reporting	Cost Optimizer and Workload Manager will track cloud costs for this account. Typical usage: master cloud accounts that are used for billing aggregation. ⚠ It is recommended that you do not add a *Reporting* account to the same tenant through different cloud groups. ⚠ Enabling a public cloud account for *Reporting* may incur expenses to retrieve cost data. These expenses are proportional to the number of configured cloud accounts and regions.
Provisioning, Reporting	Default. Account is used for both provisioning and reporting.

b. Click the**Save**button when done.

Cloud Accounts Tab

After you add cloud accounts to a cloud, they will appear in the Accounts tab for the cloud as shown in the figure below.

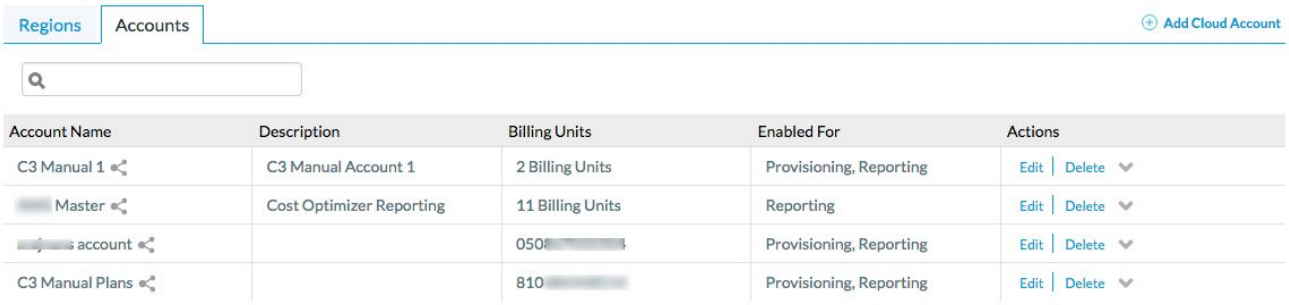

The Accounts tab contains columns for data entered when creating an account: Account Name, Description, Enabled For; and two additional columns:**Billing Units**and**Actions**.**Billing Units**is a dual function:

- If the cloud account contains only one billing unit, the ID for that billing unit is displayed.
- If the cloud account contains multiple billing units, such as an AWS master account, the number of billing units in that account is displayed followed by the text*Billing Units*.

A billing unit is the most granular level of cloud cost recording in CloudCenter Suite. The definition of a billing unit varies by a cloud provider as shown in the table below.

Cloud Provider	Billing Unit
AWS	Account ID
AzureRM	Subscription ID
Google	Project ID
IBM Cloud	Account ID
vCenter	Cloud Group Prefix - Datacenter Name
vCD	Organization Name
OpenStack	Project ID
Kubernetes	Namespace UID

The last column,**Actions**, contains links to let you edit or deleted the cloud account, ormanage instance typesfor the cloud account.

Configure an OpenStack Cloud

Configure an OpenStackCloud

Configuring an OpenStackcloud is a four-step process:

- Add an OpenStackCloud
- Add an OpenStack Region
- Configure an OpenStackRegion
- Add an OpenStack Cloud Account

Add an OpenStackCloud

To add an OpenStackcloud follow these steps.

1. Navigate to**Admin > Clouds**. This brings you to the**Clouds**page. If you, or another tenant admin in your tenant, have already added clouds to your tenant, they will be listed here.
2. Click the**Add Cloud**linkin the upper right. The**Add Cloud**dialog box is displayed.
3. Enter the**Cloud Name**and select the**cloud provider**.

   When assigning the Cloud Name, use a unique name. Ensure that you do not assign a previously-used name (including deleted cloud names) as the Cloud Name as they may cause deployments to fail intermittently.

   When done click**Next**.The second page of the **Add Clouds** dialog box, **Connectivity Settings**, appears. Set the toggle switches to configure the Cloud Connectivity Settings.

   - When adding a private VM cloud in the Workload Manager or Cost Optimizer UI, the second page of the Add Clouds dialog box, Connectivity Settings, appears with two toggles displayed:
     - **Worker VMs Directly Connect with CloudCenter Suite**
     - **VMs Directly Connect with CloudCenter Suite**
   - Setting either of these toggles to No implies you will install Cloud Remote for each region of this cloud. This also causes a third toggle to appear:**CloudCenter Suite Directly Accessible from Cloud Remote**.
   - Follow the table below for guidance on setting these toggles.

Toggle settings	Use case	Network Diagram
Cloud Endpoint Directly Accessible = Yes AND VMs Directly Connect with CloudCenter Suite = Yes	CloudCenter Suite cluster can initiate a connection to the cloud region API endpoint AND Worker VMs can initiate a connection to the CloudCenter Suite cluster Cloud Remote is not required	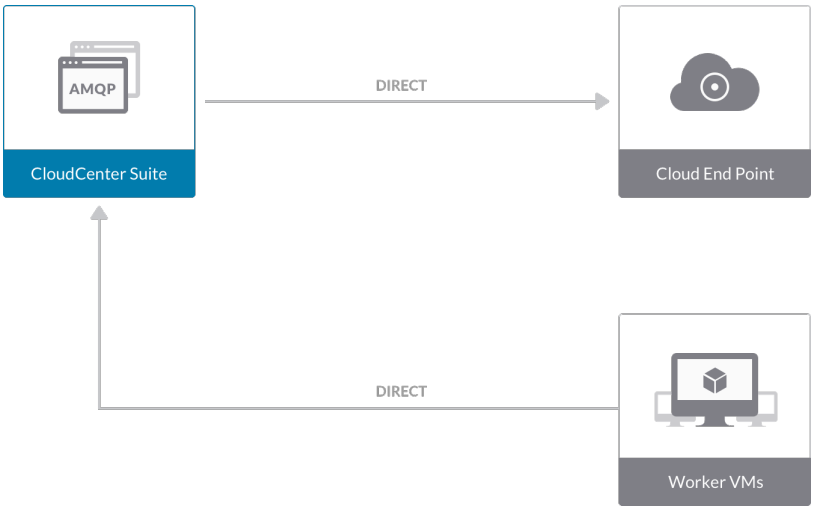

Cloud Endpoint Directly Accessible = No AND Worker VMs Directly Connect with CloudCenter Suite = No AND CloudCenter Suite Directly Accessible fromCloud Remote = Yes	CloudCenter Suite cluster cannot initiate a connection to the cloud region API endpoint AND Worker VMs cannot initiate a connection to the CloudCenter Suite cluster AND Cloud Remote can initiate the connection to the CloudCenter Suite cluster	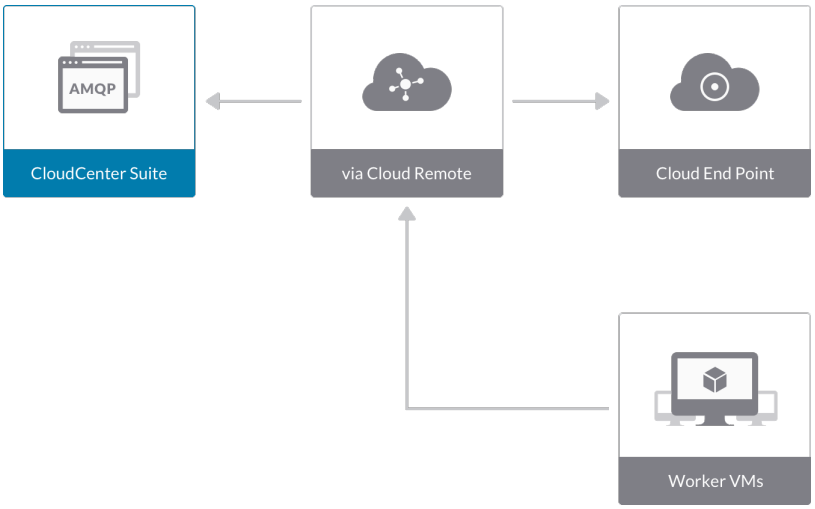
Cloud Endpoint Directly Accessible = No AND Worker VMs Directly Connect withCloudCenter Suite = No AND CloudCenter Suite Directly Accessible from Cloud Remote = No	CloudCenter Suite cluster cannot initiate a connection to the cloud region API endpoint AND Worker VMs cannot initiate a connection to the CloudCenter Suite cluster AND Cloud Remote cannot initiate the connection to the CloudCenter Suite cluster	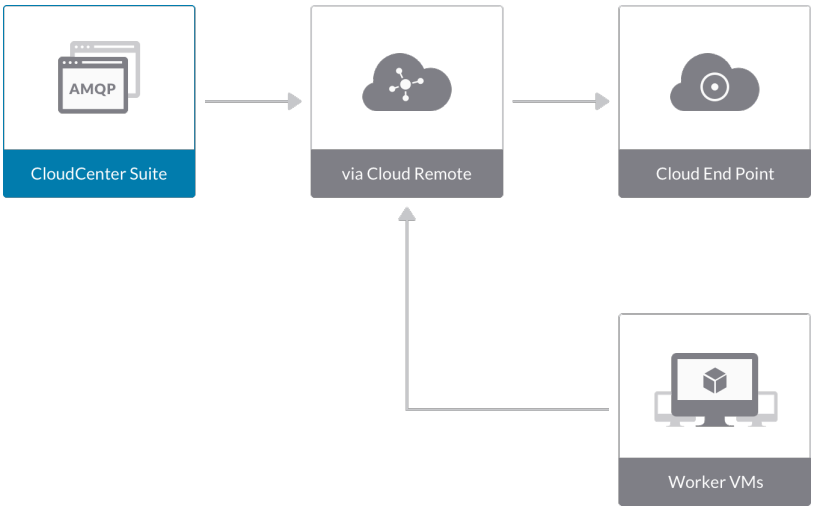
Cloud Endpoint Directly Accessible = Yes AND Worker VMs Directly Connect withCloudCenter Suite = No AND CloudCenter Suite Directly Accessible from Cloud Remote = No	CloudCenter Suite cluster can initiate a connection to the cloud region API endpoint AND Worker VMs cannot initiate a connection to the CloudCenter Suite cluster AND Cloud Remote cannot initiate the connection to the CloudCenter Suite cluster	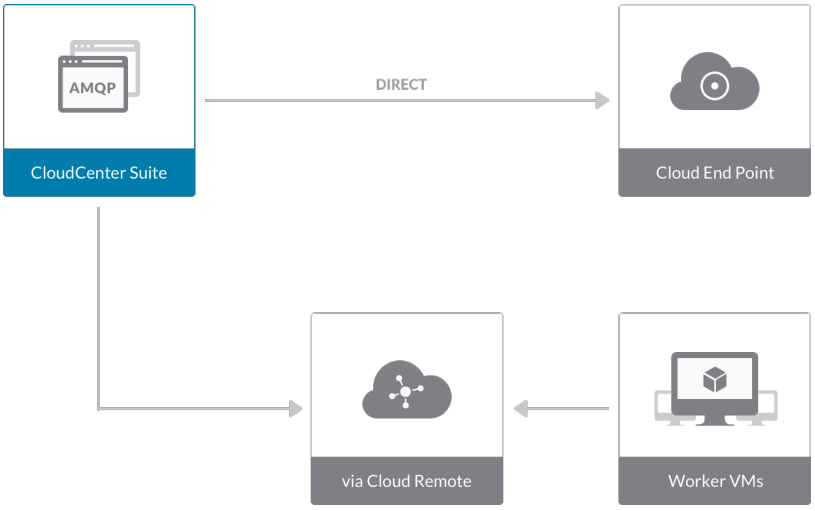

Cloud Endpoint Directly Accessible = Yes AND Worker VMs Directly Connect withCloudCenter Suite = No AND CloudCenter Suite Directly Accessible fromCloud Remote = Yes	CloudCenter Suite cluster can initiate a connection to the cloud region API endpoint AND Worker VMs need to communicate to the CloudCenter Suite cluster through Cloud Remote AND Cloud Remote can initiate the connection to the CloudCenter Suite cluster	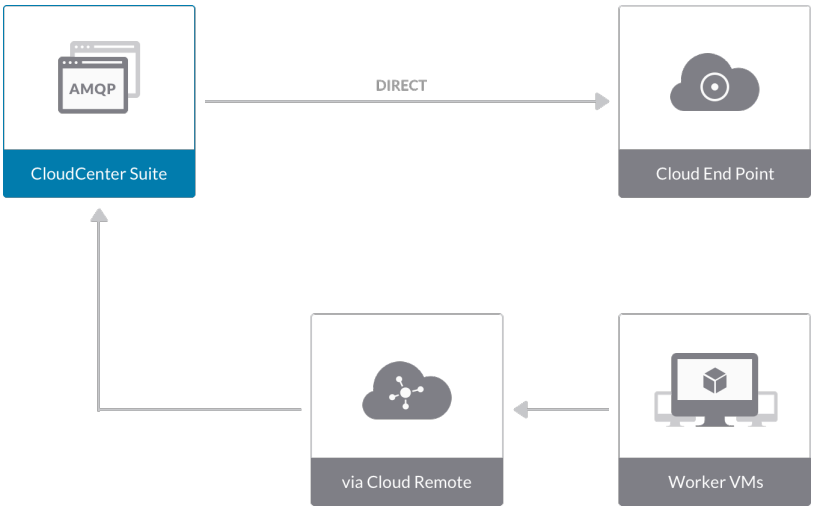

Note

The connectivity toggle settings set at the cloud level are inherited by each region you add to this cloud. However, it is possible to override these toggle settings on a per-region basis from the Regions tab for each region.

4. Click**Done**to save the configuration and close the dialog box. This brings you back to the **Clouds** page and the cloud you just created will be added to the bottom of the list on the left side of the page.

Add an OpenStack Region

After creating an OpenStack cloud, the next step is to create the first region for the cloud. Follow these steps.

1. Navigate to the **Clouds** page and select the cloud you created on the left side of the screen.
2. Click the**Add Region**button onthe right side of the screen. The **Add Region** dialog box is displayed.
3. Enter a**Region Name**and**Display Name**.

When assigning the Region Name, use a unique name. Ensure that you do not assign a previously-used name (including deleted cloud names) as the Region Name as they may cause deployments to fail intermittently.

4. Click**Save**.You are brought to the **Clouds** page with theregion you added shown on the right side of the page.

Configure an OpenStackRegion

To configure a region you added to your OpenStackcloud, perform the following steps.

1. Navigate to Clouds page:**Admin > Clouds** tofind your OpenStackcloud from the cloud list on the left half of the screen
2. Click its**Configure Cloud**link. This displays the **Regions** tab for this cloud as shown in the figure below with the**Cloud Settings** section displayed first.After you have added multiple regions to your OpenStackcloud, the **Regions** tab will show multiple individual region tabs on the left side of the screen.

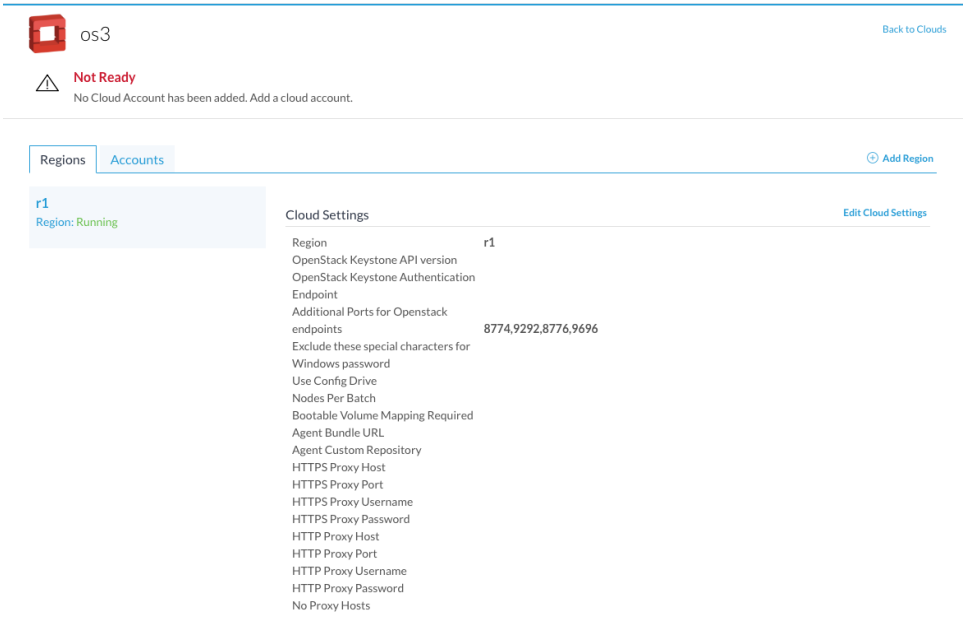

3. Click the tab of the region you want to configure.
4. Click the**Edit Cloud Settings**link in the upper right of the **Cloud Settings** section. This opens the **Configure Cloud Settings** dialog box.The **Cloud Settings** section contains fields that are unique to OpenStack and settings that are common to all cloud providers. Adjust these field values per the instructions in the following tables.

OpenStack Specific Cloud Settings

Field	Usage
Region	This is a read-only field based on the region name you entered when you created this region.
OpenStack Keystone API version	The default value is V2. Use the dropdown menu to change this to V3 if your version of OpenStack supports the V3 API.
OpenStack Keystone Authentication Endpoint	Enter the URL of your OpenStack API endpoint.
Additional Ports for OpenStack endpoints	These are pre-populated with the standard ports for communication between the OpenStack API and Workload Manager. Only change these values if you have a non-standard network configuration for OpenStack.
Use Config Drive	This is unchecked by default. Check this box if your deployments need to use configdrive.
Nodes Per Batch	This is the maximum number of VMs that can be launched simultaneously per application deployment. If left blank, the default value of 1 is applied. A value of 0 or 1 both means only one VM will be launched at a time.
Bootable Volume Mapping Required	Default means no mapping. You only need to change this field if OpenStack is configured along with a third-party infrastructure that is not visible to Workload Manager.

Agnostic Cloud Settings

Field	Usage
Exclude these special characters for Windows password	When the Workload Manager agent is installed on a Windows worker VM, a special user account, called cliqruser, is created to support RDP sessions that may be initiated by the user through the Workload Manager UI. A Workload Manager process running on the CloudCenter Suite cluster creates a random password and passes it to the agent for creating the cliqruser account. Because some Windows deployments may restrict using certain characters for Windows passwords, this field is provided to tell the Workload Manager to exclude these special characters in the generation of the password for the cliqruser account.
Agent Bundle URL	If you plan to use a local repository to host the bundle store, you need to enter the URL of the local bundle store here. Otherwise, leave blank.
Agent Custom Repository	If you plan to use alocal repository to host the package store, you need to enter the URL of the local package store here. Otherwise, leave blank.

HTTP /HTTPS proxy fields (host, username, password)	If you require VMs in your region to access public addresses through a web proxy, enter the URL and credentials of the HTTP and HTTPS proxy servers in these fields.
No Proxy Hosts	If you have specified an HTTP or HTTP proxy using the above fields, you can specify that managed VMs in the region should bypass the proxy and connect directly to certain hosts. Use this field to create a comma-separated list of IP addresses or URLs that should be accessed directly. This field is ignored if an HTTP or HTTPS proxy is not specified.

Important information on proxy settings

In CloudCenter Suite, you can specify proxy settings at the region level, as described here, and at the suite level. To understand the expected behavior when proxy settings are specified at both levels, see the subsequent *Precedence of Proxy Settings* section.

Download Configuration and Encryption Key

After saving the Region Configuration settings in the Workload Manager or Cost Optimizer UI, you can download them to your local computer and then upload them to other conditional components such as Cloud Remote .

The Configuration and Encryption key is only visible when you have configured the Cloud Remote component. Click the**Download Configuration** link in the upper right of the Region Connectivity section, as shown in the following screenshot.

Clicking**DownloadConfiguration**causestwo things to happen:

- An encrypted zip file namedartifacts.zipis downloaded by your browser. Make a note of the location of this zip file as you will need if you are using Cloud Remote .
- The Region Connectivity section header updates to display a**Copy Encryption Key**link, as shown in the following screenshot.

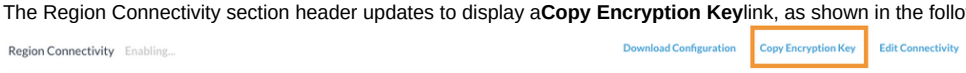

Click the**Copy Encryption Key**link to save the key to your clipboard. A success message will be displayed temporarily above the Region Connectivity section header. Make sure not to overwrite the clipboard with other data. You will need the key when you upload the configuration zip file to conditional components like Cloud Remote .

If you change the connectivity settings in the CloudCenter Suite UI and need to again download the zip file fromsoftware.cisco.com, use the automatically create a (new) encryption key, and copy the key to the clipboard by clicking the**Copy Encryption Key**link again.

Precedence of Proxy Settings

In CloudCenter Suite, you can specify HTTP and HTTPS proxy settings at the suite level as described inProxy Settings, and at the region level as described in the preceding*Agnostic Cloud Settings*section. The CloudCenter Suite cluster, Cloud Remote , and worker VMs will each use either the suite-level proxy settings or the region level proxy settings based on the conditions described below.

For Cloud Remote mode, whenever you change the region proxy settings, ensure you again download the region connectivity setting from the Workload Manager UI and reapply to Cloud Remote .

Mode	Proxy Settings Used By		
Mode (Is Cloud Remote used to communicate withcloud (or APIC) endpoint?)	CloudCenter Suite cluster to communicate with cloud (or APIC) endpoint	Cloud Remote to communicate with cloud endpoint	Worker VM to communicate with bundle store and package store
Cloud Remote Mode	NA	Region-level	Region-level
Non-Cloud Remote Mode	Suite-level	NA	Region-level

Not all clouds support all the proxy settings, and not all clouds support bothCloud Remote and non-Cloud Remote modes. OpenStack supports Cloud Remote mode and supports all proxy settings except proxy hosts innon-Cloud Remote mode.
When you are done editing the settings in the dialog box, click**Save**.

5. Determine if you need Cloud Remote for this region. Scroll down to the Region Connectivity section for the region and click on the**Configure Region**link in the upper right to open the **Configure Region** dialog box.The toggle settings should be the same as when you set them on the connectivity page of the **Add Cloud** dialog box.If all of the connectivity togglesinthe **Region Connectivity** dialog boxare set to **Yes**, then Cloud Remote is NOT needed for this cloud region. In this case,you would normally leave the region connectivity settings at their current values and continue to the next settings section.

The exceptionto this guidance is when a NAT firewall or proxy server exists between theCloudCenter Suite management cluster andworker VMs, or between the CloudCenter Suite management cluster and users that would use Workload Manager to initiate a Guacamole remote connection to a worker VM. In either of these cases, override the address fields in the Region Connectivity dialog box as explained below.

Networking Constraint	Field	Value
Worker VMs must use a proxy server or NAT firewall to access the"local" AMQP server running in the CloudCenter Suite cluster.	Worker AMQP IP Address	IP address and port number that the firewall or proxy server presents to the worker VMs on behalf of the "local" AMQP server running in the CloudCenter Suite cluster.
Users must use a proxy server or NAT firewall to access theGuacamole server running in the CloudCenter Suite cluster.	Guacamole Public IP Address and Port	IP address and port number that the firewall or proxy server presents to users on behalf of theGuacamole server running in the CloudCenter Suite cluster.
Worker VMs must use a proxy server or NAT firewall to access theGuacamole server running in the CloudCenter Suite cluster.	Guacamole IP Address and Port for Application VMs	IP address and port number that the firewall or proxy server presents to the worker VMs on behalf of the Guacamole server running in the CloudCenter Suite cluster.

6. Click**OK**to save the changes and dismiss the dialog box. You can now proceed to the next region settings section:VM Naming and IPAM Strategy. If any of theconnectivity togglesinthe **Region Connectivity** dialog box are set to No, then**you must install and configure Cloud Remote for this region**.

Configure Cloud Remote in an OpenStack Region

Configure Cloud Remote in an OpenStack region as follows.

Download and Launch the Cloud Remote Appliance in OpenStack

a. Download the Cloud Remote appliance qcow2 file fromsoftware.cisco.com.
b. Through the OpenStack console, import and launch the Cloud Remote appliance. This process is similar to importing and launching theCloudCenter Suite installer appliance for OpenStack.

> ⚠ Do not add Network Ports while launching a Cloud Remote instance in OpenStack.

c. Optional but recommended for production environments: Deploy two additional instances of the appliance to form a cluster for HA.Cloud Remote includes support for the clustering of multiple nodes. You will "add" these two additional instances to the first instance after the first instance is configured. SeeCloud Remote (Conditional)> *Scaling*for details.
d. Once the first instance of the appliance has been launched, use theOpenStack console to**note its IP public and private addresses**. You will need this information later on in order to login to the Cloud Remote web UI and to complete the Region Connectivity settings in the CloudCenter Suite Web UI. Also, note the IP addresses of any other appliances you launch.

Setup Cloud Remote Firewall Rulesfor a VM-based Cloud Region

After you deploy the Cloud Remote appliance, you will need to open various ports on each instance of the appliance. To do this, use the tools provided by the cloud provider to create a new security group for your Cloud Remote cluster; then, associate each appliance in the cluster with that security group. Use the tables below for guidance on what port rules should be added to that security group.

Port rules for a single node Cloud Remote deployment:

Port	Protocol	Source	Usage
22	TCP	Limit to address space of users needing SSH access for debugging and changing default ports	SSH
443	TCP	Limit to address space of users needing access to the Cloud Remote web UI for setup and scaling	HTTPS (Cloud Remote web UI)
8443	TCP	Limit to address space of users needing SSH or RDP access to their managed VMs	User to Guacamole
5671	TCP	Limit to address space of the managed VMs and the address of the CloudCenter Suite cluster's local AMQP service	AMQP
15671	TCP	Limit to address space of users needing web access for debugging the remote AMQP service	HTTPS (AMQP Management)
7789	TCP	Limit to address space of the managed VMs	Worker VM to Guacamole

> ⚠ The Cloud Remote web UI, User-to-Guacamole, and AMQP ports listed above are the defaults used by Cloud Remote . You may change these port numbers using the **Change Ports shell script** (see Cloud Remote (Conditional) >Custom Port Numbers (Conditional)) once the appliance is fully configured and communicating with the CloudCenter Suite cluster. If you plan to modify any of these three port numbers, update the firewall rules accordingly.

For a multi-node Cloud Remote cluster deployment, these additionalport rules should be added to the same security group used for the single node configuration:

Port	Protocol	Source
2377	TCP	<cr_sec_group> *
25672	TCP	<cr_sec_group>
7946	UDP	<cr_sec_group>
4369	TCP	<cr_sec_group>
9010	TCP	<cr_sec_group>
4789	UDP	<cr_sec_group>

* <cr_sec_group> represents the security group that all Cloud Remote nodes are joined to.

Specify AMQP and Guacamole Addresses for Supporting Cloud Remote

From the CloudCenter Suite UI, for the cloud region requiring Cloud Remote, navigate to the corresponding Regions or Details tab. Click the**Configure Region**link in the upper left of the Region Connectivity section to bring up the Configure Region dialog box. The toggle settings should be the same as when you set them on the connectivity page of the Add Cloud dialog box. You must update some of the address fields in the dialog box according to the scenarios summarized in the table below.

Toggle Settings	Field	Value
Worker VMs Directly Connect with CloudCenter = No AND CloudCenter Directly Accessible from Cloud Remote = Yes	Local AMQP IP Address	Pre-populated with the address and port number of the "local" AMQP server running in the CloudCenter Suite cluster. This address must be **accessible to Cloud Remote** . If Cloud Remote is accessing the CloudCenter Suite cluster through a user-supplied proxy server or NAT firewall, overwrite this field with the corresponding local AMQP IP address and port number provided by the user-supplied proxy server or NAT firewall and accessible to Cloud Remote .
Worker VMs Directly Connect with CloudCenter = No AND CloudCenter Directly Accessible from Cloud Remote = No	Remote AMQP IP Address	Enter <Cloud_Remote_IP>:<amqp_port>, where <Cloud_Remote_IP> = the IP address Cloud Remote which is **accessible to the CloudCenter Suite cluster**, and <amqp_port> = 5671 OR the custom AMQP port number you would later set with the **Change Ports shell script** on the Cloud Remote appliance (see Cloud Remote (Conditional) > *Custom Port Numbers (Conditional)*). If there is no user-supplied NAT firewall or proxy server between the CloudCenter Suite cluster and Cloud Remote , the IP address would be the public IP address of Cloud Remote . If there is a NAT firewall or proxy server between the CloudCenter Suite cluster and Cloud Remote , instead, enter the corresponding public IP address and port number that the firewall or proxy server presents to the internet on behalf of the "remote" AMQP server running in Cloud Remote .
Worker VMs Directly Connect with CloudCenter = No	Worker AMQP IP Address	Enter <Cloud_Remote_IP>:<amqp_port>, where <Cloud_Remote_IP> = the Cloud Remote IP address **accessible to the worker VMs**, and <amqp_port> = 5671 OR the custom AMQP port number you would later set with the **Change Ports shell script** on the Cloud Remote appliance (see Cloud Remote (Conditional) > *Custom Port Numbers (Conditional)*).
Worker VMs Directly Connect with CloudCenter = No	Guacamole Public IP and Port	Enter <Cloud_Remote_IP>:<guac_port>, where <Cloud_Remote_IP> = the Cloud Remote IP address **accessible to CloudCenter Suite users**, and <guac_port> = 8443 OR the custom Guacamole port number you would later set with the **Change Ports shell script** on the Cloud Remote appliance (see Cloud Remote (Conditional) > *Custom Port Numbers (Conditional)*).
Worker VMs Directly Connect with CloudCenter = No	Guacamole IP Address and Port for Application VMs	Enter <Cloud_Remote_IP>:<guac_port>, where <Cloud_Remote_IP> = the Cloud Remote IP address **accessible to worker VMs**, and <guac_port> = 7789

When done, click **OK** to save the setting and dismiss the dialog box.

Download Region Connectivity Settings and Upload to Cloud Remote

Cloud Remote uses the region connectivity settings set in the Workload Manager or Cost Optimizer UI. After saving the Region Configuration settings in the Workload Manager or Cost Optimizer UI, you must download them and to your local computer and then upload them to Cloud Remote as follows.

Click the**Download Configuration**link in the upper right of the Region Connectivity section, as shown in the figure below.

Clicking DownloadConfiguration causestwo things to happen:

- An encrypted zip file named**artifacts.zip**will be downloaded by your browser. Make note of the location of this zip file as you will need to upload it to Cloud Remote through the Cloud Remote web UI (see below).
- The Region Connectivity section header updates to display a **Copy Encryption Key**link, as shown in the figure below.

Click the**Copy Encryption Key**link to save the key to your clipboard. A success message will be displayed temporarily above the Region Connectivity section header. Make sure not to overwrite the clipboard with other data. You will need the key when you upload the configuration zip file to Cloud Remote .

> If you change the connectivity settings in the CloudCenter Suite UI and need to again download the zip file, a new encryption key is automatically created and can be copied to the clipboard by clicking the **Copy Encryption Key** link again.

After you have downloaded the zip file and copied the encryption key to your clipboard, login toCloud Remote web UI.

a. Open another browser tab and login to https://<Cloud Remote _ip> with thedefault credentials: admin/cisco.
b. You will immediately be required to change your password. Do so now.
c. You are now brought to the Cloud Remote home page as shown in the figure below.

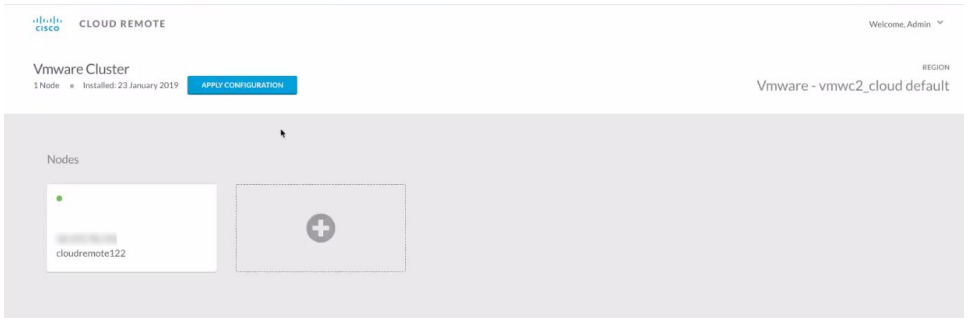

d. Click the **Apply Configuration** button in the page header. This prompts you to select a configuration file and enter the encryption key as shown in the figure below.

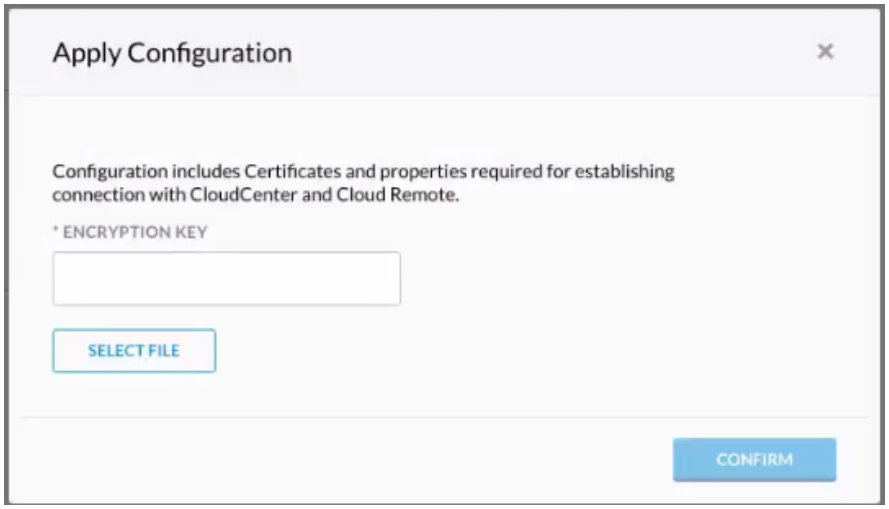

e. Paste the encryption key that was copied to the clipboard into the**Encryption Key**field in the dialog box.
f. Click**Select File**and browse tothe artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
g. Click**Confirm**.
h. Once the zip file is successfully transmitted and accepted, the Cloud Remote appliance attempts to establish communication with the CloudCenter Suite cluster and the Cloud Remote web UI home page is updated to show the name of the region it is connecting to in the upper right (see figure below).

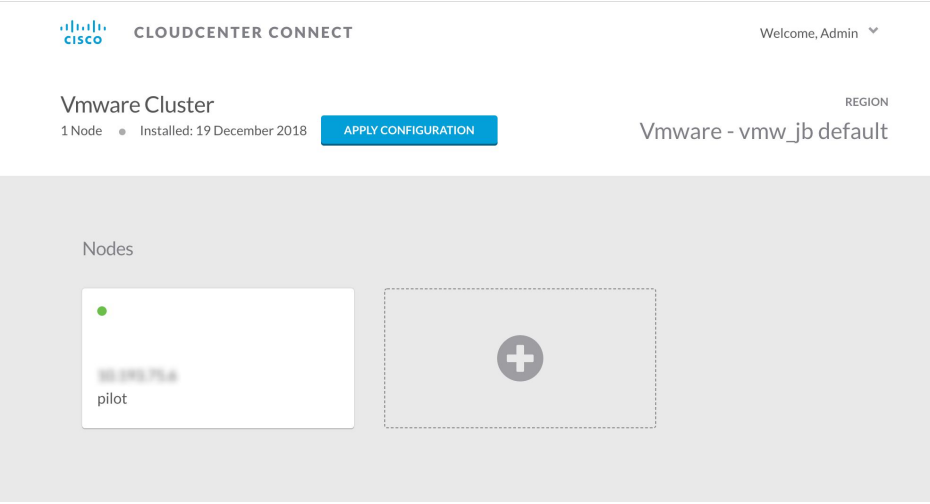

Switch your focus back to the Region Connectivity section of the target cloud region in the CloudCenter Suite web UI. The status indicator in the Region Connectivity section header will change from Not Configured to Running once connectivity between Cloud Remote and the CloudCenter Suite cluster is completely established (see figure below).

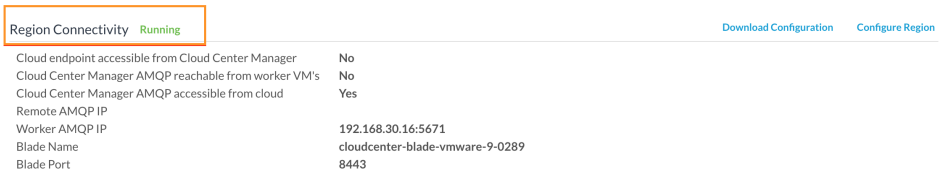

After completing these steps, Workload Manager and Cost Optimizer can use Cloud Remote for communicating with the target cloud region.

7. VM Naming and IPAM Strategy (conditional): Configure any VM naming or IPAM strategies in the Strategy section. If you leave the settings at the defaults, no IPAM strategy is applied and the default VM naming strategy is applied.

VM Naming and IPAM Strategies

Workload Manager lets you define a VM naming strategy and an IP address management (IPAM) strategy on a per cloud region basis. The VM naming strategy lets Workload Manager generate unique names for each of your VMs as they are launched. The IPAM strategy letsWorkl oad Manager generate IP addresses for each of the vNICs on each of your VMs. The VM naming strategy can reference a VM naming callout script, and the IPAM strategy requires that you specify an IP address allocation callout script andan IP address deallocation callout script. An IP address allocation callout scriptmay work in conjunction with a 3rd party IPAM server in your cloud environment.

Like External Lifecycle Actions, the scripts or commands associated with a VM naming strategy or IPAM strategy are executed in a containerized script execution engine running in the CloudCenter Suite cluster, if you are not using Cloud Remote in your region. If you are using Cloud Remote , these scripts or commands will execute in a containerized script execution engine running in Cloud Remote .

To set up the VM naming and IPAM strategies for your region, scroll down to the Strategy section of the Regions or Details tab and follow these steps.

Click the**Edit Strategy**link in the upper right to display the Edit Strategy dialog box as shown in the figure below.

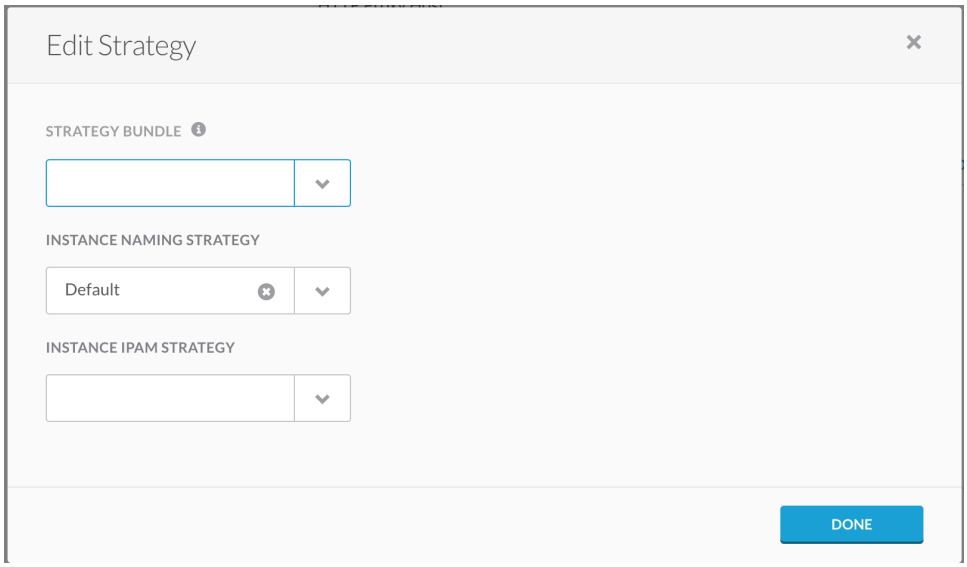

The dialog box can reveal a total of seven data entry fields as explained in the table below:

Field	Usage
Strategy Bundle	Zip file containing all of the scripts associated with your VM naming strategy or IPAM strategy. Select source from the dropdown an add the path in the new field that appears to the right. 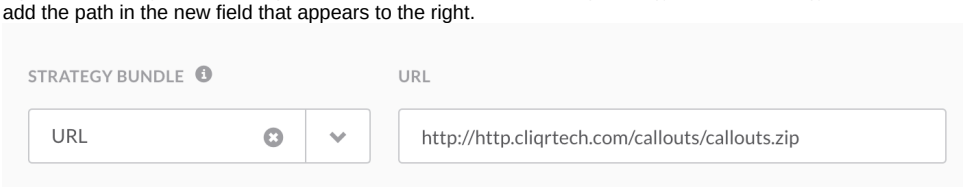
Instance Naming Strategy	Dropdown for selecting one of three naming strategies: Default, CliQr Macro Replacement, and Hostname Callout. If you select CliQ Macro Replacement, the **Node Name Config** field becomes visible and required. If you select Hostname Callout, the **Custom VM Name** field becomes visible and required. See the Instance Naming Strategy table below for details.

Node Name Config	Required if you selected "CliQr Macro Replacement" for theInstance Naming Strategyfield. Content: text string consisting of a combination of characters and numbers with embedded macro variables (see table below). The macros get translated to text strings run time.
	
	The resulting VM name string mustmeet the following requirements for job submission to succeed:
	• Only contain the following characters:
	  • Lowercase a to z
	  • 0 to 9
	  • Hyphen
	• Start withan alpha character
	• Cannot end with a hyphen

Instance IPAM Strategy	Dropdown menu with two choices: No IPAM (default), and IPAM Callouts. If you select the latter, two additional fields are displayed and become required: **IPAM Alloc Rule**and**IPAM Dealloc Rule**. (See below).
Custom VM Name	Location of your custom VM naming script if you set the**Instance Naming Strategy**field to "Custom Callout". Location options inclu "Script from bundle" if you specify a zip file in the **Strategy Bundle** field. See Guidance for Callout Scripts for guidance on developi custom VM naming scripts.
IPAM Alloc Rule	Location of your custom IPAM address allocation script if you set the**Instance IPAM Strategy**field to "IPAM Callouts". Location opti include "Script from bundle" if you specify a zip file in the **Strategy Bundle** field. See Guidance for Callout Scripts for guidance on developing custom IPAM callout scripts.
IPAM Dealloc Rule	Similar to the**IPAM Alloc Rule**above except this is where you specify the location of yourcustom IPAM address deallocation script. Guidance for Callout Scripts for guidance on developing custom IPAM callout scripts.

Instance Naming Strategy Options:

Options	Description
Default	The default method provided by Cisco is the generation of a string consisting of A predefined token prefixed with **cq jw-**
CliQr Macro Replacement	Selecting this method causes the **Node Name Config** field to be displayed and to become a required field. See the fields table above for details.
Hostname Callout	Selecting this method causes the **Custom VM Name** field to be displayed and to become a required field. See the fields table above for details.

OOB macro variables forCliQr Macro Replacement:

Macro	Returned Value
%os1%	First character of OS type string in lower case
%os2%	Firsttwo characters of OS type string in lower case
%OS1%	Firstcharacter of OS type string in upper case
%OS2%	Firsttwo character of OS type string in upper case
% RND<number>%	Fixed length random string, e.g., %RND10%, 10-character long random string ⚠ This is a mandatory field, only values in the range: %RND6% to %RND11% or just %RND% which defaults to %RND6%"
%UID%	User ID
%VID%	Vendor ID
%VM_PREFIX%	VM name prefix must be added as a global parameter in the app profile. ⚠ Only macro whose value can be user-customized.

8. External Lifecycle Actions (conditional):Specify any external lifecycle actions to be performed on all VMs launched by Workload Manager in this region.

External Lifecycle Actions Settings

Whenever you use Workload Manager to deploy a VM-based application, Workload Manager can execute external actions on VMs in the deployment as the VMs progress through these five phases of their lifecycle:

- Pre VM start
- Pre VM initialization
- PostVM initialization
- Pre VM Stop
- Post VM Stop

These lifecycle actions will execute in a script execution engine running in a container in the CloudCenter Suite cluster if no Cloud Remote appliance is configured for the region. If you have configured Cloud Remote for this region, the external lifecycle actions will executein a script execution engine running in a container in the Cloud Remote appliance.

External Lifecycle Action scripts can be defined at the service level, the application level, and the cloud region level. Use the External Lifecycle Actions section of the Region tab or Details tab (see figure below) tospecify external lifecycle actions at the region level.

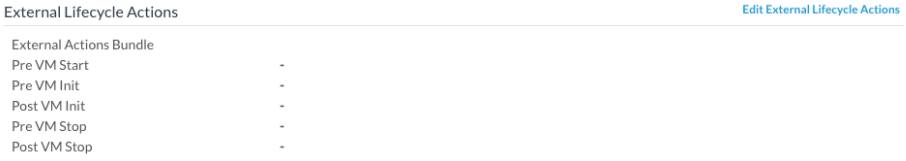

Click on the**Edit External Lifecycle Actions**link in the upper right. This displays theEdit External Lifecycle Actions dialog box as shown below.

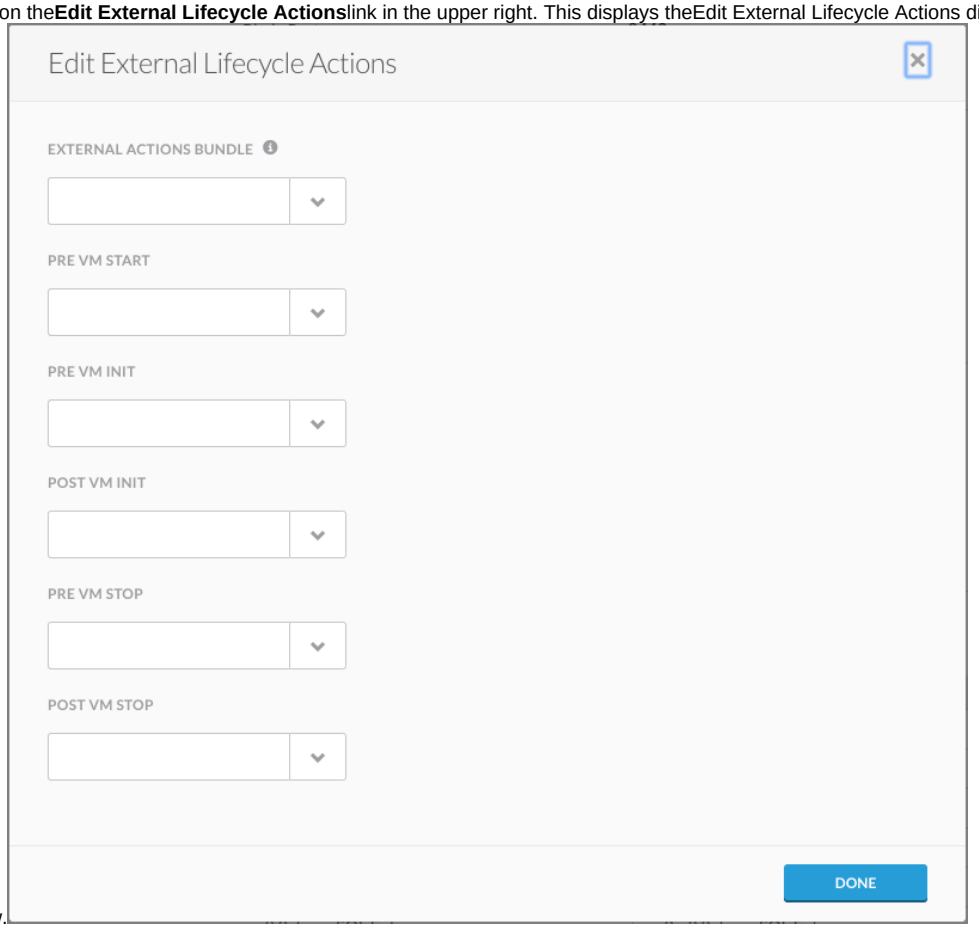

One way to specify an external lifecycle action is to use the first field,**External Actions Bundle**, to specify a zip file that contains all scripts for external lifecycle management. This zip file needs to contain a directory that contains all the scripts. Use the dropdown to select a repository location (URL, FTP, S3, etc.). This causes an additional field to be displayed to the right for entering the path to the zip file as shown in the figurebelow.

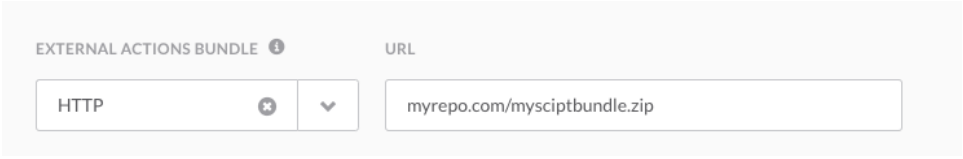

Once you specify actions bundle, for each of the individual lifecycle stage fields, select "Script from Bundle" from the dropdown. This displays an additional field to the right where you must specify the name of the script file in the bundle as shown in the figure below.

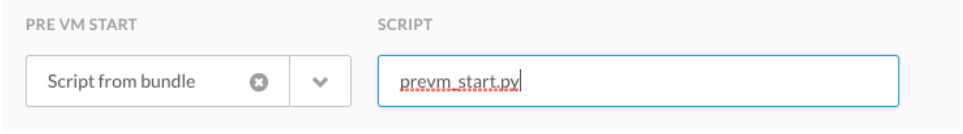

The scripts you specify can be bash scripts or Python scripts called from within a bash script.For more information on lifecycle scripts, seeDeployment Lifecycle Scripts.

9. Instance Types:For OpenStack clouds, you can sync all instance types (flavors) defined in OpenStack to CloudCenter Suite on demand. To manually sync OpenStack instance types, click the**Sync Instance Types**link in the upper right of the instances types section. Alternatively, you can manually add instance types, one by one, by clicking the**Add Instance Types**link in the upper right of the instances types sections.If you add an instance type manually, you must ensure that theinstance ID you enter in CloudCenter Suite exactly matches the corresponding flavor ID in OpenStack. Furthermore, during application deployment, the CPU, RAM, and storage parameters defined in the OpenStack flavor will override any of the corresponding parameters defined in CloudCenter Suite.

Instance Types Settings

An instance type is a performance category of a VM for VM-based clouds or a performance category of a container for container-based clouds. The parameters for an OpenStack cloud instance type is as follows:

- Display Name
- Cost Per Month per GB
- Instance type ID
- CPU
- 32- or 64-bit architecture
- RAM (MB)
- NICs
- Instance storage (GB)
- SSD support

You can sync, edit, and add instance types depending on the cloud provider as follows:

- Manual sync
- Add
- Edit

For public clouds, CloudCenter Suite auto-populates the instance types in the Instance Types section from the cloud provider. CloudCenter Suite auto-syncs this information from the cloud provider once per 24 hours. For these clouds, you cannot force sync on demand. For OpenStack clouds, you can sync all instance types (flavors) defined in OpenStack to CloudCenter Suite on demand. To manually sync OpenStack instance types, click the**Sync Instance Types**link in the upper left for the instances types section as shown in the figure below.

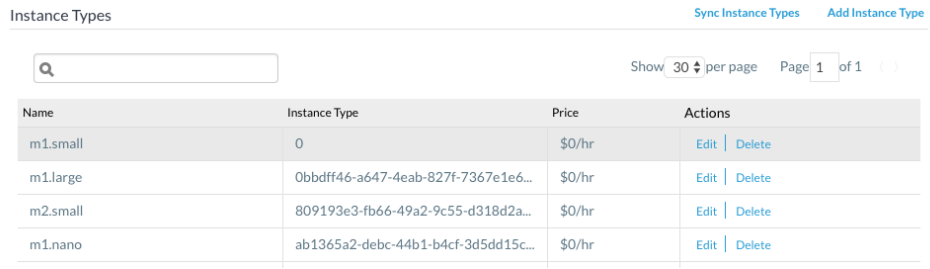

You can add new instance types for OpenStack cloud (by clicking the**Add Instance Types**link in the upper right of the section) and these instance types will control the parameters of the corresponding VMs or containers launched as part of your application. If you add an instance type within CloudCenter Suite for OpenStack, you must make sure the instance type is already defined in OpenStack, and the instance ID you enter in CloudCenter Suite must exactly match the corresponding flavor ID in OpenStack. Furthermore, during application deployment, the CPU, RAM, and storage parameters defined in the OpenStack flavor will override any of the corresponding parameters defined in CloudCenter Suite.

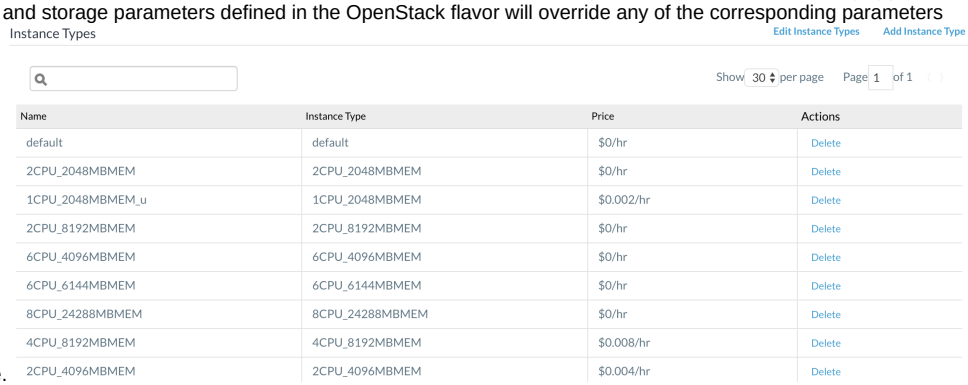

You must specify the price details in the**Price**field for OpenStack cloudwhen adding instance types. The price is used to compute resource costs, which is displayed as invoice costs in theCost Optimizer Dashboard. The**Price**column (highlighted in red in the above screenshot) displays the cost of an instance.

For all cloud providers except AWS, you can edit the fields of an instance type (except instance type ID) by clicking the**Edit**link in the Actions column in the list of instance types, else these parameters are ignored during application deployment for other clouds. The only parameter that will impact CloudCenter Suite behavior for public clouds and OpenStack is the price per hour, which is used by Workload Manager and Cost Optimizer for deployment cost calculations.

Historical Price Management

Sometimes, you want to update the historical cost to correct or update cost computation. In such cases, thecost of the affected resourcesis computed from the date and time the price is updated.You can update the prices of the instances types for the preceding three months only.

To update the price, follow these steps:

a. Click**Edit Instance Type**. The**Edit Instance Types**page appears.
b. In the**Price/Hour**column, click the**Pricing History**icon. The**EditPrice History**dialog appears.

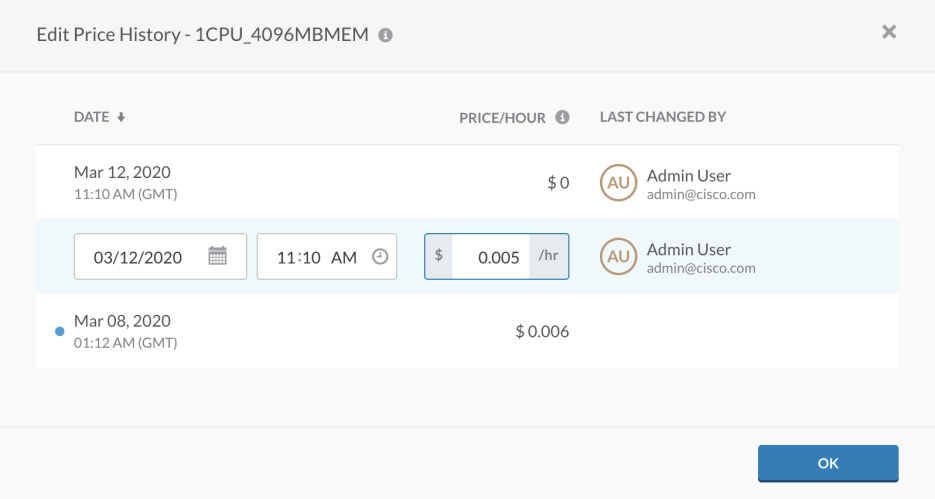

c. Specify a date and time in the appropriate fields and enter the desired price inthe**Price/Hour**field, enter the desired price.
d. Click**OK**, followed by**Done**(in the**Edit Instance Types**page) to save the changes.

After the price has been updated, the cost is recomputed for allresources(VMs and deployments, as applicable)affected by the price change. This could take a few minutes. When the recomputation is successful, you receive a notification that can be viewed in the Notification icon in the Cost Optimizer header (seeCost Optimizer Dashboard), as shown in the following sample screenshot.

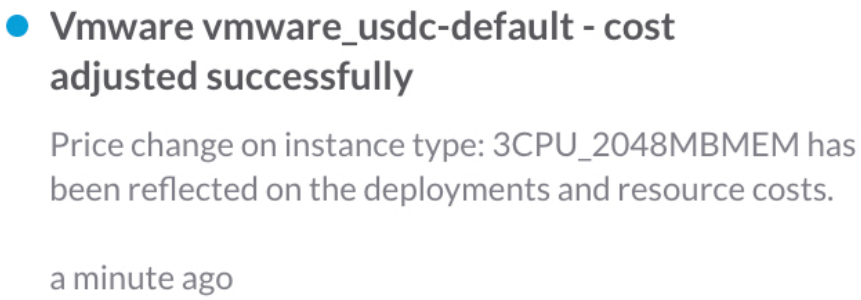

Simultaneously,the updated accruedcostof a VMis displayed on theVirtual Machinespage and the updated deployment cost is displayed on the Deployment Details page.

10. Storage Types (conditional): For private VM-based clouds like OpenStack, CloudCenter Suite uses storage types for cost tracking purposes. CloudCenter Suite creates a default storage type with zero cost.You would manually edit this storage type to enter your own cost factor. You can optionally add more storage types to your OpenStack region.

Storage Types Settings

A storage type is a categorization of storage that can be provisioned for an application during deployment. It is a construct that applies to VM-based clouds only. Storage types are displayed in the Storage Types section of the Regions or Details tab.

The storage type construct contains the following fields for all VM-based clouds:

- Display Name
- Cost Per Month per GB
- Cloud Storage Type ID
- Minimum Volume Size (GB)
- Maximum Volume Size (GB)
- Maximum Throughput (MB/s)
- Maximum IOPS
- Provisioned IOPS
- IOPS Cost Per Month

Storage types are handled by CloudCenter Suite for OpenStack cloud is as follows:

- Add
- Edit

For private VM-based clouds, such as OpenStack, CloudCenter Suite auto-creates a single defaultstoragetype with zero cost.You may manually add additional storage types by clicking the**Add Storage Type**link in the upper right of the Storage Types settings section as shown in the figure below.

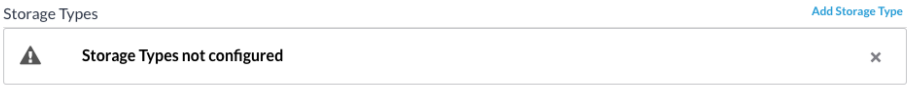

This brings up an**Add Storage Type**dialog box as shown in the figure below.

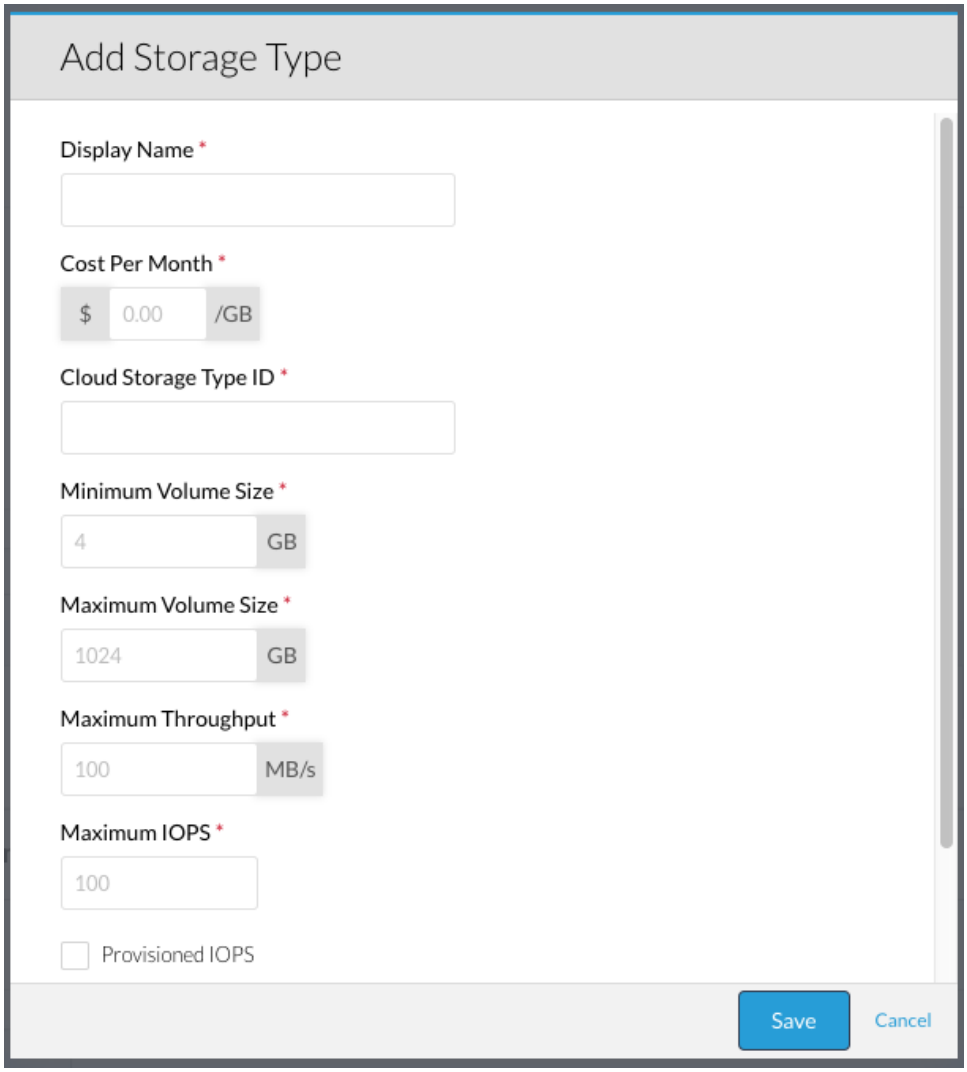

You must enter all required fields, however, the**Maximum Throughput**,**Maximum IOPS**,**Provisioned IOPS**, and**IOPS Cost Per Month**are not recognized by these cloud providers and will be ignored.

You must specify the**Cost**to compute the service costs in Cost Optimizer. The total cost is displayed in theCost Optimizer Dashboard.

Storage types for all cloud providers except AWS can be edited. From the list of storage types in the Storage Type section (see figure below), click the Edit link (in the Actions column) for that storage type.

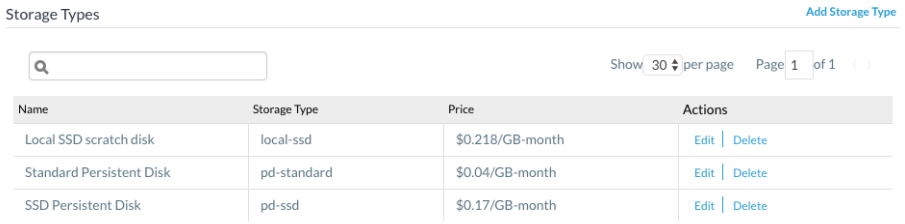

This brings up an Edit Storage Type dialog box (similar to the Add Storage Type dialog box) where you can edit all nine fields except the**Cloud Storage Type ID**field. For AzureRM and Google clouds, only changes to the pricing fields change the behavior of CloudCenter Suite: the cloud cost calculations are modified based on the new unit cost data. All other field changes are ignored.

For all cloud providers,**Minimum Volume Size**and**Maximum Volume Size**restrict the volume size you can specify at deploy time. For public cloud providers that support provisioned IOPS, the user can specify the guaranteed IOPS at deploy time.

Historical Price Management

Sometimes, you want to update the historical cost to correct or update cost computation. In such cases, the cost of the affected resources is computed from the date and time the price is updated.You can update the prices of the older storage types for the preceding three months only.

To update the price, follow these steps:

a. Click**Edit Instance Type**. The**Edit Instance Types**page appears.
b. In the**Price/Hour**column, click the**Pricing History**icon. The**EditPrice History**dialog appears.

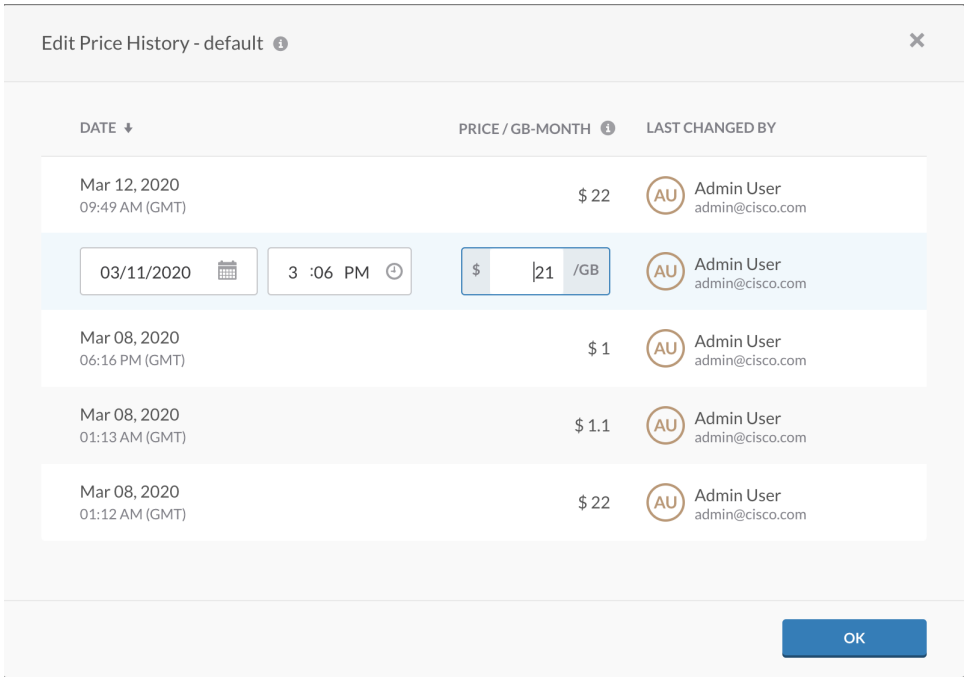

c. Specify a date and time in the appropriate fields and enter the desired price inthe**Price/Hour**field, enter the desired price.
d. Click**OK**, followed by**Done**(in the**Edit Instance Types**page) to save the changes.

After the price has been updated, the cost is recomputed for allresources(VMs and deployments, as applicable)affected by the price change. This could take a few minutes. When the recomputation is successful, you receive a notification that can be viewed in the Notification icon (seeCost Optimizer Dashboard) in the header.Simultaneously,the updated accruedcostof a VMis displayed on theStorage Volumespage and the updated deployment cost is displayed on the Deployment Details page.

11. Image Mappings: Image mappings allow services based on CloudCenter Suite logical images to be deployed using the appropriate physical image stored on the target cloud region. You must manually import these physical images into your OpenStack region and then map the appropriate CloudCenter Suite logical images to these physical images.SeeImagesfor more context.

Add an OpenStack Cloud Account

Prerequisites

Among the two OOB user roles in OpenStack admin and member-member permissions are sufficient to perform all functions in Workload Manager and Cost Optimizer. In addition, more gradual permission can be set in the configuration files of the appropriate OpenStack components per the following table.

OpenStack Module	Minimum permissions needed by Workload Manager	Minimum permissions needed by Cost Optimizer

Compute	`compute:get`
	`compute:get_all`
	`compute:get_all_tenants`
	`compute:get_instance_metadata`
	`compute:get_all_instance_metadata`
	`compute:get_all_instance_system_metadata`
	
	`compute:create`
	`compute:start`
	`compute:stop`
	`compute:reboot`
	`compute:delete`
	`compute:resize`
	`compute:attach_volume`
	`compute:detach_volume`
	
	`compute_extension:keypairs:create`
	`compute_extension:keypairs:delete`
	
	`compute:security_groups:add_to_instance`
	`compute:security_groups:remove_from_instance`
		`compute:get`
		`compute:get_all`
		`compute:get_all_tenants`
		`compute:get_instance_metadata`
		`compute:get_all_instance_metadata`
		`compute:get_all_instance_system_metadata`

Network	`get_network`
	`get_subnet`
	`network:get_all`
		`get_network`
		`get_subnet`
		`network:get_all`

Block Storage	`volume:get`
	`volume:get_all`
	
	`volume:create`
	`volume:delete`
		`volume:get`
		`volume:get_all`

Identity	`identity:list_user_projects`
	`identity:get_user`
	`identity:list_users`
	`identity:list_projects`
		`identity:list_user_projects`
		`identity:get_user`
		`identity:list_users`
		`identity:list_projects`

Image	`get_image`
	`get_images`
	
	`delete_image`
	`download_image`
	`add_image`
	`add_member`
	`delete_member`
		`get_image`
		`get_images`

Configuration Process

To add an OpenStack cloud account, follow this procedure.

1. Locate the OpenStack cloud you created on the Clouds page and click**Add Cloud Account.**This displays the Add Cloud Account dialog box as shown in the figure below.

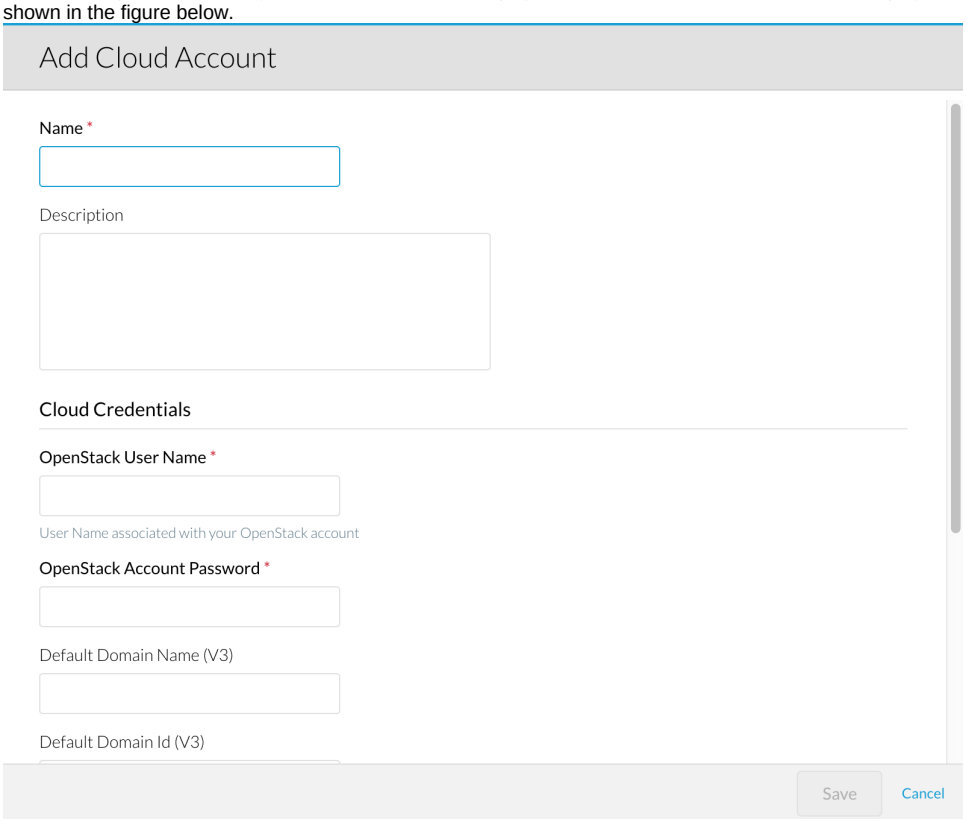

2. Assign a new cloud account**Name**.

> **Tip**
>
> The name should not contain any space, dash, or special characters.

3. Provide the OpenStack user credentials:**OpenStack User Name**and**OpenStack Account Password**.

4. Scroll the **Add Cloud Account** dialog box down to reveal the remaining four input fields as shown in the figure below.

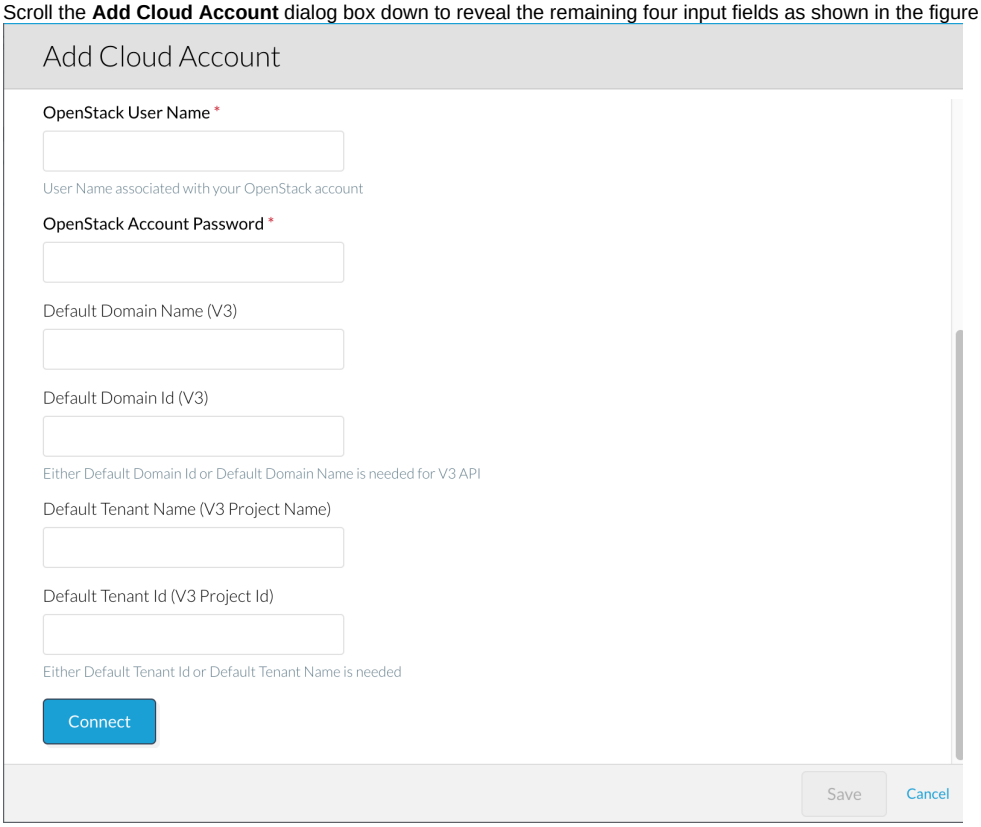

Populate these four optional fields per the table below.

Cloud Account Details	Description
Default Domain Name (V3)	These two fields are optional. When you add an OpenStack cloud account, you can choose V2 or V3 OpenStack endpoints: • Not required if you use V2 • If you use V3, provide either the default Domain ID or Default Domain Name. • The cloud region setting validates the region.
Default Domain ID (V3)	
Default Tenant Name (V3 Project Name)	Optional. The OpenStack project name.
Default Domain ID (V3 Project ID)	Optional. If set, the Default Tenant ID (OpenStack setting in CloudCenter Suite) has precedence over the Default Tenant Name.

5. Click the**Connect**button. CloudCenter Suite will now attempt to validate your account credentials.
6. After the credentials are verified, the**Connect**button changes to an**Edit**button and two new fields appear**Enable Account For**and**Enable Reporting By Org Structure.**

    a. Set the**Enable Account For**dropdown per the table below.

Value	Usage
Provisioning	Workload Manager can deploy jobs using this account.
Reporting	Cost Optimizer and Workload Manager will track cloud costs for this account. Typical usage: master cloud accounts that are used for billing aggregation.
Provisioning, Reporting	Default. Account is used for both provisioning and reporting.

    b. **For AWS and Google clouds only**: Set the**Enable Reporting By Org Structure**toggle to On to cause Cost Optimizer to import the cost hierarchy createdin the cloud provider portal. This saves the time of manually creating a comparable cost hierarchy within Cost Optimizer. SeeCost Groups Configurationfor more information on cost hierarchies in Cost Optimizer.
    c. Click the**Save**button when done.

Cloud Accounts Tab

After you add cloud accounts to a cloud, they will appear in the Accounts tab for the cloud as shown in the figure below.

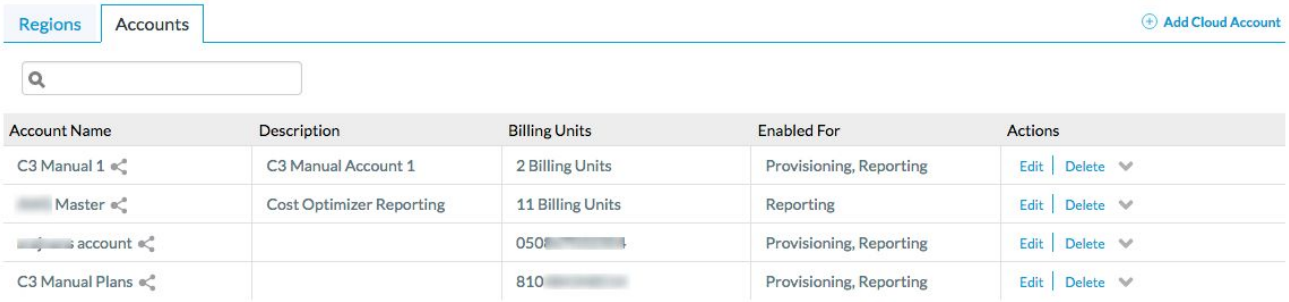

The Accounts tab contains columns for data entered when creating an account: Account Name, Description, Enabled For; and two additional columns:**Billing Units**and**Actions**.**Billing Units**is a dual function:

- If the cloud account contains only one billing unit, the ID for that billing unit is displayed.
- If the cloud account contains multiple billing units, such as an AWS master account, the number of billing units in that account is displayed followed by the text*Billing Units*.

A billing unit is the most granular level of cloud cost recording in CloudCenter Suite. The definition of a billing unit varies by a cloud provider as shown in the table below.

Cloud Provider	Billing Unit
AWS	Account ID
AzureRM	Subscription ID
Google	Project ID
IBM Cloud	Account ID
vCenter	Cloud Group Prefix - Datacenter Name
vCD	Organization Name
OpenStack	Project ID
Kubernetes	Namespace UID

The last column,**Actions**, contains links to let you edit or deleted the cloud account, ormanage instance typesfor the cloud account.

Configure an Outscale Cloud

Configure an Outscale Cloud

Configuring an Outscale cloud is a four-step process:

- Add an Outscale Cloud
- Add an Outscale Region
- Configure an Outscale Region
- Add an Outscale Cloud Account

Add an Outscale Cloud

To add an Outscale cloud follow these steps.

1. Navigate to**Admin > Clouds**. This brings you to the Clouds page. If you, or another tenant admin in your tenant, have already added clouds to your tenant, they will be listed here. Click the**Add Cloud** linkin the upper right.
2. After clicking **Add Cloud**, the Add Cloud dialog box is displayed. Enter the**cloud name**and select the**cloud provider**.
3. After clicking**Next**, the second page of the **Add Clouds** dialog box, **Connectivity Settings**, appears. Set the toggle switches to configure the **Cloud Connectivity** settings.
    - When adding a public VM cloud in the CloudCenter Suite UI, the Cloud Connectivity Settings page, the second page of the Add Cloud dialog box, appears with a single toggle displayed:**Worker VMs Directly Connect with CloudCenter Suite**.
    - Setting this toggle to No implies you will install Cloud Remote for each region of this cloud. This also causes a second toggle to appear:**CloudCenter Suite Directly Accessible from Cloud Remote**.
    - Follow the table below for guidance on setting these toggles.

Toggle settings	Use case	Diagram
Worker VMs Directly Connect with CloudCenter Suite = Yes	Unimpeded connectivity exists between the CloudCenter Suite cluster and the cloud region API endpoint AND Unimpeded connectivity exists between the CloudCenter Suite cluster and worker VMs Cloud Remote is not required	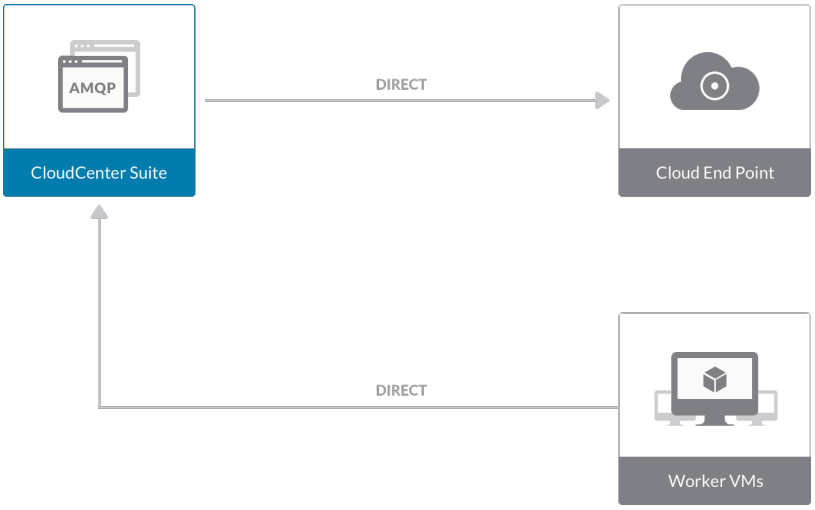

Worker VMs Directly Connect with CloudCenter Suite = No AND CloudCenter Suite Directly Accessible from Cloud Remote = Yes	Worker VMs need to communicate to the CloudCenter Suite cluster through Cloud Remote AND Cloud Remote can initiate the connection to the CloudCenter Suite cluster	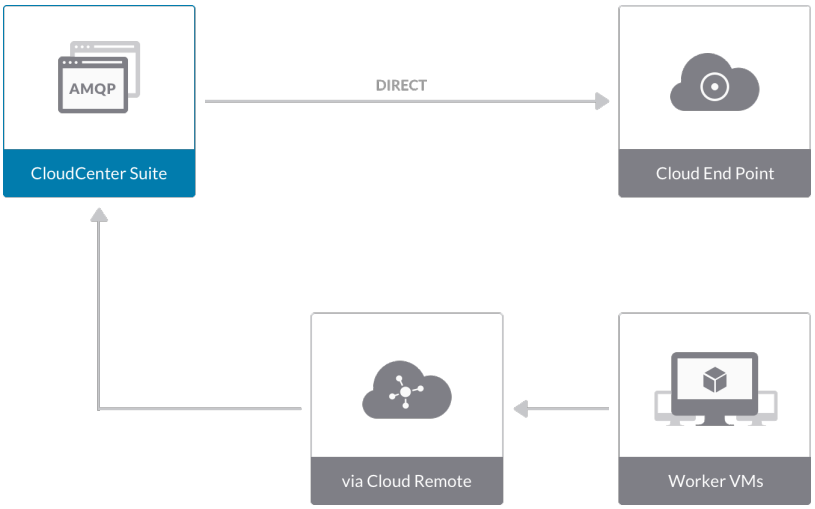
Worker VMs Directly Connect withCloudCenter Suite = No AND CloudCenter SuiteDirectly Accessible from Cloud Remote = No	Worker VMs need to communicate to the CloudCenter Suite cluster through Cloud Remote AND the CloudCenter Suite cluster cannot receive a connection initiated by Cloud Remote	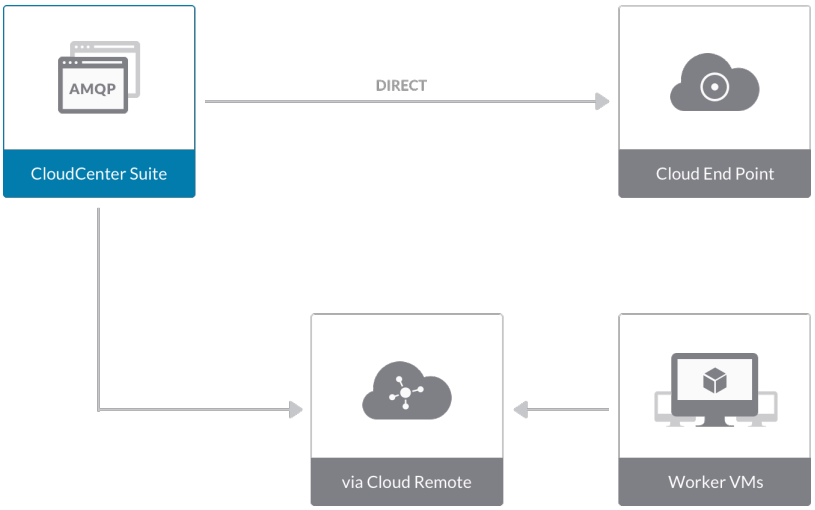

Note

The connectivity toggle settings set at the cloud level are inherited by each region you add to this cloud. However, it is possible to override these toggle settings on a per-region basis from the Regions tab for each region.

4. Click**Done**to save the configuration and close the dialog box. This brings you back to the **Clouds** page, and the cloud you just created will be added to the bottom of the list on the left side of the page.

Add an Outscale Region

After creating an Outscale cloud, the next step is to create the first region for the cloud. Follow these steps.

1. Navigate to the **Clouds** page and select the cloud you created on the left side of the screen. Then click the**Add Region**button on the right side of the screen.
2. After clicking the **Add Region** button, the Add Region dialog box is displayed. Select a region from the list and click **Save.**
3. After clicking**Save**you are brought back to the **Clouds** page with theregion you added shown on the right side of the page.

Configure an Outscale Region

To configure a region you added to your Outscale cloud, follow this procedure:

1. Navigate to Clouds page:**Admin > Clouds**. Find your Outscale cloud from the cloud list on the left half of the screen and click its**Configure Cloud**l ink. This displays the Regions tab for this cloud as shown in the figure below with theCloud Settings section displayed first.

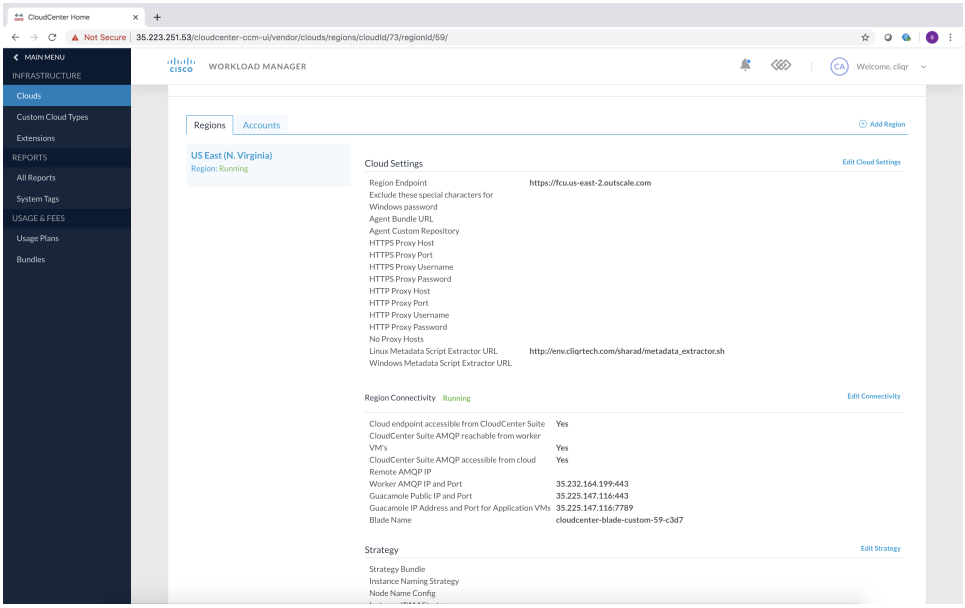

After you have added multiple regions to your Outscale cloud, the Regions tab will show multiple individual region tabs on the left side of the screen. Click the tab of the region you want to configure.

2. Click the**Edit Cloud Settings**link in the upper right of the**Cloud Settings**section. This opens the**Configure Cloud Settings**dialog box.The**Cloud Settings**section contains fields that are unique to Outscale and settings that are common to all cloud providers. Adjust these field values per the instructions in the following tables.

Outscale Specific Cloud Settings

Field	Usage
Region Endpoint	All properties mentioned in the regionMetadataProperties section in the region JSON file of the Outscale metadata package are displayed in this field.

Agnostic Cloud Settings

Field	Usage
Exclude these special characters for Windows password	When the Workload Manager agent is installed on a Windows worker VM, a special user account, called cliqruser, is created to support RDP sessions that may be initiated by the user through the Workload Manager UI. A Workload Manager process running on the CloudCenter Suite cluster creates a random password and passes it to the agent for creating the cliqruser account. Because some Windows deployments may restrict using certain characters for Windows passwords, this field is provided to tell the Workload Manager to exclude these special characters in the generation of the password for the cliqruser account.
Agent Bundle URL	If you plan to use a local repository to host the bundle store, you need to enter the URL of the local bundle store here. Otherwise, leave blank.
Agent Custom Repository	If you plan to use alocal repository to host the package store, you need to enter the URL of the local package store here. Otherwise, leave blank.
HTTP /HTTPS proxy fields (host, username, password)	If you require VMs in your region to access public addresses through a web proxy, enter the URL and credentials of the HTTP and HTTPS proxy servers in these fields.
No Proxy Hosts	If you have specified an HTTP or HTTP proxy using the above fields, you can specify that managed VMs in the region should bypass the proxy and connect directly to certain hosts. Use this field to create a comma-separated list of IP addresses or URLs that should be accessed directly. This field is ignored if an HTTP or HTTPS proxy is not specified.

Important information on proxy settings

In CloudCenter Suite, you can specify proxy settings at the region level, as described here, and at the suite level. To understand the expected behavior when proxy settings are specified at both levels, see the subsequent *Precedence of Proxy Settings* section.

Download Configuration and Encryption Key

After saving the Region Configuration settings in the Workload Manager or Cost Optimizer UI, you can download them to your local computer and then upload them to other conditional components such as Cloud Remote .

The Configuration and Encryption key is only visible when you have configured the Cloud Remote component. Click the**Download Configuration** link in the upper right of the Region Connectivity section, as shown in the following screenshot.

Clicking**DownloadConfiguration**causestwo things to happen:

- An encrypted zip file namedartifacts.zipis downloaded by your browser. Make a note of the location of this zip file as you will need if you are using Cloud Remote .
- The Region Connectivity section header updates to display a**Copy Encryption Key**link, as shown in the following screenshot.

Click the**Copy Encryption Key**link to save the key to your clipboard. A success message will be displayed temporarily above the Region Connectivity section header. Make sure not to overwrite the clipboard with other data. You will need the key when you upload the configuration zip file to conditional components like Cloud Remote .

If you change the connectivity settings in the CloudCenter Suite UI and need to again download the zip file fromsoftware.cisco.com, use the automatically create a (new) encryption key, and copy the key to the clipboard by clicking the**Copy Encryption Key**link again.

Precedence of Proxy Settings

In CloudCenter Suite, you can specify HTTP and HTTPS proxy settings at the suite level as described inProxy Settings, and at the region level as described in the*Agnostic Cloud Settings*section. The CloudCenter Suite cluster, Cloud Remote , and worker VMs will each use either the suite-level proxy settings or the region level proxy settings based on the conditions described below.

For Cloud Remote mode, whenever you change the region proxy settings, ensure you again download the region connectivity setting from the Workload Manager UI and reapply to Cloud Remote .

Mode	Proxy Settings Used By		
Mode (Is Cloud Remote used to communicate withcloud (or APIC) endpoint?)	CloudCenter Suite cluster to communicate with cloud (or APIC) endpoint	Cloud Remote to communicate with cloud endpoint	Worker VM to communicate with bundle store and package store
Cloud Remote Mode	NA	Region-level	Region-level
Non-Cloud Remote Mode	Suite-level	NA	Region-level

Not all clouds support all the proxy settings, and not all clouds support bothCloud Remote and non-Cloud Remote modes.Cloud Remote mode is not applicable to public clouds.
When you are done editing the settings in the dialog box, click**Save.**

3. Determine if you need Cloud Remote for this region.Scroll down to the **Region Connectivity** section for the region and click on the **Edit Connectivity** link (the first time) or the **Configure Region**link (subsequent times) in the upper right to open the **Configure Region** dialog box. The toggle settings should be the same as when you set them on the connectivity page of the **Add Cloud** dialog box.If all of the connectivity togglesinthe **Region Connectivity** dialog boxare set to **Yes**, then Cloud Remote is NOT needed for this cloud region. In this case,you would normally leave the region connectivity settings at their current values and continue to the next settings section. The exceptionto this guidance is when a NAT firewall or proxy server exists between theCloudCenter Suite management cluster andworker VMs, or between the CloudCenter Suite management cluster and users that would use Workload Manager to initiate a Guacamole remote connection to a worker VM. In either of these cases, override the address fields in the **Region Connectivity** dialog box as explained below.

Networking Constraint	Field	Value
Worker VMs must use a proxy server or NAT firewall to access the"local" AMQP server running in the CloudCenter Suite cluster.	Worker AMQP IP Address	IP address and port number that the firewall or proxy server presents to the worker VMs on behalf of the "local" AMQP server running in the CloudCenter Suite cluster.
Users must use a proxy server or NAT firewall to access theGuacamole server running in the CloudCenter Suite cluster.	Guacamole Public IP Address and Port	IP address and port number that the firewall or proxy server presents to users on behalf of theGuacamole server running in the CloudCenter Suite cluster.

Worker VMs must use a proxy server or NAT firewall to access theGuacamole server running in the CloudCenter Suite cluster.	Guacamole IP Address and Port for Application VMs	IP address and port number that the firewall or proxy server presents to the worker VMs on behalf of the Guacamole server running in the CloudCenter Suite cluster.

4. Click**OK**to save the changes and dismiss the dialog box. You can now proceed to the next region settings section:VM Naming and IPAM Strategy. If any of theconnectivity toggles inthe **Region Connectivity** dialog box are set to No, then you must install and configure Cloud Remote for this region.

Configure Cloud Remote in an AWS Region

Configure Cloud Remote in an AWS region as follows.

Obtain and Launch the Cloud Remote Appliance in AWS

a. Obtain the Cloud Remote shared AMI form Cisco support and launch it. Follow the same guidance for obtaining and launching theCloudCenter Suite installer appliance for AWS.
b. Optional but recommended for production environments: Deploy two additional instances of the appliance to form a cluster for HA.Cloud Remote includes support for the clustering of multiple nodes. You will "add" these two additional instances to the first instance after the first instance is configured. SeeCloud Remote (Conditional) > *Scaling*for details.
c. Once the first instance of the appliance has been launched, use your cloud console to**note its IP public and private addresses**. You will need this information later on in order to login to the Cloud Remote web UI and to complete the Region Connectivity settings in the CloudCenter Suite Web UI. Also, note the IP addresses of any other instances you launch.

Unable to render {include} The included page could not be found.

Specify AMQP and Guacamole Addresses for Supporting Cloud Remote

From the CloudCenter Suite UI, for the cloud region requiring Cloud Remote, navigate to the corresponding Regions or Details tab. Click the**Configure Region**link in the upper left of the Region Connectivity section to bring up the Configure Region dialog box. The toggle settings should be the same as when you set them on the connectivity page of the Add Cloud dialog box. You must update some of the address fields in the dialog box according to the scenarios summarized in the table below.

Toggle Settings	Field	Value
Worker VMs Directly Connect with CloudCenter = No AND CloudCenter Directly Accessible from Cloud Remote = Yes	Local AMQP IP Address	Pre-populated with the address and port number of the "local" AMQP server running in the CloudCenter Suite cluster. This address must be **accessible to Cloud Remote** . If Cloud Remote is accessing the CloudCenter Suite cluster through a user-supplied proxy server or NAT firewall, overwrite this field with the corresponding local AMQP IP address and port number provided by the user-supplied proxy server or NAT firewall and accessible to Cloud Remote .
Worker VMs Directly Connect with CloudCenter = No AND CloudCenter Directly Accessible from Cloud Remote = No	Remote AMQP IP Address	Enter <Cloud_Remote_IP>:<amqp_port>, where <Cloud_Remote_IP> = the IP address Cloud Remote which is **accessible to the CloudCenter Suite cluster**, and <amqp_port> = 5671 OR the custom AMQP port number you would later set with the **Change Ports shell script** on the Cloud Remote appliance (see Cloud Remote (Conditional) > *Custom Port Numbers (Conditional)*). If there is no user-supplied NAT firewall or proxy server between the CloudCenter Suite cluster and Cloud Remote , the IP address would be the public IP address of Cloud Remote . If there is a NAT firewall or proxy server between the CloudCenter Suite cluster and Cloud Remote , instead, enter the corresponding public IP address and port number that the firewall or proxy server presents to the internet on behalf of the "remote" AMQP server running in Cloud Remote .
Worker VMs Directly Connect with CloudCenter = No	Worker AMQP IP Address	Enter <Cloud_Remote_IP>:<amqp_port>, where <Cloud_Remote_IP> = the Cloud Remote IP address **accessible to the worker VMs**, and <amqp_port> = 5671 OR the custom AMQP port number you would later set with the **Change Ports shell script** on the Cloud Remote appliance (see Cloud Remote (Conditional) > *Custom Port Numbers (Conditional)*).
Worker VMs Directly Connect with CloudCenter = No	Guacamole Public IP and Port	Enter <Cloud_Remote_IP>:<guac_port>, where <Cloud_Remote_IP> = the Cloud Remote IP address **accessible to CloudCenter Suite users**, and <guac_port> = 8443 OR the custom Guacamole port number you would later set with the **Change Ports shell script** on the Cloud Remote appliance (see Cloud Remote (Conditional) > *Custom Port Numbers (Conditional)*).

Worker VMs Directly Connect with CloudCenter = No	Guacamole IP Address and Port for Application VMs	Enter <Cloud_Remote_IP>:<guac_port>, where <Cloud_Remote_IP> = the Cloud Remote IP address **accessible to worker VMs**, and <guac_port> = 7789

When done, click **OK** to save the setting and dismiss the dialog box.

Download Region Connectivity Settings and Upload to Cloud Remote

Cloud Remote uses the region connectivity settings set in the Workload Manager or Cost Optimizer UI. After saving the Region Configuration settings in the Workload Manager or Cost Optimizer UI, you must download them and to your local computer and then upload them to Cloud Remote as follows.

Click the**Download Configuration**link in the upper right of the Region Connectivity section, as shown in the figure below.

Clicking DownloadConfiguration causestwo things to happen:

- An encrypted zip file named**artifacts.zip**will be downloaded by your browser. Make note of the location of this zip file as you will need to upload it to Cloud Remote through the Cloud Remote web UI (see below).
- The Region Connectivity section header updates to display a **Copy Encryption Key**link, as shown in the figure below.

Click the**Copy Encryption Key**link to save the key to your clipboard. A success message will be displayed temporarily above the Region Connectivity section header. Make sure not to overwrite the clipboard with other data. You will need the key when you upload the configuration zip file to Cloud Remote .

> If you change the connectivity settings in the CloudCenter Suite UI and need to again download the zip file, a new encryption key is automatically created and can be copied to the clipboard by clicking the **Copy Encryption Key** link again.

After you have downloaded the zip file and copied the encryption key to your clipboard, login toCloud Remote web UI.

a. Open another browser tab and login to https://<Cloud Remote _ip> with thedefault credentials: admin/cisco.
b. You will immediately be required to change your password. Do so now.
c. You are now brought to the Cloud Remote home page as shown in the figure below.

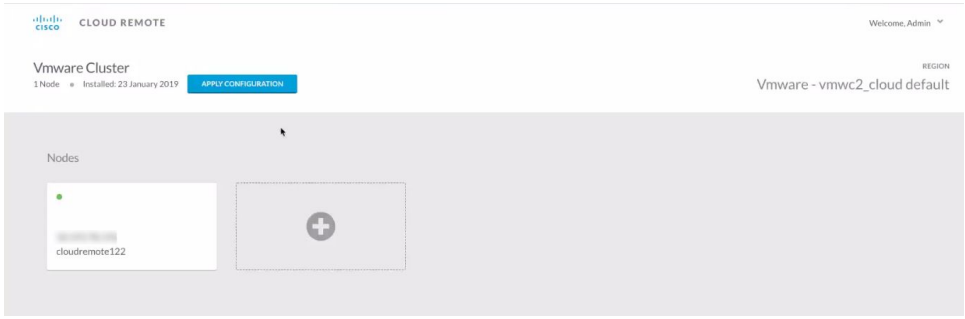

d. Click the **Apply Configuration** button in the page header. This prompts you to select a configuration file and enter the encryption key as shown in the figure below.

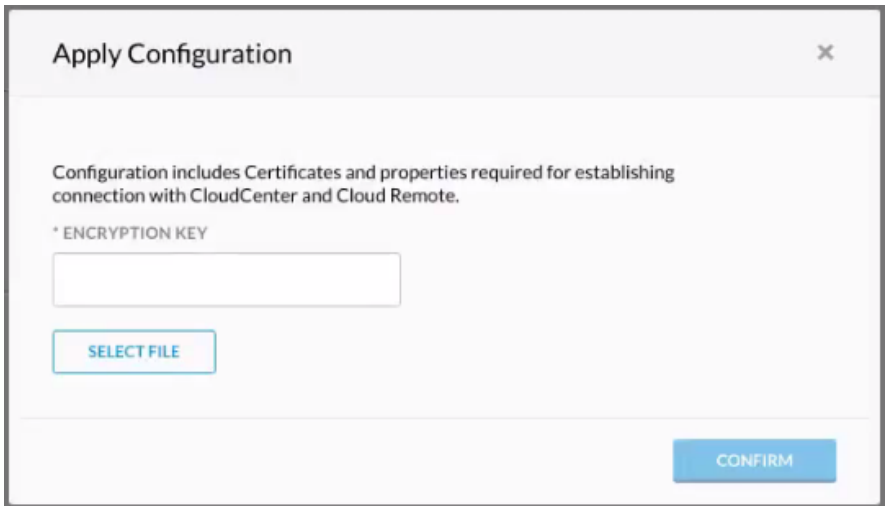

e. Paste the encryption key that was copied to the clipboard into the**Encryption Key**field in the dialog box.
f. Click**Select File**and browse tothe artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
g. Click**Confirm**.
h. Once the zip file is successfully transmitted and accepted, the Cloud Remote appliance attempts to establish communication with the CloudCenter Suite cluster and the Cloud Remote web UI home page is updated to show the name of the region it is connecting to in the upper right (see figure below).

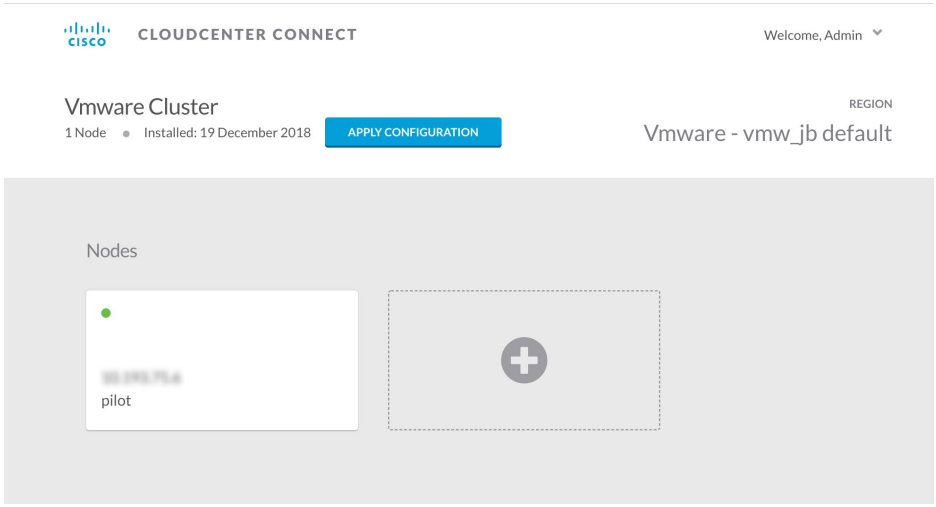

Switch your focus back to the Region Connectivity section of the target cloud region in the CloudCenter Suite web UI. The status indicator in the Region Connectivity section header will change from Not Configured to Running once connectivity between Cloud Remote and the CloudCenter Suite cluster is completely established (see figure below).

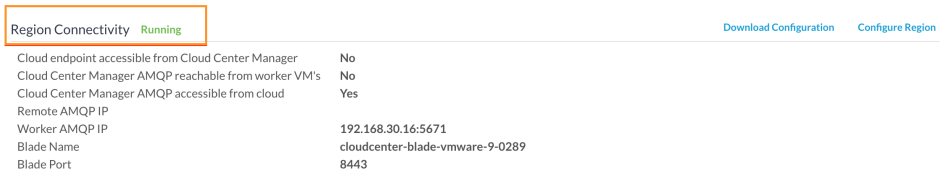

After completing these steps, Workload Manager and Cost Optimizer can use Cloud Remote for communicating with the target cloud region.

5. VM Naming and IPAM Strategy (conditional): Configure any VM naming or IPAM strategies in the Strategy section. If you leave the settings at the defaults, no IPAM strategy is applied and the default VM naming strategy is applied.

VM Naming and IPAM Strategies

Workload Manager lets you define a VM naming strategy and an IP address management (IPAM) strategy on a per cloud region basis. The VM naming strategy lets Workload Manager generate unique names for each of your VMs as they are launched. The IPAM strategy letsWorkl oad Manager generate IP addresses for each of the vNICs on each of your VMs. The VM naming strategy can reference a VM naming callout script, and the IPAM strategy requires that you specify an IP address allocation callout script andan IP address deallocation callout script. An IP address allocation callout scriptmay work in conjunction with a 3rd party IPAM server in your cloud environment.

Like External Lifecycle Actions, the scripts or commands associated with a VM naming strategy or IPAM strategy are executed in a containerized script execution engine running in the CloudCenter Suite cluster, if you are not using Cloud Remote in your region. If you are using Cloud Remote , these scripts or commands will execute in a containerized script execution engine running in Cloud Remote .

To set up the VM naming and IPAM strategies for your region, scroll down to the Strategy section of the Regions or Details tab and follow these steps.

Click the**Edit Strategy**link in the upper right to display the Edit Strategy dialog box as shown in the figure below.

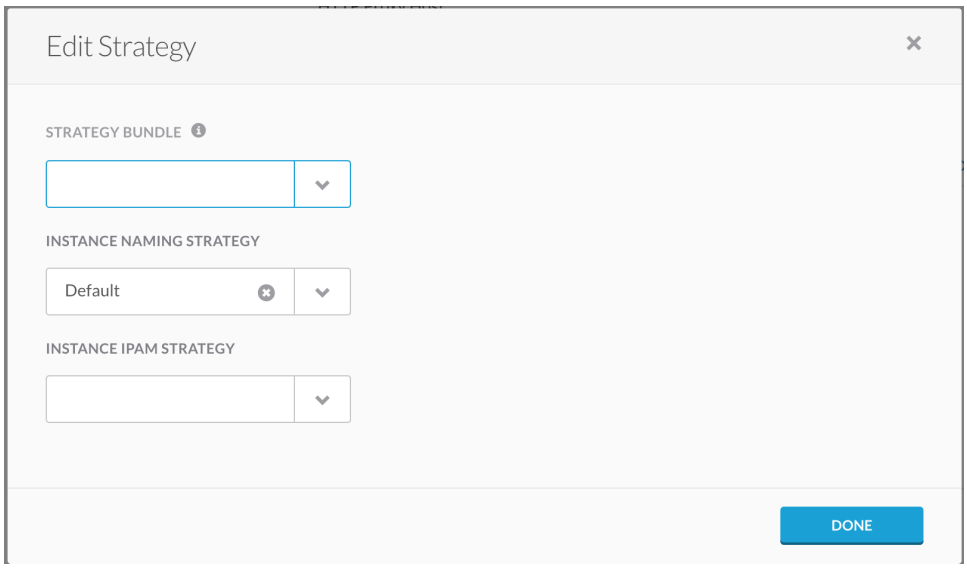

The dialog box can reveal a total of seven data entry fields as explained in the table below:

Field	Usage
Strategy Bundle	Zip file containing all of the scripts associated with your VM naming strategy or IPAM strategy. Select source from the dropdown an add the path in the new field that appears to the right. 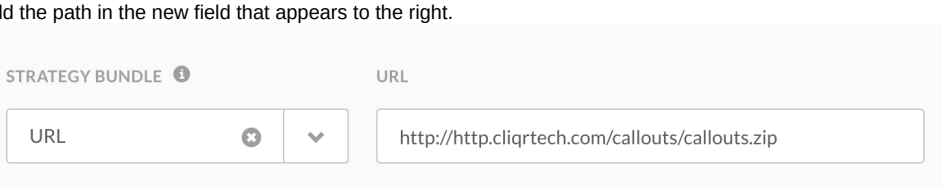
Instance Naming Strategy	Dropdown for selecting one of three naming strategies: Default, CliQr Macro Replacement, and Hostname Callout. If you select CliQ Macro Replacement, the **Node Name Config** field becomes visible and required. If you select Hostname Callout, the **Custom VM Name** field becomes visible and required. See the Instance Naming Strategy table below for details.

Node Name Config	Required if you selected "CliQr Macro Replacement" for theInstance Naming Strategyfield. Content: text string consisting of a combination of characters and numbers with embedded macro variables (see table below). The macros get translated to text string run time.
	
	The resulting VM name string mustmeet the following requirements for job submission to succeed:
	• Only contain the following characters:
	  • Lowercase a to z
	  • 0 to 9
	  • Hyphen
	• Start withan alpha character
	• Cannot end with a hyphen

Instance IPAM Strategy	Dropdown menu with two choices: No IPAM (default), and IPAM Callouts. If you select the latter, two additional fields are displayed and become required: **IPAM Alloc Rule**and**IPAM Dealloc Rule**. (See below).
Custom VM Name	Location of your custom VM naming script if you set the**Instance Naming Strategy**field to "Custom Callout". Location options inclu "Script from bundle" if you specify a zip file in the **Strategy Bundle** field. See Guidance for Callout Scripts for guidance on developi custom VM naming scripts.
IPAM Alloc Rule	Location of your custom IPAM address allocation script if you set the**Instance IPAM Strategy**field to "IPAM Callouts". Location opti include "Script from bundle" if you specify a zip file in the **Strategy Bundle** field. See Guidance for Callout Scripts for guidance on developing custom IPAM callout scripts.
IPAM Dealloc Rule	Similar to the**IPAM Alloc Rule**above except this is where you specify the location of yourcustom IPAM address deallocation script. Guidance for Callout Scripts for guidance on developing custom IPAM callout scripts.

Instance Naming Strategy Options:

Options	Description
Default	The default method provided by Cisco is the generation of a string consisting of A predefined token prefixed with **cq jw-**
CliQr Macro Replacement	Selecting this method causes the **Node Name Config** field to be displayed and to become a required field. See the fields table above for details.
Hostname Callout	Selecting this method causes the **Custom VM Name** field to be displayed and to become a required field. See the fields table above for details.

OOB macro variables forCliQr Macro Replacement:

Macro	Returned Value
%os1%	First character of OS type string in lower case
%os2%	Firsttwo characters of OS type string in lower case
%OS1%	Firstcharacter of OS type string in upper case
%OS2%	Firsttwo character of OS type string in upper case
%RND<number>%	Fixed length random string, e.g., %RND10%, 10-character long random string ⚠ This is a mandatory field, only values in the range: %RND6% to %RND11% or just %RND% which defaults to %RND6%"
%UID%	User ID
%VID%	Vendor ID
%VM_PREFIX%	VM name prefix must be added as a global parameter in the app profile. ⚠ Only macro whose value can be user-customized.

6. External Lifecycle Actions (conditional):Specify any external lifecycle actions to be performed on all VMs launched by Workload Manager in this region.

External Lifecycle Actions Settings

Whenever you use Workload Manager to deploy a VM-based application, Workload Manager can execute external actions on VMs in the deployment as the VMs progress through these five phases of their lifecycle:

- Pre VM start
- Pre VM initialization
- PostVM initialization
- Pre VM Stop
- Post VM Stop

These lifecycle actions will execute in a script execution engine running in a container in the CloudCenter Suite cluster if no Cloud Remote appliance is configured for the region. If you have configured Cloud Remote for this region, the external lifecycle actions will executein a script execution engine running in a container in the Cloud Remote appliance.

External Lifecycle Action scripts can be defined at the service level, the application level, and the cloud region level. Use the External Lifecycle Actions section of the Region tab or Details tab (see figure below) tospecify external lifecycle actions at the region level.

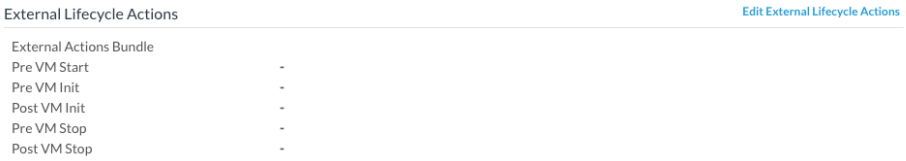

Click on the**Edit External Lifecycle Actions**link in the upper right. This displays theEdit External Lifecycle Actions dialog box as shown below.

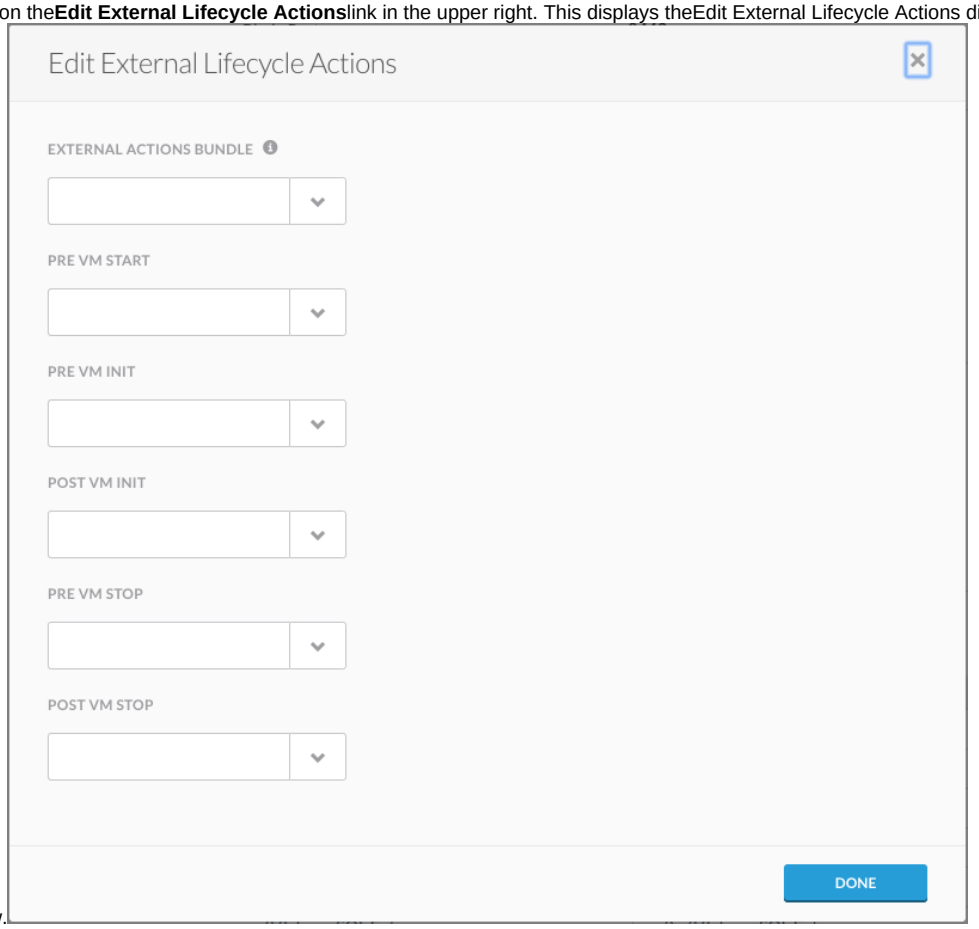

One way to specify an external lifecycle action is to use the first field,**External Actions Bundle**, to specify a zip file that contains all scripts for external lifecycle management. This zip file needs to contain a directory that contains all the scripts. Use the dropdown to select a repository location (URL, FTP, S3, etc.). This causes an additional field to be displayed to the right for entering the path to the zip file as shown in the figurebelow.

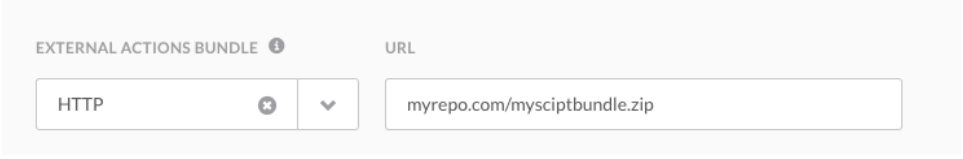

Once you specify actions bundle, for each of the individual lifecycle stage fields, select "Script from Bundle" from the dropdown. This displays an additional field to the right where you must specify the name of the script file in the bundle as shown in the figure below.

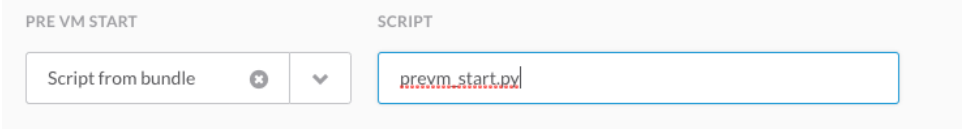

The scripts you specify can be bash scripts or Python scripts called from within a bash script.For more information on lifecycle scripts, seeDeployment Lifecycle Scripts.

7. Instance Types (informational): CloudCenter Suite automatically synchronizes instance types for public cloud regions on a daily basis. This data includes published pricing for each instance type. It is not possible to edit Outscale region instance types.

Instance Types Settings

An instance type is a performance category of a VM for VM-based clouds or a performance category of a container for container-based clouds. You can sync, edit, and add instance types depending on the cloud provider. For public clouds, CloudCenter Suite auto-populates the instance types in the Instance Types section from the cloud provider. CloudCenter Suite auto-syncs this information from the cloud provider once per 24 hours. For these clouds, you cannot force sync on demand.

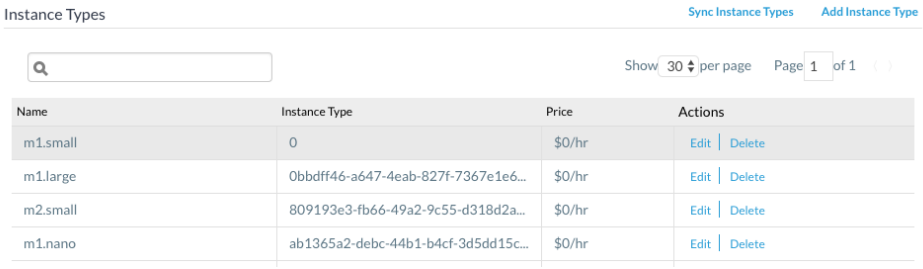

For all cloud providers except AWS, you can edit the fields of an instance type (except instance type ID) by clicking the**Edit**link in the Actions column in the list of instance types. The edits to these parameters are ignored during application deployment. The only parameter that will impact CloudCenter Suite behavior for public clouds is the price per hour, which is used by Workload Manager and Cost Optimizer for deployment cost calculations.

8. Storage Types (conditional): CloudCenter Suite automatically synchronizes storage types for public cloud regions on a daily basis. This data includes the cloud provider published pricing for each storage type. It is not possible to edit Outscale region storage types.

Storage Types Settings

A storage type is a categorization of storage that can be provisioned for an application during deployment. It is a construct that applies to VM-based clouds only. Storage types are displayed in the Storage Types section of the Regions or Details tab. The storage type construct contains the following fields for all VM-based clouds:

- Display Name
- Cost Per Month per GB
- Cloud Storage Type ID
- Minimum Volume Size (GB)
- Maximum Volume Size (GB)
- Maximum Throughput (MB/s)
- Maximum IOPS
- Provisioned IOPS
- IOPS Cost Per Month

Storage types are handled by CloudCenter Suite for various cloud providers in a fashion similar to instance types. CloudCenter Suite auto-populates the list of storage types in the Storage Types section based on information from the cloud provider. This data is stored in thepackage store and the package store auto-syncs this information from the cloud provider once per 24 hours.

Storage types for all cloud providers except AWS can be edited. From the list of storage types in the Storage Type section (see figure below), click the Edit link (in the Actions column) for that storage type.

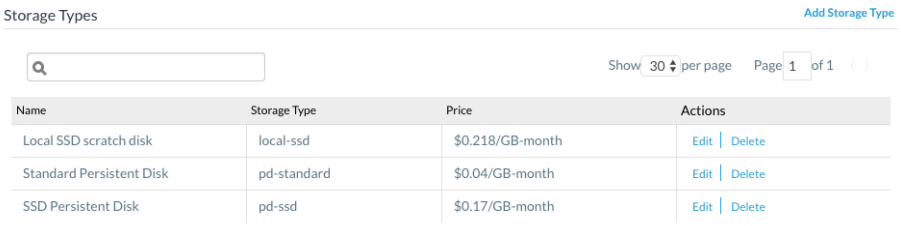

This brings up an Edit Storage Type dialog box (similar to the Add Storage Type dialog box) where you can edit all nine fields except the**Cloud Storage Type ID**field. For AzureRM and Google clouds, only changes to the pricing fields change the behavior of CloudCenter Suite: the cloud cost calculations are modified based on the new unit cost data. All other field changes are ignored.

For all cloud providers,**Minimum Volume Size**and**Maximum Volume Size**restrict the volume size you can specify at deploy time. For public cloud providers that support provisioned IOPS, the user can specify the guaranteed IOPS at deploy time.

9. Image Mappings: Image mappings allow services based on CloudCenter Suite logical images to be deployed using the appropriate physical image stored on the target cloud region. CloudCenter Suite automatically maps the OOB logical images to public cloud region physical images when you add the region to your cloud. Cisco periodically updates these mappings when new versions of OS physical images are uploaded by the cloud provider. To apply these updates to your region after it is added to your cloud, click the **Sync Image Mappings** link in the upper right of this section. If you create any custom logical images, you mustmanually import the corresponding physical images into your region and then map the corresponding logical images to these physical images. SeeImages for more context.

Add an Outscale Cloud Account

Prerequisites

Before adding an Outscale cloud account, do the following:

- Ensure the account has the minimum permissions.SeeCloud Overview> *Minimum Permissions for Public Clouds*for additional details.

Configuration Process

To add an Outscale cloud account, follow this procedure.

1. Locate your Outscale cloud on the Clouds page and click theAdd Cloud Accountlink for this cloud.This displays the Add Cloud Account dialog box, as shown below.
2. Assign a cloud account**Name**.

> **Tip**
>
> The name should not contain any space, dash, or special characters.

3. Provide the Outscale cloud credentials the credentials are the same as the properties mentioned under the cloudAccountMetadataProperties section in cloud.json file of Outscale metadata package:
    a. **Outscale Account Number**: The account number from your Outscale account.
    b. **Outscale Access Key and Secret Key**: The security credentials to access this Outscale account.
4. Click the **Connect** button. CloudCenter Suite will now attempt to validate your account credentials.
5. After the credentials are verified, the**Connect**button changes to an**Edit**button and two new fields appear, namely,**Enable Account For** and **Enable Reporting By Org Structure**,

Set the **Enable Account For** dropdown per the table below.

Value	Usage
Provisioning	Workload Manager can deploy jobs using this account.
Reporting	Cost Optimizer and Workload Manager will track cloud costs for this account. Typical usage: master cloud accounts that are used for billing aggregation. ⚠ It is recommended that you do not add a *Reporting* account to the same tenant through different cloud groups. ⚠ Enabling a public cloud account for *Reporting* may incur expenses to retrieve cost data. These expenses are proportional to the number of configured cloud accounts and regions.
Provisioning, Reporting	Default. Account is used for both provisioning and reporting.

Click the**Save**button when done.

Cloud Accounts Tab

After you add cloud accounts to a cloud, they will appear in the Accounts tab for the cloud as shown in the figure below.

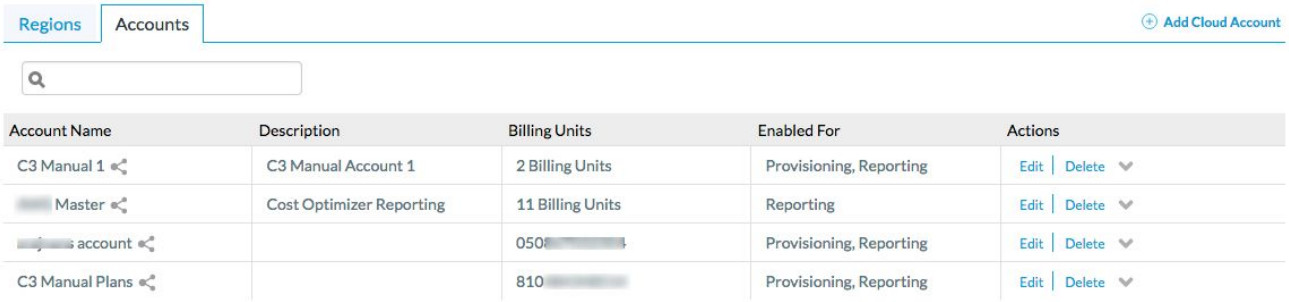

The Accounts tab contains columns for data entered when creating an account: Account Name, Description, Enabled For; and two additional columns:**Billing Units**and**Actions**.**Billing Units**is a dual function:

- If the cloud account contains only one billing unit, the ID for that billing unit is displayed.
- If the cloud account contains multiple billing units, such as an AWS master account, the number of billing units in that account is displayed followed by the text*Billing Units*.

A billing unit is the most granular level of cloud cost recording in CloudCenter Suite. The definition of a billing unit varies by a cloud provider as shown in the table below.

Cloud Provider	Billing Unit
AWS	Account ID
AzureRM	Subscription ID
Google	Project ID
IBM Cloud	Account ID
vCenter	Cloud Group Prefix - Datacenter Name
vCD	Organization Name
OpenStack	Project ID
Kubernetes	Namespace UID

The last column,**Actions**, contains links to let you edit or deleted the cloud account, ormanage instance typesfor the cloud account.

Cloud Remote

Unable to render {include} The included page could not be found.

Cloud Maintenance

Cloud Maintenance

- Sharing Cloud Regions and Cloud Accounts
- Deleting Clouds, Cloud Regions, and Cloud Accounts

Sharing Cloud Regions and Cloud Accounts

Clouds, cloud regions, and cloud accounts that are created within a tenant are automatically co-owned by all tenant admins. In Workload Manager, standard users do not have direct access to these elements for deploying workloads. Instead, users deploy workloads through an intermediary construct: the deployment environment. However, it is possible to directly share specific cloud regions and cloud accounts with subtenants as explained inTenant Management > *Manage Clouds*. Once a cloud region or cloud account is shared with a subtenant, admin users in that subtenant can use those regions and accounts for creating their own deployment environments. However, the admins in the subtenant cannot edit or delete those shared accounts or regions.

Deleting Clouds, Cloud Regions, and Cloud Accounts

Deleting clouds, cloud regions, and cloud accounts must be done in a certain sequence. Before you can delete a multi-region cloud, you must first delete all regions for that cloud. After you delete all regions, the delete icon appears for the cloud in the Clouds page. Before you can delete a region, you must first delete all cloud accounts associated with that cloud. If you attempt to delete a region when any cloud accounts are assigned to the cloud, you will get an error message as follows:

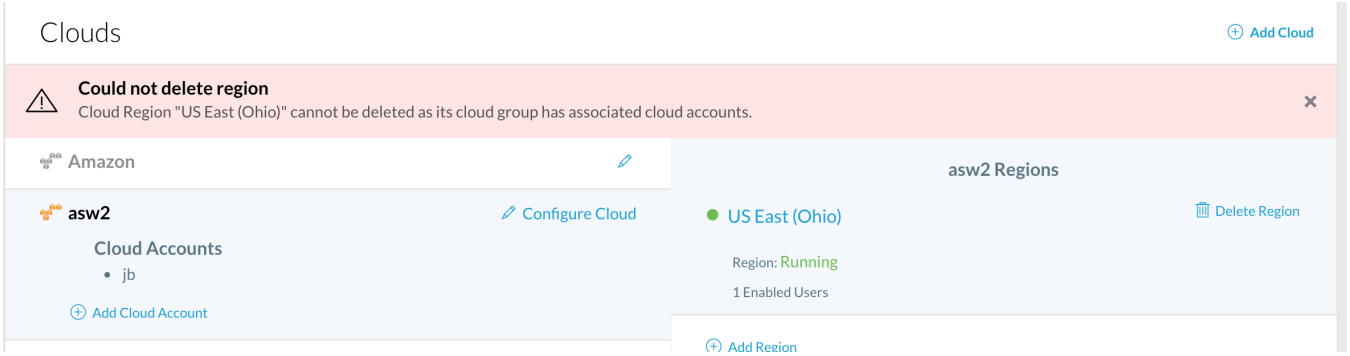

Similarly, before you can delete a single region cloudyou must first delete all cloud accounts associated with that cloud. Otherwise, you will see an error message as shown below:

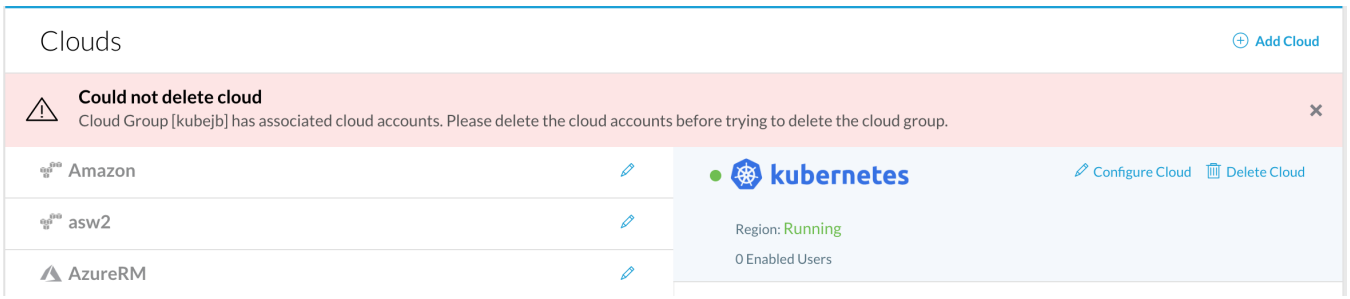

Before you can delete a cloud account you must first remove that cloud account from all deployment environmentsin which it is used.Otherwise, an error message as shown below is displayed:

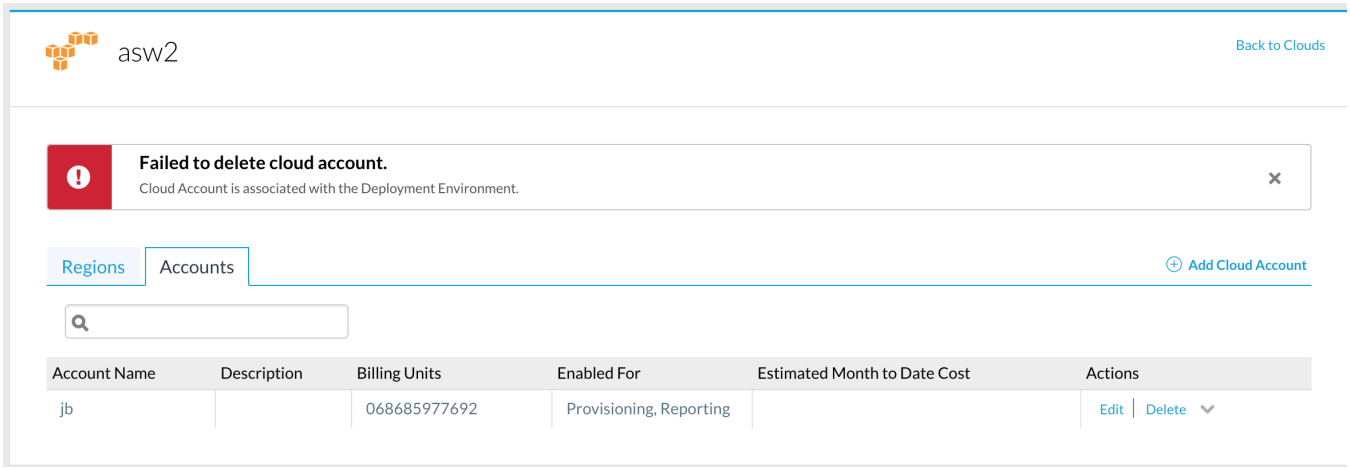

Therefore, to delete a cloud follow these steps:

1. From the Clouds page, select the cloud and click its **Configure Cloud** link which displays the page for this cloud. The page for this cloud will be displayed as shown below.

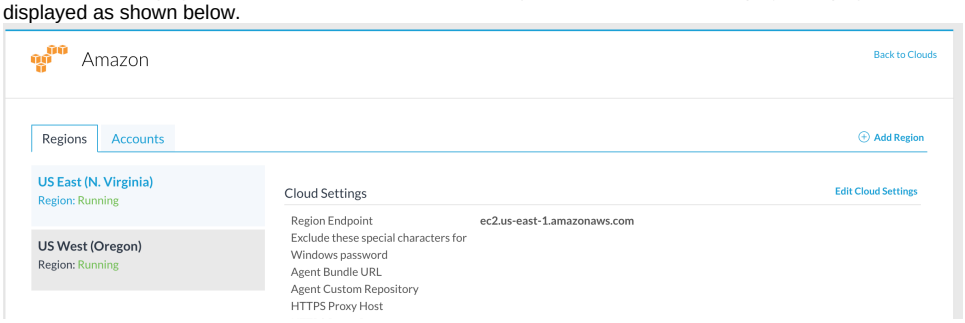

2. From the page for this cloud, select the **Accounts** tab. The Accounts tab is displayed as shown below.

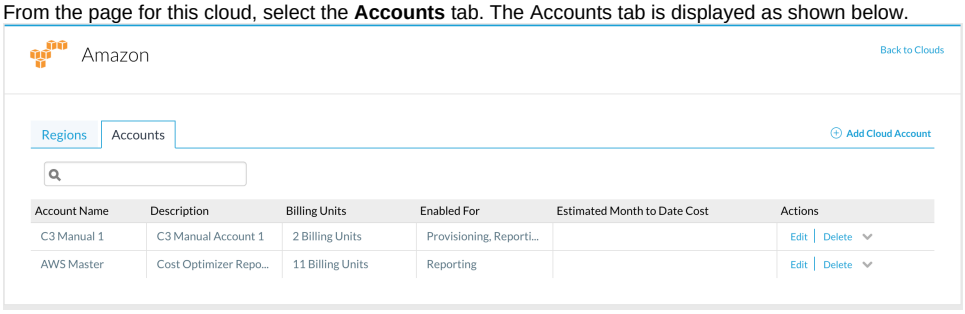

3. From the Accounts tab, delete all accounts one by one by clicking the **Delete** link in the Actions column. If an error about deployment environments appears, click on the **Main Menu >Environments** menu tab, browse the deployment environments for any references to the account, and remove the account from those deployment environments. When done, return to the Clouds page.
4. From the Clouds page, If the cloud is a single region cloud, the **DeleteCloud** link for that cloud will appear on the left side of the Clouds page, as shown in the figure below. Click the **DeleteCloud**link. You are done.

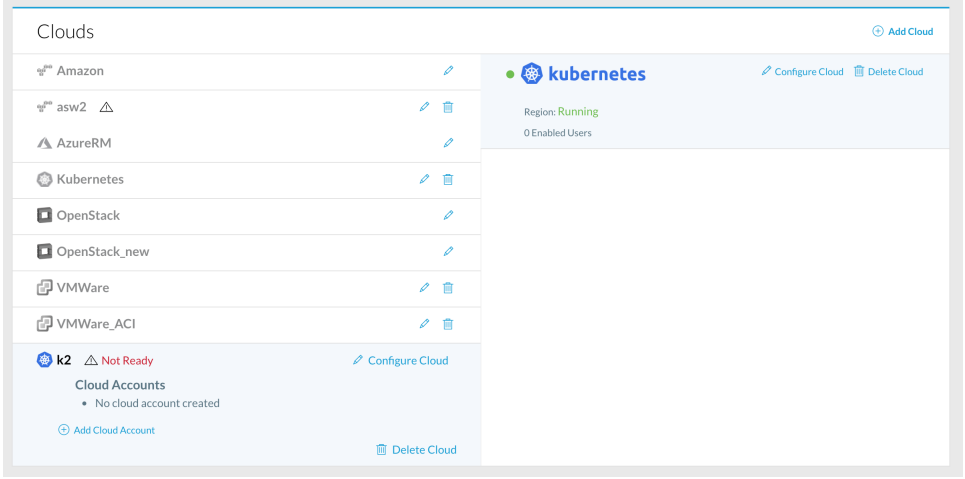

5. If the cloud is a multi-region cloud, on the left side of the Clouds page, select the cloud. This causes the regions for this cloud to be displayed on the right side of the Clouds page as shown in the figure below. For each cloud region, click on its **Delete Region** link.

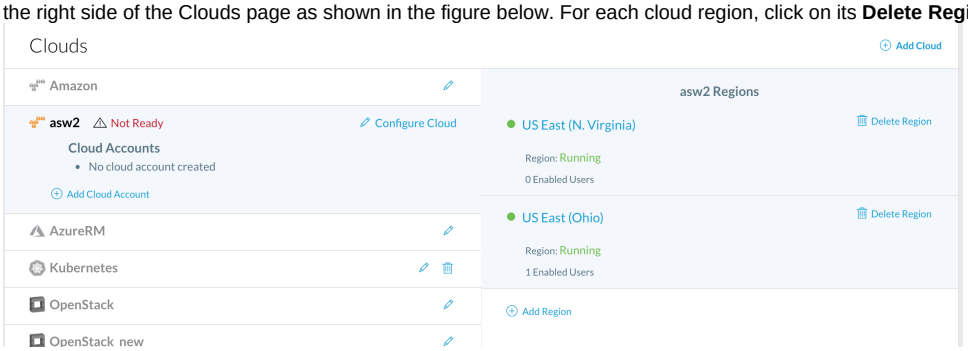

6. After you delete all cloud regions associated with a cloud, the Clouds page will appear as shown below. Click the **Delete Cloud** link for the cloud on the left side of the page. You are done.

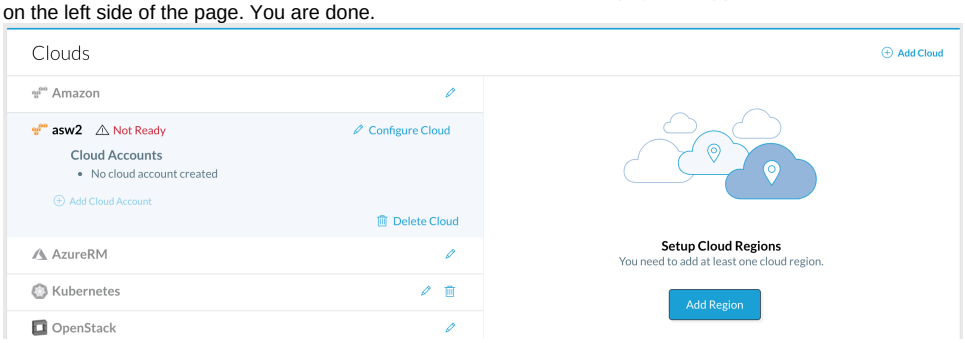

Cost Groups Configuration

Cost Groups Configuration

- Cost Groups UI
- How Do I...

Cost Groups UI

Cost Groups UI

- Overview
- Terminology
- What's in the Cost Groups UI?
- Cost Group Type
- Cost Group
- Sharing Cost Groups
- Billing Units
- Tags

Overview

After you have configured clouds for Cost Optimizer, you may set up cost groups to classify the data. Data classification helps you to distinguish and identify the data. Use the**Cost Groups**page to classify data and define your hierarchy.

Terminology

Throughout this document, you will refer to the following terms:

Term	Description
Cost Group Type	Maps to the various functions in an organization, for example, Development, HR, IT, and so on.
Cost Groups	Hierarchical structure to define your organization and distribute billing units.
Cloud Account	Credentials for logging in to a cloud provider.
Billing Units	Different entities depending on the cloud. These entities are account IDs in Amazon cloud, Project IDs in Google cloud, Subscription ID in AzureRMcloud, Datacenter name (prefixed with the cloud group) in vCenter clouds, Project ID in OpenStack cloud, and Namespace UID in Kubernetes cloud.
Budgets	Ability to allocate or reserve amounts per cloud or cost group type.
Tags	Key-value pairs associated with resources in a cloud.

What's in the Cost Groups UI?

The**Cost Groups**UI, as shown in the following sample screenshot, has the following:**Cost Group Type (Department(s))**and**Unassigned Billing Units** and **Unassigned Tags**. SeeUI Behaviorfor details onicons.

The following table explains the icons in the UI (in alphabetical order).

Icon	Description
Action	Perform action-oriented tasks Add, Delete, Rename, and Share on the Cost Group.
Add Cost Group	Add a Cost Group to a Cost Group Type.
Cost Group Type	Lists Cost Group Types set up in Cost Optimizer and add Cost Group Type.
Filter	Allows you filter for the billing units based on the specified category.
Move Billing Units	Move multiple billing units to a cost group. This icon is enabled after Billing Units are selected.
Search	Search resources based on the specified text.
Select All	Select all items displayed on the page by clicking the checkbox in the table header or by clicking the checkbox against each item.
Unassigned Billing Units	Cloud accounts that have not yet been assigned to a cost group.
Unassigned Tags	Tags associated with cloud resources that have not yet been assigned to a cost group.

Cost Group Type

A Cost Group Type is equivalent, but not restricted, to the various functions in an organization. For example, an organization might have different functions, such as Development, Finance, IT, Sales, Support, etc.Cost Optimizer ships with a seeded Cost Group Type called**Department**associated with the root tenant.

Cost Group

A Cost Group is a hierarchical structure that you define for your organization. You can have a flat or vertical cost group, depending on your need. In case of a vertical cost group, there can be as many levels (departments and subdepartments) as you desire.For instance, Development would have sub-functions such as automation, core development, testing,release team, etc. The hierarchical structure can also beimportedvia a .csv file.

Sharing Cost Groups

Sharing is the provision of an entity or servicepresent in more than one function in an organization.Any user can share an entity by using the**Share**option. This option works on the principle of ACL functionality where auser assigned to acost group can share the cost groups with other users or user groups.

You can configure sharing at the Users or Groups levels as determined by your access permission as shown in the following screenshot.

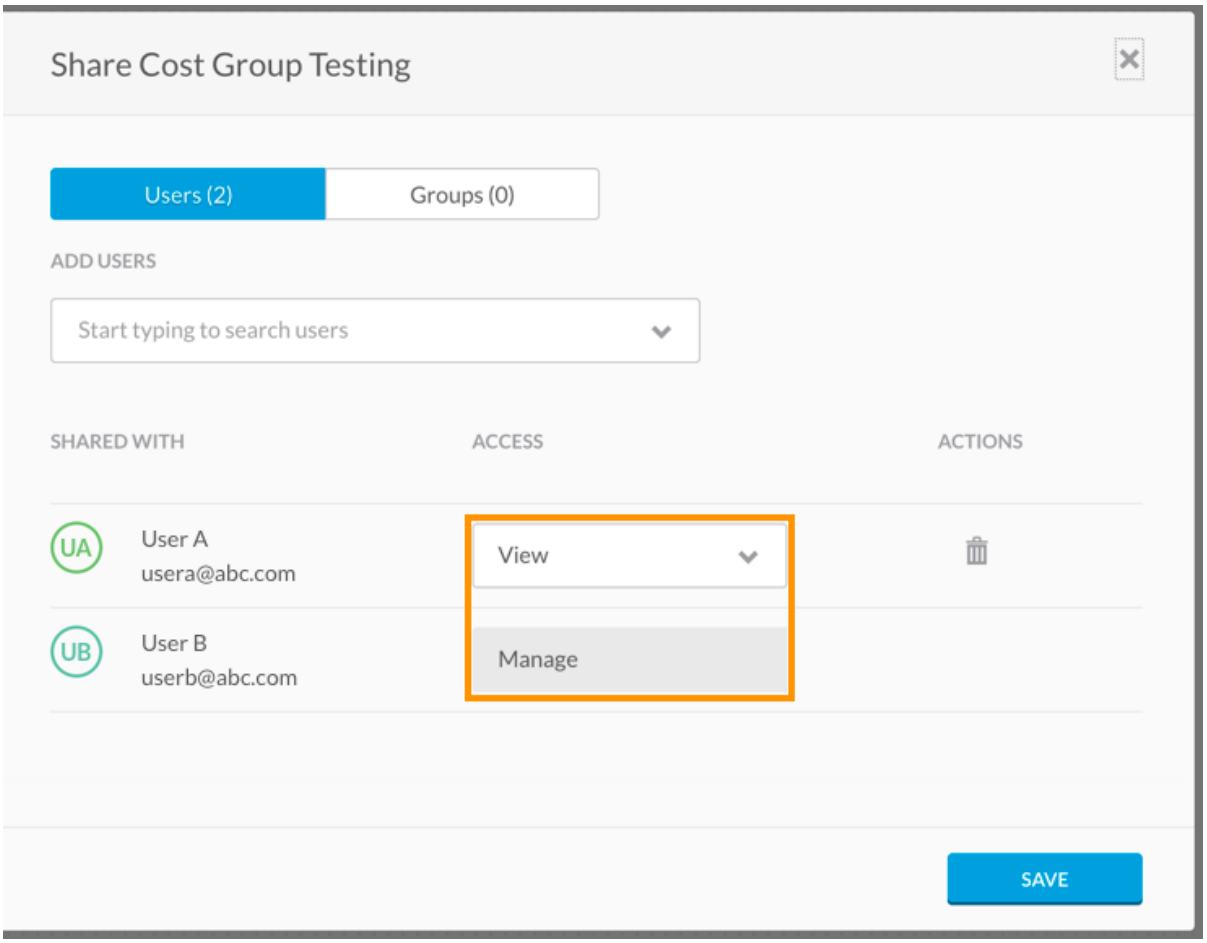

The following table identifies the access levels for Cost Optimizer.

Tab	Controls
Users	To assign specific permissions to individual users, add the users to this resource, then set permission options for each user.
Group	To assign permissions to a user group, add the user group to this resource, then set permission options.

SeeAccess and Roles> Access Control Lists (ACLs).

> When tag-based cost reports are shared, the sharing results in displaying additional cost, inventory, and recommendations for the resources associated with the cost groups.

Billing Units

Billing units are used for a cost breakdown.When validating a cloud account, billing units are automatically discovered and associated with the cloud accounts.

Tags

Tags are key-value pairs associated with cloud resources on a cloud provider. The key is mandatory and value is optional. Tags can be user-defined or system-defined.Similar to billing units, tags are also used for cost breakdown at a deeper granular level. The tags are discovered through the tag collection background process. SeeData Collection. This feature is available on AWS and Azure clouds only.

> You can choose either a billing unit associated cost group or a tags associated cost group, not both. Once created with an association, you cannot change it later, after creating the cost group.

How Do I...

How Do I...

- Add a Cost Group Type
- Add Cost Group
- Associate a Billing Unit
- Associate a Tag

Add a Cost Group Type

Perform these steps to add a cost group type:

1. Click the arrow next to the cost group type (**Department**).

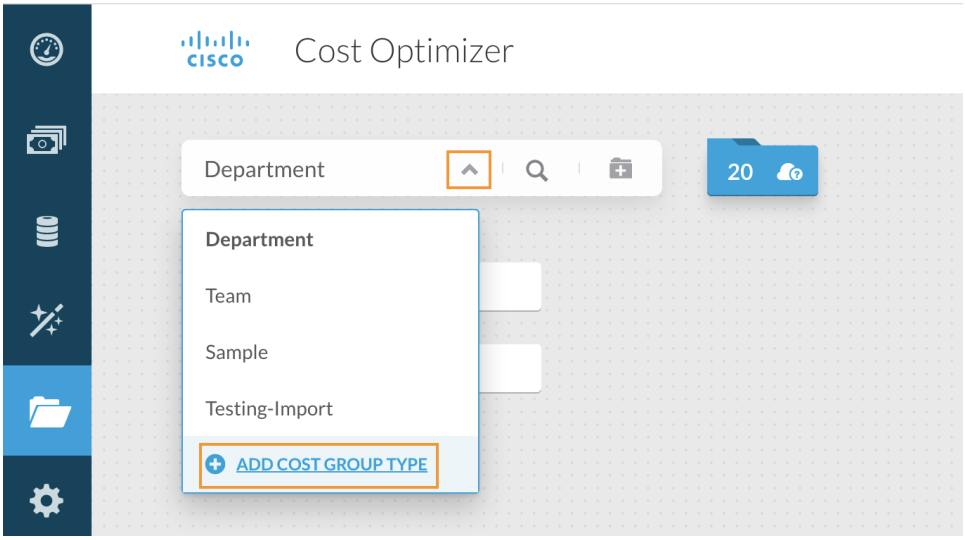

2. Choose**Add Cost Group Type**.
3. Specify a name for the cost group type, in singular and plural format in the respective fields.
4. Choose the cost report base billing units or tags.

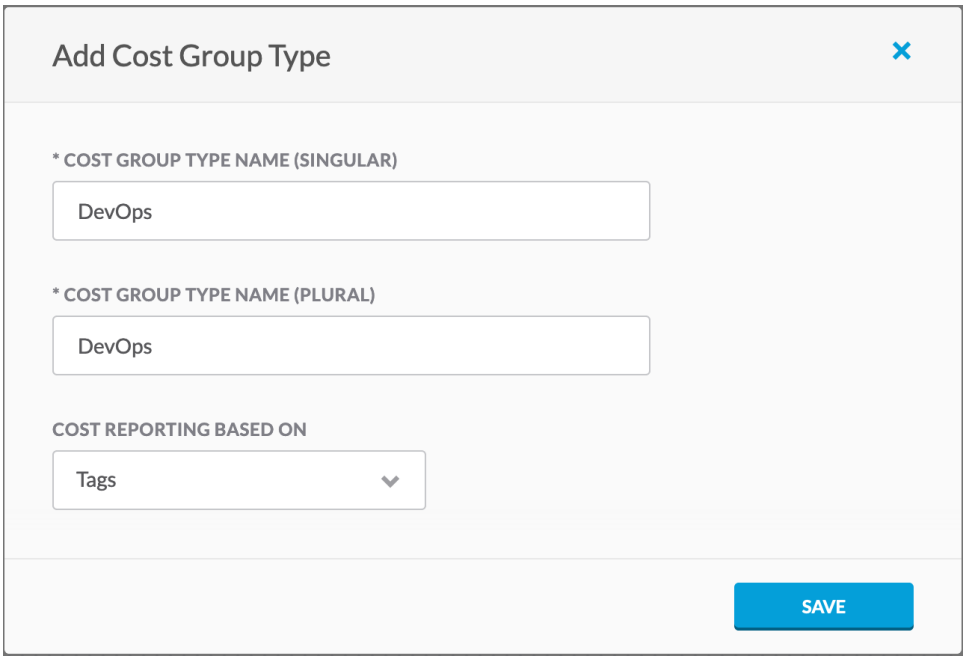

5. Click**Save**.

Add Cost Group

Perform these steps to add a cost group:

1. Click the **Add Department** icon. The **Add Department** dialog appears.
2. Specify the **Department Name** and choose the **Parent Department**.
3. Click **Done**.

Alternatively, you can import departments by uploading a .csv file. A sample .csv file is provided in the **Import Department** dialog for your reference, as shown in the sample screenshot below.

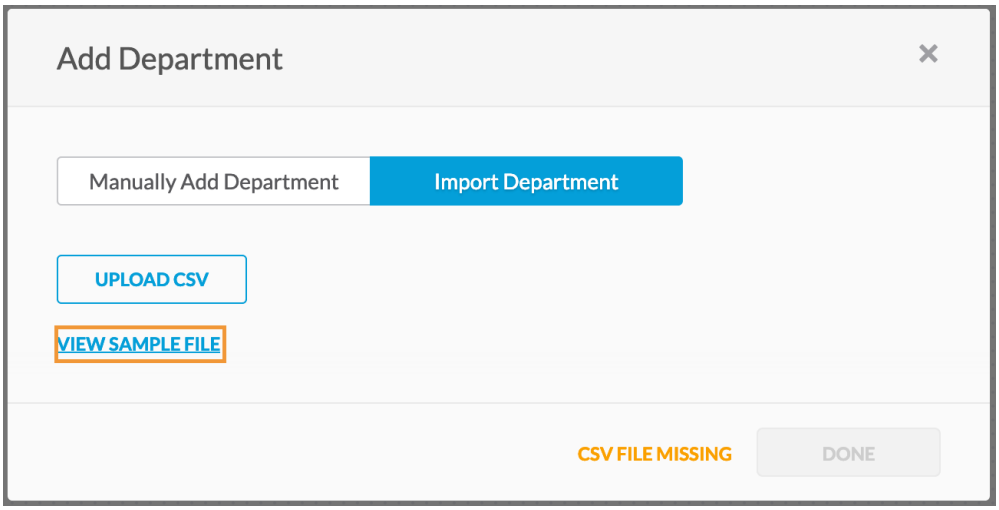

In the **Import Department** dialog, click Upload .csv and navigate to the location on your machine where the .csv file resides to import the file.

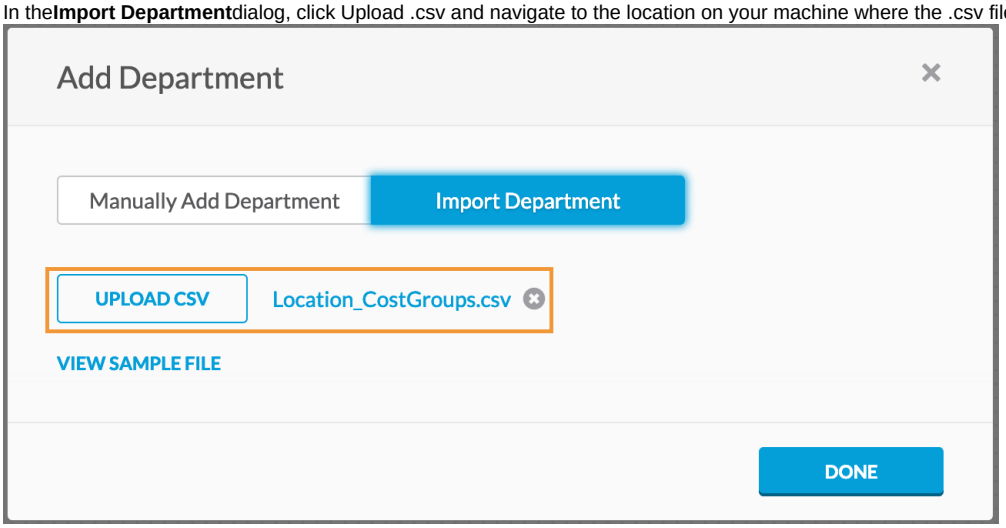

Associate a Billing Unit

You associate a billing unit to a cost group type through a cost group. To associate a billing unit, do the following in the **Billing Units** area.

1. Drag a billing unit from the Billing Unit area and drop the billing unit under a cost group type or cost group.
2. Click the **Move Billing Unit** icon adjacent to a Billing Unit.
3. To move multiple billing units, select the billing units to be moved and choose the **Move Billing Units** icon in the top right corner in the **Billing Units** area.

Associate a Tag

You associate tags that are enabled for cost reporting in Tag-Based Cost Reporting to a cost group type through a cost group. To associate a tag, do the following in the **Tags** area.

1. Drag a tag from the Tag area and drop the tag under a cost group type or cost group.
2. Click the **Move Tags** icon adjacent to a Tag.
3. To move multiple tags, select the tags to be moved and choose the **Move Tags** icon in the top right corner in the **Tags** area.

Allocate Budgets

Allocate Budgets

- Overview
- What's in the Budgets Page?
- Creating Budgets

Overview

Budgets give you the ability to allocate or reserve amounts per cloud or cost group type. Use Budgets to reserve funds utilization.Budgets can be tracked annually and quarterly andprovide the ability to track the total cost costs spent on a cloud or department and allocation of the cost among the various cloud services and billing units respectively.

Depending on the budget allocations, you receive periodic updates, known as Alerts, about budget spending. You also receive updates when your funds' utilization drops below the threshold you define or when the utilization exceeds (or are forecasted to exceed) your budgeted amount.

Use**Budgets**in the **Admin** menu to allocate budgets.

What's in the Budgets Page?

The following is a sample screenshot of the Budgets page.

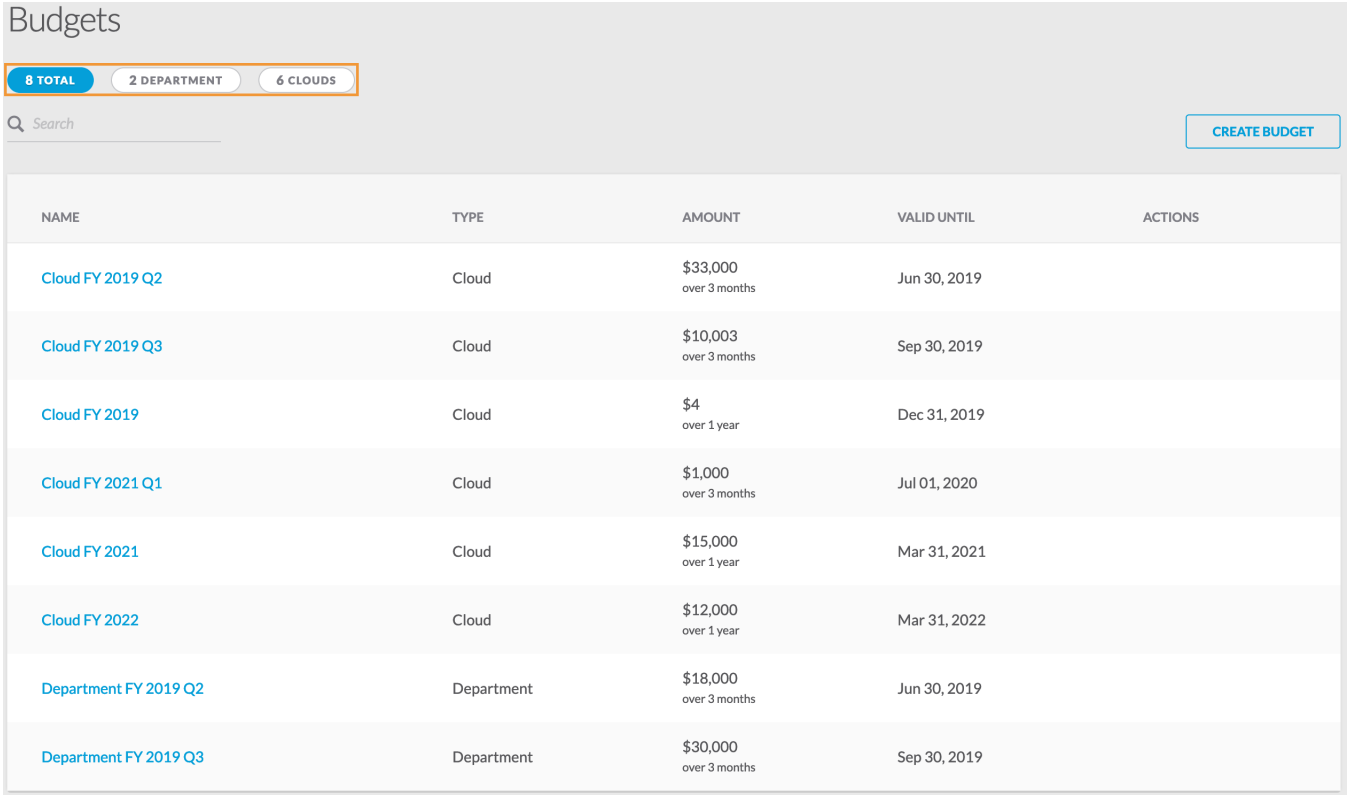

The following table explains the**Budget Summary**that is displayed at the top of the page.

Summary	Description
Total	Number of budget allocations created.
Department	Number of budget allocations assigned to Cost Group Types (Departments).
Cloud	Number of budget allocations assigned to cloud accounts.
Search	Search budgets based on the specified text.
Create Budget	Button to create a budget for a specific year.

The following explains the various aspects of the Budgets page.

Identity	Description
Name	System-generated name, which includes the Cost Group or Cloud Type and the duration the budget is being created for.
Type	Cost Group Type or Cloud the budget is allocated to.
Amount	Displayed in the denominationas defined in Suite Admin (seeCurrency Conversion).
Valid Until	Duration of the budget (end of quarter or year).
Action	Perform action-oriented tasks **Edit** or **Delete** a budget.

Creating Budgets

Use the Create Budget button to create budgets for a cloud, cost group type or department. When creating a budget, you can specify alerts specific to the budget by specifying threshold limits in the **Alert Settings** tab. You can choose to use the default threshold limits defined in the Alerts Page or enter new values specific to a budget by editing the values in the appropriate fields. The Alert settings set here override the generic threshold limits set in theAlerts Page.

By default, the **Default Alert Settings** field is toggled **ON**, which allows you to edit or modify the alert settings fields. Toggle **OFF** this field if you wish to use the values set in the Alerts Page.

Use the **Reset to Default** button to revert to the values set in the Alerts page.

Perform the following steps to create a budget.

1. Click**Create Budget.**The **New Budget** page appears.

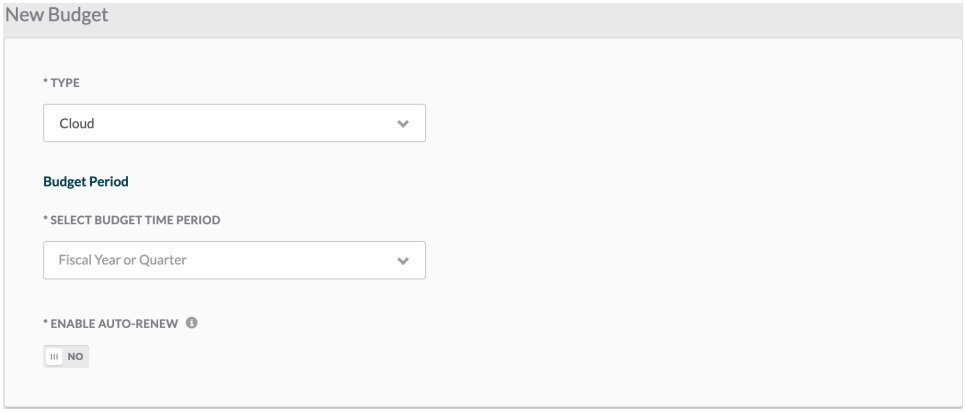

2. In the **New Budget** page, do the following:
    a. Choose the type toassign the budget to.
    b. In the**Budget Period**dropdown, specify the fiscal year or quarter for which budget is to be allocated in the**Select Budget Time Period**fiel d. The information in this dropdown is populated from**Fiscal Year**settings from theSettings Page. If you have not created a Fiscal Year, you can create the fiscal year directly in this step.
    c. Toggle **Enable Auto-Renew** to Yesto renew the budget allocations for the next year or quarter, including the remaining allocations of the current year.
3. In the **Budget Allocations** area, choose the cloud for which the amount must be allocated and enter the amount.

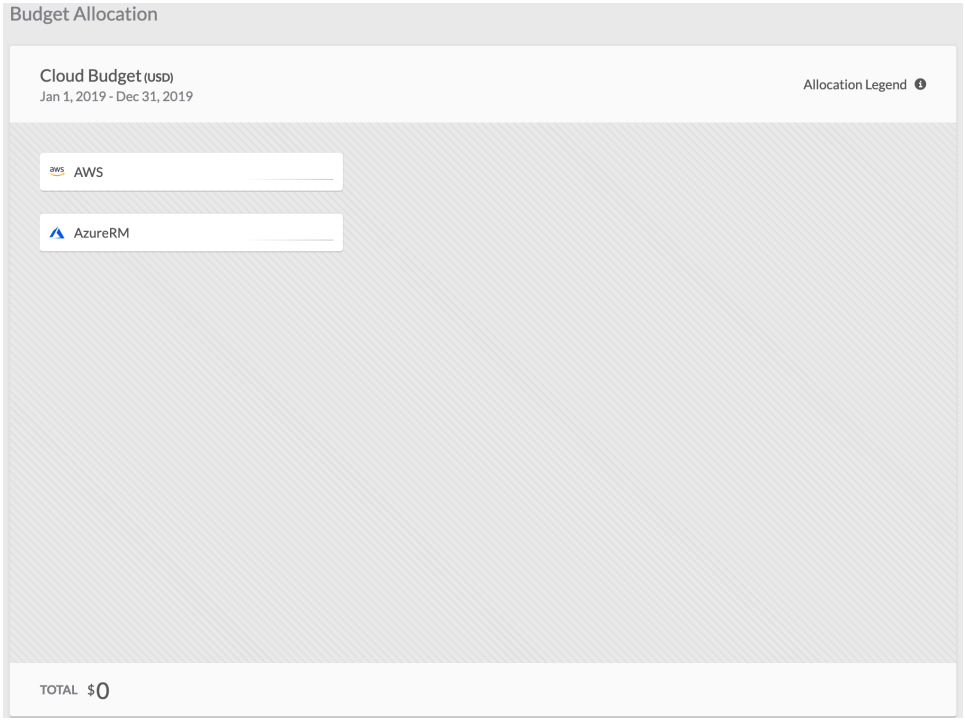

4. Navigate to **Alert Settings** taband do the following:

    a. Specify alerts specific to this budget by updating or entering the values for the following fields as appropriate:

        i. Overspending Threshold (Greater Than)
        ii. Underspending Threshold (Less Than)
        iii. Budget Threshold
        iv. Budget Recipients

    b. In the**Budget Alerts Recipients**field, choose the users or user groups who should be notified for budget-specific alerts when the thresholds are crossed. SeeAccess and Roles.

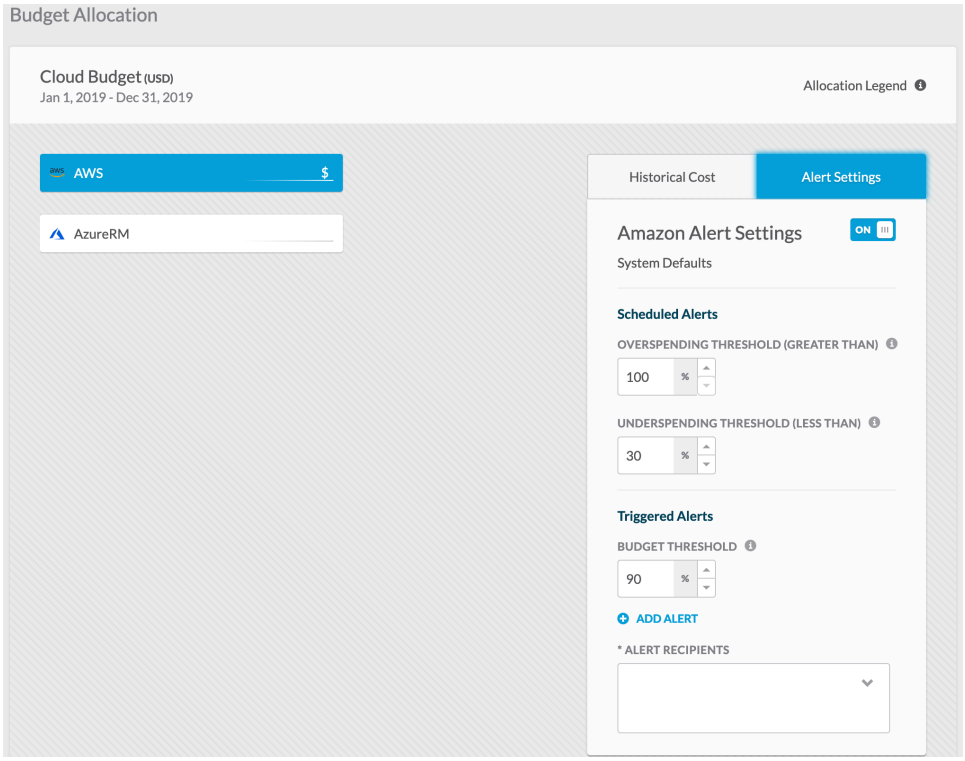

5. Optionally, you can turn off the**Alert Settings** toggle to turn off thealerts for that specific Cost Group entirely.
6. Click **Done**.

Cost Optimizer Dashboard

Cost Optimizer Dashboard

- Overview
- Who Can Access the Cost Optimizer Dashboard?
- What's in the Cost Optimizer Dashboard?
    - Potential Savings Dashlet
    - Projections Dashlet
    - Cost Dashlet
    - Inventory Dashlet
    - Budget Dashlet

Overview

The Cost Optimizer page provides a snapshot of costs incurred by the various clouds installed for an organization.

Who Can Access the Cost Optimizer Dashboard?

The*Cost Optimizer Dashboard*is visible to all users who can access Cost Optimizer. However, information is displayed according to the access levels and is thehome page for this root administrator.For example, the*Cost Optimizer Admin*(seeAccess and Roles)can view information across all cost groups, whereas a*Cost Group Owner*can view data specific to the cost group that the*Cost Group Owner*owns.

What's in the Cost Optimizer Dashboard?

The Cost Optimizer dashboarddisplays informationdepending on your roles defined for you in the system(seeAccess and Roles).The header contains the following icons:

- **Currency**Displays recommendations specific to your cost group. A sample screenshot is shown below.

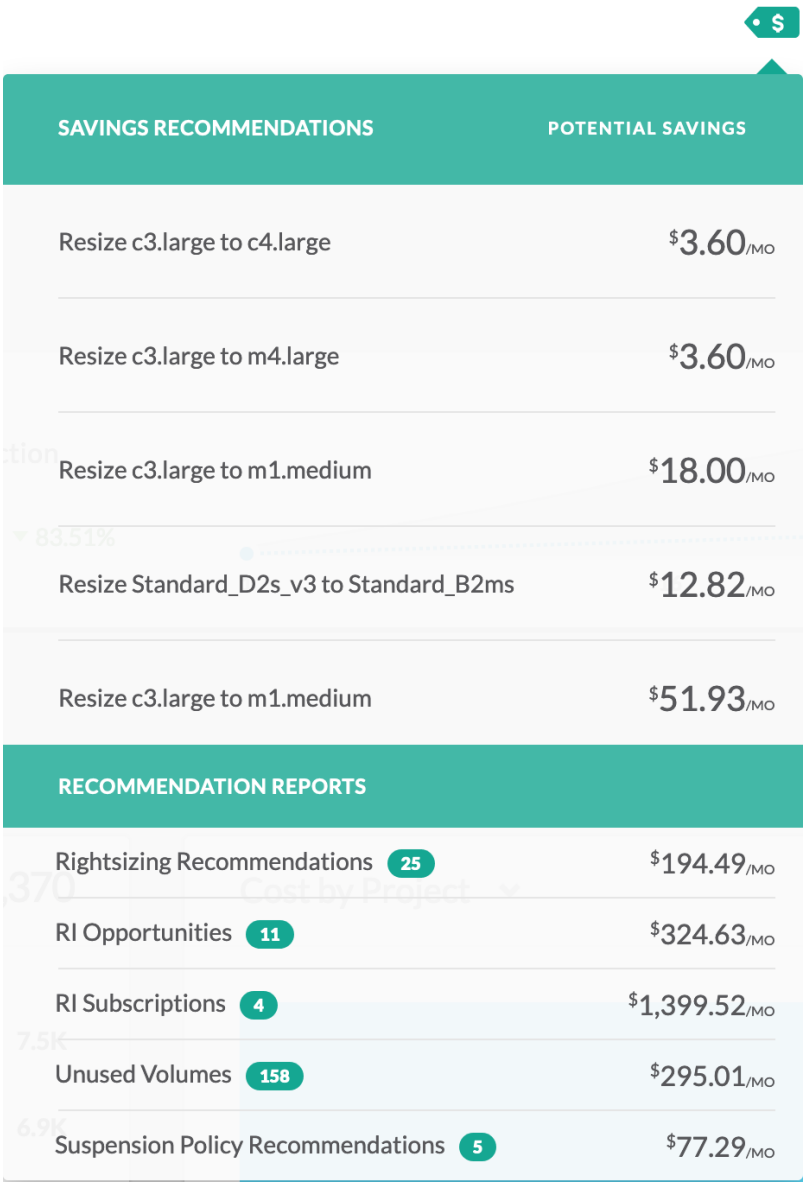
SAVINGS RECOMMENDATIONS
POTENTIAL SAVINGS
Resize c3.large to c4.large
$3.60/MO
Resize c3.large to m4.large
$3.60/MO
Resize c3.large to m1.medium
$18.00/MO
Resize Standard_D2s_v3 to Standard_B2ms
$12.82/MO
Resize c3.large to m1.medium
$51.93/MO
RECOMMENDATION REPORTS
Rightsizing Recommendations 25
$194.49/MO
RI Opportunities 11
$324.63/MO
RI Subscriptions 4
$1,399.52/MO
Unused Volumes 158
$295.01/MO
Suspension Policy Recommendations 5
$77.29/MO

- **Notifications** Displays notifications based on settings specified in theAlerts Page. Asamplescreenshotis shown below. SeeSuite Admin Dashboard> *Notifications*for additional context.

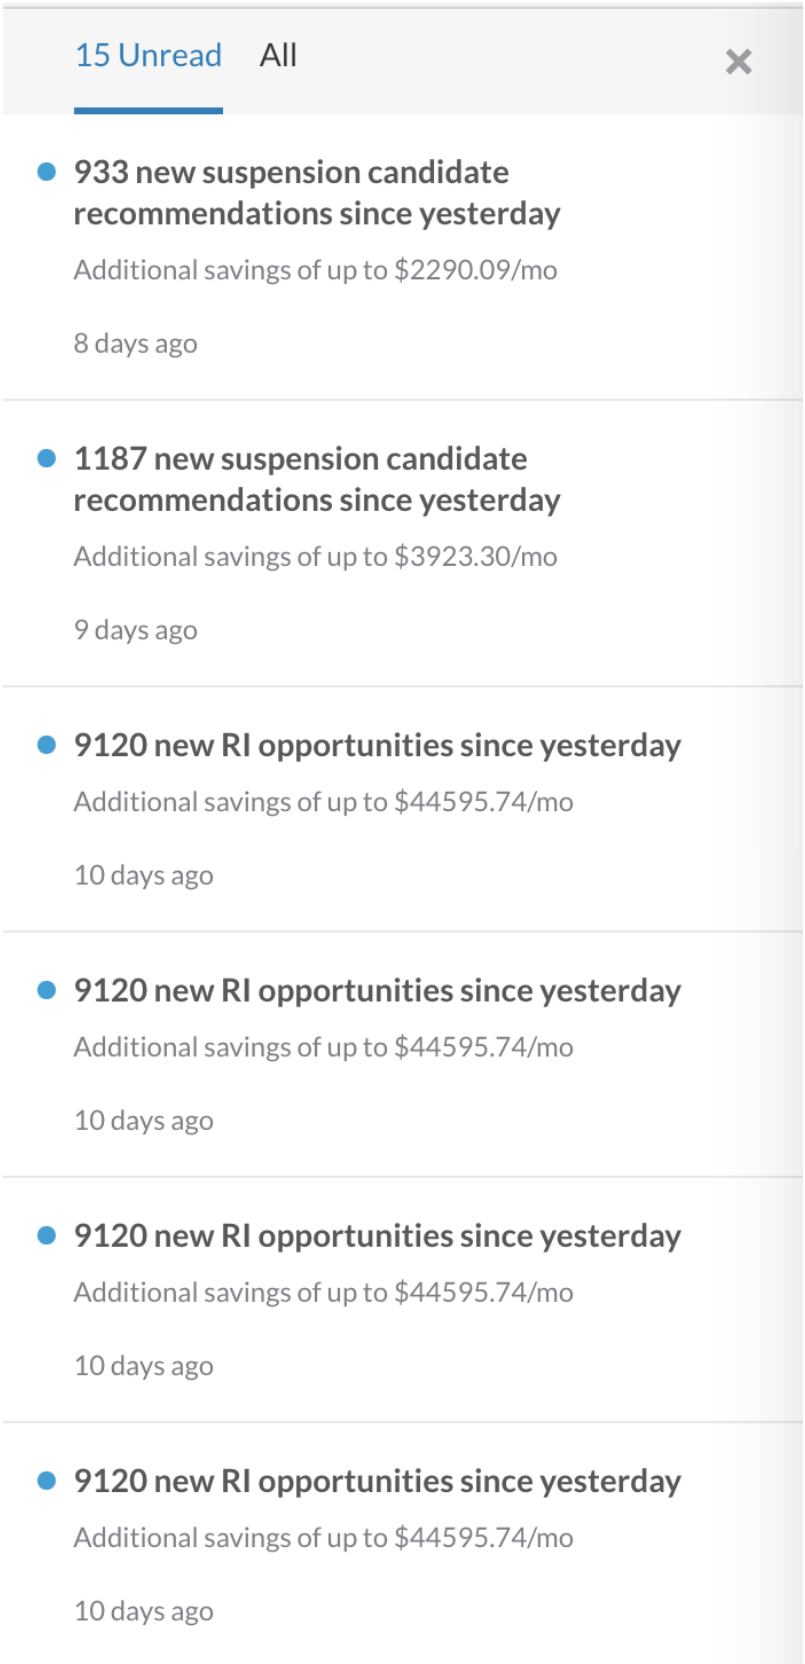

Information on the dashboard can be controlled through the widget next to**All Clouds**. This widget helps you todisplay information for all cloud groups.

In earlier releases of Cost Optimizer, information on the dashboard was classified by cloud providers.

The information is displayed asa summary of recommendations, costs, and inventory through the following:

- Potential Savings Dashlet
- Projections dashlet
- Cost dashlet
- Inventory dashlet

Potential Savings Dashlet

The Potential Savings dashlet displays the approximate, potential savings per month, in the chosen currency (see Currency), where the recommendations are listed. Click**View Details**to display costs that can be saved under the recommendations. Click on a recommendation to navigate to the page and act on the rightsizing recommendation on a resource to save and optimize costs.

This dashlet is not displayed if there are no recommendations. If the recommendations are not acted upon, the savings will decrease for the remainder of the month. Ultimately, the savings could be very minimal at the end of the month.

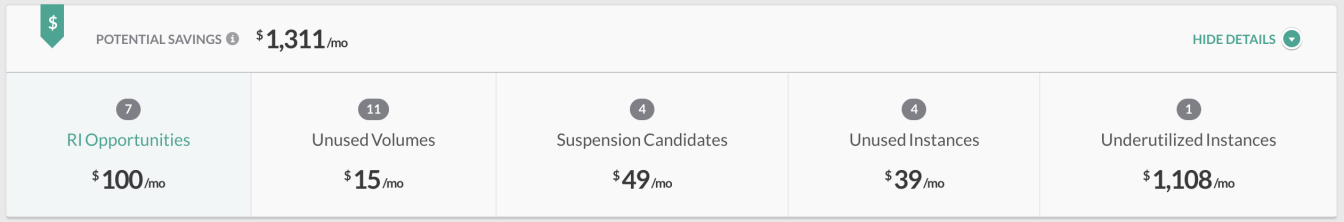

Projections Dashlet

The Projection dashlet provides information about cost projections for the month. The graph on the right denotes the cost (incurred and project for the current month) over the previous month.

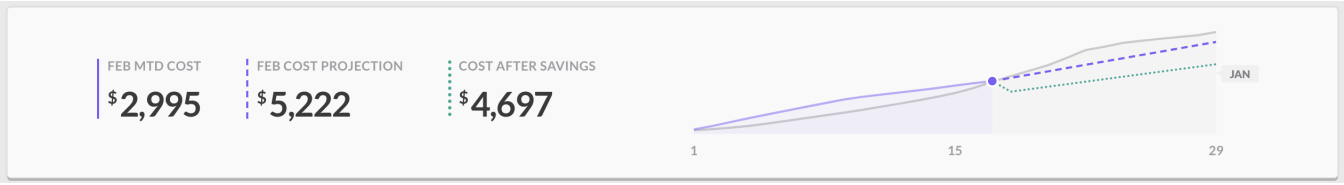

The cost after savings is calculated by using the formula is:

monthly-cost-projection* - *potential-savings-per-day** remaining-*days-in-month

The potential-savings-per-day is calculated as follows:

***potential-savings-per-month* / 30**

The ***remaining-days-in-month*** is calculated using the below formula:

remaining-days-in-month = total-number-of-days-in-month - (current-day - 1)

Using the above formulae and the cost displayed in the above screenshots, the cost after savings is deduced as follows:

$5222 - (1311/30 * (29 - 16 - 1)) =$4697.6

The following table explains the various elements (numbers) in the calculation:

Element	Description
5222	Cost projected for the month of February.
1311	Potential savings per month, taken from the Potential Savings dashlet.
30	Number of days for a month to calculate potential savings per day. This is 30 even though a month could have 31 days.
29	Number of days in February.
16	Current day of the month.

1	Previous day.

Cost Dashlet

The Cost dashlet displays cost projections and cost reports through a high-level snapshot of cost for each cloud group in Cost Optimizer.

- By default, the cost is displayed for all clouds.
- Click the widget next to**All Clouds**to view information for a specific cloud.

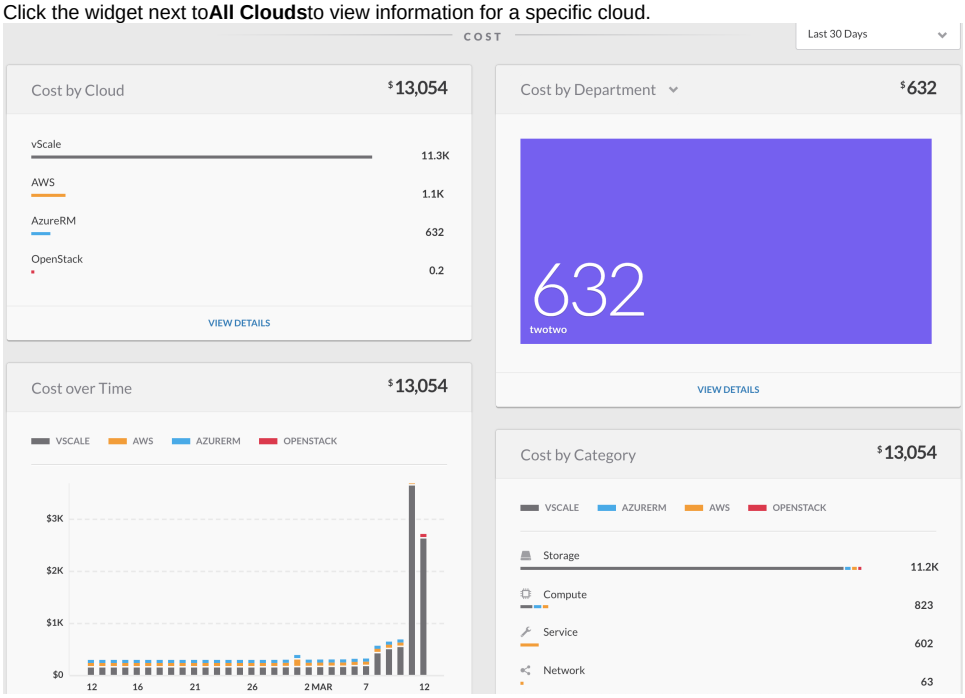

The cost projections for each cloud is displayed through the following reports:

- Cost by Cloud Group
- Cost by Department (Cost Group Type)
- Cost by Organization Hierarchy
- Cost over Time
- Cost by Category

Each report provides the ability to view in-depth details for the specific dashlet. Click**View Details**in each report to open a page that provides information specific to that report. You can also choose a date range for the reports. For more information, seeCost Reports Overview > *Date Range*.

Inventory Dashlet

This dashlet displays information about cloud inventory, such as CPU, memory, storage, etc. The information in parentheses adjacent to the header in eachdashlet indicates the update schedule. SeeData Collection.

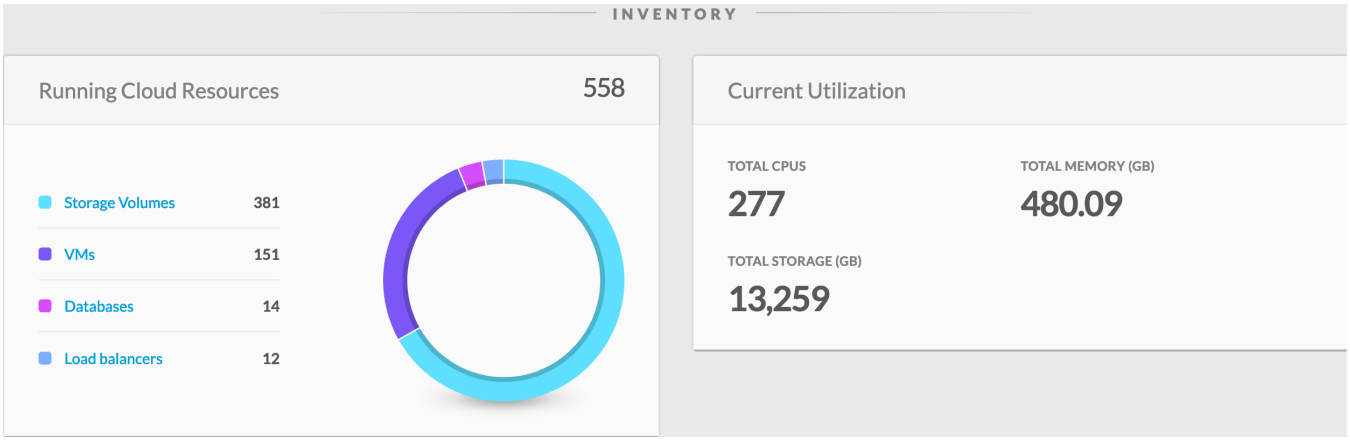

Budget Dashlet

The Budget dashlet displays thefollowing details:

- Forecasted budget projections based on the Alert Settings specified when Configuring Budgets.
- Budget reports for cost groups and cloud groups.

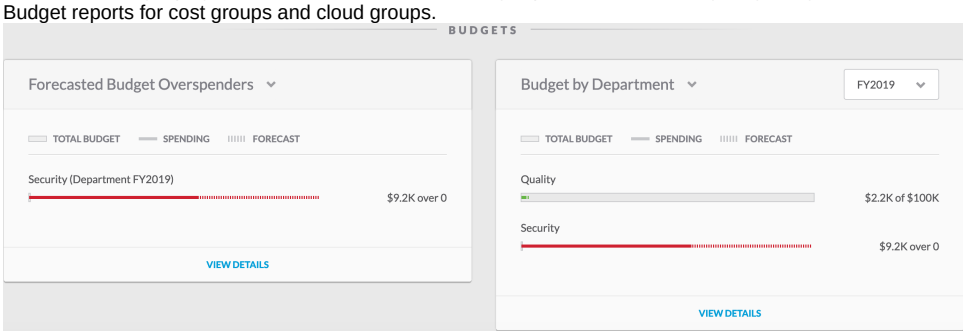

Cost Reports

Cost Reports

- Cost Reports Overview
- Cost by Cloud Group
- Cost by Category
- Invoice Report
- Cost Over Time
- Cost by Tags
- Cost by Cost Group Type
- Cost by Organization Hierarchy

Cost Reports Overview

Cost Reports Overview

- Introduction
- Cost Reports UI
- Filter
    - Advanced Options
        - Saving Filters
        - Scheduling Reports
- Date Range
    - Custom Date Range
- Cost Reports

Introduction

The **Cost Reports** page lists reports that help you analyze the data at a granular level.Cost is displayedin the currency and conversion rate as defined in Suite Admin (seeCurrency).Cost Optimizer classifies the data by cloud groups and displays them under the headings available in a dropdown menu in the **Cost Reports** page. Click**Cost**in the left tree pane to open theCost Reports UI.

Cost Reports UI

The Cost Reports displaygraphical and textual views for cost data. You can view consolidated data for all clouds or billing units you can access or specify filter criteria to view specific data that you need. The following table explains the icons specific to Cost Reports UI. Some of these icons might be displayed forsome reports only.SeeUI Behaviorfor details onicons in the UI.

Icon	Description
Filter	Allows you to filter and view cost data for one or more of the following: • Billing units • Cloud families • Cloud groups • Regions • Invoice category • Invoice tags
Download	Downloads the report in a .csv format.
Date Range	Choose a range to display the report.
Charts	Toggles graphical report display between a bar chart and a pie chart.
Schedule	Allows you to send the report via email to recipients on the fixed date.

Cost Reports displays the following:

- **Totalcost** Graphical view of costs
- **Cost percloud** Expandable textual view of costs

Filter

Use**Filter**foran in-depth analysis by furthergranularizing the data to understand accurate cost consumption.The **Filter** panel allows you to filter reports based on a set of options, thereby allowing you to drill down to the exact details that you require.

The values forbilling units, cost groups, and cost group types are autopopulated fromtheCost Groupsconfiguration and from thecloud configurationfor cloud families, cloud groups, cloud regions, and Inventory for invoice categories, cloud categories, andsubcategories.

> ⚠ The AWS Govcloud account is considered as an IAM Account on AWS master or member account. The cost for AWS Govcloud account is reported against the master or member account and the Govcloud is displayed as a region. Therefore, when a cloud is added for AWS Govcloud with a Govcloud user account, no invoice report data is populated.

Advanced Options

The advanced options in Cost Optimizer are as follows:

- Saving Filters
- Scheduling Reports

Saving Filters

You can choose to save a combination of options in the **Filter** menu for future use through the **Save Filters** feature so that you can quickly access and use the filter at a later time. To save a filter, do the following:

1. Choose the required filter options in the **Filter Panel** pane.
2. The**Save** button appears. The **Save New Filter** dialog appears.

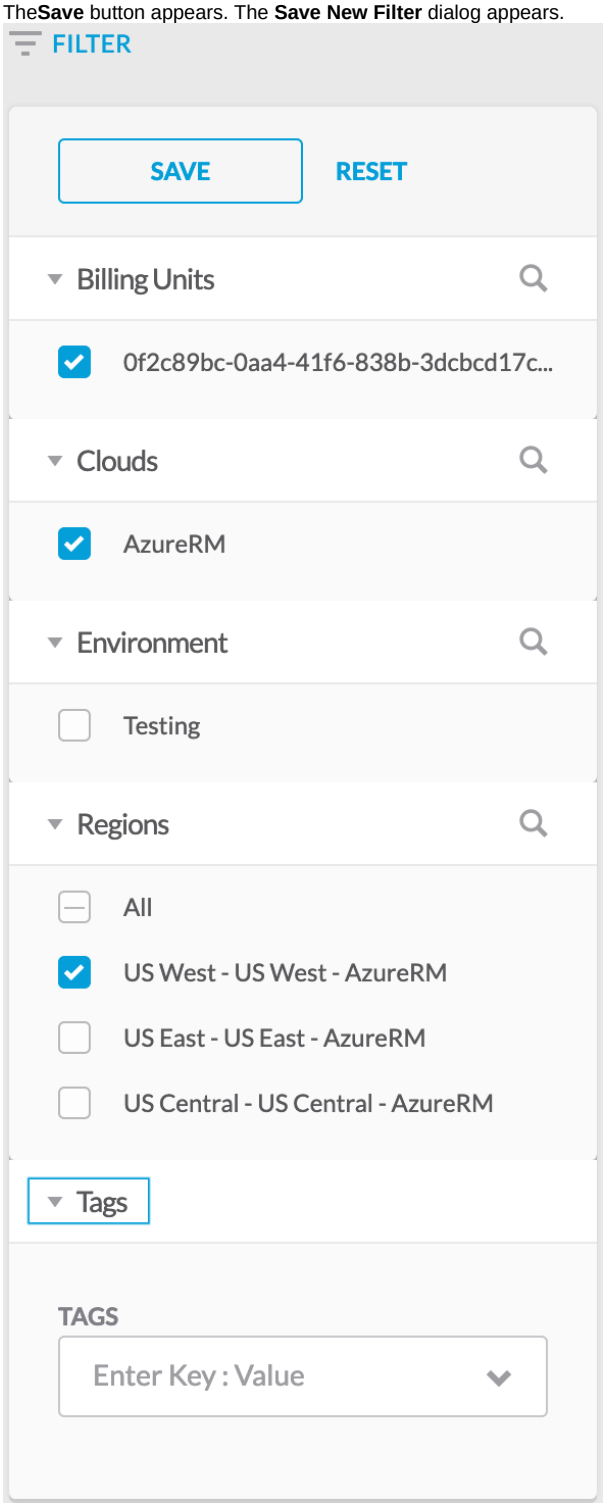

3. Specify a name for this filter and click**Save**. A status message appears indicating that the filter has been saved.

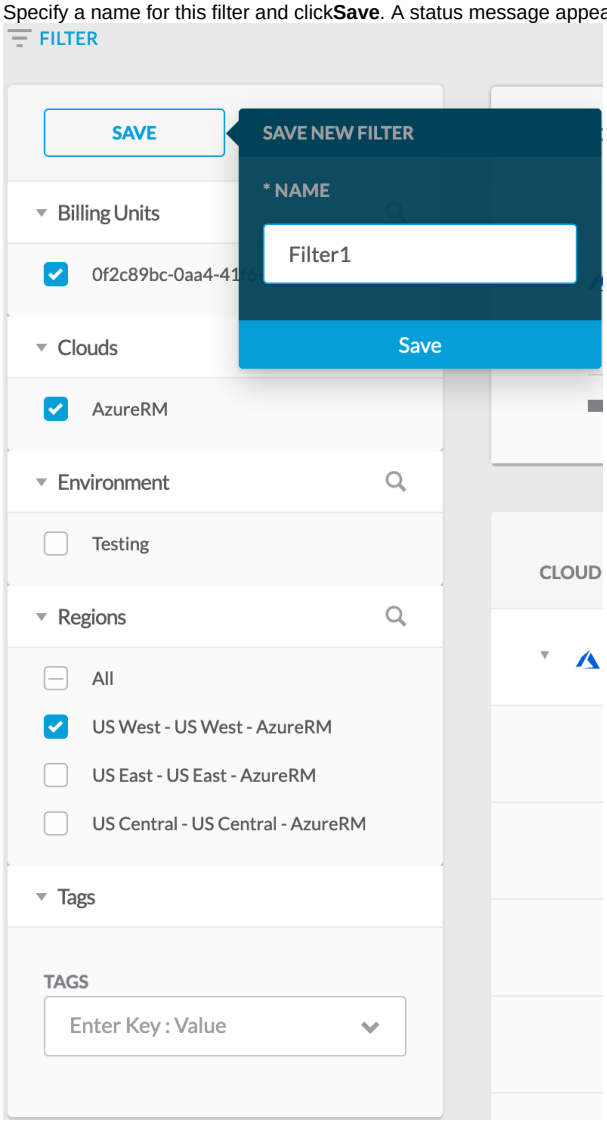

4. You can access and view the saved filters from the dropdown list.

You can also perform the following additional tasks in the **Filter** menu:

- Mark the filter as a favorite by clicking the pin icon next to the filter name.
- Remove the chosen filters by choosing the**Reset**button at any point when saving the filter.
- Delete the saved filter by clicking the **Trash** icon next to a saved filter name. Click **OK** in the **Delete Saved Filter** dialog to confirm the deletion.

Scheduling Reports

The **Scheduler** icon allows you to schedule report generation periodically on a fixed date or at intervals. To create a schedule, do the following:

1. Click the **Scheduler** icon. The **Schedule New *Report Name*** dialog appears.

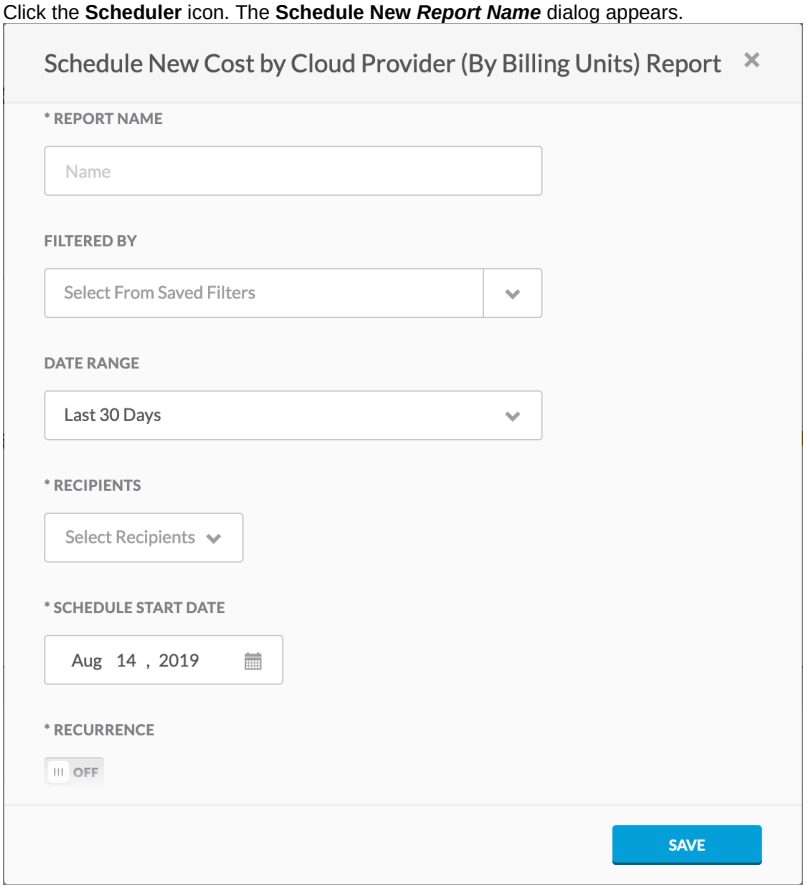

2. Do the following:

   a. Enter a name for the schedule.
   b. Choose filtering options for the schedule from the **Filtered By** field. The information in this field is populated when you save the filtering options as described in the *Advanced Filtering Options* section. You can choose to select a filter or leave the field empty.
   c. Choose the date range.
   d. Select the recipients the report must be sent to.
   e. Specify the start date.
   f. Toggle on the **Recurrence** button to send the report at intervals.
   g. In the **Repeats Every** area, specify the number of times the report must be sent to the recipients and choose the interval **Daily** or **Weekly** . if you choose Weekly, you can also specify the days of the week when the report is sent.
   h. Select the period to end the schedule. The options are:

      i. **Never** Send report forever or until the schedule is deleted.
      ii. **On** Date when the report should be sent.
      iii. **After** Number of occurrences after which the report is not scheduled.

3. Click **Save**. The report is displayed in the**Scheduled*Report Name***dialog as shown in the sample screenshot below.

Optionally, you can use the **Edit** option in the **Actions** column to make changes to the schedule. You can also delete the report using the **Delete** option.

Click **Done** to close the dialog.

The following screenshot displays a sample email format of the report.

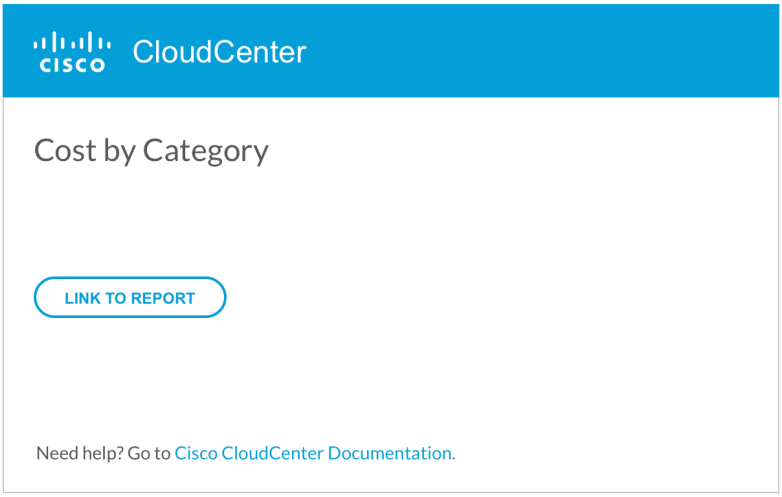

Date Range

The Date Range dropdown allows you to display costs and usage between time periods.The following table explains the date range options:

Date Range	Description
Month to Date	Report for the current month ending with today.
Last 30 Days (Default)	Report for the past 30 days ending with today.
Last Month	Report for the previous month.
3 Months	Report for the past 3 months ending with today.
6 Months	Report for the past 6 months ending with today.
1 Year	Report for a year ending with today.
CUSTOM	A custom range specified by the selected date, month, and year.

Custom Date Range

The Custom Date Range feature enables the ability to view cost reports and usage between specific date ranges. For example, you may want to ensure that a resource did not incur any costs during a specific interval (a specific week).

Choose the**CUSTOM**option in the date range dropdown to specify a range. When a date range is specified, the range appears in all reports that contain the **Date Range**dropdown. You can specify one date range only. You can choose any range between the past twelve months and the current month. When you choose a range for past months, the cost is calculated for the entire month. You must select the dates in the month to view the cost.

You can modify the range multiple times. The range is specific to a session and cannot be deleted. However, the range is removed when you log out of the application or when the application is reloaded.

The format of the dates in the dropdown is determined by the chosen range. The following table explains the formats for the ranges:

Date Range	Format in Date Range Dropdown
Greater than 1 day but less than 31 days	Days. Example: Jan. 29, 19 Feb 10, 19
Greater than 31 days but less than 31 months	Months. Example: Jan. 29, 19 Feb 26, 20

Greater than 31 months	Years. Example: Jan 29, 19 Apr 26, 21

To specify a custom date range, do the following:

1. Click the**Date Range**dropdown and choose**CUSTOM**. The**Custom Date Range**dialog appears.
2. To specify the start date, enter a date in the**From**field in the format displayed orclick the calendar icon.

   Click backward and forward arrows to choose months in a year.To choose a year, click the arrow next to the Month/Year title and select the year.

   To specify the end date, enter a date in the field in the**To**field in the format displayed orclick the calendar icon.
3. Click**Apply**.
4. Use the**Edit**icon to modify the range.

Cost Reports

The following cost reportsare available in Cost Optimizer.

- Cost by Cloud Group Cost incurred for various clouds configured in a cloud account andsegregated on billing units and tags.
- Cost by Category Cost by service categories in cloud providers.
- Invoice Reports Cost per cloud and cost segregation per region and category of supported clouds.
- Cost Over Time Cost incurred for a chosen duration.
- Cost by Tags Cost incurred by tags.
- Cost by Cost Group Type (Department) Cost incurred for a cost group type.
- Cost by Organization Hierarchy Cost associated with the various groups in an organization.

Cost by Cloud Group

Cost by Cloud Group

- Cost by Billing Units in a Cloud Group
- Cost by Tags in a Cloud Group

Cost by Billing Units in a Cloud Group

Cost by Billing Units in a Cloud Group

- Overview
- Description
  - Google Cloud Credits
- Filter

Overview

The **Cost by Billing Units** report displays the costs for one or more billing units. This option is listed when you have**Billing Units**not assigned to a **Cost Group**(See:How Do I...>*Associate a Billing Unit*).

Description

The report displays the total cost of all billing units associated with a cloud group and the cost incurred by the billing unitfor that cloud group.A summary of all running resources is displayed. Click the numbers against each resource to open the corresponding resource page. The **Group** section displays the actual cost for eachbilling unit.

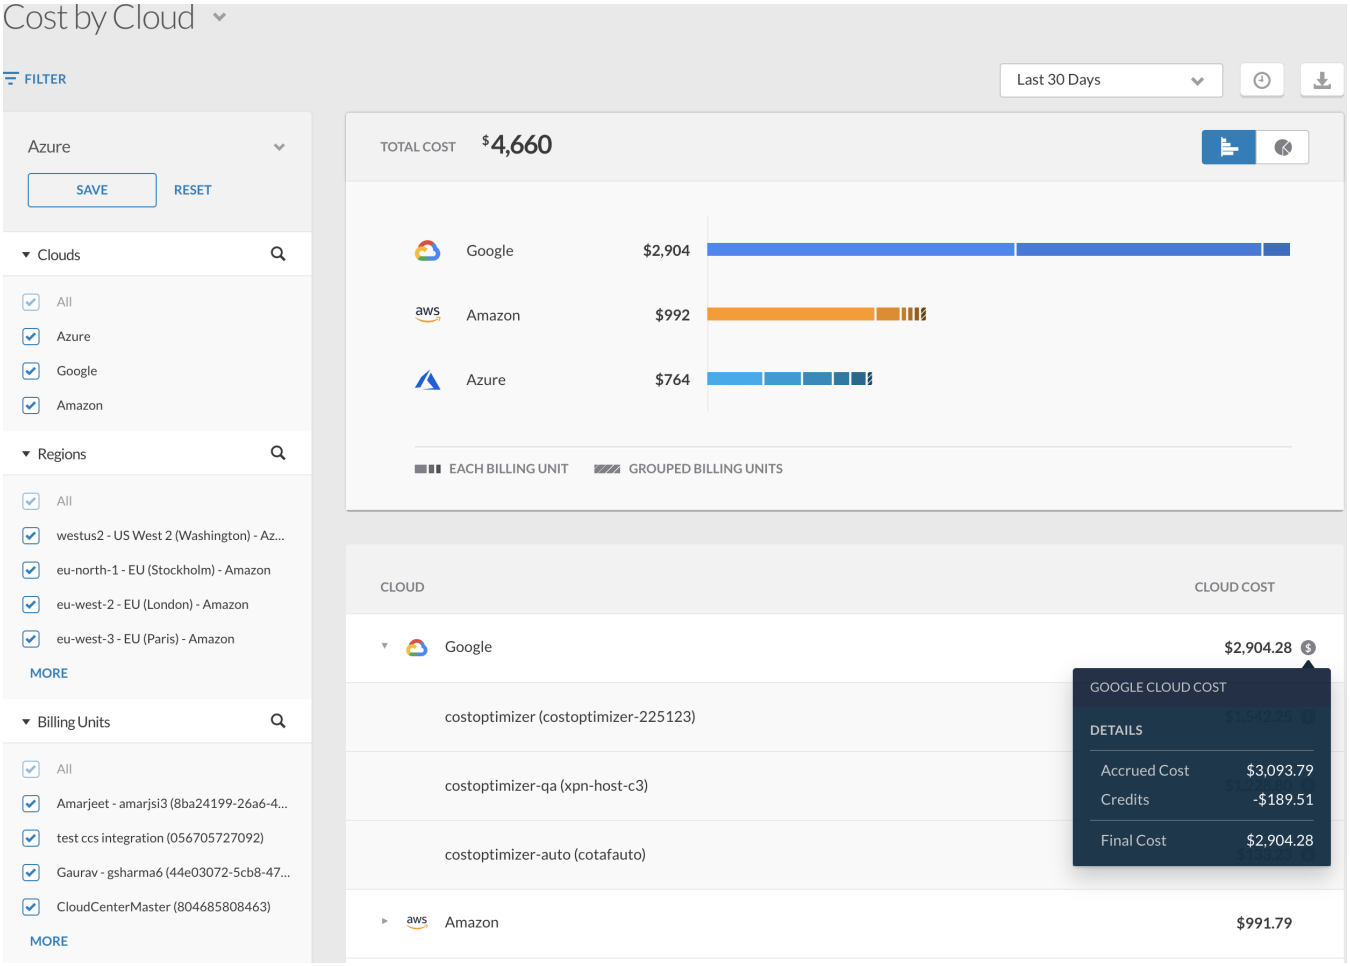

Google Cloud Credits

The total cost incurred for Google clouds is calculated by subtracting the discounts or credits incurred for the period (day, month, etc.) from the total cost. The cost is displayed in the currencyand conversion rate as defined in Suite Admin (seeCurrency). This cost computation is displayed when you hover over the currency icon adjacent to the cost information.

Filter

You can filter the report using the following options:

- Billing Units
- Clouds
- Cost Groups
- Cost GroupTypes
- Regions
- Tags

Cost by Tags in a Cloud Group

Cost by Tags in a Cloud Group

- Overview
- Description
- Filter

Overview

The Cost by Tags report displays the cost for one or more tags, for which cost is enabled inTag-Based Cost Reporting page.See:How Do I... > *Associate a Tag*.

> This report is available for AWS and Azure cloud accounts only.

Description

Thefollowing is a sample screenshot of the**Cost By Tags**report.The report displays the total cost, the cost per cloud account, and the cost per tag.The report displays the total cost of all tags associated with a cloud account and the cost incurred by the tag for that cloud account.The report also displays in different shades the cost per each tag in the cloud.Click the arrow next to the cloud account to display the tags and the costs associated with the tags in the cloud account.

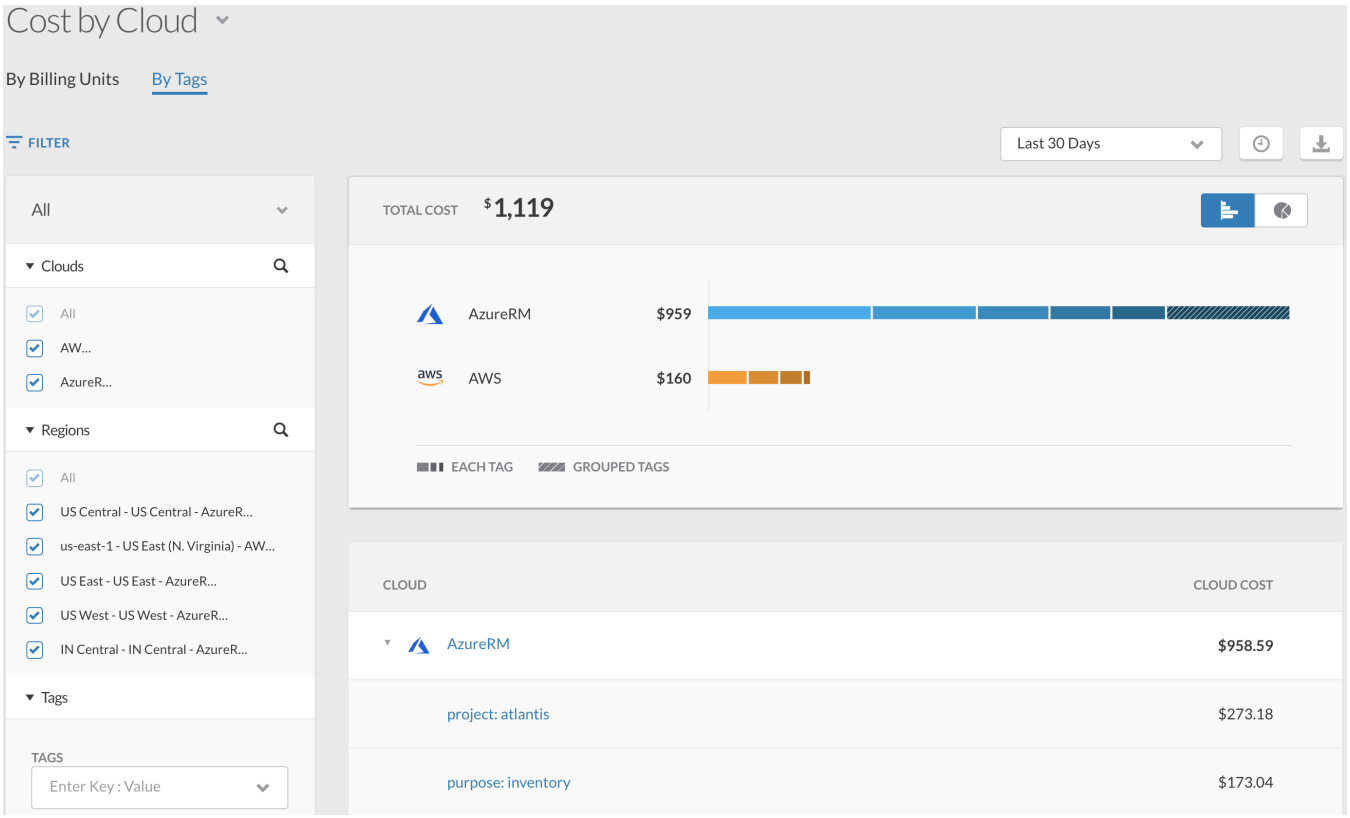

Filter

You can filter the report using the following options:

- Billing Units
- Clouds
- Cost Groups
- Cost GroupTypes
- Regions
- Tags

Cost by Category

Cost by Category

- Overview
- Description
- Filter

Overview

The **Cost by Category** report displays cost by service categories for one or more cloud groups. Service costs for cloud groups are displayed in this report. There are two types of categories inCost Optimizer:

- **Invoice**Examples: Storage, Network, Compute, etc.
- **Cloud Provider-specific** Examples: Categories, such as App Engine, Route 53, Cloudwatch, and subcategories, such as Data Transfer In or Out, Number of Requests, EBS Volume Usage.

Description

The following is a sample screenshot that displays:

- **Total Cost** Display the total cost and distribution of various categories across cloud providers.
- **Category per Cloud Provider** Display categorywhich can be expanded to reveal costs incurred by a cloud provider for each category and subcategory.

> **Other** refers to cloud categories that are not classified in Cost Optimizer.

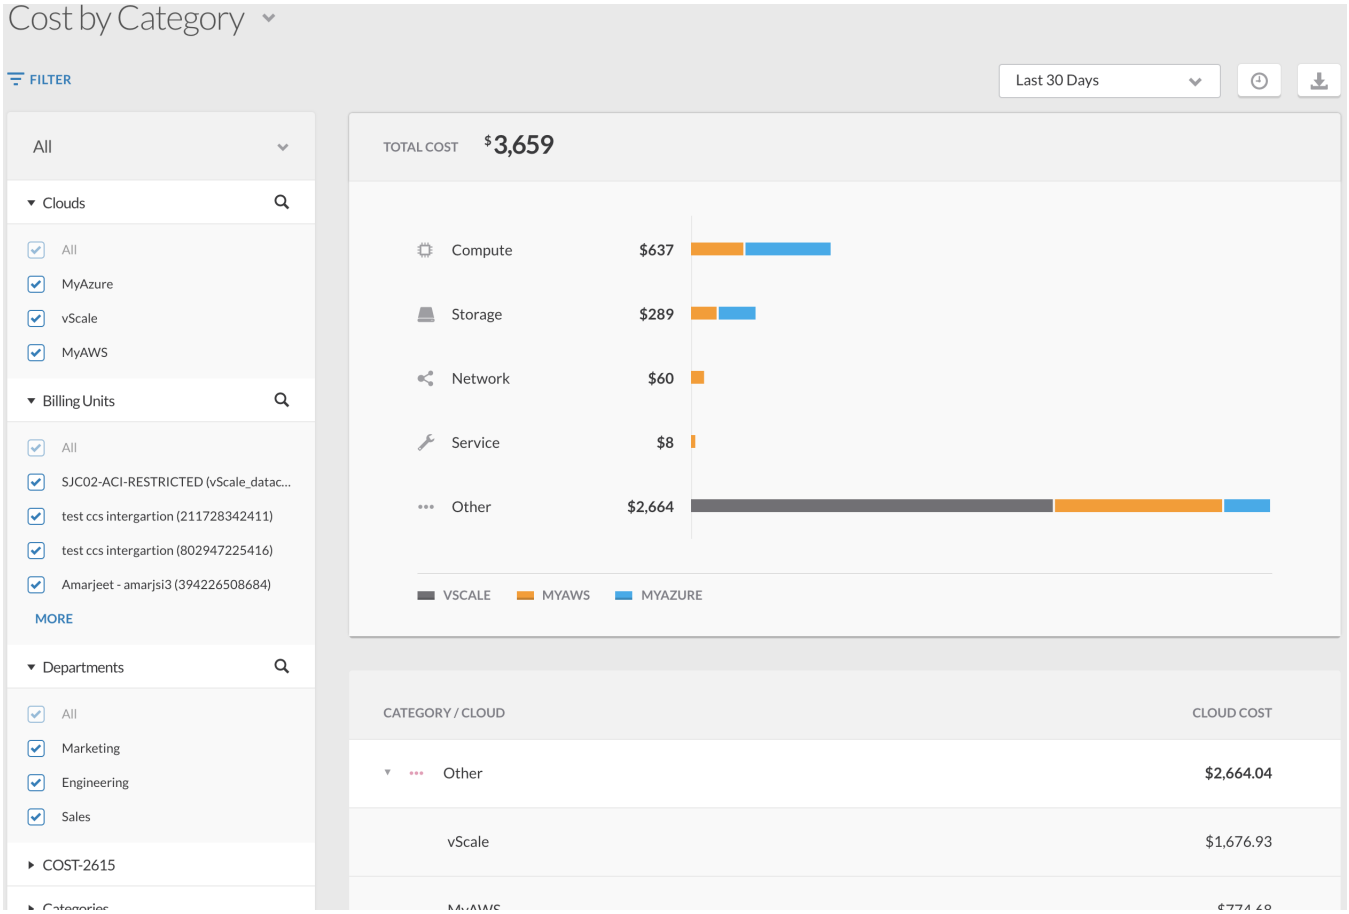

Filter

You can filter the report using the following options:

- Billing Units
- Cloud Categories (example, App Engine, Route 53, EC2)
- Cloud Subcategories (example, Data Transfer In or Out, EBS Volume)
- Cost Groups
- Cost Group Types
- Invoice Types (for example, Service, Network, Storage)
- Tags

Invoice Report

Invoice Report

- Invoice by Category
- Invoice by Region

Invoice by Category

Invoice by Category

- Overview
- Description
- Filter

Overview

The**Invoice by Category Report**displays the cost of each cloud group for the various categories, such asStorage, Network, Compute, and so on for a chosen duration. The report classifies the data to the lowest unit that can be billed.

Description

The following is a sample screenshot of the report that displays the following:

- **Total Cost** Display the total cost and distribution of various categories across cloud providers. You can choose to view the invoice category for a cloud by choosing the appropriate option from the **Select Cloud** dropdownlist.
- **Category** Display categorywhich can be expanded to display each category and subcategory.

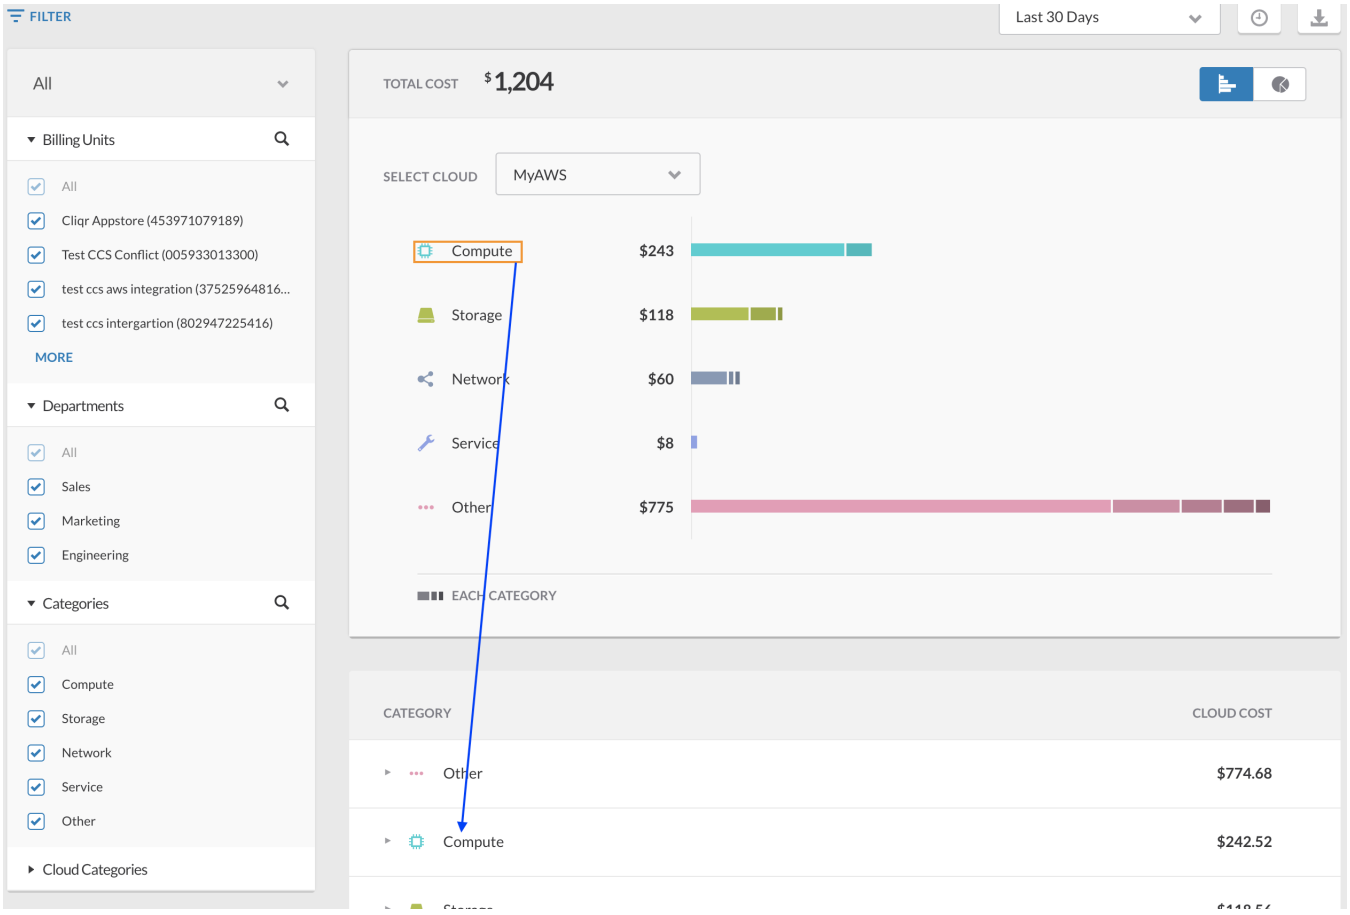

Expenses incurred by customers are categorized by the cloud provider into cloud categories and cloud subcategories. These expenses are retrieved every day and displayed in this report. Cloud providers refer to the cloud categories and cloud subcategories by using different labels, such as services, usage types, etc.

To offer a high-level view of these expenses (invoice costs), Cost Optimizer classifies the expenses into buckets or*derived*categories when the expenses are retrieved. Cost Optimizer uses classifiers to categorize expenses into the most appropriate bucket from the cloud category and cloud subcategory. When you use a new service offered by a cloud provider, the cost incurred for this new service is displayed on the report. If the existing classifiers are not matched with the new cloud category or subcategory in an appropriate bucket, the cost incurred for the new service is added in the*Other*bucket.

Filter

You can filter the report using the following options:

- Billing Units
- Cloud Categories
- Cloud Subcategories
- Cost Groups
- Cost Group Types
- Invoice Categories
- Tags

The filter panel changes according to the cloud chosen in the **Select Cloud** dropdown list. For example, if you choose GCP from the list, the filter includes **Cost Groups Types** and **Cost Groups**, in addition to the above options.

Invoice by Region

Invoice by Region

- Overview
- Description
- Filter

Overview

The**Invoice by Region Report**displays the cost of each cloud group across geographyfor a chosen duration.The data is categorized based on the service used on the cloud. The value indicates the combined value of all regions. For private and container clouds, the value is displayed for each configured region.

> This report is not available for Google cloud.

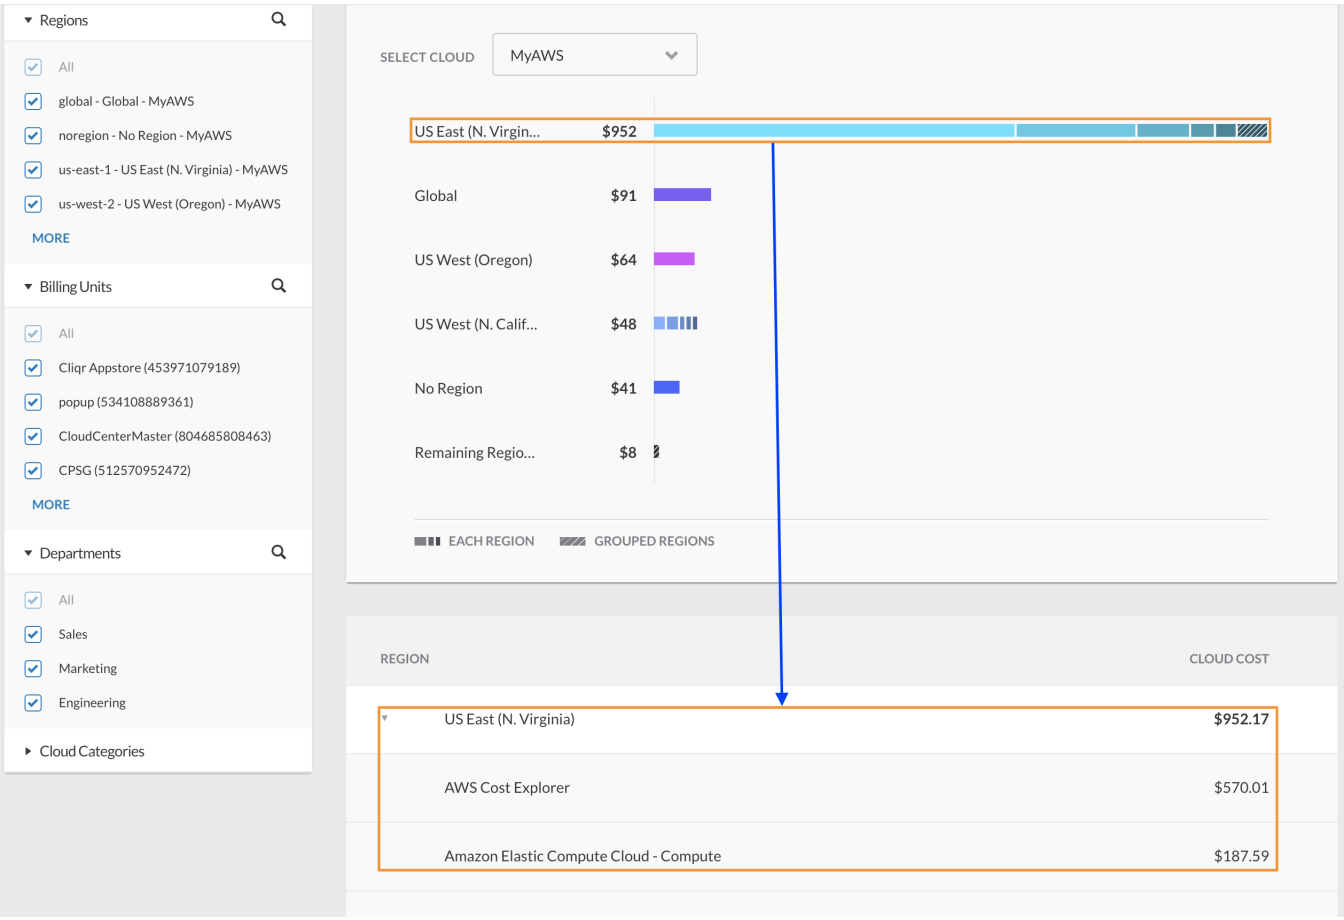

Description

The following is a sample screenshot of the report that displays the following:

- **Total Cost** Display the total cost for a cloud provider in a region. You can choose to view the cost of a cloud by choosing from the**Select Cloud** dropdownlist.
- **Region** Display cloud per region and the cloud category, such as storage, network, computer, and so forth.

> The various shades of color in the report correlate to the categories for that cloud region.

Filter

You can filter the report using the following options:

- Billing Units
- Cloud Categories
- Cost Groups
- Cost Group Types
- Regions
- Tags

> ⚠ The filter panel changes according to the cloud chosen in the **Select Cloud** dropdown list. For example, if you choose GCP from the list, the filter includes **Cost Groups Types** and **Cost Groups**, in addition to the above options.

Cost Over Time

Cost Over Time

- Overview
- Description
  - Google Cloud Credits
  - Filter

Overview

The**Cost Over Time**report shows the cost incurred for a duration that you choose from the date range dropdown for the supported clouds.

Description

The**Cost Over Time**report displays the following reports:

- **Total Cost** Displays the costs for the chosen clouds and period in a bar chart.
- **Cloud** Displays a tabular view of the costs per cloud. Expanding each cloud displays the cost incurred per each account in the cloud.

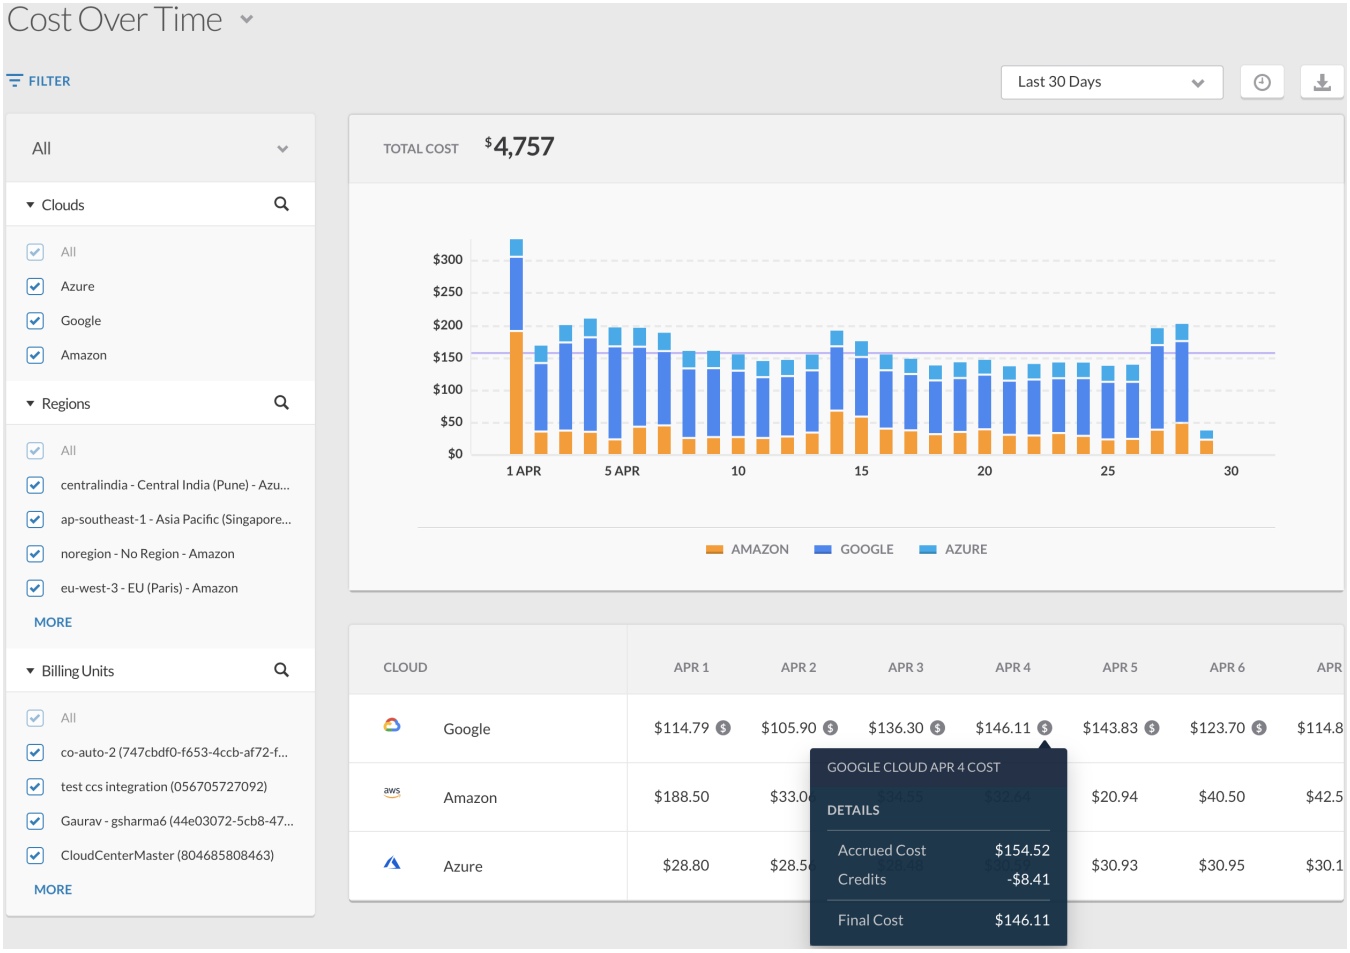

Google Cloud Credits

The total cost incurred for Google clouds is calculated by subtracting the discounts or credits incurred for the period (day, month, etc.) from the total cost. The cost is displayed in the currencyand conversion rate as defined in Suite Admin (seeCurrency). This cost computation is displayed when you hover over the currency icon adjacent to the cost information.

Filter

You can filter the report using the following options:

- Billing Units
- Cloud Families
- Clouds
- Cost Group
- Cost Group Types
- Regions
- Tags

Cost by Tags

Cost by Tags

- Overview
- Description
- Filter

Overview

The**Cost by Tags**report displays the cost incurred by tags in the absence of cost groups, provided Tag-Based Cost Reporting is enabled.

> Tag-based cost reporting is available for AWS and Azure clouds only. When tag-based cost reports are shared, the sharing results in displaying additional cost, inventory, and recommendations for the resources associated with the cost groups.

Description

The following is a sample screenshot of the Cost by Cost Group Type report that contains the following reports:

- **Total Cost** Displays total cost across tags. Click the number against the Inventory types (Virtual Machines,Storage Volumes, Services (load balancers and databases)) to open the respective pages.
- **Tag** Displaysa textual view of the costs per tag.

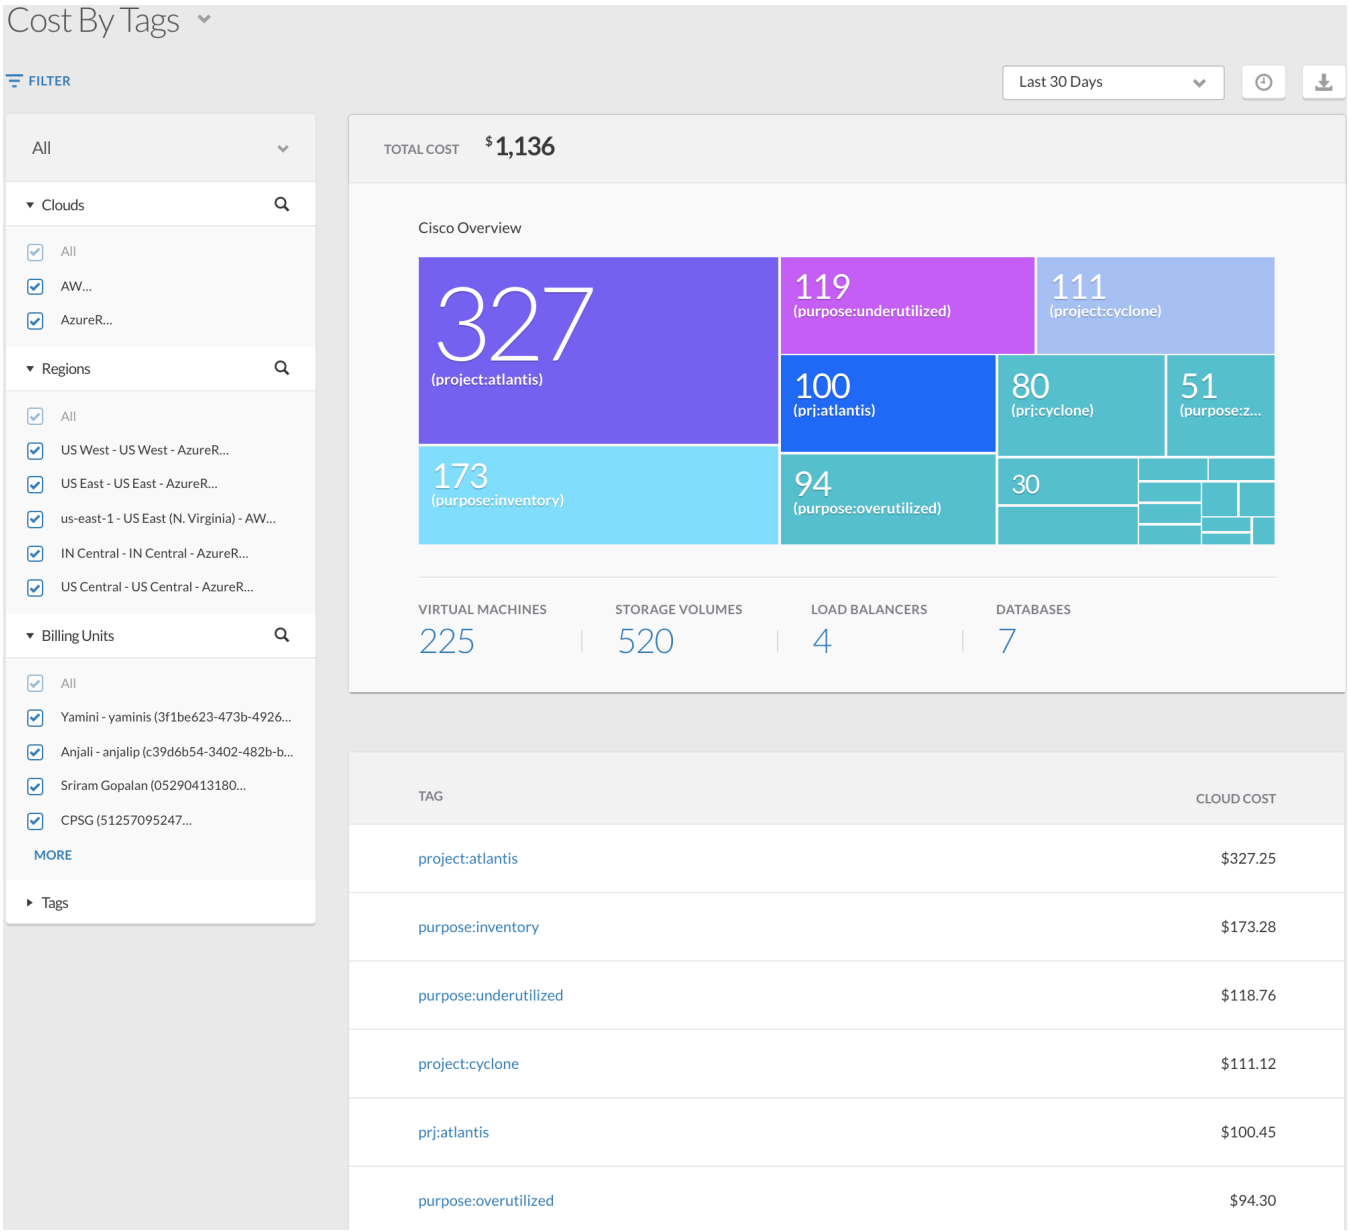

Filter

You can filter the report using the following options:

- Billing Units
- Clouds
- Regions
- Tags

Cost by Cost Group Type

Cost by Cost Group Type (Department)

- Overview
- Description
- Filter

Overview

The**Cost by Cost Group Type**report displays the cost incurred for a specific cost group type (seeCost Groups Configuration). Click the arrow and choose the cost group type to view the reportfor each cost group type.

Description

The following is a sample screenshot of the Cost by Cost Group Type report that contains the following reports:

- **Total Cost** Displays total cost across cost groups and billing units in the cost group type.You can toggle the display between cost groups and billing units associated with the cost group type. Click the number against the Inventory types (Virtual Machines,Storage Volumes, Services (load balancers and databases)) to open the respective pages.
- **Group** Displaysa textual view of the costs per cost group which can be expanded to reveal the cost per billing unit.

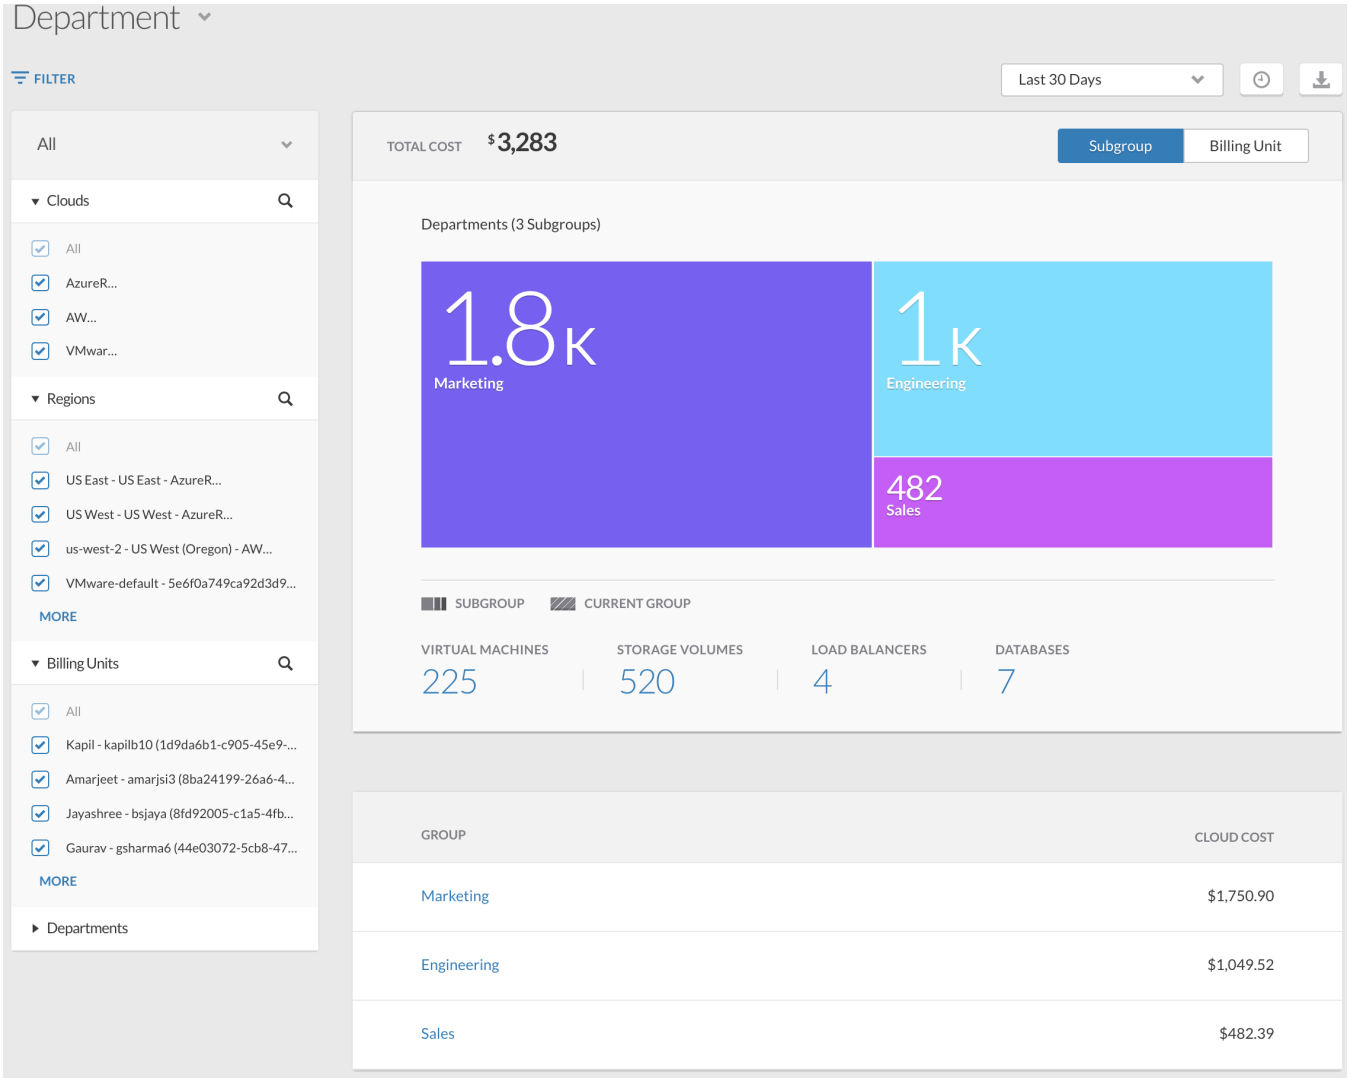

Filter

You can filter the report using the following options:

- Billing Units
- Clouds
- Cost Groups
- Regions
- Tags

Cost by Organization Hierarchy

Cost by Organization Hierarchy

- Overview
- Description
- Filter

Overview

The**Cost by Organization Hierarchy**report displays the costs incurred by theorganization. The hierarchy is created in the cloud provider portal. When configuring a cloud, change the**Enable Reporting By Org Structure**toggle to **On** to import the organization hierarchy created in the cloud provider portal into Cost Optimizer.See Configure Clouds.

Description

The report displays the cost associated with the various groups created under an organization on a cloud provider and the cost incurred by the **Engineering** group in **Master AWS** organization.

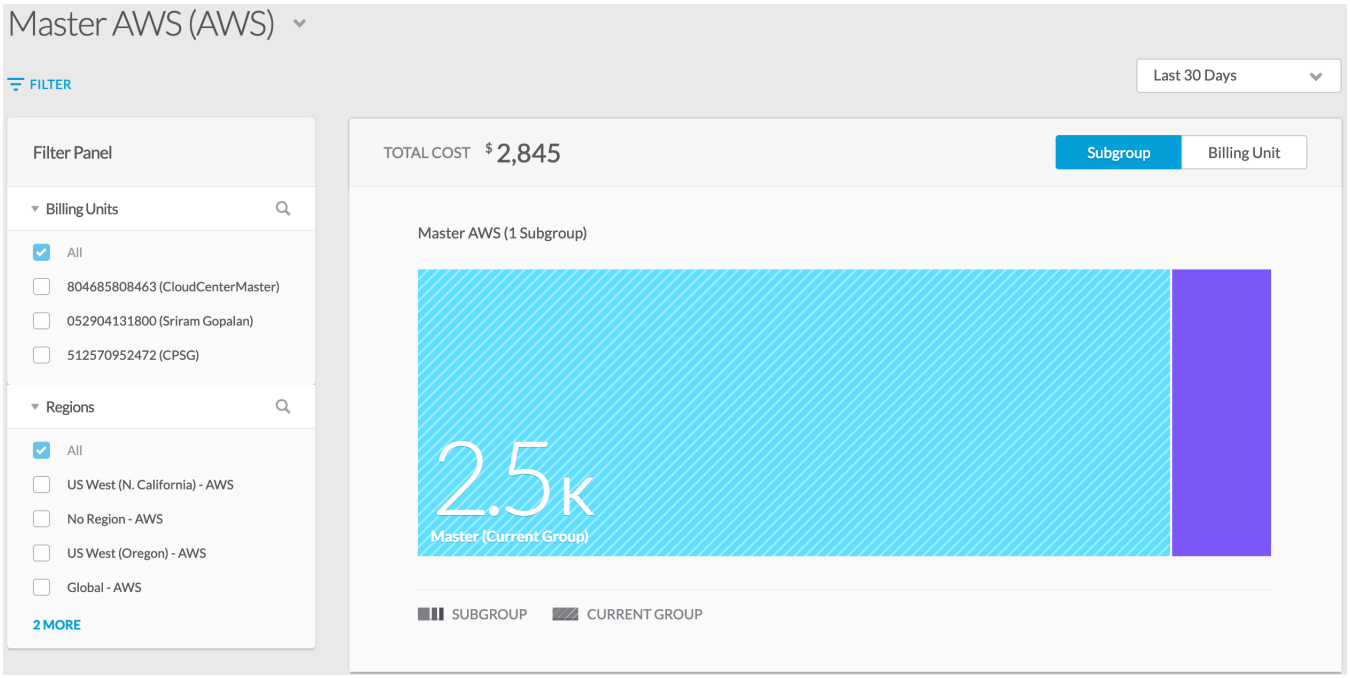

The**Group**section displays the actual cost for each group. A link on a group indicates subgroups for that group. Click**Group** link to display the individual cost for each subgroup in that group.

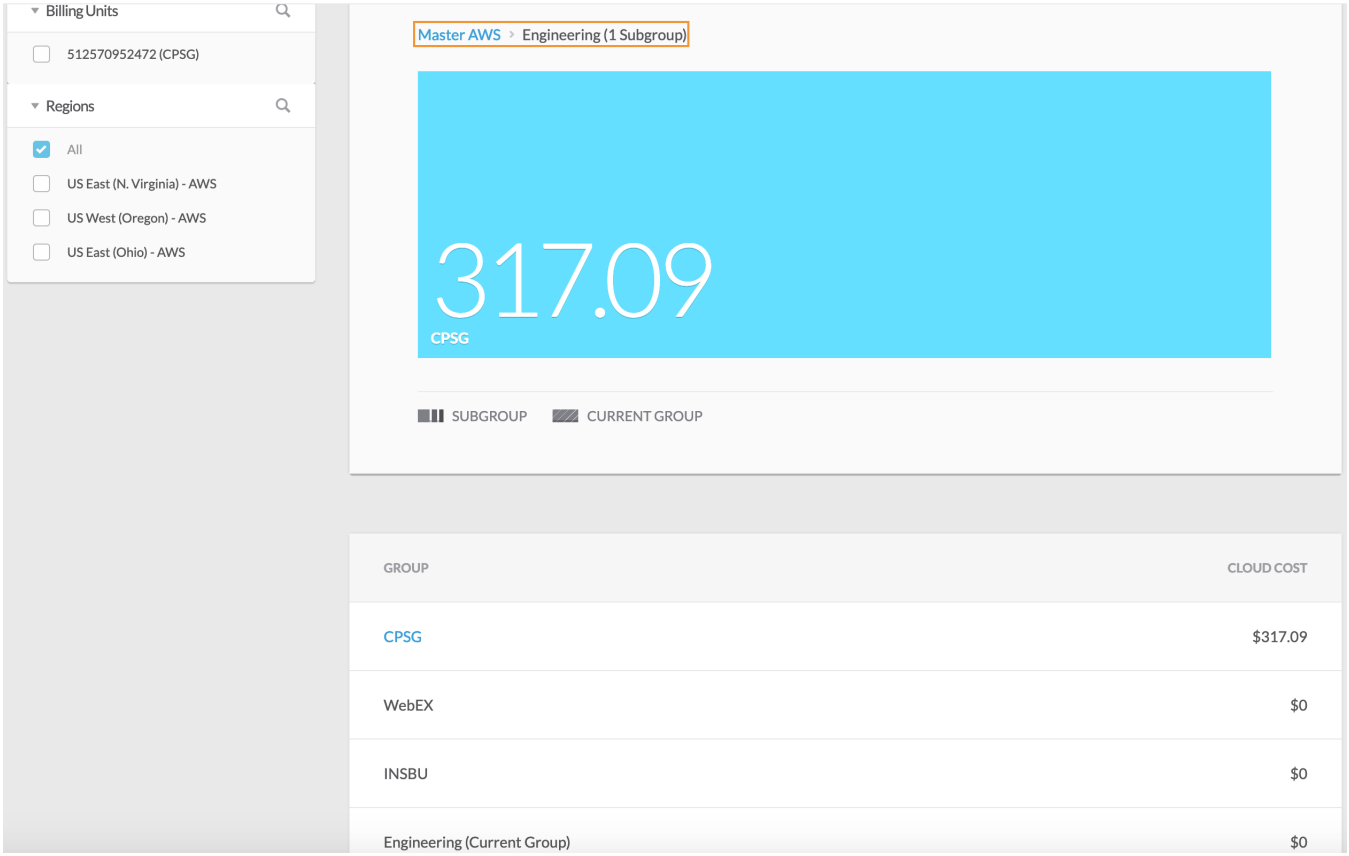

Filter

You can filter the report using the following options:

- Billing Units
- Cloud Regions
- Tags

Budget Reports

Budget Reports

- Budget Reports Overview
- Budget Overspenders
- Budget Underspenders
- Budget By Cloud
- Budget By Cost Group Type

Budget Reports Overview

Budget Reports Overview

- Overview
- Who Can Access the Page?
- What's in the Budget Dashboard?
- Budget Reports

Overview

The Budget Page provides a snapshot of the budget allocations and spending in an organization.

Who Can Access the Page?

The*Budget*pageis visible to all users who can access Cost Optimizer. However, information is displayed according to the access levels and is thehome page for this root administrator.For example, the*Cost Optimizer Admin*(seeAccess and Roles)can view information across all cost groups, whereas a*Cost Group Owner*can view data specific to the cost group that the*Cost Group Owner*owns.

What's in the Budget Dashboard?

Information in the dashboard can be controlled through the widget below the header.

The Budgets Reports displaygraphical views for budget data. You can view data for all clouds or billing units you can access or specify filter criteria to view specific data that you need. The following table explains the icons specific to Budgets Reports UI. Some of these icons might be displayed forsome reports only.SeeUI Behaviorfor details onicons in the UI.

Icon	Description
Download	Downloads the report in a .csv format.
Fiscal Year	Choose a fiscal year to display the report.
Schedule	Allows you to send the report via email to recipients on the fixed date.

Budget Reports

This widget helps you toview the information in the form of following reports that includes total allocation, total spending, forecasted spending, etc.

- Budget Overspenders
- Budget Underspenders
- Budget By Cloud
- Budget By Cost Group Type

Budget Overspenders

Budget Overspenders

- Overview
- Description

Overview

The**Budget Overspenders**report displays information about all categories clouds accounts, cost group types, departments that have exceeded the allocated budget as of date and the forecasted spending at the end of the fiscal year.

Description

Thefollowing is a sample screenshot of the **Budget Overspenders** report.

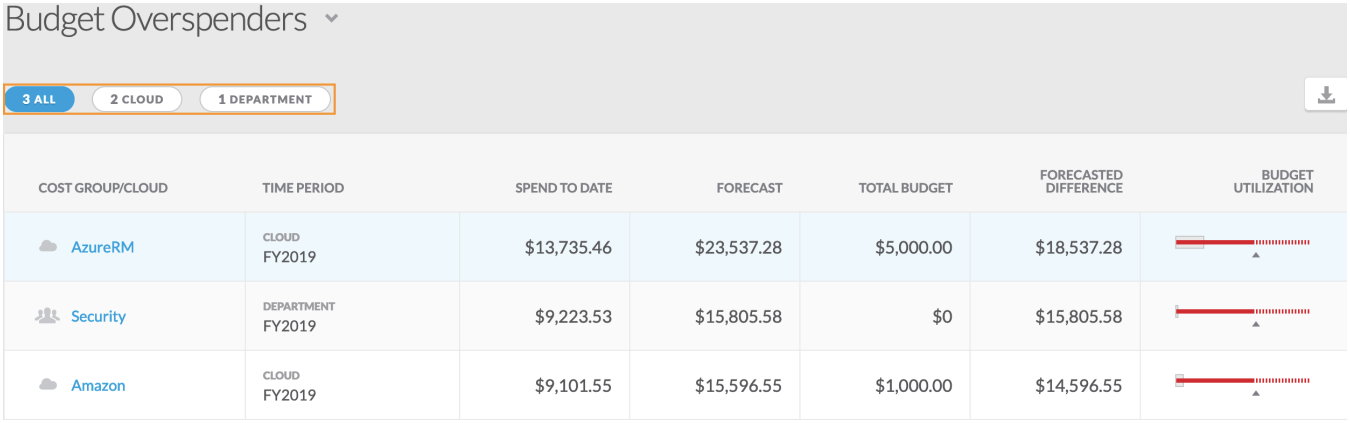

The information in the following table applies to the summary displayed at the top of the report.

Summary	Description
Total	Total number of items in the report that have exceeded the allocated budget.
Department/Cost Group Type	Number of departments or Cost Group Types that have exceeded the allocated budget.
Cloud	Number of cloud accounts that have exceeded the allocated budget.

The following table identifies various aspects of the report:

Identity	Description
Logo	Displays the cloud logo.
Cost Group /Cloud link	Displays the cloud, cost group type, or department name as a link. Click the link to open the budget report for the item. For instance, clicking on a cloud link opens the budget spending for the cloud.
Name	Fiscal year of the budget and the category to which the fiscal year is applied, cloud account, cost group type, or department.
Spend to Date	Amount spent as of date in the fiscal year.
Forecast	Based on the current spending, amount that will be spent for the remaining fiscal year.
Total Budget	Allocated budget for the category in the fiscal year.
Forecasted Difference	The difference amount for the fiscal year between the total budget and the forecasted spending amount.
Budget Utilization	Color-coded progress bar. The gray shaded box over the progress bar indicates the actual budget and the arrow indicates the budget utilization to date. • Green Spend to date and forecasted spend is within the budget. • Orange Spend to date is within budget but forecasted spend exceeds the budget. • Red Spend to date and forecasted spend exceed the budget.

Budget Underspenders

Budget Underspenders

- Overview
- Description

Overview

The**Budget Underspenders**report displays information about the clouds and cost group types that have not spent the allocated budget as of the date or if the forecasted budget is less than the total budget.

Description

Thefollowing is a sample screenshot of the**Budget Underspenders**report.

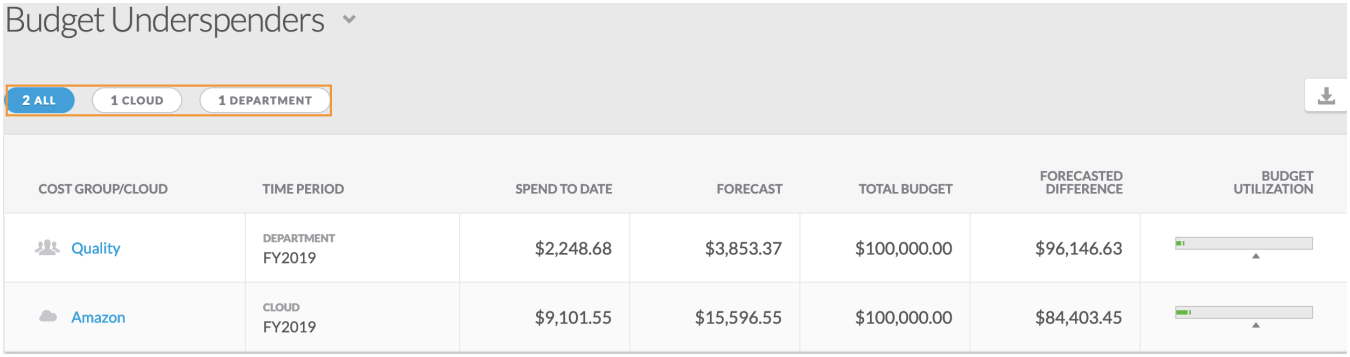

The information in the following table applies to the summary displayed at the top of the report.

Summary	Description
All	Total number of items in the report that have spent less than the allocated budget.
Cloud	Number of cloud accountsthat have spent less than the allocated budget.
Department	Number ofdepartmentsthat have spent less than the allocated budget.

The following table identifies various aspects of the report:

Identity	Description
Logo	Displays the cloud logo.
Cost Group /Cloud link	Displays the cloud, cost group type, or department name as a link. Click the link to open the budget report for the item. For instance, clicking on a cloud link opens the budget spending for the cloud.
Name	Fiscal year of the budget and the category to which the fiscal year is applied, cloud account, cost group type, or department.
Spend to Date	Amount spent as of date in the fiscal year.
Forecast	Based on the current spending, the amount that will be spent on the remaining fiscal year.
Total Budget	Allocated budget for the category in the fiscal year.
Forecasted Difference	The difference amount for the fiscal year between the total budget and the forecasted spending amount.
Budget Utilization	Color-coded progress bar. The gray shaded box over the progress bar t indicates the actual budget and the arrow indicates the budget utilization to date. • Green Spend to date and forecasted spend is within the budget. • Orange Spend to date is within budget but forecasted spend exceeds the budget. • Red Spend to date and forecasted spend exceed the budget.\

Budget By Cloud

Budget By Cloud

- Overview
- Description

Overview

The**Budget by Cloud**report displays information about the clouds providers that have not spent the allocated budget as of date.

Description

Thefollowing is a sample screenshot of the**Budget by Cloud**report. The report displays the total budget and budget for each cloud. The report also displays in different shades the budget allocated to each billing unit in the cloud.

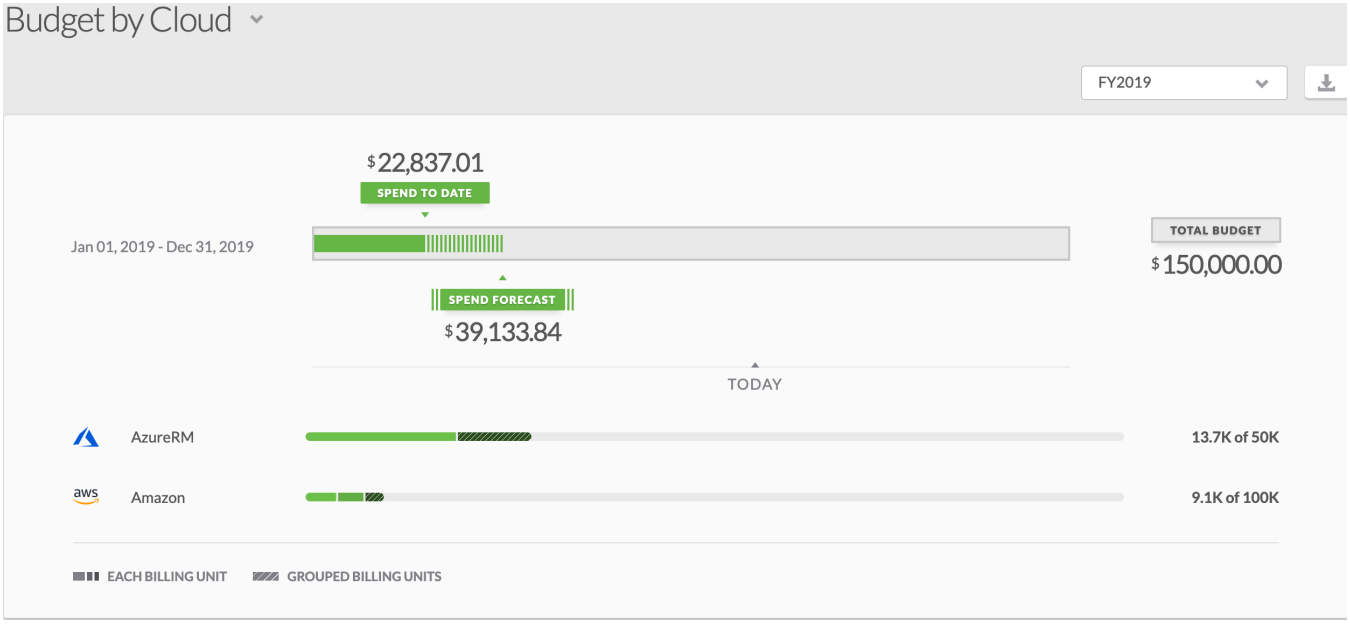

The legend for the above screenshot is as follows:

- Grey box Total budget
- Deep color (in the graph) Actual spending until the day the report was generated.
- Boxed color (in the graph) Forecasted budget

CLOUD	SPEND TO DATE	FORECAST	TOTAL BUDGET	REMAINING BUDGET
AzureRM	**$13,735.46**	$23,537.28	$50,000.00	$36,264.54
cbaba14b-e672-47d7-bb59-4a0613d6d149 (Pay-As-You-Go(Converted to EA))	$9,223.53			
0f2c89bc-0aa4-41f6-838b-3dcbcd17c166 (Microsoft Azure Enterprise)	$4,511.92			
Amazon	**$9,101.55**	$15,596.55	$100,000.00	$90,898.45
804685808463 (CloudCenterMaster)	$3,741.40			
512570952472 (CPSG)	$3,111.47			
052904131800 (Sriram Gopalan)	$2,248.68			

The following table identifies various aspects of the report:

Identity	Description
Cloud	Displays the name. Click the arrow next to the Cloud Name to display the billing units in the cloud.
Spend to Date	Amount spent by the cloud or billing unit as of date in the fiscal year.
Forecast	Based on the current spending, the amount that will be spent on the remaining fiscal year.
Total Budget	Allocated budget for the cloud in the fiscal year.
Remaining budget	Amount that would remain in the total budget at the end of the fiscal year, based on the current spending.

Budget By Cost Group Type

Budget By Cost Group Type

- Overview
- Description

Overview

The**Budget by Cost Group Type**report displays information about budgets allocated to the cost group type for a fiscal year and the spending of the allocated budget for the fiscal year.

Description

Thefollowing is a sample screenshot of the**Budget by Cost Group Type**report.

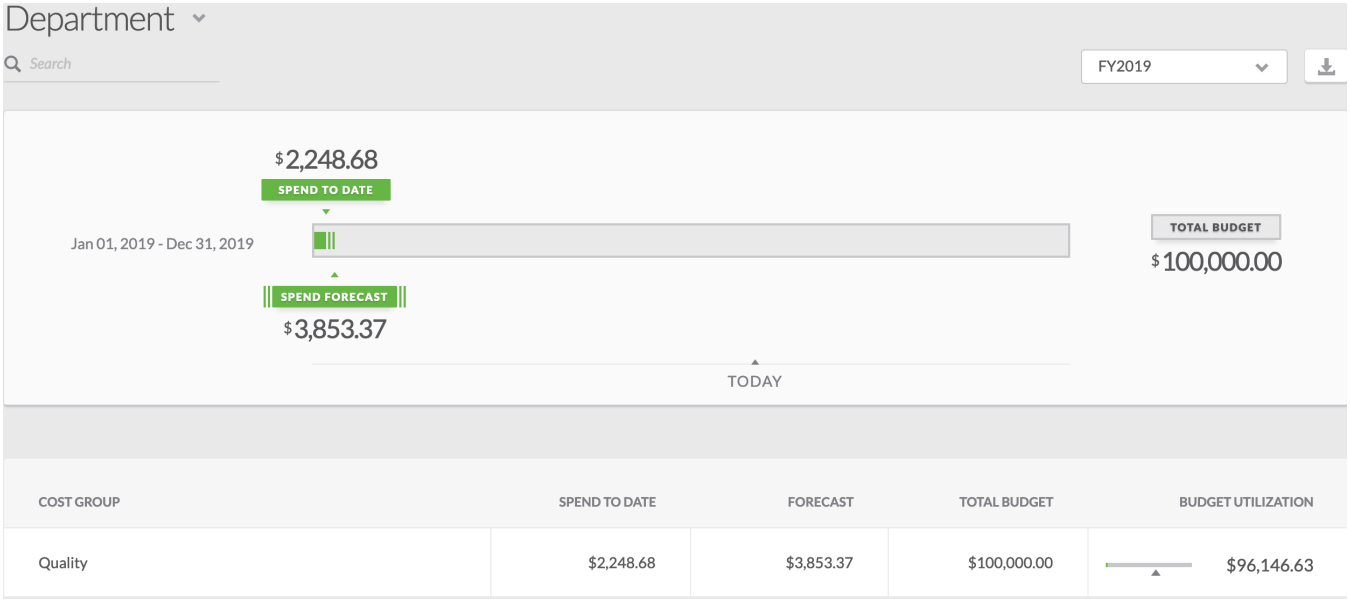

The following table identifies various aspects of the report:

Identity	Description
Cost Group	Displays the cost group type.
Spend to Date	Amount spent as of date in the fiscal year or quarter.
Forecast	Based on the current spending, the amount that will be spent on the remaining fiscal year or quarter.
Total Budget	Allocated budget for the category in the fiscal year or quarter.
Budget Utilization	Color-coded progress bar. The gray shaded box over the progress bar indicates the actual budget and the arrow indicates thebudget utilization to date. • Green Spend to date and forecasted spend is within the budget. • Orange Spend to date is within budget but forecasted spend exceeds the budget. • Red Spend to date and forecasted spend exceed the budget.

Inventory

- Inventory Overview
- Virtual Machines
- Kubernetes Workloads
- Storage Volumes
- Services
- Inventory States

Inventory Overview

Inventory Overview

- Introduction
- What's in the Inventory Pages?
- Filter
    - Advanced Options
        - Saving Filters
        - Scheduling Reports
- Inventory Types

Introduction

The**Inventory**page lists resources running onall cloudaccounts available in Cost Optimizer. A resource is a generic collection that includes instance, storage, load balancer, and database instance details.Inventory is collected for all the combinations of cloud regions and accounts at specified intervals.See Data Collectionfor details on inventory processes and their intervals.

> ⚠ Custom instance prices are not displayed as part of type metadata.

What's in the Inventory Pages?

Click**Inventory**in the left tree pane to open the**Inventory**page.The following table explains the icons in the**Inventory**UI for each of the above categories.See UI Behaviorfor details onicons in the UI.

Icon	Description
Filter	Allows you to filter data and view inventory data for one or more of the following: • Cloud Regions • Status (of resources) • CPUs • Memory GB • Billing Unit • Tags
Sort	Sort the items in the page.
Find	Find an instance of an inventory type based on specific keywords.

Filter

The**Filter**panel allows you to filter data based on a set of options, thereby allowing you to drill down to the exact details that you require.

Advanced Options

The advanced options in Cost Optimizer are as follows:

- Saving Filters
- Scheduling Reports

Saving Filters

You can choose to save a combination of options in the **Filter** menu for future use through the **Save Filters** feature so that you can quickly access and use the filter at a later time. To save a filter, do the following:

1. Choose the required filter options in the **Filter Panel** pane.

2. The**Save** button appears. The **Save New Filter** dialog appears.

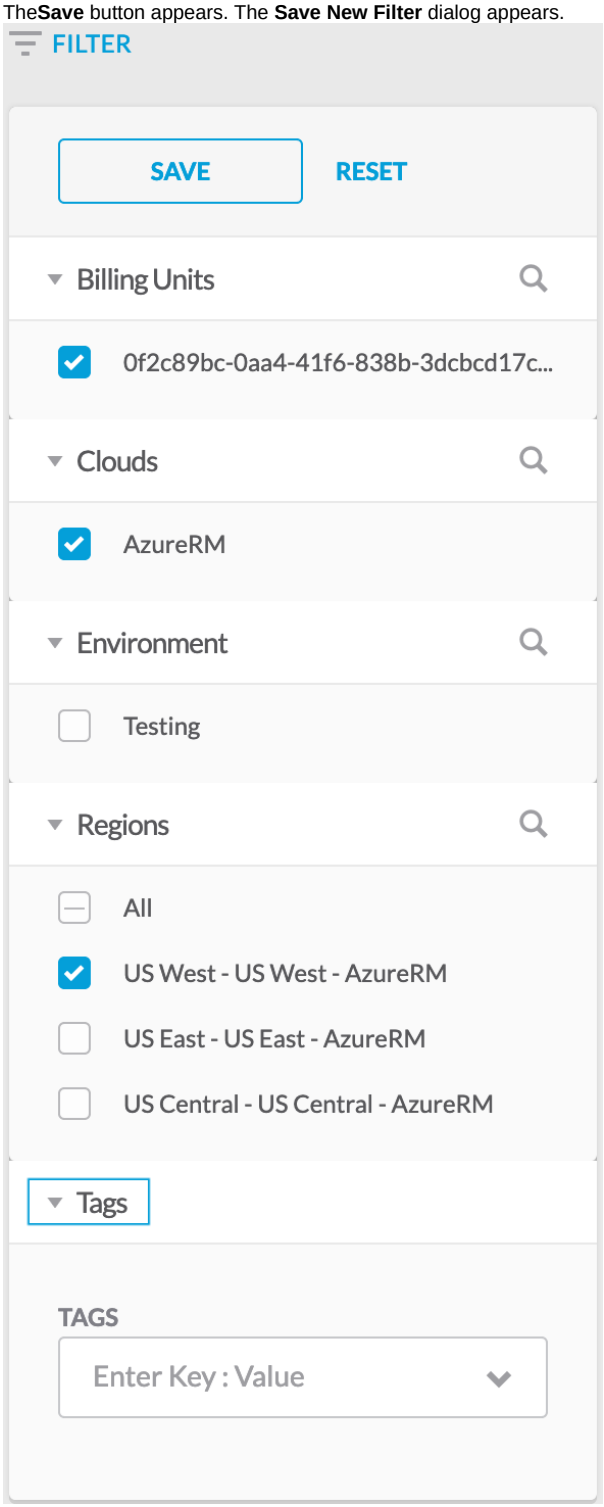

3. Specify a name for this filter and click**Save**. A status message appears indicating that the filter has been saved.

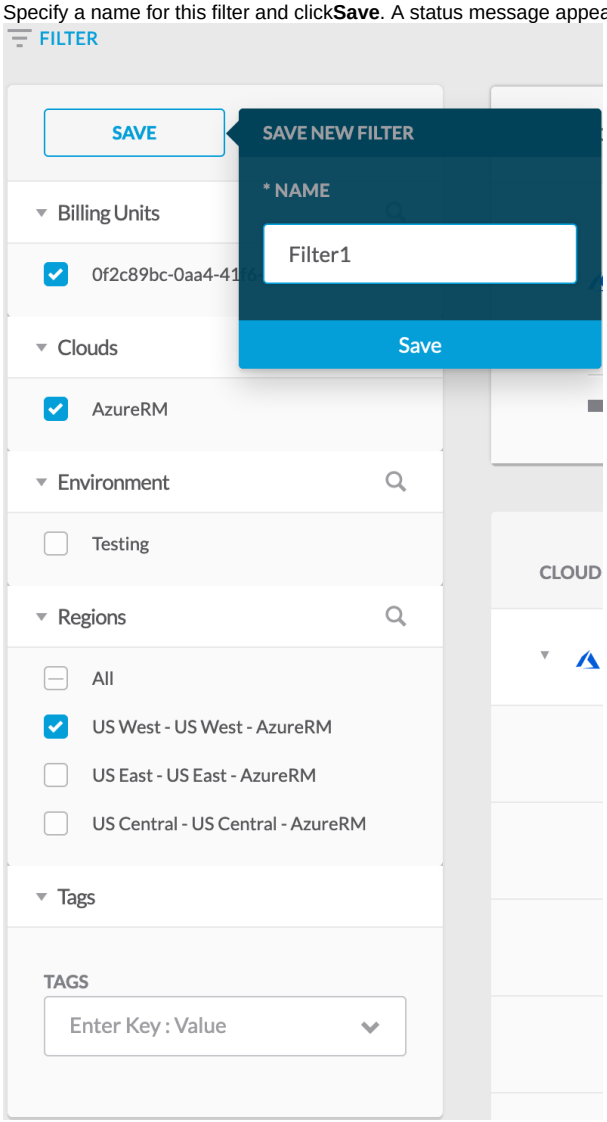

4. You can access and view the saved filters from the dropdown list.

You can also perform the following additional tasks in the **Filter** menu:

- Mark the filter as a favorite by clicking the pin icon next to the filter name.
- Remove the chosen filters by choosing the**Reset**button at any point when saving the filter.
- Delete the saved filter by clicking the **Trash** icon next to a saved filter name. Click **OK** in the **Delete Saved Filter** dialog to confirm the deletion.

Scheduling Reports

The **Scheduler** icon allows you to schedule report generation periodically on a fixed date or at intervals. To create a schedule, do the following:

1. Click the **Scheduler** icon. The **Schedule New *Report Name*** dialog appears.

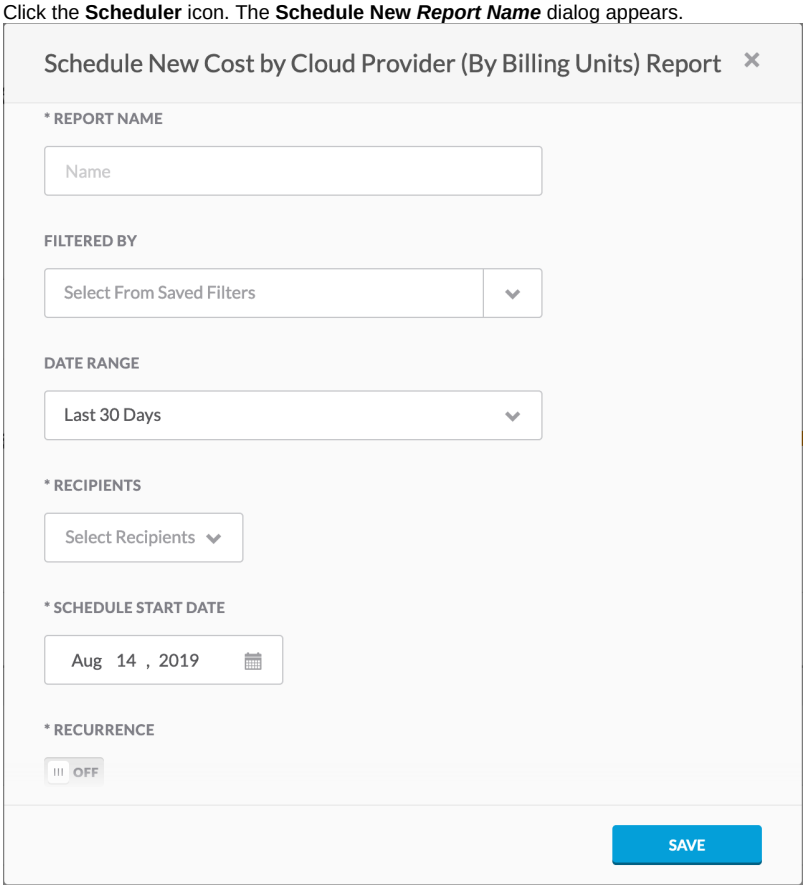

2. Do the following:

   a. Enter a name for the schedule.
   b. Choose filtering options for the schedule from the **Filtered By** field. The information in this field is populated when you save the filtering options as described in the *Advanced Filtering Options* section. You can choose to select a filter or leave the field empty.
   c. Choose the date range.
   d. Select the recipients the report must be sent to.
   e. Specify the start date.
   f. Toggle on the **Recurrence** button to send the report at intervals.
   g. In the **Repeats Every** area, specify the number of times the report must be sent to the recipients and choose the interval **Daily** or **Weekly** . if you choose Weekly, you can also specify the days of the week when the report is sent.
   h. Select the period to end the schedule. The options are:

      i. **Never** Send report forever or until the schedule is deleted.
      ii. **On** Date when the report should be sent.
      iii. **After** Number of occurrences after which the report is not scheduled.

3. Click **Save**. The report is displayed in the**Scheduled*Report Name***dialog as shown in the sample screenshot below.

Optionally, you can use the **Edit** option in the **Actions** column to make changes to the schedule. You can also delete the report using the **Delete** option.

Click **Done** to close the dialog.

The following screenshot displays a sample email format of the report.

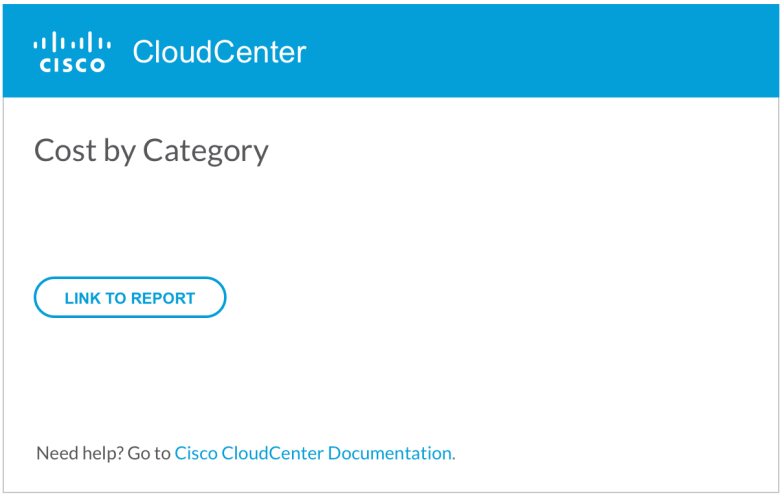

Inventory Types

The inventory types in Cost Optimizer are as follows:

- Virtual Machines
- Kubernetes Workloads
- Storage Volumes
- Services

Virtual Machines

Virtual Machines

- Overview
- Description
- Filter
- VM Details
    - Cost Breakdown

Overview

This pagedisplays the virtual machines for a cloud provider. The following is a sample screenshot of the Virtual Machines page.

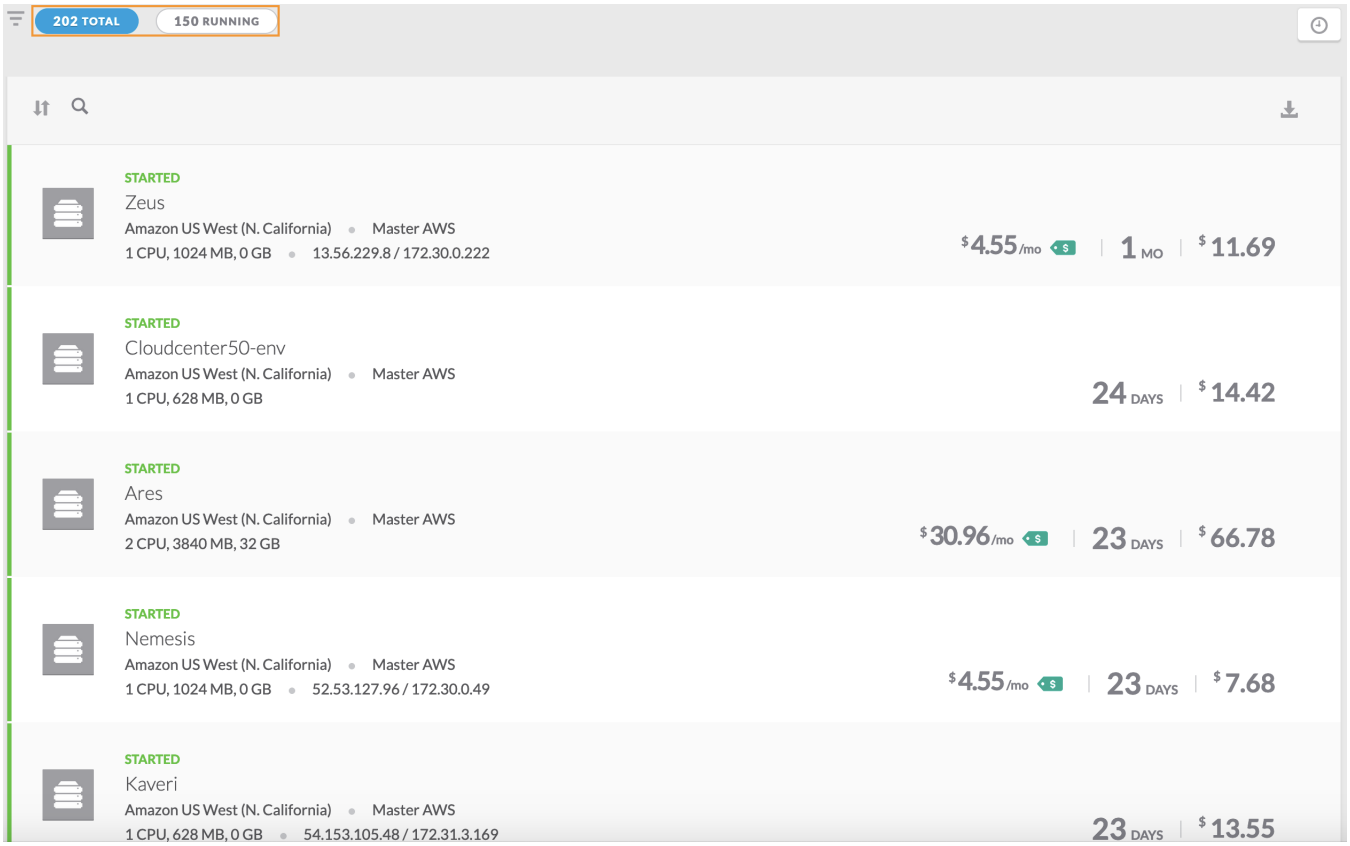

Description

Regardless of the filter settings, the information in the following table applies to the summary displayed at the top of the Virtual Machines page:

Summary	Description
Total	Total number of VMs.
Running	Total number of VMs running (billed) ***without any time restriction***. This count includes VMs that display the***ERROR***status.

The following table identifies various aspects of the Virtual Machines tab:

Identity	Screenshot and Description
Logo	Displays the OS logo.

VM details link	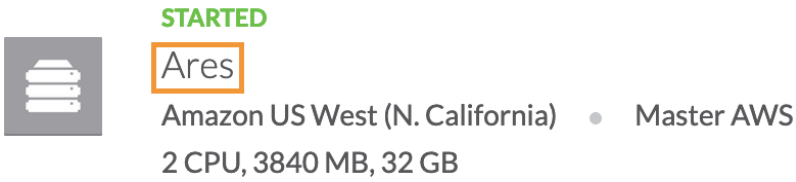 Displays the VM name as a link. Click the link to view details about the VM. For each VM, the following is displayed. • Hostname Thehostnamefor the VM, if configured. Else, thenode IDis displayed. • Cloud Region • Cloud Name • Private and public IP address of the VM.
Status	Color-coded states that identify the VM status. See Inventory Statesfor a complete list and additional details. ⚠ Terminated VMs are not displayed.
Duration	Instance runtime, in hours or minutes, the VM is in the specified state.
Cost	Calculated on on-demand prices available in the instance types. For private clouds, the cost is as decided by the admin when setting up the clouds. See Supported Datacenters and Private Clouds.

Filter

You can filter the items based on the following:

- Clouds
- Cloud Region
- Status
- CPUs
- Memory GB
- Billing Unit
- Tags
- Location
- Department
- Project

VM Details

Clicking the VM Name displays information about the VM which contains the following tabs:

- The**Details**tab (default) provides exhaustive details for the*VM.*
- The**Stats**tab provides information about resizing recommendations.

The following is a sample screenshot of the VM **Details** page.

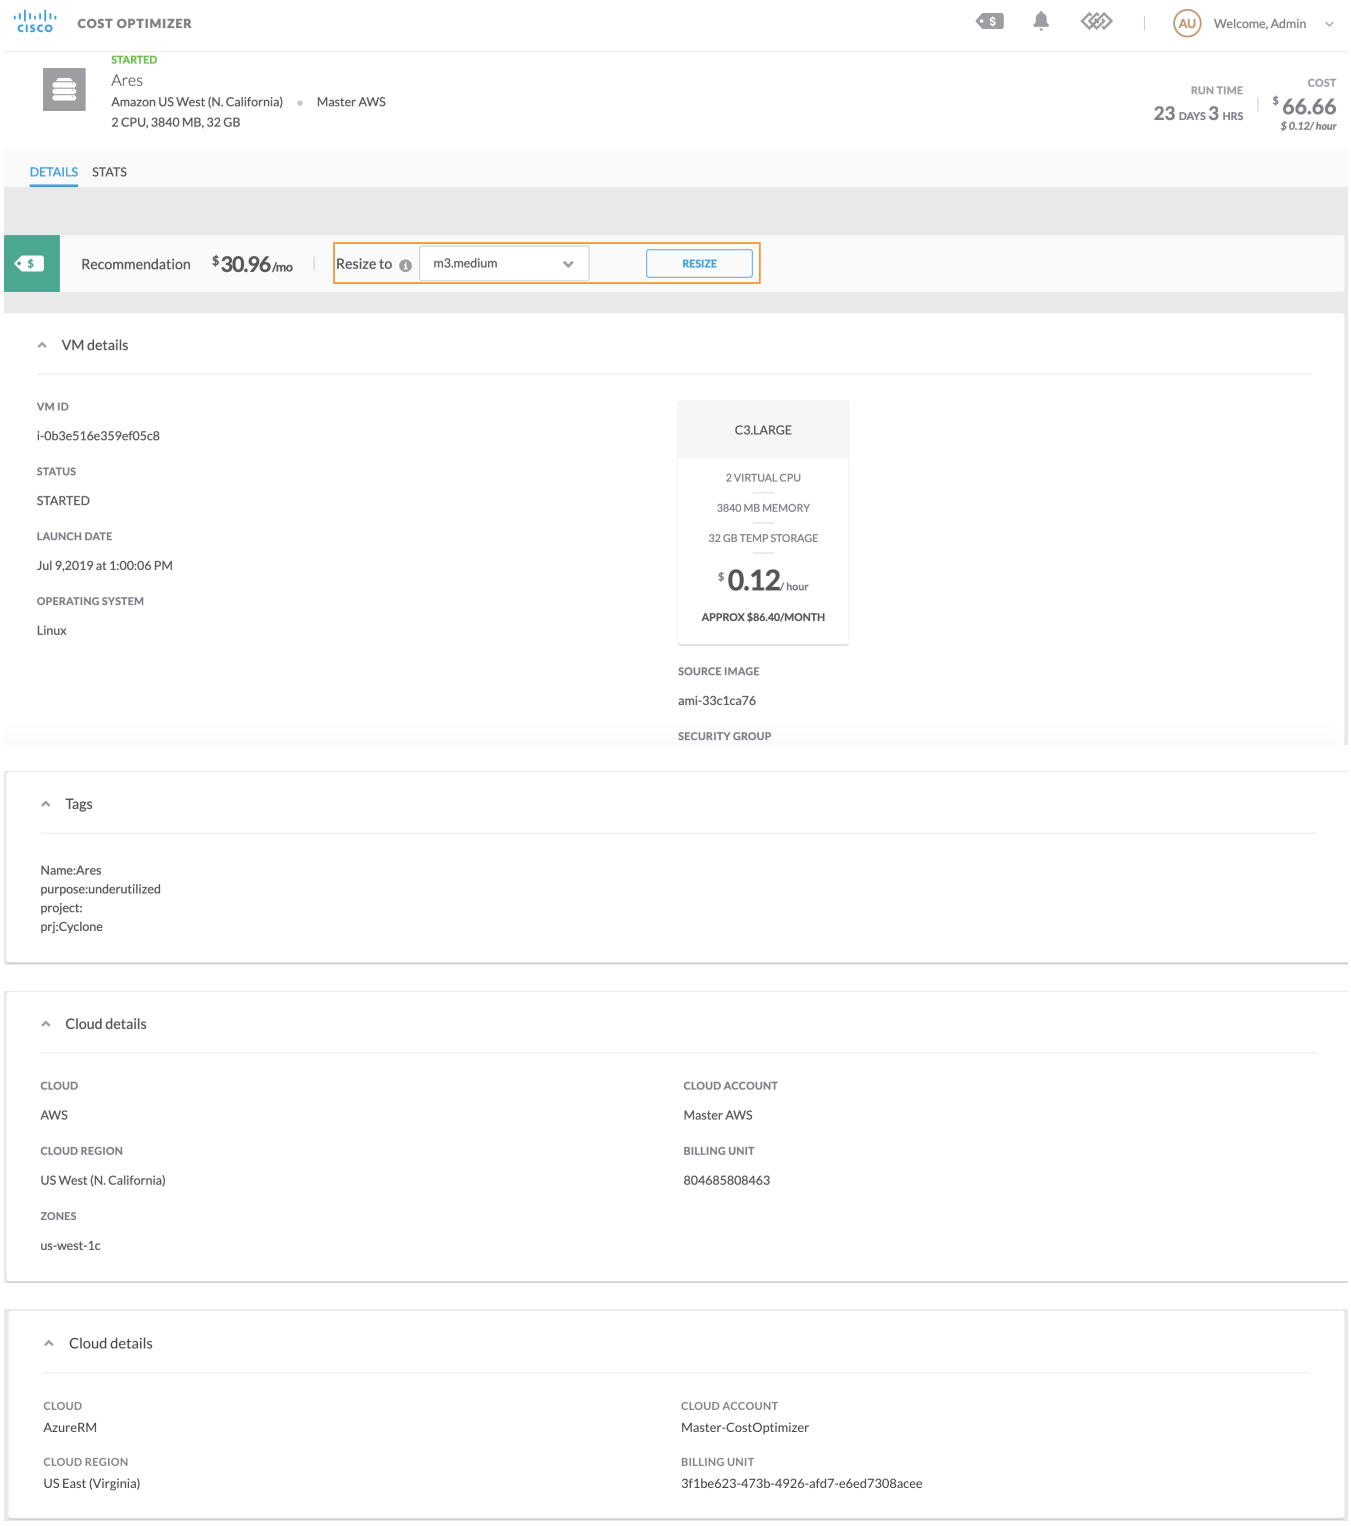
COST OPTIMIZER
Welcome, Admin
STARTED
Ares
Amazon US West (N. California) • Master AWS
2 CPU, 3840 MB, 32 GB
RUN TIME
23 DAYS 3 HRS
COST
$ 66.66
$ 0.12/ hour
DETAILS
STATS
Recommendation
$30.96/mo
Resize to
m3.medium
RESIZE
VM details
VM ID
i-0b3e516e359ef05c8
STATUS
STARTED
LAUNCH DATE
Jul 9,2019 at 1:00:06 PM
OPERATING SYSTEM
Linux
C3.LARGE
2 VIRTUAL CPU
3840 MB MEMORY
32 GB TEMP STORAGE
$ 0.12/ hour
APPROX $86.40/MONTH
SOURCE IMAGE
ami-33c1ca76
SECURITY GROUP
Tags
Name:Ares
purpose:underutilized
project:
prj:Cyclone
Cloud details
CLOUD
AWS
CLOUD ACCOUNT
Master AWS
CLOUD REGION
US West (N. California)
BILLING UNIT
804685808463
ZONES
us-west-1c
Cloud details
CLOUD
AzureRM
CLOUD ACCOUNT
Master-CostOptimizer
CLOUD REGION
US East (Virginia)
BILLING UNIT
3f1be623-473b-4926-afd7-e6ed7308acee

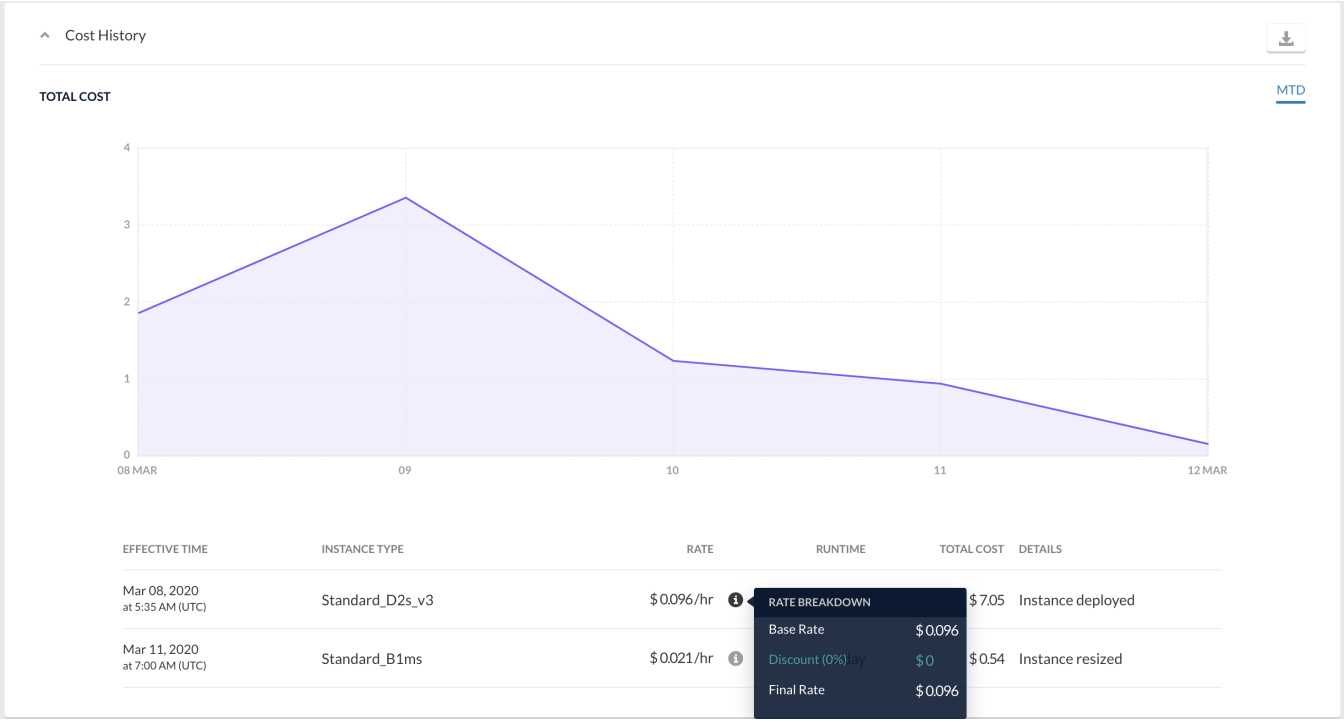

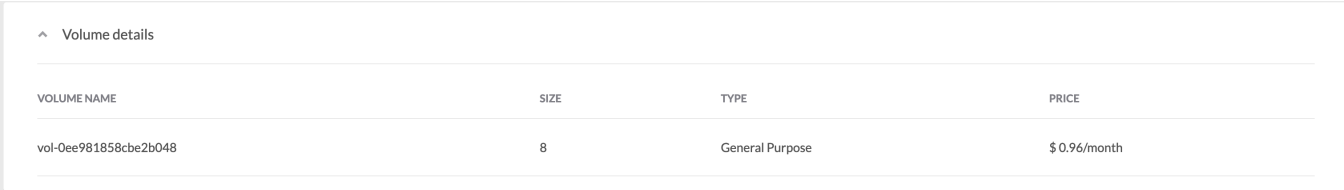

This table identifiessignificant aspects of the **Details** Tab.

Area	Identity	Description
Recommendation	Resize to	Resizing recommendation suggested by the Recommendation Engine (see Rightsizing > *Recommendation Engine*). The resizing recommendation is based on the utilization of this instance and does not affect the performance of the VM.
VM Details	VM ID	Billing unit that owns the VM.
	Status	Status of the VM. The options are as follows: • Started • Stopped • Terminated
	Source Image	Machine from which VM is launched.
	Security Group	Rules that control traffic in or out of a VM.
Tags	Name	Tag name.
	Purpose	Usage of the tag in Cost Optimizer.
	Project	Tag key-value pair.
Cloud Settings	Network	Network name issued by the cloud provider.
Cloud Details		Information about cloud provider the instance pertains to.
Cost History	Total Cost	The total cost applicable to the VM for the period. In this case, the period is month-to-date (MTD).

	Effective Time	Time when the instance was started.
	Rate	Rate, per hour, applicable to the instance includes base rate and discounts, that may be applicable. See *Cost Breakdown* section below.
	Runtime	Duration, in days, since the instance was started.
	Total Cost	Cost incurred since the instance was started.
	Details	
Volume Details	Type	Types of storage volume, varies for each cloud, for example, General Purpose, Provisioned IOPS, Throughput Optimized.
	Price	Price of volume per hour.

Cost Breakdown

Rates or prices for each resource displayed in the Inventory pages are synchronized through the Metadata Sync background process (seeData Collection)f or public clouds. In the case of private clouds, the rates areentered manually in theInstance Types Settingspage. These rates help in calculating the expenses incurred by the resources. This expense is called*Resource Cost*and is computed based on resource usage.Resource costs are based on the rate card (price) and usage (runtime use) of the resource.For private clouds, the resource costs are displayed as invoice costs (displayed in the dashboard and reports in Cost Optimizer. Discounts, if applicable, are included when computing the resource costs.

Cloud providers may revise resource costs and hence, pricing and discounts vary over time.The Resource Cost Breakdown feature provides information about the cost itemization, such as price and applicable discounts on an hourly basis, which add up to the cost for a resource at a given time. The Resource Cost Breakdown feature allows you to track the following:

- Pricing and discount rates over a period of time
- Rightsizing action history

The Cost History area in the Details page for an Inventory displays the resource cost. The following is a sample screenshot of Cost Breakdown.

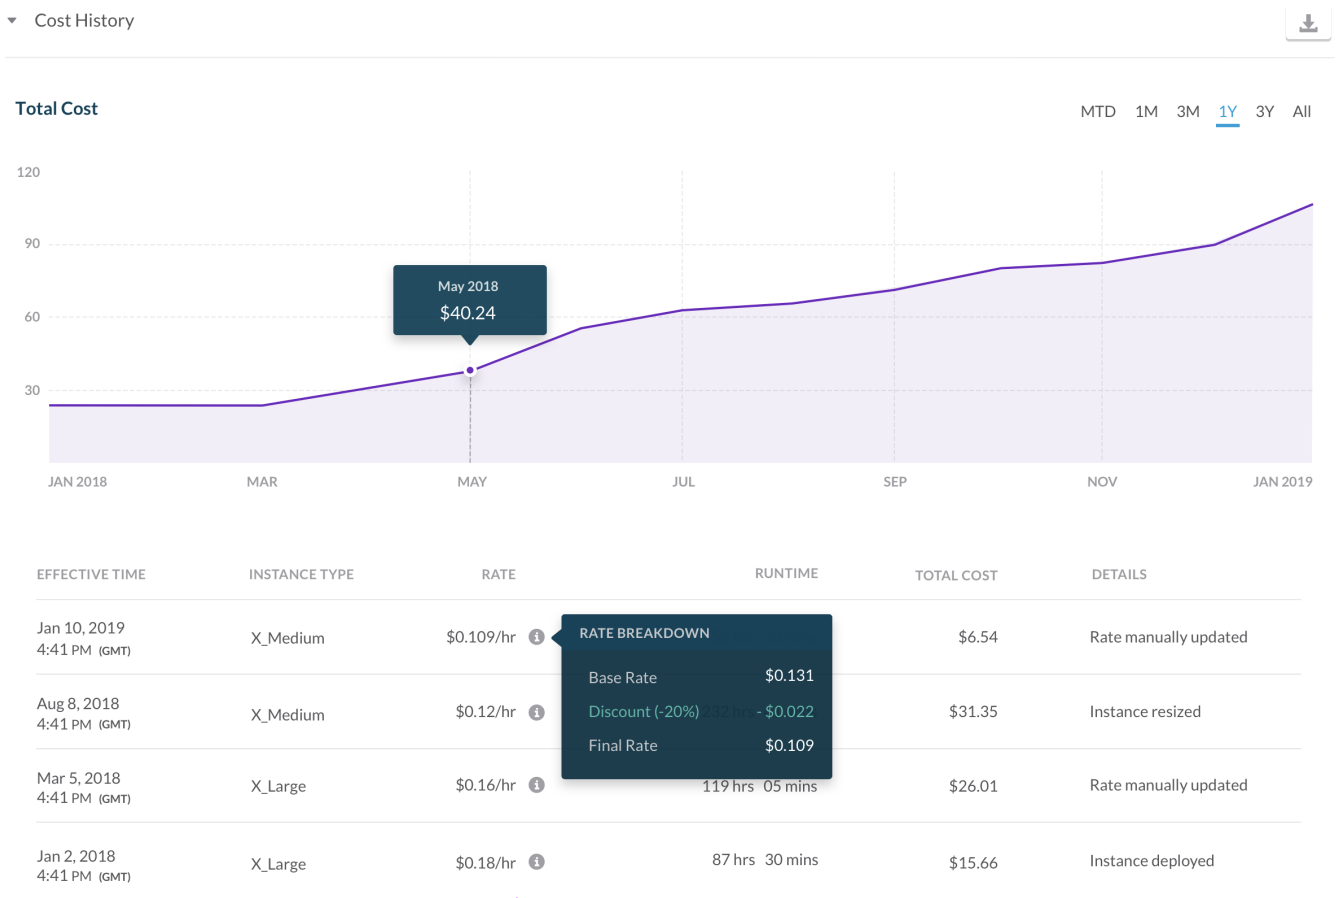

This table identifiessignificant aspects of the **Stats** Tab.A sample screenshot is shown below.

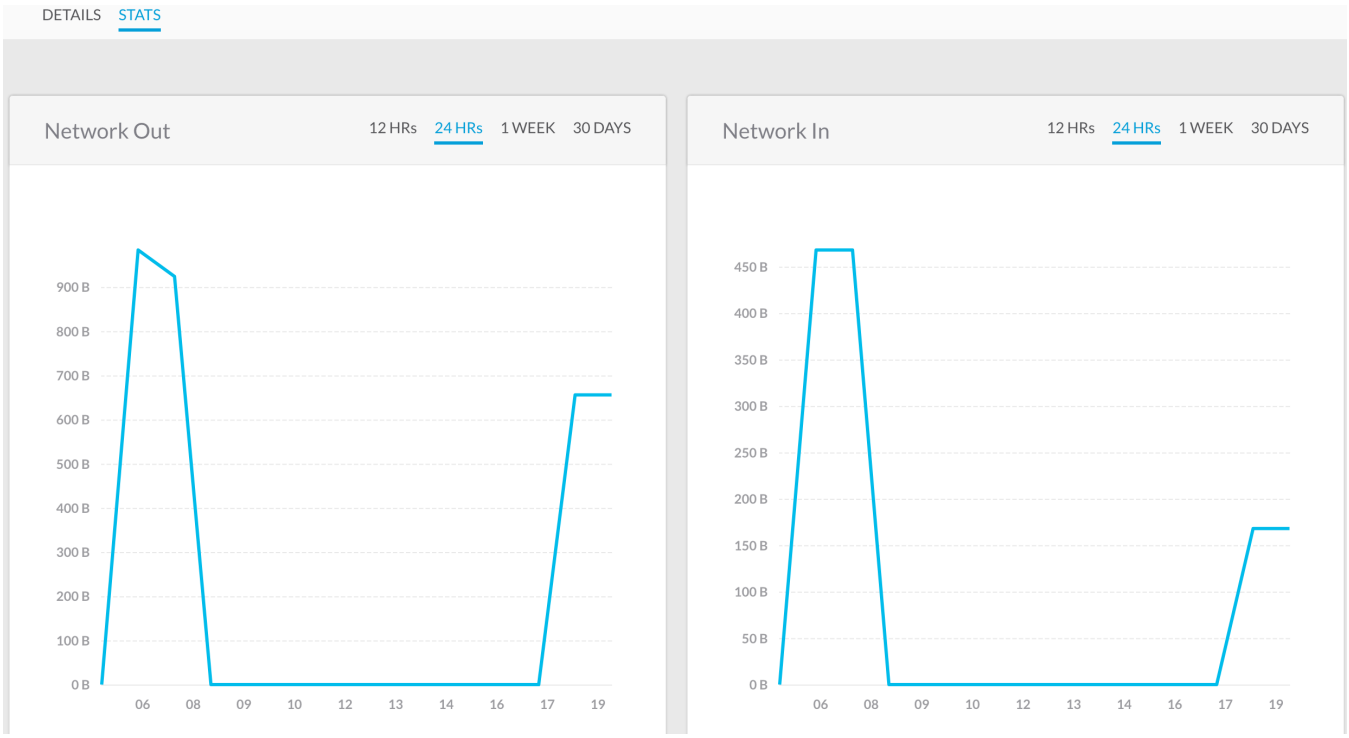

Identity	Description
Network Out	Outbound traffic in kilobytes.
Network In	Inbound traffic in kilobytes.
CPU	CPU utilization in percent.
Disk Write	Volume of data written to a disk in kilobytes.
Disk Read	Volume of data read from a disk in kilobytes.

Kubernetes Workloads

Kubernetes Workloads

- Overview
- Description
- Filter
- Kubernetes Workloads Details

Overview

To display information about Kubernetes Workloads items, choose**Kubernetes Workloads**in the**Inventory**header drop-down list.

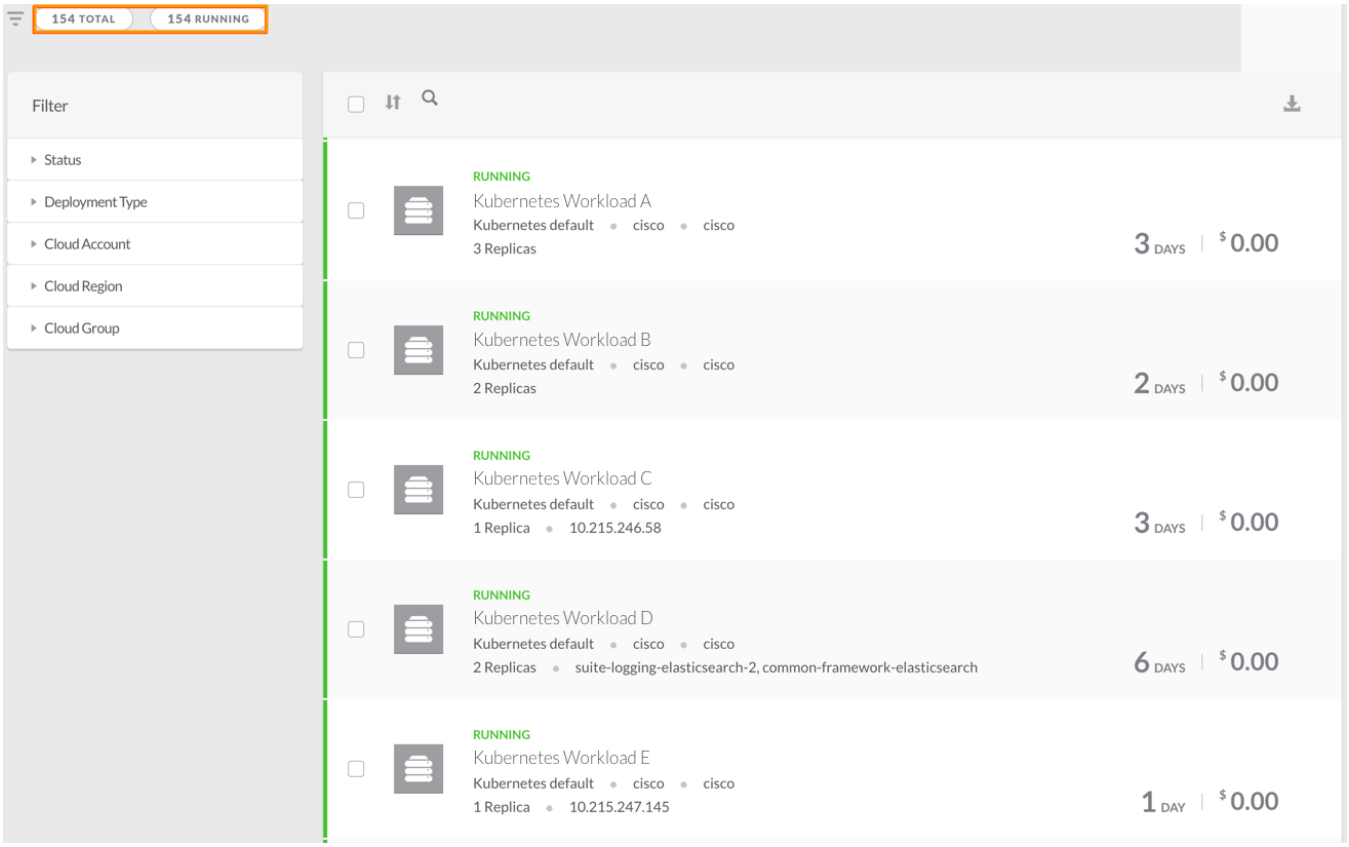

Description

Regardless of the filter settings, the information in the following table applies to the summary displayed at the top of the page:

Summary	Description
Total	Total number of Kubernetes workloads.
Running	Total number of running (billed) Kubernetes workloads***without any time restriction.***

The following table identifies various aspects of the Kubernetes Workloads tab:

Identity	Screenshot and Description

Kubernetes details link	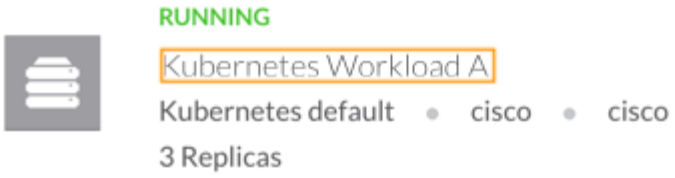 Displays the Kubernetes Workload name as a link. Click the link to view details about the Kubernetes Workload. For each workload, the following is displayed. • Hostname Thehostnamefor the workload, if configured, else thenode IDis displayed • Cloud Region • Billing Unit • Number of ReplicasPods that are running
Status	Color-coded status that identifies the Kubernetes workload status. SeeInventory Statesfor a complete list and additional details.
Duration	Runtime of the Kubernetes workload (hours or minutes).
Cost	Calculated on on-demand prices available in the instance types. For private clouds, the cost is as decided by the administrator when setting up the clouds. SeeSupported Datacenters and Private Clouds.

Filter

You can filter the items based on the following:

- Status
- Deployment Type
- Cloud Account
- Cloud Region
- Cloud Group

Kubernetes Workloads Details

Clicking the Kubernetes Workload name displays the following tabs:

- The**Details**tab (default), which provide exhaustive details for theworkload.
- The**Replicas**tab, which provides information about pod replicas.

The following is a sample screenshot of the Kubernetes**Details**tab.

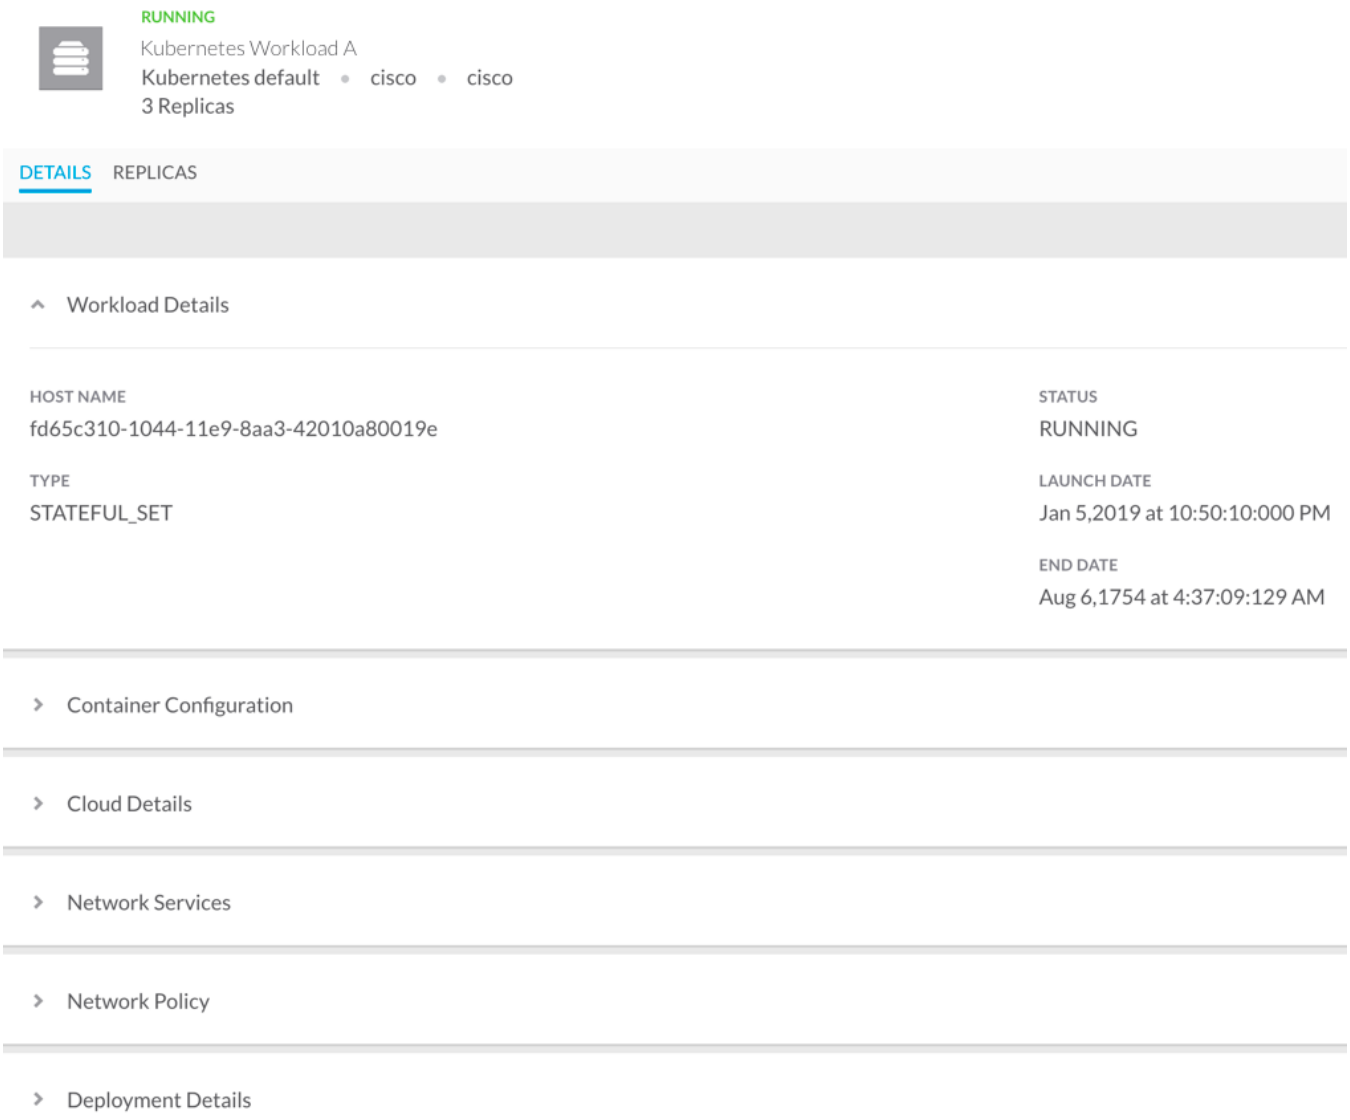

The following table explains significantitemsin the tab.

Area	Identity	Description
Workload Details	Type	Could be one of the following • Deployment • StatefulSet • DaemonSet
	Status	Workload status. The options are: • Failed • Pending • Running • Succeeded • Terminated
Container Configuration	Port/Protocol	Port and protocol to establish a connection.
	Source Namespace	Source cluster in the Kubernetes workload.
Network Services	Type	Type of IP Address assigned to the workload.
	ClusterIP	Unique internal IP address assigned to a service.

The **Deployment Details** area is not applicable to Cost Optimizer.

The following is a sample screenshot of the **Replica** tab.

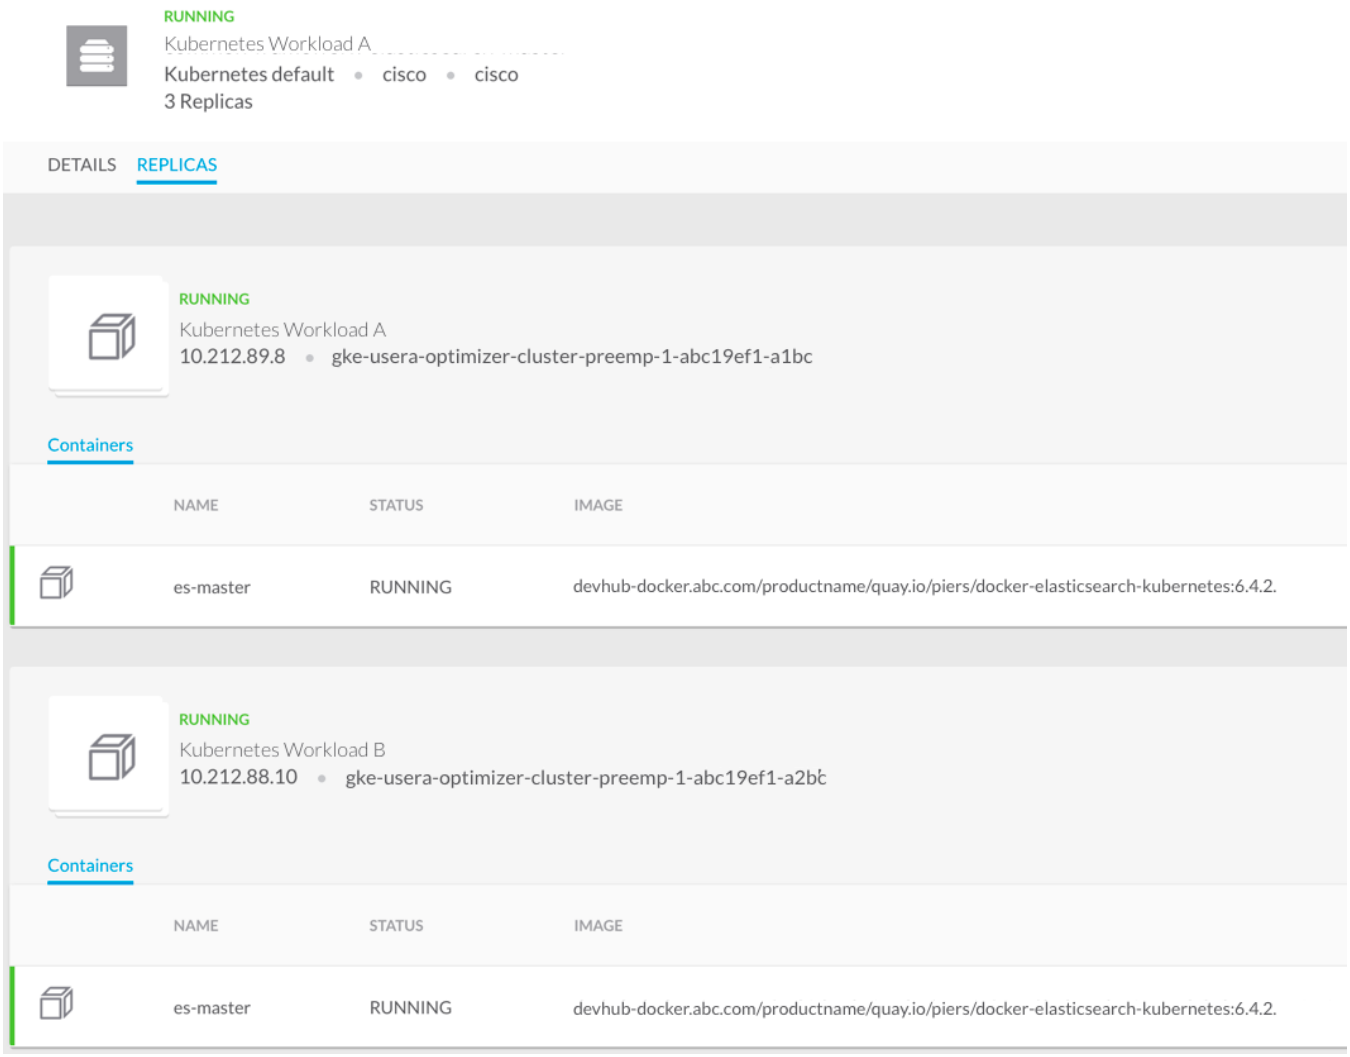

The following table explains significant aspects of the tab.

Identity	Description
Replica IP	Private IP address of the cluster.
Image	Location of the image for the Kubernetes cluster from Docker registry.

Storage Volumes

Storage Volumes

- Overview
- Filter
- Details Page

Overview

Astorage volumeis a virtual disk that provides persistent block storage space for instances. You can use storage volumes to store data and applications.

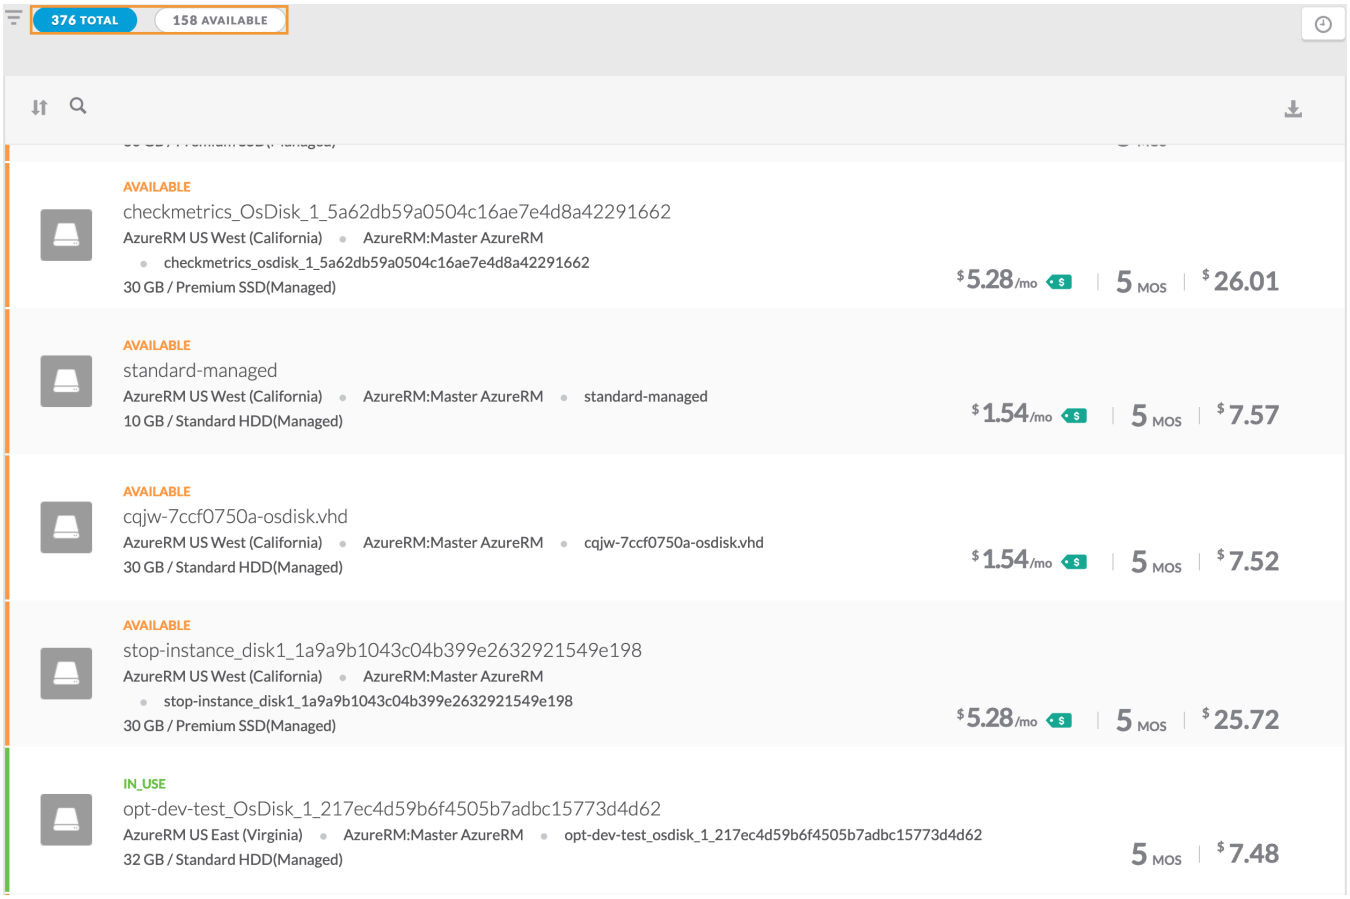

The following table identifies various aspects in the Storage Volumes page:

Identity	Screenshot and Description
Total	Total number of available storage spaces.
Running	Total number of running (billed) storage***without any time restriction.***
Status	Color-coded status that identifies the Storage Volume state. The status could be one of the following: • Available • In Use See Inventory States for a complete list and additional details.

Storage Volume link	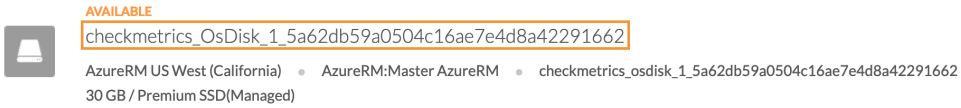 Displays the storage name as a link. Click the link to view additional information. For each volume, the following is displayed. • Storage Name • Cloud Region • Cloud Account • Volume ID • Storage space and type, for example, PD-standard, Standard Persistent Disk
Duration	Hours or minutes, the storage is in the specified state.
Cost	Cost calculated on on-demand prices.

Filter

You can filter the items based on the following:

- Clouds
- Cloud Region
- Status
- Billing Unit
- Tags
- Location
- Department
- Project

Details Page

Click the storage name link to open the Storage Details page. The following is a sample screenshot of the Storage Details page.

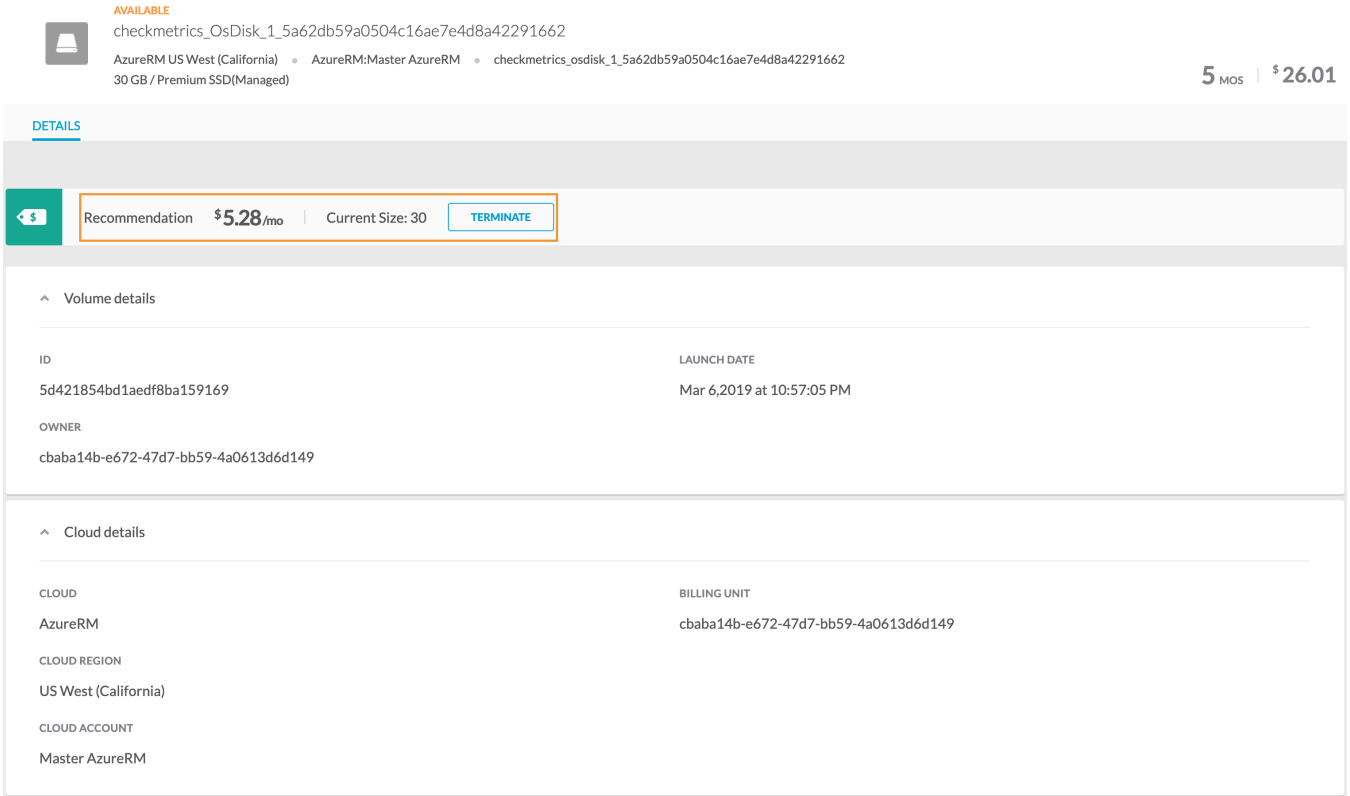

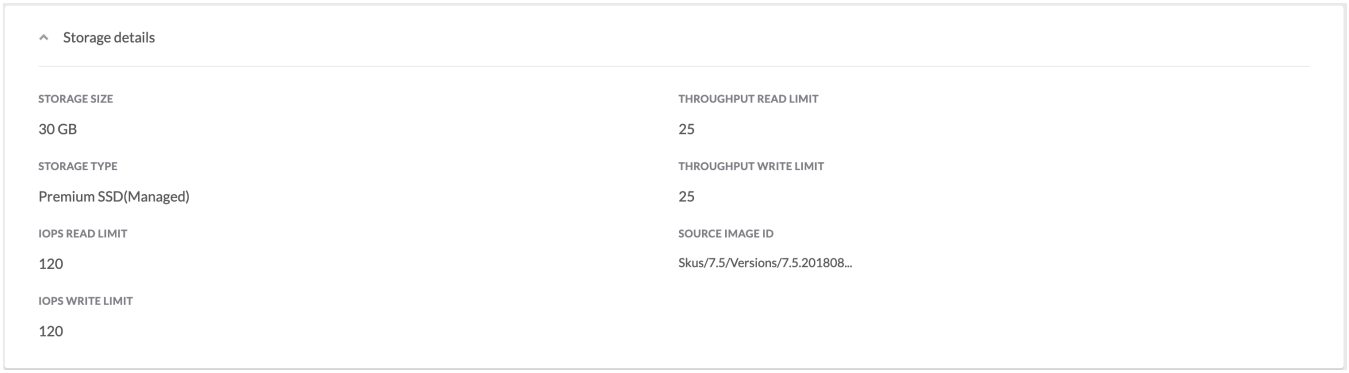

This table identifiessignificant aspects of the**Details**Tab.

Area	Identity	Description
Recommendation	Current Size	Potential savings when the recommendation is implemented.
Volume Details	ID	ID assigned by a cloud provider.
	Owner	Billing unit that owns the VM.
Storage Details	Size	Size of volume in GB.
	Type	Types of storage volume, varies for each cloud, for example,General Purpose, Provisioned IOPS, Throughput Optimized.
	IOPS Read Limit	Maximum IO (input or output) read operations per second.
	IOPS Write Limit	Maximum IO (input or output) write operations per second.
	Throughput Read Limit	Maximum data transfer rate in mebibyte (MiB) per second for read operation.
	Throughput Write Limit	Maximum data transfer rate in mebibyte (MiB) per second for write operation.
	Source Image Snapshot ID	Snapshot from which the volume was created.

Services

Services

- Overview
- Description
- Filter
- Details Page

Overview

A cloud provider offers services such as load balancer, databases, and so on. This information is displayed when you choose **Services** in the **Inventory** header drop-down list.

Description

The following is a sample screenshot of the **Services** page.

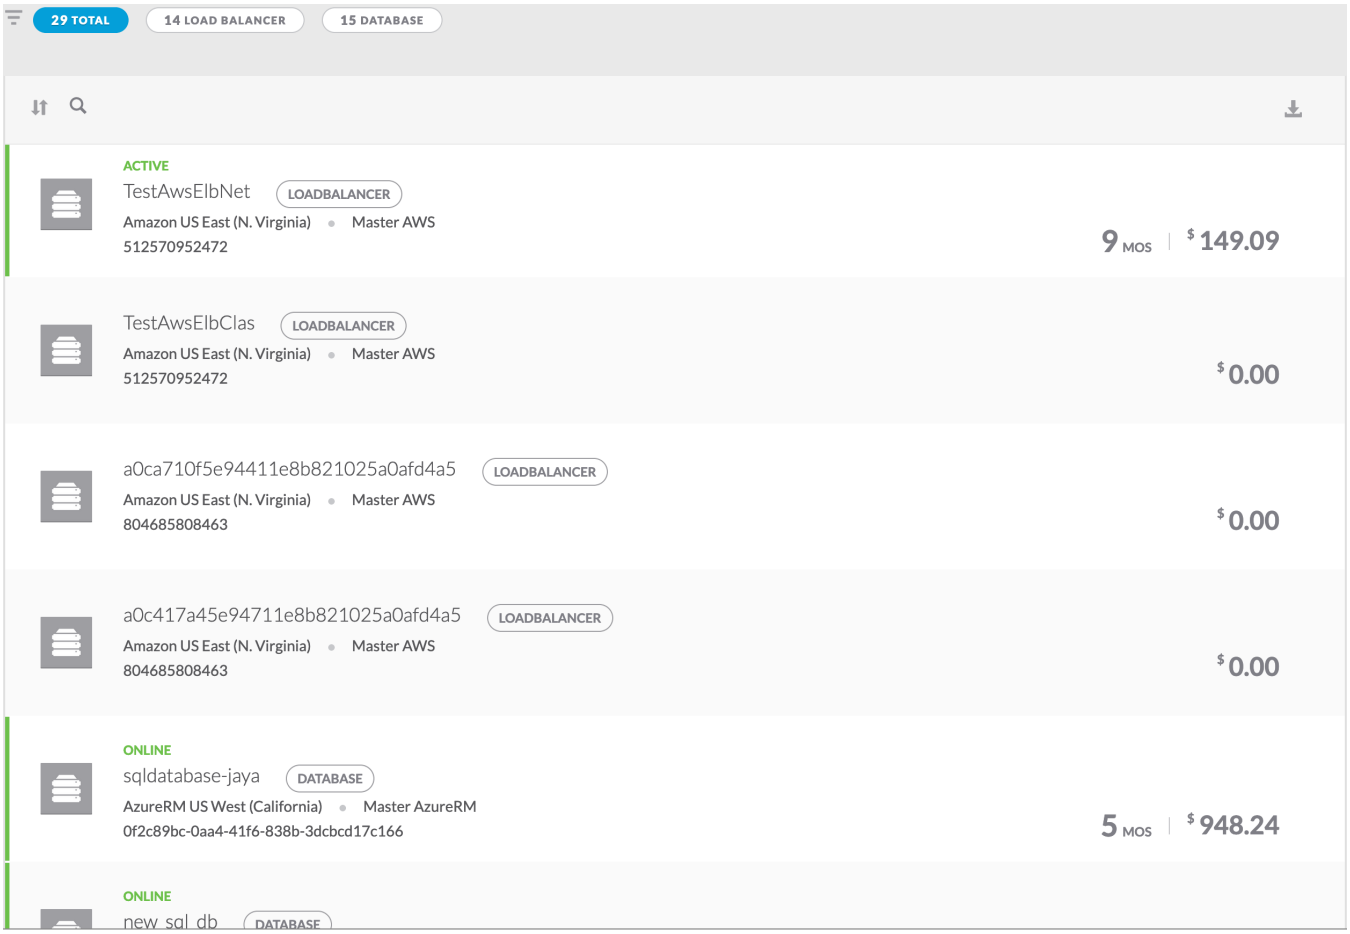

The following table identifies various aspects in the **Services** page:

Identity	Screenshot and Description
Services Header	Identifies the total number of services and the available service types.
Services Status	This could be one of the following: • Active • Available • Healthy • Unhealthy • Terminated

Service Type	Displays the type of service. This could be Loadbalancer or Database.
Services details link	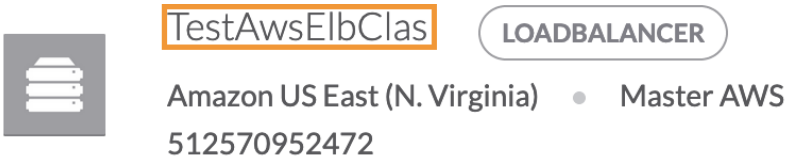 Displays the Services name as a link. Click the link to view additional information. For each service, the following is displayed. • Hostname • Cloud Region • Cloud Account Name • Billing Unit

Filter

You can filter the items based on the following:

- Clouds
- Cloud Region
- Status
- Tags
- Location
- Department
- Project

Details Page

The following is a sample screenshot of the Services Details page.

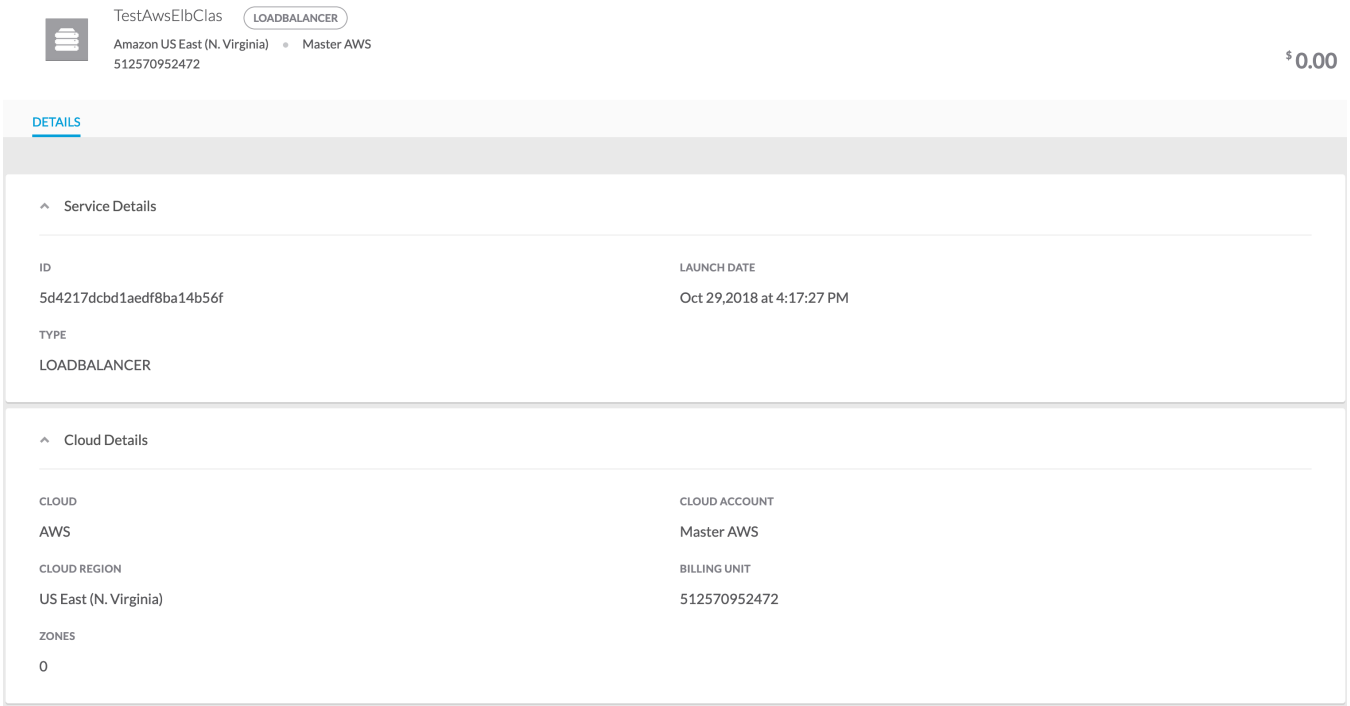

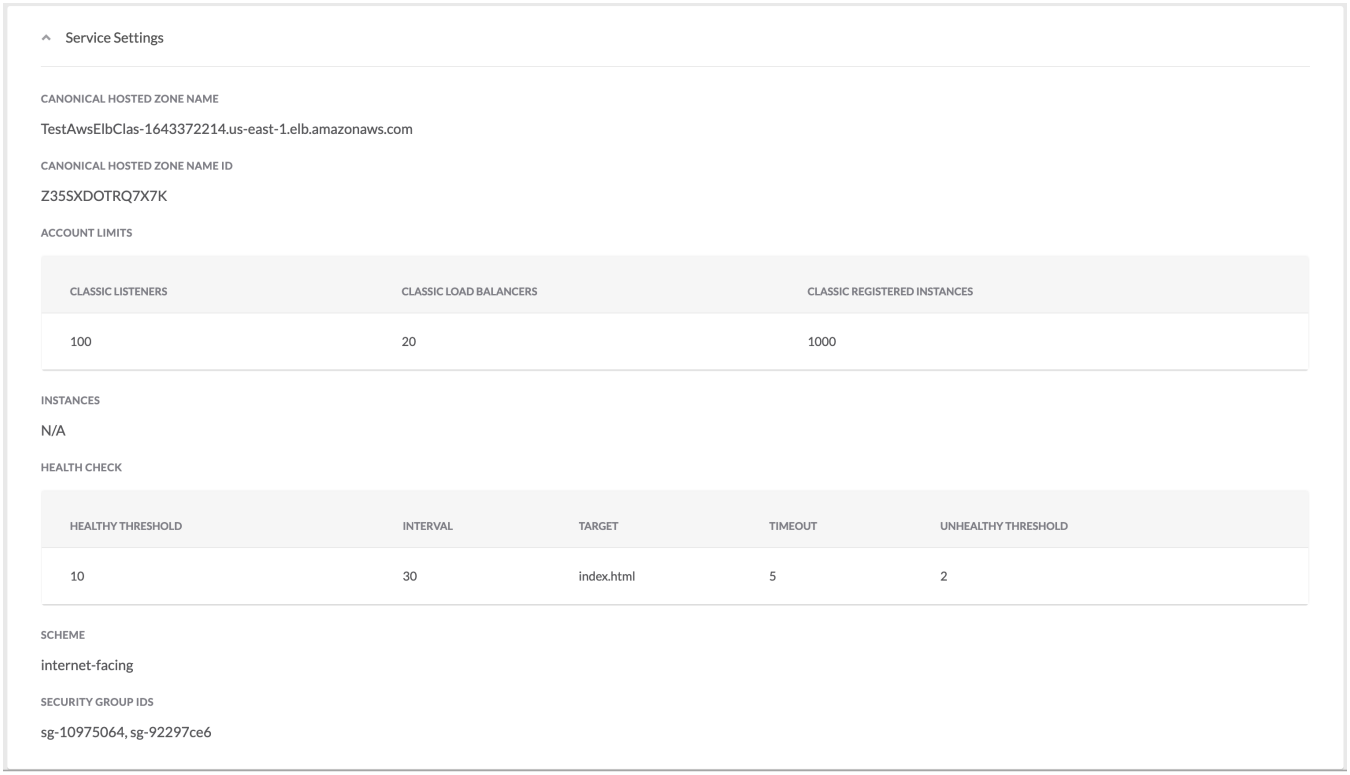

The following table identifiesvarious aspects of thepage.

> ⚠ The **Deployment Details** area does not apply to Cost Optimizer.

Area	Identity	Description
Service Details	ID	ID assigned by a cloud provider.
	Status	Status of the Service (varies for cloud providers). Options include: • Available • Active • Healthy • Unhealthy
	Type	Loadbalancer or database.
Service Settings	DNS Name	DNS name assigned by the cloud provider.
	Health Probes	Periodic requests send to the check the instance health.
	Health Check Status	Options include: • Active • Inactive
	Scheme	Type of Loadbalancer, which could be: • Internal • Internet-facing • Public • Private
	Instances	VMs used for load balancing.
	Security Group IDs	Rules that control traffic in or out of service.
	Account Limits	Loadbalancer resource limits.

	Listener Descriptions	Process that checks for connection requests.

Inventory States

Inventory States

This following table lists the states of inventory resources (Virtual Machines, Storage Volumes, Load Balancers, Database, Containers) in Cost Optimizer.

Inventory	State	Description
Virtual Machines (VM)	Error	VM is in an error state.
	Paused	VM is in an interrupted state.
	Pausing	VM is in the process of being interrupted.
	Started	VM is in a ready, rebooted, or reachable state.
	Starting	VM is in a start, reboot, or resume state.
	Stopped	VM is in a stop state.
	Suspended	VM is in suspension.
	Suspending	VM is in the process of being suspended.
	Stopping	VM is in the process of being stopped.
	Running	VM is in the start, ready, reboot, resumed, or reachable state.
	Terminating	VM is in the terminating state.
	Terminated	VM is in a clean state.
Kubernetes Workloads	Failed	Kubernetes Workload is in a failed state.
	Pending	Kubernetes Workload has been accepted by the system.
	Running	Kubernetes Workload is bound and in the start, ready, reboot, resumed, or reachable state.
	Succeeded	Kubernetes Workload have terminated in success.
	Terminated	Kubernetes Workload is in a clean state.
Storage Volume	In Use	Storage space is being utilized.
	Available	Storage space is available for use.
Services	Active	Service is in a start, ready, reboot, or resumed state.
	Available	Service is available for use.
	Healthy	Service has not undergone any issues recently.
	Unhealthy	Service underwent issues recently.
	Terminated	Service is in a clean state.

Rightsizing

Rightsizing

- Overview
- Recommendation Engine
  - Proportional Resizing
- Clouds That Support Rightsizing and Nuances
- RightsizingReport
  - Instance Utilization
  - Recommendations
    - Underutilized Tab
    - Unused Tab
    - Overutilized Tab
  - Advanced Options
    - Saving Filters
    - Scheduling Reports
- Options in the Actions Column
  - Resizing an Instance
  - Resizing using the Custom Resize Option
  - Stopping and Terminating a VM

Overview

Rightsizing is the process of recommending the use of right instance type or right resources for an application to optimize cost for an organization. Typically, instances are overprovisioned for an application. Overprovisioning of resources results in resources, such as CPU, memory, and so on,being unu sed. In turn, underutilization leads to an increase in cost you spend much more than what you should.

> ⚠ If some instances are under-provisioned, the recommendation engine recommends upsizing of instances which might not result in cost savings, but improvement in application performance.

Recommendation Engine

The recommendation engine uses analgorithmbased onCPU and memory andrecommends the right set of actions (downsizing of instances) that results in significant cost reduction without affecting the application performance. This algorithm uses the default thresholds (for CPU and memory) to arrive at downsize or upsize recommendations.

The recommendation engine works as follows:

- Thealgorithm matches the attributes (CPU, memory, network, and storage capabilities) of running a virtual machine and compares the attributes with operational metrics, such as CPU utilization to arrive at the rightsize for instance.
- The algorithm is cost-conscious and offers multiple candidate choices to resize the instance.
- The threshold limitations can be set or modified in the Settings submenu of the **Admin** menu.
- The Resize action is handled bythe Workload Manager.
- When suggesting resizing recommendations, an instance with RI opportunities is preferred for instances of the same type. This ensures maximum utilization of resources and minimum cost.

Proportional Resizing

The recommendation engine recommends instances based on the CPU or memory utilization of the instance over a period of time. Based on the factors of underutilization or overutilization, appropriate instances are identified to ensure CPU or memory ratios will be maintained approximately. If Proportional Resizing, in theSettingssubmenu of the**Admin**menu,is turned off, equal memory instances with appropriate CPU counts will be recommended.

Clouds That Support Rightsizing and Nuances

The following table describes the clouds, its resources that are supported for rightsizing and the nuances, applicable for each cloud:

Cloud	Supported Resources	Nuances
AWS	Managed and unmanaged VMs	N/A
Azure	Managed and unmanaged VMs	N/A
Google	Managed and unmanaged VMs	Cisco does not recommend custom instance sizes for Google Cloud Platform (GCP). Predefined instance types are only supported for rightsizing recommendation.

VMware vCenter	Managed and unmanaged VMs	The recommendations appear when prices for the cloud instance types are specified in the Instance Types Settings page. The prices are specified by a Cost Optimizer admin. The average values of CPU and memory metrics are used for rightsizing recommendations.

RightsizingReport

The Rightsizing report is divided into the following:

- Instance Utilization
- Recommendations

Instance Utilization

This report provides information about all running instances against the resize threshold limits, which is displayed in braces, as defined in the Settingssubm enu of the **Admin** menu.

- **Ideal Utilization** Instances running between the maximum and minimum threshold limits and for which noaction is required.
- **Unused** Instances not being used, which can be stopped or terminated.
- **Underutilized** Instances running below the minimum threshold limit and must be downsized.
- **Overutilized** Instances running above the maximum thresholdlimit and must be upsized.

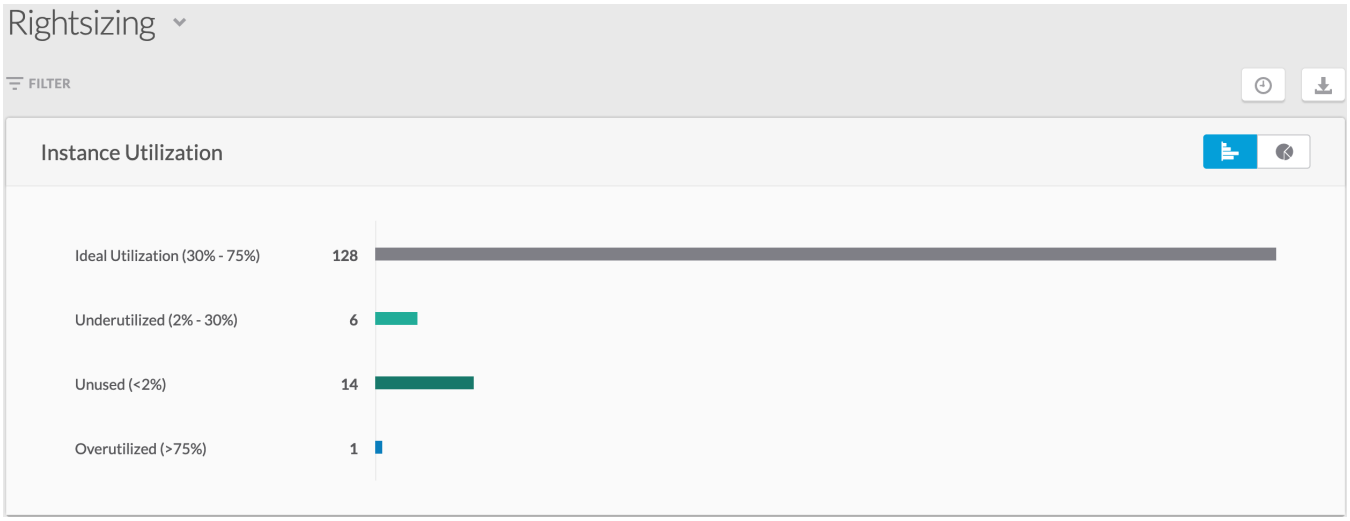

Recommendations

The Recommendations report provides detailed recommendations for all running managed and unmanaged VMs under the following tabs:

- Underutilized
- Unused
- Overutilized

> Accounts must be enabled with**PROVISIONING_REPORTING** for the rightsizing engine to offer rightsizing recommendations. Support for recommendations in accounts with other roles will be added in a future release.

The following table explains the columns in the Recommendations report.

Column Heading	Description
Recommendations	Total number of VMs on which rightsizing recommendations can be applied.
Maximum Potential Savings	Maximum savings incurred by applying the rightsizing recommendations on each VM.
Show Dismissed	Displays dismissed VMs in the recommendations report.
VM	Displays the VM name as a link. Click the link to view details about the VM.

Current Size	Current instance model type on which the VM is running.
Low/High Utilization	Actual utilization numbers (in percent) for the instances observed for a specific time.
Resize Recommendation	Instance model type to which the VM can be resized from the current size and the potential savings that can be achieved by choosing the specified instance. For VMware clouds, you can choose **Custom Resize** option that allows you to specify attributes for a user-defined instance type.
Potential Savings	Savings, based on current utilization, incurred as a result of choosing the recommended VM.
Actions	Allows you to do the following: • **Resize** Resizes to the recommended instance. • **Dismiss** Remove the instance from the recommendation report. To include the dismissed instances in the report, toggle **ON** the **Show Dismissed** icon. • **Stop** Stops the instance temporarily to restart it at a later time. • **Terminate** Shuts down the instance. ⚠ While **Resize** is available in Underutilized and Overutilized recommendations report, **Stop** and **Terminate** are available in Unused recommendations report. The **Actions** button above the table allows you to stop and terminate all or selected VMs listed on the page.

Underutilized Tab

The following is a sample screenshot of the Underutilized tab in the Recommendations report.A VM is considered underutilized if the value in the**High Utilization**field is consistently lower (for a specific time) than the value mentioned in**Min. CPU Threshold**field in theSettingssubmenu of the**Admin**menu.

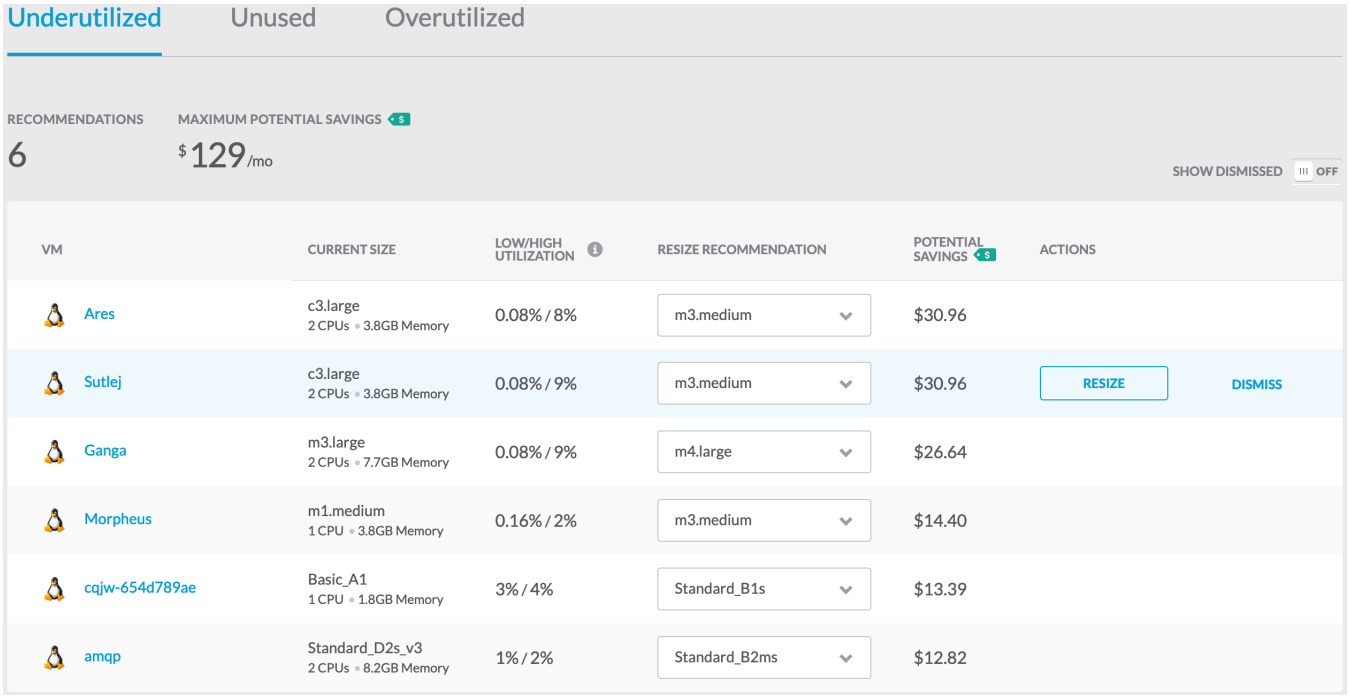

Unused Tab

An instance is termed as an unused instance if the utilization is below the terminate threshold as specified in the Rightsizing card of the Settingssubmenuwhen the Rightsize Analyzer collects the data. If the utilization is above the terminate threshold settings as specified in the Rightsizing cardwhen the rightsize analyzer (seeData Collection) runs the next day, the instance ceases to be an unused instance.

The following is a sample screenshot of the **Unused** tab in the Recommendations report.

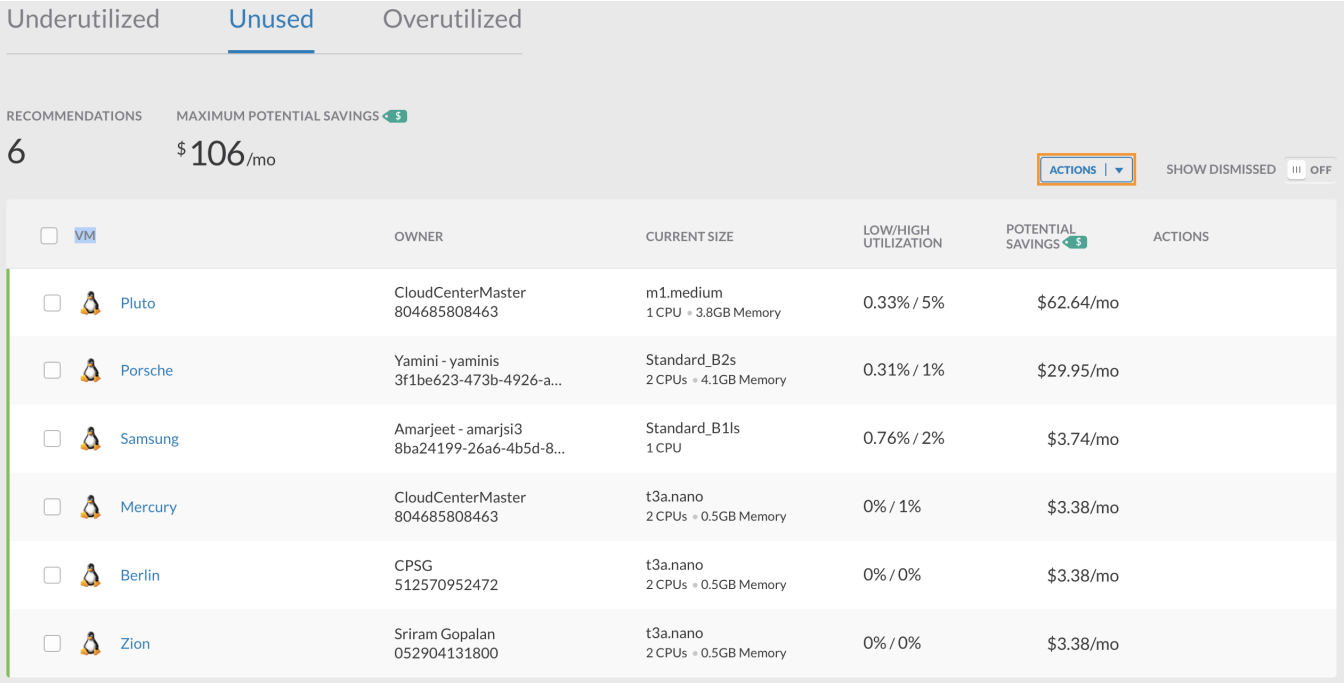

Use the **Actions** button to stop or terminate all or selected VMs listed on the page.

Overutilized Tab

The following is a sample screenshot of the **Overutilized** tab in the Recommendations report.A VM is considered overutilized if the value in the**High Utilization**field is consistently higher (for a specific time) than the value mentioned in the**Max. CPU Threshold**field in theSettingssubmenu of the**Admin**menu.

If the **ShowCost-incurring Upsize Recommendations**option in theSettings submenu is turned on, upsize recommendations for overutilized instances are provided though the recommendations do not result in potential savings.

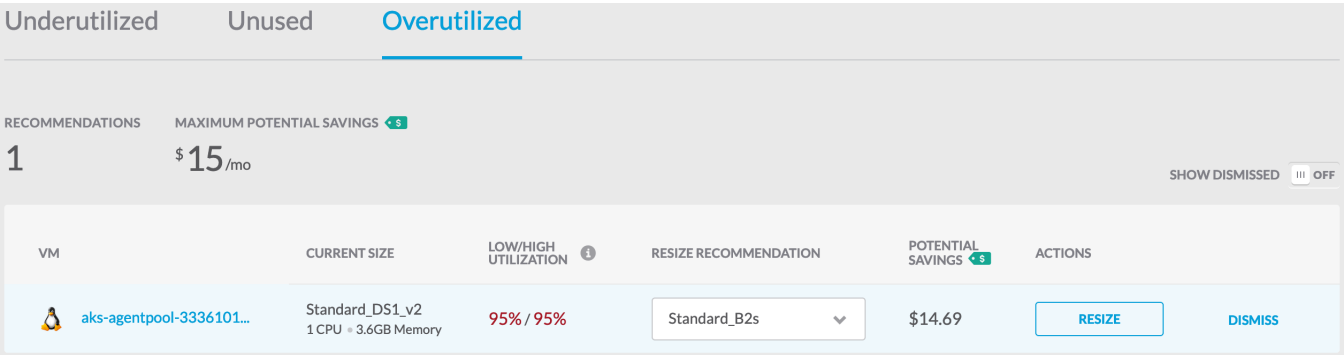

Advanced Options

You can do the following on the Rightsizing report:

- Download the report
- Save filters in the report
- Schedule a report

Saving Filters

You can choose to save a combination of options in the**Filter**menu for future use through the**Save Filters**featureso that you can quickly access and use the filter at a later time. To save a filter, do the following:

1. Choose the required filter options in the**Filter Panel**pane.

2. The**Save**button appears. The**Save New Filter**dialog appears. At this point, you can remove the selected filters by choosing the**Reset**button.

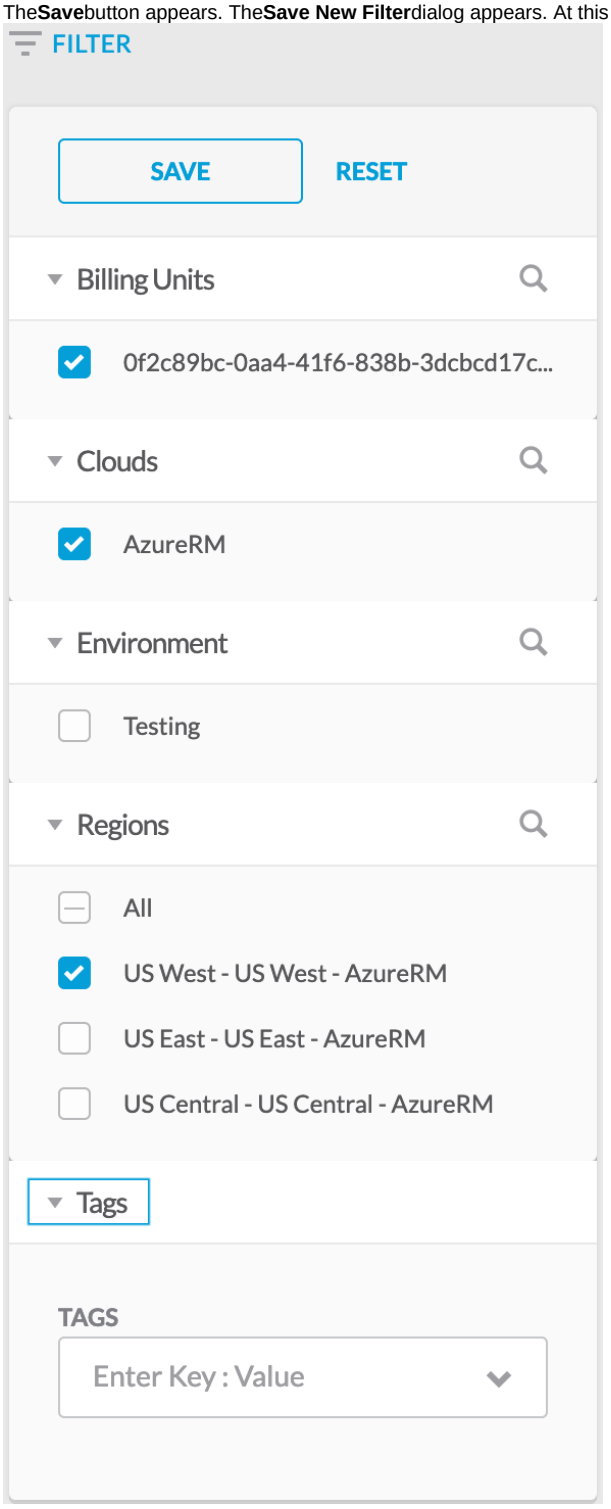

3. Specify a name for this filter and click**Save**. A status message appears indicating that the filter has been saved.

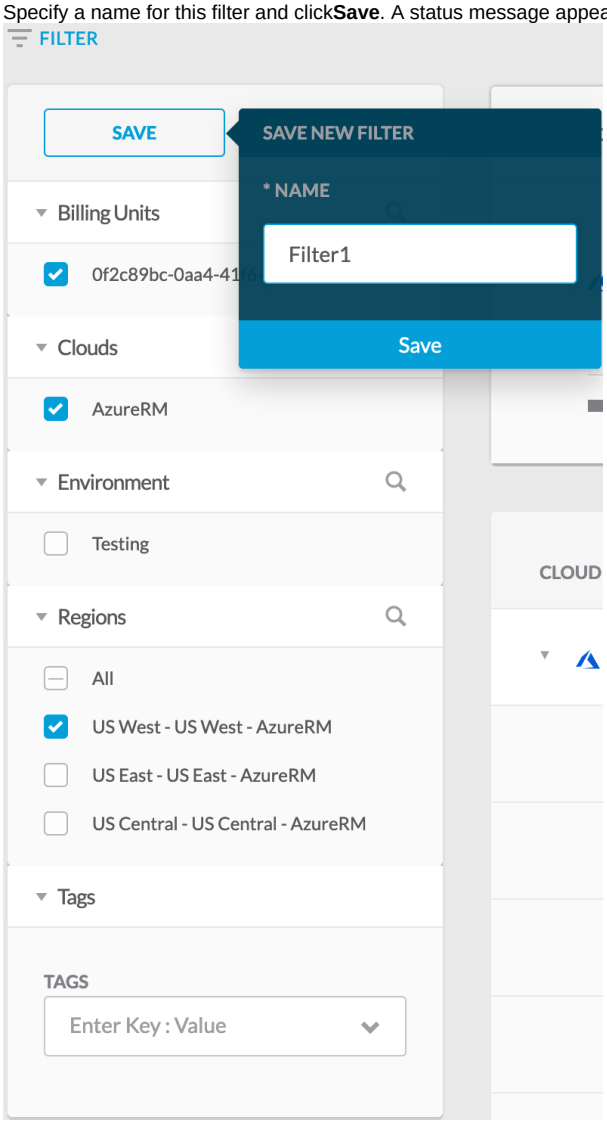

4. You can access and view the saved filters from the dropdown list.

You can mark the filter as a favorite by clicking the pin icon next to the filter name.

Scheduling Reports

The**Scheduler**icon allows you to schedule report generation periodically on a fixed date or at intervals. To create a schedule, do the following:

1. Click the**Scheduler**icon. The**Schedule**dialog appears.
2. Click**Schedule New**.

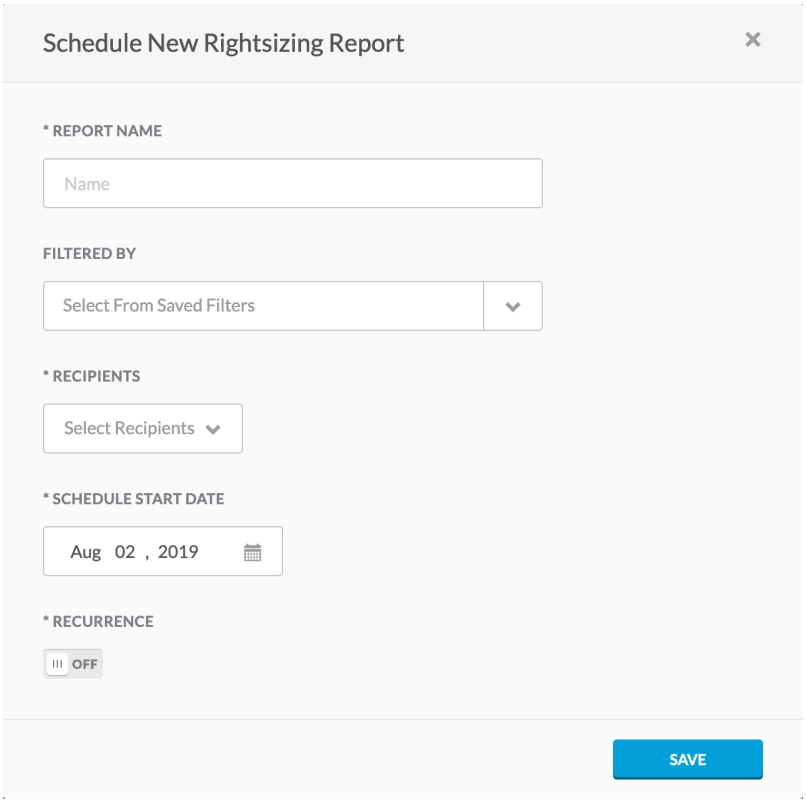

3. Do the following:

    a. Enter a name for the schedule.
    b. Choose filtering options for the schedule from the**Filtered By**field.The information in this field is populated when you save the filtering options on the page.You can choose to select a filter or leave the field empty.
    c. Choose the date range.
    d. Select the recipients the report must be sent to.
    e. Specify the start date.
    f. Toggle on the**Recurrence**button to send the report at intervals.
    g. In the**Repeats Every**area, specify the number of times the report must be sent to the recipients and choose the interval **Daily**or**Weekly**. if you choose Weekly, you can also specify the days of the week when the report is sent.
    h. Select the period to end the schedule. The options are:

        i. **Never**Send report forever or until the schedule is deleted.
        ii. **On** Date when the report should be sent.
        iii. **After** Number of occurrences after which the report is not scheduled.

4. Click**Save**. The report is displayed in the **Scheduled *Report Name*** dialog as shown in the sample screenshot below.

Optionally, you can use the**Edit**and **Delete** options in the**Actions**column to make changes to the schedule or delete the report respectively.

5. Click**Done**to close the dialog.

Options in the Actions Column

As mentioned in the table above, the**Action**button performs the following:

- Dismiss
- Resize
- Stop
- Terminate

Choose**Dismiss**to remove the recommendations for a VM from the report.

Resizing an Instance

The**Resize**button resizes a VM to the recommended instance type. If a VM in the recommendations table is anunmanagedVM, the VM must be imported first before it is resized. To resize a VM, perform the following steps:

1. Click the**Resize**button. The**Resize*VM Name***dialog appears.

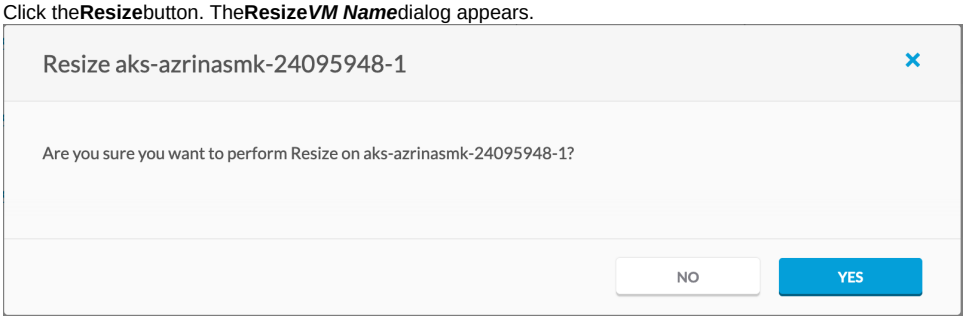

2. Click**Yes**. Adialog appears displaying the current size of the VM, the recommended size of the VM, and the potential savings as a result of resizing the instance.

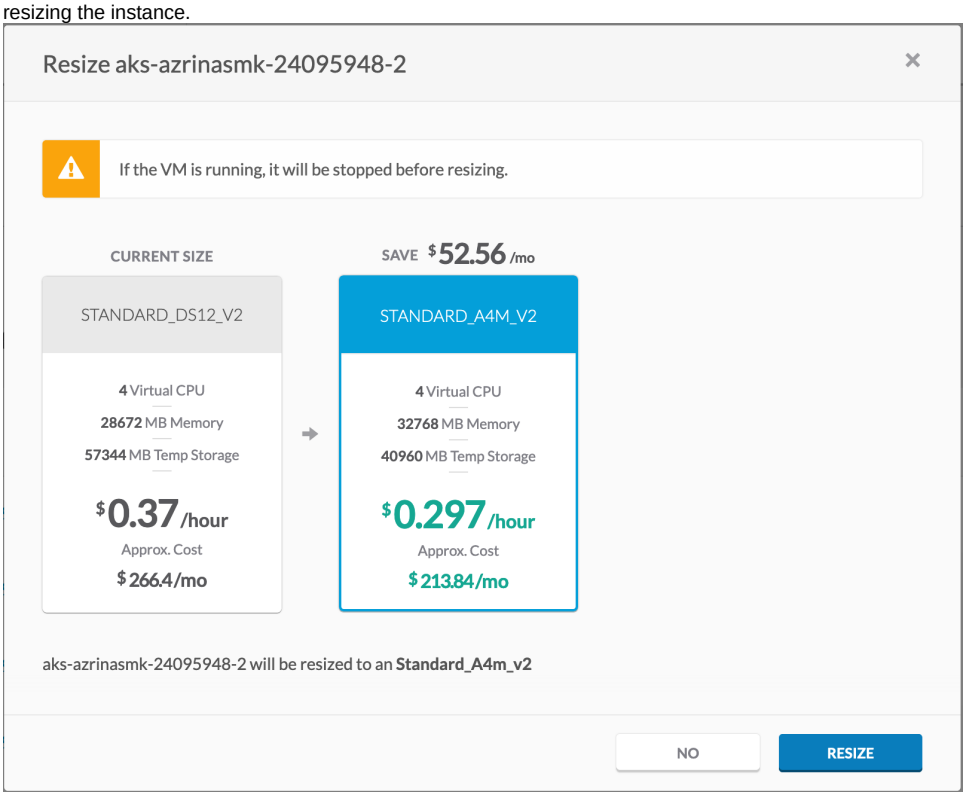

3. Click**Resize**. A spinning circle icon appears in the**Actions**column indicating that resize is in progress.
4. On completion, anotificationappears displaying that the VM has been resized.

Navigate to the**History**tab in the VM Detailspage in Workload Manager (seeVirtual Machine Management> VM Details)for acomplete history ofactionsperformed on this VM.

Resizing using the Custom Resize Option

Perform the following steps to resize an instance using the Custom Resize option:

1. Choose **Custom****Resize**option in the Recommendations dropdown. The**Resize*VM Name***dialog appears.

2. In the **Resize *VMName*** dialog, specifythe instance details.

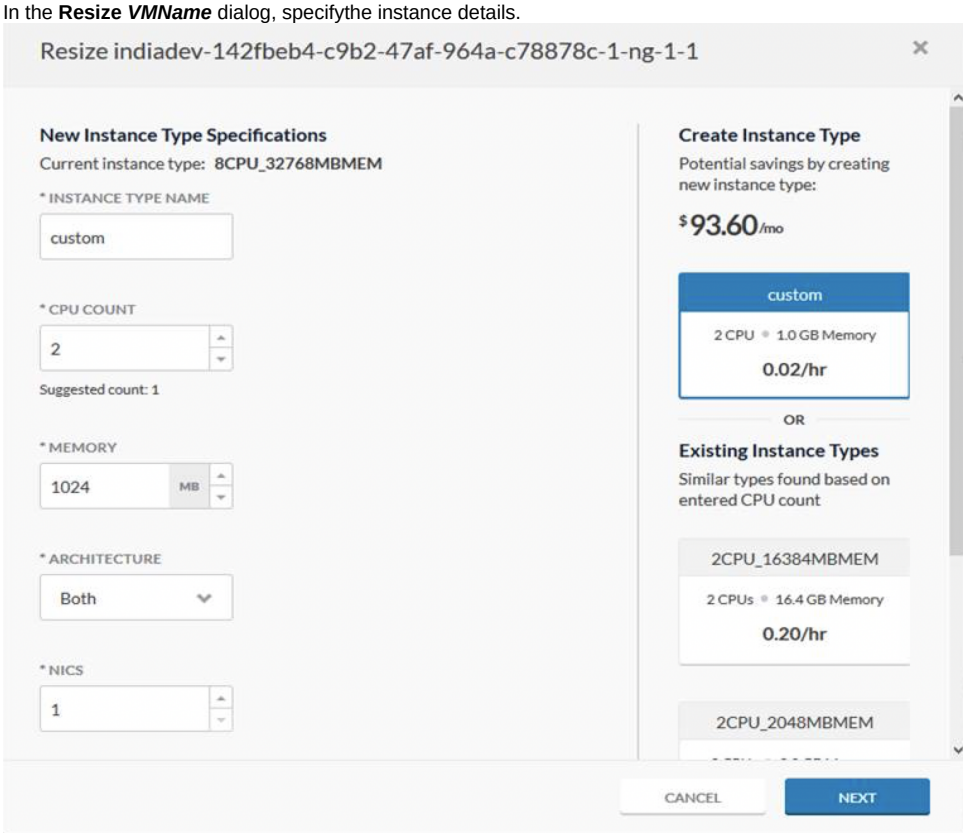

3. Click **Next** and verify the custom resize details.

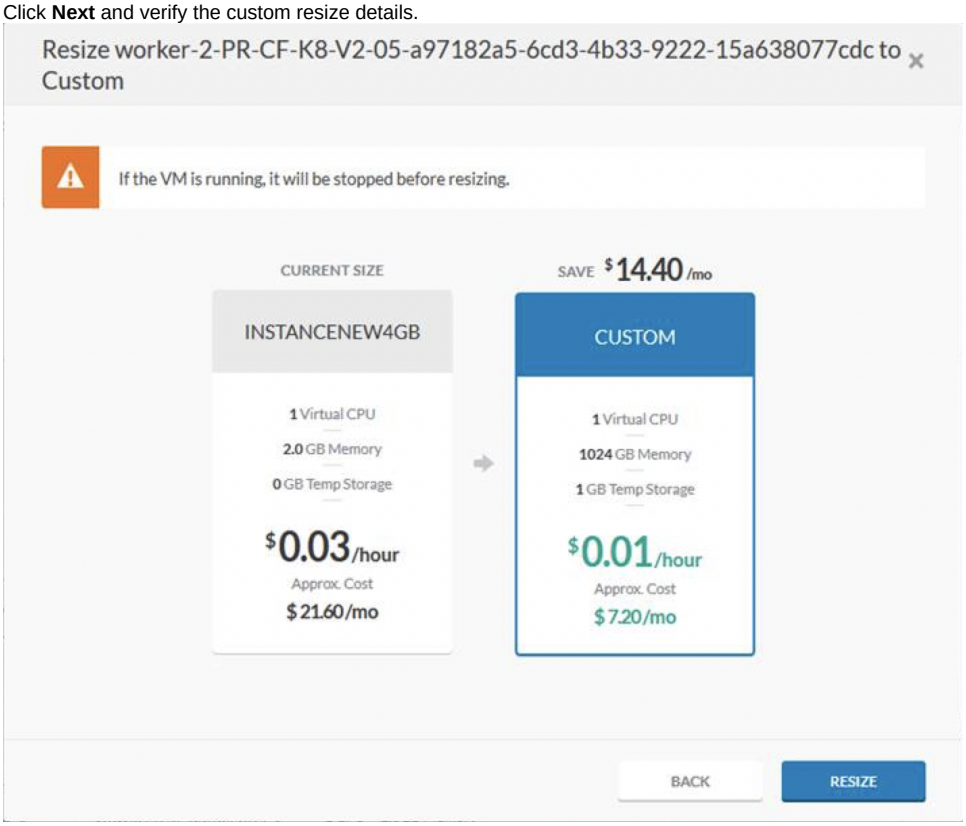

4. Click**Resize**. A spinning circle icon appears in the**Actions**column indicating that resize is in progress.
5. On completion, anotificationappears displaying that the VM has been resized.

Stopping and Terminating a VM

You can stop and terminate a VM through the following methods:

- Use the**Actions** Buttonto stop or terminate all or selected VMs by either choosing the checkbox against each VM or choosing the checkbox in the report header.
- Hover over the**Actions**column against a VM toindividually stop or terminate VMs.

When you stop a VM, the VM is shut down and you will not be charged for the usage until you plan to start it again in the near future.When you terminate a VM, the VM is shut down and permanently removed. You are not charged for the usage any longer.Therefore, you should stop a VM if you plan to start it again else, you may terminate the VM instead of stopping it to save cost.

> ⚠ If a VM in the recommendations table is anunmanaged VM, the VM must be imported first before it is stopped or terminated.

To stop or terminate all VMs in the **Unused** tab, click the**Actions**button, and choose the **Stop All** or **Terminate All** options. Irrespective of the pages the report spans, all listed VMs are stopped or terminated.

To stop or terminate a VM, perform the following steps:

1. Do one of the following:
    a. For multiple VMs,select the VMs by clicking the checkbox adjacent to a VM and click the **Actions** button.
    b. For individual VMs,hover on the**Actions**column.Click the icon to display the options **Stop**or**Terminate**and choose the appropriate option. The **Stop *VM Name*** dialog appears.

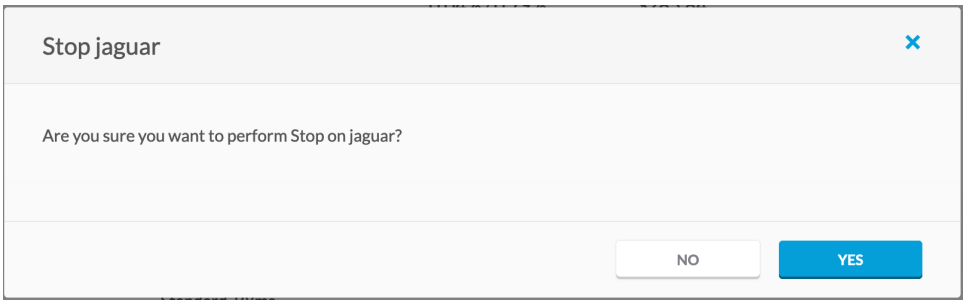

2. Click**Yes**.A spinning circle icon appears in the**Actions**column indicating that chosen action is in progress.
3. On completion, a notification appears displaying that the VM has been stopped or terminated.

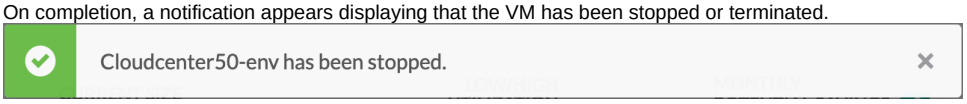

Suspension Candidates

Suspension Candidates

- Overview
- Suspension Reports
    - Advanced Options
        - Saving Filters
        - Scheduling Reports
- Suspending a VM

Overview

Suspension policies are a powerful method to conserve cloud resourcesby moving a cloud resource from Running state to Suspended state (see Inventory States) when the resource is not needed or to prevent a deployment from running during times that it should not be accessed.Thus, suspension policies help in reducing cost on cloud resources when resources are not used.Suspension policies are an everyday activity, in which Instances could be suspended for specific hours in a day for a minimum of 30 minutes. The metric collector background process (see Data Collection) collects the data for analysis based on the utilization pattern.

> The Suspension Policy Analyzer uses a python package named *numpy* as part of its machine learning code.The latest version of *numpy* package, which is 1.16.4, is used in Cost Optimizer 5.1.0.There is a known issue about loading untrusted scripts and is documented here: https://snyk.io/blog/numpy-arbitrary-code-execution-vulnerability/.The vulnerable function (*numpy.load*) is not used or invoked as part of CloudCenter Suite machine learning code.This issue will be addressed when a patch is available.

Suspension Reports

The Suspension Candidates Report lists VMs and deployments for which you can apply asuspension policy based on the specifications defined in the Suspension Candidates card in the Settingssubmenu of the**Admin**menu. You specify a schedule when a VM must be in the Running state (see Inventory States) during a certain time period every day. At other times, the deployment remains suspended.

The following is a sample screenshot of the Suspension Candidate report.

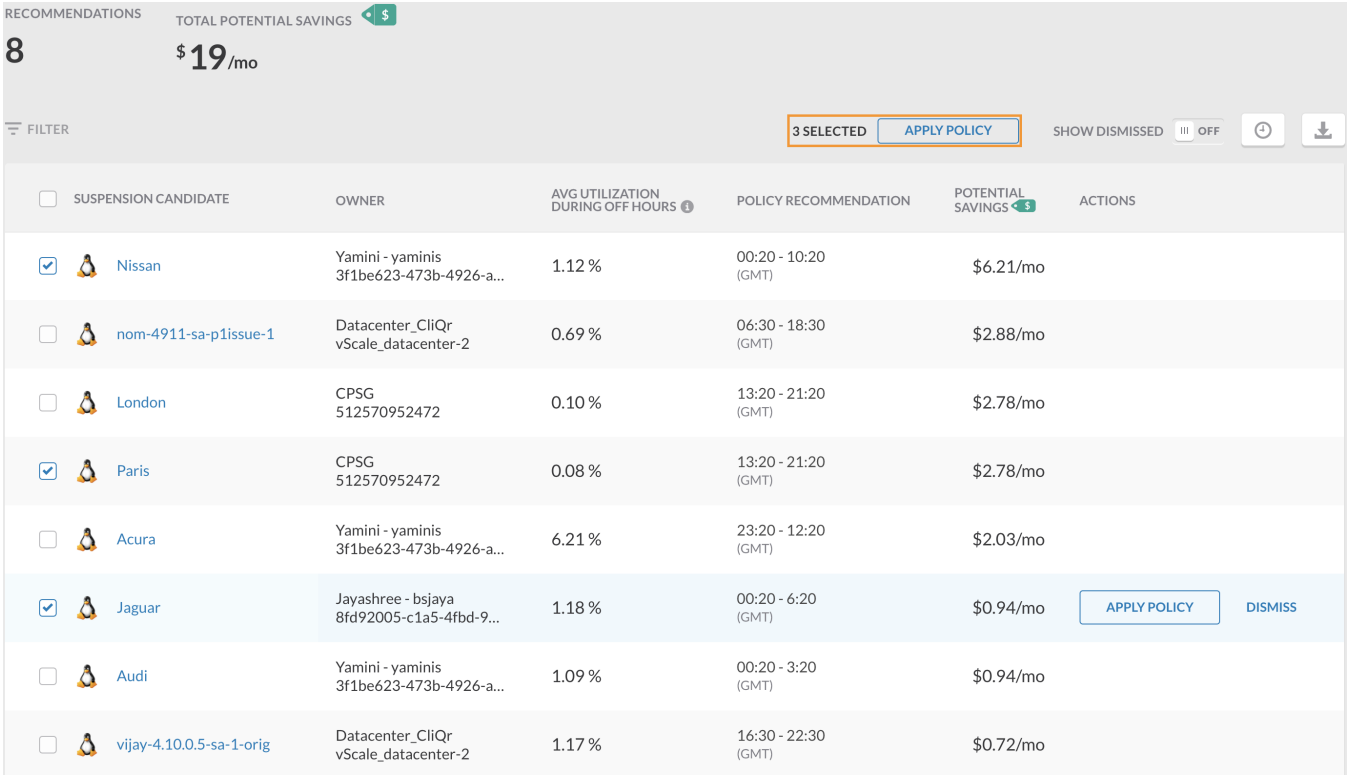

The following table explains the columns in the report.

Column Heading	Description

Recommendations	Total number of VMs on which suspension policies can be applied.
Total Potential Savings	Total savings that can be incurred by applying the suspension policy on each VM.
Show Dismissed	Displays dismissed suspension candidates in the report.
Apply Policy	Apply policies on multiple VMs listed on the page by choosing the checkbox against each VM.
Suspension Candidate	Displays the VM name as a link. Click the link to view details about the VM.
Avg Utilization During Off Hours	Actual utilization numbers (in percent) in non-peak hours.
Policy Recommendation	Suspension policy recommendation based on VM utilization pattern.
Maximum Potential Savings	Savings that will be incurred as a result of effecting the suspension policy.
Actions	Allows you to do the following: • **Suspend** Attach a suspension policy to the VM. • **Dismiss** Dismiss the recommendation.

Advanced Options

You can do the following on the Suspension Candidates report:

- Download the report
- Save filters in the report
- Schedule a report

Saving Filters

You can choose to save a combination of options in the**Filter**menu for future use through the**Save Filters**featureso that you can quickly access and use the filter at a later time. To save a filter, do the following:

1. Choose the required filter options in the**Filter Panel**pane.

2. The**Save**button appears. The**Save New Filter**dialog appears. At this point, you can remove the selected filters by choosing the**Reset**button.

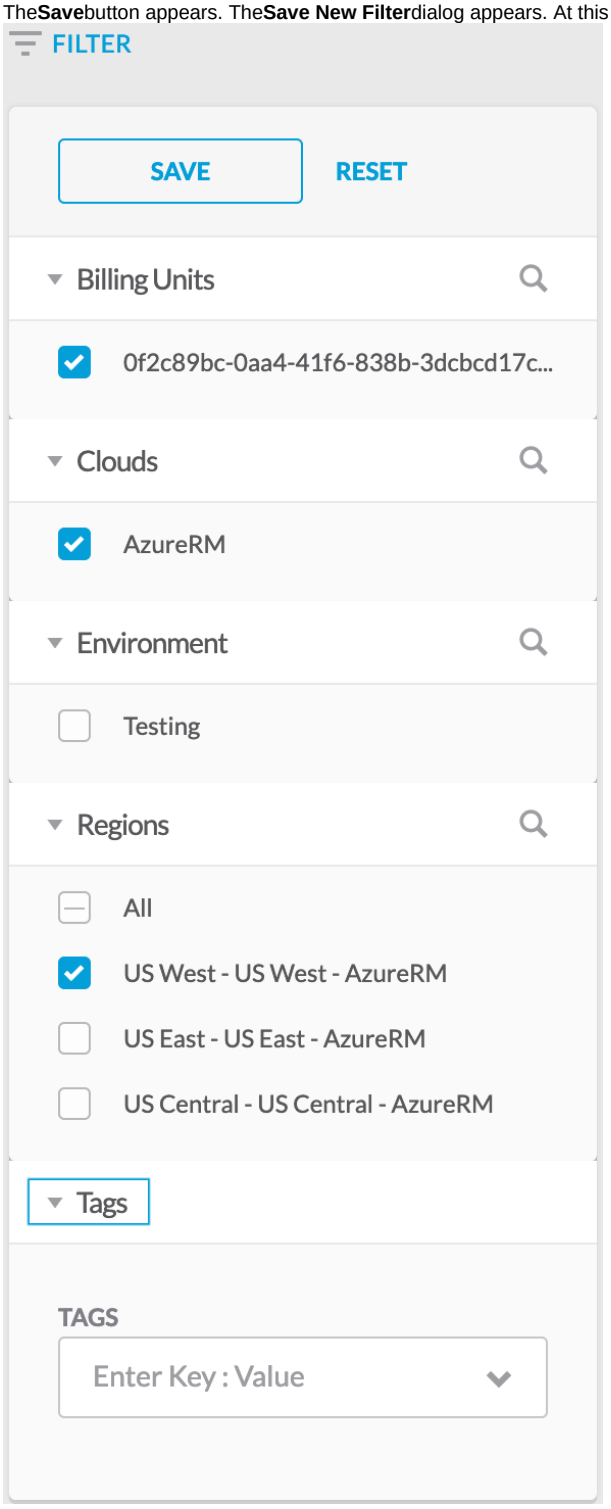

3. Specify a name for this filter and click**Save**. A status message appears indicating that the filter has been saved.

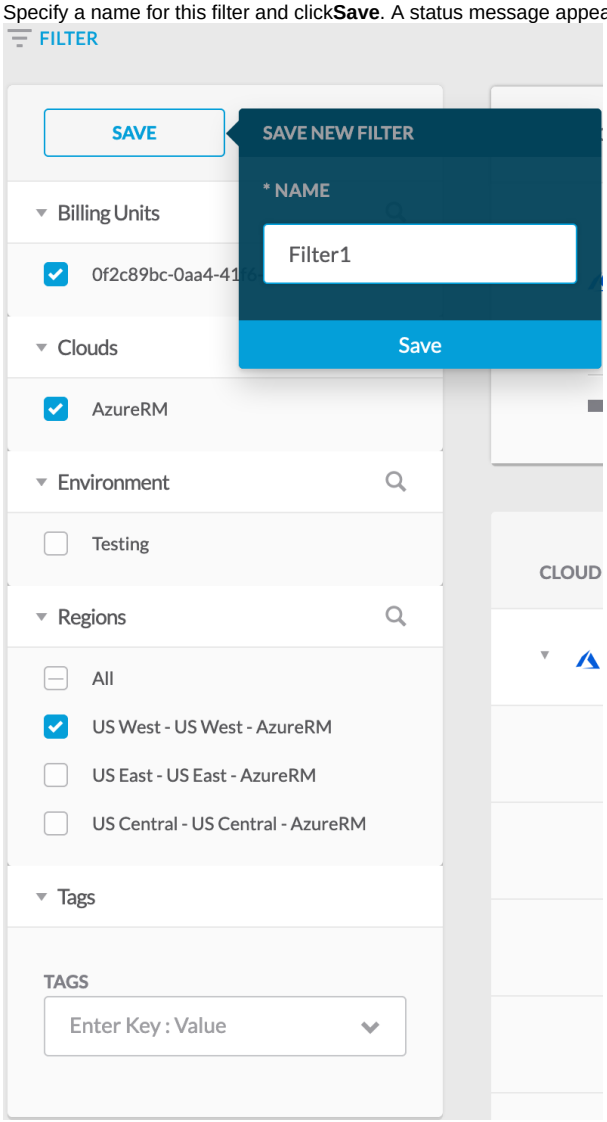

4. You can access and view the saved filters from the dropdown list.

You can mark the filter as a favorite by clicking the pin icon next to the filter name.

Scheduling Reports

The**Scheduler**icon allows you to schedule report generation periodically on a fixed date or at intervals. To create a schedule, do the following:

1. Click the**Scheduler**icon. The**ScheduleNew *Report Name*** dialog appears.

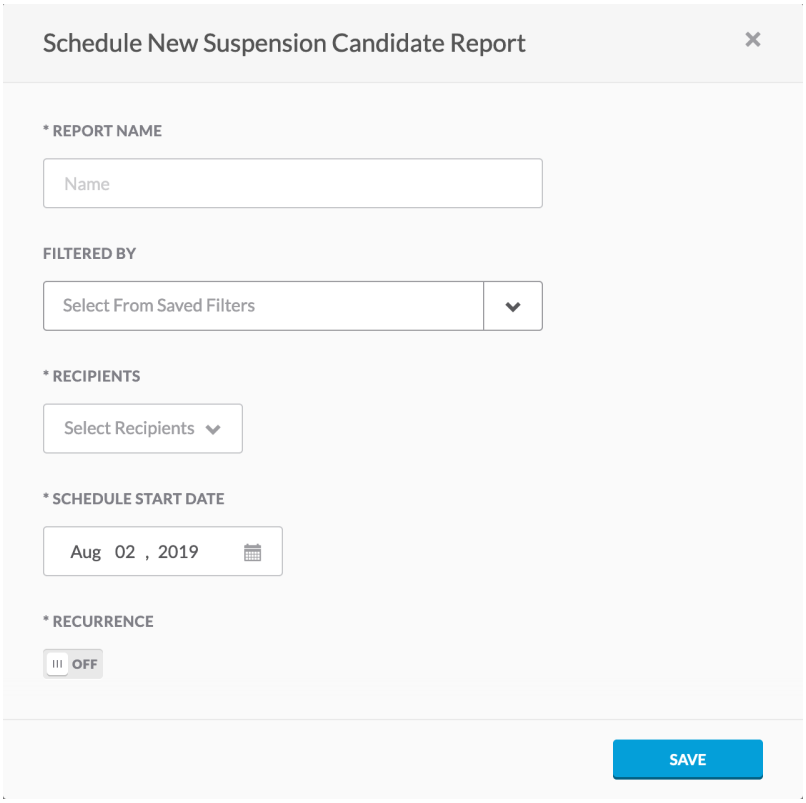

2. Do the following:

   a. Enter a name for the schedule.
   b. Choose filtering options for the schedule from the **Filtered By** field. The information in this field is populated when you save the filtering options in the page. You can choose to select a filter or leave the field empty.
   c. Choose the date range.
   d. Select the recipients the report must be sent.
   e. Specify the start date.
   f. Toggle on the **Recurrence** button to send the report at intervals.
   g. In the **Repeats Every** area, specify the number of times the report must be sent to the recipients and choose the interval **Daily** or **Weekly**. if you choose Weekly, you can also specify the days of the week when the report is sent.
   h. Select the period to end the schedule. The options are:

      i. **Never** Send report forever or until the schedule is deleted.
      ii. **On** Date when the report should be sent.
      iii. **After** Number of occurrences after which the report is not scheduled.

3. Click **Save**. The report is displayed in the **Scheduled *Report Name*** dialog as shown in the sample screenshot below.

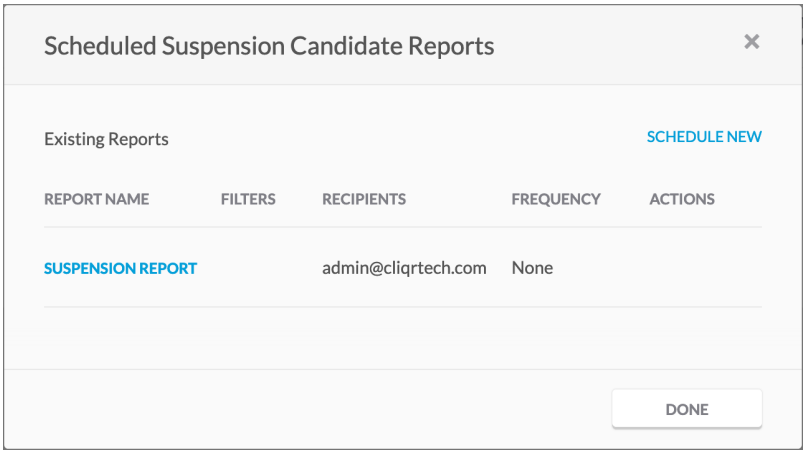

> Optionally, you can use the **Edit** and **Delete** options in the **Actions** column to make changes to the schedule or delete the report respectively.

4. Click **Done** to close the dialog.

Suspending a VM

You can suspend a VM through the following methods:

- Use the **Apply Policy** button to apply policies on all or selected VMs by either choosing the checkbox against each VM or choosing the checkbox in the report header.
- Hover over the **Actions** column against a VM toindividually suspend VMs.

You can either apply an existing suspension policy or create a new suspension policy. If you choose to use an existing suspension policy, the**Suspension Policy** field displays alist of policies, which are policies that are available for the VM or deployment in Workload Manager. Only existing suspension policy that matches the schedule and does not contain any blackout dates are listed in the**Suspension Policy** field.

> ⚠ If the VM or deployment for suspension is not available in Cost Optimizer, it must be imported first.

The following rules apply for a suspension policy:

- To apply a policy, you must be assigned to at least one Workload Manager role.
- You can create a new suspension policy only if you are part of the **WM_POLICY_MANAGER**role.

SeeOOB Groups, Roles, and Permissions for additional details about Workload Manager roles.

Perform the following steps to apply a suspension policy.

1. Do one of the following:
    a. To suspend multiple VMs, select the VMs by clicking the checkbox adjacent to a VM.
    b. To suspend an individual VM, click **Suspend** in the **Actions** column.
    The **Suspend *VM***dialog appears,requesting confirmation if you would like to import the VM into Workload Manager before it can be suspended.
2. Click **Yes**.The **Suspend *VM*** dialog with two tabs **Apply Existing Policy** and **Create New Policy** appears.
3. If you choose the**Apply Existing Policy** tab, selectan existing policy from the **Suspension Policy** drop-down list.

    If the recommended suspension schedule matches an existing policy, the *Apply Existing Policy* tab will list the matching policy schedule in the *Sus pension Policy* dropdown.

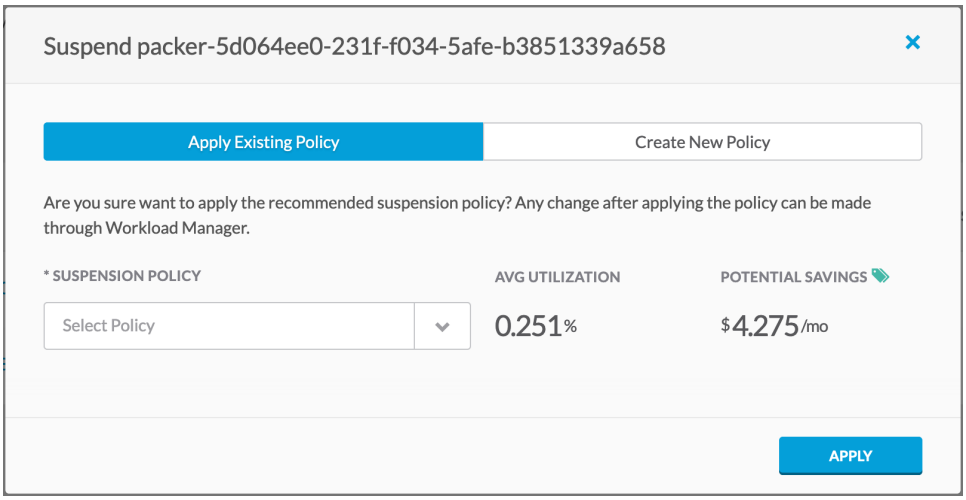

4. If you choose the**Create New Policy** tab, you can do the following:
    a. Specify a name for the suspension policy in the **Name** field.

b. Choose the duration when the VM must be suspended in the **Suspend From** and **To** fields.You may specify a schedule that is different from the recommended or an existing suspension policy.

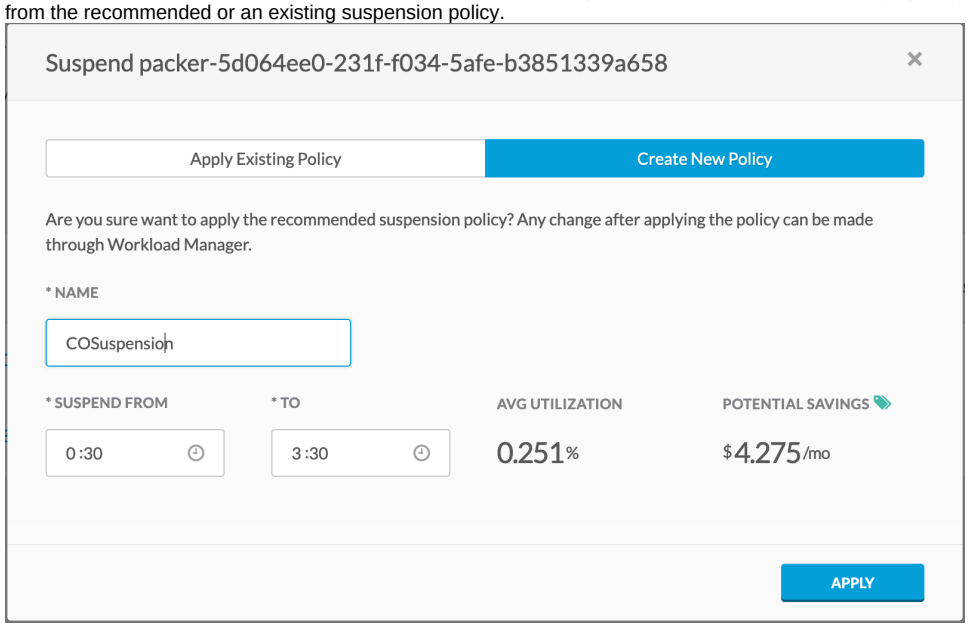

5. Click **Apply**. On completion, a dialog appears that the suspension policy has been attached.
   You can verify the successful attachment of suspension policies in the VM Details page (see Virtual Machines > *VM Details*).

After a suspension policy is applied, changes to the policy can be made in Workload Manager only, not in Cost Optimizer.

Unused Volumes

Unused Volumes

- Overview
- Unused Volumes Report
    - Advanced Options
        - Saving Filters
        - Scheduling Reports
- Terminating a Volume

Overview

Organizations may incur additional costs due to the underprovisioning or overprovisioning of storage volumes. A good solution is to find storage volumes that are not used and terminate them, thereby saving costs.

Unused Volumes Report

The Unused Volumes report displays storage volumes that are listed as **Available** in theStorage Volumes page.The following is a sample screenshot of the **Unused Volumes** report.

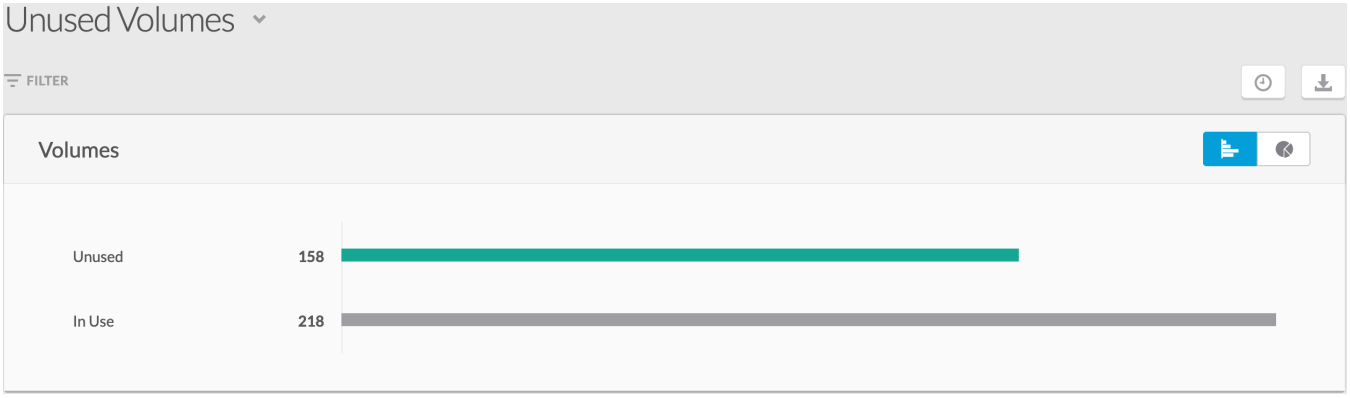

The following is a sample screenshot of the unused volumes, which are the same as the ones listed in theStorage Volumes page. The information in the following table applies to the summary displayed at the top of the page:

Summary	Description
Recommendations	Total number of volumes that are not used.
Total Potential Savings	Total savings that can be incurred by terminating all volumes.
Terminate All	Allows you to terminate all volumes listed on the page. All volumes listed in the report, irrespective of the pages, can be terminated.
Show Dismissed	Displays dismissed volumes in the report.

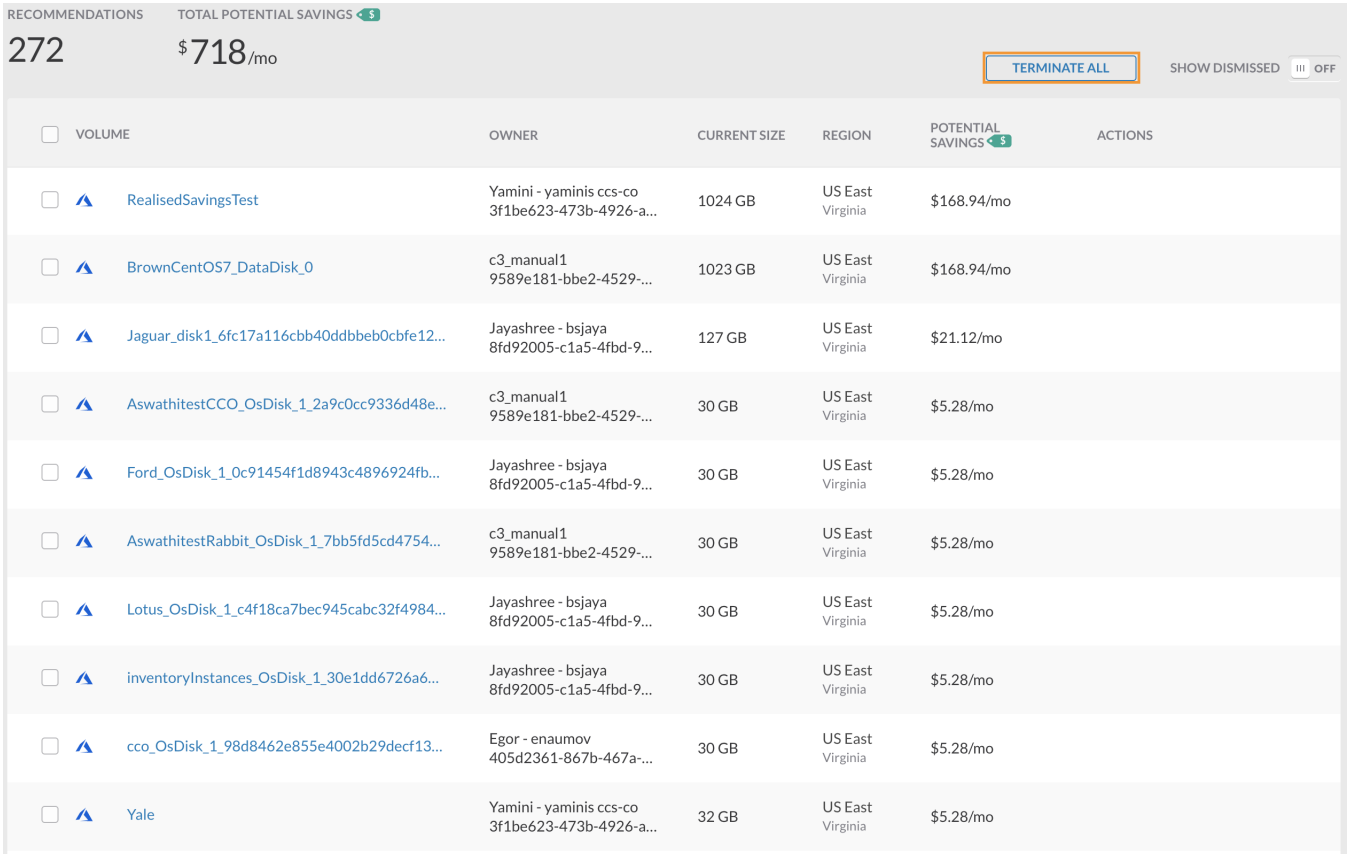

The following table explains the significantcolumns in the report.

Column Heading	Description
Volume	Displays the storage volume name as a link. Click the link to view details about the storage volume.
Current Size	The size of the volume that is not in use.
Region	Geographic location of the storage volume.
Potential Savings	Savings that will be incurred as a result of terminating the volume.
Actions	Allows you to do the following: • **Dismiss** Remove the volume from the recommendation list. • **Terminate** Deletes or terminates the volume.

Advanced Options

You can do the following on the Unused Volumes report:

- Download the report
- Save filters in the report
- Schedule a report

Saving Filters

You can choose to save a combination of options in the **Filter** menu for future use through the **Save Filters** feature so that you can quickly access and use the filter at a later time. To save a filter, do the following:

1. Choose the required filter options in the **Filter Panel** pane.

2. The**Save** button appears. The **Save New Filter** dialog appears.

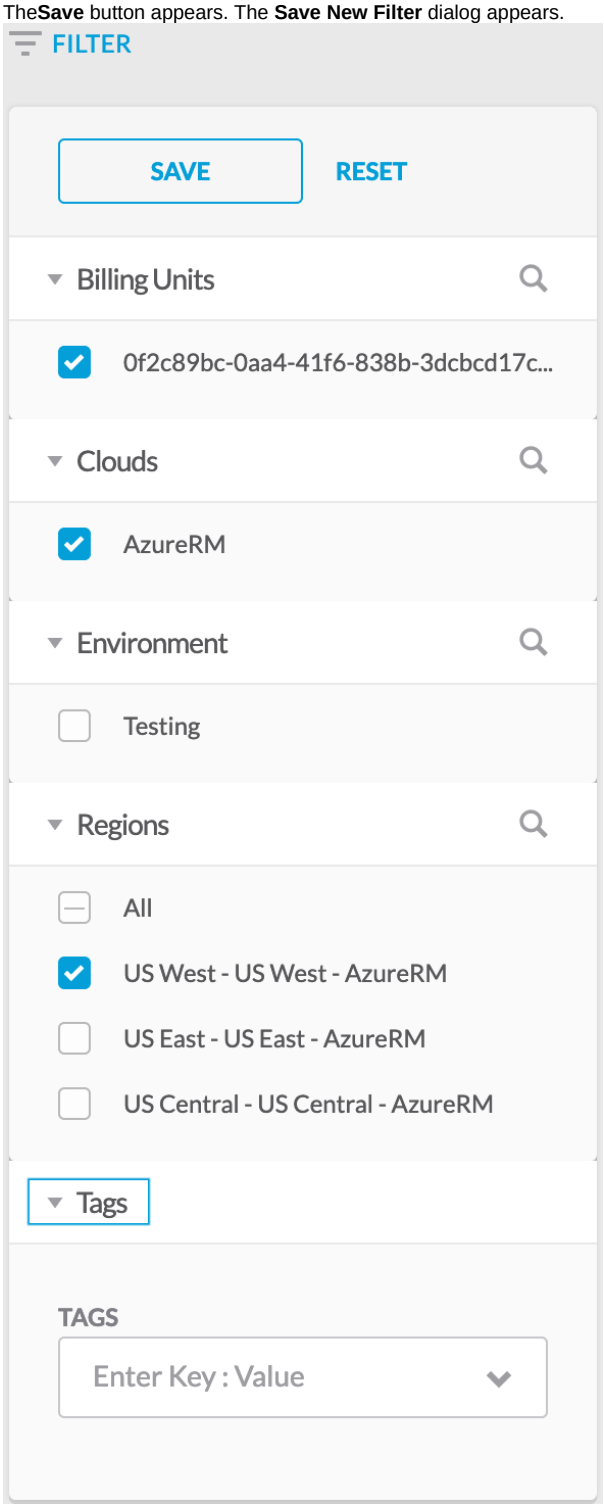

3. Specify a name for this filter and click**Save**. A status message appears indicating that the filter has been saved.

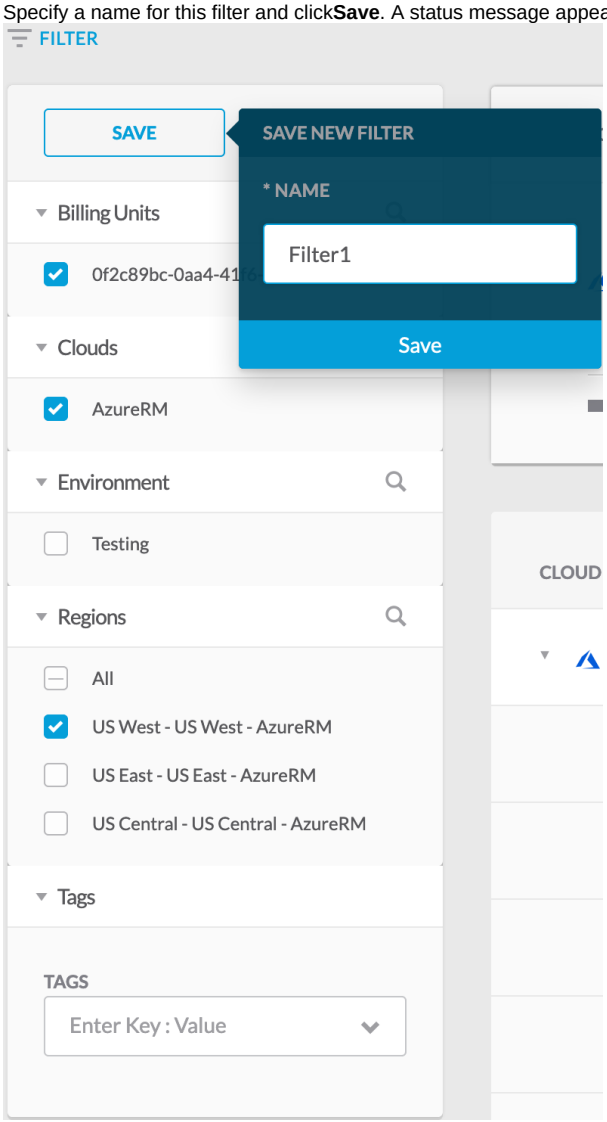

4. You can access and view the saved filters from the dropdown list.

You can mark the filter as a favorite by clicking the pin icon next to the filter name. At any point when saving the filter, youcan remove the chosen filters by choosing the**Reset**button.

Scheduling Reports

The **Scheduler** icon allows you to schedule report generation periodically on a fixed date or at intervals. To create a schedule, do the following:

1. Click the **Scheduler** icon. The **Schedule** dialog appears.
2. Click **Schedule New**.

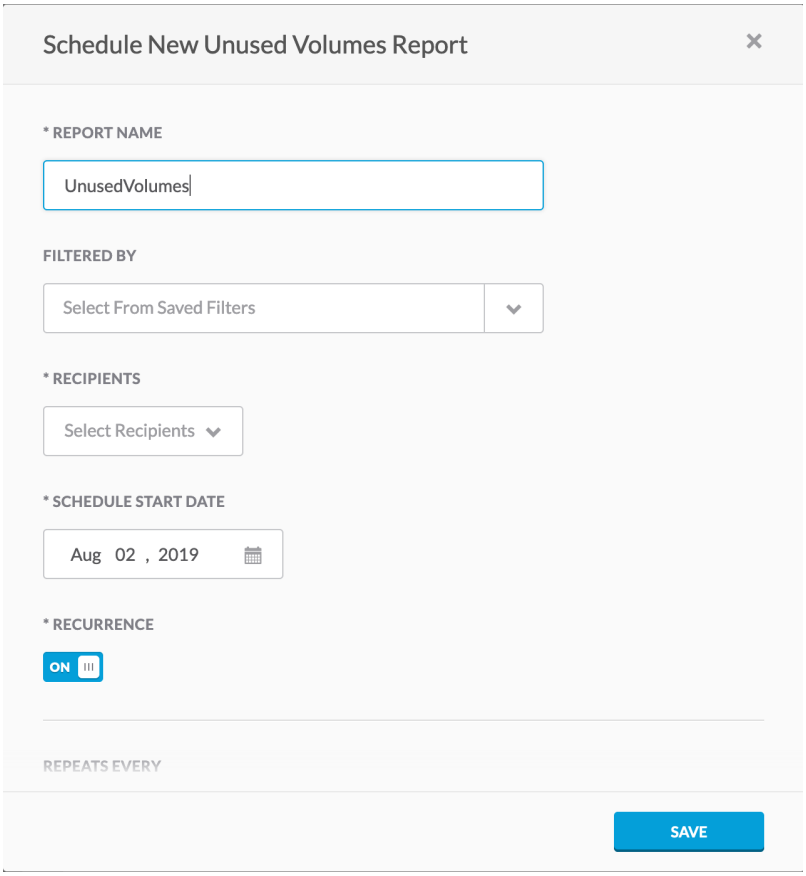

3. Do the following:

   a. Enter a name for the schedule.
   b. Choose filtering options for the schedule from the **Filtered By** field. The information in this field is populated when you save the filtering options in the page.You can choose to select a filter or leave the field empty.
   c. Choose the date range.
   d. Select the recipients the report must be sent to.
   e. Specify the start date.
   f. Toggle on the **Recurrence** button to send the report at intervals.
   g. In the **Repeats Every** area, specify the number of times the report must be sent to the recipients and choose the interval **Daily** or **Weekly** . if you choose Weekly, you can also specify the days of the week when the report is sent.
   h. Select the period to end the schedule. The options are:

      i. **Never** Send report forever or until the schedule is deleted.
      ii. **On** Date when the report should be sent.
      iii. **After** Number of occurrences after which the report is not scheduled.

4. Click**Save**.The report is displayed in the**Scheduled*Report Name***dialog as shown in the sample screenshot below.

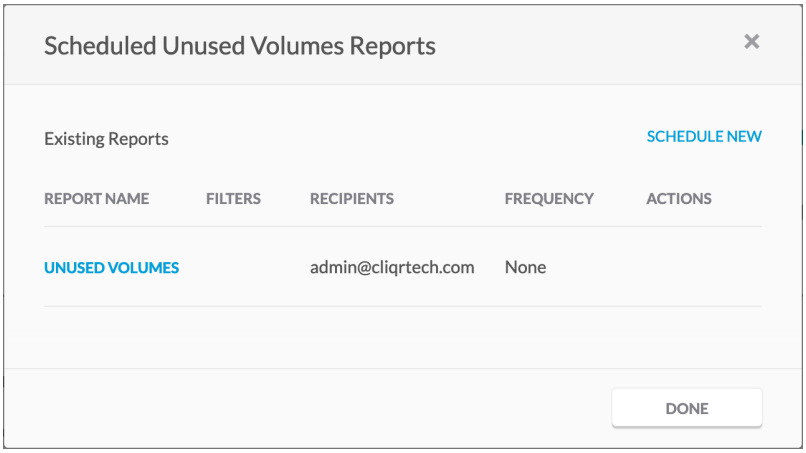

Optionally, you can use the**Edit**and **Delete** options in the**Actions**column to make changes to the schedule or delete the report respectively.

5. Click **Done** to close the dialog.

Terminating a Volume

You can stop and terminate a volume through the following methods:

- Use the**Terminate All**Buttonto terminate all or selected volumes by either choosing the checkbox against each VM or choosing the checkbox in the report header (Bulk Remedial Actions feature).
- Hover over the**Actions**column against a volume toindividually stop or terminate volumes.

As mentioned in the table above, the**Action**button performs the following:

- Dismiss
- Terminate

Choose**Dismiss**to remove a storage volume from the report.

To terminate, perform the following steps:

1. Do one of the following:
    a. For multiple volumes,select the volumes by clicking the checkbox adjacent to a volume and click the **Terminate** button.
    b. For individual volumes,hover on the**Actions**column and choose**Terminate**.
    The**Terminate*Volume Name***dialog appears.

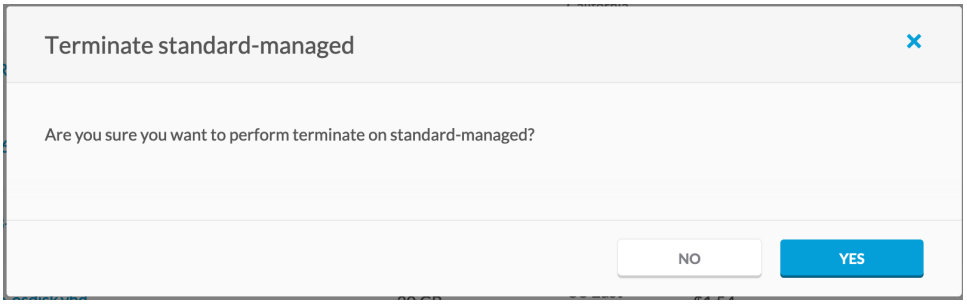

2. Click**Yes**.A spinning circle icon appears in the**Actions**column indicating that termination is in progress. On successful termination, a success message is displayed.

Reserved Instances

Reserved Instances

- Reserved Instances Overview
- RI Subscription Report
- RI Opportunities Report

Reserved Instances Overview

Reserved Instances Overview

- Introduction
- Terminology
- Payment Methods
- Reports

Introduction

> ⚠ This feature is supported on AWS EC2 only.

Reserved Instances (RI) is a method of purchasing a cloud reserve to reserve the cloud resource for a specific period.RIs offer the ability to significantly reduce instance costs over a defined term, thus, benefitting from the capacity reservation for predictable usage or workloads.

RIs offer up to 80% discount over on-demand instance costs, depending on the cloud provider, payment terms and duration.RIs is one of the most popular ways for enterprises to get great discounts on computing costs.

By combining RIs with on-demand instances, organizations can save on running costs without sacrificing reliability and flexibility.

Terminology

The following table explains the RI terminologies.

Term	Definition
Reservation	A commitment made by a customer to the cloud service provider (AWS, in this case) for using resources for a defined period. The cloud service provider in honor of this commitment offers a discount to the customer. Depending on the duration and payment terms, discounts may vary from 35% to 80%.
Utilization	The actual duration (in hours and percentage) that a RI subscription was used for a selected period. Once a reservation is purchased, you must provision or run instances of a matching type to benefit from the purchased RI hours and applicable discounts. Unused hours are not carried forward or accumulated.
Float	Transfer of utilization from one account to another.For example, ifyou reserve instance in one account and the instance is not used for an hour in that account, theusagecan be applied to another account. ⚠ To take advantage of **Float**, you must have consolidated billing enabled because **Float** is limited to the billing account and the linked accounts.
Convertible Reserved Instances	Exchange one or more RIs for another with a different configuration. There is no limit on the number of times you can exchange a RI, as long as the target Convertible Reserved Instance is of an equal or higher value than the Convertible Reserved Instances that you are exchanging.

Payment Methods

Reserved Instances can be purchased for a period of 1 year to 3 years for these resources:

- Instance Type
- Region
- Duration
- Payment Terms
- Operating System
- Availability Zone (Optional)

Foreach resource, the payment methods are as follows:

- **FullUpfront** The entire amount for the duration of the RIterm is paid in advance, providing you with a large discount. There are no per hour charges.
- **PartialUpfront** A low amount is paid to reserve the instance. A discounted hourly rate is applied for the duration of the instance.
- **NoUpfront** No upfront payment. A discounted hourly rate is applied during the duration of the instance term.

For partial or no upfront payment, the hourly cost for the RI is applied irrespective of whether the instance is running and the cost of the instance is charged to account as reservation charges.

Reports

Cost Optimizer provides a variety of out-of-the-box reports that help you to track and to manage the performance and status of RI investments in your organization. The RI reports allow youto do the following:

- Assess ROI on purchased RIs
- Discover underutilized RIs to optimize usage
- Uncover opportunities for additional RI purchases for maximum savings

RI helps in answering the following critical questions that you may have about your RIs.

- How much did I save from using RIs?
- How many RIs should I purchase?
- How can I optimize my RIs?
- What would be the additional on-demand usage that I can convert to RI to reduce my cost?
- How are my RIs performing?
- Which groups in my organization use RIs?
- Is float applicable to the groups in my organzation that use RIs?
- What is the utilization of RI across departments?
- How much of my instances are running as RIs?

This is not an exhaustive list, but a compilation of common questions that you may have on RIs.

The list of RI reports are as follows:

- RI Subscription Report
- RI Opportunities Report

RI Subscription Report

RI Subscription Report

- Overview
- RI Coverage
- RI Subscription Summary
- Advanced Options
    - Scheduling Reports
- Date Range
    - Custom Date Range
- Subscription Details
    - Group Subscription Details
    - Individual Subscription

Overview

The**RI Subscription Report** contains the following:

- RI Coverage
- RI Subscriptions

RI Coverage

This report presents the percentage of running instance hours that were covered by purchased RIs. Use this report to identify opportunities to buy matching RIs for more significant cost savings.

In the following screenshot, **Coverage 62%** indicates the number of running instance hours covered by purchased RIs, and thus, the instances benefited from RI discounts.

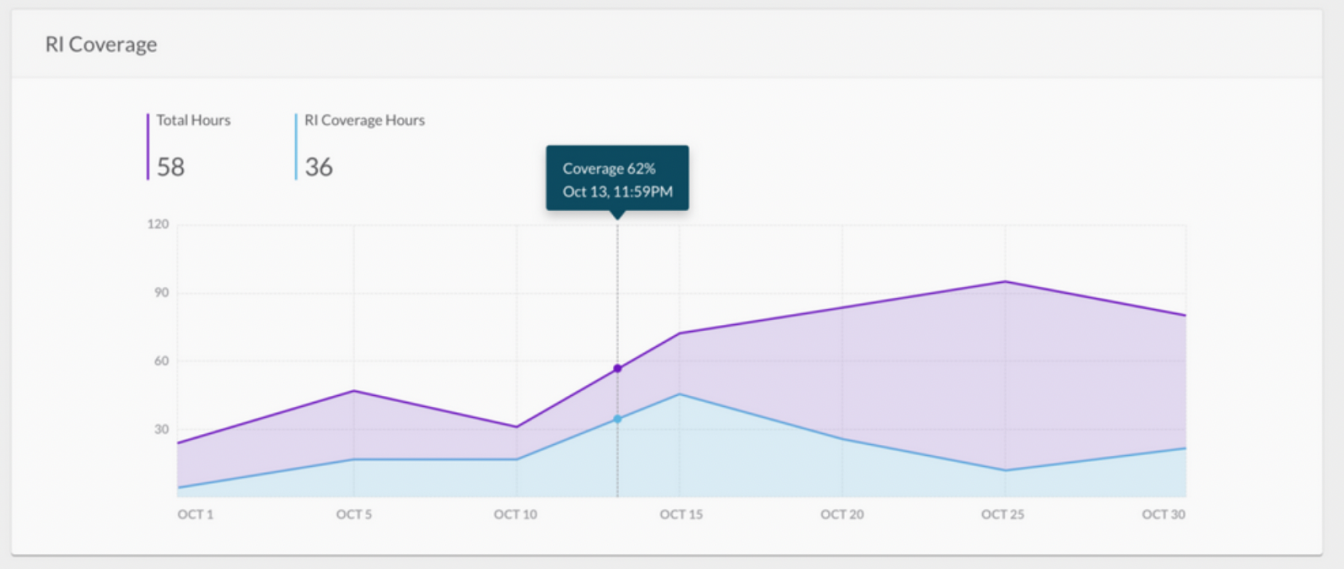

RI Subscription Summary

The following screenshot displays a sample RI Subscriptions report, which provides information about individual and group subscriptions.TheRI Utilization (**Utilization** column in the screenshot) displays the percentage of purchased RI hours that were used by matching instances over a selected period. This field helps in assessing the ROI on your RI purchases and take steps to optimize utilization, such as provisioning matching instances or switching to another instance type that is not covered by RIs or enabling float.

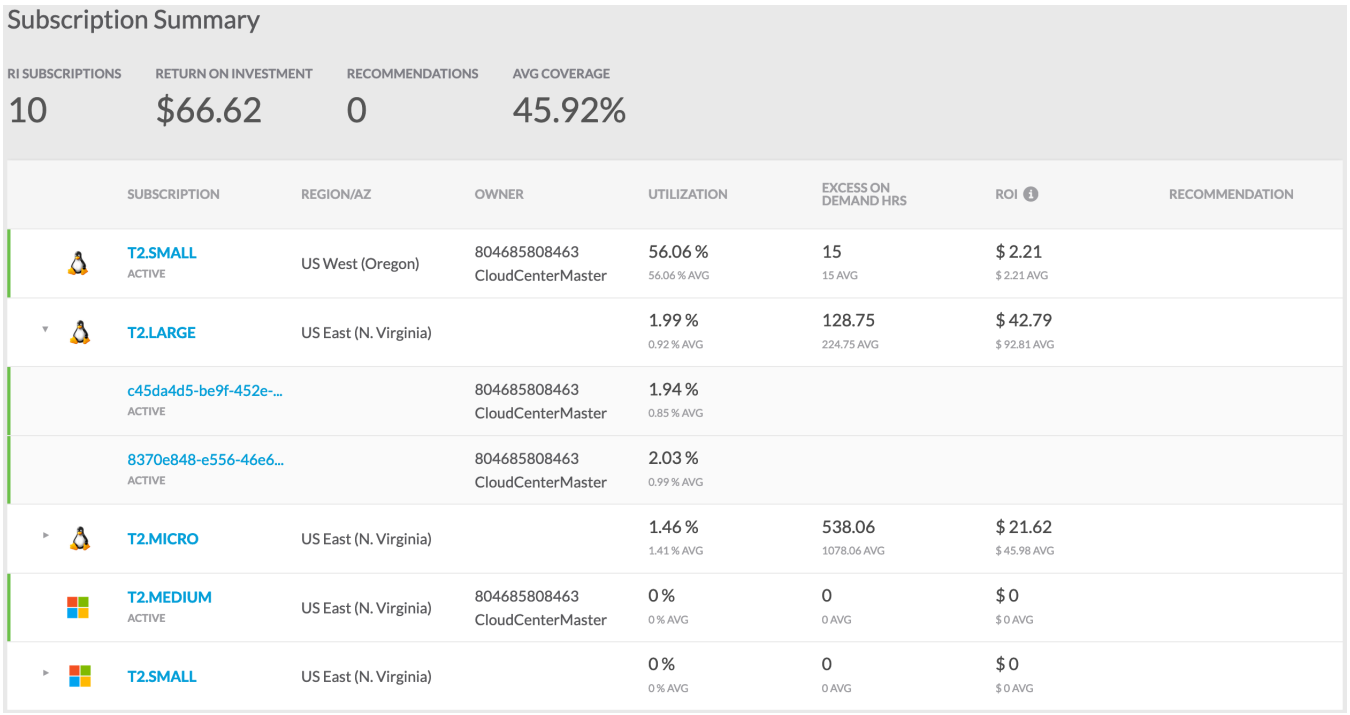

The significant fields in the report are explained in the following table.

Field	Description
RI Subscriptions	Number of RI subscriptionsfor the specified filter criteria (Accounts, Regions, Instance Types).
ROI	Average ROI achieved till date by the subscriptions for the period this report is generated. The value in smaller font indicates the ROI for the subscription.
Recommendations	Number of recommendations based on the filter criteria.
Download	Downloads the report in a .csv format.
Date Range	Choose a range to display the report.
Schedule	Allows you to send the report via email to recipients on the fixed date.
Subscription	RI Subscriptions are grouped based on the following similarities: • Instance type • Operating system type • Region
Utilization	Subscription utilization, in percentage. In the above image highlighted in green, a value of 50% indicates that on an average only half the purchased subscriptions were utilized during the period (30-day) for which the report is generated. The value in smaller font indicates the average utilization of subscription since purchase.
Excess on Demand Hrs	Displays the available, unused hours for a subscription.
ROI	ROI for RI subscriptions purchased for an instance type. The value in smaller font indicates the average ROI of subscription since purchase.

Recommendation	Recommendation for the specific instance type or subscription. The options include: • Enable Float • Purchase another subscription • Increase Utilization • Renew • Fix Payment issue

Advanced Options

Scheduling Reports

The **Scheduler** icon allows you to schedule report generation periodically on a fixed date or at intervals. To create a schedule, do the following:

1. Click the **Scheduler** icon. The **Schedule New *Report Name*** dialog appears.

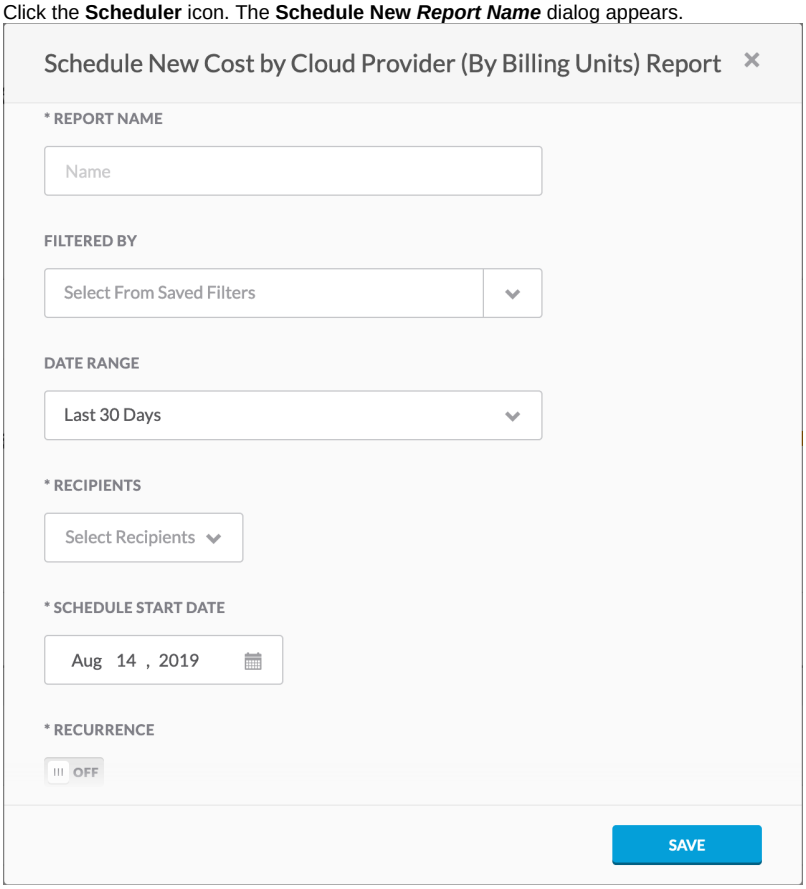

2. Do the following:
    a. Enter a name for the schedule.
    b. Choose filtering options for the schedule from the **Filtered By** field. The information in this field is populated when you save the filtering options as described in the *Advanced Filtering Options* section. You can choose to select a filter or leave the field empty.
    c. Choose the date range.
    d. Select the recipients the report must be sent to.
    e. Specify the start date.
    f. Toggle on the **Recurrence** button to send the report at intervals.
    g. In the **Repeats Every** area, specify the number of times the report must be sent to the recipients and choose the interval **Daily** or **Weekly** . if you choose Weekly, you can also specify the days of the week when the report is sent.
    h. Select the period to end the schedule. The options are:
        i. **Never** Send report forever or until the schedule is deleted.
        ii. **On** Date when the report should be sent.
        iii. **After** Number of occurrences after which the report is not scheduled.
3. Click **Save**. The report is displayed in the**Scheduled*Report Name***dialog as shown in the sample screenshot below.

Optionally, you can use the **Edit** option in the **Actions** column to make changes to the schedule. You can also delete the report using the **Delete** option.

Click **Done** to close the dialog.

The following screenshot displays a sample email format of the report.

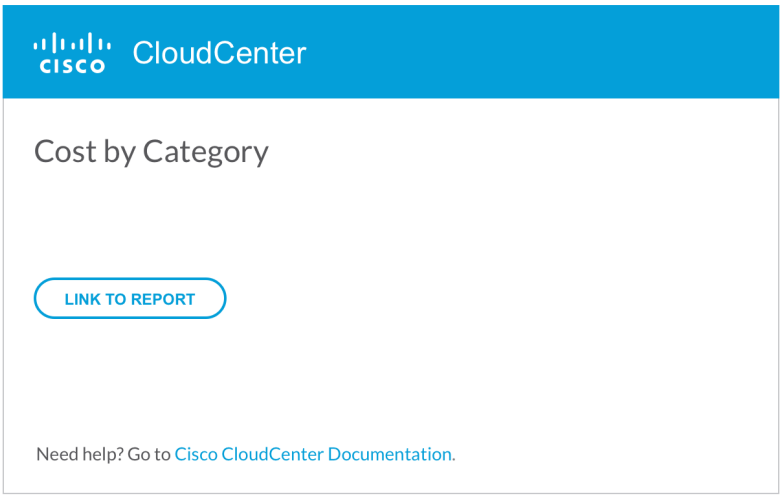

Date Range

The Date Range dropdown allows you to display costs and usage between time periods.The following table explains the date range options:

Date Range	Description
Month to Date	Report for the current month ending with today.
Last 30 Days (Default)	Report for the past 30 days ending with today.
Last Month	Report for the previous month.
3 Months	Report for the past 3 months ending with today.
6 Months	Report for the past 6 months ending with today.
1 Year	Report for a year ending with today.
CUSTOM	A custom range specified by the selected date, month, and year.

Custom Date Range

The Custom Date Range feature enables the ability to view cost reports and usage between specific date ranges. For example, you may want to ensure that a resource did not incur any costs during a specific interval (a specific week).

Choose the**CUSTOM**option in the date range dropdown to specify a range. When a date range is specified, the range appears in all reports that contain the **Date Range**dropdown. You can specify one date range only. However, you can modify the range multiple times. The date range is specific to a session and cannot be deleted, but is removed when you log out of the application or the application is reloaded.

The format of the dates in the dropdown is determined by the chosen range. The following table explains the formats for the ranges:

Date Range	Format in Date Range Dropdown
Greater than 1 day but less than 31 days	Days. Example: Jan. 29, 19 Feb 10, 19
Greater than 31 days but less than 31 months	Months. Example: Jan. 29, 19 Feb 26, 20
Greater than 31 months	Years. Example: Jan 29, 19 Apr 26, 21

To specify a custom date range, do the following:

1. Click the**Date Range**dropdown and choose**CUSTOM**. The**Custom Date Range**dialog appears.
2. To specify the start date, enter a date in the**From**field in the format displayed orclick the calendar icon.

   Click backward and forward arrows to choose months in a year.To choose a year, click the arrow next to the Month/Year title and select the year.

   To specify the end date, enter a date in the field in the**To**field in the format displayed orclick the calendar icon.
3. Click**Apply**.
4. Use the**Edit**icon to modify the range.

Subscription Details

You can drill-down the RI Subscriptions Report for the following to understand additional information:

- Group subscriptiondetails Opens when group name in the RI subscription report is clicked
- Individual subscription details (part of agroup) Opens when IDs listed under a group is clicked
- Individual subscription details (not part of agroup) Opens when IDs that are not listed under a group

Group Subscription Details

The following is a sample screenshotthat displays detailed information about group subscriptions. Click the down arrow adjacent to the logo to expand display the individual subscriptions in that group.This page enumerates the amount of savings that an organization can achieve by acting on the recommendations.

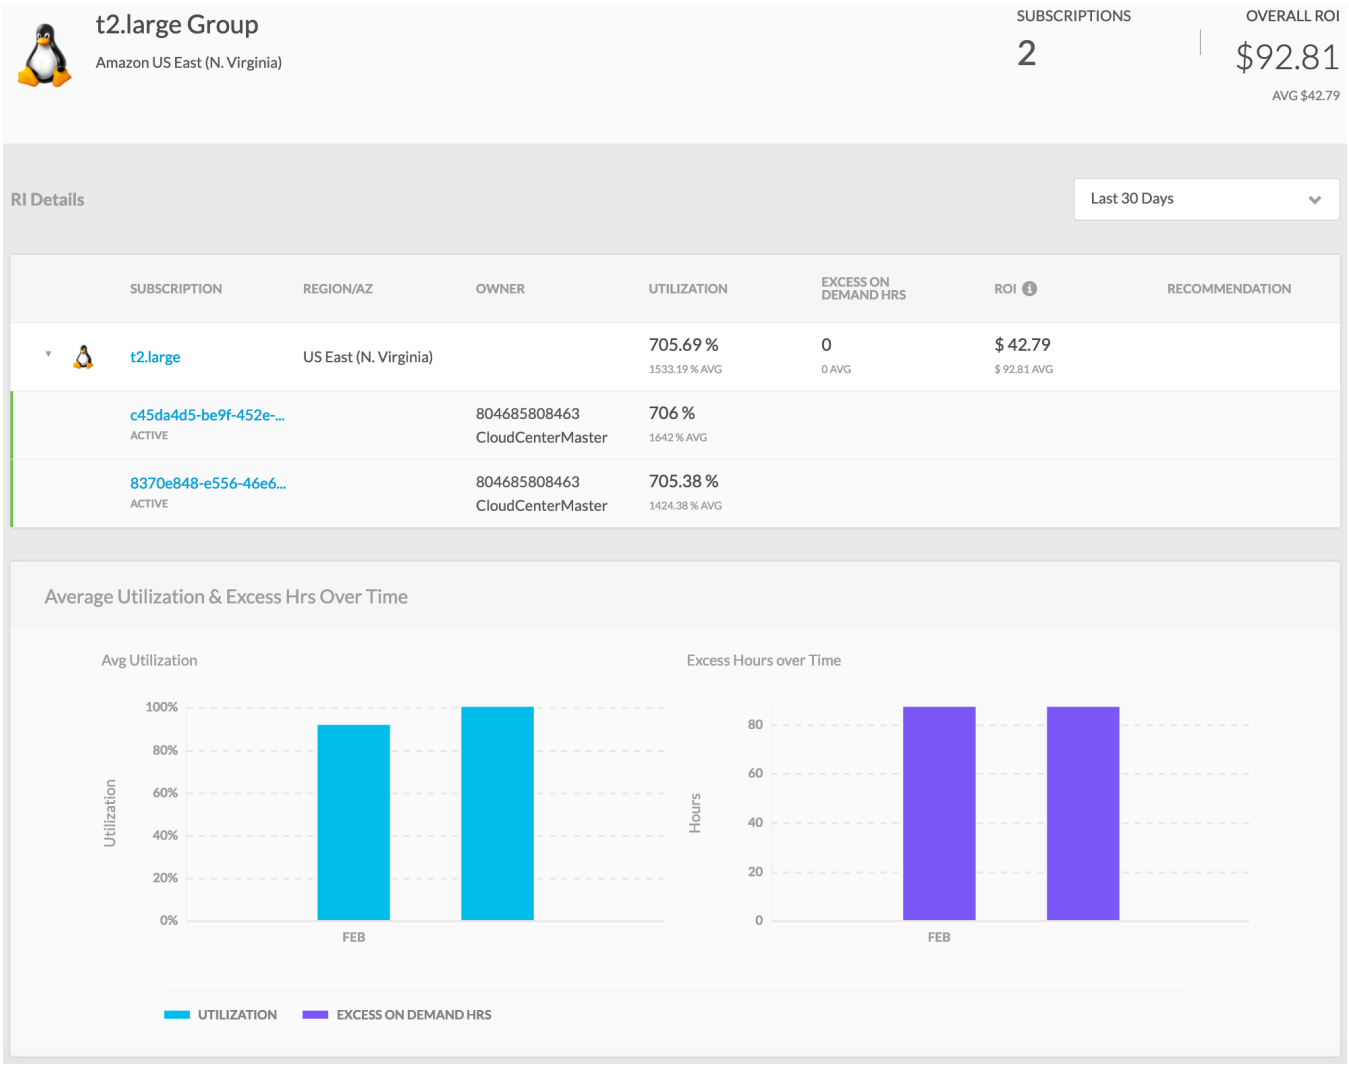

The**Average Utilization and Excess Hours over Time** report display information about the group utilization, in percentage, and the additional on-demand subscription that was purchased for the instance.

Individual Subscription

Individual subscriptions can either be included in a group or as standalone subscriptions (not part of any group). The following screenshot is an example of the individual subscription that is part of a group. This page appears when you click any individual subscription listed under a group.

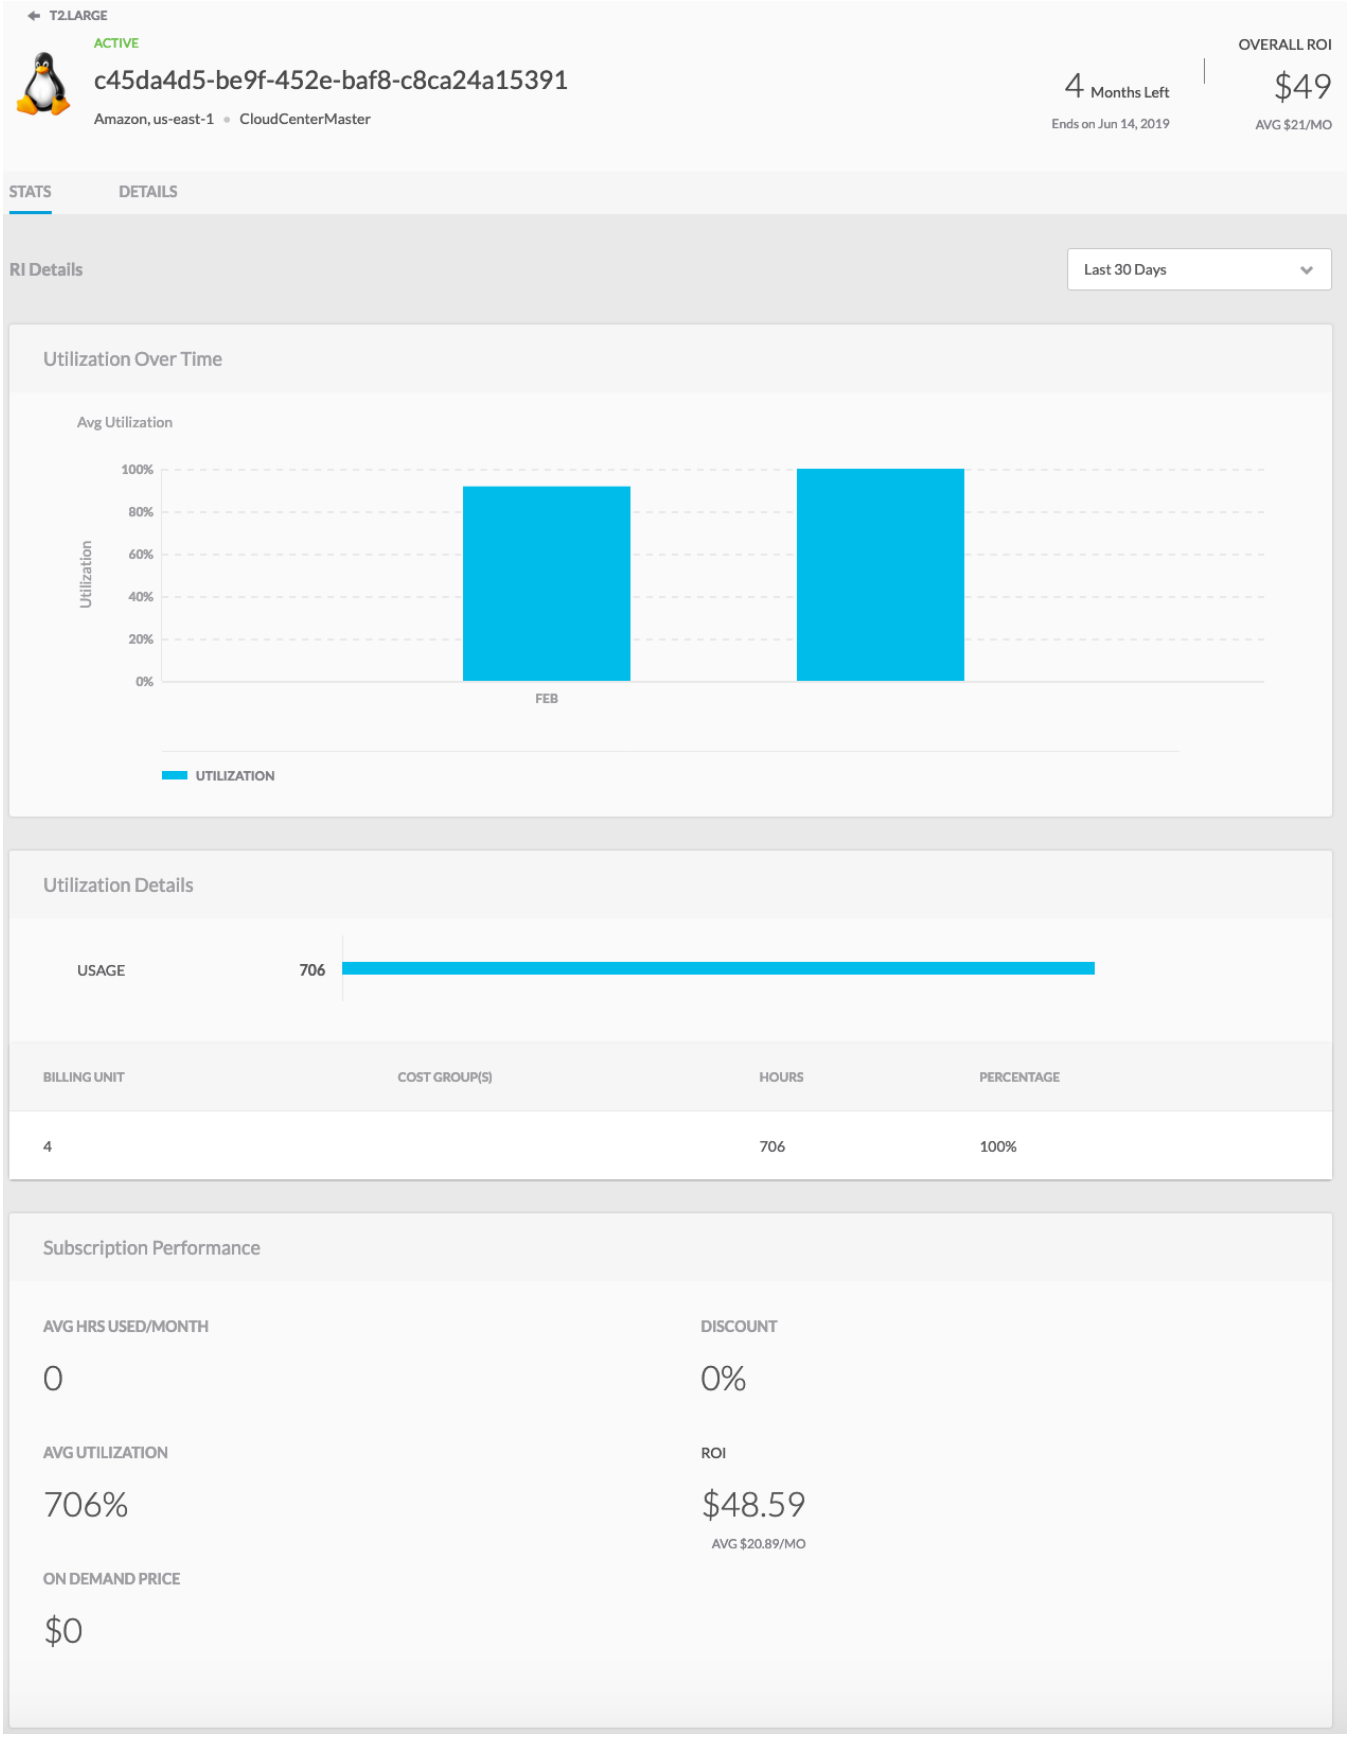

The information of an individual subscription is displayed in two tabs: **Stats** and **Details**.

The following table explains significant fields in the **Stats** tab:

Area	Field	Description
Subscription Header	Status	Color-coded state of the subscription Active or Expired. The information in the right indicates the remainder period of a subscription.
	Instance Details	• Hostname Thehostnamefor the VM, if configured, else thenode IDis displayed. • Cloud Region • Billing Unit ID
	Remainder Duration	Duration, in months or years, before the subscription expires.
	Overall ROI	Total ROI achieved by this subscription.
Utilization over Time	Date	Graphical display of past utilization for the subscription for a chosen period (from the **Date** drop-down field).
Utilization Details	Hours	Graphical display of instance utilization against various billing units associated with the instance.
Subscription Performance	Average Hours Used/Month	Number of instances available as of date.
	Discount	Discount applicable to the instance if the recommended subscription type is purchased.
	Average Utilization	Instance utilization, in percent, for the chosen period.
	ROI	Total ROI achieved by the subscription and the average ROI achieved per month.
	On-Demand Price	On-Demand price for the instance if the subscription had not been purchased. See Reserved Instances Overview > Payment Methods.
	Average Excess On-Demand Hours per Month	On-demand purchases made in addition to RI subscription purchase. ⚠ This field appears for an individual subscription not part of any group.

The **Details** tab presents information about the subscription scope, term, type.

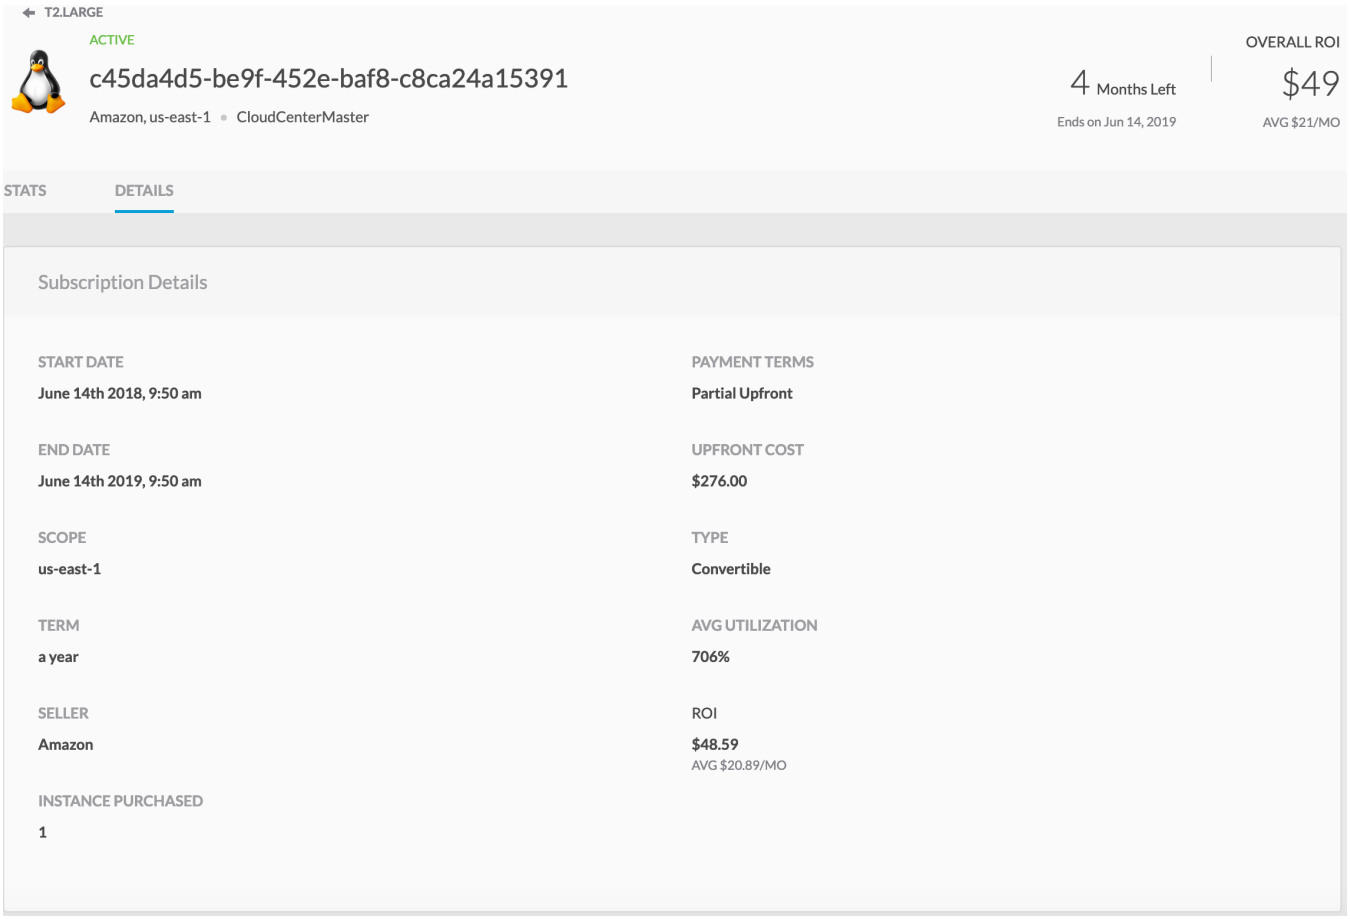
T2.LARGE
ACTIVE
c45da4d5-be9f-452e-baf8-c8ca24a15391
Amazon, us-east-1 • CloudCenterMaster
4 Months Left
Ends on Jun 14, 2019
OVERALL ROI
$49
AVG $21/MO
STATS
DETAILS
Subscription Details
START DATE
June 14th 2018, 9:50 am
END DATE
June 14th 2019, 9:50 am
SCOPE
us-east-1
TERM
a year
SELLER
Amazon
INSTANCE PURCHASED
1
PAYMENT TERMS
Partial Upfront
UPFRONT COST
$276.00
TYPE
Convertible
AVG UTILIZATION
706%
ROI
$48.59
AVG $20.89/MO

RI Opportunities Report

RI Opportunities Report

- Overview
- Purchase Report
- Savings Report
    - Scheduling Reports

Overview

The **RI Opportunities Report** provides informationabout the number of new RIs that must be purchased and the potential savings that can be achieved as a result of the purchase. This report contains the following:

- Purchase report
- Savings report

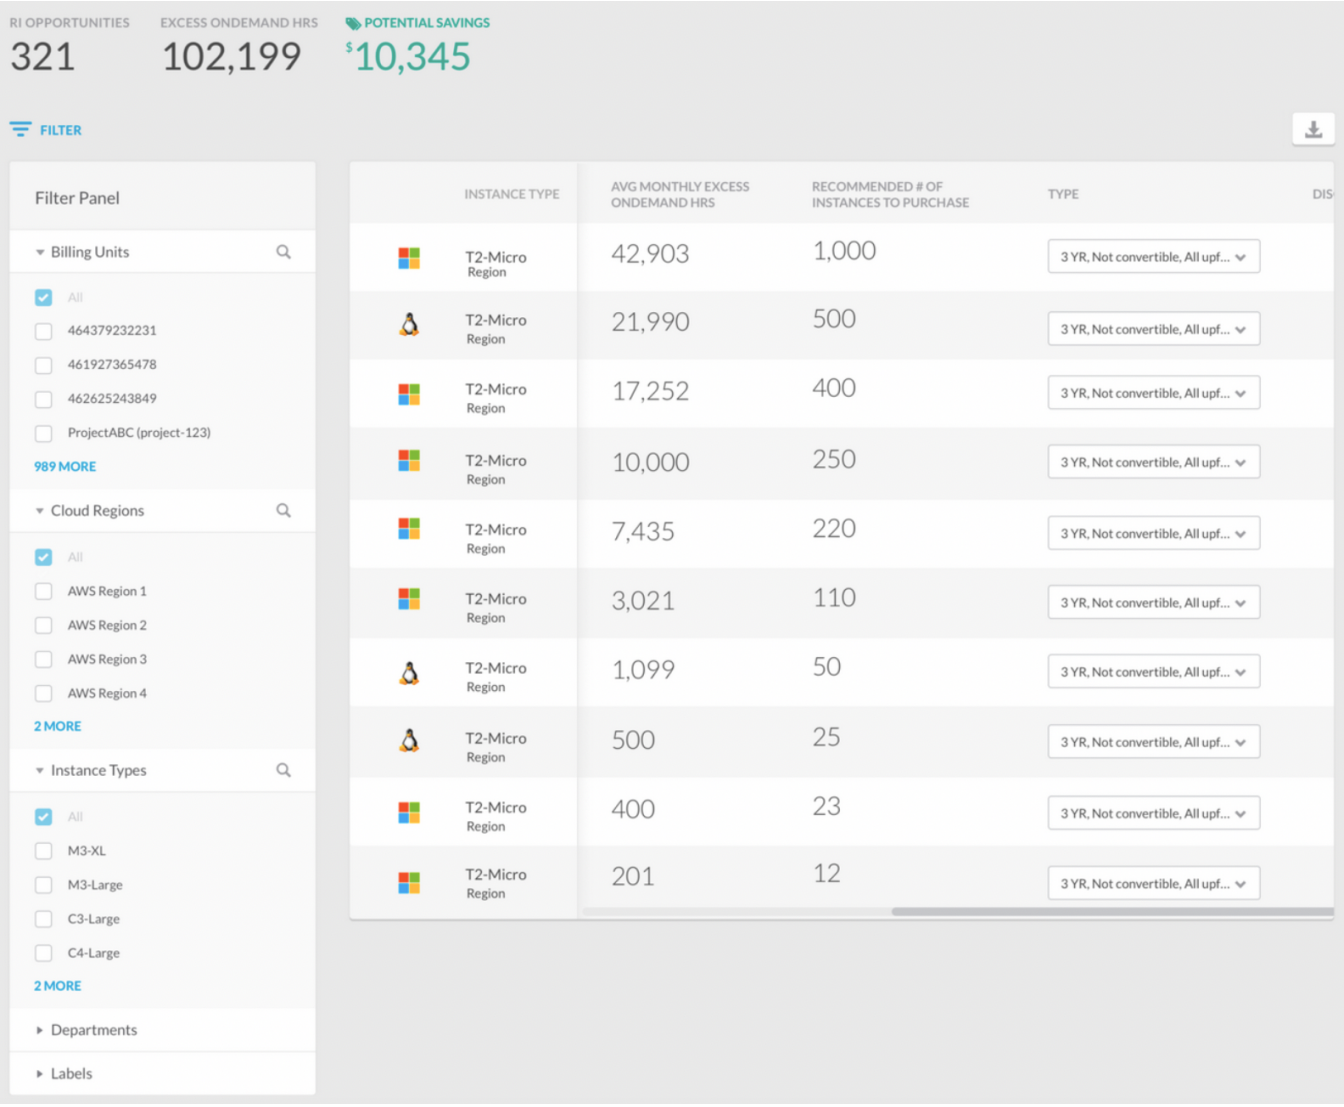

Purchase Report

This report appears as a header and provides an overview of the total RIs to be purchased, the total number of on-demand hours (across instance types), and potential savings.

Savings Report

This report presents the following data in a tabular format.You can filter this table using the **Filter** option.

- Individual instance types to be purchased
- Number of on-demand hours that the instance types are currently running
- Recommended instances to be purchased
- Type of RIs that can be purchased. This is a dropdown list. Depending on your selection, the savings and discount (in the **Discounts** field) vary.

The following table explains the significant aspects of the Savings Report.

Aspect	Description
RI Opportunities	The opportunities present across the organization.
Excess On Demand Hours	Instance hours that ran at on-demand rates, not covered by RI subscription.
Potential Savings	Savings that can be achieved by moving to RI.
Schedule	Allows you to send the report via email to recipients on the fixed date.
Download	Downloads the report in a .csv format.
Average Monthly Excess On Demand Hours	Number of on-demand hours that were purchased each month for an instance.
Type	Recommended instance type and payment method.
Discount	Discount that can be availed as of date if the recommended RI is chosen.

Scheduling Reports

The **Scheduler** icon allows you to schedule report generation periodically on a fixed date or at intervals. To create a schedule, do the following:

1. Click the **Scheduler** icon. The **Schedule New *Report Name*** dialog appears.

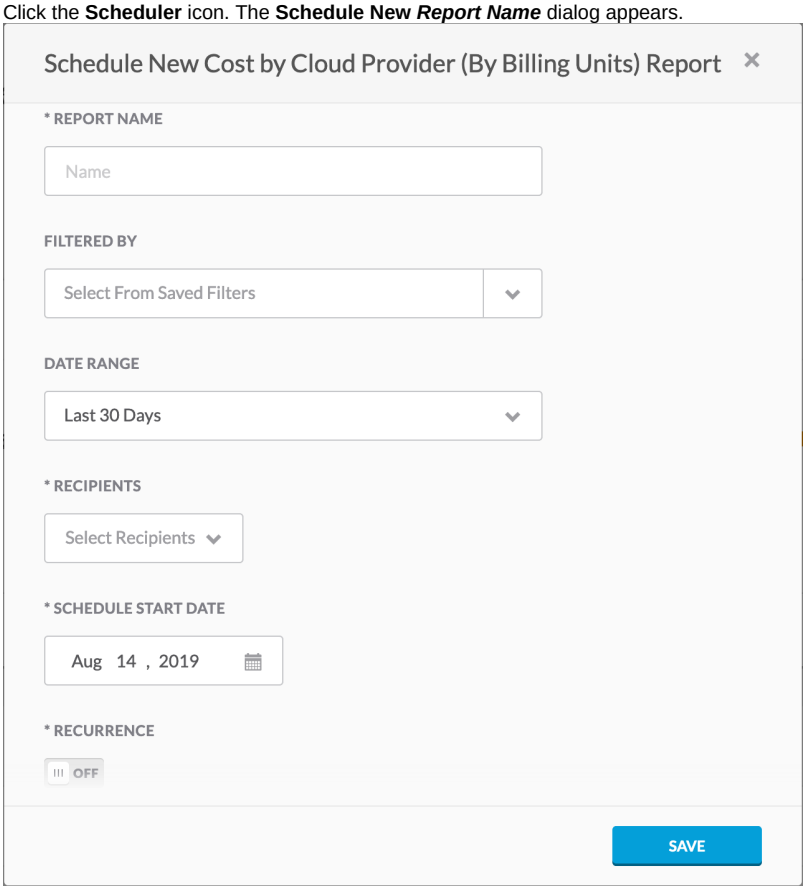

2. Do the following:
    a. Enter a name for the schedule.
    b. Choose filtering options for the schedule from the **Filtered By** field. The information in this field is populated when you save the filtering options as described in the *Advanced Filtering Options* section. You can choose to select a filter or leave the field empty.
    c. Choose the date range.
    d. Select the recipients the report must be sent to.
    e. Specify the start date.

f. Toggle on the **Recurrence** button to send the report at intervals.
g. In the **Repeats Every** area, specify the number of times the report must be sent to the recipients and choose the interval **Daily** or **Weekly** . if you choose Weekly, you can also specify the days of the week when the report is sent.
h. Select the period to end the schedule. The options are:

   i. **Never** Send report forever or until the schedule is deleted.
   ii. **On** Date when the report should be sent.
   iii. **After** Number of occurrences after which the report is not scheduled.

3. Click **Save**. The report is displayed in the**Scheduled*Report Name***dialog as shown in the sample screenshot below.

Optionally, you can use the **Edit** option in the **Actions** column to make changes to the schedule. You can also delete the report using the **Delete** option.

Click **Done** to close the dialog.

The following screenshot displays a sample email format of the report.

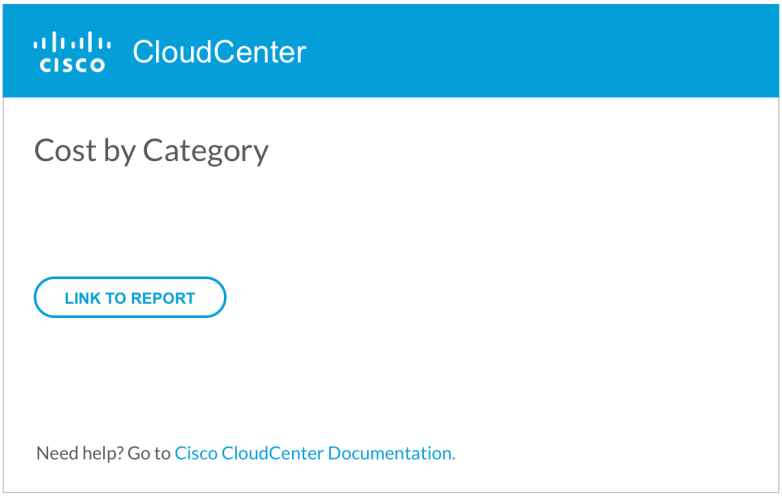

Administration

Administration

- Admin Tasks in Cost Optimizer
- Settings Page
- Data Collection
- Alerts Page
- Tag-Based Cost Reporting

Admin Tasks in Cost Optimizer

Admin Tasks in Cost Optimizer

Overview

Navigate to the**Admin**menu toperform the following administrative tasks:

- Configure Clouds- Set up clouds and cloud accounts.
- Define Settings- Enable recommendations.
- Allocate Budgets- Set up budgets for a financial year or quarter.
- Define Alerts- Send notifications to specific users or user groups.
- Tag-Based Cost Reporting - Enable cost reporting for AWS and Azure tags.

The availability of the**Admin** menu is subject to roles and permissions; it is visible to administrators only. For more information, seeAccess and Roles.

Settings Page

Settings Page

- Rightsizing Card
- Suspension Candidates Card
- RI Opportunities Card
- Historical Collection Card
- Metrics Collection Card
- Fiscal Year Card

Rightsizing Card

The **Rightsizing card**allows you to define rightsizing thresholds in Cost Optimizer.

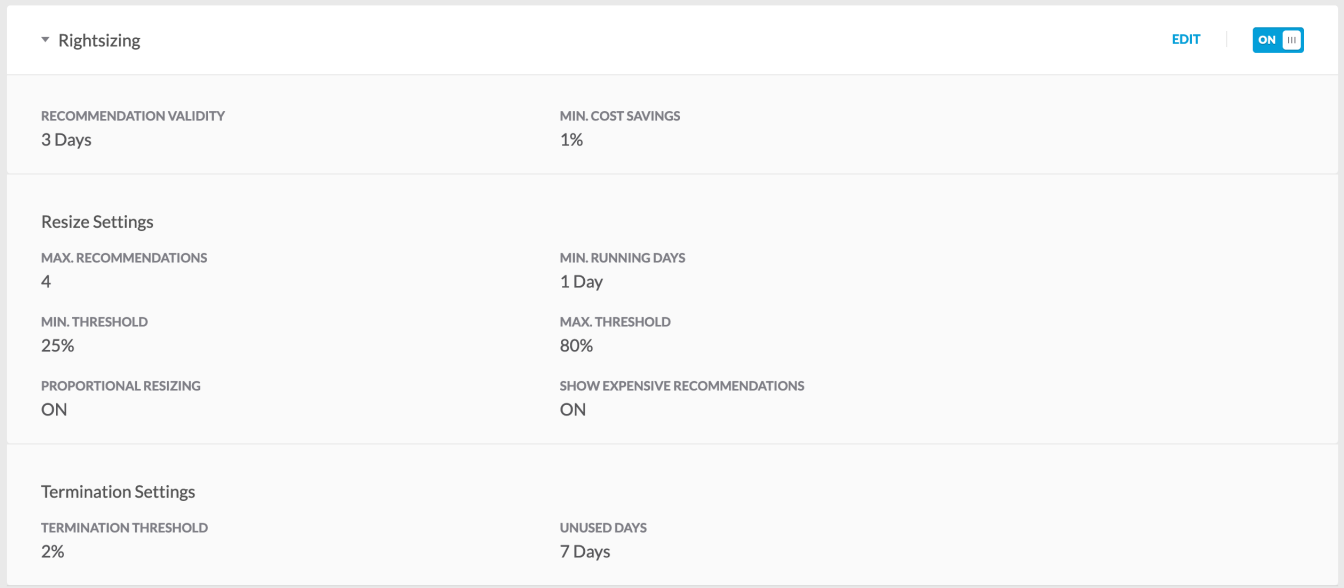

The following table explains the fields in the **Rightsizing** card of the**Settings** tab.

Field	Description
Recommendation Validity	Duration, in days, a recommendation is valid or visible in Cost Optimizer.
Min Cost Savings	Minimum savings offered by recommended downsize alternatives. Instances that do not meet this criterion are not recommended.
Resize Settings	
Max Recommendations	Desired number of alternate recommendations.
Min Running Days	Duration, in days, an instance must be run for rightsizing analysis and recommendations.
Min Threshold	Downsizing recommendation. If a CPU utilization is below the threshold specified in this field, small instances are recommended.
Max Threshold	Upsizing recommendation. If a CPU utilization crosses the threshold specified in this field, large instances arerecommended.
Proportional Resizing	Resizes the memory in proportion to CPU usage.
Show Expensive Recommendations	Enables upsizing recommendation. Toggling this option provides recommendations for overutilized instances (see Rightsizing > *Recommendations > Overutilized Tab*) even though the recommendation does not result in saving costs.
Termination Settings	
Termination Threshold	Termination recommendation. If the utilization is lesser than the threshold specified in this field, it is recommended that the instance be stopped or terminated.
Unused Days	Duration, in days, an instance must be in an unused state to be recommended to stop or terminate the instance.

Suspension Candidates Card

The**Suspension Candidates card**allows you to define suspension thresholds in Cost Optimizer.

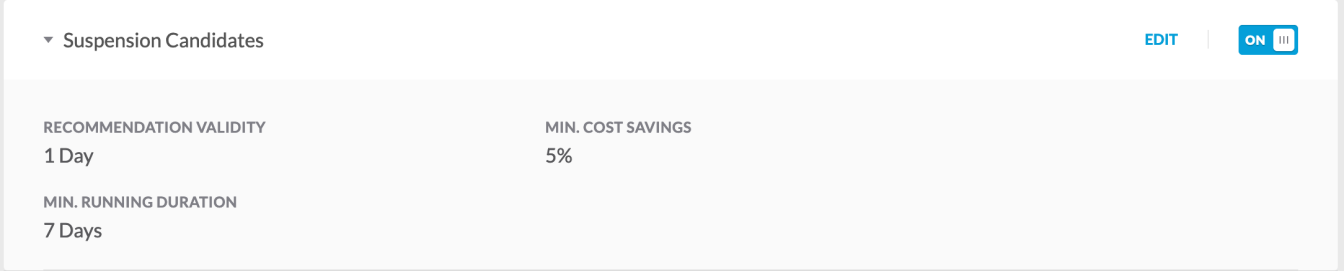

The following table explains the fields in the**Suspension Candidates** card.

Field	Description
Recommendation Validity	Duration, in days, a recommendation is valid or visible in Cost Optimizer.
Min. Cost Savings	Minimum savings offered by suspending a resource. Resources that do not meet this criterion are not recommended.
Min. Running Duration	Duration, in days, an instance should be running. If the instance is running below the threshold specified in this field, it is recommended that the instance is suspended.

RI Opportunities Card

Use the **RI Opportunities** card toenable or disable reserved instances recommendations.

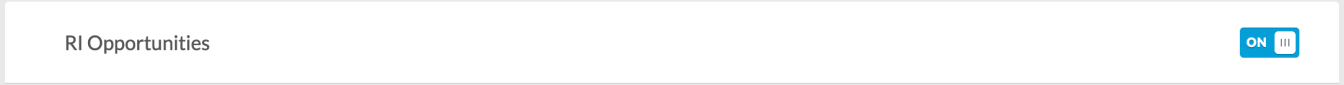

Historical Collection Card

The **Historical Collection** card allows you to specify in the **Number of Days** field the duration, in days,for which historical data must be collected.

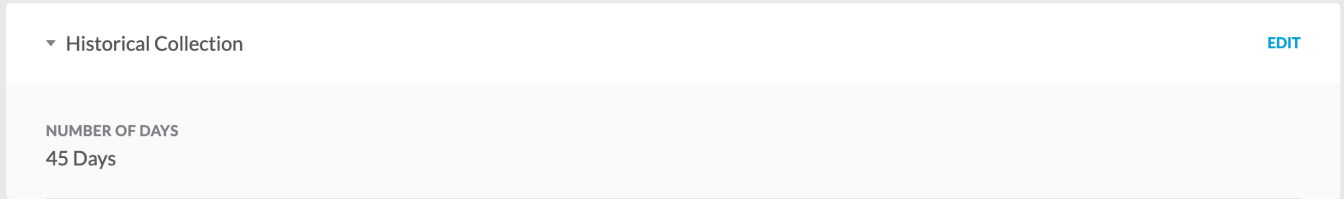

The default is 60 days. To change the value, click **Edit** and enter a value between *0* and *100* days in the **Number of Days** field.

Metrics Collection Card

The**Metrics Collection**card allows you to specify in the**Number of Days**field the duration, in days,for which metrics data must be collected.

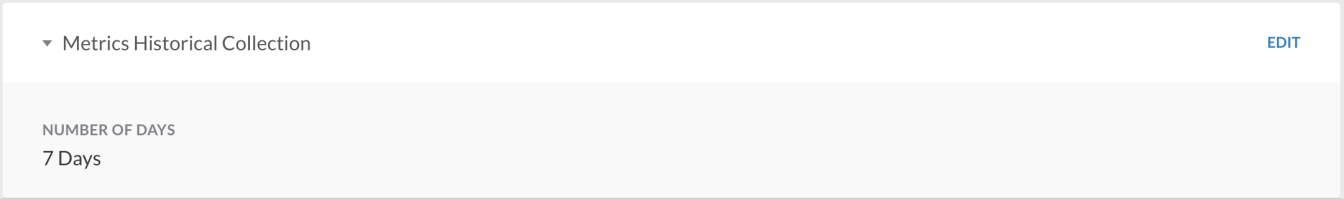

The default is 7 days. To change the value, click**Edit**and enter a value between*0*and*100*days in the**Number of Days**field.

Fiscal Year Card

Use the **Fiscal Year** card to define a financial year and quarters in the financial year for budget allocation.

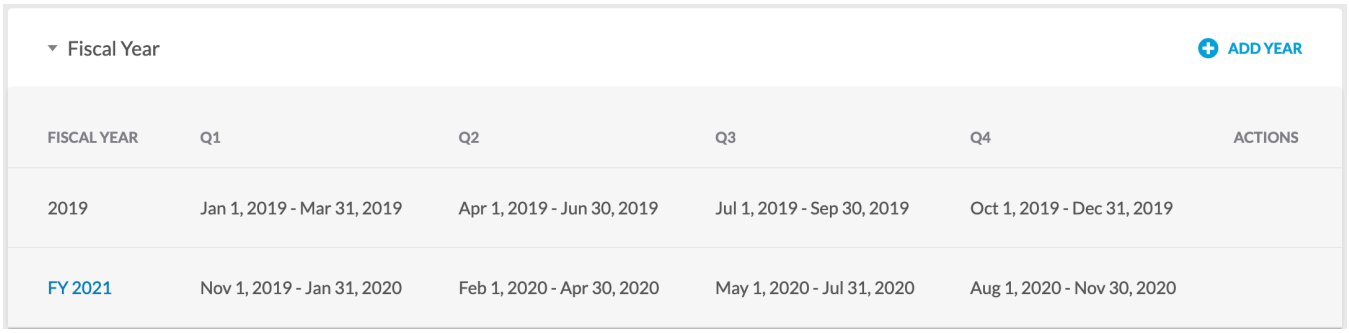

An Optimizer Admin only can define a fiscal year. SeeAccess and Roles.

Perform the following steps to add a fiscal year:

1. Click **Add Year**. The **Fiscal Year** dialog appears.

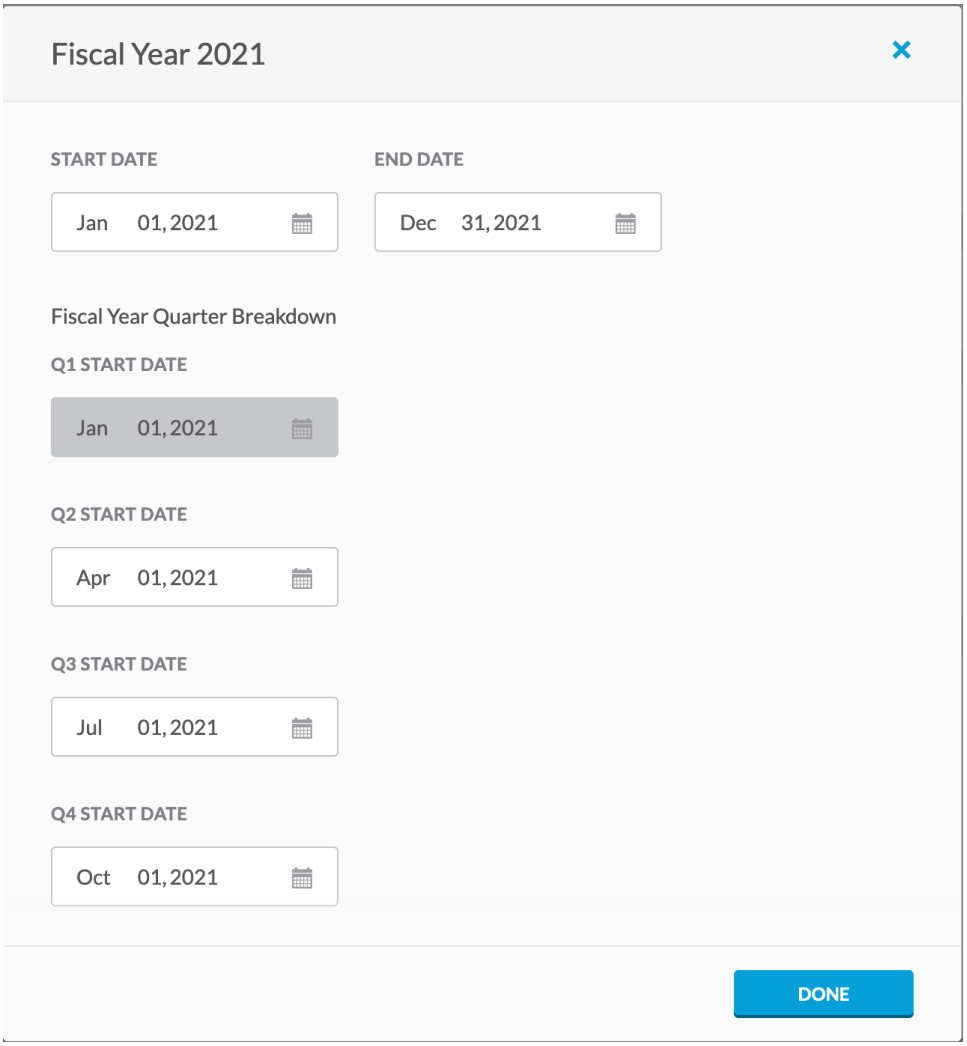

2. Select an appropriate value in the **Start Date** and **End Date** for the fiscal year and accept the quarter breakdown.
3. Click **Done**.

Data Collection

Data Collection

Cost Optimizer runs background processes to collect data from a cloud provider for reporting and analysis. The processes are scheduled at specific intervals to connect with a cloud provider to receive the latest information. The following table lists the background processes and corresponding schedules.

Process	Description	Schedule
Inventory Collector	Inventory details from a cloud provider (Virtual Machines, Storage Volumes, Load Balancers, Database, Containers)	Every 15 minutesstarting at 0, 15, 30, and 45 minutes of the hour
Cost Calculator	Resource cost	Every hour
Metrics Collector	Basic metrics (CPU)	Every 30 minutes at 20 and 50 minutes of the hour
	Advanced metrics (Network and IOPS-related)	Every 6 hours
Invoice Aggregation	Billing information from a cloud provider	Daily at 04:00 hours GMT
Rightsize Analyzer	Analyses resources to identify resource wastage	Daily at 01:00 hours GMT
Reservation Analyzer	Analyzes all AWS RI resources	Daily at02:00 hours GMT
Metadata Sync	Cloud provider metadata information for public clouds (regions, zones, instance types, rate card, and, so on)	Daily at00:00hours GMT
Tag Sync	Fetch tags from all clouds. For AWS, collects tags which are enabled in AWS console.	Daily at 01:00 hours GMT
Unused Volumes Analyzer	Fetch details of unused volumes	Daily at 01:00 hours GMT
Suspension Policy Analyzer	Fetch details of suspension policies	Daily at 03:00 hours GMT

Alerts Page

Alerts Page

- Overview
- Budget Alerts
- Trend Alerts

Overview

The **Alerts** tab allows you to send notifications to specified users or user groups when the threshold limits cross the limits as mentioned in the page. You canmodify the thresholds by using the **Edit** button in the header. Notifications are sent via the SMTP settings in Suite Admin. For more information about SMTP settings, seeEmail Settings.

Budget Alerts

Budget alerts compare expenses against the budget allocatedin the current quarter.Budget Alerts are applicable globally at a tenant level.These settings can be overridden on a per cost group level when creating a budget through the**Alert Settings**tab(seeAllocate Budgets> *Creating a budget*).

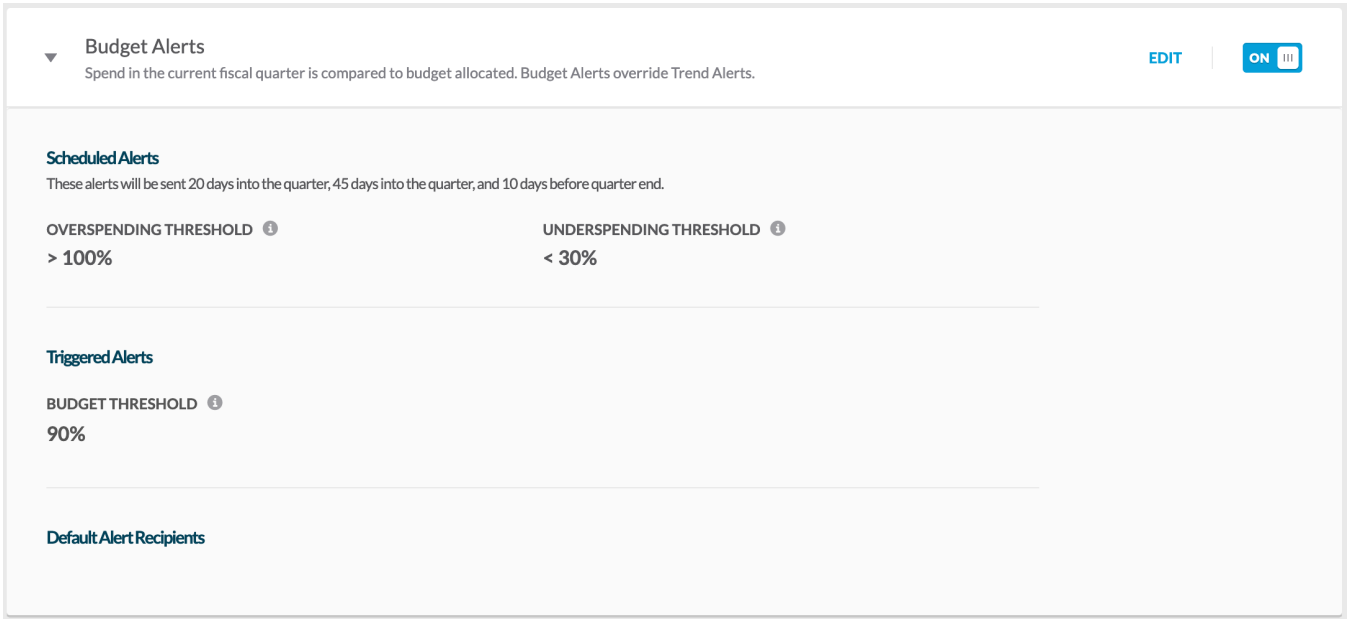

Trend Alerts

Trend Alerts compare the expenditure and cost in the current quarter against the last quarter.

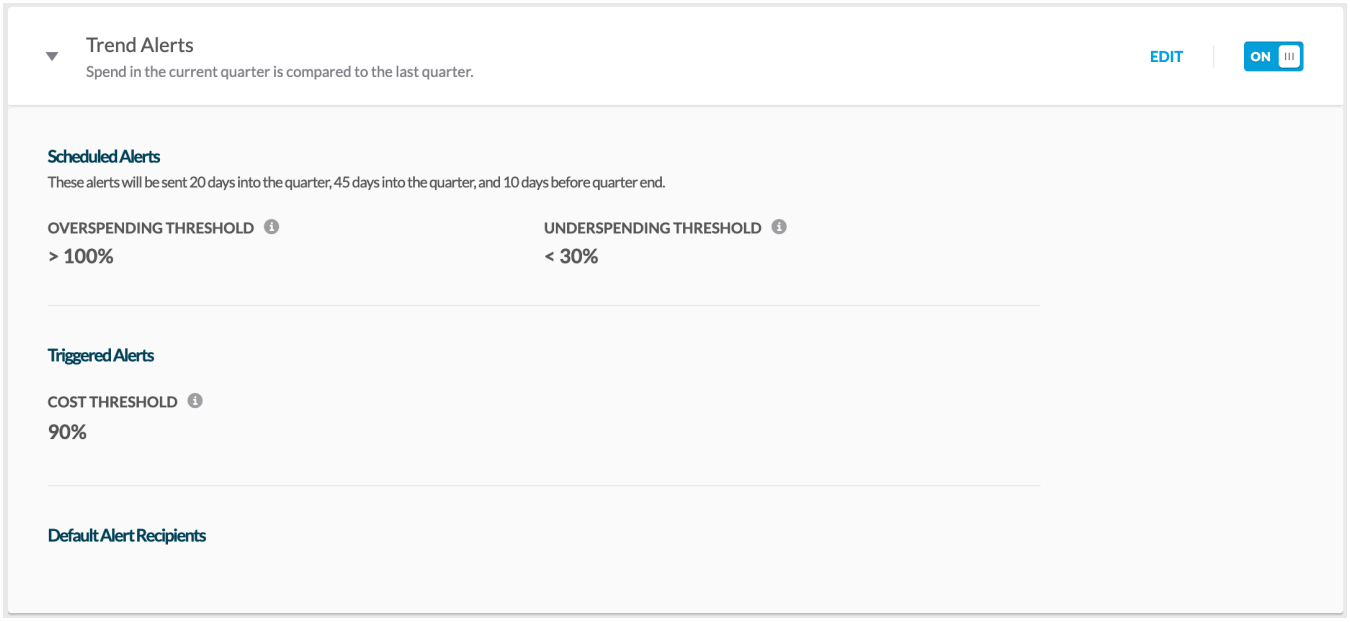

The following table explains the significant fields in the **Alerts** page.

Field	Description
Overspending Threshold	A cost group is considered an overspender if its forecasted total expenditure for a duration exceeds this threshold. An alert is sent at specific intervals about overspending cost groups.
Underspending Threshold	A cost group is considered an underspender if its forecasted total expenditure for a duration is below this threshold. An alert will be sent at specific intervals about underspending cost groups.
Triggered Alerts	An alert is sent the day the cost group's expenditure reaches the defined cost threshold.
Default Recipients	Persons who receive alert notifications.

Tag-Based Cost Reporting

Tag-Based Cost Reporting

- Overview
- Enable Cost Reporting

Overview

Tags are key-value pairs associated with cloud resources on a cloud provider. The key is mandatory and value is optional. Tags can be user-definedor system-defined.Similar to billing units, tags are also used for cost breakdown at a deeper granular level.Use the Tag-Based Cost report to enable cost reporting for tags in the associated cloud accounts. Tag-based cost reporting is disabled by default.

> ⚠ Tag-based cost reporting is available for AWS and Azure clouds only. You can share tag-based cost groups. When tag-based cost reports are shared, the sharing results in displaying additional cost, inventory, and recommendations for the resources associated with the cost groups.

The tags are fetched automatically through theTag Syncbackground process that runs every day at 01:00 hours GMT. SeeData Collection.

Enable Cost Reporting

You may incur an expense on your billing account (for example, on AWS the expenses are recorded as Cost Explorer expenses) when tag-based cost reporting is enabled.By enabling this feature and by enabling the tags.To enable tag-based cost reporting, do the following:

- Set the toggle to **ON** the button at the top of the page.
- Enable individual tags (pertaining to AWS and Azure clouds) to fetch the invoices.

When tag-based cost reporting is toggled **ON** (at the top of the page), you would notice thatindividual Azure tags are OFF by default, whereas few individual AWS tags are enabled by default. The AWS tags that are On by default are the incurring costs tags. To enable tag-based cost reporting for Azure tags, you must set the toggle to**ON**in the**Cost Reporting**column against the tag for which you want to display the cost. Additionally, you can set the toggle to**ON**for AWS tags, as appropriate.

If the individual tags are not toggled to ON, no invoices are fetched.Invoices are collected through the Metrics Collectorbackground process (see Data Collection) 5 minutes after the individual tags are enabled. Invoices are fetched for tags that are enabled in the 5-minute window. If tags are enabled after 5 minutes, the invoices for these tags will be in the next invoice collection schedule. However, it is recommended that you enable AWS tags 24 hours in advance.

Individual tags are listed in the **Unassigned Tags** list for association with a cost group.SeeHow Do I...>*Associate a Tag* for additional context.

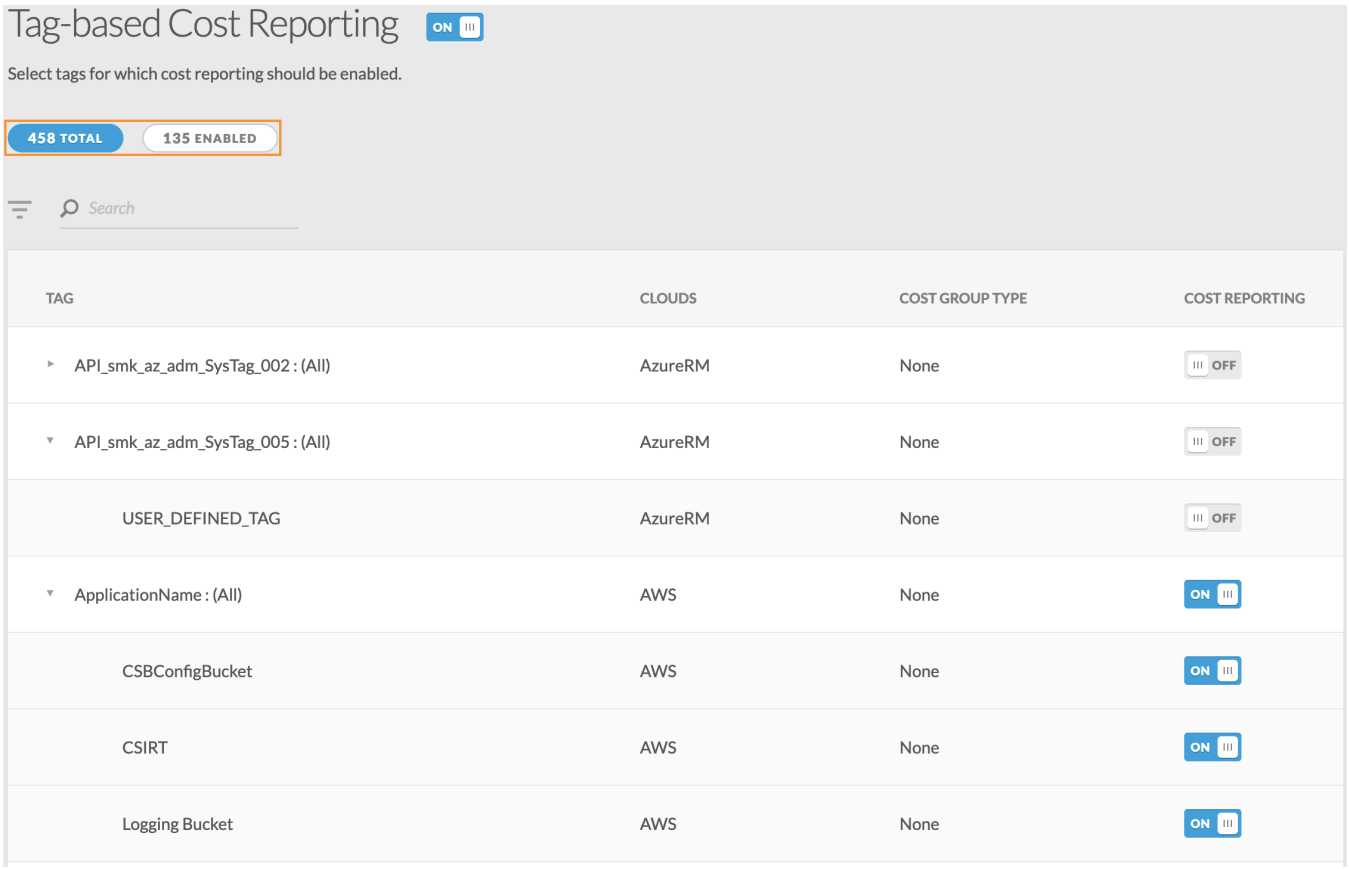

The following table identifies various aspects of the page.

Summary	Description
All	Number of tags the Tag Sync background process fetched from the cloud provider.
Enabled	Number oftags that have been enabled for cost reporting.
Filter	Filter tags on cloud providers.
Search	Search for tags from the list.
Tag	Tag name.
Cloud	Cloud provider that the tag belongs to.
Cost Group Type	Cost group type the tag is assigned to. See:How Do I...>*Associate a Tag*.
Cost Reporting	Toggle this option does the following: • Display tags in the Unassigned Tags area in the Cost Groups page (see Cost Groups Configuration). These tags can be associated with a cost group. • Display the cost report for this tag in the Cost by Tags in a Cloud Group report. • Lists the tag in the various filtering panel in Cost Optimizer.

Troubleshooting

Troubleshooting

- Cost Optimizer Troubleshooting
- Scheduling MongoDB

Cost Optimizer Troubleshooting

Cost Optimizer Troubleshooting

- Adding a CloudAccount
- Costs for Private Clouds are not Displayed
- Cost by Organization Hierarchy Report is not Displayed
- Incorrect Numbers in Cost by Cost Group Type Report
- Inventory Types not Displayed
- No Rightsizing Recommendations are Displayed
- Kubernetes Troubleshooting

Adding a CloudAccount

If you are unable to add a cloud account, ensure that the credentials are valid andvalidate that the user or role assigned with the credentials has the correct permissions (see Cloud Overview > *Minimum Permissions for Public Clouds*). If the cloud account is a cloud master account, all child accounts must have AWS IAM role as *Optimizer*. This role must have the same permissions as described in the *Minimum Permissions for Public Clouds* section.

Costs for Private Clouds are not Displayed

If the Cost Optimizer Dashboard does not display costs for private clouds (vCenter, OpenStack, and Kubernetes), verify the following:

- Price is specified in the **Price** field when adding instance types.
- Costis entered in the **Cost** field whenadding storage types.

SeeInstance Types Settings andStorage Types Settingsfor additional context.

Cost by Organization Hierarchy Report is not Displayed

If the Cost by Organization Hierarchy report (applies toMasterAWSand GCP accounts only) is not visible, ensure thatthe**Enable Reporting By Org Structure**istoggled to**On**when adding a cloud. SeeConfigure an AWS Cloudfor more details. You must set the toggle to**On**to cause Cost Optimizer to import the cost hierarchy created in the cloud provider portal.

Incorrect Numbers in Cost by Cost Group Type Report

If the Cost byCost Group Type (Department) reportdisplays incorrect costs for a specific cost group type, verify that the billing units are mapped correctly to the cost group type.

Inventory Types not Displayed

Inventory is only collected for regions that are explicitly added while setting up clouds. If you do not see expected inventory types (VMs, Kubernetes Workloads, Storage Volumes, and Services), verify that regions you added contain the inventory on the cloud.

No Rightsizing Recommendations are Displayed

Rightsizing recommendations are governed by the Settings tab in the Admin area. Review the values in this tab, in particular, review the value set for the **Min. Running Days** field.

Kubernetes Troubleshooting

Based on the error message that you see in the UI, you could perform basic troubleshooting steps if you have access to both the Kubernetes setup and to the CloudCenter Suite:

Issue	Error Reference Location
Errors returned by the Kubernetes cluster	Go to the Kubernetes dashboard and look for the event messages and login to the pod that you created for the CloudCenter Suite.
Kubernetes cluster API interaction issues	Login to Kibana (Monitor Modules > *View Logs in Kibana*) and look for error messages in logs with the text "cloudcenter-blade".

Orchestration or lifecycle issues	Login to Kibana (Monitor Modules > *View Logs in Kibana*) and look for error messages in logs with the text "cloudcenter-cco". You may find the following warning message in the Kubernetescloudcenter-cco logs you can safely ignore this message as it does not impactproduct functionality. `WARNING!!! The linux bootstrap URL might be valid: http://build-rel.cliqr.com/..../bootstrap-cliqr-init.sh. If Workload Manager cannot access the file, all deployments would fail!`
Model, manage, deploy issues	Login to Kibana (Monitor Modules > *View Logs in Kibana*) and look for error messages in logs with the text "cloudcenter-ccm-backend" or "cloudcenter-cloud-setup".

For additional details, refer to the following documents:

- Container Clouds
- Configure a Kubernetes Cloud

Scheduling MongoDB

Scheduling MongoDB

- Introduction
- Affinity and Toleration
- Running the MongoDB on aNew Node

Introduction

MongoDB is a shared component in Workload Manager and Cost Optimizer modules of CloudCenter Suite.Kubernetes schedules the MongoDB pod as any other pod to share resources (CPU, memory) with other pods. However, in a large setup and over a period of time, MongoDB might want to consume additional resources but could be limited by its peer pods.

This section provides guidance on how to configure the Kubernetes cloud environment to run the MongoDB pod on a new node.

Affinity and Toleration

The following values are definedon the MongoDB pod.

- **Toleration**
  - Key:cloudcenter/dedicated
  - Value:cloudcenter-mongodb
- **Affinity**
  - nodeaffinity:preferredDuringSchedulingIgnoredDuringExecution
- **Node Label**
  - Key: cloudcenter/purpose
  - Value: cloudcenter-mongodb

Running the MongoDB on aNew Node

Perform the following steps to run the MongoDB pod on a new node.

1. Add a new node and label it.

   ```
   kubectl label node NAME cloudcenter/purpose=cloudcenter-mongodb
   ```

2. Apply a taint to assign the pod to MongoDB as shown in the below example.

   ```
   kubectl taint node -l cloudcenter/purpose=cloudcenter-mongodb cloudcenter/dedicated=cloudcenter-mongodb:
   NoSchedule
   ```

3. Delete the pod to restart the MongoDB on this node.

   ```
   kubectl delete po cloudcenter-shared-cloudcenter-mongodb-0
   ```

⚠ It is not recommended to run the above steps on Amazon EKS.

Refer to these links for additional context:

- Assigning Pods to Nodes
- Taints and Tolerations

Cost Optimizer API

Cost Optimizer API

- API Overview
- API Authentication
- API Key
- Base URI Format
- HTTP Status Codes
- CSRF Token Protection
- API Permissions
- Synchronous and Asynchronous Calls
- Cost and Inventory Calls 5.5.0
- Recommendation Calls 5.5.0
- Cost Groups Calls 5.5.0
- Tags Collector Calls 5.5.0
- Cloud Setup Calls 5.5.0
- Remedial Actions on Recommendations for Cost Optimizer 5.5.0

API Overview

CloudCenter Suite API Overview

- Overview
- CloudCenter Suite API Version
- Date Format
- HTTPS Request Methods
- Response Schema
- Resource URL and ID
- Pagination
    - Pagination Request Attributes
    - Pagination Response Attributes
- Sorting
- Searching
- HTTP Location URL
- Who Can Use CloudCenter SuiteAPIs?

Overview

The payloads for the CloudCenter Suite APIs are visible in the API documentation section for each module.

CloudCenter Suite API Version

CloudCenter Suite APIs provide support for theCloudCenter Suite modules: Suite Admin API, Workload Manager API, Action Orchestrator API, and Cost Optimizer API.

The User, Groups, and Tenant APIs are part of the Suite Admin and each API using these services have an additional prefix in the URI. The payloads for the CloudCenter Suite APIs are visible in the API documentation section for each module.

The v2 APIs, where available, provide structured responses with minimum details and provides links for nested resources as well as improved search, sort, and pagination filters.

Date Format

The CloudCenter Suite API date and time values are formatted inUnix timeto the millisecond level.The APIs are agnostic to dates and time zones.

HTTPS Request Methods

CloudCenter Suite APIs support the following request methods:

- **GET**: To query or view the server information based on a CloudCenter Suite deployment
- **PUT**: To replace the entire object for update operations
- **POST**: To perform a CloudCenter Suitetask or creating the resource
- **DELETE**: To remove specific aspects of the CloudCenter Suite deployment

Response Schema

CloudCenter APIs issue responses for all APIs using both JSON and XML formats. You can set the response format by sending the appropriate Content-Type request headers:

- JSON (Default)

```
Content-Type: application/json Accept: application/json
```

- XML

```
Content-Type: application/xml Accept: application/xml
```

- CSV (Only for Reports)

The CSV format only applies to report-based APIs

```
Content-Type: application/csv Accept: text/csv
```

Resource URL and ID

For each API request, you see two common attributes displayed in the API response:

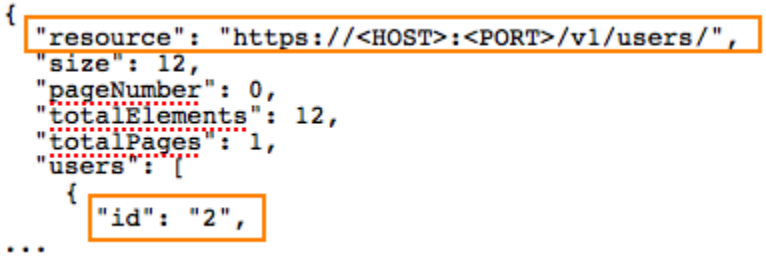

- The**resource**URL: A unique URL that provides access to the requested*CloudCenter Suite Resource*.
- The POST and PUT API calls additionally provide an**id**attribute for each new*CloudCenter Suite Resource*.

Pagination

The pagination information differs based on the API version:

- **v1 APIs**: The GET (view or list) APIs support pagination by default. CloudCenter Suite APIs use the following attributes to provide paginated results:

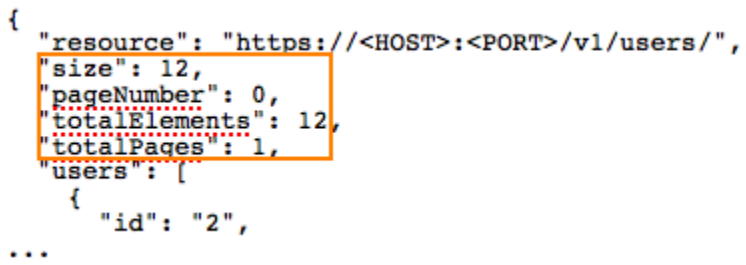

- **v2 APIs**: Requires the*page*and*size*attributes for any request. The default size for v2 APIs now list 50 records by default.

Pagination Request Attributes

page

- **Description:**The total number of pages in for the API listing.
  - Default = 0
  - If**size=0**, then the*page*value is ignored.
  - If not specified (**page=0&size=20**), the defaultsize(default = 20) value displays the first 20 elements, which is equal to one page
  - If you specify both the page and the size values, the following applies:

If you specify...	...then
size=21	Elements numbered 21 - 40 entities are displayed, which is equal to 2 pages
page=0 (or not specified)	The first set of 20 elements in the list, elements 1 to 20 are displayed
page=1	The second set of 20 elements in the list, elements 21 to 40 are displayed
page=2	The third set of 20 elements in the list (the third page). if the page does not have more than 10 elements, then only those 10 elements are displayed.
page=1&&size=10	A set of 10 elements, Elements 11 to 20 are displayed
page=1&&size=20	A set of 20 elements, Elements 21 to 40 are displayed
page=2&&size=10	A set of 10 elements, Elements 21 to 30 are displayed

- **Type**: Integer

size

- **Description**: Total number of records that any list page should contain.The default is:
  - v1 APIs = 20records
  - v2 APIs = 50 records
- **Type**: Integer

Pagination Response Attributes

- v1 APIs:
  pageResource
  - **Description**: Identifies the pagination information for each resource
  - **Type**: Sequence of attributes for v1 APIs

size (see above)
pageNumber • **Description**: The page number that the client wants to fetch. Page numbers start with 0 (default). • **Type**: Integer
totalElements • **Description**: The number resources that an API call returns • **Type**: Long
totalPages • **Description**: The number of pages in a response • **Type**: Integer

- v2 APIs:
  pageResource
  - **Description**: Identifies the pagination information for each resource
  - **Type**: Sequence of attributes for v2 APIs

resource • **Description**: Unique URL to access this resource. • **Type**:String
size (see above)

pageNumber
- **Description**: The page number that the client wants to fetch. Page numbers start with 0 (default).
- **Type**: Integer

totalPages
- **Description**: The number of pages in a response
- **Type**: Integer

jobs
- **Description**: Array of JSON objects that use**jobs**as the key.
- **Type**: Array of JSON objects

previousPage
- **Description**: A resource link to the previous page.
- **Type**: URI as a string

nextPage
- **Description**: A resource link to the following page.
- **Type**: URI as a string

lastPage
- **Description**: A resource link to the last page.
- **Type**: URI as a string

Sorting

- **v1 APIs**: All list APIs support sorting by default and use the query-string parameters to provide sorted results with a comma-separated set of property names.
  - Sorting Order:
    - Ascending order: Default when you specify the property.
    - Descending order: Append a dash to the property.
  - Example:
    - **sort=id,name**:Sort by ID property in ascending order and then sort by name property in ascending order.
    - **sort=id,name-,description**: Sort by ID property in ascending order, then sort by name property in descending order, and finally sort by description in ascending order.
  - Property name validation: Property names in sort parameters are validated. For example, APIs that return a list of users can sort only on properties exposed by the user object as sortable.
  - The following example displays the use of sorting and pagination attributes in the same API request.

```
curl -k -X GET -H "Accept: application/json" -u cliqradmin:D3DD6F7874E6B26B "https://test.cliqr.com/v1/users?
detail=true&page=0&size=30&sort=firstName

> GET /v1/users?detail=true&page=0&size=30&sort=firstName HTTP/1.1
> Authorization: Basic Y2xpcXJhZG1pbipEM0RENkY3ODc0RTZCMiZC
> User-Agent: curl/7.37.1
> Host: test.cliqr.com
> Accept: application/json
>
< HTTP/1.1 200 OK
```

- **v2 APIs**: Requires the*sort*attributes for any request.

  sort
  - **Description**: Sorts API responses based on the format specified.
  - **Type**: String
    - Sorting order:
      - Ascending order = ASC
      - Descending order = DESC
    - Default: Sort criteria is based on*startTime*and DESC order.
    - Format: sort=[*attribute*,*order*]
    - Example: [endTime,ASC]
    - Sorting attributes:

      id
      - **Description**:Unique,system-generatedidentifier for this *resource*.
      - **Type**: String

status

- **Description**:Status of the operation. Seethe*APIs for the relevant module* to view a list of all job operations.
- **Type**: Enumeration

Enumeration	Description
SUBMITTED	The operation has been submitted
RUNNING	The operation is currently in progress
SUCCESS	The operation succeeded
FAIL	The operation failed

startTime/endTime

- **Description**: Start/Endtime for this resource.Unix epoch time in milliseconds.
- **Type**:
  - v1 APIs = Long
  - v2 APIs = Epoch time as a String

totalCost

- **Description**: Identifies the total cost per hour of the job for billing purposes. See theCost Optimizer APIsection to view additional details.
- **Type**: Float

nodeHours

- **Description**: The number of VM hoursfor this resource. See theCost Optimizer APIsection to view additional details.
- **Type**: Float

name

- **Description**:The name assigned for this*CloudCenter Suite Resource*.Valid characters are letters, numbers, underscores, and spaces.
- **Type**: String

deploymentEntity.name

- **Description**: Identifies evolving resource details about the deployment. The deploymentEntity attribute uses the*deploymentEntity.name*format, where**.name**isa searchvalue for deploymentEntity and deploymentEntity itself is a JSON object.

  Instead of placing the deployment name at the top level search and adding numerous query parameters, this format allows for nested search results. The top level **name** is the job name and deploymentEntity.**name** is the deployment name.

- **Type**: JSON objects

favoriteCreationTime

- **Description**: If the job was configured as a favorite job, then this attribute identifies the time when this configuration took place. Seethe*Favorite Deployments*section for the relevant release for additional context.
- **Type**: Epoch time as a String

Searching

This attribute is only available for v2 APIs.

search

- **Description**: Searches API responses based on the format specified.
- **Type**: String
  - Format:search=[field, searchType,*SearchExpression1*,*SearchExpression2*]
  - Example: search =[startTime, gt, 01/01/2016]
  - Search Expressions:
    - *pattern*: Provide a pattern using the format provided in the*searchTypes*table below.
    - searchTypes

searchType	Format
eq	==
ne	!=
el	LIKE *pattern*%
fl	LIKE %*pattern*

eln	NOT LIKE *pattern*%
fln	NOT LIKE %*pattern*
fle	LIKE %*pattern*%"
gt	> *searchValue*
lt	< *searchValue*
ge	>= *searchValue*
le	<= *searchValue*
gtlt	>*searchValue* &&*searchValue*
gtelt	>=*searchValue* && <*searchValue*
gtlte	>*searchValue* && <=*searchValue*
gtelte	>=*searchValue* && <=*searchValue*
emp	Empty string
noemp	Not Empty string
nu	Null value
nn	Not Null Value

- searchValue:

searchValue	SearchType Availability
id	eq
startTime	eq, nu, gtlt
endTime	eq, nu, nn, gtlt
totalCost	eq, gt, ge, le, gtlt, gtlte, gtelte, gtelt
favoriteCreationTime	eq, nu, ,nn gtlt
jobStatusMessage	el, eln, fl, fln, fle, nn, emp, noemp
nodeHours	eq, gt, ge, le, gtlt, gtlte, gtelte, gtelt
name	eq, nn, eln, fle, fln, el, emp, noemp, fl
description	eq, nn, eln, fle, fln, el, emp, noemp, fl
deploymentEntity.name	eq, nn, eln, fle, fln, el, emp, noemp, fl
ownerEmailAddress	eq
cloudFamily	eq, nu
status	eq, nu

HTTP Location URL

The HTTP Status code and the Location URL (highlighted in blue in the following example) is provided in the Response Header when Create*resource*API calls are successful:

```
curl -k -X POST -H "Content-Type: application/json" -H "Accept: application/json"
cliqradmin:D3DD6F7874E6B26B https://test.cliqr.com/v1/users -d '{
    "firstName": "User 02",
    "lastName": "Cliqr",
    "password": "cliqr",
    "emailAddr": "user.02@cliqr.com",
    "companyName": "Cliqr, Inc",
    "phoneNumber": "14085467899",
    "externalId": "",
    "tenantId": 1
}'

> POST /v1/users HTTP/1.1
> Authorization: Basic Y2xpcXJhZGlpbjpEM0RENkY3ODc0RTZCMjZC
> User-Agent: curl/7.37.1
> Host: test.cliqr.com
> Content-Type: application/json
> Accept: application/json
> Content-Length: 217
>
< HTTP/1.1 201 Created
< Server: Apache-Coyote/1.1
< Set-Cookie: JSESSIONID=0E85227543C66D55E06449582091C2B4; Path=/; Secure; HttpOnly
< osmosix_content: true
< X-Frame-Options: SAMEORIGIN
< Pragma: no-cache
< Expires: Thu, 01 Jan 1970 00:00:00 GMT
< Cache-Control: no-cache
< Cache-Control: no-store
< Location: https://test.cliqr.com/v1/users/12
< Content-Type: application/json;charset=UTF-8
< Transfer-Encoding: chunked
< Vary: Accept-Encoding
< Date: Fri, 07 Aug 2015 20:59:18 GMT
```

Who Can Use CloudCenter SuiteAPIs?

Both admins and users can use CloudCenter Suite REST APIs.

Your login credentials determine if you are an admin (platform (root), tenant admin, or co-admin) or a user. If you do not have the requiredPermission Controllevel to access any*resource*,you receive the HTTP 403 status error mentioned in theHTTP Status Codessection.

Back to:

- Suite Admin API
- Workload Manager API
- Action Orchestrator API
- Cost Optimizer API

API Authentication

API Authentication

- Overview
- Authentication Format in CURL Requests
- Successful Authentication
- Session Timeout Length

Overview

CloudCenter Suite APIs require the following authentication details for each API call:

- Username
- APIaccess key

> ⚠ The authentication HTTP header is not required when making standalone REST API calls using the username/API Key credentials.

Authentication Format in CURL Requests

Standalone CURL Request Example:

```
curl -H "Accept:application/json" -H "Content-Type:application/json" -u writer:BED74F4D9BFE0DA0 -X GET
https://<HOST>:<PORT>/v1/users/27
```

In this CURL request example:

- **writer1**is theusername
- **BED74F4D9BFE0DA0**is the APIaccessKey

Yourtenant administratorcan retrieve the username and API access key from the UI. SeeAPI Key for additional details.

Successful Authentication

On successful authentication, CloudCenter Suite sends a browser cookie to maintain the authentication session. The cookie forwards the information to the server for each API call so you do not need to authenticate each time you make an API call. If you do not want to maintain cookies in your browser, you can send the authentication information for each API request. Once authenticated, you can begin making API calls.

Session Timeout Length

The CloudCenter Suite authentication session times out after 15 minutes. If you use a REST client to make API calls by authenticating through the UI's, this session timeout applies to the REST client as well.

However, if you add and save the REST client authentication headers or if you issue CURL commands with the authentication details, you can circumvent the session timeout restriction.

Back to:

- Suite Admin API
- Workload Manager API
- Action Orchestrator API
- Cost Optimizer API

API Key

Generate API Key

- Overview
- UI Process to Generate Your Own API Key
- UI Process to Generate API Key for Another User
- API Process to Generate a New API Key

Overview

You need an**API key**to use CloudCenter Suite APIs. Suite administrators or tenant administrator (for their respective tenants) can generate/regenerate an API keyby using the Suite Admin UI or the**user_api_key**API call.

UI Process to Generate Your Own API Key

To generate the API key from the UI for yourself, follow this procedure:

1. Navigate to the Suite Admin Dashboard and click youraccount profile dropdown.
2. Click the **Generate API Key**link to generate a new API key.
3. Click **Yes** to replace the API key. You can now use this key to make REST API calls as listed in the Usage Example in the following screenshot.

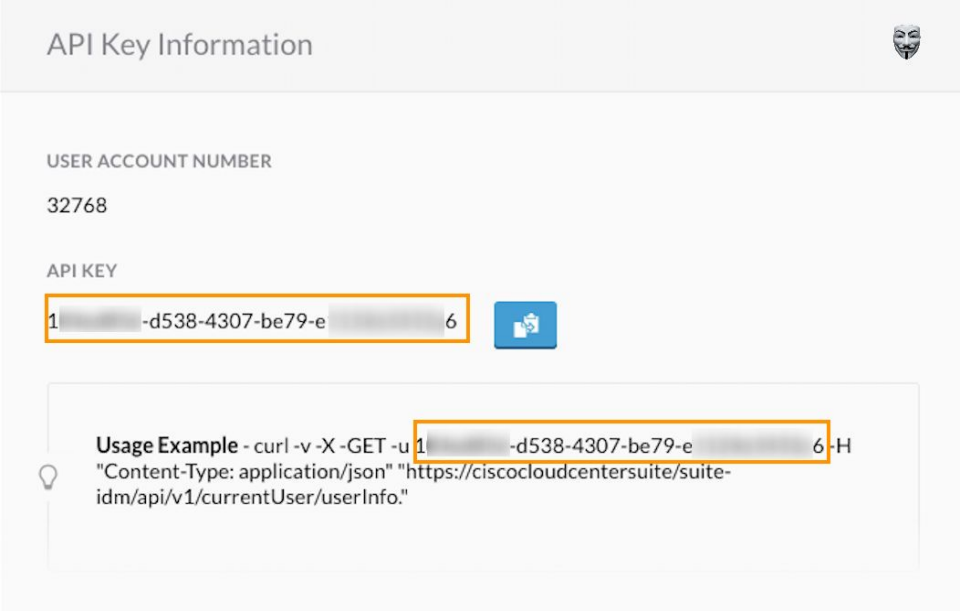

UI Process to Generate API Key for Another User

To generate the API key from the UI for another user, follow this procedure:

1. Navigate to theSuite Admin Dashboard> **Users**.
2. Search for the required user and select**Generate API Key**from the Actions dropdown for this user as displayed in the following screenshot.

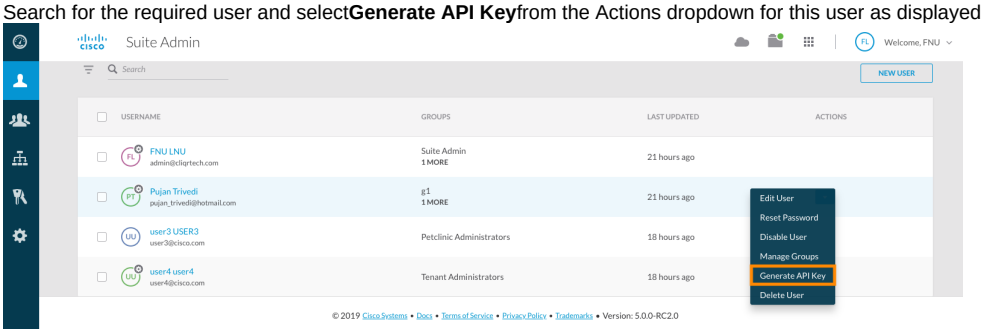

3. Click the**Generate API Key**link to generate a new API key.This user can now make REST API calls using the new API key.

API Process to Generate a New API Key

To generate the API key using the Suite Admin API call, follow this procedure:

1. Issue thePassword Service API Calls >**/api/v1/users/{userId}/user_api_key**API POSTcall to generate/regenerate the API key for yourself or for any other user.

   ```
   POST https://host-port/suite-password/api/v1/users/1/user_api_key
   ```

2. Retrieve the*apiKey*from the response for this API.

   ```
   {
      "userId":1,
      "apiKey":"1.......-d538-4307-be79-e.........6",
      "accountNumber":"32768"
   }
   ```

3. Use this*apiKey*to make REST API calls.

Back to:

- Suite Admin API
- Workload Manager API
- Action Orchestrator API
- Cost Optimizer API

Base URI Format

Base URI Format

- Overview
- Host Name
- Port Usage
- API Version
- Parameters
- Parameter Types

Overview

The base URI format is **https:// <host>:<port>/**

Host Name

The host is generally represented as <HOST> in all CloudCenter APIs. It represents the IP address or the DNS name.

The host differs based on your DNS or IP address and port usage.

Port Usage

The port is generally represented at <PORT> in all CloudCenter APIs. It represents the port used to connect to theCCO server for the API connection. The <PORT> in the REST endpoint is***optional***. You can decide if you want to use the port for each API call. All CloudCenter API requests and responses display <PORT> in all examples.

```
curl -H "Accept:application/json" -H "Content-Type:application/json" -u \
cloudcenteradmin:40E45DBE57E35ECB -X GET https://<HOST>:<PORT>/...
```

If you do not specify the port, then API requests default to Port 443 for an HTTPS connection when accessing CloudCenter Suite REST APIs.

API Version

The CloudCenter Suite 5.0.0 API version can be v1 or v2 as applicable. The version is identified for each API, where applicable.

Parameters

Parameters used to make the API call are displayed after the APIs and are called out after the description.

Parameter Types

Attribute Type	Description
String	Any combination of characters. Maximum of 255 characters.
Integer	A whole number value. Restricted to 32-bit values.
Long	A whole number value. Restricted to 64-bit values.
Float	A number with or without a decimal point. Displayed as a string in the response.
Boolean	A logical true or false value. May be passed to API requests as true or false or 1 or 0.
Enumeration	A predefined list of values, for example, STANDARD or TENANT describes the possible values for each type. Only listed values are permitted, other values result in an error.
JSON Object	A method to parse JavaScript Object Notation (JSON) and return the object value to which a specified name is mapped.

Name-Value Pair	A name-value pair whereeach element is an attribute-value pair.
Array	A sequential collection of like elements corresponding to the element's data type. The type of the array is determined by the types of the elements (can be String, Integer, Name-Value Pair Type)
Perms List	Lists the permissions for a specific user if the user is logged in. An empty response is also indicative of the resource not being currently supported.
Metadata	Metadata information associated with the cloud provider.

Back to:

- Suite Admin API
- Workload Manager API
- Action Orchestrator API
- Cost Optimizer API

HTTP Status Codes

Unable to render {include} The included page could not be found.

CSRF Token Protection

CSRF Token Protection

- Overview
- The 403 Forbidden Error for Some APIs
- Setting the CSRF Token
- Retrieving the CSRF Token
- Using the CSRF Token

Overview

Cisco provides CSRF protection for all API calls. When an API call is made by you or the CloudCenter Suite, be aware that a CSRF token is required for the following scenarios:

- If the request method is **POST**, **PUT**, or **DELETE**
  and
- If the request **Content-Type** is not **application/json**

For example, the following functions require the CSRF token:

- Suite Admin Resource Management Service API Calls that use the following functions:
  - Company logo upload
  - User avatar upload
- Workload Manager API Calls that use the following functions
  - Application profiles
  - Logo upload
  - Services logo upload
  - Import applications
  - Cloud account management API calls
  - DELETE calls that change the database contents

The 403 Forbidden Error for Some APIs

If the CSRF token is missing or incorrect, you will see a 403 error due to the CSRF token protection.

If you see this error, you must first set the CSRF token in the request header for the affected API.

Setting the CSRF Token

To set a CSRF token, add **X-CSRF-TOKEN** to the header name (case sensitive, all uppercase).

Retrieving the CSRF Token

To obtain the CSRF token, follow this procedure.

1. You must first pass authentication. SeeAPI Authentication for details.
2. Once authenticated, use one of the following APIs to retrieve the CSRF token from the response body (**csrfToken**attribute). See Authentication Service API Calls for details.
   a. Login API (/suite-auth/login)
   b. Token Refresh API (/suite-auth/api/v1/token)
   c. CSRF Token API (/suite-auth/api/v1/csrfToken)

Using the CSRF Token

See the following request for examples of using a CSRF Token.

Java Rest Client Example

```
WebResource.Builder builder = webResource.type(MediaType.APPLICATION_JSON).header("X-CSRF-TOKEN", "<TOKEN>");
```

Python Example

```
headers = {'content-type': 'application/json', 'X-CSRF-TOKEN': '<TOKEN>'}

requests.delete(url, headers = headers, verify=False)

requests.post(url, json=jobJson, headers = headers, verify=False)
```

Where **<TOKEN>** is retrieved as specified in the *Retrieving the CSRF Token* section above.

Back to:

- Suite Admin API
- Workload Manager API
- Action Orchestrator API
- Cost Optimizer API

API Permissions

API Permissions Allowed Roles

- Overview
- Current User Permissions
- Suite Level Permissions
- Workload Manager Roles
- Action Orchestrator Roles
- Cost Optimizer Roles

Overview

Each API identifies the permissions and roles required to execute that API call. Permissions for each API are governed by Role-Based Access Control (RBAC) as explained inUnderstand Roles and user-level as explained inUnderstand User Levels.

Current User Permissions

Users can find their permission level by executing the**GET /suite-idm/api/v1/currentUser/userInfo**API listed in theIDM Service API Calls > *User Controller* section.

Suite Level Permissions

Based on the current user's permissions the Suite Admin APIs display enumerations for the **Allowed Role(s)** described in the following table.

Allowed Role(s) Enumeration	Description
SUITE_ADMIN	The initial administrator described in Initial Administrator Setup. This user can perform the following tasks: • Module Lifecycle Management • Manage Clusters
SUITE_TENANT_ADMIN	The tenant administrator set up as part of the root tenant configuration described in Manage Tenants. This user can perform the following tasks: • Manage sub-tenants • Create, update, and delete sub-tenant users (including createTenantWithAdmin atomic operation) • Tenant resource management including Email Settings, Branding Information, and so forth
SUITE_USER	Any user added to the CloudCenter Suite. A newly-added user can only view the Suite Admin Dashboard, if not assigned to a group.
SUITE_USER_ADMIN	ASUITE_ADMINcan promote any SUITE_USER to the SuiteAdministratorgroup as described in Create and Assign Groups. This user can perform the following tasks: • Manage users and groups • Create, update, delete users and groups • Assign roles to users and groups • Manage passwords for users
SUITE_OUTOFBOX_USER	ASUITE_ADMINcan promote any SUITE_USER to be a SUITE_OUTOFBOX_USER, which basically implies that this user has been added to one or more OOB Suite Admin Groups.
SUITE_RESET_PASSWORD	Users withSUITE_ADMINpermissions and/or SUITE_TENANT_ADMIN for this tenant as described in Create and Manage Users > *User Actions.* This user can perform the following tasks: • Edit any user's profile by changing the first/middle/last name and email • Configure metadata details • Configure groups • Reset password • Disable a user

Workload Manager Roles

SeeOOB Groups, Roles, and Permissions for details.

Action Orchestrator Roles

SeeAction Orchestrator Roles for details.

Cost Optimizer Roles

SeeAccess and Roles for details.

Back to:

- Suite Admin API
- Workload Manager API
- Action Orchestrator API
- Cost Optimizer API

Synchronous and Asynchronous Calls

Synchronous and Asynchronous Calls

- Overview
- Synchronous
- Asynchronous
  - Call States
  - Operation ID Availability

Overview

CloudCenter Suite APIs support both synchronous and asynchronous calls. Some APIs return data in the response body and others will only return an HTTP status. For example, CloudCenter DELETE calls return a**Status204 No Content**after deleting the*resource*in the background.

Synchronous

Synchronous APIs indicate that the program execution waits for a response to be returned by the API. The execution does not proceed until the call is completed. The real state of the API request is available in the response.

Asynchronous

Asynchronous APIs do not wait for the API call to complete. Program execution continues, and until the call completes, you can issue GET requests to review the state after the submission, during the execution, and after the call completion. Use the**Get Operation Status**API to retrieve the status of an asynchronous operation.

As asynchronous calls may take some time to complete, they returnHTTP Status Codesresponses containing information with anHTTP Status Code, which allows you to retrieve the progress, status, response, and other information for the call.

After submitting an asynchronous API call:

1. Retrieve the resource URL from theHTTP Status Codes.
2. Use this location URL and query the system using GET calls. While the call is in progress and you issue the GET request, you get additional details of the operation being performed. These details are only available while the operation is in various states of execution (RUNNING, SUCCESS, FAILED).
3. When the asynchronous API call completes successfully, issue a GET request to view the SUCCESS state and theresourceURL for this operation.

Call States

In the following example of aCreate Cloud AccountAPI:

- The various states of execution (RUNNING, SUCCESS, FAILED) are highlighted in corresponding colors
- The first and last GET requests are in bold to show the sequence of events

```
Location: https://test.cliqr.com/v1/operationStatus/f503c52a-d13b-4b62-840d-0faa22ccbb78

{ "operationId": "f503c52a-d13b-4b62-840d-0faa22ccbb78", "status": "RUNNING", "msg": "Updating
Image permissions...", "progress": 50, "timestamp": 1438850245522, "additionalParameters": null,
"operationHistory": [ ], "subtaskResults": null, "resourceUrl": "https://test.cliqr.com/v1/
operationStatus/f503c52a-d13b-4b62-840d-0faa22ccbb78" }
curl 'https://test.cliqr.com/v1/operationStatus/f503c52a-d13b-4b62-840d-0faa22ccbb78' -H 'Accept:
application/json'

{ "status": "RUNNING", "msg": "Updating Image permissions...", "resource": "https://
test.cliqr.com", "additionalParameters": [] }
…
curl 'https://test.cliqr.com/v1/operationStatus/f503c52a-d13b-4b62-840d-0faa22ccbb78' -H 'Accept:
application/json'

{ "status": "RUNNING", "msg": "Saving cloud account...", "resource": "https://test.cliqr.com/
https://test.cliqr.com/v1/operationStatus/f503c52a-d13b-4b62-840d-0faa22ccbb78",
"additionalParameters": [ ] }
curl 'https://test.cliqr.com/v1/operationStatus/f503c52a-d13b-4b62-840d-0faa22ccbb78' -H 'Accept:
application/json'

{ "status": "SUCCESS", "msg": "Cloud Account is saved successfully.", "resource": "https://
test.cliqr.com/https://test.cliqr.com/v1/operationStatus/f503c52a-d13b-4b62-840d-0faa22ccbb78",
"additionalParameters": [ ] }
```

Operation ID Availability

Operation IDs (displayed below the Location URL in the above image) allow you to query the status of asynchronous APIs and are only available for a brief period as identified in the following table:

Operation ID Availability	Description
5 minutes	The Operation ID is available for five minutes if the operation completes (regardless of success or failure).
1 hour	The Operation ID is available for one hour if the operation times out and does not complete.

Back to:

- Suite Admin API
- Workload Manager API
- Action Orchestrator API
- Cost Optimizer API

Cost and Inventory Calls 5.5.0

Cost and Inventory Calls for Cost Optimizer 5.5.0

Download the Swagger file fromhttps://www.cisco.com/content/dam/en/us/td/docs/cloud-systems-management/cloudcenter-suite/Cost-Optimizer/55_CO_Swagger.zip.

Back to: **Cost Optimizer API**

Recommendation Calls 5.5.0

RecommendationCalls for Cost Optimizer 5.5.0

Download the Swagger file fromhttps://www.cisco.com/content/dam/en/us/td/docs/cloud-systems-management/cloudcenter-suite/Cost-Optimizer/55_CO_Swagger.zip.

Back to: **Cost Optimizer API**

Cost Groups Calls 5.5.0

Cost GroupsCalls for Cost Optimizer 5.5.0

Download the Swagger file fromhttps://www.cisco.com/content/dam/en/us/td/docs/cloud-systems-management/cloudcenter-suite/Cost-Optimizer/55_CO_Swagger.zip.

Back to: **Cost Optimizer API**

Tags Collector Calls 5.5.0

Tags Collector Calls for Cost Optimizer 5.5.0

Download the Swagger file fromhttps://www.cisco.com/content/dam/en/us/td/docs/cloud-systems-management/cloudcenter-suite/Cost-Optimizer/55_CO_Swagger.zip.

Back to: **Cost Optimizer API**

Cloud Setup Calls 5.5.0

Cloud SetupCalls for Cost Optimizer 5.5.0

Download the Swagger file fromhttps://www.cisco.com/content/dam/en/us/td/docs/cloud-systems-management/cloudcenter-suite/Cost-Optimizer/55_CO_Swagger.zip.

Back to:Cost Optimizer API

Remedial Actions on Recommendations for Cost Optimizer 5.5.0

Remedial Actions on Recommendationsfor Cost Optimizer 5.5.0

Download the Swagger file fromhttps://www.cisco.com/content/dam/en/us/td/docs/cloud-systems-management/cloudcenter-suite/Cost-Optimizer/55_CO_Swagger.zip.

Back to:Cost Optimizer API