



CloudCenter Workload Manager 5.4 Documentation

First Published: July 20, 2020

Last Modified: December 9, 2020

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

1. Workload Manager 5.4 Home	4
1.1 Release Notes	5
1.1.1 Workload Manager 5.4.3	6
1.1.2 Workload Manager 5.4.2	8
1.1.3 Workload Manager 5.4.1	10
1.1.4 Workload Manager 5.4.0	12
1.2 What Is Supported?	18
1.2.1 Public Clouds	19
1.2.2 Datacenters and Private Clouds	23
1.2.3 Container Clouds	24
1.2.4 Conditional Component Appliance Images	25
1.2.5 OOB Logical Images	27
1.2.6 OOB Services	28
1.2.6.1 Supported OOB Services	29
1.2.6.2 Apache Service	31
1.2.6.3 Chef Service	32
1.2.6.4 Docker Service	33
1.2.6.5 MySQL Service	36
1.2.6.6 Nginx Service	37
1.2.6.7 Puppet Service	38
1.2.6.8 Tomcat Service	39
1.2.6.9 Varnish Service	40
1.2.7 Dynamic Bootstrapping	41
1.2.8 OOB Application Templates	43
1.2.9 OOB Groups, Roles, and Permissions	44
1.2.10 Permission Control	45
1.2.11 Understand ACLs	50
1.3 Getting Started	55
1.3.1 Workload Manager Overview	56
1.3.2 Architecture	57
1.3.3 Dashboard and Menus	62
1.3.4 Next Steps for Administrators	64
1.3.5 Next Steps for Standard Users	65
1.4 Install Conditional Components	66
1.4.1 Installation Overview	67
1.4.2 Resource Requirements	68
1.4.3 Cloud Remote (Conditional)	69
1.4.4 Local Repo Appliance (Conditional)	111
1.4.5 Worker (Conditional)	115
1.4.5.1 Management Agent (Worker)	116
1.4.5.2 Options to Install the Worker	118
1.4.5.3 Install Worker on a Linux Image	119
1.4.5.4 Install Worker on a Windows Image	121
1.4.6 Local Package Store (Conditional)	126
1.4.7 Local Bundle Store (Conditional)	128
1.5 Clouds	129
1.5.1 Cloud Overview	130
1.5.2 Configure a Cloud End-to-End	135
1.5.2.1 Configure a vCenter Cloud	136
1.5.2.2 Configure an AWS Cloud	154
1.5.2.3 Configure an AzureRM Cloud	168
1.5.2.4 Configure a Google Cloud	189
1.5.2.5 Configure an OpenStack Cloud	206
1.5.2.6 Configure a Kubernetes Cloud	226
1.5.2.7 Configure a vCD Cloud	246
1.5.2.8 Configure an IBM Cloud	260
1.5.2.9 Configure an Outscale Cloud	273
1.5.3 Cloud Maintenance	285
1.5.4 Manage Instance Types	288
1.6 Images	292
1.6.1 Images Overview	293
1.6.2 Images Page	294
1.6.3 Regions Tab Image Settings Section	302
1.7 Other Admin Functions	304
1.7.1 All Reports	305
1.7.1.1 Reports Overview	306
1.7.1.1.1 Advanced Filters	312
1.7.1.1.2 Time Period Filter	316
1.7.1.2 Usage Summary Report	317
1.7.1.3 Application Deployments Report	318
1.7.1.4 Running VM History Report	319
1.7.2 System Tags	321
1.7.3 Usage Plans and Fees	323
1.7.3.1 Financial Overview	324
1.7.3.2 Plan Configuration	326
1.7.3.3 Bundle Configuration	328
1.7.3.4 Workload Manager Cost and Fees	329
1.7.4 Tenant Management	330
1.8 Services	333

1.8.1 Service Administration	334
1.8.2 Custom Service Definition	336
1.8.3 External Service	343
1.8.4 Container Service	351
1.8.5 Service Lifecycle Actions	357
1.8.6 Service Versions and Usage	361
1.8.7 Service Export/Import	365
1.9 Deployment Environments	367
1.9.1 Deployment Environments Overview	368
1.9.2 Environments Page	369
1.9.3 Create a Deployment Environment	370
1.9.4 Define Resource Placement	379
1.9.5 Define Resource Validation	386
1.9.6 Cloud-Specific Configurations	388
1.9.6.1 AWS Configurations	389
1.9.6.2 AzureRM Configurations	393
1.9.6.3 Google Cloud Configurations	397
1.9.6.4 OpenStack Configurations	398
1.9.6.5 vCenter Configurations	402
1.9.7 Availability Sets and Zones	404
1.9.8 Multi-Site, Multi-Account Deployment Considerations	406
1.10 Application Profiles	411
1.10.1 Application Profile Overview	412
1.10.2 Application Profiles Page	413
1.10.3 Application Tasks	415
1.10.4 Model Applications	419
1.10.4.1 Application Profile Creation Considerations	420
1.10.4.2 Create an Application Profile	421
1.10.4.3 New Application Profile	423
1.10.4.4 Topology Modeler	438
1.10.4.5 Container Placement Groups	443
1.10.4.6 Understand Application Tier Properties	445
1.10.4.7 Deployment Parameters	451
1.10.4.8 Deployment Lifecycle Scripts	455
1.10.4.9 Multiple Volumes	468
1.10.4.10 Security and Firewall Rules	473
1.10.4.11 IP Address Allocation	476
1.10.4.12 Parameters and Macros	479
1.10.4.12.1 Using Parameters	480
1.10.4.12.2 Pre-Defined Parameters	483
1.10.4.12.3 Parameter Substitution	492
1.10.4.12.4 Configuration Files	495
1.10.4.13 Application Using Imported Profile	496
1.10.4.14 Application Using App Package	497
1.10.4.15 Application Using Puppet or Chef	498
1.10.5 Deploy Applications	500
1.10.5.1 Deploy an Application	501
1.10.5.1.1 Simplified Network Settings	511
1.10.5.2 JSON for the Deploy API Call	518
1.10.6 Benchmark an Application	521
1.11 Policy Management	522
1.12 Artifact Repository	532
1.13 Actions Library	536
1.13.1 Actions Library Overview	537
1.13.2 Actions Library Page	540
1.13.3 Create a Custom Action	542
1.13.4 Manage Actions	552
1.13.5 Create Snapshot	553
1.14 Manage Deployments and VMs	556
1.14.1 Deployment, VM, and Container States	557
1.14.2 Virtual Machine Management	565
1.14.3 Deployments List Page	592
1.14.4 Deployment Details Page	593
1.14.5 Terminate Protection	607
1.14.6 Track Cloud Costs	611
1.14.7 Project and Phase Management	613
1.14.8 Migrating Applications	624
1.14.9 Container Tier Rolling Updates	628
1.15 Advanced Configuration Topics	629
1.15.1 Guidance for Callout Scripts	630
1.15.2 Advanced Configuration using Kubernetes ConfigMaps	636
1.15.3 Limiting Acceptance of Self-Signed Certificates for HTTPS Web Services	638
1.15.4 Update Self-Signed Certificates for Guacamole	639
1.15.5 Adjusting Action Message Purge Frequency	641
1.15.6 Password-Based Authentication	642
1.15.7 Custom Docker Image for Scripts	643
1.16 Integrations	645
1.16.1 Workload Manager - Action Orchestrator Integration	646
1.16.2 Extensions	650

1.16.2.1 ACI Extensions	651
1.16.2.2 ServiceNow Extensions	660
1.16.2.3 ACI Multi-Site Extensions	681
1.16.3 Jenkins Integration	688
1.16.4 Service Libraries	693
1.17 Troubleshoot	695
1.17.1 Obtaining Your SSH Private Key	696
1.17.2 Bundle Download Failure	697
1.17.3 Kubernetes Troubleshooting	698
1.17.4 Agent Troubleshooting	700
1.17.5 Troubleshooting Parameters	702
1.18 Workload Manager API	705
1.18.1 API Overview	706
1.18.2 API Authentication	712
1.18.3 API Key	713
1.18.4 Base URI Format	715
1.18.5 HTTP Status Codes	717
1.18.6 CSRF Token Protection	718
1.18.7 API Permissions	720
1.18.8 Synchronous and Asynchronous Calls	722
1.18.9 CCM Calls 5.4.0	724
1.18.10 Cloud Setup Calls 5.4.0	725
1.18.11 Manual API Documentation	726

Workload Manager 5.4 Home

CloudCenter Workload Manager 5.4 Documentation

Cisco released the following Workload Manager releases:

- [Workload Manager 5.4.0](#) was released on July 30, 2020
- [Workload Manager 5.4.1](#) was released on October 1, 2020
- [Workload Manager 5.4.2](#) was released on November 9, 2020
- [Workload Manager 5.4.3](#) was released on December 9, 2020

Search

[Install Worker on a Linux Image](#)

updated Jan 25, 2021

[view change](#)

[Install Worker on a Windows Image](#)

updated Jan 25, 2021

[view change](#)

[Workload Manager 5.4.3](#)

updated Jan 25, 2021

[view change](#)

Release Notes

Workload Manager Release Notes

- [Workload Manager 5.4.3](#)
- [Workload Manager 5.4.2](#)
- [Workload Manager 5.4.1](#)
- [Workload Manager 5.4.0](#)

Workload Manager 5.4.3

Workload Manager 5.4.3 Release Notes

- [Release Date](#)
- [Installation](#)
- [Upgrade](#)
- [UI](#)
- [API](#)
- [Documentation](#)
- [Resolved Issues](#)

First Published: December 9, 2020

Updated:

- December 14, 2020: Updated the API section to ensure technical accuracy.
- January 25, 2020: Updated the *Documentation* section to include a list of pages that were updated.

No updates

You can upgrade to Workload Manager 5.4.3 via the Suite Admin UI. See [Update Module](#) for additional details.



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

The recommended upgrade path is to upgrade from Workload Manager 5.1.4 or 5.2.1 or 5.3 to Workload Manager 5.4.

- If you are using Workload Manager 5.1.1 and earlier releases, upgrade the installer to Workload Manager 5.1.2 before upgrading to Workload Manager 5.1.4, and then to Workload Manager 5.4.
- If you are using Workload Manager 5.1.2 or Workload Manager 5.1.3, upgrade to Workload Manager 5.1.4 and then to Workload Manager 5.3.0 along with a Suite Admin version that is earlier than 5.2.2.
- If you are in a version older than Workload Manager 5.1.4, upgrade to Workload Manager 5.1.4 before upgrading to Workload Manager 5.4 (as identified in the [Update Module](#) section).
- If you are using Workload Manager 5.3.x or earlier, you can directly upgrade to Workload Manager 5.4.x with the above documents ways.
- New Workload Manager 5.3.x installations are not available.

The following table lists the validated Workload Manager scenarios:

Scenario	From Workload Manager Version	To Workload Manager Version	Upgrade Path Versions					
			Source Cluster			Target Cluster		
			Suite Installer	Suite Admin	Workload Manager	Suite Installer	Suite Admin	Workload Manager
1	Workload Manager 5.1.4	Workload Manager 5.4.3 (this release)	Suite Admin 5.1.1	Suite Admin 5.1.2	Workload Manager 5.1.4 Workload Manager 5.2.2	Suite Admin 5.2.3	Suite Admin 5.2.3	Workload Manager 5.4.3 (this release)
2	Workload Manager 5.3.1	Workload Manager 5.4.3 (this release)	Suite Admin 5.2.0	Suite Admin 5.2.0	Workload Manager 5.3.1	Suite Admin 5.2.0	Suite Admin 5.2.3	Workload Manager 5.4.3 (this release)

The Services related UI pages use the new v2 APIs to ensure that the payload is light when viewing the APIs using Browser Tools.

As of Workload Manager 5.4.3, all Version 1 (v1) APIs are deprecated. Moving forward, only Version 2 (v2) APIs will be supported. In order to ensure backwards compatibility, v1 APIs will remain available though unsupported. Continued support for existing scripts and integrations will require migrating to v2 APIs.

The following documentation changes were implemented in Workload Manager 5.4.3:

- [Install Worker on a Linux Image](#) (updated to point to [OOB Logical Images](#) for a list of servers)
- [Install Worker on a Windows Image](#) (updated to point to [OOB Logical Images](#) for a list of servers)

The following issue was resolved/addressed in Workload Manager 5.4.3:

- **CSCvw65290:** At 2:00 AM each day, a scheduler runs in the cloud-setup pods to soft delete and obfuscate customer-sensitive data in clouds, regions, accounts, and [applications](#) along with dependent resources like instance and storage types [of all the tenants](#). If the suite-idm pod goes down during this time [or if the certs have expired and pods are unable to communicate](#), the current cloud-setup service may have cleaned up the mentioned resources for all tenants.
Resolution: Workload Manager 5.4.3 includes a fix to address this issue and even if the suite-idm service goes down [or the certs have expired](#), it no longer cleans up the resources for any tenant.

- **CSCvw50832:** The VM List page does not list any VMs and remains blank. However, the corresponding API functions as designed and displays the VM list.
Resolution: Workload Manager 5.4.3 includes a fix to address this issue and the VMs list displays as designed even via the UI.

Workload Manager 5.4.2

Workload Manager 5.4.2 Release Notes

- [Release Date](#)
- [Installation](#)
- [Upgrade](#)
- [Architecture](#)
- [Clouds](#)
- [Cloud SDK](#)
- [Services, Applications, and Deployments](#)
- [Administration and Governance](#)
- [Policy Management](#)
- [Security](#)
- [UI](#)
- [API](#)
- [Integrations](#)
- [Deprecated Functions](#)
- [Limitations](#)
- [Known Issues](#)
- [Resolved Issues](#)

First Published: November 9, 2020

No updates

You can upgrade to Workload Manager 5.4.2 via the Suite Admin UI. See [Update Module](#) for additional details.



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

The recommended upgrade path is to upgrade from Workload Manager 5.1.4 or 5.2.1 or 5.3 to Workload Manager 5.4. If you are in a version older than Workload Manager 5.1.4, first upgrade to Workload Manager 5.1.4 before upgrading to Workload Manager 5.4 (as identified in the [Update Module](#) section).

- If you are using Workload Manager 5.1.1 and earlier releases, be sure to upgrade to Workload Manager 5.1.2 before upgrading to Workload Manager 5.1.4, and then to Workload Manager 5.4.
- If you are using Workload Manager 5.1.2 or Workload Manager 5.1.3, you must first upgrade to Workload Manager 5.1.4 and then to Workload Manager 5.4.

No updates

No updates

Clouds that were not supported by the CloudCenter Suite can be made available as the Workload Manager architecture supports plug and play as listed in [Cloud SDK Home](#).

No updates

No updates

No updates

No updates

No updates

No updates

No updates

Workload Manager will also be deprecating support for the following clouds in an upcoming release: VMware vCloud Director 8.1 and VMware vCloud Director 9.1. See [Datacenters and Private Clouds](#) for a list of supported clouds.

No updates

No updates

The following issue was resolved/addressed in Workload Manager 5.4.2:

- **CSCvu57722:** A user requested an enhancement to improve the node cost handler performance as the handler performance degrades over time with increases in Resource and ResourceCost collection sizes.
Resolution: Workload Manager 5.4.2 includes a fix to ensure that handler performance keeps up with increases in Resource and ResourceCost collection sizes.
- **CSCvw28541:** A user encountered a failure when deploying in a Kubernetes environment.
Resolution: Workload Manager 5.4.2 includes a fix to ensure that Kubernetes-based deployments complete as designed.

- **CSCvw15387:** When using Workload Manager 5.3.1 to deploy a single container in a Kubernetes environment, a user encountered a failure with the container first being created and then removed.
Resolution: Workload Manager 5.4.2 includes a fix to ensure that the Kubernetes container is not removed and the deployment completes as designed.

Workload Manager 5.4.1

Workload Manager 5.4.1 Release Notes

- [Release Date](#)
- [Installation](#)
- [Upgrade](#)
- [Architecture](#)
- [Cloud SDK](#)
- [Clouds](#)
- [Administration and Governance](#)
- [Security](#)
- [Deprecated Functions](#)
- [Documentation](#)
- [Known Issues](#)
- [Resolved Issues](#)

First Published: October 1, 2020

Updated:

- October 13, 2020:
 - Updated the *Upgrade* section to include dependencies for the local repository appliance.
 - Updated the *Deprecated* section to provide support notification.
- November 6, 2020: Updated the *Documentation* section to include a list of pages that were updated.

Cloud Remote:

- Manual installers are available to install Cloud Remote on any supported cloud.
- See [Cloud Remote \(Conditional\)](#) for additional details.
- See the [Upgrade an Existing Cloud Remote Installation](#) section for details on upgrading this component.

You can upgrade to Workload Manager 5.4.1 via the Suite Admin UI. See [Update Module](#) for additional details.

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

The recommended upgrade path is to upgrade from Workload Manager 5.1.4 or 5.2.1 or 5.3 or 5.4.0 to Workload Manager 5.4. If you are in a version older than Workload Manager 5.1.4, first upgrade to Workload Manager 5.1.4 before upgrading to Workload Manager 5.4.1 (as identified in the [Update Module](#) section).

- If you are using Workload Manager 5.1.1 and earlier releases, be sure to upgrade to Workload Manager 5.1.2 before upgrading to Workload Manager 5.1.4, and then to Workload Manager 5.4.1.
- If you are using Workload Manager 5.1.2 or Workload Manager 5.1.3, you must first upgrade to Workload Manager 5.1.4 and then to Workload Manager 5.4.1.

Workload Manager and Cloud Remote have been upgraded to use RabbitMQ 3.8.5. We recommend upgrading both. *This is a change!*

If you upgrade your local repository appliance to use Workload Manager 5.4.1, verify the following dependencies:

1. Verify if the local repository appliance has the bundle store listed under the `/repo/` folder as **release-5.4.1**.
2. If release-5.4.1 is not listed in this folder, download the `ccs-bundle_artifacts-5.4.1-20201001.zip` from software.cisco.com.
3. Create the `/repo/release-5.4.1/` folder in the repository appliance.
4. Unzip the bundle contents under the `/repo/release-5.4.1/bundle` folder in the repository appliance.

Clouds that were not supported by the CloudCenter Suite can be made available as the Workload Manager architecture supports plug and play as listed in [Cloud SDK Home](#).

When using Ubuntu 14 and 18.04 in VMware environments, you must create the worker images using the Workload Manager 5.4.1 worker installers.

Installing the Management Agent on **imported VMs** is an alternate option for VMs that do not require the capability to launch applications but do require some basic Workload Manager functionality like performing platform actions.

Workload Manager 5.4.1 provides secure communication when installing the Management Agent on imported VMs. The procedure to use this improved functionality differs from the procedure provided in previous releases. To use this new functionality, you must download the ZIP file which contains the certificates required to communicate with RabbitMQ.

See [Virtual Machine Management > Install the Management Agent on a Linux VM > Process > Effective Workload Manager 5.4.1, provides secure communication when installing the Management Agent on imported VMs](#) for additional details on this functionality.

The Workload Manager communicates with the VMs it manages, using a RabbitMQ service. The following changes enhance the security of the service communication between the VM and RabbitMQ:

- In earlier releases, the Worker VMs used to connect to the RabbitMQ using a default non-configurable username and password. Effective Workload Manager 5.4.1 this username and password will be automatically generated for each region.
- In Workload Manager 5.4.1, the RabbitMQ service enforces verification using certificates for VMs that connect to it. This change prevents security issues such as man-in-the-middle attacks. The VMs are provided with a CA certificate chain and a public/private key pair as files that are used to establish connection to the RabbitMQ service.

- The automatically generated values and files are provided to VMs using the following approaches:
 - VMs deployed using Workload Manager: The Workload Manager injects the autogenerated values during creation.
 - VMs imported into Workload Manager:
 - In case of automatic installation, the Workload Manager injects the autogenerated values during agent push.
 - In case of manual installation, you must provide a ZIP file, the agentartifacts.zip, to the installer script.
- The agentartifacts.zip archive, used for manual agent installation on imported VMs will contain encrypted files.
- See [Virtual Machine Management](#) for additional details.
- Workload Manager will be deprecating support for worker image and removing the image from software.cisco.com in an upcoming release. However, you can create your own worker images using the procedure provided in the [Install Worker on a Linux Image](#).
- Workload Manager will also be deprecating support for the following clouds in an upcoming release: VMware vCloud Director 8.1 and VMware vCloud Director 9.1. See [Datacenters and Private Clouds](#) for a list of supported clouds.

The following documentation changes were implemented in Workload Manager 5.4.1:

- [Virtual Machine Management](#) (added a new section to *Install the Management Agent on Imported VMs*)
- [Local Repo Appliance \(Conditional\)](#) (changed the *Configure Cloud Settings* > Agent Bundle URL to *release-<release_version>*)
- [Configure a vCenter Cloud](#) (added a note on the cloud name assignment)
- [Configure a vCD Cloud](#) (added a note on the cloud name assignment)
- [Configure an AWS Cloud](#) (added a note on the cloud name assignment)
- [Configure an AzureRM Cloud](#) (added a note on the cloud name assignment)
- [Configure a Google Cloud](#) (added a note on the cloud name assignment)
- [Configure a Kubernetes Cloud](#) (added a note on the cloud name assignment)
- [Configure an IBM Cloud](#) (added a note on the cloud name assignment)
- [Configure an OpenStack Cloud](#) (added a note on the cloud name and region name assignments)

No known issues.

The following issue was resolved/addressed in Workload Manager 5.4.1:

- **CSCvt60783:** The **Send CMDB update** toggle switch remained disabled in the Cisco CloudCenter Suite – ServiceNow Integration.
Resolution: Workload Manager 5.4.1 includes a fix to ensure that these actions trigger an update to ServiceNow CMDB.
- **CSCvv81491:** A Workload Manager instance in an OpenStack environment did not display all unmanaged VMs from the OpenStack cloud.
Resolution: Workload Manager 5.4.1 includes a fix to ensure that all unmanaged VMs displayed in this environment.
- **CSCvv19731:** After creating a plan with a bundle limit and assigning the plan and bundle limit to a tenant, if the bundle is assigned again or updated under the CloudCenter Suite tenant options, then the previously-assigned value is use as the limit and the new value is not updated.
Resolution: Workload Manager 5.4.1 includes a fix to ensure that the updated value is used in re-assignments or updated configurations.

Workload Manager 5.4.0

Workload Manager 5.4.0 Release Notes

- [Release Date](#)
- [Installation](#)
- [Upgrade](#)
 - [Upgrade to Python Version 2.7.5 or Later](#)
- [Architecture](#)
- [Clouds](#)
- [Cloud SDK](#)
- [Services, Applications, and Deployments](#)
 - [The Create New Boot Volume Toggle](#)
 - [The /etc/hosts file](#)
 - [The Managed VMs CSV File](#)
 - [Availability Zones](#)
 - [Services Enhancements](#)
 - [Application Profile Changes](#)
- [Administration and Governance](#)
- [Policy Management](#)
- [Security](#)
- [UI](#)
- [API](#)
 - [New APIs](#)
 - [Updated APIs](#)
 - [Deprecated APIs](#)
- [Integrations](#)
- [Deprecated Functions](#)
- [Limitations](#)
- [Documentation](#)
- [Known Issues](#)
- [Resolved Issues](#)

First Published: July 30, 2020

- **Out-of-box Supported Images:**
 - In earlier Workload Manager releases, SUSE Linux was not supported for VMware and OpenStack environments. Effective Workload Manager 5.4.0, Cisco supports the following images for both VMware and OpenStack environments. *This is a change from earlier releases.*
 - SUSE Linux Enterprise 12
 - SUSE Linux Enterprise 15
 - Effective Workload Manager 5.4.0, Cisco also supports RHEL 8.x.
 - See [OOB Logical Images](#) for a complete list of supported images.
- **Custom Docker Image:**
 - Refers to the bundle store and contains the agent and service bundles.
 - This bundle is now downloadable from [software.cisco.com](#). *This is a change from earlier releases.*
 - See the [Local Bundle Store \(Conditional\)](#) and [Custom Docker Image for Scripts](#) for additional details.
- **Cloud Remote:**
 - Manual installers are available to install Cloud Remote on any supported cloud.
 - See [Cloud Remote \(Conditional\)](#) for additional details.
 - See the [Upgrade an Existing Cloud Remote Installation](#) section for details on upgrading this component.



CentOS 6/RHEL 6 deployments only work for VMware environments. All other cloud environments must update their images to Python 2.7.5 or later before upgrading to Workload Manager 5.4.

You can upgrade to Workload Manager 5.4.0 via the Suite Admin UI. See [Update Module](#) for additional details.



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

The recommended upgrade path is to upgrade from Workload Manager 5.1.4 or 5.2.1 or 5.3 to Workload Manager 5.4. If you are in a version older than Workload Manager 5.1.4, first upgrade to Workload Manager 5.1.4 before upgrading to Workload Manager 5.4 (as identified in the [Update Module](#) section).

- If you are using Workload Manager 5.1.1 and earlier releases, be sure to upgrade to Workload Manager 5.1.2 before upgrading to Workload Manager 5.1.4, and then to Workload Manager 5.4.0.
- If you are using Workload Manager 5.1.2 or Workload Manager 5.1.3, you must first upgrade to Workload Manager 5.1.4 and then to Workload Manager 5.4.0.

Upgrade to Python Version 2.7.5 or Later

A bootstrap timeout error can occur for deployments on AWS/Google/Azure RM/OpenStack clouds when the following set of conditions coincide:

- Linux Images used for deployment have Python installed with a version < 2.7.5. Usually CentOS 6 and RHEL 6 flavors have this lower version.
- HTTPS Bundle URL is configured for the region.

To work around the issue, you have one of the following options:

- Upgrade Python on these images from Version 2.6.6 to Version 2.7.5 or later, and map to the newer AMIs.
- Change the bundle URL to a HTTP-based URL in the Regions page. For example, <http://cdn.cliqr.com/release-5.4.0/bundle>.

No updates

The Workload Manager continues to support the following cloud families:

- AWS
- AzureRM
- IBM Cloud
- Google (GCP)
- VMware
 - vCenter
 - vCloud Director
- OpenStack
- Outscale
- Kubernetes

Clouds that were not supported by the CloudCenter Suite can be made available as the Workload Manager architecture supports plug and play as listed in [Cloud SDK Home](#).

This section described changes to the multiple areas.

The Create New Boot Volume Toggle

OpenStack environments can now use the **Create New Boot Volume** toggle to set the Root Disk Volume size in both the Environments or Deployments pages.

See [Multiple Volumes](#) > *OpenStack Type Nuances* and *UI Configuration Options* for additional details.

The /etc/hosts file

The /etc/hosts file was modified by the Workload Manager in certain scenarios:

- Adding the following line to the /etc/hosts file to Linux VMs/deployments.
127.0.0.1 *hostname*

```
# cat /etc/hosts127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4
::1 localhost localhost.localdomain localhost6 localhost6.localdomain6
127.0.0.1 cqjw-681bc0be3
```

- Appending the hostname to the following line:
127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4 *hostname*

```
# cat /etc/hosts127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4 cqjw-681bc0be3
::1 localhost localhost.localdomain localhost6 localhost6.localdomain6
```

Effective CloudCenter 5.4.0, the code that used to modify the /etc/hosts file has been removed and these file will remain intact for VMs/deployments.

The Managed VMs CSV File

When you download the CSV file for a Managed VM, you may not see the subtle change in functionality in the UI as the link to download the CSV file has not changed. However, when you download this file, you will see additional parameters (global and deployment-specific parameters) included in this file. *This is a change from earlier releases.*

This change does not impact unmanaged VMs or their corresponding CSV files.

See the [Virtual Machine Management](#) > Download CSV File for Managed VMs section for additional details.

Availability Zones

When configuring Availability Zones, you can select multiple Availability Zones in the Environment and lock the selection so users can select any/all/none out of the locked Availability Zones. See [Availability Sets and Zones](#) > *OpenStack* for additional details.

Services Enhancements

The new Services page introduces the following functions:

- **Pagination for Listed Services:**
 - Displays 20 services by default.
 - You can display additional services at will (clicking the More option) instead of all the services being loaded by default in this page.
- **Search by Name Filter:**
 - As a limited number of services are displayed in this page (until you click the More option), the Search function allows you to search for a service by name.
 - The search feature is listed below the Services page title.
- **Configuring Services:**
 - Along with a completely new UI experience, users will continue providing the same data as earlier releases, with the following differences:
 - The Service Type is now available as a dropdown list, but the functionality remains the same.
 - New fields like Version allow the Workload Manager to understand your service addition specifications.
 - The mandatory fields are validated: for example, provide a Version number, add at least one image, specify a default image, add a logo file, and so forth as displayed in the following image. This validation ensures that you add meaningful services in the Workload Manager.
 - See [Custom Service Definition](#) for additional details.
- **Lifecycle Actions Based on Groups:**
 - Lifecycle Actions are configured in a separate tab, but the functionality remains the same. — this is not a mandatory field for every service and this field is not validated.
 - You can continue to add multiple hooks for each service type.
 - All lifecycle actions are now grouped and displayed in a dropdown list when configuring this field in the UI.
 - The order and grouping for lifecycle actions for each service type is relevant when importing and merging services.
 - See [Service Lifecycle Actions](#) for additional details.
- **Parameters:**
 - The parameters are configured in a separate tab, but the functionality remains the same.
 - The following new parameter types are now available for service configuration:
 - JSON
 - YAML
 - The Parameter Name and Display Name are mandatory fields and are validated.
 - The String type parameters now accepts Regular Expression as a value, and the Workload Manager validates (when saving an application using this value) the provided value against this regular expression.
 - See [Using Parameters](#) for additional details.
- **Import and Export Services:**
 - Similar to application profiles, you can now import and export services.
 - **Export:**
 - Users can selectively export one or more services to easily share them with other CloudCenter Suite users.
 - Exported services include service versions including the ability to deprecate services.
 - View (read only) service permissions are required to use this option.
 - Located in the Actions dropdown menu, the Export function allows you to export a selected service as a .zip file that is saved to your local device.
 - Click the checkbox to the left of each service to multi-select services.
 - See [Service Export/Import > Exporting Services](#) for additional details.
 - **Import:**
 - View and Modify (read and write) service permissions are required to optimally use this option. With read-only permissions, service-related information that is shared from your parent tenant are not imported and details in relationship to your own tenant are imported as you do not have write permission for this service.
 - Located in the Actions dropdown menu, the Import function allows you to import a service as a .zip file to any other environment setup or a different tenant.
 - See [Service Export/Import > Importing Services](#) for additional details.
- **Merge and re-map Service Versions:**
 - When importing a service you also have the option to merge services by switching the toggle button to On.
 - Depending on the repository options and lifecycle actions used in the merged service, you can re-identify (or re-map) each repository and action to the same or different repository and action respectively for your merged setup.
 - See [Service Export/Import > Merge Duplicate Services during a Service Import](#) for additional details.
- **Multi-Select Services:**
 - When exporting or deleting service(s), you have the option to select multiple services.
 - See [Service Export/Import > Select Multiple Services when Exporting](#) for additional details.
- **View Usage:**
 - You can view the application profile or deployment where as service is being used.
 - See [Service Versions and Usage > Version Usage](#) for additional details.
- **Service Versions:**

- Similar to application profiles, you can now create versions for a service.
- Each service can have multiple versions.
- Existing services will continue to work as Version 1. No additional work is required by users.
- See [Service Versions and Usage](#) for additional details.

Application Profile Changes

In earlier releases, you were able to map a particular application profile to a service. Effective CloudCenter Suite 5.4.0, you may have multiple versions for a service. Consequently an Application Profile must be mapped to a service version. Additionally:

- When you edit/create a Service within an application profile:
 - An additional dropdown list to select a service version as displayed in the following image.
 - You must map the application profiled to the required service version, this configuration is required and validated.
 - See [Understand Application Tier Properties](#) for additional details.
- The Service category options are only loaded/expanded on demand (when you click a category).
 - In earlier releases, all category options were loaded each time you accessed the Topology Modeler page. *This is a change from earlier releases.*
 - See [Topology Modeler](#) for additional details.
- The service details are only fetched on demand (when you drag and drop a service category option).
 - In earlier releases, all the details were displayed for each tier. *This is a change from earlier releases.*
 - The last updated version is used by default for the Service Version. If you have multiple service versions, you can switch between versions for each application profile. See [Create an Application Profile](#) for additional details.
- The Deployment Parameters also include the two additional parameter types:
 - JSON and YAML.
 - See [Using Parameters](#) for additional details.

No updates

No updates

No updates

The Services related UI pages use the new v2 APIs to ensure that the payload is light when viewing the APIs using Browser Tools.

Workload Manager 5.4.0 includes the following new and updated APIs.

New APIs

The payload for the v2 service API calls requires a version number.

- DELETE — `/api/v1/tenants/{tenantId}/services?serviceIds=v1,v2`
- DELETE: `/api/v1/tenants/{tenantId}/services/{id}/versions/{versionId}`
- GET: `/api/v1/tenants/{tenantId}/services/{id}/versions/{versionId}`
- GET: `/api/v1/tenants/{tenantId}/services/webserviceParams/{serviceId}`
- GET: `/api/v1/tenants/{tenantId}/services/search?name=serviceName`
- POST: `/api/v1/tenants/{tenantId}/services/{serviceId}/versions`
- GET: `/api/v1/tenants/{tenantId}/services/export`
- POST: `/api/v1/tenants/{tenantId}/services/import`
- GET: `/api/v1/tenants/{tenantId}/services/{serviceId}/versions/{versionName}/usage`
- GET: `/api/v1/tenants/{tenantId}/services/actionsMetadata?serviceType=value`
- POST: `/api/batch`
- See [CCM Calls 5.4.0 > service-controller](#) for additional details.

Updated APIs

- The APIs that changed as a result of the service enhancements in this release have a v2 API implementation:
 - GET: `/api/v1/tenants/{tenantId}/services?size=20&sort=displayName&order=asc&parentService=true` (v1 marked for deprecation)
 - GET: `/api/v1/tenants/{tenantId}/services/{id}` (v1 marked for deprecation)
 - POST: `/api/v1/tenants/{tenantId}/services` (v1 marked for deprecation)
 - PUT: `/api/v1/tenants/{tenantId}/services/{id}` (v1 marked for deprecation)
 - DELETE: `/api/v1/tenants/{tenantId}/services/{id}`
 - See [CCM Calls 5.4.0 > service-controller](#) for additional details.
- You can use an optional parameter, `includeJobParameters=true`, to include deployment-specific parameters when downloading a CSV file for managed VMs.
 - By default the value for this parameters is `false`.
 - If set to true, the API call additionally returns all the global and deployment-time parameters for the corresponding VM.
 - See the [CCM Calls 5.4.0 > virtual-machines-controller](#) API calls for additional details.

Deprecated APIs

In the above list, some APIs are marked for deprecation. Going forward, only the v2 APIs will be used for Service management and configuration.

The [Service Management API Calls](#) v1 version is still available to ensure backward compatibility, but are not used with the new features. They will remain available in case you are using them in any existing scripts. This allows you to migrate to the new APIs.

The ServiceNow app *Integration – CloudCenter Suite v4.2.0* includes the following enhancements and updates:

- Integration and tenant configuration settings have been consolidated for easier setup. *This is a change from earlier releases.*
- A ServiceNow Requested Item record is now created regardless of the deployment request being initiated from the ServiceNow Service Portal or directly from CloudCenter Suite. In previous releases, a ServiceNow Request was created for deployments initiated from CloudCenter Suite and a Request with Requested Item was created for deployments initiated from the ServiceNow Service Portal. *This is a change from earlier releases.*
- When VMs are imported into CloudCenter Suite, their details and status can now be added to the ServiceNow CMDB. This can be enabled or disabled from the *Tenant Configuration* module.
- If a deployment in CloudCenter Suite remains in a *Starting* state for 2 hours, the deployment workflow in ServiceNow is terminated. This prevents ServiceNow Requests from remaining in an *Open* state.
- If a new deployment request is canceled in CloudCenter Suite prior to being approved in ServiceNow, the Request record in ServiceNow is marked as *Closed*.
- The Approval request now contains a *Zone* value for OpenStack environments.
- The Virtual Machine table in CMDB now displays the *Start Date*, *Terminated On*, and *VM running hours* fields
- The Storage Volumes table in CMDB now the *Storage Type* field.
- A new *info* variable of RITM contains additional details about requested action, job or VM, which can be used in workflows to automate approvals, display on the form, or define catalog tasks.
- Resolved an issue where validating the owner API key in ServiceNow created a new record in the Tenants Mapping module.
- Resolved an issue where certain Day 2 actions resulted in multiple ServiceNow approval requests. For example, a resize action may require a VM to be stopped, resize and then restarted. This action now creates a single approval request.
- See [ServiceNow Extensions](#) for additional details.
- **CSCvu66141:** Previous versions of Workload Manager supported integration with [Arcus](#). However, this announcement is to provide notification that Arcus is deprecated, and that [Workload Manager 5.3.1](#) is the last release with integration support.
- The *worker1* and *worker_full* modes are not supported effective CloudCenter 4.10.0.9. The *worker_basic* mode is the only supported mode when installing the worker.
- [Microsoft announced](#) the End of support for Windows Server 2008 and Windows Server 2008 R2. Consequently, Workload Manager will be deprecating support for this server and removing image support in the corresponding UI pages and APIs in this release. This includes the removal of Windows 2008 references in the following pages:
 - [OOB Logical Images](#)
 - [Supported OOB Services](#)
 - [Dynamic Bootstrapping](#)
 - [Install Worker on a Windows Image](#)
 - [Images Page](#)
 - [Virtual Machine Management](#)

No updates

The following documentation changes were implemented in Workload Manager 5.4.0:

- [Cloud Remote \(Conditional\)](#) (added a use case to the *Navigating Cloud Remote through Proxy* section)
- [Define Resource Placement](#) (added a note to the overview section and two commonly used variables to the table)

Workload Manager 5.4.0 has the following known issues:

- The Azure blade uses Version 1.10 of Windows custom extensions. However, the Regions setting page for existing regions may display versions as visible in the following image.

The screenshot shows the 'Cloud Settings' page for the 'US West (California)' region. The page is divided into two main sections: 'Regions' and 'Accounts'. The 'Regions' section is active, showing the region name and a status of 'Region: Enabling...'. The 'Accounts' section is inactive. The 'Cloud Settings' section is visible, showing a list of settings for the 'AZURE' environment. The settings include: Azure Environment (AZURE), Linux Custom Script Extension Version (2.1), Windows Custom Script Extension Version (1.1), Linux Diagnostics Extension Version (2.0), and Windows Diagnostics Extension Version (1.6). There are also options to 'Exclude these special characters for Windows password', 'Delete Boot Diagnostic Logs On VM', and 'Termination Agent Bundle URL'.

- If user scripts print a password for an [External Service](#), this password may be visible in the Deployment logs. Ensure that you are not printing sensitive passwords as part of a script output.

The following issue was resolved/addressed in Workload Manager 5.4.0:

- **CSCvu54099:** When the Default Deployment Environment feature is used in an Application Profile, the deployment screen disappears for a few seconds after clicking the **Deploy** button.
Resolution: Workload Manager 5.4.0 includes a fix to ensure that this issue is fixed and the deployment screen displays as designed.
- **CSCvu40570:** When instance types for the VMs were launched during the deploy flow, the instance type details were not displayed immediately. Additionally, standard users were unable to see these details.
Resolution: Workload Manager 5.4.0 includes a fix to ensure that the instance type details were displayed immediately and that standard users can see the same instance type details as an admin user.

- **CSCvs38738:** A user was unable to load a Action Orchestrator workflow from the Workload Manager application profile configuration page. During addition the system performed the necessary redirects in the Web Browser and then displayed the following error message: *Loading... Create AO.*
Resolution: Workload Manager 5.4.0 includes a fix to ensure that the Action Orchestrator - Workload Manager communication was successful, even with multiple proxy entries as this issue was a result of the bypass proxy list in the Suite Admin configuration settings.

What Is Supported?

What Is Supported by Workload Manager?


- Public Clouds
- Datacenters and Private Clouds
- Container Clouds
- Conditional Component Appliance Images
- OOB Logical Images
- OOB Services
 - Supported OOB Services
 - Apache Service
 - Chef Service
 - Docker Service
 - MySQL Service
 - Nginx Service
 - Puppet Service
 - Tomcat Service
 - Varnish Service
- Dynamic Bootstrapping
- OOB Application Templates
- OOB Groups, Roles, and Permissions
- Permission Control
- Understand ACLs

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-box Workload 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)

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 8.1
	VMware vCloud Director 9.1
VMware vCenter	VMware vCenter 6.0
	VMware vCenter 6.5
	VMware vCenter 6.7
OpenStack	OpenStack Newton
	OpenStack Mitaka
	OpenStack Pike
	OpenStack Queens

To compute costs in Cost Optimizer, you must specify the compute 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. See [Images](#) for more context.

Container Clouds

Supported Container Clouds

- [Overview](#)
- [Requirements](#)
- [Upstream Support and Capability](#)

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.

Kubernetes cloud configurations require:

- [Kubernetes](#) version 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)

Workload Manager supports *upstream* Kubernetes setups. *Upstream* refers to any bare Kubernetes setup like Google 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 Manager dynamically deploys applications based on the Workload Manager application 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. See [Container 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. See [Container Service > Deploying a Container Service > Firewall Rules](#) for details.

See [Container Tier Rolling Updates](#) for details.

Conditional Component Appliance Images

Conditional Component Appliance Images

- [Appliance Availability](#)
- [Cisco Provided Appliance File Names](#)
 - [AzureRM Appliances](#)
 - [OpenStack Appliances](#)
 - [vCenter Appliances](#)
- [Supported OS Types for User-Built Appliances](#)

Cisco provides images for the conditional Cloud Remote appliance, Pre-bootstrapped Worker appliance, and the Local Repo appliance as described in the following table.

Cloud	Workload Manager Components			Distribution
	Cloud Remote	Local Repo ¹	Pre-bootstrapped Worker ²	
AWS	Yes	Yes	No	Shared AMI
AzureRM	Yes	Yes	No	Download VHD images (software.cisco.com)
OpenStack	Yes	Yes	No	Download QCOW2 images
vCenter	Yes	Yes	Yes	Download OVA images
Google	Yes	No	No	Shared VMI

¹ The Local Repo virtual appliance supports three subcomponents: Bundle Store, Package Store, and Docker Registry. See [Local Repo Appliance \(Conditional\)](#) for additional details.

² The Worker images are built using CentOS 6. See [Worker \(Conditional\)](#) for additional details.

When the appliances are provided as downloadable files at software.cisco.com, they are named using the conventions summarized below

AzureRM Appliances

Component	Name
Cloud Remote	azure-cc-centos7-cloudremote-<release.tag>-YYYYMMDD.0.zip
Local Repo	azure-cc-centos7-repo-<release.tag>-YYYYMMDD.0.zip

OpenStack Appliances

Component	Name
Cloud Remote	openstack-cc-centos7-cloudremote-<release.tag>-YYYYMMDD.0.qcow2
Local Repo	openstack-cc-centos7-repo-<release.tag>-YYYYMMDD.0.qcow2

vCenter Appliances

Component	Name
Cloud Remote	vmware-cc-centos7-cloudremote-<release.tag>-YYYYMMDD.0.OVA
Worker	vmware-cc-centos7-worker1-<release.tag>-YYYYMMDD.0.ova
Local Repo	vmware-cc-centos7-repo-<release.tag>-YYYYMMDD.0.ova

The following base OS images may be used to create the user-built appliance using the Cisco provided installer programs.

Workload Manager User-Built Appliance	Supported OS

	Public Clouds and Datacenters and Private Clouds
Local Bundle Store	CentOS7
Local Package Store	<ul style="list-style-type: none">• RHEL7• CentOS7• Ubuntu1404
Pre-bootstrapped Worker	See Management Agent (Worker) for supported Linux and Windows OS versions
Cloud Remote Artifact	CentOS7 (see Cloud Remote (Conditional) for additional details) Name: ccs-cloudremote-artifacts-<release.tag>-YYYYMMDD.0.zip



We recommend that you only use the Installer method for unique installation scenarios (for example, if you use custom OS images).

The installer files are contained in the artifacts.zip file which can be downloaded from software.cisco.com.

OOB Logical Images

OOB Logical Images

Workload Manager supports many popular OS images for Linux and Microsoft Windows VMs. These are represented by logical images that are translated to physical cloud images at deploy time. The correlation between the logical image and physical cloud image is defined in a process called image mapping (see [Images Overview](#)). Workload Manager includes the following logical images:

- CentOS 6
- CentOS 7.x
- RHEL 6.x
- RHEL 7.x
- RHEL 8.x
- SUSE Linux Enterprise 12
- SUSE Linux Enterprise 15[†]
- Ubuntu 14.04
- Ubuntu 16.04
- Ubuntu 18.04
- Windows Server 2012
- Windows Server 2012 with MSSQL 2012
- Windows Server 2016
- Windows Server 2019*

* Windows Server 2019 is not supported OpenStack.

[†] The image for SUSE Linux Enterprise 15 is sourced as an OOB image – however, this image is not mapped to any service. Once you upgrade your setup to Workload Manager 5.4, you (with a Service Manager role) must map this image to a service before you start using it. To map the image to a service, edit the required service and then add the SUSE Linux Enterprise 15 image to this service. See [Custom Service Definition \(Edit Service\)](#) for additional details. This step ensures that the SUSE Linux Enterprise 15 image is displayed in the base images dropdown when editing an [Application Profile](#).

OOB Services


OOB Services

- [Supported OOB Services](#)
- [Apache Service](#)
- [Chef Service](#)
- [Docker Service](#)
- [MySQL Service](#)
- [Nginx Service](#)
- [Puppet Service](#)
- [Tomcat Service](#)
- [Varnish Service](#)


Supported OOB Services


Supported Out-of-Box Services

- [Application VM Support](#)
- [PaaS Support](#)
- [Commercial Support](#)

 Cisco does not provide out-of-box image mapping for [Datacenters and Private Clouds](#). Once you set up image mapping (see [Map Images](#)) for application VMs, Workload Manager out-of-box services (listed in this page) are automatically installed and displayed in the Workload Manager UI. See [Application Tier Properties](#) for details on configuring these services.


The following table lists the Out-of-Box (OOB) service categories that are included with the Workload Manager module. They are based on CentOS 6.x unless otherwise specified.

 The OOB services included with the Workload Manager module are merely examples that serve as a point of reference. If you plan to use these in your environment, be sure to review them and ensure that they meet your operational and security standards.

Service Type	Service	Version	License
Orchestration	Chef	12.6 (latest)	Apache License 2.0
	Puppet	4.3 (latest)	Apache License 2.0
Message Bus	ActiveMQ	5.8.0	Apache License 2.0
	RabbitMQ	3.5.1	Mozilla Public License
Web Server	Apache 2	2.2.15	Apache License 2.0
	Geronimo3	3.0.0	Apache License 2.0
	IIS	7 & 8	Windows
	Jetty	9.4.7	Apache License 2.0
	Ruby	1.8.7 and 1.9.3	Ruby, GPLv2, FreeBSD
	Ruby on Rails	2.3.14	MIT License
	 Ruby on Rails is only supported on Ubuntu12.04		
	Tomcat 6	6.0.43	Apache License 2.0
	Tomcat 7	7.0.59	Apache License 2.0
Frontend Cache	Varnish	4.0.4	BSD License
Load Balancer	HAProxy	1.5.4	GPLv2
	Nginx	1.0.15	FreeBSD License
Backend Cache	Memcached	1.4.4	Revised BSD License
Database	MySQL	5.6.27	GPLv2
	SQL Server	2012	Windows
NoSQL Database	Cassandra	2.0.17	Apache License 2.0
	MongoDB	2.6.7	GNU AGPL v3.0
Others ?? Jagadeesh ??	Docker	1.7.1	Apache License 2.0
	Jmeter	2.12	Apache License 2.0

The **Services** menu has moved out of the **Admin** menu and is now available in the Workload Manager Main Menu. Management of services can be done by users belonging to the Service Managers group.

When installing the CloudCenter platform, the following services are dynamically displayed (out-of-box) for each cloud:

No.	CloudCenter Supported Services	Supported Clouds
1	Relational Database Service (RDS)	AWS (latest version) <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> The <code>vpcId</code> and <code>dbSubnetGroup</code> deployment parameters are added to the RDS Service and you can pass these values at deployment time. See External Service for additional context.</div>
2	Elastic Load Balancing (ELB)	AWS (latest version)

Cisco additionally provides commercial services that are not included in the above categories due to licensing issues. Some examples include, but are not restricted to, Microsoft SQL Server, IBM WebSphere, Oracle WebLogic, Oracle Database, and WSO2 Application Server. CloudCenter customers and partners can directly model these services or contact your CloudCenter Suite account executive or partner for additional support.

Apache Service

Apache Service

- [Overview](#)
- [Guidelines](#)
- [General Settings](#)
- [Remaining Sections](#)

See [Understand Application Tier Properties](#) for general nuances and details for each field in the **Properties** pane when configuring a service.

This page only identifies the DIFFERENCES, dependencies, and best practices for this service.

Can run as a single instance or as part of a cluster.

Properties	Description
Scaling Policy	See Policy Management for additional context.

See [Understand Application Tier Properties](#)

Chef Service

Chef Service

- [Overview](#)
- [Guidelines](#)
- [General Settings](#)
- [Clustering Support](#)
- [Remaining Sections](#)

See [Understand Application Tier Properties](#) for general nuances and details for each field in the **Properties** pane when configuring a service.

This page only identifies the DIFFERENCES, dependencies, and best practices for this service.

The following table identifies the terminology used when associating a [Chef](#) service.

Chef Service	Associated Terminology
Repository Dependency	Server
Service Association	Client
Configuration Reference	Recipes

The Chef Server is available as a repository type. See [Artifact Repository](#) for a complete list of available types for repositories.



Workload Manager does not allow you to provision or manage the Chef Servers either directly or indirectly. Be sure to *pre-configure* these servers before adding either as repositories in Workload Manager.

The service script installs, configures, starts, and stops the Chef Client. Each service requires the information to setup the Chef Client and configure the **Chef Recipe** and **environment name**.

Properties	Description
Minimum Number of Nodes	Default = 1. The minimum number of nodes within each tier in use to ensure manual or automatic scaling.
Maximum Number of Nodes	Default = 2. The maximum number of nodes within each tier in use to ensure manual or automatic scaling.
Chef Server	Required. The configured repository dependency for the Chef server.
Chef Organization	Required. The organization to which this Chef Server belongs.
Chef Recipe	Required. Identifies the configuration reference for this service.
Chef Environment	Required. The service script invokes the Chef Client and registers itself to the selected Chef Server. <div style="border: 1px solid green; padding: 5px; margin-top: 10px;"> To associate the Chef repository (when you Project and Phase Management) with the Workload Manager deployment environment, use the %DEP_ENV_NAME% macro. When a user provides this macro, the macro is replaced with the environment name dynamically at runtime. When users configure the environment, they have the option to provide additional granular parameters to complete this configuration. See Pre-Defined Parameters and Using Parameters for additional context. </div>

See [Service Administration > Custom Service Clustering](#) for details on using parameters that are required to allow a Chef-defined tier to scale.

See [Understand Application Tier Properties](#).

See [Applications Using Puppet and Chef](#) for additional details.

Docker Service

Docker Service

- [Overview](#)
- [Guidelines](#)
- [General Settings Properties](#)
- [Connect to Docker Containers](#)
- [Remaining Sections](#)
- [Dedicated Docker Containers](#)

See [Understand Application Tier Properties](#) for general nuances and details for each field in the **Properties** pane when configuring a service.

This page only identifies the DIFFERENCES, dependencies, and best practices for this service.

Workload Manager provides the Docker integration as a custom service and enables you to import and launch Docker containers on any Workload Manager supported VM-based cloud. Users can drag and drop the Docker service into the Topology Modeler graphical workspace and create a topology with single or multiple Docker containers.

The screenshot displays the 'Edit "Two-Container Wordpress" Application Profile' window. The interface is divided into several sections:

- Services:** A list of services including Frontend Cache, Load Balancer, and Web Server. Under Web Server, there are options for Apache2, Geronimo3, IIS, and Jetty.
- Topology Modeler:** A central workspace where a 'docker_...' service is connected to a 'DB' service. The DB service is configured with 'CPU: 1', 'Memory: 1GB', and 'Storage: 10GB'.
- Properties:** A pane on the right showing configuration options for the selected service.
 - General Settings:** Includes a field for 'Default Volume Size *' set to '10 GB' and a note about persistent data storage for database files.
 - Docker Container *:** Shows the 'Container Profile' set to 'tutummysql'.
 - Service Initialization:** A section for defining initialization steps.

This allows enterprises to create a mixed topology of applications using Docker containers and services and launch them directly on VMs thus providing the following benefits:

- Enterprises can use Workload Manager for governance as well as container management.
- The IT group can oversee the Docker container usage on a per department basis and charge users for services.

When you add the Docker service in the Topology Modeler, you must specify the Docker container details in General Settings Properties pane. Click **Add** to add additional Docker containers and specify additional parameters for each Docker container.

Add Docker Container

Container Profile*

Unique name for Docker container

Docker Image*

Select Docker image from Docker hub, Github, Repo or Storage Please type in full path. Your file will be downloaded from an unknown repository

Port Mapping

Hostport:DockePort mappings seperated by comma


Docker Parameters

Additional Docker command parameters

Additional Commands

Commands for applications to be run in Docker

When you add a Docker container, you can specify the following parameters:

Service Parameter	Description
Container Profile	The Docker container profile's unique name.
Docker Image	The Docker Image that you can import from the repositories, Workload Manager-supported storage, Docker Hub, or GitHub. Docker images can also be automatically pulled from a GitHub repository.
Port Mapping	You can add multiple Hostport:DockePort mapping, separated by commas.
Docker Parameters	Assign additional Docker-specific parameters to this profile.
Additional Commands	Parameters to execute additional scripts, commands, or programs inside the Docker image.
IP Address	Static IP address assigned to the Docker container. <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  For containers on different VMs to communicate, ensure that the TCP/UDP 6783 ports are open on the relevant VMs. </div>
Linked Containers	You can add multiple ContainerName:AliasName container lists, separated by commas.

Once you add additional containers, you see the container name displayed in the General Settings Properties pane. Similarly, you can manage multiple containers for a single VM in the General Settings tab.

When you launch an application with the Docker container, the corresponding job status messages for each run are displayed on a per-instance basis in the **Runs** page. You can access this information by selecting the **Show Deployments** option in the Application dropdown list.

Workload Manager provides the ability to model application with services that are running on the VM and can connect to a Docker container. This real world enterprise use case, addresses scenarios where some services are not available in Docker but are required to launch the application.

"NT-docker-rds-AWS-3107_run_1" Deployment Details Auto refresh every 30 seconds

Name	NT-docker-rds-AWS-3107_run_1	Start Time	2015-07-31 05:24:55
Version	1.0	End Time	N/A
Deployment Environment	Amazon	Cloud	Amazon US West (Oregon)
Status	JobStarting	Status Message	N/A
Aging Policy	1hr_suspend	Last Update Time	2015-07-31 05:24:55
Promoted from	-	Deployment Initiated by	Sriram (sriram@clqr.com)
Approved/Rejected by	-	Approval Requested on	N/A
Approved/Rejected on	N/A	Approval Status	-
Approval/Rejection Comment	-		
Terminate Protection	Disabled		

DB

Status	JobRunning
Status Message	N/A
Start Time	2015-07-31 05:25:00
End Time	N/A
Cloud Name	Amazon US West (Oregon)
Instance Type	db.m1.medium (410 GB Local Storage, 1 CPU, 3750 MB Memory)
VPC Identifier	vpc-312eac54
Scaling Policy	N/A
Security Profiles	N/A
Storage	N/A
Storage IP	N/A

Associate Tags

Enter tag name

▶ (running) rdsmysql-3360.cckxv74syde.us-west-2.rds.amazonaws.com

Thus, Workload Manager's service additions provide the flexibility to:

- Model standalone Docker containers.
- Add multiple Docker containers running on single VM.
- Use services/other applications to connect to several Docker containers.

See [Understand Application Tier Properties](#).

See [Local Bundle Store \(Conditional\)](#) for installation and configuration details.

MySQL Service

MySQL Service

- [Overview](#)
- [Guidelines](#)
- [General Settings](#)
- [Remaining Sections](#)

See [Understand Application Tier Properties](#) for general nuances and details for each field in the **Properties** pane when configuring a service.

This page only identifies the DIFFERENCES, dependencies, and best practices for this service.

The Tomcat service tier has a dependency on the MySQL service tier. The MySQL tier must be invoked before the Tomcat tier as the database must have an IP address to configure any database connection.

Properties	Description
Root Password	The password for the MySQL instance.
DB Setup Script	The SQL file script. The location of the SQL file to be loaded into the MySQL instance. Specify the file from the relative path % rootPath%.

See [Understand Application Tier Properties](#)

Nginx Service

Nginx Service

- [Overview](#)
- [Guidelines](#)
- [General Settings](#)
- [Service Initialization](#)
- [Remaining Sections](#)

See [Understand Application Tier Properties](#) for general nuances and details for each field in the **Properties** pane when configuring a service.

This page only identifies the DIFFERENCES, dependencies, and best practices for this service.

Runs on a single VM (not a cluster).

Nginx is treated as a load balancer and fields like number of nodes, App Package, App Config files, Deploy Folder are not applicable for this service.

Manually written script is required based on the application being in the pre-start or post-start mode for this service.

See [Understand Application Tier Properties](#).

Puppet Service

Puppet Service

- [Overview](#)
- [Guidelines](#)
- [General Settings](#)
- [Remaining Sections](#)

See [Understand Application Tier Properties](#) for general nuances and details for each field in the **Properties** pane when configuring a service.

This page only identifies the DIFFERENCES, dependencies, and best practices for this service.

The following table identifies the terminology used when associating a [Puppet](#) service.


Puppet Service	Associated Terminology
Repository Dependency	Primary server
Service Association	Agent
Configuration Reference	Manifests

The Puppet primary server is available as a repository type. See [Artifact Repository](#) for a complete list of available types for repositories.



Workload Manager does not provision or manage the Puppet primary server either directly or indirectly. These servers must be pre-configured before adding either as repositories in Workload Manager.

The service script installs, configures, starts, and stops the Puppet Agent. Each service requires the information to setup the Puppet Agent and configure the **Puppet role** and **environment name**.

Properties	Description
Minimum Number of Nodes	Default = 1. The minimum number of nodes within each tier in use to ensure manual or automatic scaling.
Maximum Number of Nodes	Default = 2. The maximum number of nodes within each tier in use to ensure manual or automatic scaling.
Puppet Servers	Required. Select the configured Puppet server repository (for example, Puppet primary server).
Puppet Role	Required. Configure the application using the configured Puppet role.
Puppet Environment	Required. The service script leverages the function of Puppet Agent to configure the environment and deploy the application. <div style="border: 1px solid green; padding: 5px; margin-top: 10px;">  To associate the Puppet repository (when you Manage Projects and Phases) with the Workload Manager deployment environment, use the <code>%DEP_ENV_NAME%</code> macro. When a user provides this macro, the macro is replaced with the environment name dynamically at runtime. When users configure the environment, they have the option to provide additional granular parameters to complete this configuration. See Pre-Defined Parameters and Using Parameters for additional context. </div>
Agent Run Interval	Required. Identifies the frequency at which the service script must leverage the function of Puppet Agent when deploying the application.

See [Understand Application Tier Properties](#).

See [Model Applications Using Puppet and Chef](#) for additional details.

Tomcat Service

Tomcat Service

- [Overview](#)
- [Guidelines](#)
- [General Settings](#)
- [Remaining Sections](#)

See [Understand Application Tier Properties](#) for general nuances and details for each field in the **Properties** pane when configuring a service.

This page only identifies the DIFFERENCES, dependencies, and best practices for this service.

The Tomcat service tier has a dependency on the MySQL service tier. The MySQL tier must be invoked before the Tomcat tier as the database must have an IP address to configure any database connection.

	Properties	Description
1	App Run-time	Required and available for the Tomcat service. Identifies the JDK version number (JDK 6 or JDK 7)
2	App Package	Must be a .war file. Application package file (the binaries for the web server). The file is in relative path from http://env.cliqrtech.com/ . See Application Using App Package for additional context.
3	App Config files	Optional. Where the property file for the Workload Manager system token replacement occurs during runtime.

See [Understand Application Tier Properties](#).

Varnish Service

Varnish Service

- [Overview](#)
- [Guidelines](#)
- [General Settings](#)
- [Service Initialization](#)
- [Remaining Sections](#)

See [Understand Application Tier Properties](#) for general nuances and details for each field in the **Properties** pane when configuring a service.

This page only identifies the DIFFERENCES, dependencies, and best practices for this service.

By default, Varnish is always a single instance.

Varnish is a front-end cache system and fields like number of nodes, App Package, App Config files, Deploy Folder, Scaling Policy are not applicable for this service.

Manually written script is required based on the application being in the pre-start or post-start mode for this service.

See [Understand Application Tier Properties](#).

Dynamic Bootstrapping

Dynamic Bootstrapping

- [Overview](#)
- [Workload Manager Detection](#)
- [Cloud Support](#)
- [HTTPS Dependencies](#)

The CloudCenter Suite platform enables enterprises to use the *init script through user data* option to dynamically bootstrap custom cloud images.

Subsequently, the [Management Agent \(Worker\)](#) communicates with the bundle store and package store and installs the remaining components.

Images built on any OS version listed in [OOB Logical Images](#), allow Workload Manager to detect if the agent is missing on such VM images and automatically pushes the agent to the VMs at provisioning time.

The Workload Manager management agent can be dynamically installed on Application VMs launched from images if they do not have the management agent installed.

Indicator	Description
Yes	If Dynamic Bootstrapping is supported, you have two options: <ul style="list-style-type: none"> • Use the public image that is already mapped or map to a publicly-available image. • Build your own custom image with the worker installed (see Management Agent (Worker) > <i>Using the Worker Installer Executable</i>).
No	If Dynamic Bootstrapping is not supported (for example, VMware), then you can build your own custom image with the worker installed (see Management Agent (Worker) > <i>Using the Worker Installer Executable</i>).

The following table shows the clouds and images for which dynamic bootstrapping is supported.

Cloud Name	AWS	Outscale	AzureRM	Google	OpenStack	IBM Cloud
Windows 2012			Yes ²			No
Windows 2012 with MSSQL			Yes ³			No
Windows 2016	Yes	Yes	Yes	No	Yes	Yes
CentOS 6	Depends on the setup	Depends on the setup	Openlogic	Yes	Yes	Yes
CentOS 7	Yes	Yes	Openlogic	Yes	Yes	Yes
RHEL 6	Yes	Yes	Yes	Yes	No	Yes
RHEL 7	Yes	Yes	Yes	Yes	No	Yes
Ubuntu14	Yes	Yes	Yes	Yes	Yes	No
Ubuntu16 ¹	Yes	Yes	Yes	No	Yes	No
Ubuntu18	Yes	Yes	Yes	No	Yes	No

¹ The default Ubuntu 16.04 image from cloud providers uses Python 3. However, the CloudCenter Suite platform expects a dynamically bootstrapped VM to use Python 2. Ensure to install Python 2 in any VM that uses this version of Ubuntu. This Python 2 requirement does not apply to worker images on application VMs.

² Windows with cloud-init (set to automatically run the user data as a script (default behavior)).

³ Image mappings for **Windows Server 2012 with MSSQL 2012** have been removed in AzureRM as these images do not support cloud-init anymore. To use these services, create your own custom image and add the mapping.



Google Cloud Nuances

Windows Bootstrapping does not work on default public images due to the lack of an administrator user for Google cloud. As a result, cloud-Init (bootstrap) scripts are not executed on these instances.



OpenStack Nuances

Linux images used for dynamic bootstrapping in OpenStack must have the [net-tools package](#) preinstalled.

The Bundle Store configuration procedure defaults to using the HTTP protocol. If you prefer to use HTTPS to ensure a secure connection, adhere to the following requirements:

- Pre-install certificates on the Worker image.
- Verify your cloud dependencies. For example, if your cloud is running a Python script to dynamically bootstrap a Linux VM, be aware that the Linux Worker image uses Python Version 3.4 or later.



The Worker image (see [Worker \(Conditional\)](#)) requires Python Version 3.4 or later, to use Python scripts for dynamic bootstrapping purposes.

Administrators need to assign explicit privileges to the **cliqruser** role if additional software must be installed.

As part of our Security Hardening, the umask settings for all Workload Manager components is set to 077. As a result, you must set the unmask rule to 022 to install any additional software.

- Additionally, the Python script requires that you install openssl-devel lib to support HTTPS certificate validation.

OOB Application Templates

Supported Out-of-Box Application Templates

The Workload Manager module includes these ready-to-use templates for the building N-Tier application profiles:

- Generic
- Apache Web
- Java Web
- PHP Web
- Windows.Net Web
- Ruby on Rails Web

OOB Groups, Roles, and Permissions

OOB Groups, Roles, and Permissions

- [OOB Groups and their Roles](#)
- [OOB Roles and their Permissions](#)
- [Creating a Read-Only User for Workload Manager](#)

Workload Manager comes with several OOB groups, each group contains one or more roles, and each role has its own set of permissions.

Workload Manager users may be assigned to one or more groups. Each group, in turn, may contain one or more roles, where each role gives the user certain permissions. The following table summarizes the OOB groups for Workload Manager and their associated roles.

OOB Group	Associated Roles
Workload Manager admins	WM_ADMIN, SUITE_TENANT_ADMIN, SUITE_USER_ADMIN
Workload Manager Standard User	WM_USER
Deployment Environment Managers	WM_ENVIRONMENT_MANAGER
Application Architects	WM_APPLICATION_ARCHITECT
Project Managers	WM_PROJECT_MANAGER
Policy Managers	WM_POLICY_MANAGER
Dev Ops Users	WM_DEV_OPS
Service Managers	WM_SERVICE_MANAGER
Image Managers	WM_IMAGE_MANAGER

The following table summarizes the OOB roles for Workload Manager and their associated permissions sorted by increasing levels of permissions.

Role	Associated Permissions
WM_USER	Deploy applications, benchmark applications, view own deployments, view own VMs
WM_ENVIRONMENT_MANAGER	Create and manage deployment environments and policies
WM_APPLICATION_ARCHITECT	Create and manage application profiles
WM_PROJECT_MANAGER	Create and manage projects
WM_POLICY_MANAGER	Create and manage policies
WM_DEV_OPS	Create and manage custom on-demand and lifecycle actions
WM_SERVICE_MANAGER	Create and manage services
WM_IMAGE_MANAGER	Create and manage images
WM_ADMIN	All of the above permissions plus: create and manage clouds and cloud accounts

To create a read-only user, add the user ID as comma-separated values to the key **read.only.user.ids=** into cloudcenter-manager configmap. The CCM service will restart after you edit the configmap. The user with read-only permission can only operate GET call in CCM.

Permission Control

Permission Control

- [Role-Based Permissions](#)
- [Resource-Based Permissions](#)
 - [Deployment Permissions](#)
 - [Deployment Environment Permissions](#)
 - [Extension Permissions](#)
 - [Application Profile Permissions](#)
 - [Repository Permissions](#)
 - [Service Permissions](#)
 - [Actions Library Permissions](#)
 - [Image Permissions](#)
 - [Temporary Permission to Launch an Image](#)
- [Tenant Owner Permission Nuances](#)
- [Project and Phase Permissions](#)

Role-based permissions are a set of permissions that can be individually configured for each role. Roles are assigned to groups and users are assigned to groups and inherit those roles. This process is handled in the Suite Admin as described in [User Tenant Management](#). Workload Manager comes with its own [OOB Groups, Roles, and Permissions](#).

Resource-based permissions control how users, members of user groups and, in some cases, tenants associated with a resource can share the resource and perform related activities.

Resource-based permissions are available to resource owners, users who created the resource, and users who are permitted to share the resource. These users can grant permissions to other users.

Deployment Permissions

The deployment owner is always associated with a deployment and can:

- Manage web SSH/VNC access to a deployment VM
- Control which other users have access to deployment VMs



Only the deployment owner can control permissions and cannot provide manage permissions to any other user – *no other user can control permissions for this deployment.*

From the *Share Popup* (see [Understand ACLs](#)) for a deployment, the deployment owner (referred to as *owner*) can control permissions for a deployment, as summarized in the table below.

Permission	Description
Access	Controls whether other users/groups/tenants have SSH or VNC access to VMs in this deployment. Access permission is only effective for users/groups/tenants that also have Access or Manage privilege for the deployment environment associated with this deployment.


Deployment Environment Permissions

The tenant administrator can:

- Manage who has access to the deployment environment.
- Control which other users have access to the deployments in this environment.
- Deploy applications to or promote applications from this environment.
- Approve the deployments of applications to the environment.
- Share the deployment environment with users who are directly under the tenant owner – these users can manage the environment, if they have inherited **deployment environment permissions** based on a role configuration. Users further down this tenant hierarchy can only view the environment, if shared, in read-only mode.

Additionally, **All users in your (my) tenant** can control deployment environment permissions as described in the following table:

Permission	Deployment Environment Implications	Description
------------	-------------------------------------	-------------

Deploy To	A member of your tenant has permission to deploy applications to this deployment environment.	<p>This permission is used to provide permission to a user to deploy in this deployment environment.</p> <p>All users in the tenant with the Deployment Environment permission enabled in their role automatically have permission to manage all environments in the tenant.</p> <p>Conversely, users outside the tenant can no longer be given permission to modify or manage any environment in the tenant.</p> <p>You can restrict environment availability deployment permissions for individual users within and outside the tenant and for groups within the tenant, by clicking the Deploy To checkbox for those users/groups – these users/groups will inherit read-only access to all policies and tags specified in that deployment environment.</p>
User's Deployments	Identifies permission for deployments launched by you (the user) in this deployment environment	<p>Controls the activities that users can perform on deployments that they started in this deployment environment.</p> <ul style="list-style-type: none"> • None: The user or member of your tenant and/or sub-tenant cannot view deployments – even if this user owns the deployment. • Access: The user or members of your tenant and/or sub-tenant can view deployments • Manage: The user or members of your tenant and/or sub-tenant can manage deployments, including view, start, suspend, reboot, resume, upgrade, and terminate deployments.
Others' Deployments	Identifies permission for deployments launched by other users in this deployment environment	<p>Controls the activities that users or members of user groups can perform on deployments that other users started in this deployment environment.</p> <ul style="list-style-type: none"> • None: The user or member of your tenant and/or sub-tenant cannot view deployments – even if this user owns the deployment. • Access: The user or members of your tenant and/or sub-tenant can view deployments • Manage: The user or members of your tenant and/or sub-tenant can manage deployments, including view, start, suspend, reboot, resume, upgrade, and terminate deployments.
Promote From	A member of your tenant has permission to promote a running deployment from this deployment environment to another deployment environment.	<p>If both deployment settings (User's Deployments and Others' Deployments) are set to None for this user or users within a tenant, then this setting is <i>greyed</i> out and you will not be able to check this box as these viewers will not be able to view the deployment, and hence cannot promote it!</p> <p>Be sure to provide Access permission for either of these settings if you want to allow this user to promote deployments.</p> <div style="border: 1px solid #f0e68c; padding: 5px; margin-top: 10px;">  When you create a Deployment Environment and share it with a user without checking the Promote from option, be aware that the Migrate/Promote From action <i>will not be available</i> when this user deploys an application that uses this deployment environment. </div>
Authorized Approver	A member of your tenant has permission to authorize approvals for a deployment.	<p>Allows a user to approve the start of a deployment in the environment, if approval is required. By providing this permission, you are essentially authorizing this user to be an admin for the deployments within your deployment environment.</p> <p>If a user's deployment requires approval and the user does not have Authorized Approver permission, then the deployment must be approved by someone else before it being deployed.</p>

Extension Permissions

The Workload Manager administrator is always associated with an Extension and can:

- Manage who has access to the Extension
- Control which other users have access to the Extension
- Deploy applications to or promote applications using these Extensions
- Approve the deployments of applications using these Extension

Administrators can control permissions for an Extension as described in the *Share Popup* (see [Understand ACLs](#)). The following table describes the permission options.

Permission Options	Description
--------------------	-------------

Access	Controls permissions to users, groups, and tenants when using an Extension. <ul style="list-style-type: none"> • View: The user or member of a user group can only view the Extension but cannot make changes. • Modify: The user or member of a user group can make changes to this Extension. • Manage: The user or member of a user group can share, edit, or delete this Extension.
--------	---

Application Profile Permissions

Application profile permissions define certain activities that a user can perform with the application profile.

From the *Share Popup* (see [Understand ACLs](#)) for an application profile, the application owner (referred to as *owner*) of the can control permissions for an application profile:

- **Owner:**
 - The author who created an application or application profile is the *owner*, and by default, manages all permissions for this application.
 - The owner must explicitly assign access or deploy permissions to any user, admin, group, or sub-tenant. See [Application Profiles](#) for additional context.



By default, the tenant admin does not have any permission to view/modify/manage/deploy an application profile created by any user within this admin's tenant.

The owner must explicitly assign share or deploy permissions to the admin.

Only admins with appropriate permissions can access permitted applications or application profiles.

- **User:** The owner must explicitly assign access or deploy permissions. Only permitted users can access applications or application profiles.



By default, only the application profile owner can assign permissions for any user, admin, group, or tenant.

The following table describes the application profile permissions options.

Permission	Description
Access	Controls the activities that users or members of user groups can perform for this application profile. <ul style="list-style-type: none"> • View: The user or member of a group/tenant can only view this application profile but cannot modify, share, or delete it. • Modify: The user or member of a group/tenant can edit this application profile, but cannot share or delete it. • Manage: The user or member of a group/tenant can view, edit, share, and delete this application profile.
Deploy	Allows a user or member of a user group to benchmark and deploy this application profile. Without the app profile being shared with a user, the user cannot promote or migrate deployments as he does not own that app profile.

From the Publish option for an application profile, a tenant administrator can control the [permissions](#) for an application profile when publishing it to a marketplace as described in the following table. These permissions control access to the application profile after it is imported from the marketplace by a subscribing user. The following table describes these permission options.

Permission	Description
Imported App Permissions	Permissions for the imported application profile. <ul style="list-style-type: none"> • None: A subscribing user with appropriate privileges user can benchmark and deploy this application profile • View: A subscribing user can view application profile details, and, with appropriate privileges, can benchmark and deploy this application profile • Modify: A subscribing user can edit application profile details, and, with appropriate privileges, can benchmark and deploy this application profile
Can be shared	Allows subscribing user to share this application profile with other users.

Repository Permissions

Repository permissions define certain activities that users can perform with repositories. You can control the permissions for a repository as described in the *Share Popup* (see [Understand ACLs](#)). The following table describes the permission options.

Permission	Description
------------	-------------

View	The user, members of a user group, or tenant can only see this repository but cannot modify, share, or delete it.
Modify	The user, members of a user group, or tenant can edit this repository.
Manage	The user, members of a user group, or tenant can edit or delete this repository.

Each tenant and users within a tenant can only view shared repositories specific to their tenant (or as permitted by their admin). See [Artifact Repository](#) for additional context.

Service Permissions

Service permissions define certain activities that users can perform with custom services. You can control the permissions for a custom service as described in the *Share Popup* (see [Understand ACLs](#)). The following table describes the permission options.

Permission	Description
View	The user, members of a user group, or tenant can see this service but cannot modify, share, or delete it.
Modify	The user, members of a user group, or tenant can edit this service.
Manage	The user, members of a user group, or tenant can edit or delete this service.

Each tenant and users within a tenant can only view services specific to their tenant (or as permitted by their admin). See [Topology Modeler > Supported OOB Services](#) or [Services](#) for additional context.

Actions Library Permissions

Custom actions permissions define certain actions that users can perform. You can control the permissions for a custom action. The following table describes the permission options.

Permission	Description
View	The user or members of a user group can view this custom action but cannot make changes to, share, or delete the custom action. Users who only have <i>View</i> permissions on these actions cannot toggle the Enable (default) or Disable action in the Actions Library page.
Modify	The user or members of a user group can edit this custom action and toggle the Enable (default) or Disable action in the Actions Library page but cannot share or delete it.
Manage	The user or members of a user group can edit this custom action and toggle the Enable (default) or Disable action in the Actions Library page, share it, and delete it.



If you create a custom action and share it, be aware that the permissions for the application profile to which this action is attached must also be in the correct share state for shared users to run this action. You must either create the application profile or share the application profile with these users and assign *modify* or *manage permissions*.

Each tenant and users within a tenant can only view/modify custom actions specific to their tenant (or as permitted by their admin). See [Actions Library](#) for additional context.

Image Permissions

The Share popup lets you assign one of the following permissions to share an image as described in *the Share Popup* (see [Understand ACLs](#)). The following table describes the permission options.

Permission	Description
View	The user, members of a user group, or tenant can see this image but cannot modify, share, or delete it.
Modify	The user, members of a user group, or tenant can edit this image.
Manage	The user, members of a user group, or tenant can edit or delete this image.

Each tenant and users within a tenant can only view shared images specific to their tenant (or as permitted by their admin).

Only permitted users can add images. See [Manage Images](#) or [Image Permissions](#) for additional context.

Temporary Permission to Launch an Image



The Grant and Revoke Image Permission option appears for OpenStack and Cisco clouds only.

The Grant and Revoke Image Permission option in the Add Cloud Mapping window lets you set up temporary permission to allow any user to launch the image in an OpenStack or Cisco cloud. To set up this permission, check the **Grant and Revoke Image Permission** box, and then choose the cloud account that owns this image from the **Image Owner Cloud Account** dropdown menu that appears.

The following table identifies the permission nuances for each resource and their associated API settings

Resource	Permission Can Be Assigned To	Tenant Owner Permission	API objectType Enumeration	API permsList Enumeration
Application profiles	<ul style="list-style-type: none"> Tenant co-admins Users within a tenant 	Always have this permission	APP	CREATE_APP
Global, aging and scaling policies			POLICY	CREATE_POLICY
Deployment environments			DEPLOYMENT_ENVIRONMENT	CREATE_DEPLOYMENT_ENVIRONMENT
Application profile templates	Tenant owners		APP_PROFILE	CREATE_APP_PROFILE
Cloud groups		Without this permission (even for a cloud group assigned by their parent tenant), sub-tenants cannot: <ul style="list-style-type: none"> Create new cloud groups Add new cloud regions to existing cloud groups Configure an existing cloud region different from their parent tenant 	CLOUD	CREATE_CLOUD
Cloud accounts		Without this permission (even for a cloud account assigned by their parent tenant), sub-tenants cannot create new cloud accounts	CLOUD_ACCOUNT	CREATE_CLOUD_ACCOUNT

Projects are only displayed in the Project Owner's dashboard. Even if other users are added to a project, the project is only displayed in the users dashboard after the project is **published**.

Users can perform the functions that the following table describes based on assigned privileges:

Permission	Description
View	The user or members of a user group can only view this resource.
Modify	The user or members of a user group can Edit phases.
Manage	The user or members of a user group can edit, turn it on or off, share, and delete this resource.

All applications are apart of the project:

- The application is **not shared** with a user – The User cannot see the application listed when clicking the **Add Deployment** link.
- A user does not have **Deploy** privilege for the application – The **Add Deployment** link is disabled.

All deployment environments are part of a project:

- A user does not have **Deploy To** privilege – The **Add Deployment** link is disabled.
- A user's deployment environment privileges determine access, as described in the following table

Deployment Environment Privilege	Description
None	The Add Deployment link is disabled.
Access	Running deployments are not visible.
Manage	<ul style="list-style-type: none"> Running deployments are visible Cannot perform any job action
Manage, Promote from	<ul style="list-style-type: none"> Running deployments are visible Perform any job action <i>except</i> the Promote action
Manage, Promote from, Deploy to	<ul style="list-style-type: none"> Running deployments are visible Perform any job action

See [Project and Phase Management](#) for additional context.

Understand ACLs

Understand ACLs

- [Overview](#)
- [ACL-Managed Resources](#)
- [Default Permissions for ACL Resources](#)
- [ACL Manage Permissions](#)
- [Permission Categories](#)
- [Default ACL Resource Permissions](#)
- [UI and API Differences](#)
- [UI Configuration](#)

Access Control Lists (ACLs) allow you to modify/view permissions for an API resource. Resources are identified using a unique ID and corresponding resource name. Not all resources are supported by the ACL function. See the *ACL-Managed Resources* section below for the list of supported resources.

Any user with administration permissions (*perms*) on a resource can view/modify the ACL for that resource using the *ACL Management APIs* APIs:

- [View ACL Resource Details](#)
- [Update ACL Resource Details](#)

The following table identifies the resources that are supported by the ACL function along with the corresponding pages that provide additional information for the resource. This information is identical to the *resourceName* attribute used by the Workload Manager APIs.

Enumeration	Description
POLICY	See <i>Policy Management > Scaling Policies</i> or <i>Aging Policies</i>
ACTION_POLICY	See <i>Policy Management > Action Policies</i>
DEPLOYMENT_ENVIRONMENT	See <i>Deployment Environment</i>
APPLICATION	See <i>Application Profile</i> <ul style="list-style-type: none"> • Model an Application • Model an Application by Importing the Profile
REPOSITORY	See <i>Share Artifact Repositories</i>
CLOUD_ACCOUNT	See <i>Configure Cloud(s)</i>
SYSTEM_TAG	See <i>System Tags</i>
SECURITY_PROFILE	See <i>Security and Firewall Rules</i>
SERVICE	See <i>Manage Services</i>
CUSTOM_ACTION	See <i>Policies > Custom Actions</i>
PROJECT	See <i>Projects</i>
IMAGE	See <i>Manage Images</i>
DISTRIBUTED_JOB	See <i>Deployment Environments > Sharing Deployments</i>
EXTENSION	See <i>Extensions</i>
ACI_EXTENSION	See <i>ACI Extensions</i>
SERVICE_NOW_EXTENSION	See <i>ServiceNow Extensions</i> (Effective CloudCenter Legacy 4.8.2)
ACTION	See <i>Actions Library</i> (Effective CloudCenter Legacy 4.8)
VIRTUAL_MACHINE	See <i>VM Management</i> (Effective CloudCenter Legacy 4.8)
AGING_POLICY	See <i>Policies</i> (Effective CloudCenter Legacy 4.8.2)
SUSPENSION_POLICY	

Permissions are tightly controlled by Workload Manager and not all permissions are applicable to all resources. You will receive validation errors in the following cases:

- When you apply a permission that is not applicable to a particular resource.

For example, `move_in` and `move_out` are only applicable to deployment environments. If you apply either of these two strings to any other resource, you will receive a validation error.

- When you apply a random string that is not listed in the `perms` array. For example, if you assign your own permission value like `readwrite`, you will receive a validation error.
- Imported VMs do not have any default permissions.

As this information is identical to the `perms` attribute used by the Workload Manager APIs, the same information is included here.

<i>resourceName</i> Permissions Enumeration	read write create delete administration	execute	move_in move_out access approve authorize	manage	notify
POLICY	Yes	Yes			
ACTION_POLICY	Yes				
PUBLISHED_APP	Yes				
DEPLOYMENT_ENVIRONMENT	Yes	Yes	Yes	Yes	
APPLICATION	Yes	Yes			
REPOSITORY	Yes				
CLOUD_ACCOUNT	Yes				
SYSTEM_TAG	Yes				
SECURITY_PROFILE	Yes				
SERVICE	Yes				
CUSTOM_ACTION	Yes			Yes	
PROJECT	Yes				Yes
IMAGE	Yes				
MANAGE_EXPORT	Yes	Yes			
MANAGE_IMPORT	Yes	Yes			

Permissions are divided into the categories that the following table describes:

Permission Category (<i>id</i> and <i>perms</i>)	Description
Users	<ul style="list-style-type: none"> • Users can be a tenant admin, promoted admin, co-admin, root admin, or a standard user. • Each user is identified by the User ID, and permissions are granted to the User ID for the specified resource. • Set the <code>perms</code> array with the specific permissions (enumerations) to be provided for this user.
User Groups	<ul style="list-style-type: none"> • A collection of one or more users. If a group consists of 10 users, all 10 users will be assigned permissions based on the group permissions. • Each group is identified by the Group ID, and permissions are granted to the Group ID for the specified resource. • If you add future users, the added users receives the same permission that the other 10 users already have by virtue of being a member of this group. If you delete a user from this group, the deleted user no longer has permissions for this group.
Tenant Users	<ul style="list-style-type: none"> • All users belonging to this tenant will be assigned permission. • Each tenant is identified by the Tenant ID, and permissions are granted to the Tenant ID for the specified resource.
Tenant & Sub-Tenants	<ul style="list-style-type: none"> • Extends tenant level permissions to all sub-tenants at any level below (down the hierarchy) this tenant-level and all users within each of those sub-tenants as well. • Each tenant is identified by the Tenant ID, and permissions are granted to the Tenant ID for the specified resource.

The following default permissions are automatically granted to ACL resources after each resource is created.

- User permissions are granted to the user who created the resource.
- Tenant permissions are granted to:
 - All users of the tenant to which the logged-in user belongs.

- All users in sub-tenant hierarchy starting at tenant of the user who created the resource.



If not specified for Vendor and Tenant then default permissions **are not available at that level**.

The following table describes ACL resource permissions.

Permissions	User	Vendor	Tenant
POLICY	<ul style="list-style-type: none"> • read • write • create • delete • administration • execute 		
ACTION_POLICY	<ul style="list-style-type: none"> • read • write • create • delete • administration 		
PUBLISHED_APP	<ul style="list-style-type: none"> • read • write • create • delete • administration 	READ	
DEPLOYMENT_ENVIRONMENT	<ul style="list-style-type: none"> • read • write • create • delete • administration • execute • move_in • move_out • approve • authorize • manage 		
APPLICATION	<ul style="list-style-type: none"> • read • write • create • delete • administration • execute 		
REPOSITORY	<ul style="list-style-type: none"> • read • write • create • delete • administration 		
CLOUD_ACCOUNT	<ul style="list-style-type: none"> • read • write • create • delete • administration 		
SYSTEM_TAG	<ul style="list-style-type: none"> • read • write • create • delete • administration 	READ	

SECURITY_PROFILE	<ul style="list-style-type: none"> • read • write • create • delete • administration 	READ	
SERVICE	<ul style="list-style-type: none"> • read • write • create • delete • administration 	READ	READ
CUSTOM_ACTION	<ul style="list-style-type: none"> • read • write • create • delete • administration 		
PROJECT	<ul style="list-style-type: none"> • read • write • create • delete • administration • notify 		
IMAGE	<ul style="list-style-type: none"> • read • write • create • delete • administration 		
NODE_STATUS_DEFINITION	<ul style="list-style-type: none"> • read • write • create • delete • administration • execute 		

ACL Configuration differences between the UI and API:

- **UI** – If you have a complicated hierarchy with multiple permission combinations in a tenant hierarchy, then the UI only displays permission for the current level. Permissions for parent and child tenants will not be visible to the logged in user.
- **API** – API users can view or modify permissions for all levels, regardless of this user's level in the tenant hierarchy. Only prerequisite is that the logged in user has administration *perms* on this resource.

If you are the tenant owner, you can provide any permission to the sub-tenant organization and all its users at the same time.

When providing access to Tenant and Sub-Tenant users, access the Share popup for the required service (Workload Manager UI > **Admin** > **Services** > *MyService* > **Share** dropdown), click the **Tenants** tab in the popup, and check the **My Tenants & Sub-Tenants** check box to provide access to the entire hierarchy.

You also have the option to select just one tenant (if you want to give just one tenant, but not their sub-tenants, and provide access to just that tenant.

Share Users (3) Groups (1) **Tenants (2)**

This resource will be shared with all users in the shared Tenant's hierarchy.

Invite tenants to share this resource

Tenants	Access
<input checked="" type="checkbox"/> My Tenant & Sub-Tenants	View
Marketing	Manage View Modify Manage ✓

Save Cancel

Getting Started

Getting Started with Workload Manager

- [Workload Manager Overview](#)
- [Architecture](#)
- [Dashboard and Menus](#)
- [Next Steps for Administrators](#)
- [Next Steps for Standard Users](#)

Workload Manager Overview

Overview

Workload Manager lets end users model multi-tier applications, deploy those application in various public and private clouds, and manage those applications throughout their lifecycle. You can also import and manage VMs deployed outside of Workload Manager.

Workload Manager gives administrators fine grained control over which features and cloud resources are available to users and how those cloud resources are consumed.

First, become familiar with Workload Manager's capabilities by reading these user guide sections:

- [Architecture](#)
- [Dashboard and Menus](#)
- [What is Supported?](#)


Then, follow [Next Steps for Administrators](#) or [Next Steps for Standard Users](#) depending on your role.

Architecture

Architecture

- [Basic Install Architecture](#)
- [Full Install Architecture](#)
- [Port Requirements](#)
 - [Without Cloud Remote](#)
 - [With Cloud Remote](#)

After installing Workload Manager from the Suite Admin, if your CloudCenter Suite Kubernetes cluster can receive connections from public internet addresses, you have what you need to use all of Workload Manager's core features with VM-based public clouds. This includes deploying and managing workloads, and importing and managing VMs launched outside of Workload Manager. Once you deploy a workload in a public cloud, or import VMs, you will have deployed all of the component needed to create the basic install architecture as shown in the figure below.



Basic Install Architecture – VM-Based Target Cloud

The basic install architecture consists of four components

- Manager
- Agent
- Cisco-hosted bundle store
- Cisco-hosted package store

The **manager** component is the main component of Workload Manager. It consists of services running within pods in the CloudCenter Suite cluster. Some of these services are common framework services used by all modules, some are specific to Workload Manager alone, and some are shared between Workload Manager and Cost Optimizer.

One function of the manager is communicating 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.

A second function of the manager is to communicate directly with the VMs running your workloads. This is only possible when those VMs have the second Workload Manager component installed: the **agent**. VMs with the agent installed are called **worker VMs**.

The agent gives you additional control of your VMs by allowing you to execute commands or run scripts from within the VM. These can be scripts that run at certain points in the VM's lifecycle, such as a script to install and launch a service at startup, or they can be actions that are executed on-demand via the Workload Manager UI. See [Actions Library Overview](#) for more info.


Prior to the agent being downloaded and started on a VM, a set of scripts and installer packages must be installed on the VM first. This set of prerequisite software is collectively known as the **worker**. Once installation of the worker is complete, the first time the VM is started, a script in the worker downloads and starts the agent executable file.

The worker can be installed on a VM in one of three ways:

- When you launch a VM-based workload to a public cloud using Workload Manager, installation of the worker happens automatically in a process called dynamic bootstrapping. When Workload Manager issues the API call to the cloud endpoint to launch the VM, it passes as user data a bootstrap script that downloads and installs prerequisite software for the agent, and then downloads and starts the agent executable.
- If your target VM-based cloud does not support dynamic bootstrapping, or if you prefer not to use dynamic bootstrapping, the alternative is to use "pre-bootstrapped" images for your VM-based services. These are OS images with worker software pre-installed. See [Management Agent \(Worker\)](#) for more details.
- You can also install the agent on VMs in your cloud that were not launched through the Workload Manager in a process called VM import. See [Virtual Machine Management](#) for more details.

The last two components of the basic install architecture are the Cisco-hosted software repositories: the **bundle store** at <http://cdn.cliqr.com/cloudcenter-<version>/bundle> (do not add a slash at the end of the URL) and the **package store** at <http://repo.cliqrtech.com>. The bundle store contains the scripts used to install the worker software, the latest version of the **agent**, and scripts that run within the worker VM for launching and controlling the service that should run in that VM. The package store contains the install packages for the worker software and the install packages for the Workload Manager OOB services, as well as the public cloud instance types, storage types and image mapping.

For Kubernetes target clouds, there are no worker VMs and all control of the container-based workloads is through the Kubernetes API. The basic install architecture relative to Kubernetes target clouds is summarized in the figure below.



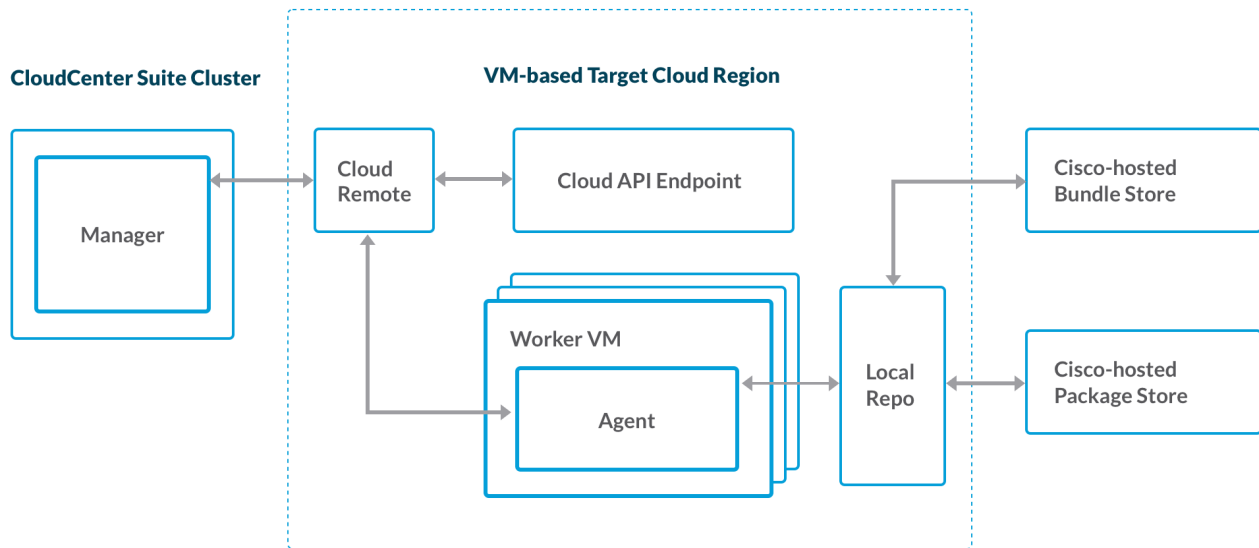
Basic Install Architecture – Kubernetes Target Cloud

Since your workloads are deployed in Kubernetes containers, there are no workers and no need to access the Cisco-hosted bundle store and package store. Instead, your target Kubernetes cloud must allow access to the public Docker hub for downloading the public Docker image files referenced in your containerized workloads.

The basic install architecture has a key limitation: it assumes that the manager and all of 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 Workload Manager. For VM-based clouds you will need to install two additional components:

- [Cloud Remote](#)
- [Local Repo Appliance](#)

The full install architecture for VM-based cloud regions is shown in the figure below. Be aware that if you use Cloud Remote, you only need access in one direction to/from the CloudCenter Suite as Cloud Remote handles the communication in the other direction.



Full Install Architecture – VM-Based Target Cloud

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 which manages a collection of containerized services. As such, it can be deployed as a single VM and later scaled to a cluster of VMs.

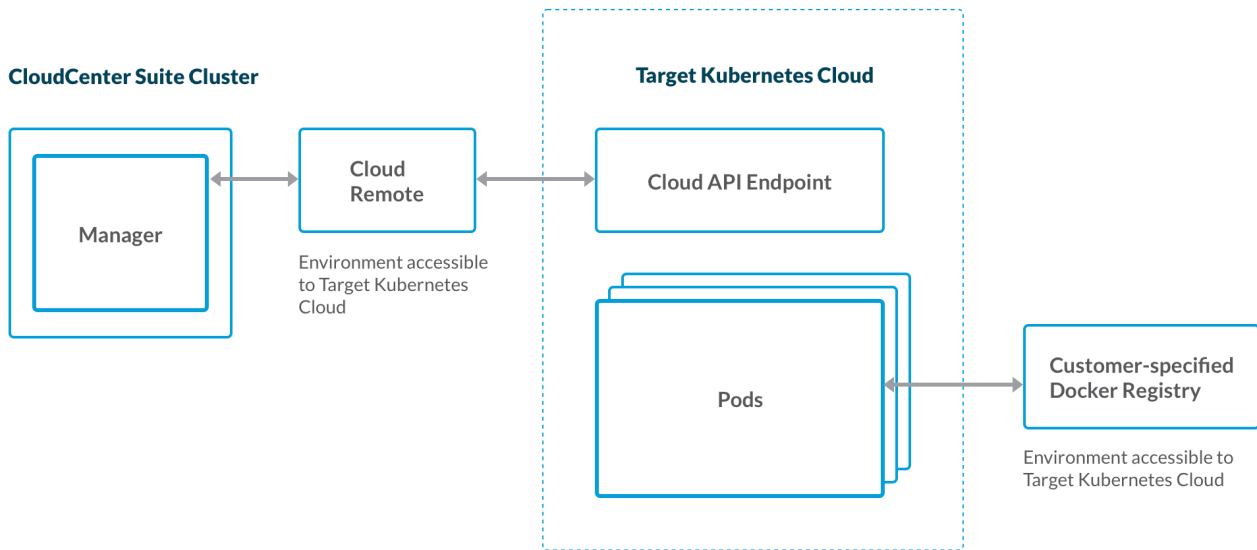
For VM-based cloud regions, Cloud Remote performs the following functions:

- Proxies communications between the manager and the cloud API endpoint (also used by Cost Optimizer).
- Executes external scripts on the workload VMs (even those without the management agent) to support external lifecycle actions.
- Proxies communications for user SSH/RDP sessions with worker VMs.
- Proxies communication between the manager and the worker VMs to support internal lifecycle actions, internal on demand actions, and reporting of workload status.

Note: If the manager component cannot accept inbound connections from public addresses, you will need to install Cloud Remote in all VM-based target regions that are not within the same network as your manager.

The **local repo appliance** is also delivered as a virtual appliance that you import to your target VM-based cloud region. The local repo appliance can be configured to support both a local bundle store and a local package store. The local repo appliance must have periodic internet access in order to sync with the master bundle store and package store hosted by Cisco. Cisco also provides scripts for creating your local repo appliance on the Linux

The full install architecture for Kubernetes target clouds is shown in the following figure. Be aware that if you use Cloud Remote, you only need access in one direction to/from the CloudCenter Suite as Cloud Remote handles the communication in the other direction.



Full Install Architecture – Kubernetes Target Cloud

For Kubernetes target clouds, you would install the Cloud Remote appliance in an environment in the same network as the target Kubernetes cloud. In this case, Cloud Remote perform two functions:

- Proxies the API calls from the manager to the cloud API endpoint.
- Executes external scripts on the workload pods to support external lifecycle actions.

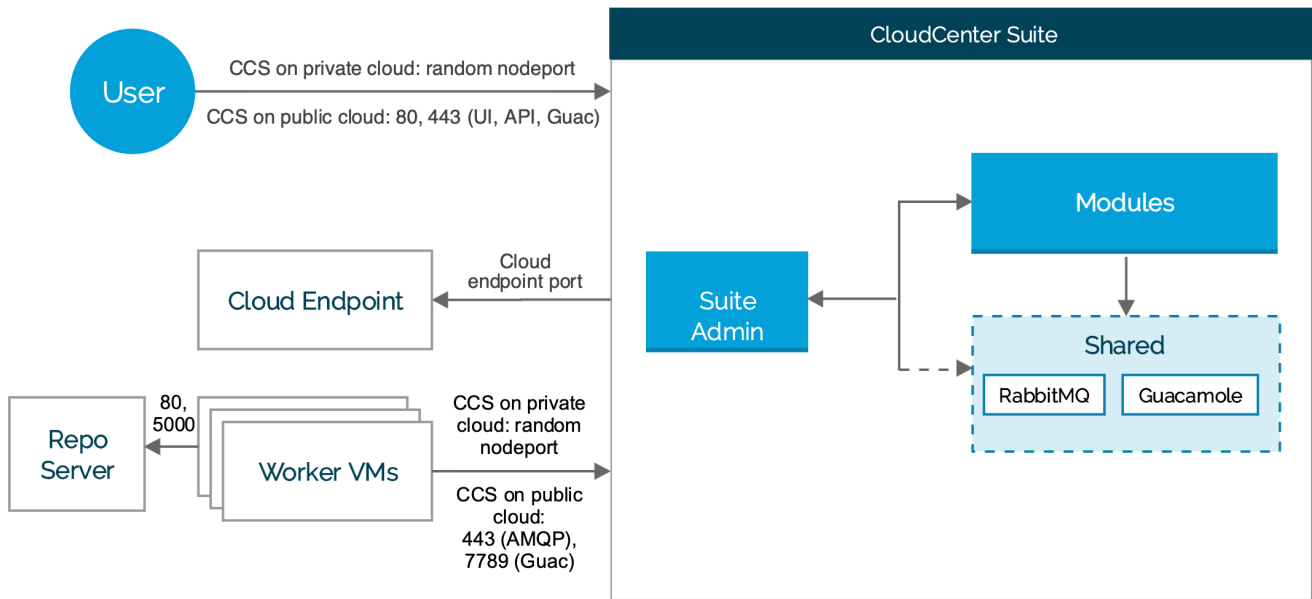
For Kubernetes clouds that do not have outbound internet access, you will also need to install your own Docker registry on a VM in the same network as your target Kubernetes cloud. You will need to populate that registry with all of the public and private Docker images used by the containerized services in your workload. (See Docker.io user documentation for more on setting up your own Docker registry).

The following images identify the ports that must be open for Workload Manager.

Without Cloud Remote

The following image identifies the ports that must be open for Workload Manager.


CCS with Full Cloud Connectivity (Workload Manager)



With Cloud Remote

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

CCS with Cloud Remote for Workload Manager



* 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 to <https://kubernetes.io/docs/concepts/services-networking/service/> for additional context.

Dashboard and Menus

Dashboard and Menus

- [Main Dashboard](#)
- [Main Menu](#)
- [Administrator Menu](#)

When you first login to Workload Manager, you are directed to the main dashboard with the main menu tabs on the left side of the screen. The dashboard and menu tabs you see will vary depending upon the groups and roles your user account is associated with. See [OOB Groups, Roles, and Permissions](#) for more context.

The main dashboard summarizes information about your deployments and imported VMs. In addition, if your user account includes the WM_ADMIN role, the dashboard will let you display information about the deployments and imported VMs for all users in your tenant.

The main dashboard has five dashlets:

- **Summary:** This shows total cloud costs and VM hours for the past 30 days and the count of currently running VMs. If you are logged in as an administrator, there is dropdown in the upper left of the dashlet that lets you toggle between data for just your own workloads or all workloads managed by users in your tenant. Toggle buttons in the upper right let you aggregate the data by cloud or application. Links in this dashlet point to more detailed cost and usage reports. See [Reports Overview](#) for more information on these reports.
- **My Plan Usage:** This displays the name of the usage plan assigned to your user id. Usage plans restrict how much a user can spend on cloud resources in terms of a dollar amount or VM hours. Workload Manager includes a single out of box unlimited usage plan. As an administrator, you can create additional usage plans. (see [Usage Plans and Fees](#) for more details).
- **Cloud Status:** Once you as an administrator have set up cloud regions in Workload Manager (see [Clouds](#) for details), the status of those cloud regions will be displayed here.
- **My Recent Deployments:** Once you deploy applications, they will be listed here.
- **Favorite deployments:** You can mark a deployment as a "favorite" via the deployments list page. See [Virtual Machine Management](#) for more details.
- **Notifications:** When any important Workload Manager generated messages are produced, they are listed here.

The main menu has the following tabs:

- **Dashboard:** This shows the main dashboard as described above.
- **App Profiles:** This shows the Application Profiles page. From here, you can perform the following actions:
 - Create a multi-tiered application profile using the drag and drop topology modeler
 - Deploy an application to a particular deployment environment
 - Share your application profile with other users
 - Export and import application profiles
 - Jump to a list of all deployments associated with an application profile
 - Benchmark an application
- **Deployments:** This shows the Deployments List page. From here, you can perform the following actions:
 - View summary information on each deployment (run time, costs)
 - Perform on-demand actions on a deployment (suspend, resume, terminate)
 - Drill down to details of a deployment and perform manual scaling of individual tiers and change which policies are applied to the deployment.
- **Virtual Machines:** This shows the Virtual Machines List page. This page has tabs for managed VMs and unmanaged VMs:
 - Unmanaged VMs are VMs launched using one of the cloud accounts associated with you as a user, but which are not running the management agent. These are typically VMs you launched outside of Workload Manager. You can select any of these VMs and import them. When you do this, the management agent will be installed on the selected VMs and they will become managed VMs.
 - Managed VMs are VMs that are running the management agent. They include VMs that were formerly unmanaged and then imported, and VMs belonging to applications launched through Workload Manager. You can perform various OOB on-demand actions on these VMs (suspend, resume, terminate) as well as custom on-demand actions.
- **Environments** (visible only to users with at least one of these roles: WM_ADMIN, WM_ENVIRONMENT_MANAGER): Displays the Deployment Environments list page. When you deploy an application, you do so by first selecting a deployment environment. The deployment environment describes what cloud regions an application can be deployed in, with which possible cloud accounts, and with which possible scaling, aging, suspension, and security policies. Administrators create, edit and delete deployment environments and share them with selected users.
- **Services** (visible only to users with at least one of these roles: WM_ADMIN, WM_SERVICE_MANAGER): Displays the Services List page. Services are the building blocks of multi-tier applications. Workload Manager comes with many OOB services, including VM-based, container-based, and external services. You can also create your own custom services here.
- **Benchmarks:** Displays the list of benchmark test result reports for the application benchmark jobs that you specified and ran from the Application Profiles page.
- **Repositories:** Displays the Repositories List page. Here you can add a repository location for storing scripts used by your application during its various lifecycle phases.
- **Policies** (visible only to users with at least one of these roles: WM_ADMIN, WM_ENVIRONMENT_MANAGER): Displays the lists of policies organized by tabs. From here you can create, edit, share and delete policies. The five types of policies are as follows:
 - **Event:** Let's you define commands that are executed when a cloud region or a deployment changes state.
 - **Scaling:** Let's you define the conditions when a VM-based tier of a deployment will automatically scale up or scale down.
 - **Aging:** Let's you define a time limit or cost limit for a deployment, after which the deployment will be automatically terminated.
 - **Suspension:** Let's you define the time periods when the VM-based tiers of a deployment will be automatically suspended.
 - **Security:** Let's you define firewall rules that can be applied to a deployment as a whole or to particular tiers of a deployment.

- **Images** (visible only to users with at least one of these roles: WM_ADMIN, WM_IMAGE_MANAGER): Displays the Images List page. This is a list of OS base images that are mapped to the actual cloud provider images for each cloud region. Workload Manager includes several OOB images. For many public cloud providers, Workload Manager automatically applies the appropriate mapping of these OOB logical images to cloud provider images. You must set up your own image mapping for private clouds. You can also add your own custom images.
 - **Actions Library** (visible only to users with at least one of these roles: WM_ADMIN, WM_DEV_OPS): Displays The Actions Library page which is where you can define and share custom actions of two types:
 - On-demand Actions. These action can be commands, scripts or invocations of web services. They are made available for use via dropdown menus in the pages where you can view deployment and VM status.
 - Lifecycle Actions. These action can be commands or scripts. They are made available for use via dropdown menus in the pages where you specify your application profiles and service definitions.
 - **Projects** (visible only to users with at least one of these roles: WM_ADMIN, WM_PROJECT_MANAGER): Displays the list of projects created by you or shared with you. A project lets you manage the lifecycle of your application by specifying different phases. In each phase, the application will be deployed using the deployment environment and usage plan specified for the phase.
- Admin** (visible only to users with the WM_ADMIN role): Displays the administrator menu on the left side of the screen and the Clouds Page on the right side for the screen (see below).

The administrator menu has the following tabs:

- **Main Menu:** Brings you back to the main menu on the left side of the screen and the main dashboard on the right side of the screen.
- **Clouds.** This shows the clouds that have been configured for your tenant. From here, you can create, edit and delete clouds. A cloud requires at least one cloud account and at least one cloud region.
 - Cloud accounts are the cloud provider user credentials needed to launch instances in a cloud region
 - Cloud regions correspond to geographical regions for public clouds. Some private clouds can have multiple regions, but VMware vCenter clouds and Kubernetes container clouds only have one region. Cloud region parameters for VM-based clouds include the cloud API endpoint address, VM naming and IPAM scripts, storage types, instance types, image mappings, and lifecycle actions.
- **Extensions:** Displays the Extensions List page list. From here you can add or edit one of the two supported extension types: ACI and ServiceNow.
- **All Reports:** Brings you to the Usage Summary report. Using the dropdown next to the report name, you can also view these two other reports: Application Deployments and Running VM History.
- **System Tags:** Displays the System Tags List page. Here you can add system tags, which are text strings, and share them with all subtenants. Once defined, a system tag can be associated with an application profile as a whole, a tier in an application profile, or a deployment environment.
- **Usage Plans:** Displays the Usage Plans List page. A [usage plan](#) determines how much total cloud resources a user can consume. You can add plans that limit consumption based on cost or VM-hours. A usage plan, can refer to a [bundle](#) (see below). Usage plans are assigned to users through the [Suite Admin UI](#).
- **Bundles:** Displays the Bundles List page which is where you can define a limit on cloud consumption based on budgeted cost or VM-hours used. A bundle is not directly assigned to users. Instead, it may be assigned to a usage plan (see above) which is in turn can be assigned to a users.

Next Steps for Administrators

Next Steps for Administrators

As a tenant administrator you need to get your Workload Manager system ready to support users. This involved multiple steps, some of which may or may not apply depending on your circumstances.

- Set up [users and subtenants](#) using the Suite Admin UI. In order to do this you will need to have the suite admin role. If you don't, contact your suite admin to assign that role to you. Ensure that users for your tenant are entered into the suite, they are assigned to your tenant, and they have the appropriate roles (see [OOB Groups, Roles, and Permissions](#)). Create any subtenants as needed (see [Tenant Management](#) and [Manage Tenants](#)).
- Use the [Architecture](#) section as a guide to determine if you want to or need to install any additional components such as [Cloud Remote](#), [local repo appliance](#), or a Docker registry for any of your target clouds. If yes, install and configure those components per the install instructions.
- Create and configure the clouds and associated cloud regions and cloud accounts as explained in [Clouds](#). This may involve defining custom storage types and instance types for your clouds.
- Determine if you want or need [pre-bootstrapped images](#) for any of your VM-based clouds. If yes, create those images using the Cisco provided installer tools, import those images to all of your cloud regions where they are needed, and [map](#) the corresponding logical images to those pre-bootstrapped images.
- For private VM-based clouds, [map](#) Workload Manager's [OOB logical images](#) to the physical images imported to your private clouds. These physical images may be [pre-bootstrapped](#) images that you created and imported.
- Optional. Create aging, suspension, security, and scaling [policies](#) to be applied to user deployments.
- Create [deployment environments](#) for your users to allow them to deploy their applications to certain cloud regions using certain cloud accounts.
- Create additional [usage plans and bundles](#). The root tenant has an unlimited usage plan by default that applies to all users in that tenant. As tenant administrator you can create additional plans and bundles within Workload Manager that limit cloud spend. As a suite administrator, you can then assign a plan to a subtenant or to individual users in your tenant (see [Tenant Management](#) and [Manage Tenants](#)).

Your users will now have everything they need to create applications using the Workload Manager [OOB services](#), and deploy those applications using the cloud regions, cloud accounts, deployment environment, and policies you set up. You may also consider creating custom VM-based or container-based [services](#) and share them with your users.

Next Steps for Standard Users

Next Steps for Standard Users

As a standard user, you will have the ability to model, deploy and manage your applications. To get started, follow these steps.

- [Model an application](#). We recommend starting simple: a single-tier VM-based application consisting of CentOS OOB service, for example.
- [Deploy your application](#).
- [Manage your application deployment](#).
- [Manage individual VMs](#).

Once you are comfortable with the above steps, consider trying these optional steps.

- Specify a [custom repository](#) that contains artifacts for your application.
- [Benchmark your application](#).
- Create a CI/CD workflow for your application using [Projects](#).
- [Import VMs](#) that you launched outside of Workload Manager so that they may be managed within Workload Manager.
- Create custom on-demand and lifecycle [actions](#) that can be applied to your VMs.

Install Conditional Components

Install Conditional Components

- [Installation Overview](#)
- [Resource Requirements](#)
- [Cloud Remote \(Conditional\)](#)
- [Local Repo Appliance \(Conditional\)](#)
- [Worker \(Conditional\)](#)
 - [Management Agent \(Worker\)](#)
 - [Options to Install the Worker](#)
 - [Install Worker on a Linux Image](#)
 - [Install Worker on a Windows Image](#)
- [Local Package Store \(Conditional\)](#)
- [Local Bundle Store \(Conditional\)](#)

Installation Overview

Installation Overview

- [Basic Workload Manager Installation](#)
- [Supplemental Installations](#)
- [Module Update Considerations](#)

The basic Workload Manager installation is initiated through the [Suite Admin](#) and provides all of the functionality needed to configure your environment and let users model, deploy, and manage applications. However, the basic installation is sufficient only in certain environments as summarized in the [Architecture](#) section. Specifically, the following prerequisites must be met for the basic installation to be sufficient:

1. All VM-based cloud regions can access the Cisco-hosted bundle store and package store.
2. The CloudCenter Suite cluster can initiate communication with the cloud API endpoints for all target cloud regions.
3. The Workload Manager agent running on VMs in all of the VM-based target cloud regions can initiate communications with the CloudCenter Suite cluster.
4. All VM-based cloud regions support dynamic bootstrapping: the ability to inject a bootstrap script to a VM dynamically upon launch of that VM.
5. All target Kubernetes clouds have access to the public Docker registry at docker.io.

If any of these conditions is not true, you will need to perform the supplemental installations as explained below.

Since complete functionality with the basic installation has several requirements, plan for the supplemental installations based on the table below.

Condition motivating supplemental installation	Supplemental component	Where deployed or installed	How Obtained	Install Instructions
The Cisco-hosted package store or bundle store are not accessible from your VM-based target cloud region OR You want to reduce latency associated with downloading artifacts from the package store or bundle store	For AWS, AzureRM, and OpenStack: Local Repo Appliance OR For Google: A Linux VM with the bundle store and package store manually installed by you	Deployed as a VM in the target VM-based cloud region.	Local Repo Appliance: Obtain the appliance image as explained in Conditional Component Appliance Images Bundle store and package store installers: Download the appropriate installer script contained in the artifacts.zip file at software.cisco.com	For the Local Repo Appliance: Local Repo Appliance (Conditional) For installing the bundle store manually: Local Bundle Store (Conditional) For installing the package store manually: Local Package Store (Conditional)
The CloudCenter Suite cluster cannot initiate communications with a cloud region API endpoint OR Workload Manager agents in your VM-based cloud region cannot initiate communications with the CloudCenter Suite cluster	Cloud Remote	For VM-based cloud regions: Deployed as one or more VMs in the target cloud. For Kubernetes clouds: Deployed as one or more VMs accessible from the target Kubernetes cloud.	Obtain the appliance image as explained in Conditional Component Appliance Images	See Cloud Remote (Conditional)
Your VM-based cloud region does not support dynamic bootstrapping OR You want to deploy a VM based on your own custom logical image	For CentOS 6: CentOS 6 Worker Image For all other base OSes: A physical image created by you using the Linux or Windows agent installer script	Imported to or shared with your VM-based cloud region	For CentOS 6: Obtain the worker image as explained in Conditional Component Appliance Images For all other base OSes: Download the appropriate installer script contained in the artifacts.zip file at software.cisco.com	For creating your own worker image: see Management Agent (Worker)
Your target Kubernetes cloud cannot initiate connections to the internet	Local Docker Registry containing the public Docker images used by your container-based application tiers	Installed on a VM accessible from the target Kubernetes cluster.	The registry is part of Docker. The Docker images are available from Docker Hub	https://docs.docker.com/registry/

When updating the Workload Manager module, be aware that the update occurs over the course of 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 users are encouraged to keep this potential loss of connectivity in mind before applying Workload Manager updates.

The design of Workload Manager is tightly coupled with the design of Cost Optimizer and Action Orchestrator, therefore it is recommended that you maintain all three modules at the same minor release level.

Resource Requirements

Resource Requirements

- [Overview](#)
- [Cloud Remote Resource Requirements](#)
- [Local Repo Appliance Resource Requirements](#)

Resource requirements for the for Workload Manager's manager component are included in the [resource requirements for the CloudCenter Suite cluster as a whole](#). This section specifies the resource requirements for

- Cloud Remote appliance
- Local repo appliance

[Cloud Remote](#) is required in each cloud region where a direct connection between the CloudCenter Suite cluster and the cloud region API endpoint or between the CloudCenter Suite cluster and the managed VMs is not possible. The Cloud Remote appliance should be deployed using the following resources:

- 2 CPU, 8 GB memory, 30 GB storage

The [local repo appliance](#) should be deployed using the following resources:

- 2 CPU, 8 GB memory, 50 GB storage

Cloud Remote (Conditional)

Cloud Remote

- [Overview](#)
- [Install Cloud Remote](#)
 - [Configure Cloud Remote in a vCenter Region](#)
 - [Configure Cloud Remote in a vCenter Region for a Kubernetes Cloud](#)
 - [Configure Cloud Remote in an OpenStack Region](#)
 - [Configure Cloud Remote in an OpenStack Region for a Kubernetes Cloud](#)
 - [Configure Cloud Remote in an AWS Region](#)
 - [Configure Cloud Remote in an AWS Region for a Kubernetes Cloud](#)
 - [Cloud Remote for AzureRM](#)
 - [Configure Cloud Remote in an AzureRM Region for a Kubernetes Cloud](#)
 - [Configure Cloud Remote in a Google Region](#)
 - [Configure Cloud Remote in a Google Region for a Kubernetes Cloud](#)
- [The Cloud Remote Artifacts](#)
- [Static IP Address Usage](#)
- [Using Non-Conflicting Networks](#)
- [Install Cloud Remote on a Custom CentOS 7 VM](#)
- [Upgrade an Existing Cloud Remote Installation](#)
- [Scaling](#)
- [Custom Port Numbers \(Conditional\)](#)
- [Navigating Cloud Remote through Proxy](#)
 - [Proxy Service on the Cloud Remote Instance](#)
 - [Proxy Service on the CloudCenter Suite Cluster](#)
- [Troubleshooting Cloud Remote Issues](#)

The Cloud Remote component is deployed on a per cloud region basis if communication between the CloudCenter Suite cluster and the target cloud region is restricted. More specifically, it is needed when:

- Communication between the CloudCenter Suite cluster and the API endpoint of your private cloud region is restricted or
- Communication between the CloudCenter Suite cluster and worker VMs in your VM-based cloud region is restricted

When Cloud Remote is used to support communications with a VM-based cloud region, it is installed as a virtual appliance launched in that region. When it is used to support communications with a Kubernetes cloud, it is installed as a virtual appliance in a network accessible from that Kubernetes cloud.

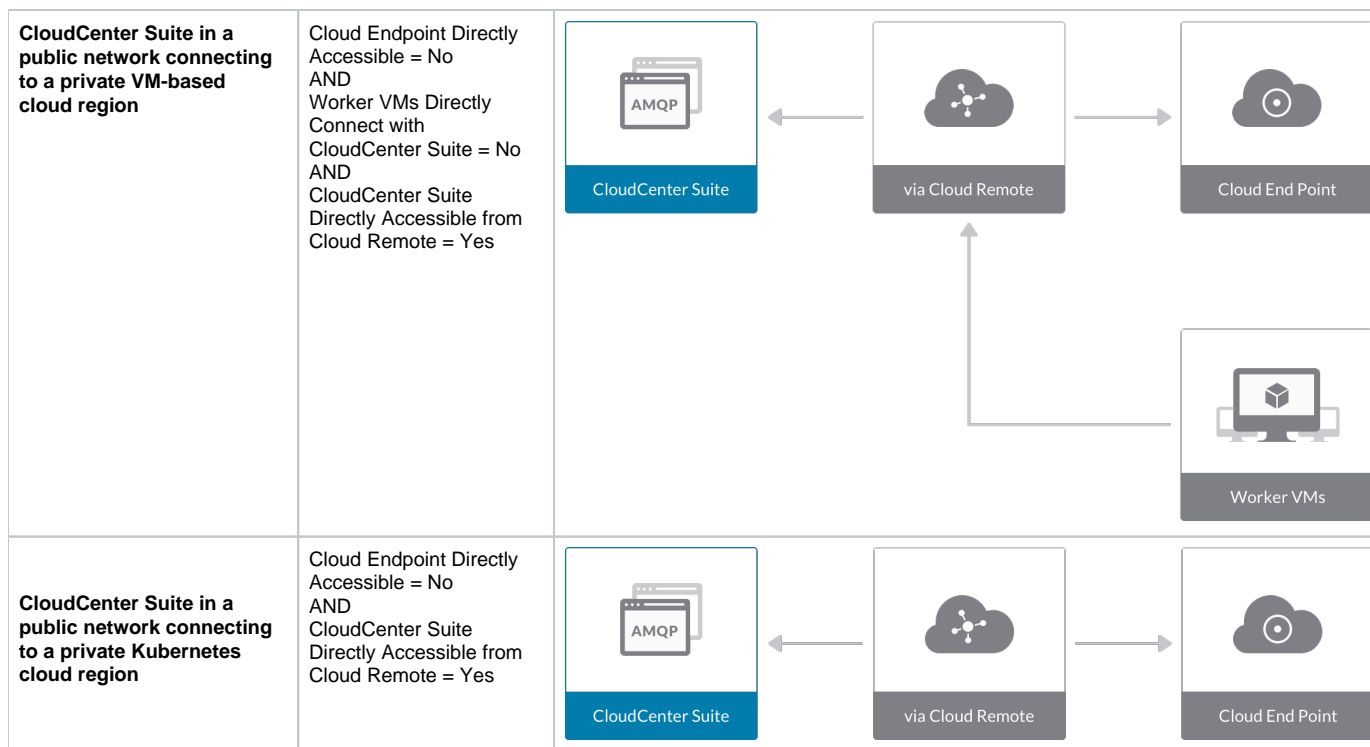
Cloud Remote can run as a single appliance or it can [scale up](#) to multiple appliances working as a single cluster.

Cloud Remote includes the following services running as containers:

- AMQP server for communicating with the CloudCenter Suite cluster and with worker VMs
- Script execution engine for executing external lifecycle action scripts
- Proxy server for communicating with the script execution engine and the cloud API endpoint
- Guacamole server for encapsulating SSH or RDP sessions to worker VMs in a browser window

Some typical network configurations involving Cloud Remote are as follows:

Use Case	Network Connectivity	Network Diagram
CloudCenter Suite in a private network connecting to a public cloud region	Worker VMs Directly Connect with CloudCenter Suite = No AND CloudCenter Suite Directly Accessible from Cloud Remote = No	<p>The diagram illustrates the network connectivity for the use case. It shows four main components: CloudCenter Suite (top left), Cloud End Point (top right), Worker VMs (bottom right), and a path labeled 'via Cloud Remote' (bottom center). A direct arrow labeled 'DIRECT' connects CloudCenter Suite to Cloud End Point. Another arrow connects CloudCenter Suite to the 'via Cloud Remote' path, which then connects to Worker VMs.</p>



The remaining sections describe how to acquire and configure Cloud Remote, and how to scale Cloud Remote.

Cloud Remote is installed as a virtual appliance obtained from Cisco. The procedure to obtain, launch and configure Cloud Remote depends on:

- The VM-based cloud in which Cloud Remote will be deployed.
and
- The overall networking constraints of the CloudCenter Suite cluster and the target cloud region.

Prior to installing Cloud Remote, make sure you have already added the cloud to CloudCenter Suite, and if a multi-region cloud, you added the first region. Then, use one of the following procedures corresponding to where Cloud Remote will be deployed and whether it will be used to support VM-based workloads in that cloud region or Kubernetes container workloads in a Kubernetes cloud hosted in that 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

1. From your local computer, download the Cloud Remote appliance OVA from software.cisco.com.
2. 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.
3. 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.
4. 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.
5. Complete the fields for Name and location, Host / Cluster, Resource Pool, Storage, and Disk Format appropriate for your environment.
6. 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.
7. 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.
8. Confirm your settings and click Finish to launch the VM.
9. 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.
10. Once the first instance of the appliance has been launched, use the vSphere client 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) > <i>Custom Port Numbers (Conditional)</i>). 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) > <i>Custom Port Numbers (Conditional)</i>).
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) > <i>Custom Port Numbers (Conditional)</i>).
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.

Region Connectivity Running [Download Configuration](#) [Configure Region](#)

Clicking Download Configuration 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.

Region Connectivity Enabling... [Download Configuration](#) [Copy Encryption Key](#) [Edit Connectivity](#)


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.



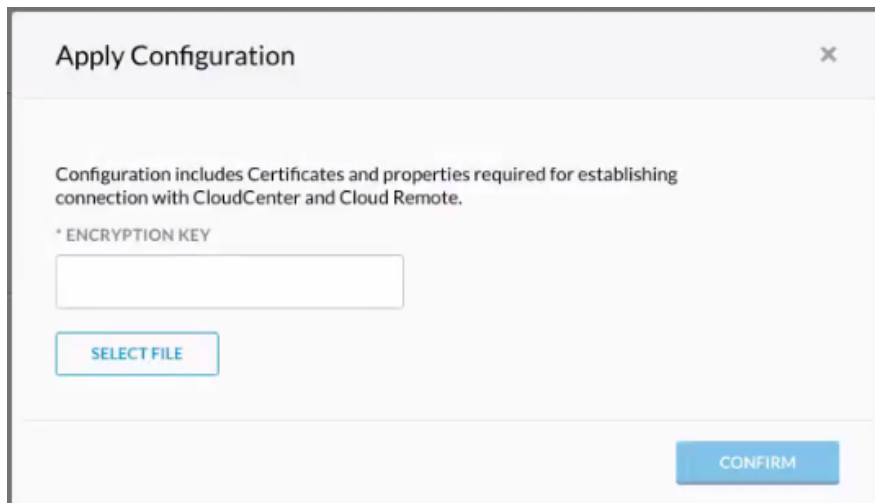
If you change the connectivity settings in the CloudCenter Suite UI and need to again download the zip file, a new encryption key is automatically created and can be copied to the clipboard by clicking the **Copy Encryption Key** link again.

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

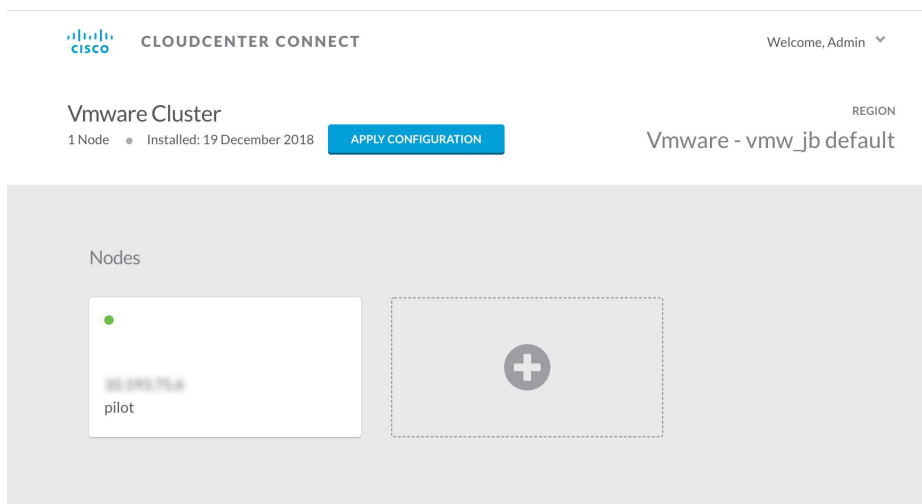
1. Open another browser tab and login to https://<Cloud Remote_ip> with the default credentials: admin/cisco.
2. You will immediately be required to change your password. Do so now.
3. You are now brought to the Cloud Remote home page as shown in the figure below.



4. 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.



5. Paste the encryption key that was copied to the clipboard into the **Encryption Key** field in the dialog box.
6. Click **Select File** and browse to the artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
7. Click **Confirm**.
8. 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).

Region Connectivity Running[Download Configuration](#) [Configure Region](#)

Cloud endpoint accessible from Cloud Center Manager	No
Cloud Center Manager AMQP reachable from worker VM's	No
Cloud Center Manager AMQP accessible from cloud	Yes
Remote AMQP IP	
Worker AMQP IP	192.168.30.16:5671
Blade Name	cloudcenter-blade-vmware-9-0289
Blade Port	8443

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

1. From your local computer, download the Cloud Remote appliance OVA from software.cisco.com.
2. 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.
3. 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.
4. 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.
5. Complete the fields for Name and location, Host / Cluster, Resource Pool, Storage, and Disk Format appropriate for your environment.
6. 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.
7. 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.
8. Confirm your settings and click Finish to launch the VM.
9. 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.
10. Once the first instance of the appliance has been launched, use the vSphere client 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.

Region Connectivity Running Download Configuration [Configure Region](#)

Clicking Download Configuration 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.


Region Connectivity Enabling... [Download Configuration](#) Copy Encryption Key [Edit Connectivity](#)

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.

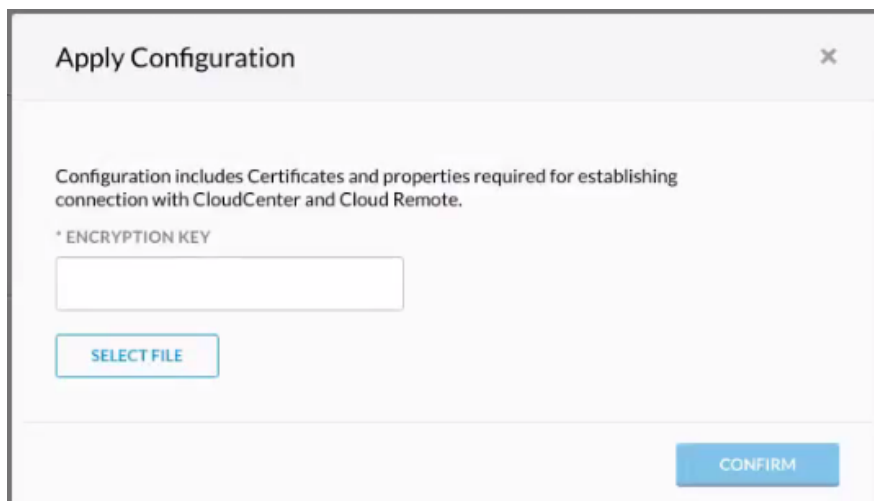
 If you change the connectivity settings in the CloudCenter Suite UI and need to again download the zip file, a new encryption key is automatically created and can be copied to the clipboard by clicking the **Copy Encryption Key** link again.

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

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

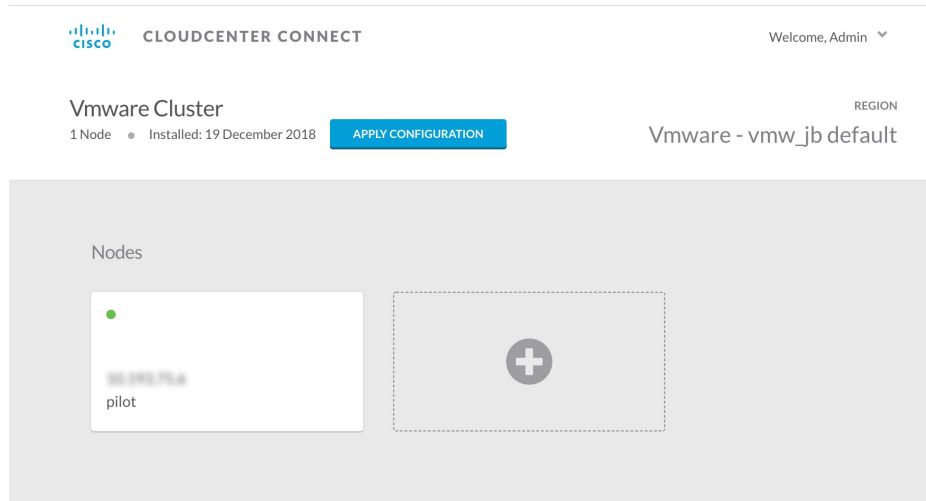


4. Click the **Apply Configuration** button in the page header. This prompts you to select a configuration file and enter the encryption key as shown in the figure below.



5. Paste the encryption key that was copied to the clipboard into the **Encryption Key** field in the dialog box.
6. Click **Select File** and browse to the artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
7. Click **Confirm**.

8. 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).

Region Connectivity	Running	Download Configuration	Configure Region
Cloud endpoint accessible from Cloud Center Manager	No		
Cloud Center Manager AMQP reachable from worker VM's	No		
Cloud Center Manager AMQP accessible from cloud	Yes		
Remote AMQP IP			
Worker AMQP IP	192.168.30.16:5671		
Blade Name	cloudcenter-blade-vmware-9-0289		
Blade Port	8443		

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

Configure Cloud Remote in an OpenStack region as follows.

Download and Launch the Cloud Remote Appliance in OpenStack

1. Download the Cloud Remote appliance qcow2 file from software.cisco.com.
2. Through the OpenStack console, import and launch the Cloud Remote appliance. This process is similar to importing and launching the [Cloud Center Suite installer appliance for OpenStack](#).



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

3. 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.
4. 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.

Setup Cloud Remote Firewall Rules for 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 additional port 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) > <i>Custom Port Numbers (Conditional)</i>).
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) > <i>Custom Port Numbers (Conditional)</i>).
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 Download Configuration 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 to Cloud Remote web UI.

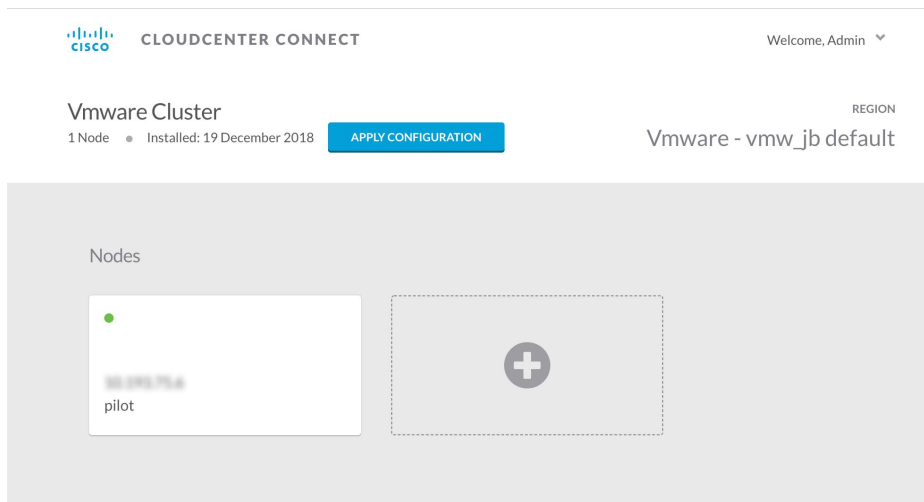
1. Open another browser tab and login to https://<Cloud_Remote_ip> with the default credentials: admin/cisco.
2. You will immediately be required to change your password. Do so now.
3. You are now brought to the Cloud Remote home page as shown in the figure below.



4. 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.



5. Paste the encryption key that was copied to the clipboard into the **Encryption Key** field in the dialog box.
6. Click **Select File** and browse to the artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
7. Click **Confirm**.
8. 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).

Region Connectivity	Running	Download Configuration	Configure Region
Cloud endpoint accessible from Cloud Center Manager	No		
Cloud Center Manager AMQP reachable from worker VM's	No		
Cloud Center Manager AMQP accessible from cloud	Yes		
Remote AMQP IP			
Worker AMQP IP	192.168.30.16:5671		
Blade Name	cloudcenter-blade-vmware-9-0289		
Blade Port	8443		

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

1. Download the Cloud Remote appliance qcow2 file from software.cisco.com.
2. Through the OpenStack console, import and launch the Cloud Remote appliance. This process is similar to importing and launching the [Cloud Center Suite installer appliance for OpenStack](#).



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

3. 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.
4. 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.

Region Connectivity Running Download Configuration [Configure Region](#)

Clicking Download Configuration 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.

Region Connectivity Enabling... Download Configuration Copy Encryption Key [Edit Connectivity](#)


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 to Cloud Remote web UI.


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



4. 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.



5. Paste the encryption key that was copied to the clipboard into the **Encryption Key** field in the dialog box.
6. Click **Select File** and browse to the artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
7. Click **Confirm**.
8. 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).

Region Connectivity Running		Download Configuration Configure Region
Cloud endpoint accessible from Cloud Center Manager	No	
Cloud Center Manager AMQP reachable from worker VM's	No	
Cloud Center Manager AMQP accessible from cloud	Yes	
Remote AMQP IP		
Worker AMQP IP	192.168.30.16:5671	
Blade Name	cloudcenter-blade-vmware-9-0289	
Blade Port	8443	

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

Configure Cloud Remote in an AWS region as follows.

Obtain and Launch the Cloud Remote Appliance in AWS

1. Obtain the Cloud Remote shared AMI from Cisco support and launch it. Follow the same guidance for obtaining and launching the [CloudCenter Suite installer appliance for AWS](#).
2. 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.
3. 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) > <i>Custom Port Numbers (Conditional)</i>).
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) > <i>Custom Port Numbers (Conditional)</i>).
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 Download Configuration 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.

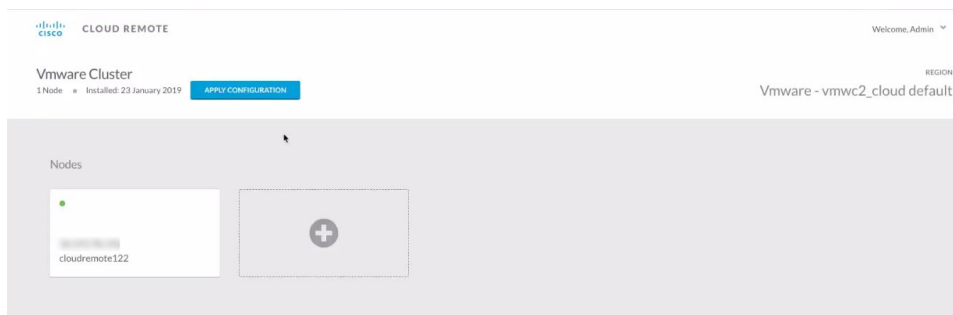


Click the **Copy Encryption Key** link to save the key to your clipboard. A success message will be displayed temporarily above the Region Connectivity section header. Make sure not to overwrite the clipboard with other data. You will need the key when you upload the configuration zip file to Cloud Remote.


 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 to Cloud Remote web UI.


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



4. 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.



5. Paste the encryption key that was copied to the clipboard into the **Encryption Key** field in the dialog box.
6. Click **Select File** and browse to the artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
7. Click **Confirm**.
8. 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).

Region Connectivity Running		Download Configuration Configure Region
Cloud endpoint accessible from Cloud Center Manager	No	
Cloud Center Manager AMQP reachable from worker VM's	No	
Cloud Center Manager AMQP accessible from cloud	Yes	
Remote AMQP IP		
Worker AMQP IP	192.168.30.16:5671	
Blade Name	cloudcenter-blade-vmware-9-0289	
Blade Port	8443	

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 region to support a Kubernetes target cloud as follows.

Obtain and Launch the Cloud Remote Appliance in AWS

1. Obtain the Cloud Remote shared AMI from Cisco support and launch it. Follow the same guidance for obtaining and launching the [CloudCenter Suite installer appliance for AWS](#).
2. 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.
3. 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 (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.

Region Connectivity Running [Download Configuration](#) [Configure Region](#)

Clicking Download Configuration 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.

Region Connectivity Enabling... [Download Configuration](#) [Copy Encryption Key](#) [Edit Connectivity](#)


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 to Cloud Remote web UI.


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



4. 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.



5. Paste the encryption key that was copied to the clipboard into the **Encryption Key** field in the dialog box.
6. Click **Select File** and browse to the artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
7. Click **Confirm**.
8. 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).

Region Connectivity Running		Download Configuration Configure Region
Cloud endpoint accessible from Cloud Center Manager	No	
Cloud Center Manager AMQP reachable from worker VM's	No	
Cloud Center Manager AMQP accessible from cloud	Yes	
Remote AMQP IP		
Worker AMQP IP	192.168.30.16:5671	
Blade Name	cloudcenter-blade-vmware-9-0289	
Blade Port	8443	

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

1. Download the Cloud Remote appliance for AzureRM as a zip file from software.cisco.com and then unzip it to reveal the VHD file.
2. 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 the [CloudCenter Suite installer appliance for AzureRM](#).



You must use the AzureRM CLI to perform this upload.

3. 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.
4. Once the first instance of the appliance has been launched, use the AzureRM 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) > <i>Custom Port Numbers (Conditional)</i>). 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) > <i>Custom Port Numbers (Conditional)</i>).
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) > <i>Custom Port Numbers (Conditional)</i>).
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.

Region Connectivity Running [Download Configuration](#) [Configure Region](#)

Clicking Download Configuration 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.

Region Connectivity Enabling... [Download Configuration](#) [Copy Encryption Key](#) [Edit Connectivity](#)


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.



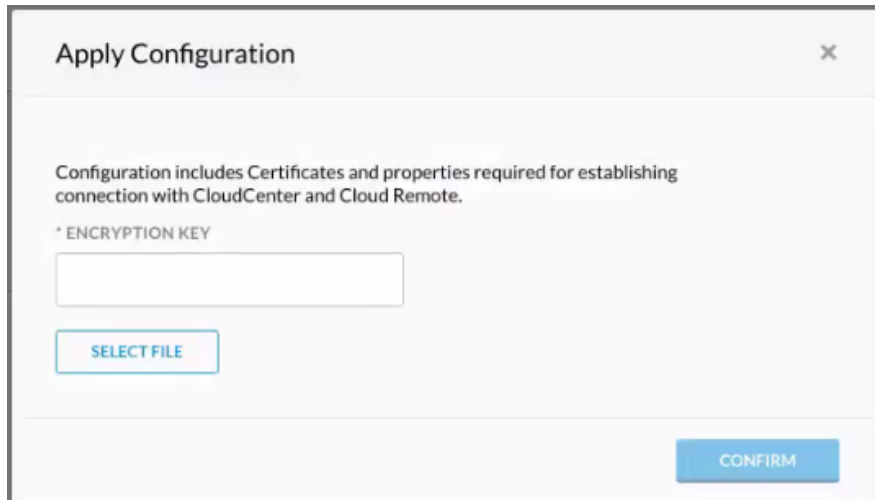
If you change the connectivity settings in the CloudCenter Suite UI and need to again download the zip file, a new encryption key is automatically created and can be copied to the clipboard by clicking the **Copy Encryption Key** link again.

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


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



- 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.



- Paste the encryption key that was copied to the clipboard into the **Encryption Key** field in the dialog box.
- Click **Select File** and browse to the artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
- Click **Confirm**.
- 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).

Region Connectivity Running[Download Configuration](#)[Configure Region](#)

Cloud endpoint accessible from Cloud Center Manager	No
Cloud Center Manager AMQP reachable from worker VM's	No
Cloud Center Manager AMQP accessible from cloud	Yes
Remote AMQP IP	
Worker AMQP IP	192.168.30.16:5671
Blade Name	cloudcenter-blade-vmware-9-0289
Blade Port	8443

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 region to support a Kubernetes target cloud as follows.

Download and Launch the Cloud Remote Appliance in AzureRM

1. Download the Cloud Remote appliance for AzureRM as a zip file from software.cisco.com and then unzip it to reveal the VHD file.
2. 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 the [CloudCenter Suite installer appliance for AzureRM](#).



You must use the AzureRM CLI to perform this upload.

3. 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.
4. Once the first instance of the appliance has been launched, use the AzureRM 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.

Region Connectivity **Running** [Download Configuration](#) [Configure Region](#)

Clicking Download Configuration 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.

Region Connectivity **Enabling...** [Download Configuration](#) [Copy Encryption Key](#) [Edit Connectivity](#)

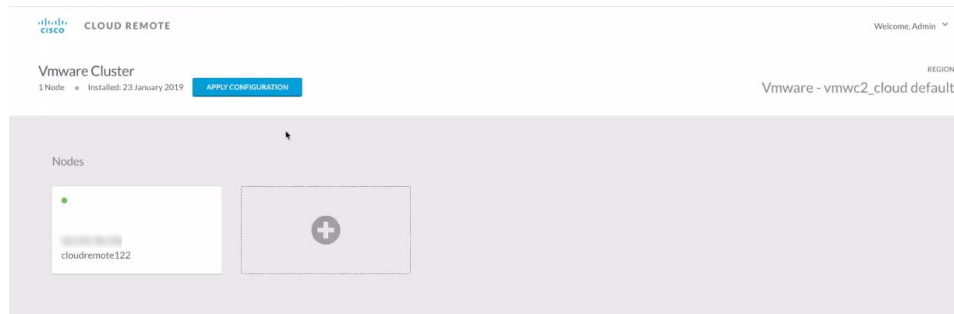
Click the **Copy Encryption Key** link to save the key to your clipboard. A success message will be displayed temporarily above the Region Connectivity section header. Make sure not to overwrite the clipboard with other data. You will need the key when you upload the configuration zip file to Cloud Remote.



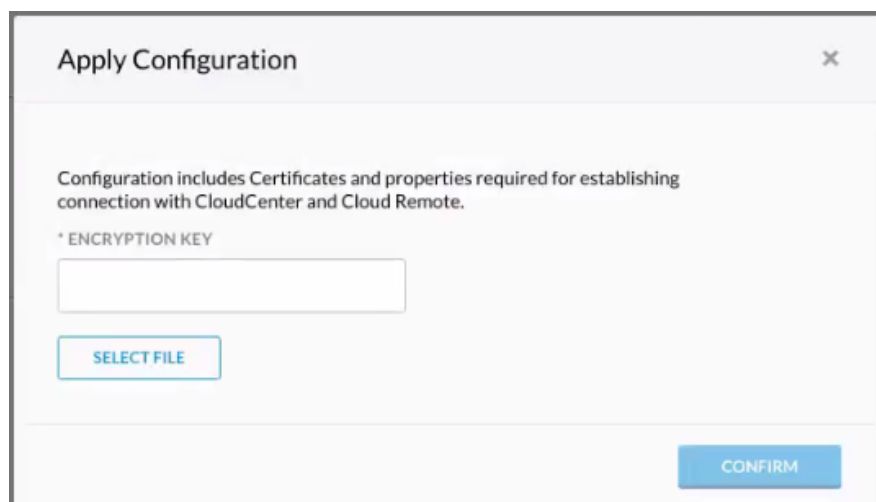
If you change the connectivity settings in the CloudCenter Suite UI and need to again download the zip file, a new encryption key is automatically created and can be copied to the clipboard by clicking the **Copy Encryption Key** link again.

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


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



4. 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.



5. Paste the encryption key that was copied to the clipboard into the **Encryption Key** field in the dialog box.
6. Click **Select File** and browse to the artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
7. Click **Confirm**.
8. 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).

Region Connectivity Running	Download Configuration	Configure Region
Cloud endpoint accessible from Cloud Center Manager	No	
Cloud Center Manager AMQP reachable from worker VM's	No	
Cloud Center Manager AMQP accessible from cloud	Yes	
Remote AMQP IP		
Worker AMQP IP	192.168.30.16:5671	
Blade Name	cloudcenter-blade-vmware-9-0289	
Blade Port	8443	

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 Google Region

Configure Cloud Remote in a Google region as follows.

Obtain and Launch the Cloud Remote Appliance in Google

1. Request the Cloud Remote shared VMI form Cisco support by opening a [CloudCenter Support case](#). In your request, specify the following details:
 - a. Your GCP account number
 - b. Your GCP project ID number
 - c. Your CloudCenter Suite version
 - d. Your Customer ID (CID)
 - e. Your customer name
 - f. Specify if your setup is in production or for a POC
 - g. Your Contact Email
2. After you open a case, your support case is updated with the shared VMI ID. **Proceed to the next step only after your support case is updated with the VMI ID.**
3. Navigate to the GCP dashboard and search for the VMI ID name provided in the [CloudCenter Support case](#) in the list of images for your project.
4. Launch an instance using the shared VMI.
 - a. Click on the image name. This takes you to the page for the image

The screenshot shows the Google Cloud Platform console interface. At the top, the header includes 'Google Cloud Platform', a user profile 'cliqimages', and a search icon. Below the header, the left sidebar shows the 'Compute Engine' menu with options like 'VM instances', 'Instance groups', 'Instance templates', 'Sole tenant nodes', 'Disks', 'Snapshots', 'Images' (highlighted), 'TPUs', and 'Committed use discounts'. The main content area is titled 'Images' and shows a single image entry with the following details:

- Instance Name:** ccs-google-cloudremote-5-0-0-20190204-1
- Description:** Pilot VM image
- Labels:** None
- Creation time:** Feb 4, 2019, 11:20:40 AM
- Encryption type:** Google managed
- Equivalent REST:** Equivalent REST

At the top of the image details, there are three action buttons: 'EDIT', 'DELETE', and 'CREATE INSTANCE'. The 'CREATE INSTANCE' button is highlighted with a blue background and a white cursor is pointing at it.

- b. Click on Create Instance to display the Instance properties page

Name ⓘ

instance-2

Region ⓘ **Zone** ⓘ

us-west1 (Oregon) us-west1-a


Machine type
Customize to select cores, memory and GPUs.

1 vCPU 3.75 GB memory [Customize](#)

Container ⓘ

Deploy a container image to this VM instance. [Learn more](#)

Boot disk ⓘ

 **New 30 GB standard persistent disk**
Image
ccs-google-cloudremote-5-0-0-2019020... [Change](#)

Identity and API access ⓘ

Service account ⓘ

Compute Engine default service account

Access scopes ⓘ

Allow default access
 Allow full access to all Cloud APIs
 Set access for each API

Firewall ⓘ

Add tags and firewall rules to allow specific network traffic from the Internet

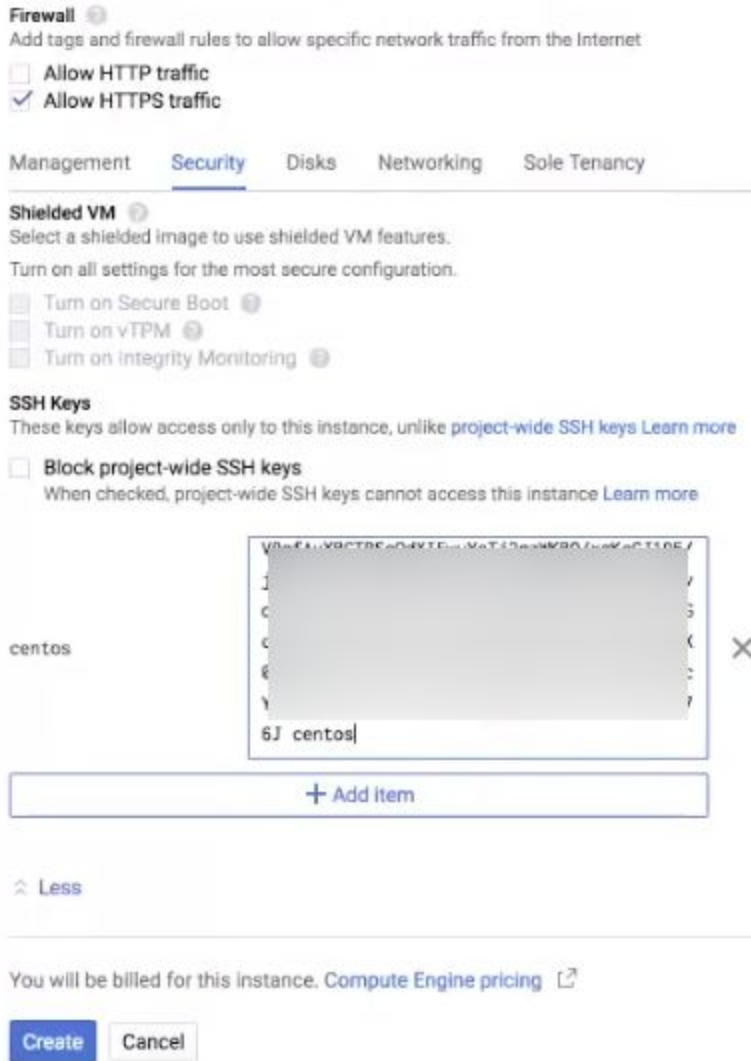
Allow HTTP traffic
 Allow HTTPS traffic

Management, security, disks, networking, sole tenancy

You will be billed for this instance. [Compute Engine pricing](#)

c. Complete these fields:

- i. Instance name
- ii. Region and zone
- iii. Machine type: select 2 vCPU, 7.5 GB RAM
- iv. Click the checkbox to allow HTTPS access
- v. 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.



- d. Click Create to launch the instance.
- 5. 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.
- 6. Once the first instance of the appliance has been launched, use the GCP 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) > <i>Custom Port Numbers (Conditional)</i>). 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) > <i>Custom Port Numbers (Conditional)</i>).
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) > <i>Custom Port Numbers (Conditional)</i>).
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.

Region Connectivity Running [Download Configuration](#) [Configure Region](#)

Clicking Download Configuration 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.

Region Connectivity Enabling... [Download Configuration](#) [Copy Encryption Key](#) [Edit Connectivity](#)


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.



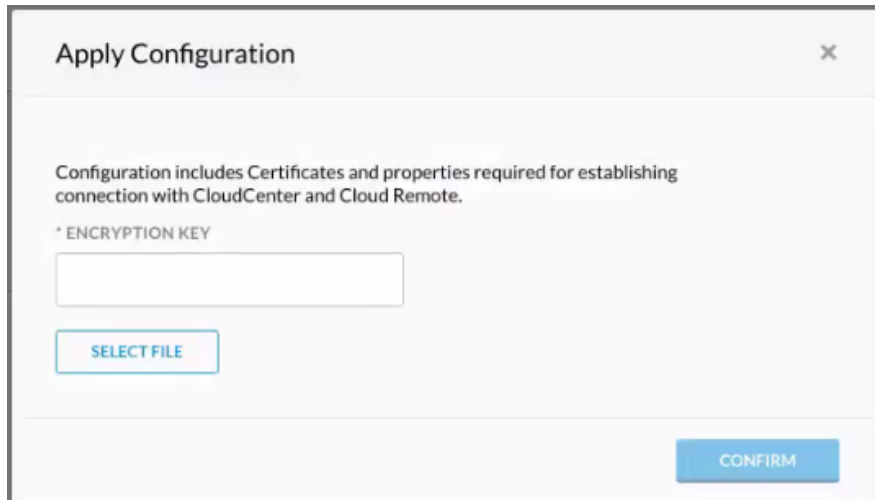
If you change the connectivity settings in the CloudCenter Suite UI and need to again download the zip file, a new encryption key is automatically created and can be copied to the clipboard by clicking the **Copy Encryption Key** link again.

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


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



- 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.



- Paste the encryption key that was copied to the clipboard into the **Encryption Key** field in the dialog box.
- Click **Select File** and browse to the artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
- Click **Confirm**.
- 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).

Region Connectivity Running[Download Configuration](#) [Configure Region](#)

Cloud endpoint accessible from Cloud Center Manager	No
Cloud Center Manager AMQP reachable from worker VM's	No
Cloud Center Manager AMQP accessible from cloud	Yes
Remote AMQP IP	
Worker AMQP IP	192.168.30.16:5671
Blade Name	cloudcenter-blade-vmware-9-0289
Blade Port	8443

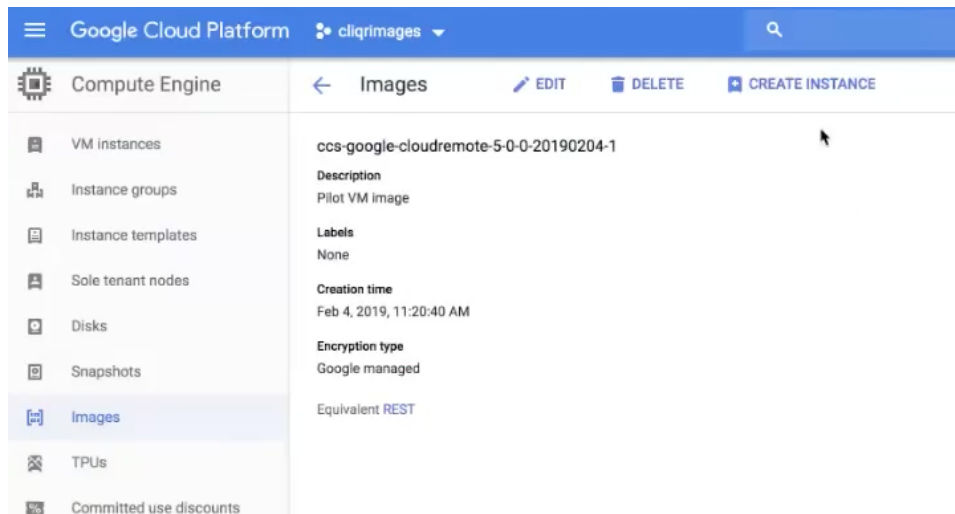
After completing these steps, Workload Manager and Cost Optimizer can use Cloud Remote for communicating with the target cloud region.

Configure Cloud Remote in a Google Region for a Kubernetes Cloud

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

Obtain and Launch the Cloud Remote Appliance in Google

- Request the Cloud Remote shared VMI form Cisco support by opening a [CloudCenter Support case](#). In your request, specify the following details:
 - Your GCP account number
 - Your GCP project ID number
 - Your CloudCenter Suite version
 - Your Customer ID (CID)
 - Your customer name
 - Specify if your setup is in production or for a POC
 - Your Contact Email
- After you open a case, your support case is updated with the shared VMI ID. **Proceed to the next step only after your support case is updated with the VMI ID.**
- Navigate to the GCP dashboard and search for the VMI ID name provided in the [CloudCenter Support case](#) in the list of images for your project.
- Launch an instance using the shared VMI.
 - Click on the image name. This takes you to the page for the image



- Click on Create Instance to display the Instance properties page

Name ⓘ

instance-2

Region ⓘ **Zone** ⓘ

us-west1 (Oregon) us-west1-a


Machine type
Customize to select cores, memory and GPUs.

1 vCPU 3.75 GB memory [Customize](#)

Container ⓘ

Deploy a container image to this VM instance. [Learn more](#)

Boot disk ⓘ

 New 30 GB standard persistent disk
Image
ccs-google-cloudremote-5-0-0-2019020... [Change](#)

Identity and API access ⓘ

Service account ⓘ

Compute Engine default service account

Access scopes ⓘ

Allow default access
 Allow full access to all Cloud APIs
 Set access for each API

Firewall ⓘ

Add tags and firewall rules to allow specific network traffic from the Internet

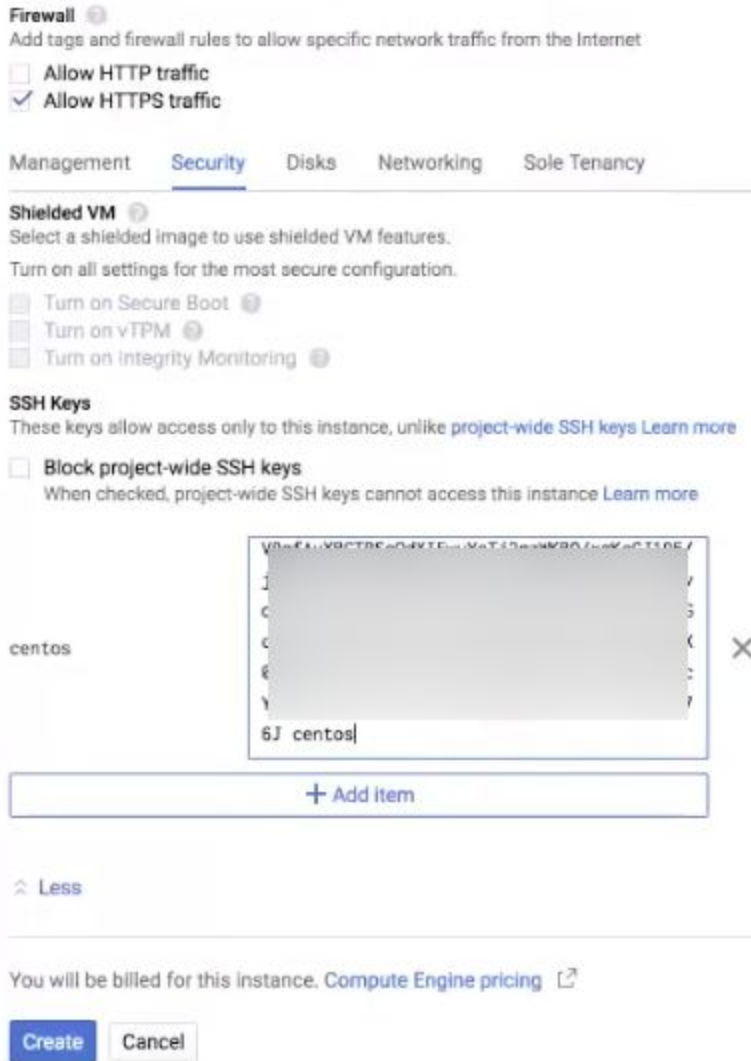
Allow HTTP traffic
 Allow HTTPS traffic

Management, security, disks, networking, sole tenancy

You will be billed for this instance. [Compute Engine pricing](#)

c. Complete these fields:

- i. Instance name
- ii. Region and zone
- iii. Machine type: select 2 vCPU, 7.5 GB RAM
- iv. Click the checkbox to allow HTTPS access
- v. 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.



- d. Click Create to launch the instance.
- 5. 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.
- 6. Once the first instance of the appliance has been launched, use the GCP 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	<p>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)).</p> <p>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.</p> <p>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.</p>
--	------------------------	---

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 Download Configuration 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 to Cloud Remote web UI.


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



4. 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.



5. Paste the encryption key that was copied to the clipboard into the **Encryption Key** field in the dialog box.
6. Click **Select File** and browse to the artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
7. Click **Confirm**.
8. 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).

Region Connectivity	Running	Download Configuration	Configure Region
Cloud endpoint accessible from Cloud Center Manager	No		
Cloud Center Manager AMQP reachable from worker VM's	No		
Cloud Center Manager AMQP accessible from cloud	Yes		
Remote AMQP IP			
Worker AMQP IP	192.168.30.16:5671		
Blade Name	cloudcenter-blade-vmware-9-0289		
Blade Port	8443		

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

The Cloud Remote artifacts mentioned in [Conditional Component Appliance Images](#) is called **ccs-cloudremote-artifacts-*<release.tag>*-YYYYMMDD.0.zip** and contains the following items:

- Installer script – Only applicable for IBM Cloud and vCD Cloud.
- Upgrade script – Applicable for all supported clouds.
- The proxy service script for the CloudCenter Suite cluster – Applicable for all supported clouds.

The items from this artifact are used in the procedures provided in this section.

To use a static IP address with Cloud Remote, follow this procedure.

1. SSH into the Cloud Remote VM.
2. Set the static (private) IP address using the following commands.

```
export HOST_IP=<static IP>
/opt/cisco/pilot/builds/
cd pilot_XXXX
cd bin
./bootstrap.sh
```



If multiple pilot_XXXX folder versions exist, use the following examples to identify the **latest, major version**:

- pilot_5.1.2-20191015.1
- pilot_5.1.2-20200111.1 > this is the latest pilot folder based on major version and date

The Cloud Remote internal network uses the 10.10.0.0/16 network range of IP addresses. If the Cloud Remote VM needs to be deployed in the same network range (10.10.0.0/16), then you must change the internal network range to another non-conflicting range.

To change the Cloud Remote internal network range, follow this procedure.

1. SSH into the Cloud Remote VM.
2. Issue the following commands:

```
cd /opt/cisco/pilot/builds/
cd pilot_XXXXXXX
cd docker
vi pilot_base.yml
# a. Search for 10.10.0.0/16 #Change this line to appropriate non-conflicting range
# b. Save and quit

cd ../bin/
./bootstrap.sh
```

Verify the following requirements to run the installer script on a custom CentOS 7 VM:

- This procedure is only applicable to CentOS 7 VMs.
- The VM should have 2 CPUs, 8GB Memory, and 30G storage.
- Run **yum update** on the VM.
- Run the following commands to update the kernel:

```
sudo rpm --import https://www.elrepo.org/RPM-GPG-KEY-elrepo.org
sudo rpm -Uvh http://www.elrepo.org/elrepo-release-7.0-3.el7.elrepo.noarch.rpm
sudo yum --disablerepo='*' --enablerepo='elrepo-kernel' list available
sudo yum --enablerepo=elrepo-kernel -y install kernel-ml
sudo grub2-set-default 0
sudo grub2-mkconfig -o /boot/grub2/grub.cfg
sudo reboot
```

To install Cloud Remote in your custom CentOS system, follow this procedure.



This procedure is only applicable for IBM Cloud and vCD.

1. Locate the Cloud Remote installer script (available in the Cloud Remote artifact mentioned in the section above) at software.cisco.com and copy it to a directory in your Cloud Remote instance.
2. Establish a terminal session to the Cloud Remote instance and navigate to the directory containing the installer script.
3. Run the following commands from the Cloud Remote command prompt.


```
[root@centos7cpsgcore ~]# ./cloudRemote5.1.0.bin
Verifying archive integrity... All good.
Uncompressing cloud remote 5.1.0 installer 100%
Usage: ./INSTALLER_FILE -- [--host-ip 'PRIVATE NETWORK IP ADDRESS']
example: ./cloudRemote5.1.0-20190614.0.bin -- --host-ip '1.2.3.4'    >>> Please note the extra --
before --host-ip
[root@centos7cpsgcore ~]#
```

4. Confirm the successful execution of the script.

To upgrade Cloud Remote (script available in the Cloud Remote artifact file mentioned in the section above) in your Workload Manager or Cost Optimizer system, follow this procedure for each instance of Cloud Remote.


1. Locate the Cloud Remote upgrade script at software.cisco.com and copy it to a directory in your Cloud Remote instance.
2. Establish a terminal session to the Cloud Remote instance and navigate to the directory containing the upgrade script.
3. Run the following commands from the Cloud Remote command prompt.

```
chmod +x UPGRADE_FILE
sudo ./ UPGRADE_FILE
```

4. Confirm the successful execution of the script.

After your initial Cloud Remote instance is launched and configured, it is recommended that you can add two additional nodes to form a cluster. When scaling up or down it is recommended not to run your cluster continuously with only two nodes. Follow this procedure:

1. Deploy a new instance of the appliance in the same network as the first appliance. Record its IP address. Alternatively, if you have another instance of Cloud Remote that you launched previously but stopped, restart that instance.
2. At the home page of the Cloud Remote web UI for the initial instance, click the tile with the plus icon. After clicking the plus icon, the tile will change and show an **Add IP** field as shown in the figure below. Enter the address of your newly launched (or restarted) instance in this field and then click **Done**.



Your new instance will become part of the cluster. There is no need to login to the new instance to set configuration. The cluster can be managed through the first instance's Web UI.

You can scale down the cluster in two steps:

1. From the Cloud Remote web UI home page, take note of the IP address of the node you want to remove from the cluster. Then remove it by hovering over its tile and clicking the trash icon.
2. Login to the cloud console for your target cloud and find the VM with the IP address of the node you just removed from the cluster. Stop that VM.

If firewall settings prevent you from using standard port numbers for HTTPS, AMQP, and Guacamole protocols, you can specify custom port numbers for those protocol using a **Change Ports shell script** that is included in the Cloud Remote appliance. Otherwise, Cloud Remote will use the standard port numbers as shown in the table below.

Service	Default Port
HTTPS (web UI)	443
AMQP (RabbitMQ)	5671
Guacamole	8443



- The Guacamole service is only needed for user access to VM-based deployments. Therefore, there is no need to create a custom port number for the Guacamole service if this Cloud Remote cluster is used to support connectivity to a Kubernetes target cloud.
- Only run the script after you have downloaded the artifacts.zip file (mentioned in the section above) from the region connectivity settings section of the Regions tab in the Workload Manager or Cost Optimizer UI, and then uploaded that file to Cloud Remote through the Cloud Remote web UI. In addition, if you later need to upload a new artifacts.zip file to Cloud Remote, the custom port settings will be erased and you will need to run the Change Ports script again.

Follow these steps to run the script:

1. Establish an ssh session to master (initial) Cloud Remote instance.
2. Navigate to the directory: `/opt/cisco/pilot/builds/<pilot folder>/bin`
3. Run the shell script:

```
changeports.sh
```

4. You are first prompted to see if you want to change the web UI port number. Type **Y** or **N**.
 - a. If you enter **Y**, you are prompted for:
 - i. Current port number. Type any number and then ENTER.
 - ii. New port number. Type the new port number and then ENTER. The script will attempt to change the port number on this node and then on all other nodes in your Cloud Remote cluster. When done, you are prompted whether you want to change the value of the next port.
 - b. If you enter **N**, you are prompted whether you want to change the value of the next port.
5. When you are prompted for the RabbitMQ port number, type **Y** and enter the old and then new port numbers as above, or type **N**, whichever is appropriate.
6. When you are prompted for the Guacamole port number, type **Y** and enter the old and then new port numbers as above, or type **N**, whichever is appropriate. If the target cloud is a Kubernetes cloud, the Guacamole server is not used and you would, therefore, enter **N**.



Be sure to verify that your proxy can access Cloud Remote's Port 5671 (RabbitMQ). If you've changed Cloud Remote's RabbitMQ port to 443, then the proxy must be able to access Cloud Remote's Port 443.

If your proxy restricts outbound ports, then you must configure Cloud Remote's's RabbitMQ port to *one of the accessible ports (usually 443)* using the **changeports.sh** script as listed in the *Custom Port Numbers (Conditional)* section.

The Cloud Remote can communicate with the CloudCenter Suite server by using the Cisco proxy to access outbound environments. You can enable direct connectivity between CloudCenter Suite and Cloud Remote using a script that is included with the Cloud Remote artifact file mentioned in the section above. This allows you to avoid using the Cisco proxy for external communications when using the CloudCenter Suite.

This section directly relates to the setting when you specify the *AMQP and Guacamole Addresses for Supporting Cloud Remote* or when you specify the *AMQP Addresses for Supporting Cloud Remote for a Kubernetes Cloud*. This setting is highlighted in the following screenshots for a private (screenshot on the left) and private (right screenshot on the right) clouds:

Configure Region ✕

IS CLOUD END POINT DIRECTLY ACCESSIBLE?

NO

SHOULD WORKER VMS DIRECTLY CONNECT WITH CLOUDCENTER SUITE?

NO

IS CLOUDCENTER SUITE DIRECTLY ACCESSIBLE FROM YOUR CLOUD REMOTE?

YES

LOCAL AMQP IP ADDRESS

WORKER AMQP IP ADDRESS

GUACAMOLE PUBLIC IP AND PORT

GUACAMOLE IP ADDRESS AND PORT FOR APPLICATION VMS

Add Cloud ✕

Connectivity Settings

SHOULD WORKER VMS DIRECTLY CONNECT WITH CLOUDCENTER SUITE?

NO

IS CLOUDCENTER SUITE DIRECTLY ACCESSIBLE FROM YOUR CLOUD REMOTE?

YES

Diagram below is based on selections above

Depending on the environment, users may need the proxy service to be on the Cloud Remote or the CloudCenter Suite cluster.

When you click **Edit connectivity** and in the *Configure a Region* popup, you have answered as follows:

- **Should Worker VMs directly connect with CloudCenter Suite?** is set to **NO**
- **AND**
- **Is CCS directly accessible from your cloud remote?** is also set to **NO**
- **THEN:**

- You are indicating that the communication between the CloudCenter Suite and the Cloud Remote is going through an AMQP instance.
- **After configuring Cloud Remote**, you must return to this popup and add the IP of your Cloud Remote server.

Proxy Service on the Cloud Remote Instance

For this scenario, the CloudCenter Suite resides on one cloud (for example, VMware datacenter/Private cloud) and the Cloud Remote resides on another cloud (for example, GKE/SaaS/Public cloud). When you configure the region for a cloud in this scenario and you toggle the **Is CloudCenter Suite Directly Accessible from Your Cloud Remote** setting to **Yes**, then this setting is indicative of the CCS to Cloud Remote communication going through an AMQP instance.

To enable the proxy service on the Cloud Remote instance, follow this procedure.

1. Establish an SSH session to the master (initial) Cloud Remote instance.
2. Navigate to the directory: `/opt/cisco/pilot/builds/<pilot folder>/bin` folder. For example:

```
cd /opt/cisco/pilot/builds/pilot_5.1.0-PILOTVERSION/bin/config_crproxy.bin
```

3. SSH into the Cloud Remote instance and run the CR proxy installer that is located in the directory that you set in Step 2 above.
4. Here are the sample usage and output.

You have now enabled the proxy service on the Cloud Remote instance. You can verify the connectivity in the region settings Connectivity section as displayed in the following screenshot.

Region Connectivity **Running**

Cloud endpoint accessible from CloudCenter Suite	No
CloudCenter Suite AMQP reachable from worker VM's	No
CloudCenter Suite AMQP accessible from cloud	Yes
Local AMQP IP	[REDACTED]
Worker AMQP IP	[REDACTED]
Guacamole Public IP and Port	[REDACTED]
Guacamole IP Address and Port for Application VMs	[REDACTED]
Blade Name	cloudcenter-cloud-blade-[REDACTED]

Proxy Service on the CloudCenter Suite Cluster

For this scenario, the CloudCenter Suite resides on one cloud (for example, GKE/SaaS/Public cloud) and the Cloud Remote resides on another cloud (for example, VMware datacenter/Private cloud). When you configure the region for a cloud in this scenario and you toggle the **Is CloudCenter Suite Directly Accessible from Your Cloud Remote** setting to **No**, then this setting is indicative of the communication between the CloudCenter Suite and the Cloud Remote going through an AMQP instance.

To enable the proxy service on the CloudCenter Suite cluster, follow this procedure.

1. Make sure KUBECONFIG environment variable is set. The user must have the applicable permissions to create Kubernetes services and deployments.

```
kubectl get svc
```

```
#The above command should return all the services in your Cisco CloudCenter Suite cluster.
```

2. Locate and download the `ccs-cloudremote-artifacts-<release>-<date>.zip` for the Cloud Remote version that you are using from software.cisco.com.
3. Locate and copy the `config_k8scrproxy.bin` file from the `ccs-cloudremote-artifacts-<release>-<date>.zip` file to a directory in your CloudCenter Suite instance, and execute it.
4. Here are the sample usage and output.

```

CISCO-M-K192:crproxy cisco$ ./config_k8scrproxy.bin
Verifying archive integrity... 100% All good.
Uncompressing Proxy for cloudremote in K8S cluster 100%
Usage:
./config_k8scrproxy.bin
-- --namespace 'K8S NAMESPACE' --region-id 'CLOUD REGION ID'
--proxy-host 'PROXY HOST' --proxy-port 'PROXY PORT' --target-amqp-host
'CLoud REMOTE IP' --target-amqp-port 'CLOUD REMOTE AMQP PORT'
[--docker-image-url 'DOCKER IMAGE URL of CRPROXY' --proxy-user 'PROXY
USERNAME' --proxy-passwd 'PROXY PASSWORD']
if option --docker-image-url is not provided, predefined image will be used
No Authentication example: ./config_k8scrproxy.bin -- --namespace cisco --region-id 28 --proxy-host
proxy.example.com --proxy-port 80 --target-amqp-host 1.2.3.4 --target-amqp-port 443
With Authentication and non-default docker image url example: ./config_k8scrproxy.bin -- --namespace
cisco --region-id 28 --proxy-host proxy.example.com --proxy-port 80 --target-amqp-host 1.2.3.4 --target-
amqp-port 443 --proxy-user 'user' --proxy-passwd 'password' --docker-image-url devhub.example.com
/crproxy:latest

CISCO-M-K192:crproxy cisco$ ./config_k8scrproxy.bin -- --namespace cisco --region-id 28 --proxy-host
proxy.example.com --proxy-port 80 --target-amqp-host 1.2.3.4 --target-amqp-port 443 --docker-image-url
dockerhub.cisco.com/cloudcenter-dev-docker/custom/cloudcenter/crproxy:latest
Verifying archive integrity... 100% All good.
Uncompressing Proxy for cloudremote in K8S cluster 100%
cisco 28 dockerhub.cisco.com/cloudcenter-dev-docker/custom/cloudcenter/crproxy:latest proxy.example.com
80 1.2.3.4 443
service "cloudcenter-blade-crproxy-28" deleted
deployment.extensions "cloudcenter-blade-crproxy-28" deleted
service "cloudcenter-blade-crproxy-28" created
deployment.apps "cloudcenter-blade-crproxy-28" created
cloudcenter-blade-crproxy-28
ClusterIP xx.xxx.xx.xxx <none>
12850/TCP 0s
socat TCP4-LISTEN:12850,reuseaddr,fork PROXY:proxy.example.com:1.2.3.4:443,proxyport=80

```

5. In this sample procedure, the CloudCenter Suite is configured to use `<cloudcenter-blade-crproxy-28:12850>` proxy pod. You must now set up connectivity between Cloud Remote and the CloudCenter Suite cluster:
 - a. Login to the CloudCenter Suite and navigate to the corresponding *cloud Region* page.
 - b. Click the **Edit Connectivity** link.
 - c. Set the value of the **Remote AMQP IP** field to `cloudcenter-blade-crproxy-<regionId>:12850`.
 - d. Download and apply the configuration to the Cloud Remote and wait for the **Region Connectivity** status to change to **Running**.
6. You have now enabled the proxy pod service on the CloudCenter Suite instance. You can verify the connectivity in the region settings Connectivity section as displayed in the following screenshot.

Region Connectivity **Running**

Cloud endpoint accessible from CloudCenter Suite	Yes
CloudCenter Suite AMQP reachable from worker VM's	No
CloudCenter Suite AMQP accessible from cloud	No
Remote AMQP IP	cloudcenter-cloud-blade-amazon- XXXXXXXXXX
Worker AMQP IP	XXXXXXXXXX
Guacamole Public IP and Port	XXXXXXXXXX
Guacamole IP Address and Port for Application VMs	XXXXXXXXXX
Blade Name	cloudcenter-blade-amazon- XXXXXXXXXX

- **Issue:** When you install Cloud Remote, you may sometimes see the following issues:
 - The Cloud Remote UI does not render even after a long time.
 - The Cloud Remote installer continues to poll after the installation.

Workaround: In both situations, follow this procedure to address the issue.

1. Run the following command to verify if the *Pilot/Babl* container is crashing.

```
docker ps
```

2. If it is crashing, run the following command.

```
docker service update --health-interval=30s --health-retries=1000 pilot_babl
```

3. This command can take up to 5 minutes to complete. After applying the configuration, if the Pilot/RabbitMQ container continues to crash, run the following additional command.

```
docker service update --health-interval=30s --health-retries=1000 pilot_rabbitmq
```

- **Issue:** The network connection is slow when using Cloud Remote.

Workaround: Try changing the health interval timeout period:

```
docker service update --health-interval=5m --health-start-period=10m --health-timeout=10m  
pilot_remoteproxy
```


Local Repo Appliance (Conditional)

Local Repo Appliance

- [Overview](#)
- [Deploy the Appliance](#)
- [Configure the Appliance](#)
- [Configure Cloud Settings](#)
- [Periodically Update the Bundle Store](#)

The Workload Manager local repo appliance is based on CentOS 7 and has an Apache web service and a version of the package store and bundle store preinstalled. You would deploy this appliance to your target VM-based cloud if either of these two conditions is true:

- Your workload VMs cannot access the Cisco-hosted package store and bundle store.
- You want to reduce the latency associated with downloading files from the package store or bundle store.

There are four tasks associated with the installing and using the local repo appliance:

- Deploy the appliance appropriate to your cloud type
- Configure the appliance using the repo wizard script
- Configure the Cloud Settings section of the associated region with the appliance's package store and bundle store URLs.
- Periodically update the bundle store



If using an unhardened version of CentOS 7 for the local repo appliance is not permitted or desired in your environment, you can build your own local repo appliance on your own Linux OS by following the instructions in [Local Package Store \(Conditional\)](#) and [Local Bundle Store \(Conditional\)](#).

The local repo appliance comes in different form factors corresponding to the following cloud types:

Cloud Type	Appliance Form Factor
vCenter	OVA file downloaded from http://software.cisco.com
OpenStack	qcow2 file downloaded from http://software.cisco.com
AzureRM	zip file downloaded from http://software.cisco.com
AWS	private image shared with your AWS account upon your request
GCP	private image shared with your GCP account upon your request

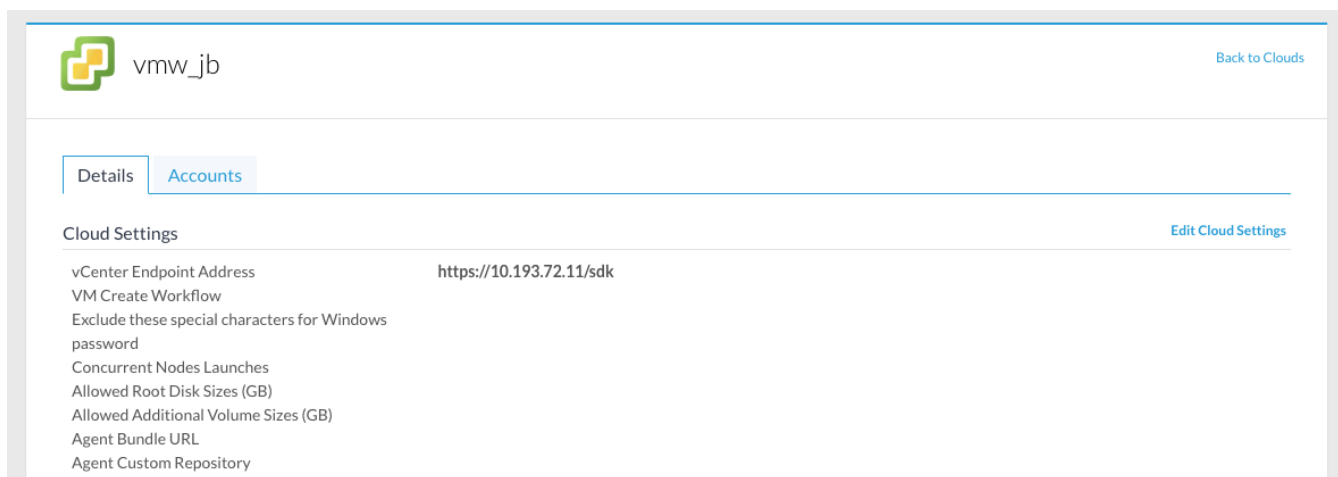
Deploy the appliance per the convention for your cloud region. Note the VM's IP address.

1. Once your appliance is powered on, establish an SSH session to it using IP address you noted above.
2. From the appliance command prompt, run the repo configuration shell script at `/usr/bin/repo_config_wizard.sh`. This invokes a text UI. Dismissing the welcome message displays the configuration menu. The menu has three choices: Proxy_Settings, Repo_Syncup, and Exit.
3. Selecting Proxy_Setting brings up a new menu allowing you to specify a SOCKS proxy URL. Enter the address of your proxy server if you have one.
4. Selecting Repo_Syncup displays a confirmation message. Agreeing to the confirmation message causes the script to sync the local package store the latest package store at repo.cliqrtech.com.



The local repo appliance is configured with a cron job that will automatically attempt to sync with the package store at <http://repo.cliqrtech.com> every day at midnight. It is, therefore, necessary that the local repo appliance has at least occasional internet access.

In order for your workload VMs to access the local repo appliance, the URLs of the appliance's package store and bundle store must be entered into the appropriate fields in the Cloud Settings section of the corresponding region. For a vCenter cloud, this section is at the top of the Details tab (see figure below). For all other clouds this section is displayed in the Regions tab after selecting the appropriate region.



vmw_jb [Back to Clouds](#)

Details **Accounts**

Cloud Settings [Edit Cloud Settings](#)

vCenter Endpoint Address	https://10.193.72.11/sdk
VM Create Workflow	
Exclude these special characters for Windows password	
Concurrent Nodes Launches	
Allowed Root Disk Sizes (GB)	
Allowed Additional Volume Sizes (GB)	
Agent Bundle URL	
Agent Custom Repository	

Click the **Edit Cloud Settings** link in the upper right to bring up the Cloud Settings dialog box as shown below. Note that the fields in this dialog box will vary based on the cloud type.

Configure Cloud Settings

vCenter Endpoint Address *

VM Create Workflow

Exclude these special characters for Windows password

Concurrent Nodes Launches

Allowed Root Disk Sizes (GB)

Allowed Additional Volume Sizes (GB)

Agent Bundle URL

Agent Custom Repository

For the **Agent Custom Repository** field enter the address of the local repo appliance:

```
http://<local _repo_ip_address>
```

For the **Agent Bundle Store** field enter the Agent Custom Repository URL followed by "/" and the directory name of the bundle store. To get the directory name of the bundle store, connect to the appliance's console, navigate to the http root directory, which is /repo/, and perform a directory listing. Look for directory with a name in the following format:

```
release-<release_version>
```

The bundle store directory will be found here in a subdirectory named "bundle". In this case, the Agent Custom Repository URL you should use would be in the format:

```
http://<local _repo_ip_address>/release-<release_version>/bundle
```



If you later download a newer bundle store to the appliance and place it in a new directory, you will need to update the Agent Bundle URL for this region with a URL that points to the new directory.

The local repo appliance automatically syncs the package store with `repo.ciqrtech.com`. However, the bundle store contents on the appliance must be updated manually. You would do this whenever a new version of the bundle store is posted on `software.cisco.com` which is typically every minor software release cycle. It is recommended that you create a new directory for the new bundle store under `/repo/` and name that directory using a format that includes the release number, for example, `release-<release_number>`.

After creating the new directory, download the bundle store corresponding to your cloud type to the new directory. Workload Manager has bundle stores for these cloud types: AzureRM, OpenStack, vCenter. These bundle stores are listed at `software.cisco.com` using the following naming convention: `<cloud_type>-cc-bundle_artifacts.zip`.

After downloading the bundle store zip file, unzip the contents and delete the original zip file.

Make sure to update the **Agent Bundle Store** URL field in the Cloud Settings dialog box as explained above.

Worker (Conditional)

Worker (Conditional)

- [Management Agent \(Worker\)](#)
- [Options to Install the Worker](#)
- [Install Worker on a Linux Image](#)
- [Install Worker on a Windows Image](#)

Management Agent (Worker)

Management Agent (Worker)

- [About the Worker](#)
- [Management Agent Tasks](#)
- [FAQs](#)

To deploy VMs and run applications on VMs that use the CloudCenter platform (called *worker* VMs), you need a *Management Agent* (software installed on the worker to communicate with the Workload Manager) to be installed on each worker VM.

When the VM first boots up, the worker communicates with the bundle store to download the latest version of the agent and then starts the agent. Once the Management Agent is started, it needs access to the Cisco hosted bundle store (<http://cdn.cliqr.com>) or to your locally-installed custom bundle store to download and start the service.

The *Worker image* provided or shared by Cisco contains the necessary software installed to enable agent installation even if the cloud does not support the `cloud-init` functionality.

See [Options to Install the Worker](#) for details on how you can install the worker.

The Management Agent communicates with the CloudCenter Suite cluster, either directly or through [Cloud Remote](#), and receives instructions to perform the following tasks:

- Complete application deployment tasks
- Perform provisioning tasks (for example running configuration scripts)
- Run custom cleanup scripts (for example to de-provision or shutdown applications)
- Enforce policies (for example, to reconfigure middle ware service during auto-scaling)
- Collect system metrics based on policy requirements
- Monitor data, provide status updates, and keep alive system heartbeats.

The Management agent (also called the agent) can be used in two modes.

VM Mode	Description
Imported	<ul style="list-style-type: none"> • Refers to the non-dynamic bootstrapping mode. • You must install the agent (add custom code to the VM) when the image is launched as a VM. • The agent can be installed on VMs that have been <i>imported</i> into CloudCenter. • Use it to run custom actions on Imported VMs which were not deployed via the Workload Manager. • It is an alternate option for VMs that do not require the capability to launch applications but do require some basic CloudCenter functionality like performing platform actions. • If installed, the Virtual Machines page and the VM Details page display the icon and version. • See Virtual Machine Management > Management Agent for additional details.
Deployed by Workload Manager	<ul style="list-style-type: none"> • Refers to the dynamic bootstrapping mode. • You do not need to install the agent as it is handled by Cisco. A custom image is not required by common public clouds as these clouds are capable of supporting <code>cloud-init</code>. However, if you prefer to use custom images in these clouds, you can. • Runs on workers created as part of applications deployed by the Workload Manager. • In this mode, it has all available capabilities listed in the <i>About the Worker</i> and <i>Management Agent Tasks</i> sections above. • The Cloud SDK functionality, introduced with Workload Manager 5.2, provides details on this functionality – if you are using the Cloud SDK solution.

1. Question: Is the Management Agent required?

Answer: You do not need to install either the Management Agent in the following cases:

- If a service is defined as *agentless* (*without an agent*, see [Custom Service Definition](#) for additional details), then the applications running these services do not require an agent to be installed.
- If you don't run custom scripts on imported VMs, you can continue use these VMs without installing the Management Agent.

2. Question: How is the Management Agent different from the AgentLite?

Answer: The legacy management agent used in CloudCenter Platform 4.x is no longer available. AgentLite was the interim lightweight agent that was used in CloudCenter Suite 5.0. Effective CloudCenter Suite 5.0, the term **Management Agent** refers to the agent in the CloudCenter Suite. The term *AgentLite* is no longer used in documentation. If you have a legacy agent installed, you can always update it to the CloudCenter Suite 5.x Management Agent.

3. Question: Why is the Agent installed on Worker VM as part of Application deployment?

Answer: An agent is required to support on-demand actions and lifecycle actions defined in the Actions Library, service definition and application profile

4. Question: How does the worker get installed on a Workload Manager deployed VM?

Answer: Two ways:

- a. *Create a pre-bootstrapped image:* Use the worker installer for [Linux](#) or [Windows](#), depending on your base OS, to manually create an image with the worker fully installed. Cisco also provides a Centos 6 pre-bootstrapped image as an appliance. These images are used in clouds that do not support dynamic bootstrapping.
- b. *Dynamic Bootstrapping Injection:* When a VM is launched in out-of-box clouds that support the cloud-init functionality then Workload Manager refers to this bootstrapping process as dynamic bootstrapping. In clouds that support dynamic bootstrapping, the Workload Manager specifies a script that must be run on the VM when it is booted for the first time. This script on execution installs the Management Agent on the Worker VM.

5. Question: Should I use the Custom Image or use Dynamic Bootstrapping?

Answer: This depends on your environment! Here are the details to help you make your decision:

- a. **Custom Image:** See Question 2 above.
- b. **Dynamic Bootstrapping:** Workload Manager installs the agent and its dependencies dynamically by automatically pushing the updated agent to each VM at the time of provisioning.

6. Question: What services are created by the worker? Why are these services created?

Answer: The following services are created when you install the agent:

- a. *CliQr install Service:* This file is a part of the agent install process for the Windows agent. It is the first service that sets the stage to download the agent bundle and kickoff the agent startup process.
- b. *Agent Service:* Platform-independent service of the management agent. It establishes the AMQP connection and processes information from the CloudCenter Suite cluster.
- c. *Sigar Service:* Platform-independent for data collection and reporting.

7. Question: If I want to run an image with the agent, is it necessary to have access to the Bundle Store?

Answer: All worker VMs when first launched must communicate with the bundle store to install the latest version of the agent. If the publicly available Cisco-hosted bundle store at <http://cdn.cliqr.com/cloudcenter-/bundle> (no trailing slash at the end of the URL) is not accessible to the worker VM, you must install the [Local Repo Appliance \(Conditional\)](#).

8. Question: How does the agent communicate with the CloudCenter Suite cluster?

Answer: Once the agent is installed, it needs to connect with the AMQP service in the CloudCenter Suite cluster. If the worker VM cannot initiate this connection due to firewall restrictions, you must install [Cloud Remote \(Conditional\)](#).

9. Question: Which files are created as part of the agent installation process?

Answer: The list of file created as part of the agent installation process differs for Linux and Windows installations. See the [Deployment Lifecycle Scripts > Lifecycle Action Script Definition](#) for additional context.

10. Question: Why is cliqruser permission required to run some scripts?

Answer: By default, key-based authentication is configured using cliqruser permission – this user refers to the OS user in the Application VM (Worker). See the [Deployment Lifecycle Scripts > Lifecycle Action Script Definition](#) for script-specific details on when cliqruser permission is required.

11. Question: Is the authorized_keys folder automatically deleted?

Answer: Yes. After a VM launch the Workload Manager cleans up the authorized keys specified in the `/root/.ssh/authorized_keys` file for all OS types (except Ubuntu) during the node initialization phase. For Ubuntu OS, the Workload Manager removes the `authorized_keys` from the `/home/ubuntu/.ssh/` folder. After removing the `authorized_keys` file, the agent injects the specified or auto-generated keys associated with a `cliqruser`. The `cliqruser` credentials can be used to log into the VM.

If you do not want the `authorized_keys` file to be automatically deleted, be sure to set the following flag in the worker image:

```
touch /etc/opt/.cloudcenter_do_not_delete_authorized_keys && chmod 444 /etc/opt/.cloudcenter_do_not_delete_authorized_keys
```

12. Question: Why does the agent program retain files in the C:\temp and C:\ directories?

Answer: Some deployment scripts like `resumeScript` or `reboot` or `restore` take effect based on information retained in these directories. These scripts require information from those files to function as designed. Each script, the level at which it is defined, the script download location, the user running the script, and the location from which the script is run in provided in detail in the [Deployment Lifecycle Scripts > Lifecycle Action Script Definition](#) section for both Linux and Windows environments.

13. Question: Does the worker installer for Windows install any open source tools on Windows workers?

Answer: No, the Windows worker installer does not install any open source tools on Windows. It does, however, install `ccc_unarchiver`, a proprietary tool to open TAR files.

14. Question: How are application services different from the services created by the agent installation process?

Answer: The agent services and application services are independent of each other.

- a. **Agent Services:** Specific to the agent installer. When you install the agent, the Agent installer installs some services to help with the installation and maintenance of the management agent. These files include the CliQr Installer Service, Agent Service, and other services.
- b. **Application Services:** Specific to application deployment. Once you set up image mapping (see [Map Images](#)) for each application VM, the Workload Manager module dynamically makes some out-of-box services available in the Workload Manager UI Topology Modeler page. See [OOB Services](#) for a complete list of these services.

Options to Install the Worker

Options to Install the Worker

- [Overview](#)
- [Using Dynamic Bootstrapping](#)
- [Using the CentOS 6 Pre-bootstrapped Image](#)
- [Using the Worker Installer Executable](#)

Depending on your cloud environment and application requirements there are three options to install the worker installed on the VMs you want to launch through Workload Manager:

- Using Dynamic Bootstrapping
- Using the CentOS 6 Pre-bootstrapped Image
- Using the Worker Installer Executable

The worker can be dynamically installed on VMs launched in most clouds that support an init-string passed as user data. For a matrix of which cloud providers and logical base images are supported see [Dynamic Bootstrapping](#). This is the simplest approach to get the worker and agent installed on you deployed VMs and is recommended unless you have special requirements.

- CloudCenter Suite includes a pre-bootstrapped CentOS 6 image for the clouds listed in [Conditional Component Appliance Images](#).
- After you obtain the image (received a shared version or import the image to your cloud), you must make sure the Workload Manager CentOS 6 OOB logical image is properly [mapped](#) to the physical image via its Image ID.
- If your application services can run on CentOS 6, this may be a convenient option.



Due to licensing restriction, a pre-bootstrapped virtual appliance is not available for custom Windows images. You are required to use the instructions in [Install Worker on a Windows Image](#).

The most flexible option for installing the worker is to run the worker installer appropriate for your OS:

- [Install Worker on a Linux Image](#)
- [Install Worker on a Windows Image](#)

See [Management Agent](#) for additional details.

Install Worker on a Linux Image

Install Worker on a Linux Image

- [Overview](#)
- [Containerized OOB Services](#)
- [Cloud Nuances](#)
- [Installation Process](#)
- [Identifying the OS Type](#)
- [Successful Installation](#)
- [Map a Logical Image to the Pre-bootstrapped Image](#)

Use the Workload Manager installer binary *worker_installer.bin* to create Workload Manager-enabled Linux images on different clouds and regions.

See [OOB Logical Images](#) for a list of logical images for Linux and other servers.



When you modify a pre-built image (for example, when performing an OS update, installing a new tool, upgrading the Management Agent, and so forth), be aware that the pre-built image already has Workload Manager Tools installed.

When updating an image that already has Workload Manager Tools installed, you must delete this file (*/usr/local/osmosix/etc/hostid*) prior to taking the final snapshot.



Use this procedure to also install Workload Manager Tools on CentOS 7 and RHEL-7.

When you manually install the [Management Agent \(Worker\)](#) on CentOS 7, RHEL 7, or Ubuntu 14, and select the "worker1" install option, the following Out-Of-Box (OOB) Workload Manager services will automatically run inside a Docker container inside the worker VM:

- [Apache 2](#)
 - [Memcached](#)
 - [MongoDB](#)
 - [MySQL](#)
 - [Nginx](#)
 - [Tomcat 6](#)
 - [Tomcat 7](#)
- To use VMware VMs, you should have already installed VMware tools on the Linux machine.
1. Install CliQr Tools on the application VM.
 2. Run the following commands after the installation is completed:

```
rm -f /var/lib/waagent/ovf-env.xml
waagent -deprovision+user
```

3. Stop the VM.
4. Capture the image.

To acquire and run the Linux worker installer, follow this procedure:

- Download the artifacts.zip archive from software.cisco.com, unzip the archive, and identify the installer package: *worker_installer.bin*.
- SSH into the application VM instance using the key pair that you used to launch the VM and go to the */tmp* directory.

```
sudo -i cd /tmp
```

- Copy the installer package from you computer to this directory.
- Change permissions to allow execution of the installer, specify a local package store (if necessary), and run the installer using the syntax shown below:

```
chmod 755 worker_installer.bin

#Set the following only if a local store is setup
export CUSTOM_REPO=<http://local_package_store IP>

./worker_installer.bin <ostype> <cloudtype> worker_basic

#See Syntax
```

ostype = rhel6, rhel7, oel6, centos6, centos7, ubuntu10, ubuntu1204, ubuntu1404 (see the section below to identify your OS type)

cloudtype = amazon, azure, google, openstack, vmware, and custom



If your cloud is one of the **custom** Cloud Type, you will be prompted to enter the location of your userdata and metadata extractor scripts. Contact your custom cloud provider for the script locations.

worker_basic = The *worker_basic* option installs all components necessary to download, start, and support the agent.

- Clean up and exit the VM instance:

```
rm worker_installer.bin rm ~/.ssh/authorized_keys exit
```

- If your image will be launched in an environment without internet access, add a qualifier on download and install the net-tools package: <https://sourceforge.net/projects/net-tools/>
- Save the image.

To verify or identify the OS type, follow this procedure:

1. In your terminal, issue one of the following commands (based on your Linux implementation):

- a. **\$ cat /etc/issue**

```
Ubuntu 14.04.2 LTS \n \l
```

or

- b. **\$ cat /etc/*-release**

```
CentOS release 6.3 (Final)
```

2. Based on the response to one of these commands, you can identify if your OS type.

After you run the installer commands, the installation results are displayed on the screen as follows:

- **Success scenarios:** Identifies a list of successfully installed components in green text.
- **Failure scenarios:** Provides a path to the log file that provides details of each failure.

In order to use the new pre-bootstrapped image in deployments, you must set up the appropriate [image mapping](#).

Install Worker on a Windows Image

Install Worker on a Windows Image

- [Overview](#)
- [Cloud Nuances](#)
- [Installation Process](#)

Use the Workload Manager installer executable `cliqr_installer.exe` to create Workload Manager-enabled Windows images on different clouds and regions.

See [OOB Logical Images](#) for a list of logical images for Windows and other servers.

You must first install VMware tools on the VM.



This procedure may differ based on your cloud and your Windows version. This procedure provides a point of reference to install the `WINDOWS_WORKER_OS_VERSION` image.

To install the worker on a Windows image, follow this procedure:

1. Download the `artifacts.zip` file from software.cisco.com and unzip it to obtain the installer package (`cliqr_installer.exe`).



You should only run the `cliqr_installer.exe` package once on a clean image. Do not attempt to run the installer twice on the same image.

2. Launch a base Windows image from the Cloud Provider Console. The following image is an example for OpenStack.

Image Name	Type	Status	Visible	Public	RAW	Actions
windows-2008-r2-std_2014-11-14-v6	Image	Active	Yes	No	RAW	Launch More
windows-2012-r2-std_2014-11-14-v6	Image	Active	Yes	No	RAW	Launch More
trusty-server-cloudimg-amd64-disk1.img1414403314246	Image	Active	Yes	No	RAW	Launch More
ubuntu-precise-server-cloudimg-amd64-disk1.2014-09-27	Image	Active	Yes	No	RAW	Launch More

3. Workload Manager requires PowerShell 4.0 or above on Windows 2016. Verify that you are running the correct version of PowerShell:

```
PS C:\> $PSVersionTable.PSVersion

Major Minor Build Revision
-----
4 0 -1 -1
```


4. Verify that the `C:\PROGRA~1` path is resolvable to `C:\Program Files`.

```
dir "C:\PROGRA~1"
```

Otherwise, you must create the corresponding link:

```
mklink /J "C:\PROGRA~1" "C:\Program Files"
```

5. Download CloudCenter Tools as directed by [CloudCenter Support](#).




6. Go to the command prompt window and run the downloaded file:


C:\<path to the file>\cliqr_installer.exe /CLOUDTYPE=openstack /CLOUDREGION=default
 (The other cloud options are: amazon, azure, google, openstack, vmware, custom)



If your cloud is one of the *custom* Cloud Type, you will be prompted to enter the location of your userdata and metadata extractor scripts. Contact your custom cloud provider for the script locations.



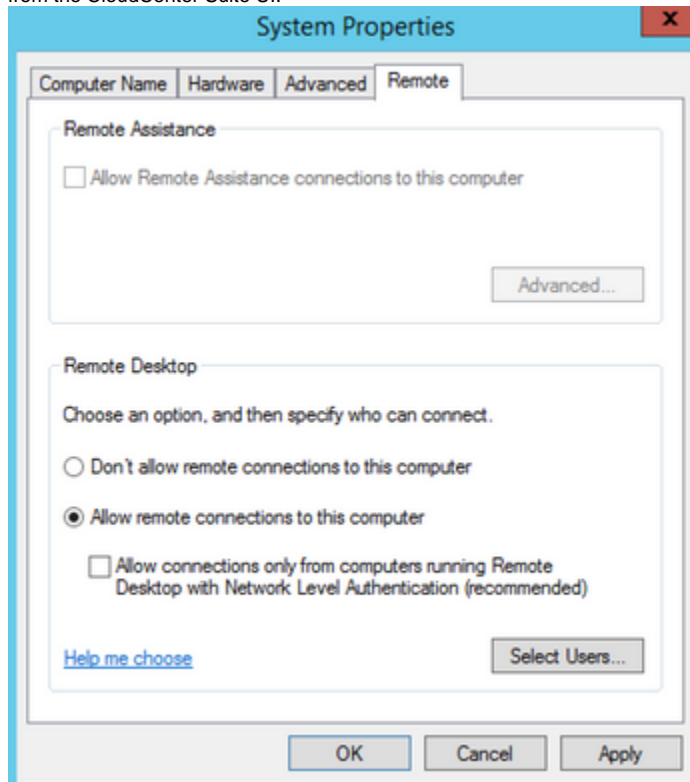
7. Run the Installer. You may optionally add the IIS installation (depends on your deployment).



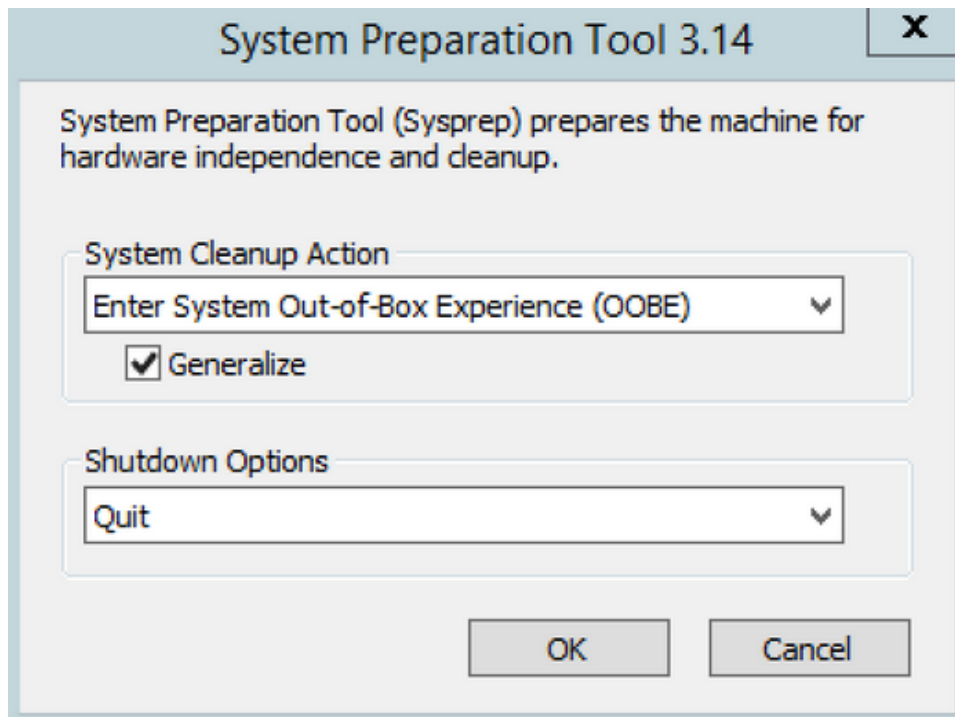
- Recently, AWS has changed its metadata script execution behavior, and custom images will not work out-of-box
- To execute the scripts while launching the image, the following extra step is required *if* configuring a Windows 2016 application VM image on AWS.
 - Before creating this image run the following scripts PowerShell.

```
C:\ProgramData\Amazon\EC2-Windows\Launch\Scripts\InitializeInstance.ps1 -Schedule
```

- Reference:
 - <https://stackoverflow.com/questions/26158411/amazon-ec2-custom-ami-not-running-bootstrap-user-data>
 - <http://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/ec2launch.html#ec2launch-config>
- The following extra steps are required *if* configuring Windows 2012 R2 or Windows 2016 application VM image in AzureRM.
 - a. Open Server Manager and start/open the **Add Roles and Features Wizard**.
 - b. Click **Next** at the **Before you begin** pane.
 - c. Select **Role-based or feature-based installation** on the **Installation type** pane and click **Next**.
 - d. Select the **Select a server from the server pool** option in the **Server Selection** pane.
 - e. In the Server Roles pane, select **Web Server (IIS)**, and add all roles under Application Development.
 - f. Click **Next** to proceed to the **Features** pane.
 - g. In the **Internet Information Services** pane, expand **Web Management Tools**, and select **IIS Management Console**.
 - h. Select **Add Features** and select:
 - i. .NET Framework 3.5 Features, add .NET Framework 3.5
 - ii. .NET Framework 4.5 Features, add ASP.NET 4.5
- Remove the installer.exe file.
- Go to **Control Panel > System and Security > System > Allow Remote Access**. Uncheck the check box to **Allow connections only from computers running Remote Desktop with Network Level Authentication**. This action allows you to RDP into the VM from the CloudCenter Suite UI.



- In the file explorer go to C:\Windows\System32\Sysprep. Double click on **sysprep**. Make sure you choose the options displayed in the image below. Once sysprep is done, your RDP session will terminate



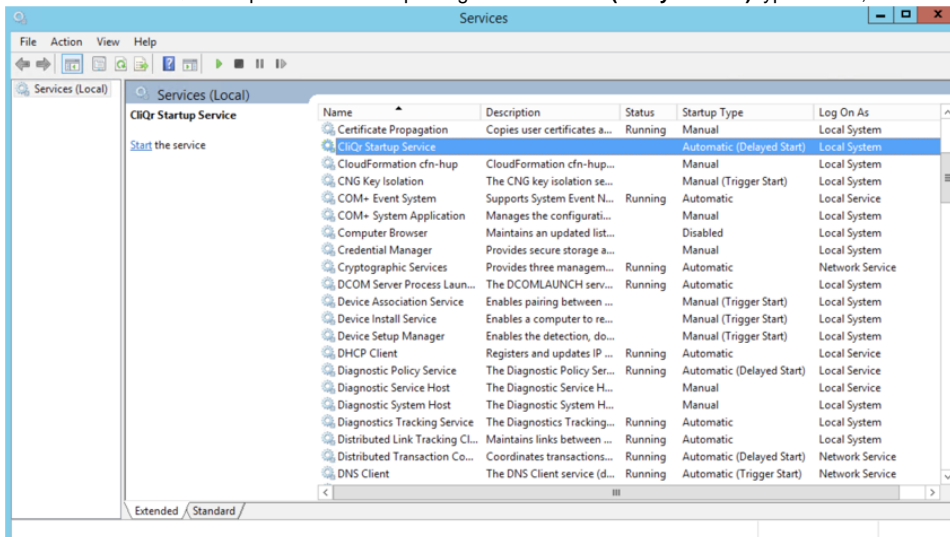
- Go to the Azure Management UI and shutdown the application VM image.

8. After installing the agent, verify the following post-installation information.



If this information remains unverified, be aware that the agent may not be successfully installed.

Ensure that the CliQr Startup Service starts up using the **Automatic (Delayed Start)** type – if not, set it to **Automatic(Delayed Start)**



9. Capture an image of the Worker VM.

10. After creating the image, use the Image ID to map the corresponding Workload Manager logical image to this physical image (Workload Manager UI > Admin > Images > Manage Cloud Mapping > Add/Edit Mapping). See [Images Page](#) for additional context).

Local Package Store (Conditional)

Local Package Store (Conditional)



1. The local package store can be created on your own base OS image using the *installer* file listed below.
2. If running the local package store on a CentOS 7 distribution is acceptable in your environment, consider deploying the [Local Repo Appliance](#), instead.
3. This component is required if worker VMs in the region do not have continuous access to the internet.

To configure a local package store, complete the following procedure.



Internet Access!

The package store requires internet access to be able to periodically synchronize OS packages from the default package store (<http://repo.ciqrtech.com>).

1. SSH into the VM instance using the key pair that you used to launch the VM.
2. Login and become the ROOT user.

```
sudo -i
```

3. Change to the location where you want to download the installer.
4. Download the `repo_installer.bin` file from software.cisco.com and save to the `/tmp` folder.
5. Run the repository (repo) installer using the following commands:

```
chmod 755 repo_installer.bin
./repo_installer.bin <ostype> <cloudtype> repo
```

Where:

<ostype>= centos6, centos7, rhel6, rhel7, ubuntu1404, ubuntu1604

<cloudtype>= *amazon, azurearm, azurepack, google, opsource, openstack, softlayer, vmware, vcd* (run the `./core_installer.bin help` command for a complete list)

For example:

```
./repo_installer.bin centos7 amazon repo
```

6. Invoke the repo config wizard and configure the basic properties.

Config Wizard Path

```
/usr/bin/repo_config_wizard.sh
```

This invokes a text UI. Dismissing the welcome message displays the configuration menu. The menu has three choices: `Proxy_Settings`, `Repo_Syncup`, and `Exit`.

7. Selecting `Proxy_Setting` brings up a new menu allowing you to specify a SOCKS proxy URL. Enter the address of your proxy server if you have one.
8. Selecting `Repo_Syncup` displays a confirmation message. Agreeing to the confirmation message causes the script to sync the local package store the latest package store at <http://repo.ciqrtech.com>.
9. Exit the Repo wizard.
10. Run the following command to verify if the repo sync is complete and the files are downloaded.

```
du -sh /repo/
5.6G  /repo/

#Approximately 5 GB or greater value is downloaded from the repo. Once the files are downloaded the repo
metadata is built for all the repositories.
```


11. From the CloudCenter Suite UI: *Admin > Clouds > Configure cloud > Regions / Details tab > Edit Cloud Settings*, set the **Agent Custom Repository** field to "https://" followed by the IP address of the local package store VM.



The local package store should be synced periodically with repo.cliqtech.com. Instead of doing this manually, consider creating a cron job to do this once every 24 hours.

Local Bundle Store (Conditional)

Local Bundle Store (Conditional)



1. The local bundle store can be created on your own base OS image using the *installer* file listed below.
2. If running the local bundle store on a CentOS 7 distribution is acceptable in your environment, consider deploying the [Local Repo Appliance](#), instead.
3. This component is required if worker VMs in the region do not have continuous access to the internet.

To configure the local bundle store, follow this procedure.

1. Set up the HTTP server.



This setup assumes Apache2 on a CentOS server. If you use a different OS/HTTP server, adjust the following commands accordingly.

2. Locate the document root of the HTTP server

- a. Change directory to **`/etc/httpd/conf`**
- b. Check `httpd.conf` for `site-available/default` files.
- c. Locate the ***DocumentRoot*** in one of these configuration files. Typically, it will be either `/var/www` or `/var/www/html`.

3. Change directory to ***DocumentRoot*** directory.

```
cd <DocumentRoot>
```

4. Create a directory to reflect the CloudCenter release you are installing (for example, 4.8.0) and create a bundle directory under the release folder level.

```
mkdir release-<CloudCenter Version>
cd release-<CloudCenter Version>
mkdir bundle
```

5. Change to the bundle directory.

```
cd bundle
```

6. Copy or download the `bundle_artifacts.zip`
7. Unzip the `bundle_artifacts.zip` file.

```
unzip bundle_artifacts.zip
```

8. From the CloudCenter Suite UI: *Admin > Clouds > Configure cloud > Regions / Details tab > Edit Cloud Settings*, set the **Agent Bundle Store** field as explained in [Local Repo Appliance \(Conditional\)](#) > Configure cloud settings.

Periodically update the local bundle store as explained in [Local Repo Appliance \(Conditional\)](#) > Periodically Update the Bundle Store.

See [Dynamic Bootstrapping](#) > HTTPS Dependencies for additional context.

Clouds

Clouds

- [Cloud Overview](#)
- [Configure a Cloud End-to-End](#)
- [Cloud Maintenance](#)
- [Manage Instance Types](#)

Cloud Overview

Cloud Overview

- [Overview](#)
- [Scope of a Cloud Region](#)
- [Minimum Permissions for Public Clouds](#)

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](#).





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 vCD – clouds, 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	<ul style="list-style-type: none"> • Accounts • Sub-Accounts • Identity and Access Management (IAM)
VMware vCenter	vCenter instance	<ul style="list-style-type: none"> • Datacenter • Clusters • Resource pools • Accounts • Datastores • Datastore clusters
VMware vCloud Director	vCD instance	<ul style="list-style-type: none"> • Datacenter • Clusters • Resource pools • Accounts • Datastores • Datastore clusters
Azure RM	Geographical Region	<ul style="list-style-type: none"> • Networks • Cloud services • Accounts
Google Cloud	Geographical Region	<ul style="list-style-type: none"> • Projects • Accounts
IBM Cloud	Geographical Region	<ul style="list-style-type: none"> • Accounts

OpenStack	Logical Region	<ul style="list-style-type: none"> • Tenants • Networks • Accounts
Kubernetes	Kubernetes cluster	<ul style="list-style-type: none"> • Accounts • Namespaces • VPCs • IAM policies
Outscale	Geographical Region	<ul style="list-style-type: none"> • Accounts • Sub-Accounts • Identity and Access Management (IAM)

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, see <https://docs.aws.amazon.com/awscostexplorer/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*	<i>Cost management reader</i>	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*	<i>Billing reader</i>	storage.objects.get,list storage.buckets.get,list
		 AWS Cost Explorer must be enabled to view AWS-specific costs on Cost Optimizer.		
Cost Optimizer and Workload Manager	Collect VMs and volumes	ec2:DescribeAvailabilityZones ec2:DescribeAddresses ec2:DescribeInstances ec2:DescribeVolumes ec2:DescribeTags tag:getTagKeys tag:getTagValues	VM: <i>VM contributor</i> Volume: <i>Reader</i>	compute.instances.get,list compute.disks.get,list
		 The <i>Reader</i> role must be offered because no built-in role is provided.		
		 <ul style="list-style-type: none"> • 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 are required for tag-based reporting and only applicable to Cost Optimizer. 		
Cost Optimizer	Collect PAAS services	rds:Describe* elasticloadbalancing:Describe*	SQL Server and SQL database: <i>SQL Server contributor</i> MySQL and PostgreSQL Server: <i>Reader</i>	cloudsql.databases.get,list cloudsql.instances.get,list compute.forwardingRules.get,list compute.targetPools.get,list
		 The <i>Reader</i> role must be offered because no built-in role is provided.		

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
Cost Optimizer and Workload Manager	Collect data for AWS member account	To allow a primary account to collect data on behalf of member accounts, the following is necessary: <ul style="list-style-type: none"> 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> <p>Effective from Cost Optimizer 5.4, the AWS <i>OrganizationAccountAccessRole</i> is supported on member accounts.</p> <p>Alternatively, you can create a role named Optimizer and do the following to the new role:</p> <ul style="list-style-type: none"> 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>	N/A	
Workload Manager	Manage VMs and volumes	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	Offer the <i>Italicized</i> roles to create, modify, or delete: <ul style="list-style-type: none"> NICs, Public IPs and security group: <i>Network Contributor</i> Diagnostics: <i>Storage Account Contributor</i> Unmanaged data disk: <i>Storage Account Contributor</i> Managed data disks: <i>Owner</i> VMs with managed data disks: <i>Owner</i> VMs with unmanaged data disks and diagnostic logs: <i>Virtual Machine Contributor</i>, <i>Network Contributor</i>, and <i>Storage Account Contributor</i> VMs with no data disks: <i>Virtual Machine Contributor</i> and <i>Network Contributor</i> <div style="border: 1px solid orange; padding: 5px; margin-top: 10px;"> In some cases, the <i>Owner</i> role must be offered because no built-in role is provided. </div>	Use the pre-defined <i>Project Editor</i> 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.networks.get,list,use compute.projects.get

	ec2:DeleteNetworkInterface		compute.regions.get
	ec2:DeleteSecurityGroup		compute.subnetworks.get,list,use,useExternalIp
	ec2:DeleteSnapshot		compute.zones.get
	ec2:DeleteTags		
	ec2:DeleteVolume		iam.serviceaccounts.get,list
	ec2:DescribeAccountAttributes		
	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 a Cloud End-to-End

Configure a Cloud End-to-End

- [Configure a vCenter Cloud](#)
- [Configure an AWS Cloud](#)
- [Configure an AzureRM Cloud](#)
- [Configure a Google Cloud](#)
- [Configure an OpenStack Cloud](#)
- [Configure a Kubernetes Cloud](#)
- [Configure a vCD Cloud](#)
- [Configure an IBM Cloud](#)
- [Configure an Outscale Cloud](#)

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](#)

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.

* VCENTER API ENDPOINT ⓘ

NEXT



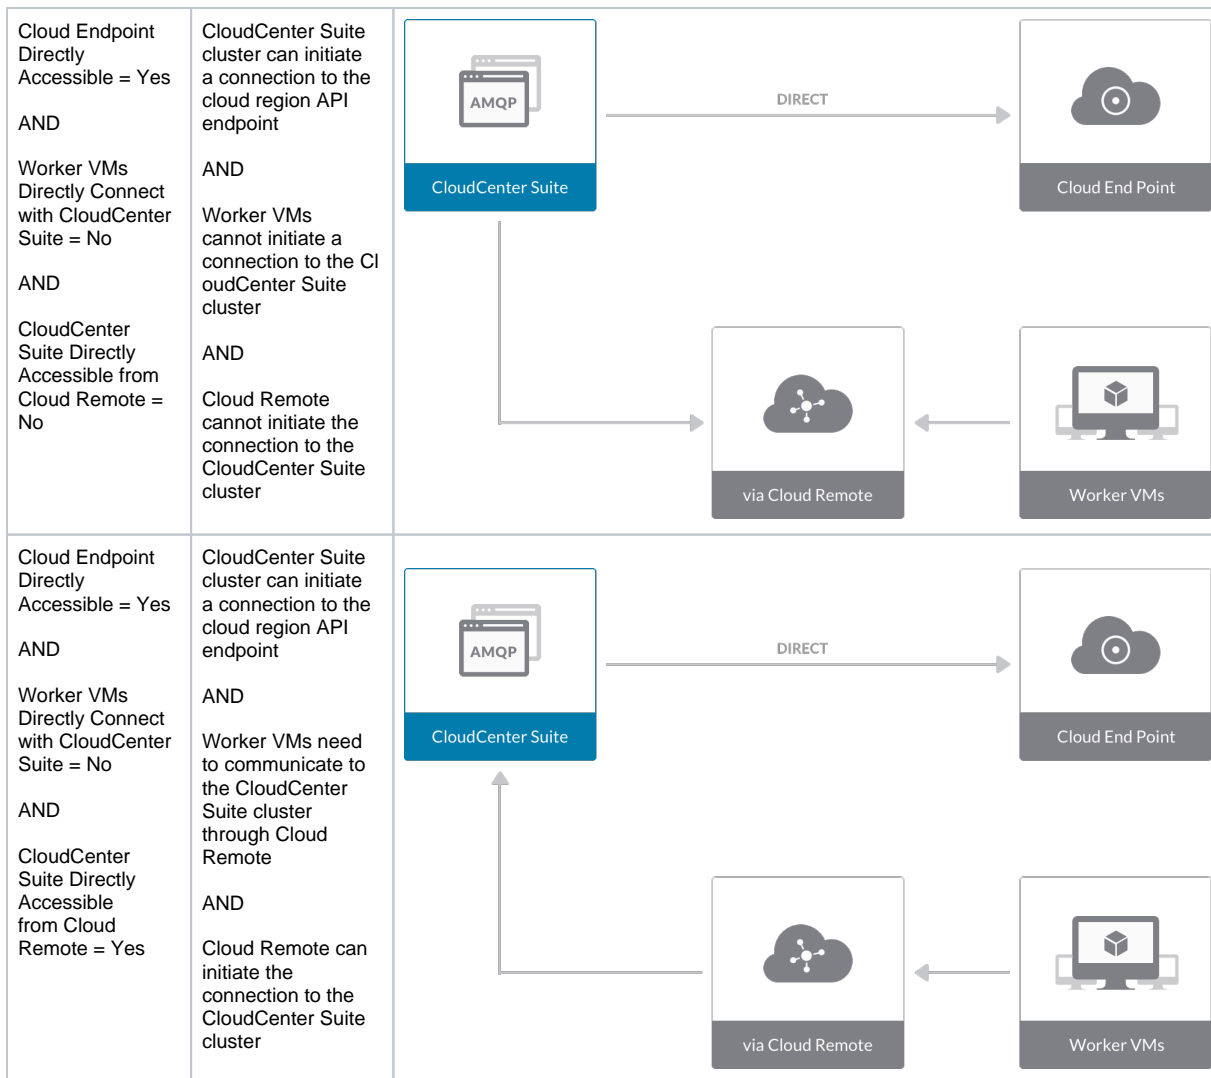
Note

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
-----------------	----------	-----------------

<p>Cloud Endpoint Directly Accessible = Yes</p> <p>AND</p> <p>VMs Directly Connect with CloudCenter Suite = Yes</p>	<p>CloudCenter Suite cluster can initiate a connection to the cloud region API endpoint</p> <p>AND</p> <p>Worker VMs can initiate a connection to the CloudCenter Suite cluster</p> <p>Cloud Remote is not required</p>	<p>The diagram shows three main components in boxes: 'CloudCenter Suite' (with an AMQP icon), 'Cloud End Point' (with a cloud icon), and 'Worker VMs' (with a computer icon). A horizontal arrow labeled 'DIRECT' points from CloudCenter Suite to Cloud End Point. A vertical arrow labeled 'DIRECT' points from Worker VMs up to CloudCenter Suite.</p>
<p>Cloud Endpoint Directly Accessible = No</p> <p>AND</p> <p>Worker VMs Directly Connect with CloudCenter Suite = No</p> <p>AND</p> <p>CloudCenter Suite Directly Accessible from Cloud Remote = Yes</p>	<p>CloudCenter Suite cluster cannot initiate a connection to the cloud region API endpoint</p> <p>AND</p> <p>Worker VMs cannot initiate a connection to the CloudCenter Suite cluster</p> <p>AND</p> <p>Cloud Remote can initiate the connection to the CloudCenter Suite cluster</p>	<p>The diagram shows three main components in boxes: 'CloudCenter Suite' (with an AMQP icon), 'via Cloud Remote' (with a cloud icon), and 'Cloud End Point' (with a cloud icon). A horizontal arrow points from 'via Cloud Remote' to CloudCenter Suite. Another horizontal arrow points from 'via Cloud Remote' to Cloud End Point. A vertical arrow points from Worker VMs up to 'via Cloud Remote'.</p>
<p>Cloud Endpoint Directly Accessible = No</p> <p>AND</p> <p>Worker VMs Directly Connect with CloudCenter Suite = No</p> <p>AND</p> <p>CloudCenter Suite Directly Accessible from Cloud Remote = No</p>	<p>CloudCenter Suite cluster cannot initiate a connection to the cloud region API endpoint</p> <p>AND</p> <p>Worker VMs cannot initiate a connection to the CloudCenter Suite cluster</p> <p>AND</p> <p>Cloud Remote cannot initiate the connection to the CloudCenter Suite cluster</p>	<p>The diagram shows three main components in boxes: 'CloudCenter Suite' (with an AMQP icon), 'via Cloud Remote' (with a cloud icon), and 'Cloud End Point' (with a cloud icon). A horizontal arrow points from CloudCenter Suite to 'via Cloud Remote'. Another horizontal arrow points from 'via Cloud Remote' to Cloud End Point. A vertical arrow points from Worker VMs up to 'via Cloud Remote'.</p>



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.

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.

- 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: <ul style="list-style-type: none"> "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	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.
Windows Max Sockets	The platform attempts to use the maximum number of sockets during deployment as well as when resizing instance types. During an application deployment: <ul style="list-style-type: none"> 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 does not use the Max Sockets setting when 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 a local 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 HTTPS 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.

Region Connectivity Running Download Configuration [Configure Region](#)

Clicking **Download Configuration** causes two things to happen:

- An encrypted zip file named artifacts.zip is 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.

Region Connectivity Enabling... [Download Configuration](#) Copy Encryption Key [Edit Connectivity](#)

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 from software.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 in [Proxy 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 with cloud (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 both Cloud 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**.

- 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 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.

The exception to this guidance is when a NAT firewall or proxy server exists between the CloudCenter Suite management cluster and worker 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 the Guacamole 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 the Guacamole server running in the CloudCenter Suite cluster.
Worker VMs must use a proxy server or NAT firewall to access the Guacamole 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.
- If any of the connectivity toggles in the 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

- From your local computer, download the Cloud Remote appliance OVA from software.cisco.com.
- 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.
- 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.
- 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.
- Complete the fields for Name and location, Host / Cluster, Resource Pool, Storage, and Disk Format appropriate for your environment.
- 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.
- 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.
- Confirm your settings and click Finish to launch the VM.
- 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.

- j. Once the first instance of the appliance has been launched, use the vSphere client 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) > <i>Custom Port Numbers (Conditional)</i>). 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) > <i>Custom Port Numbers (Conditional)</i>).
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) > <i>Custom Port Numbers (Conditional)</i>).
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 Download Configuration 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.

Region Connectivity Enabling... [Download Configuration](#) [Copy Encryption Key](#) [Edit Connectivity](#)


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 to Cloud Remote web UI.

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



- 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.



- Paste the encryption key that was copied to the clipboard into the **Encryption Key** field in the dialog box.
- Click **Select File** and browse to the artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
- Click **Confirm**.
- 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).

Region Connectivity	Running	Download Configuration	Configure Region
Cloud endpoint accessible from Cloud Center Manager	No		
Cloud Center Manager AMQP reachable from worker VM's	No		
Cloud Center Manager AMQP accessible from cloud	Yes		
Remote AMQP IP			
Worker AMQP IP	192.168.30.16:5671		
Blade Name	cloudcenter-blade-vmware-9-0289		
Blade Port	8443		

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

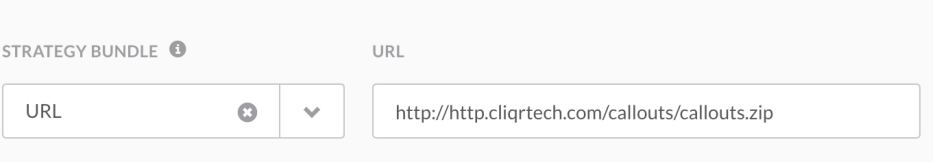
Workload Manager lets you define a VM naming strategy and an IP address management (IPAM) strategy on a per cloud region basis. The VM naming strategy lets Workload Manager generate unique names for each of your VMs as they are launched. The IPAM strategy lets Workload Manager generate IP addresses for 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 and an IP address deallocation callout script. An IP address allocation callout script may 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 and 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 CliQr 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 the Instance Naming Strategy field. 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 must meet the following requirements for job submission to succeed: <ul style="list-style-type: none"> Only contain the following characters: <ul style="list-style-type: none"> Lowercase a to z 0 to 9 Hyphen Start with an 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 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 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 options 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 your custom IPAM address deallocation script. See 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 for CliQr Macro Replacement:

Macro	Returned Value
%os1%	First character of OS type string in lower case
%os2%	First two characters of OS type string in lower case
%OS1%	First character of OS type string in upper case
%OS2%	First two 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
- Post VM 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 execute in 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) to specify external lifecycle actions at the region level.

External Lifecycle Actions	Edit External Lifecycle Actions
External Actions Bundle	
Pre VM Start	-
Pre VM Init	-
Post VM Init	-
Pre VM Stop	-
Post VM Stop	-

Click on the **Edit External Lifecycle Actions** link in the upper right. This displays the Edit 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 figure below.

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, see [Deployment Lifecycle Scripts](#).

- 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.

Instance Types [Sync Instance Types](#) [Add Instance Type](#)

Show 30 per page Page 1 of 1

Name	Instance Type	Price	Actions
m1.small	0	\$0/hr	Edit Delete
m1.large	0bbdff46-a647-4eab-827f-7367e1e6...	\$0/hr	Edit Delete
m2.small	809193e3-fb66-49a2-9c55-d318d2a...	\$0/hr	Edit Delete
m1.nano	ab1365a2-debc-44b1-b4cf-3d5dd15c...	\$0/hr	Edit Delete

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.

Instance Types [Edit Instance Types](#) [Add Instance Type](#)

Show 30 per page Page 1 of 1

Name	Instance Type	Price	Actions
default	default	\$0/hr	Delete
2CPU_2048MBMEM	2CPU_2048MBMEM	\$0/hr	Delete
1CPU_2048MBMEM_u	1CPU_2048MBMEM	\$0.002/hr	Delete
2CPU_8192MBMEM	2CPU_8192MBMEM	\$0/hr	Delete
6CPU_4096MBMEM	6CPU_4096MBMEM	\$0/hr	Delete
6CPU_6144MBMEM	6CPU_6144MBMEM	\$0/hr	Delete
8CPU_24288MBMEM	8CPU_24288MBMEM	\$0/hr	Delete
4CPU_8192MBMEM	4CPU_8192MBMEM	\$0.008/hr	Delete
2CPU_4096MBMEM	2CPU_4096MBMEM	\$0.004/hr	Delete

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

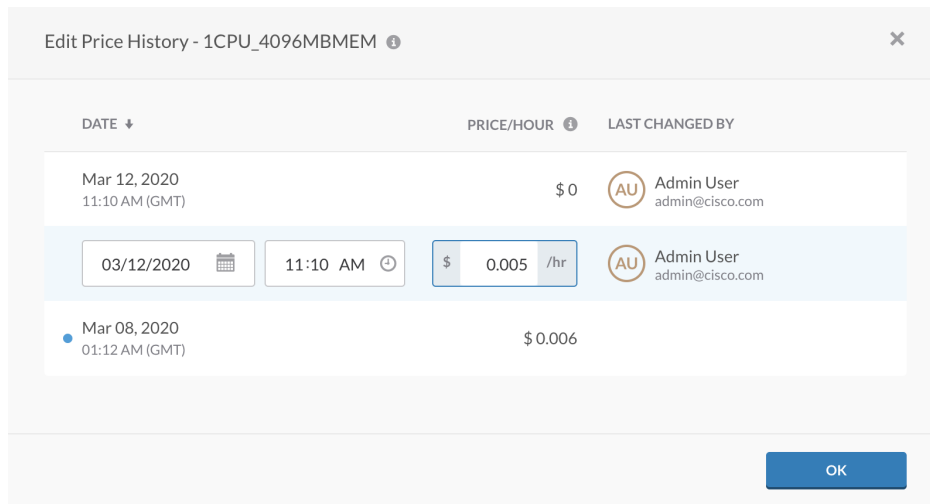
For all cloud providers except AWS, you can edit the fields of an instance type (except instance type ID) by clicking the **Edit** link in the Actions column in the list of instance types. For vCenter cloud changing the values of parameters such as CPU and RAM will impact the properties of the VMs or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

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



- c. Specify a date and time in the appropriate fields and enter the desired price in the **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 all resources (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 (see [Cost Optimizer Dashboard](#)), as shown in the following sample screenshot.

● VMware vmware_usdc-default - cost adjusted successfully

Price change on instance type: 3CPU_2048MBMEM has been reflected on the deployments and resource costs.

a minute ago

Simultaneously, the updated accrued cost of a VM is displayed on the [Virtual Machines](#) page 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 the package 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 default storage type 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.

The 'Add Storage Type' dialog box contains the following fields and controls:

- Display Name ***: A text input field.
- Cost Per Month ***: A numeric input field with a '\$' symbol on the left and '/GB' on the right. The value is '0.00'.
- Cloud Storage Type ID ***: A text input field.
- Minimum Volume Size ***: A numeric input field with a 'GB' unit on the right. The value is '4'.
- Maximum Volume Size ***: A numeric input field with a 'GB' unit on the right. The value is '1024'.
- Maximum Throughput ***: A numeric input field with a 'MB/s' unit on the right. The value is '100'.
- Maximum IOPS ***: A numeric input field. The value is '100'.
- Provisioned IOPS**

At the bottom right of the dialog box are two buttons: 'Save' (blue) and 'Cancel' (light blue).

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 the [Cost 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.

Storage Types [Add Storage Type](#)

Search: Show 30 per page Page 1 of 1

Name	Storage Type	Price	Actions
Local SSD scratch disk	local-ssd	\$0.218/GB-month	Edit Delete
Standard Persistent Disk	pd-standard	\$0.04/GB-month	Edit Delete
SSD Persistent Disk	pd-ssd	\$0.17/GB-month	Edit Delete

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:

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

Edit Price History - default ✕

DATE ↓	PRICE / GB-MONTH ⓘ	LAST CHANGED BY
Mar 12, 2020 09:49 AM (GMT)	\$ 22	Admin User admin@cisico.com
<input type="text" value="03/11/2020"/> <input type="text" value="3 :06 PM"/>	<input type="text" value="\$ 21"/> <input type="text" value="/GB"/>	Admin User admin@cisico.com
Mar 08, 2020 06:16 PM (GMT)	\$ 1	Admin User admin@cisico.com
Mar 08, 2020 01:13 AM (GMT)	\$ 1.1	Admin User admin@cisico.com
Mar 08, 2020 01:12 AM (GMT)	\$ 22	Admin User admin@cisico.com

[OK](#)

- Specify a date and time in the appropriate fields and enter the desired price in the **Price/Hour** field, enter the desired price.
- 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 all resources (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 (see [Cost Optimizer Dashboard](#)) in the header. Simultaneously, the updated accrued cost of a VM is displayed on the [Storage Volumes](#) page and the updated deployment cost is displayed on the Deployment Details page.

- 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. See [Images](#) for more context.

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 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 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.

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.

Account Name	Description	Billing Units	Enabled For	Actions
C3 Manual 1	C3 Manual Account 1	2 Billing Units	Provisioning, Reporting	Edit Delete
Master	Cost Optimizer Reporting	11 Billing Units	Reporting	Edit Delete
Account		050	Provisioning, Reporting	Edit Delete
C3 Manual Plans		810	Provisioning, Reporting	Edit Delete

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, or [manage instance types](#) for 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](#)

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** link in 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	<p>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</p> <p>Cloud Remote is not required</p>	<p>The diagram illustrates a direct connection path. On the left, a box labeled 'CloudCenter Suite' contains an 'AMQP' icon. An arrow labeled 'DIRECT' points from the CloudCenter Suite to a box on the right labeled 'Cloud End Point' with a cloud icon. Below this, another box labeled 'Worker VMs' with a computer icon has an arrow labeled 'DIRECT' pointing back up to the CloudCenter Suite.</p>



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.

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 the region you added shown on the right side of the page.

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 the Cloud Settings section displayed first.

The screenshot shows the 'aws1' cloud account configuration page. At the top, there is a 'Back to Clouds' link. Below it, a 'Not Ready' warning indicates that no cloud account has been added. The main content area is divided into 'Regions' and 'Accounts' tabs. The 'US East (Ohio)' region is selected, and its 'Cloud Settings' are displayed. The settings include a 'Region Endpoint' set to 'ec2.us-east-2.amazonaws.com' and a list of other fields like 'Agent Bundle URL', 'Agent Custom Repository', and various proxy settings.

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.

- 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 a local 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 HTTPS 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.



Region Connectivity Running Download Configuration Configure Region

Clicking **Download Configuration** causes two things to happen:

- An encrypted zip file named artifacts.zip is 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.



Region Connectivity Enabling... Download Configuration Copy Encryption Key Edit Connectivity

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 from software.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 in [Proxy 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 with cloud (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 both Cloud Remote and non-Cloud Remote modes. Cloud Remote mode is not applicable to AWS but supports all proxy sections in non-Cloud Remote mode.

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

- 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 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. The exception to this guidance is when a NAT firewall or proxy server exists between the CloudCenter Suite management cluster and worker 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 the Guacamole 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 the Guacamole server running in the CloudCenter Suite cluster.
Worker VMs must use a proxy server or NAT firewall to access the Guacamole 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.

- If any of the connectivity toggles in the **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

- Obtain the Cloud Remote shared AMI from Cisco support and launch it. Follow the same guidance for obtaining and launching the [CloudCenter Suite installer appliance for AWS](#).
- 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.
- 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 Download Configuration 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 to Cloud Remote web UI.

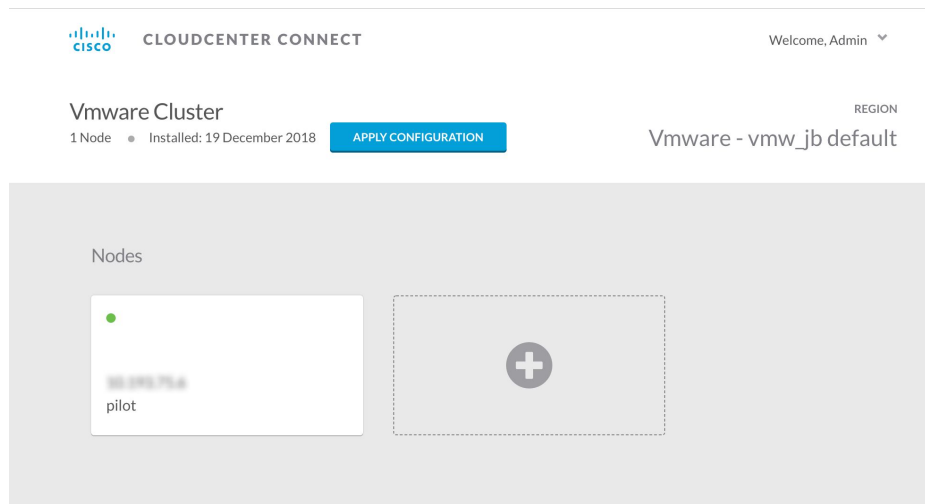
- Open another browser tab and login to https://<Cloud Remote_ip> with the default credentials: admin/cisco.
- You will immediately be required to change your password. Do so now.
- You are now brought to the Cloud Remote home page as shown in the figure below.




- 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.



- Paste the encryption key that was copied to the clipboard into the **Encryption Key** field in the dialog box.
- Click **Select File** and browse to the artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
- Click **Confirm**.
- 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.

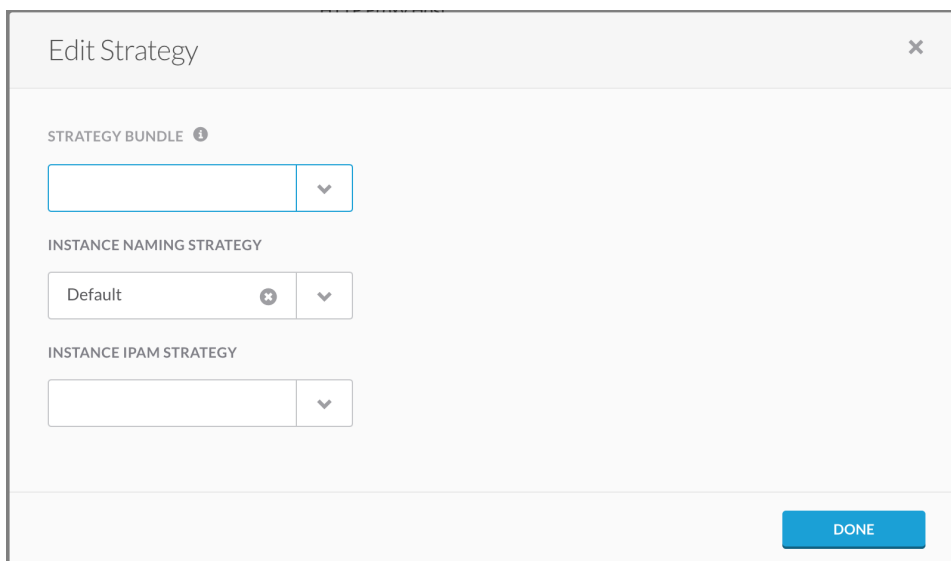
VM Naming and IPAM Strategies

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 lets Workload 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 and an IP address deallocation callout script. An IP address allocation callout script may work in conjunction with a 3rd party IPAM server in your cloud environment.

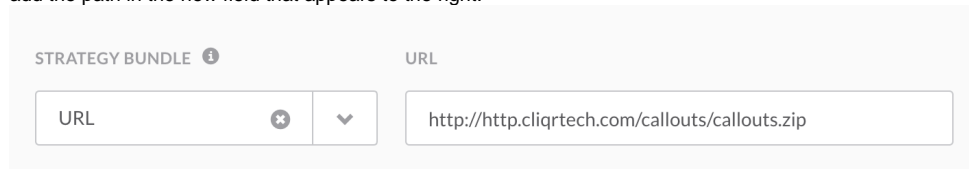
Like External Lifecycle Actions, the scripts or commands associated with a VM naming strategy or IPAM strategy are executed in a containerized script execution engine running in the CloudCenter Suite cluster, if you are not using Cloud Remote in your region. If you are using Cloud Remote, these scripts 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 and 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 CliQr 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 the Instance Naming Strategy field. 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 must meet the following requirements for job submission to succeed: <ul style="list-style-type: none"> • Only contain the following characters: <ul style="list-style-type: none"> • Lowercase a to z • 0 to 9 • Hyphen • Start with an 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 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 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 options 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 your custom IPAM address deallocation script. See 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 for CliQr Macro Replacement:

Macro	Returned Value
%os1%	First character of OS type string in lower case
%os2%	First two characters of OS type string in lower case
%OS1%	First character of OS type string in upper case
%OS2%	First two 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
- Post VM 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 execute in 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) to specify external lifecycle actions at the region level.

External Lifecycle Actions	Edit External Lifecycle Actions
External Actions Bundle	
Pre VM Start	-
Pre VM Init	-
Post VM Init	-
Pre VM Stop	-
Post VM Stop	-

Click on the **Edit External Lifecycle Actions** link in the upper right. This displays the Edit 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 figure below.

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, see [Deployment 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.

Instance Types [Sync Instance Types](#) [Add Instance Type](#)

Show 30 per page Page 1 of 1

Name	Instance Type	Price	Actions
m1.small	0	\$0/hr	Edit Delete
m1.large	0bbdff46-a647-4eab-827f-7367e1e6...	\$0/hr	Edit Delete
m2.small	809193e3-fb66-49a2-9c55-d318d2a...	\$0/hr	Edit Delete
m1.nano	ab1365a2-debc-44b1-b4cf-3d5dd15c...	\$0/hr	Edit Delete



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 must manually import the corresponding physical images into your region and then map the corresponding logical images to these physical images. See [Images](#) for more context.

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 the Add Cloud Account link for this cloud. This displays the Add Cloud Account dialog box,

The screenshot shows the 'Add Cloud Account' dialog box with the following fields and labels:

- Name ***: A text input field.
- Description**: A text area input field.
- Cloud Credentials**: A section header.
- AWS Email Address ***: A text input field with the example 'name@example.com'. Below it is the text 'Email address associated with your AWS account'.
- AWS Account Number ***: A text input field with the example 'your account number'. Below it is the text '12-digit number located at the top of your AWS account profile'.
- AWS Access Key ***: A text input field with the example 'your key'. Below it is the text '20 character key located in your security credentials'.
- AWS Secret Access Key ***: A text input field with the example 'your secret key'. Below it is the text '40 character key located in your security credentials'.
- At the bottom right are 'Save' and 'Cancel' buttons.

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, see <https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/billing-reports-gettingstarted-s3.html>.

Add Cloud Account

Email address associated with your AWS account

AWS Account Number *

12-digit number located at the top of your AWS account profile

AWS Access Key *

20 character key located in your security credentials

AWS Secret Access Key *

40 character key located in your security credentials

Billing

S3 Bucket Region

S3 Bucket Name

Report Path Prefix

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. <div style="margin-top: 10px; border: 1px solid #ccc; padding: 5px; background-color: #fff9c4;"> It is recommended that you do not add a <i>Reporting</i> account to the same tenant through different cloud groups. </div> <div style="margin-top: 10px; border: 1px solid #ccc; padding: 5px; background-color: #fff9c4;"> Enabling a public cloud account for <i>Reporting</i> may incur expenses to retrieve cost data. These expenses are proportional to the number of configured cloud accounts and regions. </div>
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/awscostexplorer/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.

Account Name	Description	Billing Units	Enabled For	Actions
C3 Manual 1	C3 Manual Account 1	2 Billing Units	Provisioning, Reporting	Edit Delete
Master	Cost Optimizer Reporting	11 Billing Units	Reporting	Edit Delete
Account		050	Provisioning, Reporting	Edit Delete
C3 Manual Plans		810	Provisioning, Reporting	Edit Delete

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, or [manage instance types](#) for the cloud account.

Configure an AzureRM Cloud

Configure an AzureRM Cloud

Configuring an AzureRM cloud is a four-step process:

- [Add an AzureRM Cloud](#)
- [Add an AzureRM Region](#)
- [Configure an AzureRM Region](#)
- [Add an AzureRM Cloud Account](#)

To add an AzureRM cloud, follow these steps.

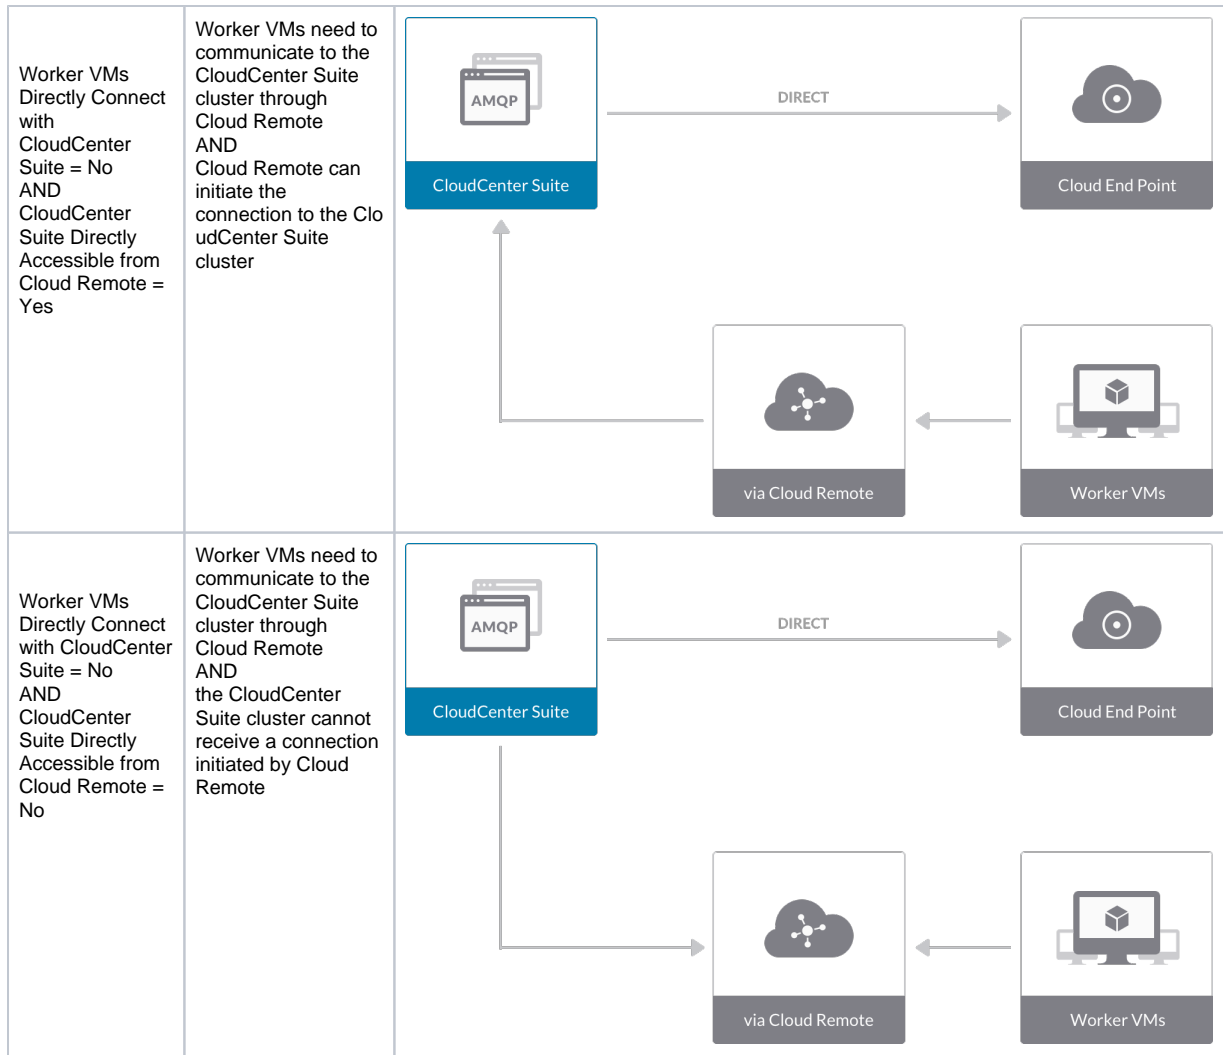
1. Navigate to **Admin > Clouds**. This brings you to the Clouds page. If you, or another tenant admin in your tenant, have already added clouds to your tenant, they will be listed here. Click the **Add Cloud** link in the upper right.
2. Click **Add Cloud**. The **Add Cloud** dialog box is displayed.
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	<p>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</p> <p>Cloud Remote is not required</p>	<p>The diagram illustrates direct connectivity. On the left is a box labeled 'CloudCenter Suite' containing an 'AMQP' icon. On the right is a box labeled 'Cloud End Point' with a cloud icon. A horizontal arrow labeled 'DIRECT' points from the CloudCenter Suite to the Cloud End Point. Below the CloudCenter Suite box is a box labeled 'Worker VMs' with a computer icon. A horizontal arrow labeled 'DIRECT' points from the Worker VMs to the CloudCenter Suite.</p>



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.

After creating an AzureRM 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 on the 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 the region you added shown on the right side of the page.

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 the Cloud 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.

- 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 a local 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 HTTPS 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.

Region Connectivity Running Download Configuration Configure Region

Clicking **Download Configuration** causes two things to happen:

- An encrypted zip file named artifacts.zip is 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.

Region Connectivity Enabling... Download Configuration Copy Encryption Key Edit Connectivity

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 from software.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 in [Proxy 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 with cloud (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 both Cloud Remote and non-Cloud Remote modes. Cloud Remote mode is 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**.

- 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 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.

The exception to this guidance is when a NAT firewall or proxy server exists between the CloudCenter Suite management cluster and worker 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 the Guacamole 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 the Guacamole server running in the CloudCenter Suite cluster.
Worker VMs must use a proxy server or NAT firewall to access the Guacamole 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 the connectivity toggles in the **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

- Download the Cloud Remote appliance for AzureRM as a zip file from software.cisco.com and then unzip it to reveal the VHD file.
- 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 the [CloudCenter Suite installer appliance for AzureRM](#).



You must use the AzureRM CLI to perform this upload.

- 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.
- Once the first instance of the appliance has been launched, use the AzureRM 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 Download Configuration 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.

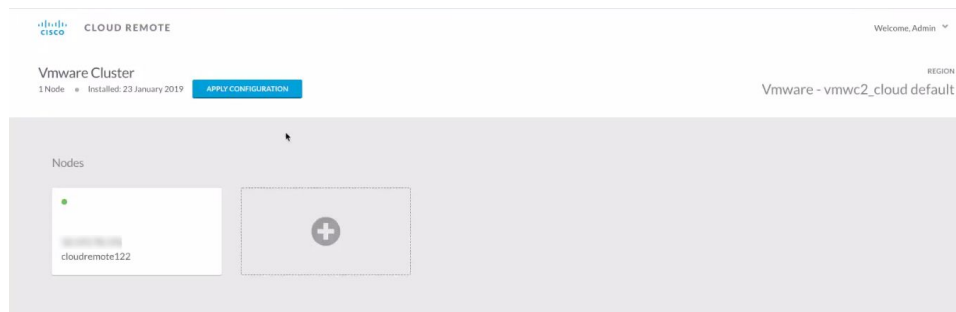


Click the **Copy Encryption Key** link to save the key to your clipboard. A success message will be displayed temporarily above the Region Connectivity section header. Make sure not to overwrite the clipboard with other data. You will need the key when you upload the configuration zip file to Cloud Remote.

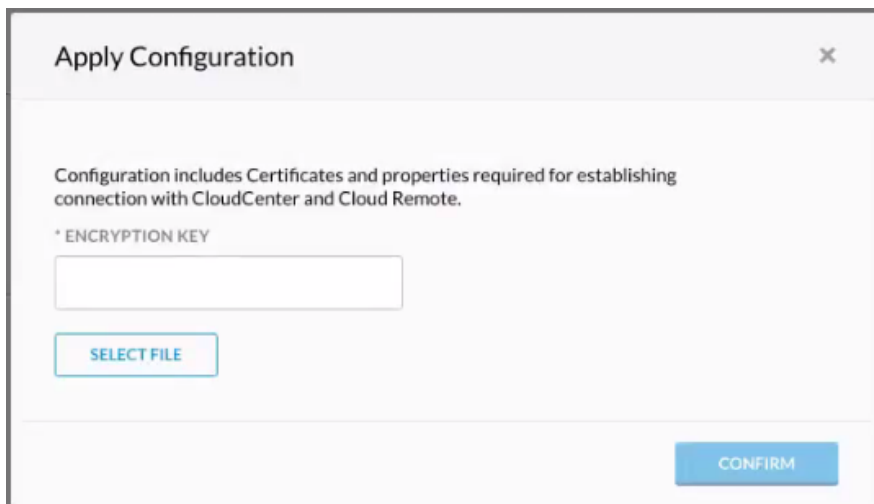
If you change the connectivity settings in the CloudCenter Suite UI and need to again download the zip file, a new encryption key is automatically created and can be copied to the clipboard by clicking the **Copy Encryption Key** link again.

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

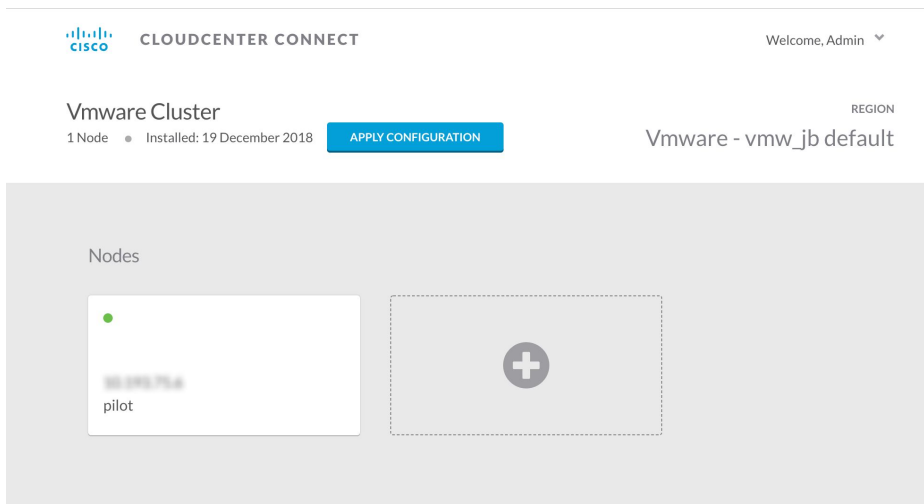
- Open another browser tab and login to https://<Cloud_Remote_ip> with the default credentials: admin/cisco.
- You will immediately be required to change your password. Do so now.
- You are now brought to the Cloud Remote home page as shown in the figure below.




- 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 to the 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.

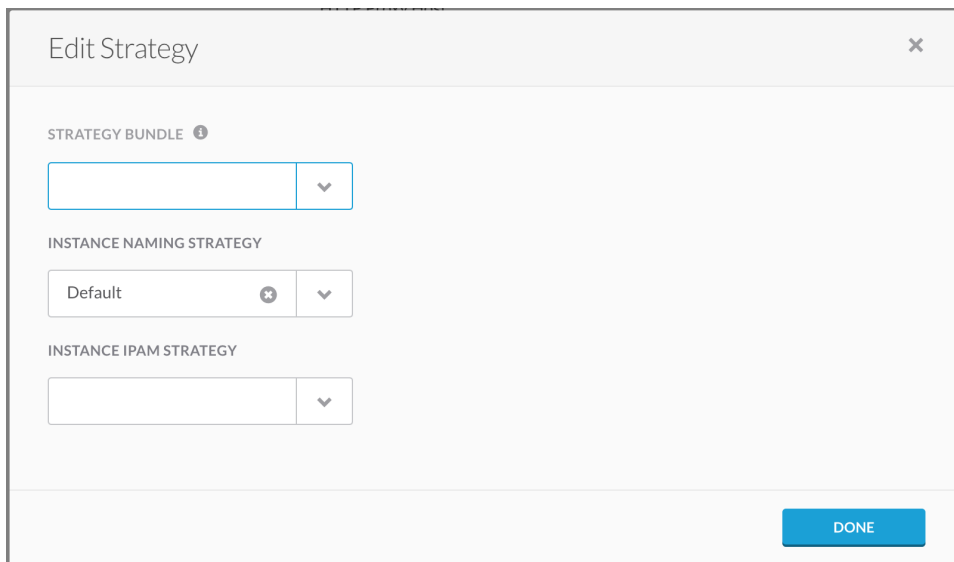
VM Naming and IPAM Strategies

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 lets Workload 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 and an IP address deallocation callout script. An IP address allocation callout script may work in conjunction with a 3rd party IPAM server in your cloud environment.

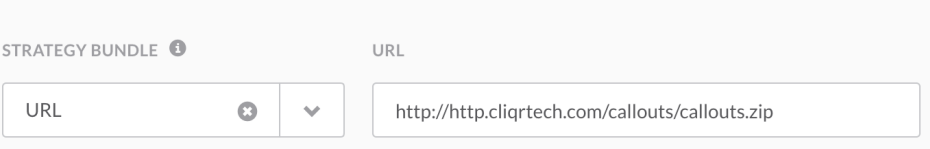
Like External Lifecycle Actions, the scripts or commands associated with a VM naming strategy or IPAM strategy are executed in a containerized script execution engine running in the CloudCenter Suite cluster, if you are not using Cloud Remote in your region. If you are using Cloud Remote, these scripts 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 and 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 CliQr 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 the Instance Naming Strategy field. 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 must meet the following requirements for job submission to succeed: <ul style="list-style-type: none"> Only contain the following characters: <ul style="list-style-type: none"> Lowercase a to z 0 to 9 Hyphen Start with an 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 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 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 options 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 your custom IPAM address deallocation script. See 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 for CliQr Macro Replacement:

Macro	Returned Value
%os1%	First character of OS type string in lower case
%os2%	First two characters of OS type string in lower case
%OS1%	First character of OS type string in upper case
%OS2%	First two character of OS type string in upper case
%RND<number>%	Fixed length random string, e.g., %RND10%, 10-character long random string <div style="border: 1px solid #ffc107; padding: 5px; margin-top: 5px;">  This is a mandatory field, only values in the range: %RND6% to %RND11% or just %RND% which defaults to %RND6% </div>
%UID%	User ID
%VID%	Vendor ID
%VM_PREFIX%	VM name prefix must be added as a global parameter in the app profile. <div style="border: 1px solid #ffc107; padding: 5px; margin-top: 5px;">  Only macro whose value can be user-customized. </div>

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
- Post VM 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 execute in 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) to specify external lifecycle actions at the region level.

External Lifecycle Actions		Edit External Lifecycle Actions
External Actions Bundle	-	
Pre VM Start	-	
Pre VM Init	-	
Post VM Init	-	
Pre VM Stop	-	
Post VM Stop	-	

Click on the **Edit External Lifecycle Actions** link in the upper right. This displays the Edit 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 figure below.

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, see [Deployment Lifecycle Scripts](#).

- 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.

- 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 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 the package 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.

Add Storage Type

Display Name *

Cost Per Month *

\$ 0.00 /GB

Cloud Storage Type ID *

Minimum Volume Size *

4 GB

Maximum Volume Size *

1024 GB

Maximum Throughput *

100 MB/s

Maximum IOPS *

100

Provisioned IOPS

Save Cancel

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 the [Cost 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.

Storage Types [Add Storage Type](#)


Show 30 per page Page 1 of 1

Name	Storage Type	Price	Actions
Local SSD scratch disk	local-ssd	\$0.218/GB-month	Edit Delete
Standard Persistent Disk	pd-standard	\$0.04/GB-month	Edit Delete
SSD Persistent Disk	pd-ssd	\$0.17/GB-month	Edit Delete

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 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 must manually import the corresponding physical images into your region and then map the corresponding logical images to these physical images. See [Images](#) for more context.

 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 valid [Windows 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 Manager Portal**, 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 the AzureRM 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.



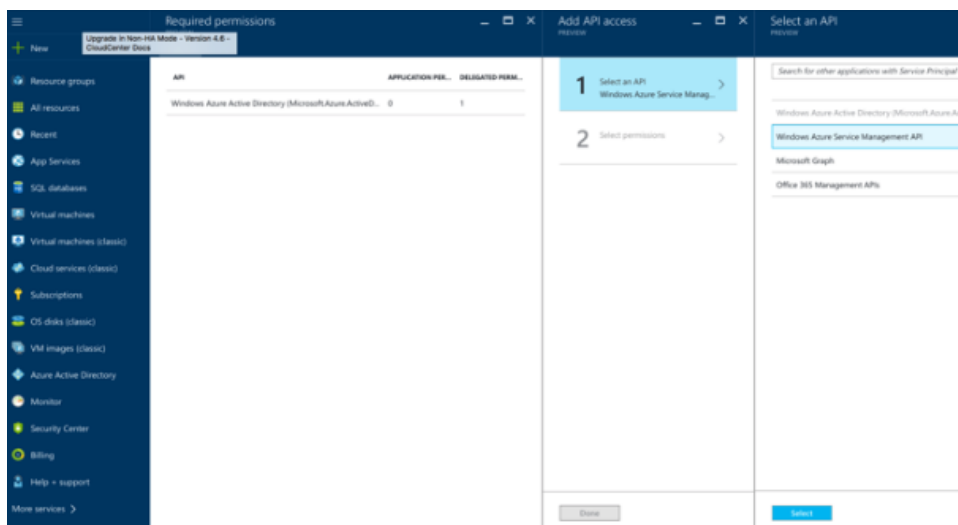
Note down the Application ID; it is required to create a Cloud Account in CloudCenter – this is the Client ID.

If you prefer to use *Certificate-Based Authentication*, see the related bullet further in this section.


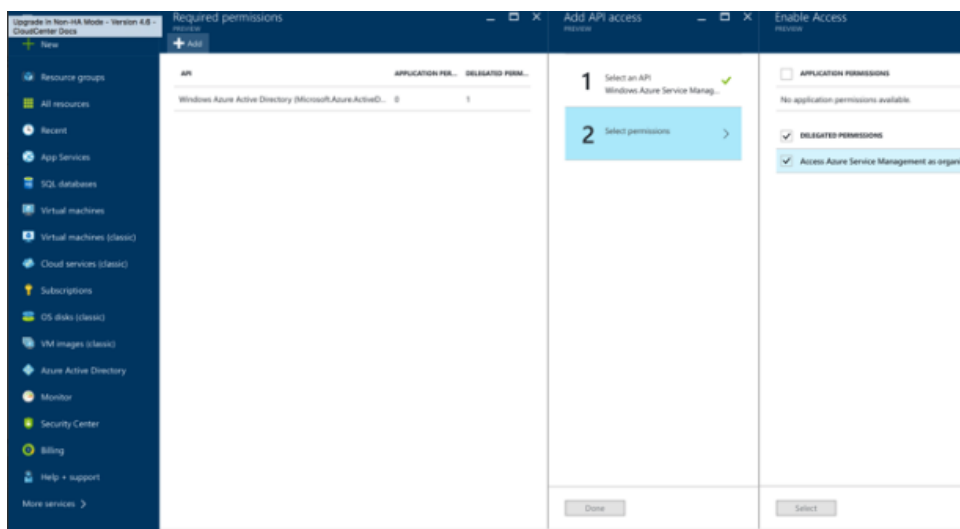
4. Click **All Settings**.
5. Select **Required Permission** under API Access and click **Add**. See [Cloud Overview](#) > *Minimum Permissions for Public Clouds* for additional details.



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**.



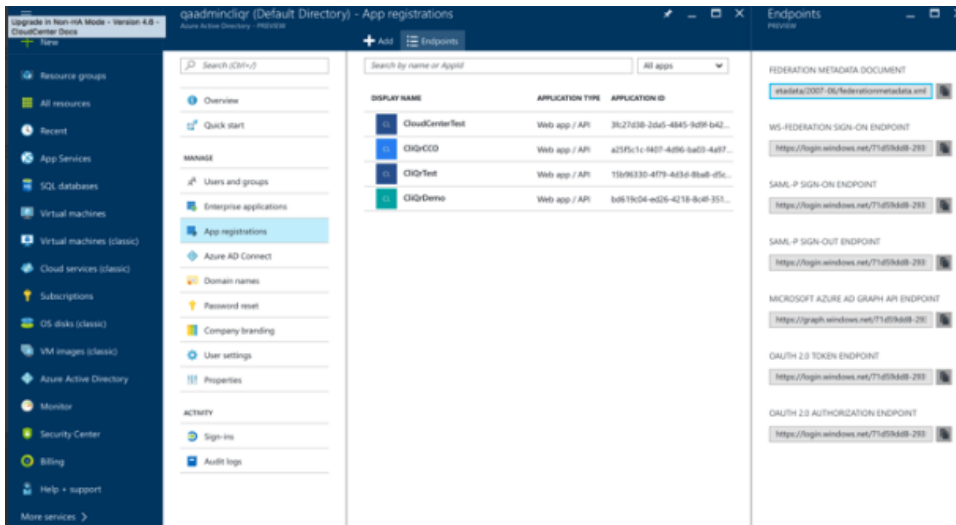
Note down the key after you click save – this key cannot be retrieved later from the portal, and it is used by the Workload Manager as the Client Key when creating the cloud account.



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 options – You can create either type using the *openssl* command from the command prompt of any Linux system:
 - A self-signed certificate: See the following example.

i 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
```

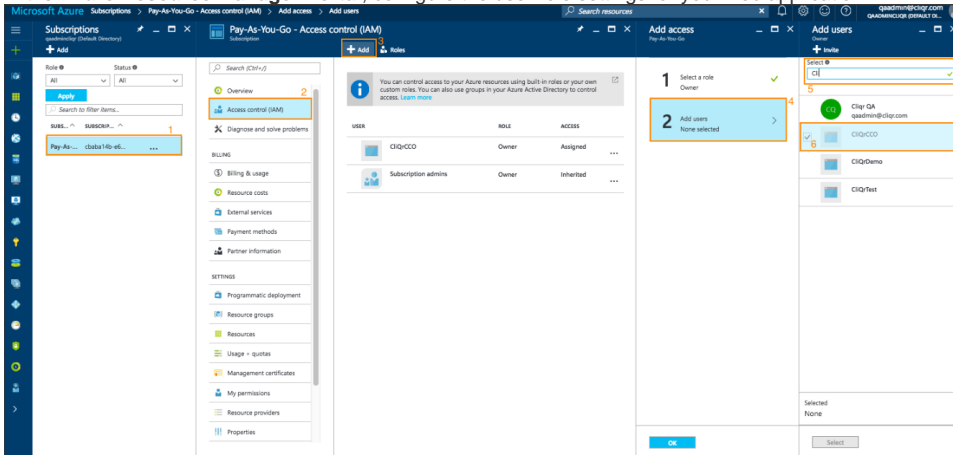
- Provide a password to this command when prompted.
- A Certificate Authority (CA) signed certificate – Generate a key and CSR, send/receive the certificate.csrfile(s) to the signature authority, convert the signed-certificate.pem to PKCS 12format, and provide a password to this command when prompted.

i 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 at <https://www.base64encode.org/>.
- Enter the base64 formatted certificate, and the export password used to create the PKCS formatted certificate, in the corresponding fields in the Workload Manager Add 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.

i 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 is the 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 AzureRM for 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 click the **Add Cloud Account** link. The Add Cloud Account dialog box displays, as shown in the figure below:

Add Cloud Account

Name *

Description

Cloud Credentials

Azure Login ID *

Azure Subscription ID *

Tenant ID *

Client ID *

Save Cancel

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 UUID identified 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 as identified in the *Certificate-Based Authentication* bullet in the *Prerequisites* section above.
 - ii. **Password** – Enter the password used to create the certificate as identified 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**, **Offer Id**, **EA Enrollment Number**, and **EA API Access Key**, as shown in the figure below. For information on setting up billing information, see <https://docs.microsoft.com/en-us/azure/billing/billing-enterprise-api>.

Add Cloud Account

Use Cert Based Auth OFF

Client Key *

Billing

Region Info

Offer Id

EA Enrollment Number

EA API Access Key

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** and navigate 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. <div style="border: 1px solid #ffc107; padding: 5px; margin: 5px 0;"> It is recommended that you do not add a <i>Reporting</i> account to the same tenant through different cloud groups. </div> <div style="border: 1px solid #ffc107; padding: 5px; margin: 5px 0;"> Enabling a public cloud account for <i>Reporting</i> may incur expenses to retrieve cost data. These expenses are proportional to the number of configured cloud accounts and regions. </div>
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.

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.

Account Name	Description	Billing Units	Enabled For	Actions
C3 Manual 1	C3 Manual Account 1	2 Billing Units	Provisioning, Reporting	Edit Delete
Master	Cost Optimizer Reporting	11 Billing Units	Reporting	Edit Delete
s account		050	Provisioning, Reporting	Edit Delete
C3 Manual Plans		810	Provisioning, Reporting	Edit Delete

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, or [manage instance types](#) for the cloud account.

Configure a Google Cloud

Configure a Google Cloud

Configuring a Google cloud is a four-step process:

- [Add a Google Cloud](#)
- [Add a Google Region](#)
- [Configure a Google Region](#)
- [Add a Google Cloud Account](#)

To add a Google cloud follow these steps.

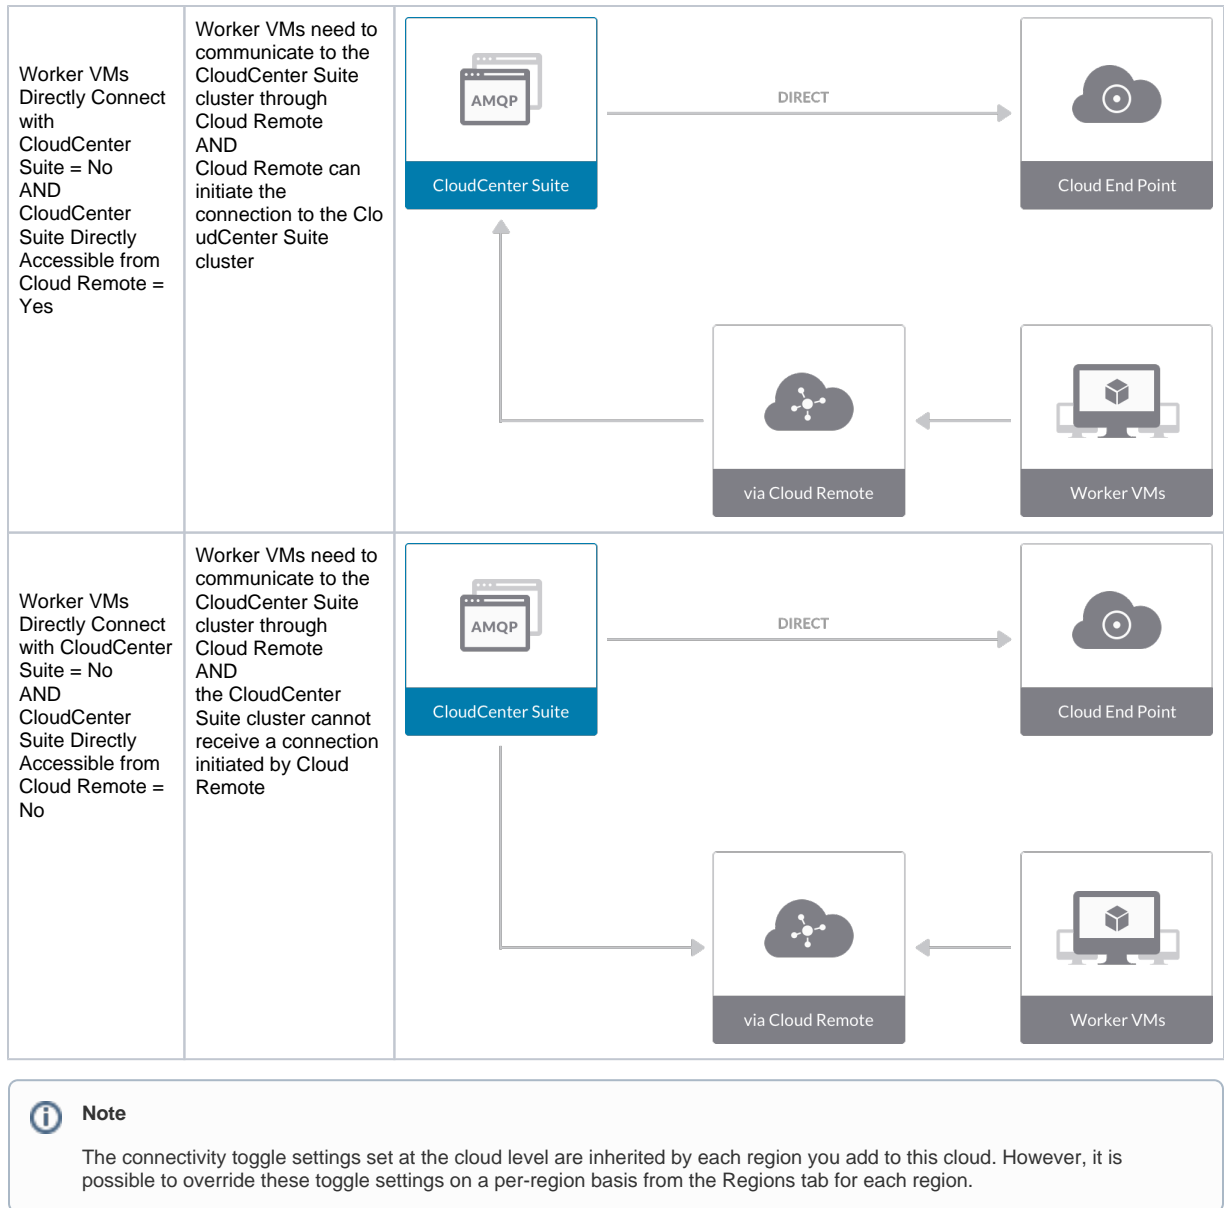
1. Navigate to **Admin > Clouds**. This brings you to the Clouds page. If you, or another tenant admin in your tenant, have already added clouds to your tenant, they will be listed here.
2. Click the **Add Cloud** link in the upper right. The Add Cloud dialog box is displayed.
3. Enter the **Cloud Name** 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. 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	<p>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</p> <p>Cloud Remote is not required</p>	<pre> graph LR subgraph CCS [CloudCenter Suite] AMQP[AMQP] end subgraph CE [Cloud End Point] Cloud[Cloud] end subgraph WVM [Worker VMs] VM[Worker VMs] end CCS -- DIRECT --> CE WVM -- DIRECT --> CCS </pre>



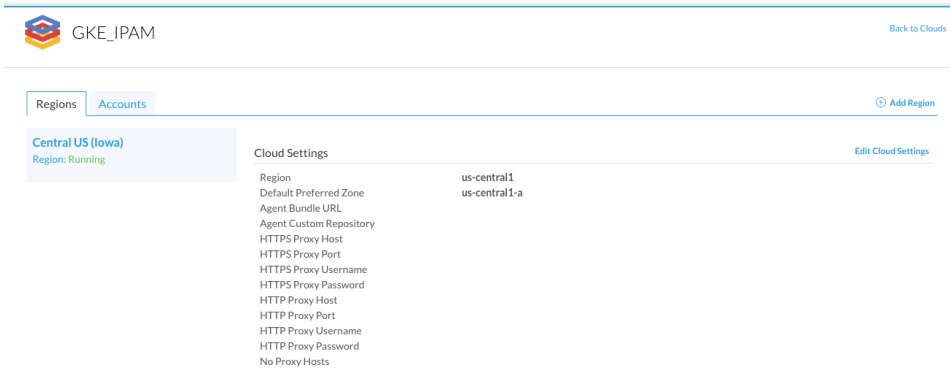
5. Click **Done** to save the configuration and close the dialog box. This brings you back to the Clouds page and the cloud you just created will be added to the bottom of the list on the left side of the page.

After creating a 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 the region you added shown on the right side of the page.

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 the Cloud 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.

- 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 a local 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 HTTPS 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 **Download Configuration** causes two things to happen:

- An encrypted zip file named artifacts.zip is 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 from software.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 in [Proxy 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 with cloud (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 both Cloud 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 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 all region connectivity settings at their current values and continue to the next settings section.

The exception to this guidance is when a NAT firewall or proxy server exists between the CloudCenter Suite management cluster and worker 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 the Guacamole 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 the Guacamole server running in the CloudCenter Suite cluster.
Worker VMs must use a proxy server or NAT firewall to access the Guacamole 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 the connectivity toggles in the 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 support case is updated with the VMI ID.**
- c. Navigate to the GCP dashboard and search for the VMI ID name provided in the [CloudCenter 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

Name ⓘ

instance-2

Region ⓘ **Zone** ⓘ

us-west1 (Oregon) ▼ us-west1-a ▼


Machine type
Customize to select cores, memory and GPUs.

1 vCPU ▼ 3.75 GB memory [Customize](#)

Container ⓘ

Deploy a container image to this VM instance. [Learn more](#)

Boot disk ⓘ

 New 30 GB standard persistent disk image
ccs-google-cloudremote-5-0-0-2019020... [Change](#)

Identity and API access ⓘ

Service account ⓘ

Compute Engine default service account ▼

Access scopes ⓘ

Allow default access
 Allow full access to all Cloud APIs
 Set access for each API

Firewall ⓘ

Add tags and firewall rules to allow specific network traffic from the Internet

Allow HTTP traffic
 Allow HTTPS traffic

⌵ [Management, security, disks, networking, sole tenancy](#)

You will be billed for this instance. [Compute Engine pricing](#) ↗

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. See [Cloud Remote \(Conditional\) > Scaling](#) for details.
- f. Once the first instance of the appliance has been launched, use the GCP 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 Download Configuration 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 to Cloud Remote web UI.


- Open another browser tab and login to https://<Cloud_Remote_ip> with the default credentials: admin/cisco.
- You will immediately be required to change your password. Do so now.
- 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 to the 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).

Region Connectivity Running		Download Configuration Configure Region
Cloud endpoint accessible from Cloud Center Manager	No	
Cloud Center Manager AMQP reachable from worker VM's	No	
Cloud Center Manager AMQP accessible from cloud	Yes	
Remote AMQP IP		
Worker AMQP IP	192.168.30.16:5671	
Blade Name	cloudcenter-blade-vmware-9-0289	
Blade Port	8443	

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

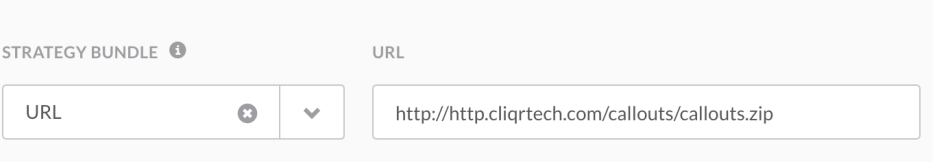
Workload Manager lets you define a VM naming strategy and an IP address management (IPAM) strategy on a per cloud region basis. The VM naming strategy lets Workload Manager generate unique names for each of your VMs as they are launched. The IPAM strategy lets Workload Manager generate IP addresses for 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 and an IP address deallocation callout script. An IP address allocation callout script may 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 and 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 CliQr 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	<p>Required if you selected "CliQr Macro Replacement" for the Instance Naming Strategy field. 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.</p> <p>The resulting VM name string must meet the following requirements for job submission to succeed:</p> <ul style="list-style-type: none"> Only contain the following characters: <ul style="list-style-type: none"> Lowercase a to z 0 to 9 Hyphen Start with an alpha character Cannot end with a hyphen
Instance IPAM Strategy	<p>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).</p>
Custom VM Name	<p>Location of your custom VM naming script if you set the Instance Naming Strategy field to "Custom Callout". Location options incl "Script from bundle" if you specify a zip file in the Strategy Bundle field. See Guidance for Callout Scripts for guidance on developing custom VM naming scripts.</p>
IPAM Alloc Rule	<p>Location of your custom IPAM address allocation script if you set the Instance IPAM Strategy field to "IPAM Callouts". Location options 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.</p>
IPAM Dealloc Rule	<p>Similar to the IPAM Alloc Rule above except this is where you specify the location of your custom IPAM address deallocation script. See Guidance for Callout Scripts for guidance on developing custom IPAM callout scripts.</p>

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 for CliQr Macro Replacement:

Macro	Returned Value
%os1%	First character of OS type string in lower case
%os2%	First two characters of OS type string in lower case
%OS1%	First character of OS type string in upper case
%OS2%	First two character of OS type string in upper case
%RND<number>%	<p>Fixed length random string, e.g., %RND10%, 10-character long random string</p> <div style="border: 1px solid #ccc; padding: 5px; background-color: #fff9c4;">  This is a mandatory field, only values in the range: %RND6% to %RND11% or just %RND% which defaults to %RND6% </div>
%UID%	User ID
%VID%	Vendor ID
%VM_PREFIX%	<p>VM name prefix must be added as a global parameter in the app profile.</p> <div style="border: 1px solid #ccc; padding: 5px; background-color: #fff9c4;">  Only macro whose value can be user-customized. </div>

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
- Post VM 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 execute in 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) to specify external lifecycle actions at the region level.

External Lifecycle Actions		Edit External Lifecycle Actions
External Actions Bundle		
Pre VM Start	-	
Pre VM Init	-	
Post VM Init	-	
Pre VM Stop	-	
Post VM Stop	-	

Click on the **Edit External Lifecycle Actions** link in the upper right. This displays the Edit 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 figure below.

EXTERNAL ACTIONS BUNDLE ⓘ	URL
HTTP ⓘ	myrepo.com/myscriptbundle.zip

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.

PRE VM START	SCRIPT
Script from bundle ⓘ	prevm_start.py

The scripts you specify can be bash scripts or Python scripts called from within a bash script. For more information on lifecycle scripts, see [Deployment 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.

Instance Types [Sync Instance Types](#) [Add Instance Type](#)

Search:

Show 30 per page Page 1 of 1

Name	Instance Type	Price	Actions
m1.small	0	\$0/hr	Edit Delete
m1.large	0bbdff46-a647-4eab-827f-7367e1e6...	\$0/hr	Edit Delete
m2.small	809193e3-fb66-49a2-9c55-d318d2a...	\$0/hr	Edit Delete
m1.nano	ab1365a2-debc-44b1-b4cf-3d5dd15c...	\$0/hr	Edit Delete



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 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 must manually import the corresponding physical images into your region and then map the corresponding logical images to these physical images. See [Images](#) for more context.



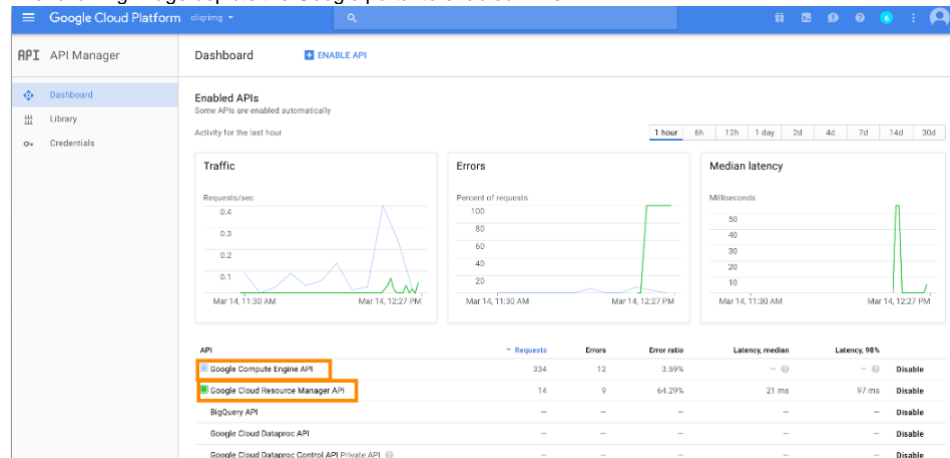
Be aware that these screenshots may change based on the Google Cloud platform changes. They are provided in this section as a point of reference.

Prerequisites

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

- A valid [Google Cloud Platform account](#) with *Project Owner* permissions
- If using the Shared VPC network feature, you also required Shared VPC Admin permissions (see <https://cloud.google.com/vpc/docs/provisioning-shared-vpc> for 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: abcdefghijklmnopqrstuvwxyz-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. See [Cloud 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.

Create private key for "[REDACTED]"

Downloads a file that contains the private key. Store the file securely because this key can't be recovered if lost.

Key type

JSON
Recommended

P12
For backward compatibility with code using the P12 format

CANCEL
CREATE

- 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 in the **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:

Add Cloud Account

Name *

Description

Cloud Credentials

GCP Email Address *

Email address associated with your GCP account

Service Account JSON file *

Choose File No file chosen

Billing

Bucket Name

Save
Cancel

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 <i>Prerequisites</i> section)

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

Add Cloud Account

Cloud Credentials

GCP Email Address *

name@example.com

Email address associated with your GCP account

Service Account JSON file *

Choose File No file chosen

Billing

Bucket Name

Report Prefix

Connect



Save
Cancel



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.

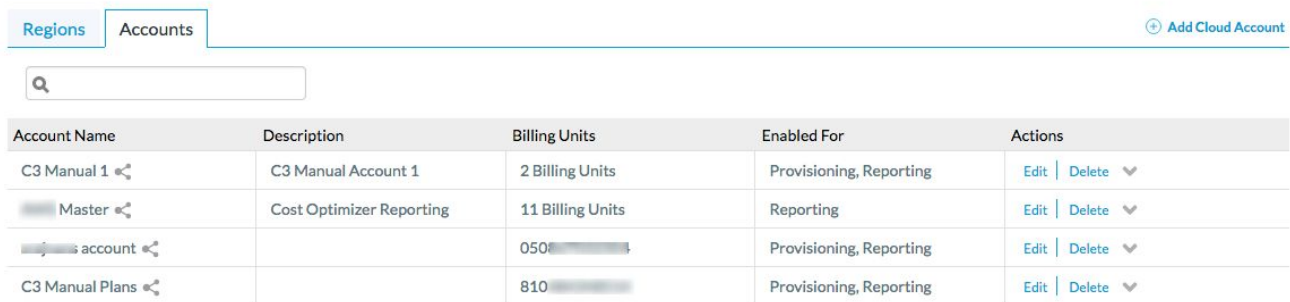
Value	Usage
Provisioning	Workload Manager can deploy jobs using this account.

Reporting	<p>Cost Optimizer and Workload Manager will track cloud costs for this account. Typical usage: master cloud accounts that are used for billing aggregation.</p> <div style="border: 1px solid #ccc; padding: 5px; margin: 5px 0;"> <p> It is recommended that you do not add a <i>Reporting</i> account to the same tenant through different cloud groups.</p> </div> <div style="border: 1px solid #ccc; padding: 5px; margin: 5px 0;"> <p> Enabling a public cloud account for <i>Reporting</i> may incur expenses to retrieve cost data. These expenses are proportional to the number of configured cloud accounts and regions.</p> </div>
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.

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.



Account Name	Description	Billing Units	Enabled For	Actions
C3 Manual 1	C3 Manual Account 1	2 Billing Units	Provisioning, Reporting	Edit Delete
Master	Cost Optimizer Reporting	11 Billing Units	Reporting	Edit Delete
account		050	Provisioning, Reporting	Edit Delete
C3 Manual Plans		810	Provisioning, Reporting	Edit Delete

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, or [manage instance types](#) for the cloud account.

Configure an OpenStack Cloud

Configure an OpenStack Cloud

Configuring an OpenStack cloud is a four-step process:

- [Add an OpenStack Cloud](#)
- [Add an OpenStack Region](#)
- [Configure an OpenStack Region](#)
- [Add an OpenStack Cloud Account](#)

To add an OpenStack 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**.



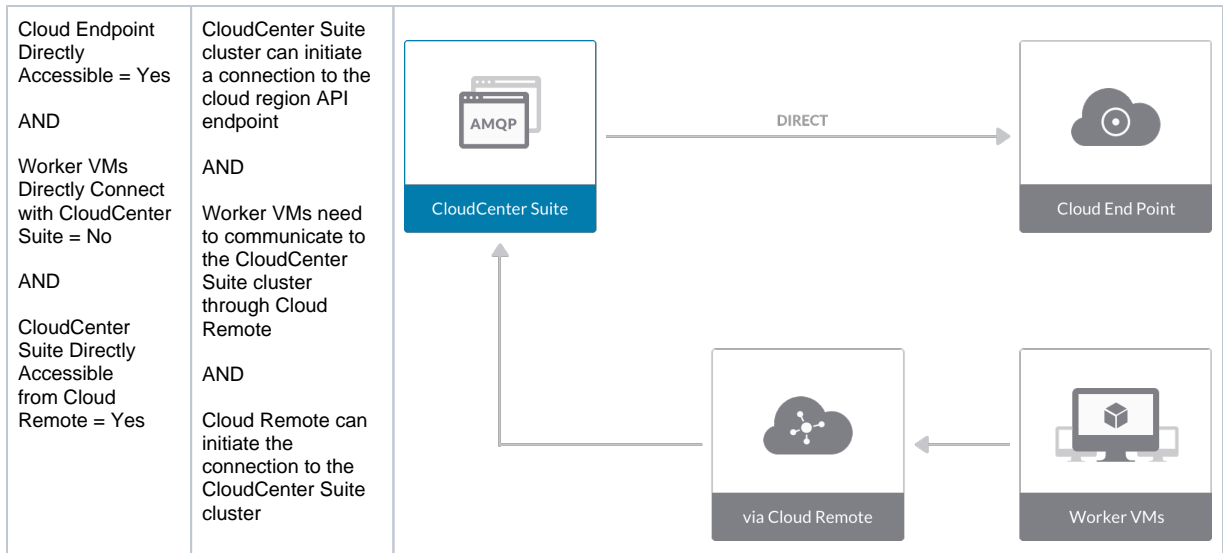
When assigning the Cloud Name, use a unique name. Ensure that you do not assign a previously-used name (including deleted cloud names) as the Cloud Name as they may cause deployments to fail intermittently.

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	<p>The diagram shows a central box labeled 'CloudCenter Suite' containing an 'AMQP' icon. Two arrows labeled 'DIRECT' point from this box to two separate boxes on the right. The top box is labeled 'Cloud End Point' and contains a cloud icon with a plus sign. The bottom box is labeled 'Worker VMs' and contains a computer monitor icon with a cube on it.</p>

<p>Cloud Endpoint Directly Accessible = No</p> <p>AND</p> <p>Worker VMs Directly Connect with CloudCenter Suite = No</p> <p>AND</p> <p>CloudCenter Suite Directly Accessible from Cloud Remote = Yes</p>	<p>CloudCenter Suite cluster cannot initiate a connection to the cloud region API endpoint</p> <p>AND</p> <p>Worker VMs cannot initiate a connection to the CloudCenter Suite cluster</p> <p>AND</p> <p>Cloud Remote can initiate the connection to the CloudCenter Suite cluster</p>	<pre> graph LR CC[CloudCenter Suite] CR[Cloud Remote] CE[Cloud End Point] W[Worker VMs] W --> CR CR --> CC CR --> CE </pre>
<p>Cloud Endpoint Directly Accessible = No</p> <p>AND</p> <p>Worker VMs Directly Connect with CloudCenter Suite = No</p> <p>AND</p> <p>CloudCenter Suite Directly Accessible from Cloud Remote = No</p>	<p>CloudCenter Suite cluster cannot initiate a connection to the cloud region API endpoint</p> <p>AND</p> <p>Worker VMs cannot initiate a connection to the CloudCenter Suite cluster</p> <p>AND</p> <p>Cloud Remote cannot initiate the connection to the CloudCenter Suite cluster</p>	<pre> graph LR CC[CloudCenter Suite] CR[Cloud Remote] CE[Cloud End Point] W[Worker VMs] W --> CR CR --> CE </pre>
<p>Cloud Endpoint Directly Accessible = Yes</p> <p>AND</p> <p>Worker VMs Directly Connect with CloudCenter Suite = No</p> <p>AND</p> <p>CloudCenter Suite Directly Accessible from Cloud Remote = No</p>	<p>CloudCenter Suite cluster can initiate a connection to the cloud region API endpoint</p> <p>AND</p> <p>Worker VMs cannot initiate a connection to the CloudCenter Suite cluster</p> <p>AND</p> <p>Cloud Remote cannot initiate the connection to the CloudCenter Suite cluster</p>	<pre> graph LR CC[CloudCenter Suite] CR[Cloud Remote] CE[Cloud End Point] W[Worker VMs] W --> CR CR --> CC CC -- DIRECT --> CE </pre>



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.

- 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.

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

- Navigate to the **Clouds** page and select the cloud you created on the left side of the screen.
- Click the **Add Region** button on the right side of the screen. The **Add Region** dialog box is displayed.
- 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.

- Click **Save**. You are brought to the **Clouds** page with the region you added shown on the right side of the page.

To configure a region you added to your OpenStack cloud, perform the following steps.

- Navigate to Clouds page: **Admin > Clouds** to find your OpenStack cloud from the cloud list on the left half of the screen
- 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 OpenStack cloud, the **Regions** tab will show multiple individual region tabs on the left side of the screen.

The screenshot shows the 'Accounts' tab for region 'r1'. The 'Cloud Settings' section is expanded, displaying the following fields:

- Region: r1
- OpenStack Keystone API version
- OpenStack Keystone Authentication Endpoint
- Additional Ports for Openstack endpoints: 8774,9292,8776,9696
- Exclude these special characters for Windows password
- Use Config Drive
- Nodes Per Batch
- Bootable Volume Mapping Required
- Agent Bundle URL
- Agent Custom Repository
- HTTPS Proxy Host
- HTTPS Proxy Port
- HTTPS Proxy Username
- HTTPS Proxy Password
- HTTP Proxy Host
- HTTP Proxy Port
- HTTP Proxy Username
- HTTP Proxy Password
- No Proxy Hosts

- Click the tab of the region you want to configure.
- 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 a local 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 HTTPS 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.

Region Connectivity Running Download Configuration [Configure Region](#)

Clicking **Download Configuration** causes two things to happen:

- An encrypted zip file named artifacts.zip is 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.

Region Connectivity Enabling... Download Configuration Copy Encryption Key [Edit Connectivity](#)

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 from software.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 in [Proxy 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 with cloud (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 both Cloud Remote and non-Cloud Remote modes. OpenStack supports Cloud Remote mode and supports all proxy settings except proxy hosts in non-Cloud Remote mode.

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

- 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 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.

The exception to this guidance is when a NAT firewall or proxy server exists between the CloudCenter Suite management cluster and worker 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 the Guacamole 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 the Guacamole server running in the CloudCenter Suite cluster.
Worker VMs must use a proxy server or NAT firewall to access the Guacamole 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.

6. If any of the connectivity toggles in the **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 from software.cisco.com.
- b. Through the OpenStack console, import and launch the Cloud Remote appliance. This process is similar to importing and launching the [CloudCenter 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.

Setup Cloud Remote Firewall Rules for 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 additional port 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 Download Configuration 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 to Cloud Remote web UI.

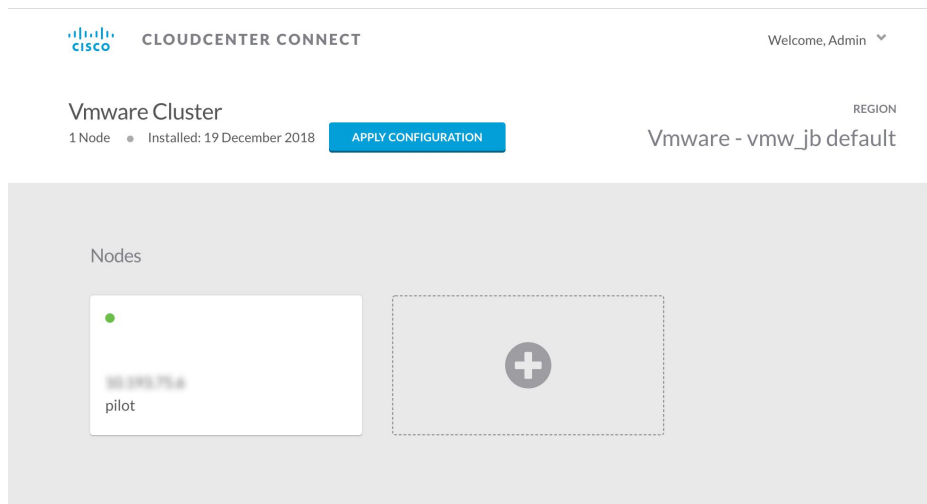
- Open another browser tab and login to https://<Cloud Remote_ip> with the default credentials: admin/cisco.
- You will immediately be required to change your password. Do so now.
- You are now brought to the Cloud Remote home page as shown in the figure below.




- 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.



- Paste the encryption key that was copied to the clipboard into the **Encryption Key** field in the dialog box.
- Click **Select File** and browse to the artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
- Click **Confirm**.
- 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

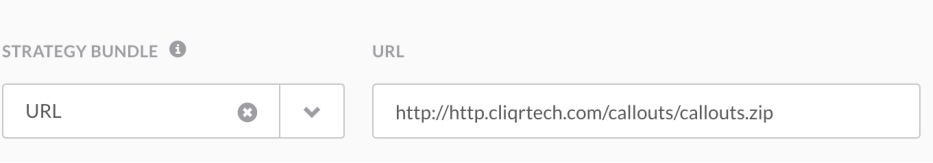
Workload Manager lets you define a VM naming strategy and an IP address management (IPAM) strategy on a per cloud region basis. The VM naming strategy lets Workload Manager generate unique names for each of your VMs as they are launched. The IPAM strategy lets Workload Manager generate IP addresses for 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 and an IP address deallocation callout script. An IP address allocation callout script may 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 and 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 CliQr 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 the Instance Naming Strategy field. 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 must meet the following requirements for job submission to succeed: <ul style="list-style-type: none"> Only contain the following characters: <ul style="list-style-type: none"> Lowercase a to z 0 to 9 Hyphen Start with an 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 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 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 options 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 your custom IPAM address deallocation script. See 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 for CliQr Macro Replacement:

Macro	Returned Value
%os1%	First character of OS type string in lower case
%os2%	First two characters of OS type string in lower case
%OS1%	First character of OS type string in upper case
%OS2%	First two 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
- Post VM 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 execute in 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) to specify external lifecycle actions at the region level.

External Lifecycle Actions	Edit External Lifecycle Actions
External Actions Bundle	
Pre VM Start	-
Pre VM Init	-
Post VM Init	-
Pre VM Stop	-
Post VM Stop	-

Click on the **Edit External Lifecycle Actions** link in the upper right. This displays the Edit 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 figure below.

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, see [Deployment Lifecycle Scripts](#).

- 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 the instance 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.

Instance Types [Sync Instance Types](#) [Add Instance Type](#)

Search:

Show 30 per page Page 1 of 1

Name	Instance Type	Price	Actions
m1.small	0	\$0/hr	Edit Delete
m1.large	0bbdff46-a647-4eab-827f-7367e1e6...	\$0/hr	Edit Delete
m2.small	809193e3-fb66-49a2-9c55-d318d2a...	\$0/hr	Edit Delete
m1.nano	ab1365a2-debc-44b1-b4cf-3d5dd15c...	\$0/hr	Edit Delete

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

Instance Types [Edit Instance Types](#) [Add Instance Type](#)

Search:

Show 30 per page Page 1 of 1

Name	Instance Type	Price	Actions
default	default	\$0/hr	Delete
2CPU_2048MBMEM	2CPU_2048MBMEM	\$0/hr	Delete
1CPU_2048MBMEM_u	1CPU_2048MBMEM	\$0.002/hr	Delete
2CPU_8192MBMEM	2CPU_8192MBMEM	\$0/hr	Delete
6CPU_4096MBMEM	6CPU_4096MBMEM	\$0/hr	Delete
6CPU_6144MBMEM	6CPU_6144MBMEM	\$0/hr	Delete
8CPU_24288MBMEM	8CPU_24288MBMEM	\$0/hr	Delete
4CPU_8192MBMEM	4CPU_8192MBMEM	\$0.008/hr	Delete
CloudCenter Suite. 2CPU_4096MBMEM	2CPU_4096MBMEM	\$0.004/hr	Delete

You must specify the price details in the **Price** field for OpenStack cloud when adding instance types. The price is used to compute resource costs, which is displayed as invoice costs in the [Cost 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, the cost of the affected resources is 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:

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

DATE	PRICE/HOUR	LAST CHANGED BY
Mar 12, 2020 11:10 AM (GMT)	\$ 0	Admin User admin@cisco.com
03/12/2020 11:10 AM	\$ 0.005 /hr	Admin User admin@cisco.com
Mar 08, 2020 01:12 AM (GMT)	\$ 0.006	

- c. Specify a date and time in the appropriate fields and enter the desired price in the **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 all resources (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 (see [Cost Optimizer Dashboard](#)), as shown in the following sample screenshot.

● VMware vmware_usdc-default - cost adjusted successfully

Price change on instance type: 3CPU_2048MBMEM has been reflected on the deployments and resource costs.

a minute ago

Simultaneously, the updated accrued cost of a VM is displayed on the [Virtual Machines](#) page 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 default storage type 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.

Storage Types [Add Storage Type](#)

⚠
Storage Types not configured
×

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

Add Storage Type

Display Name *

Cost Per Month *

\$

/GB

Cloud Storage Type ID *

Minimum Volume Size *

GB

Maximum Volume Size *

GB

Maximum Throughput *

MB/s

Maximum IOPS *

Provisioned IOPS

Save
Cancel

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 the [Cost 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.

Storage Types [Add Storage Type](#)

Show 30 per page Page 1 of 1

Name	Storage Type	Price	Actions
Local SSD scratch disk	local-ssd	\$0.218/GB-month	Edit Delete
Standard Persistent Disk	pd-standard	\$0.04/GB-month	Edit Delete
SSD Persistent Disk	pd-ssd	\$0.17/GB-month	Edit Delete

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:

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

DATE	PRICE / GB-MONTH	LAST CHANGED BY
Mar 12, 2020 09:49 AM (GMT)	\$ 22	Admin User admin@cisco.com
03/11/2020 3 :06 PM	\$ 21 /GB	Admin User admin@cisco.com
Mar 08, 2020 06:16 PM (GMT)	\$ 1	Admin User admin@cisco.com
Mar 08, 2020 01:13 AM (GMT)	\$ 1.1	Admin User admin@cisco.com
Mar 08, 2020 01:12 AM (GMT)	\$ 22	Admin User admin@cisco.com

- Specify a date and time in the appropriate fields and enter the desired price in the **Price/Hour** field, enter the desired price.
- 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 all resources (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 (see [Cost Optimizer Dashboard](#)) in the header. Simultaneously, the updated accrued cost of a VM is displayed on the [Storage Volumes](#) page and the updated deployment cost is displayed on the Deployment Details page.

- 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. See [Images](#) for more context.

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	<pre> 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 </pre>	<pre> compute:get compute:get_all compute:get_all_tenants compute:get_instance_metadata compute:get_all_instance_metadata compute:get_all_instance_system_metadata </pre>
Network	<pre> get_network get_subnet network:get_all </pre>	<pre> get_network get_subnet network:get_all </pre>
Block Storage	<pre> volume:get volume:get_all volume:create volume:delete </pre>	<pre> volume:get volume:get_all </pre>
Identity	<pre> identity:list_user_projects identity:get_user identity:list_users identity:list_projects </pre>	<pre> identity:list_user_projects identity:get_user identity:list_users identity:list_projects </pre>
Image	<pre> get_image get_images delete_image download_image add_image add_member delete_member </pre>	<pre> get_image get_images </pre>

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.

Add Cloud Account

Name *

Description

Cloud Credentials

OpenStack User Name *

User Name associated with your OpenStack account

OpenStack Account Password *

Default Domain Name (V3)

Default Domain Id (V3)

Save Cancel

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.

Add Cloud Account

OpenStack User Name *

User Name associated with your OpenStack account

OpenStack Account Password *

Default Domain Name (V3)

Default Domain Id (V3)

Either Default Domain Id or Default Domain Name is needed for V3 API

Default Tenant Name (V3 Project Name)

Default Tenant Id (V3 Project Id)

Either Default Tenant Id or Default Tenant Name is needed

Connect

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: <ul style="list-style-type: none"> • 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 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.

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.

Account Name	Description	Billing Units	Enabled For	Actions
C3 Manual 1	C3 Manual Account 1	2 Billing Units	Provisioning, Reporting	Edit Delete
Master	Cost Optimizer Reporting	11 Billing Units	Reporting	Edit Delete
account		050	Provisioning, Reporting	Edit Delete
C3 Manual Plans		810	Provisioning, Reporting	Edit Delete

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, or [manage instance types](#) for the cloud account.

Configure a Kubernetes Cloud


Configure a Kubernetes Cloud

Configuring a Kubernetes cloud is a three-step process:


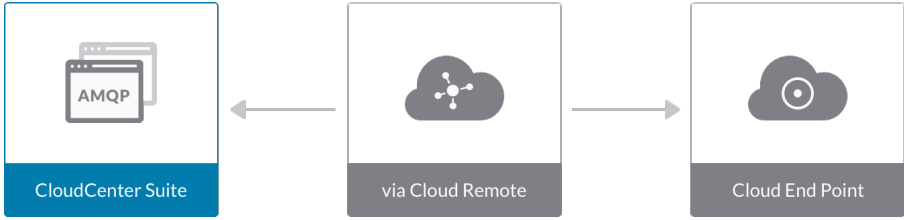
- [Add a Kubernetes Cloud](#)
- [Configure a Kubernetes Region](#)
- [Add a Kubernetes Cloud Account](#)


To add a Kubernetes 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 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 **Next**.
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	 <p>The diagram shows a box labeled 'CloudCenter Suite' containing an 'AMQP' icon. An arrow points from this box to a box labeled 'Cloud End Point' containing a cloud icon with a central dot.</p>
Cloud Endpoint Directly Accessible = No AND CloudCenter Suite Directly Accessible from Cloud 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	 <p>The diagram shows three boxes. On the left is 'CloudCenter Suite' with an 'AMQP' icon. In the middle is a box labeled 'via Cloud Remote' with a cloud icon and a network diagram. On the right is 'Cloud End Point' with a cloud icon and a central dot. Arrows show a bidirectional connection between 'CloudCenter Suite' and 'via Cloud Remote', and a connection from 'via Cloud Remote' to 'Cloud End Point'.</p>



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.

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

1. Navigate to Clouds page: **Admin > Clouds**. Find your newly created Kubernetes 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
Kubern etes cluster API Endpoi nt	This field is set to the value you set for the API endpoint when you created this Kubernetes cloud. You can edit it here but should only do so if the API endpoint address of your Kubernetes cloud has changed since you added it to CloudCenter Suite.
API version override	This tells CloudCenter 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: <ul style="list-style-type: none"> • Example 1: Secret: custom_api_version;Service: custom_api_version;PersistentVolumeClaim: custom_api_version;NetworkPolicy: custom_api_version;Pod: custom_api_version;Deployment: custom_api_version • Example 2: PersistentVolumeClaim: custom_api_version;NetworkPolicy: custom_api_version;Pod: custom_api_version;Deployment: custom_api_version • Example 3: PersistentVolumeClaim: custom_api_version;NetworkPolicy: custom_api_version
Names pace(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 the connectivity toggles in the 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 region to 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 support case is updated with the VMI ID.**
- c. Navigate to the GCP dashboard and search for the VMI ID name provided in the [CloudCenter 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

Name ⓘ

instance-2

Region ⓘ **Zone** ⓘ

us-west1 (Oregon) ▼ us-west1-a ▼


Machine type
Customize to select cores, memory and GPUs.

1 vCPU ▼ 3.75 GB memory [Customize](#)

Container ⓘ

Deploy a container image to this VM instance. [Learn more](#)

Boot disk ⓘ

 New 30 GB standard persistent disk image
ccs-google-cloudremote-5-0-0-2019020... [Change](#)

Identity and API access ⓘ

Service account ⓘ

Compute Engine default service account ▼

Access scopes ⓘ

Allow default access
 Allow full access to all Cloud APIs
 Set access for each API

Firewall ⓘ

Add tags and firewall rules to allow specific network traffic from the Internet

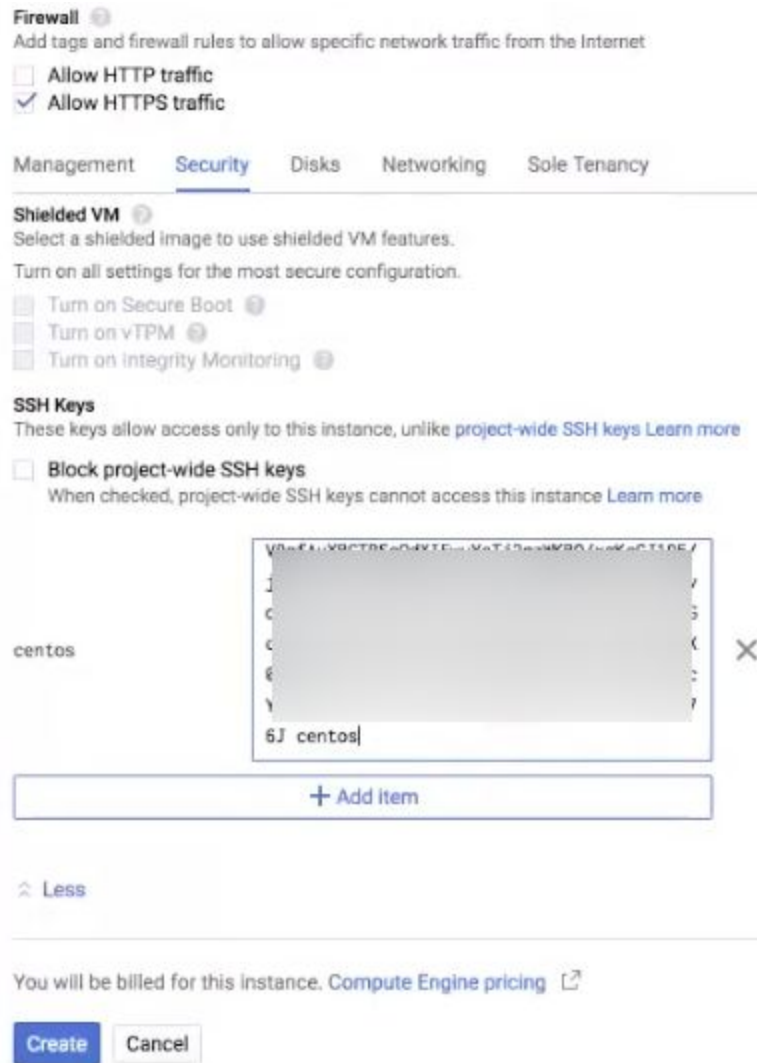
Allow HTTP traffic
 Allow HTTPS traffic

⌵ [Management, security, disks, networking, sole tenancy](#)

You will be billed for this instance. [Compute Engine pricing](#) ↗

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. See [Cloud Remote \(Conditional\) > Scaling](#) for details.
- f. Once the first instance of the appliance has been launched, use the GCP 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 Download Configuration 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 to Cloud Remote web UI.


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



- 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 to the 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).

Region Connectivity Running		Download Configuration Configure Region
Cloud endpoint accessible from Cloud Center Manager	No	
Cloud Center Manager AMQP reachable from worker VM's	No	
Cloud Center Manager AMQP accessible from cloud	Yes	
Remote AMQP IP		
Worker AMQP IP	192.168.30.16:5671	
Blade Name	cloudcenter-blade-vmware-9-0289	
Blade Port	8443	

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 from software.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. See [Cloud Remote \(Conditional\)](#) > *Scaling* for details.
- j. Once the first instance of the appliance has been launched, use the vSphere client 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) > <i>Custom Port Numbers (Conditional)</i>). 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 Download Configuration 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 to Cloud Remote web UI.


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



- 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.



- Paste the encryption key that was copied to the clipboard into the **Encryption Key** field in the dialog box.
- Click **Select File** and browse to the artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
- Click **Confirm**.
- 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).

Region Connectivity Running		Download Configuration Configure Region
Cloud endpoint accessible from Cloud Center Manager	No	
Cloud Center Manager AMQP reachable from worker VM's	No	
Cloud Center Manager AMQP accessible from cloud	Yes	
Remote AMQP IP		
Worker AMQP IP	192.168.30.16:5671	
Blade Name	cloudcenter-blade-vmware-9-0289	
Blade Port	8443	

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 region to support a Kubernetes target cloud as follows.

Obtain and Launch the Cloud Remote Appliance in AWS

- Obtain the Cloud Remote shared AMI from Cisco support and launch it. Follow the same guidance for obtaining and launching the [CloudCenter Suite installer appliance for AWS](#).
- 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.
- 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 (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 Download Configuration 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 to Cloud Remote web UI.

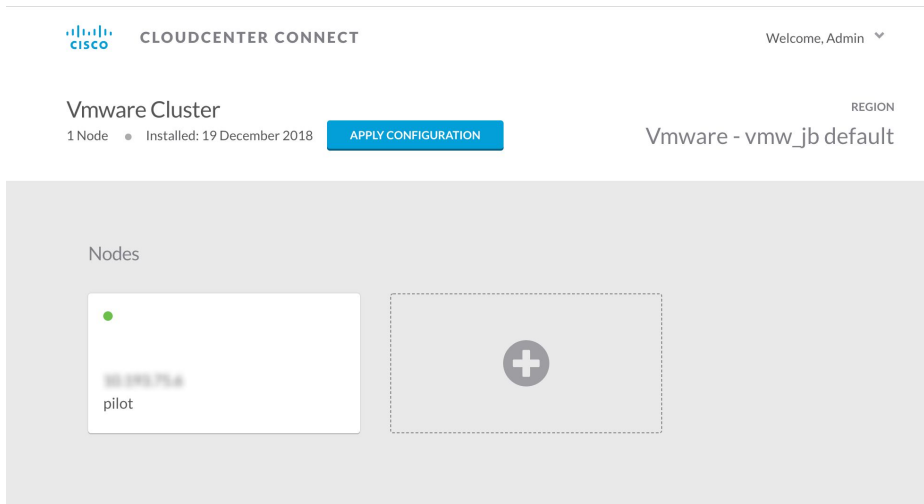
- Open another browser tab and login to https://<Cloud Remote_ip> with the default credentials: admin/cisco.
- You will immediately be required to change your password. Do so now.
- You are now brought to the Cloud Remote home page as shown in the figure below.




- 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.



- Paste the encryption key that was copied to the clipboard into the **Encryption Key** field in the dialog box.
- Click **Select File** and browse to the artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
- Click **Confirm**.
- 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 region to support a Kubernetes target cloud as follows.

Download and Launch the Cloud Remote Appliance in AzureRM

- Download the Cloud Remote appliance for AzureRM as a zip file from software.cisco.com and then unzip it to reveal the VHD file.
- 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 the [CloudCenter Suite installer appliance for AzureRM](#).

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

- 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.
- Once the first instance of the appliance has been launched, use the AzureRM 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 Download Configuration 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 to Cloud Remote web UI.


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



- 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 to the 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).

Region Connectivity Running		Download Configuration Configure Region
Cloud endpoint accessible from Cloud Center Manager	No	
Cloud Center Manager AMQP reachable from worker VM's	No	
Cloud Center Manager AMQP accessible from cloud	Yes	
Remote AMQP IP		
Worker AMQP IP	192.168.30.16:5671	
Blade Name	cloudcenter-blade-vmware-9-0289	
Blade Port	8443	

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 from software.cisco.com.
- b. Through the OpenStack console, import and launch the Cloud Remote appliance. This process is similar to importing and launching the [CloudCenter 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.

Region Connectivity Running Download Configuration Configure Region

Clicking Download Configuration 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.

Region Connectivity Enabling... Download Configuration Copy Encryption Key Edit Connectivity

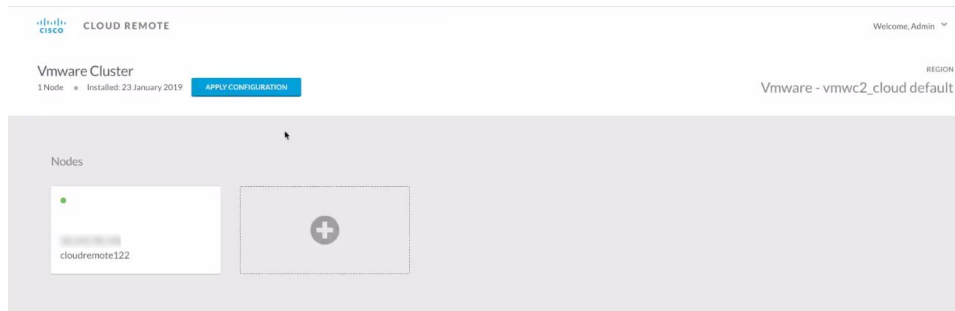
Click the **Copy Encryption Key** link to save the key to your clipboard. A success message will be displayed temporarily above the Region Connectivity section header. Make sure not to overwrite the clipboard with other data. You will need the key when you upload the configuration zip file to Cloud Remote.



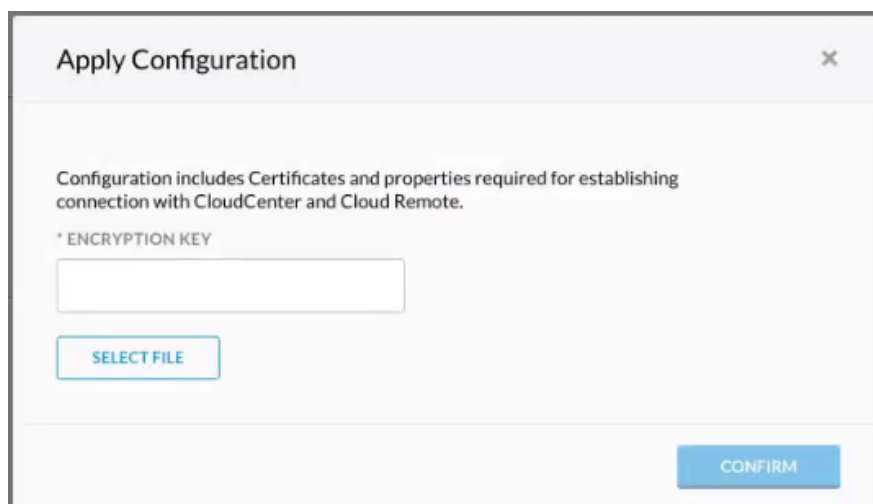
If you change the connectivity settings in the CloudCenter Suite UI and need to again download the zip file, a new encryption key is automatically created and can be copied to the clipboard by clicking the **Copy Encryption Key** link again.

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


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



- 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.



- Paste the encryption key that was copied to the clipboard into the **Encryption Key** field in the dialog box.
- Click **Select File** and browse to the artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
- Click **Confirm**.
- 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).

Region Connectivity Running		Download Configuration Configure Region
Cloud endpoint accessible from Cloud Center Manager	No	
Cloud Center Manager AMQP reachable from worker VM's	No	
Cloud Center Manager AMQP accessible from cloud	Yes	
Remote AMQP IP		
Worker AMQP IP	192.168.30.16:5671	
Blade Name	cloudcenter-blade-vmware-9-0289	
Blade Port	8443	

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

- 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.

Instance Types Edit Instance Types Add Instance Type			
Show 30 ↓ per page Page 1 of 1			
Name	Instance Type	Price	Actions
default	default	\$0/hr	Delete
2CPU_2048MBMEM	2CPU_2048MBMEM	\$0/hr	Delete
1CPU_2048MBMEM_u	1CPU_2048MBMEM	\$0.002/hr	Delete
2CPU_8192MBMEM	2CPU_8192MBMEM	\$0/hr	Delete
6CPU_4096MBMEM	6CPU_4096MBMEM	\$0/hr	Delete
6CPU_6144MBMEM	6CPU_6144MBMEM	\$0/hr	Delete
8CPU_24288MBMEM	8CPU_24288MBMEM	\$0/hr	Delete
4CPU_8192MBMEM	4CPU_8192MBMEM	\$0.008/hr	Delete
2CPU_4096MBMEM	2CPU_4096MBMEM	\$0.004/hr	Delete

You must specify the price details in the **Price** field for Kubernetes cloud when adding instance types. The price is used to compute resource costs, which is displayed as invoice costs in the [Cost 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.

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 the [Kubernetes 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 the [Kubernetes web UI](#) and 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:

Name ↕	Type	Age ↕
cisco-token-9ptfm	kubernetes.io/service-account-token	2 months
wordpress	Opaque	2 months
mysql-pass	Opaque	2 months
mysql	Opaque	3 months
default-token-j0qlx	kubernetes.io/service-account-token	3 months

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

Details	
Name:	default-token-j0qlx
Namespace:	default
Annotations:	kubernetes.io/service-account.name: default kubernetes.io/service-account.uid: 67aaf3de-f668-11e7-8478-42010a8a0107
Creation Time:	2018-01-11T00:43 UTC
Type:	kubernetes.io/service-account-token
Data	
	ca.crt: 1119 bytes
	namespace: 7 bytes
	token: 846 bytes

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 from the Kubernetes Web UI assures this to be true.

- *The **kubect!** 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 the CloudCenter 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 <i>Prerequisites</i> 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 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.

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.

Account Name	Description	Billing Units	Enabled For	Actions
C3 Manual 1	C3 Manual Account 1	2 Billing Units	Provisioning, Reporting	Edit Delete
Master	Cost Optimizer Reporting	11 Billing Units	Reporting	Edit Delete
s account		050	Provisioning, Reporting	Edit Delete
C3 Manual Plans		810	Provisioning, Reporting	Edit Delete

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, or [manage instance types](#) for the cloud account.

Configure a vCD Cloud

Configure a vCD Cloud

Configuring a vCD cloud is a four-step process:

- [Add a vCD Cloud](#)
- [Configure a vCD Region](#)
- [Add a vCD Cloud Account](#)

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**, then 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 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	<p>The network diagram illustrates a central 'CloudCenter Suite' box containing an 'AMQP' icon. Two 'DIRECT' arrows originate from this box: one points to a 'Cloud End Point' box (represented by a cloud icon) and the other points to a 'Worker VMs' box (represented by a computer monitor icon). Both boxes are labeled 'DIRECT' above the arrows.</p>

<p>Cloud Endpoint Directly Accessible = No</p> <p>AND</p> <p>Worker VMs Directly Connect with CloudCenter Suite = No</p> <p>AND</p> <p>CloudCenter Suite Directly Accessible from Cloud Remote = Yes</p>	<p>CloudCenter Suite cluster cannot initiate a connection to the cloud region API endpoint</p> <p>AND</p> <p>Worker VMs cannot initiate a connection to the CloudCenter Suite cluster</p> <p>AND</p> <p>Cloud Remote can initiate the connection to the CloudCenter Suite cluster</p>	<p>The diagram shows three main components in boxes: 'CloudCenter Suite' (with an AMQP icon), 'via Cloud Remote' (with a cloud icon), and 'Cloud End Point' (with a cloud icon). A bidirectional arrow connects 'CloudCenter Suite' and 'via Cloud Remote'. A bidirectional arrow connects 'via Cloud Remote' and 'Cloud End Point'. Below these, 'Worker VMs' (with a computer icon) has an arrow pointing up to 'via Cloud Remote'.</p>
<p>Cloud Endpoint Directly Accessible = No</p> <p>AND</p> <p>Worker VMs Directly Connect with CloudCenter Suite = No</p> <p>AND</p> <p>CloudCenter Suite Directly Accessible from Cloud Remote = No</p>	<p>CloudCenter Suite cluster cannot initiate a connection to the cloud region API endpoint</p> <p>AND</p> <p>Worker VMs cannot initiate a connection to the CloudCenter Suite cluster</p> <p>AND</p> <p>Cloud Remote cannot initiate the connection to the CloudCenter Suite cluster</p>	<p>The diagram shows three main components in boxes: 'CloudCenter Suite' (with an AMQP icon), 'via Cloud Remote' (with a cloud icon), and 'Cloud End Point' (with a cloud icon). A bidirectional arrow connects 'CloudCenter Suite' and 'via Cloud Remote'. A bidirectional arrow connects 'via Cloud Remote' and 'Cloud End Point'. Below these, 'Worker VMs' (with a computer icon) has an arrow pointing up to 'via Cloud Remote'.</p>
<p>Cloud Endpoint Directly Accessible = Yes</p> <p>AND</p> <p>Worker VMs Directly Connect with CloudCenter Suite = No</p> <p>AND</p> <p>CloudCenter Suite Directly Accessible from Cloud Remote = No</p>	<p>CloudCenter Suite cluster can initiate a connection to the cloud region API endpoint</p> <p>AND</p> <p>Worker VMs cannot initiate a connection to the CloudCenter Suite cluster</p> <p>AND</p> <p>Cloud Remote cannot initiate the connection to the CloudCenter Suite cluster</p>	<p>The diagram shows three main components in boxes: 'CloudCenter Suite' (with an AMQP icon), 'via Cloud Remote' (with a cloud icon), and 'Cloud End Point' (with a cloud icon). A direct arrow labeled 'DIRECT' connects 'CloudCenter Suite' to 'Cloud End Point'. Below these, 'Worker VMs' (with a computer icon) has an arrow pointing up to 'via Cloud Remote', which then has an arrow pointing to 'CloudCenter Suite'.</p>




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.

- 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.

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

- 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.



- 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.
- 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**.

- 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 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.

The exception to this guidance is when a NAT firewall or proxy server exists between the CloudCenter Suite management cluster and worker 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 the Guacamole 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 the Guacamole server running in the CloudCenter Suite cluster.
Worker VMs must use a proxy server or NAT firewall to access the Guacamole 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.
- If any of the connectivity toggles in the 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

- 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.

- 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. See [Cloud 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) > <i>Custom Port Numbers (Conditional)</i>). 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) > <i>Custom Port Numbers (Conditional)</i>).
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) > <i>Custom Port Numbers (Conditional)</i>).
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 Download Configuration 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 to Cloud Remote web UI.


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



- 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.



- Paste the encryption key that was copied to the clipboard into the **Encryption Key** field in the dialog box.
- Click **Select File** and browse to the artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
- Click **Confirm**.
- 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).

Region Connectivity Running		Download Configuration Configure Region
Cloud endpoint accessible from Cloud Center Manager	No	
Cloud Center Manager AMQP reachable from worker VM's	No	
Cloud Center Manager AMQP accessible from cloud	Yes	
Remote AMQP IP		
Worker AMQP IP	192.168.30.16:5671	
Blade Name	cloudcenter-blade-vmware-9-0289	
Blade Port	8443	

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 lets Workload 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 and an IP address deallocation callout script. An IP address allocation callout script may 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 and add the path in the new field that appears to the right.				
	<table border="1" style="width: 100%;"> <tr> <td style="width: 50%;">STRATEGY BUNDLE ⓘ</td> <td style="width: 50%;">URL</td> </tr> <tr> <td> <input type="text" value="URL"/> ⓘ </td> <td> <input type="text" value="http://http.cliqrtech.com/callouts/callouts.zip"/> </td> </tr> </table>	STRATEGY BUNDLE ⓘ	URL	<input type="text" value="URL"/> ⓘ	<input type="text" value="http://http.cliqrtech.com/callouts/callouts.zip"/>
STRATEGY BUNDLE ⓘ	URL				
<input type="text" value="URL"/> ⓘ	<input type="text" value="http://http.cliqrtech.com/callouts/callouts.zip"/>				


Instance Naming Strategy	Dropdown for selecting one of three naming strategies: Default, CliQr Macro Replacement, and Hostname Callout. If you select CliQr 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 the Instance Naming Strategy field. 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 must meet the following requirements for job submission to succeed: <ul style="list-style-type: none"> Only contain the following characters: <ul style="list-style-type: none"> Lowercase a to z 0 to 9 Hyphen Start with an 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 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 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 options 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 your custom IPAM address deallocation script. See 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 for CliQr Macro Replacement:

Macro	Returned Value
%os1%	First character of OS type string in lower case
%os2%	First two characters of OS type string in lower case
%OS1%	First character of OS type string in upper case
%OS2%	First two 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%	<p>VM name prefix must be added as a global parameter in the app profile.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  Only macro whose value can be user-customized. </div>
-------------	---

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
- Post VM 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 execute in 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) to specify external lifecycle actions at the region level.

External Lifecycle Actions		Edit External Lifecycle Actions
External Actions Bundle		
Pre VM Start	-	
Pre VM Init	-	
Post VM Init	-	
Pre VM Stop	-	
Post VM Stop	-	

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

Edit External Lifecycle Actions


EXTERNAL ACTIONS BUNDLE ⓘ

▼

PRE VM START

▼

PRE VM INIT

▼

POST VM INIT

▼

PRE VM STOP

▼

POST VM STOP

▼

DONE

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 figure below.

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, see [Deployment Lifecycle Scripts](#).

- 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.

Instance Types [Sync Instance Types](#) [Add Instance Type](#)

Search: Show 30 per page Page 1 of 1

Name	Instance Type	Price	Actions
m1.small	0	\$0/hr	Edit Delete
m1.large	0bbdff46-a647-4eab-827f-7367e1e6...	\$0/hr	Edit Delete
m2.small	809193e3-fb66-49a2-9c55-d318d2a...	\$0/hr	Edit Delete
m1.nano	ab1365a2-debc-44b1-b4cf-3d5dd15c...	\$0/hr	Edit Delete

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, the cost of the affected resources is 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:

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

DATE	PRICE/HOUR	LAST CHANGED BY
Mar 12, 2020 11:10 AM (GMT)	\$ 0	Admin User admin@cisco.com
03/12/2020	11:10 AM	\$ 0.005 /hr
Admin User admin@cisco.com		
Mar 08, 2020 01:12 AM (GMT)	\$ 0.006	

OK

- c. Specify a date and time in the appropriate fields and enter the desired price in the **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 all resources (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 (see [Cost Optimizer Dashboard](#)), as shown in the following sample screenshot.

● VMware vmware_usdc-default - cost adjusted successfully

Price change on instance type: 3CPU_2048MBMEM has been reflected on the deployments and resource costs.

a minute ago

Simultaneously, the updated accrued cost of a VM is displayed on the [Virtual Machines](#) page 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 the package 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.

Storage Types [Add Storage Type](#)

Search: Show 30 per page Page 1 of 1

Name	Storage Type	Price	Actions
Local SSD scratch disk	local-ssd	\$0.218/GB-month	Edit Delete
Standard Persistent Disk	pd-standard	\$0.04/GB-month	Edit Delete
SSD Persistent Disk	pd-ssd	\$0.17/GB-month	Edit Delete

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:

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

Edit Price History - default ✕

DATE ↓	PRICE / GB-MONTH ⓘ	LAST CHANGED BY
Mar 12, 2020 09:49 AM (GMT)	\$ 22	Admin User admin@cisco.com
<input type="text" value="03/11/2020"/>	<input type="text" value="3 :06 PM"/>	<input type="text" value="\$ 21"/> <input type="text" value="/GB"/>
Mar 08, 2020 06:16 PM (GMT)	\$ 1	Admin User admin@cisco.com
Mar 08, 2020 01:13 AM (GMT)	\$ 1.1	Admin User admin@cisco.com
Mar 08, 2020 01:12 AM (GMT)	\$ 22	Admin User admin@cisco.com

[OK](#)

- Specify a date and time in the appropriate fields and enter the desired price in the **Price/Hour** field, enter the desired price.
- 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 all resources (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 (see [Cost Optimizer Dashboard](#)) in the header. Simultaneously, the updated accrued cost of a VM is displayed on the [Storage Volumes](#) page and the updated deployment cost is displayed on the Deployment Details page.

- 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. See [Images](#) for more context.

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 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 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.

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.

Account Name	Description	Billing Units	Enabled For	Actions
C3 Manual 1	C3 Manual Account 1	2 Billing Units	Provisioning, Reporting	Edit Delete
Master	Cost Optimizer Reporting	11 Billing Units	Reporting	Edit Delete
account		050	Provisioning, Reporting	Edit Delete
C3 Manual Plans		810	Provisioning, Reporting	Edit Delete

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, or [manage instance types](#) for 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 Cloud Region](#)
- [Configure an IBM Cloud Region](#)
- [Add an IBM Cloud Cloud Account](#)

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** link in 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**.

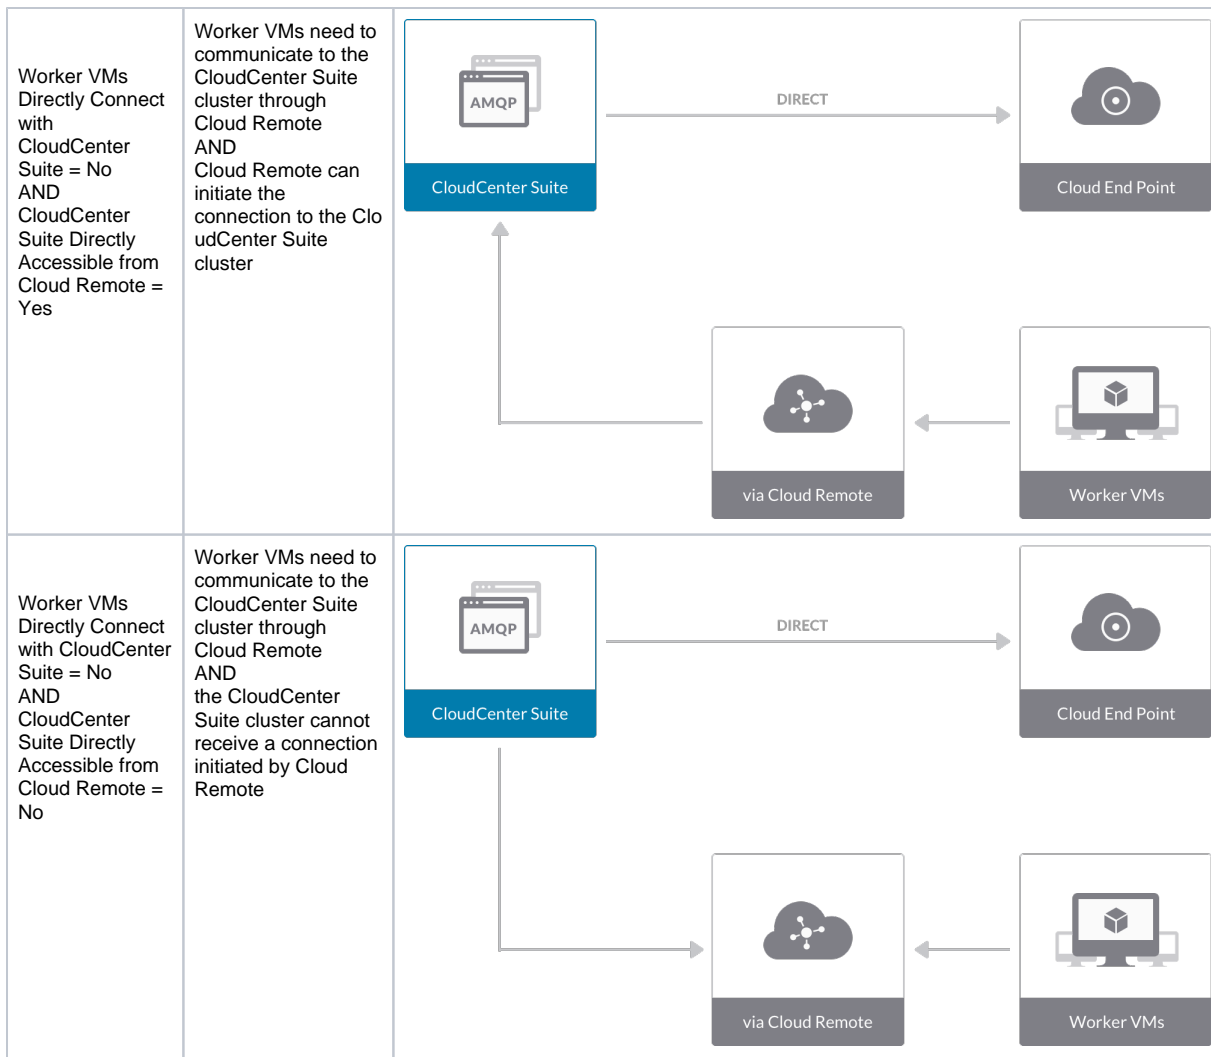


When assigning the Cloud Name, use a unique name. Ensure that you do not assign a previously-used name (including deleted cloud names) as the Cloud Name as they may cause deployments to fail intermittently.

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	<p>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</p> <p>Cloud Remote is not required</p>	<p>The diagram illustrates a central box labeled 'CloudCenter Suite' containing an 'AMQP' icon. Two arrows labeled 'DIRECT' point from this central box to two separate boxes on the right. The top box is labeled 'Cloud End Point' and contains a cloud icon with a play button. The bottom box is labeled 'Worker VMs' and contains a computer monitor icon with a cube on it.</p>



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.

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 on the 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 the region you added shown on the right side of the page.

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.



- Click the tab of the region you want to configure.
- 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**.

- 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 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.

The exception to this guidance is when a NAT firewall or proxy server exists between the CloudCenter Suite management cluster and worker 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 the Guacamole 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 the Guacamole server running in the CloudCenter Suite cluster.
Worker VMs must use a proxy server or NAT firewall to access the Guacamole 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.

- If any of the connectivity toggles in the 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

- 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. See [Cloud 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) > <i>Custom Port Numbers (Conditional)</i>). 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) > <i>Custom Port Numbers (Conditional)</i>).
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) > <i>Custom Port Numbers (Conditional)</i>).
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 Download Configuration 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.




Region Connectivity Enabling... Download Configuration **Copy Encryption Key** Edit Connectivity

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.

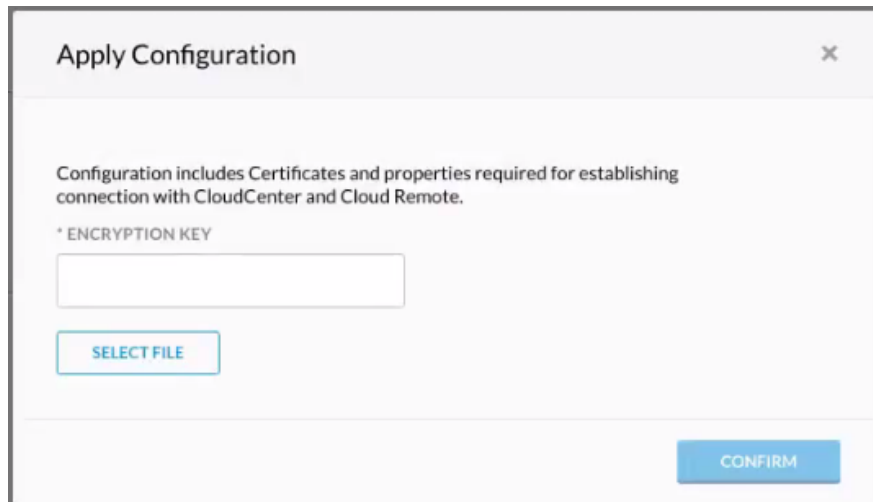
 If you change the connectivity settings in the CloudCenter Suite UI and need to again download the zip file, a new encryption key is automatically created and can be copied to the clipboard by clicking the **Copy Encryption Key** link again.

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

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



- 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.



Apply Configuration [X]

Configuration includes Certificates and properties required for establishing connection with CloudCenter and Cloud Remote.

* ENCRYPTION KEY

[Text Field]

SELECT FILE

CONFIRM

- Paste the encryption key that was copied to the clipboard into the **Encryption Key** field in the dialog box.
- Click **Select File** and browse to the artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
- Click **Confirm**.
- 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).

Region Connectivity Running		Download Configuration Configure Region
Cloud endpoint accessible from Cloud Center Manager	No	
Cloud Center Manager AMQP reachable from worker VM's	No	
Cloud Center Manager AMQP accessible from cloud	Yes	
Remote AMQP IP		
Worker AMQP IP	192.168.30.16:5671	
Blade Name	cloudcenter-blade-vmware-9-0289	
Blade Port	8443	

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

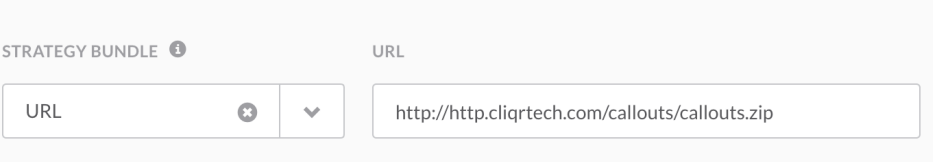
Workload Manager lets you define a VM naming strategy and an IP address management (IPAM) strategy on a per cloud region basis. The VM naming strategy lets Workload Manager generate unique names for each of your VMs as they are launched. The IPAM strategy lets Workload Manager generate IP addresses for 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 and an IP address deallocation callout script. An IP address allocation callout script may 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 and 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 CliQr 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 the Instance Naming Strategy field. 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 must meet the following requirements for job submission to succeed: <ul style="list-style-type: none"> Only contain the following characters: <ul style="list-style-type: none"> Lowercase a to z 0 to 9 Hyphen Start with an 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 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 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 options 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 your custom IPAM address deallocation script. See 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 for CliQr Macro Replacement:

Macro	Returned Value
%os1%	First character of OS type string in lower case
%os2%	First two characters of OS type string in lower case
%OS1%	First character of OS type string in upper case
%OS2%	First two 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
- Post VM 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 execute in 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) to specify external lifecycle actions at the region level.

External Lifecycle Actions	Edit External Lifecycle Actions
External Actions Bundle	
Pre VM Start	-
Pre VM Init	-
Post VM Init	-
Pre VM Stop	-
Post VM Stop	-

Click on the **Edit External Lifecycle Actions** link in the upper right. This displays the Edit 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 figure below.

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, see [Deployment Lifecycle Scripts](#).

- 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.

Instance Types [Sync Instance Types](#) [Add Instance Type](#)

Search:

Show 30 per page Page 1 of 1

Name	Instance Type	Price	Actions
m1.small	0	\$0/hr	Edit Delete
m1.large	0bbdff46-a647-4eab-827f-7367e1e6...	\$0/hr	Edit Delete
m2.small	809193e3-fb66-49a2-9c55-d318d2a...	\$0/hr	Edit Delete
m1.nano	ab1365a2-debc-44b1-b4cf-3d5dd15c...	\$0/hr	Edit Delete

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 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 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 the package 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.

Add Storage Type

Display Name *

Cost Per Month *

\$ 0.00 /GB

Cloud Storage Type ID *

Minimum Volume Size *

4 GB

Maximum Volume Size *

1024 GB

Maximum Throughput *

100 MB/s

Maximum IOPS *

100

Provisioned IOPS

Save Cancel

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 the [Cost 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.

Storage Types [Add Storage Type](#)

Show 30 per page Page 1 of 1

Name	Storage Type	Price	Actions
Local SSD scratch disk	local-ssd	\$0.218/GB-month	Edit Delete
Standard Persistent Disk	pd-standard	\$0.04/GB-month	Edit Delete
SSD Persistent Disk	pd-ssd	\$0.17/GB-month	Edit Delete

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 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 must manually import the corresponding physical images into your region and then map the corresponding logical images to these physical images. See [Images](#) for more context.

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.

The screenshot shows a dialog box titled "Add Cloud Account". It has a light gray header with the title. Below the header, there are four main sections:

- Name ***: A text input field with a red asterisk indicating it is required.
- Description**: A larger text area for optional details.
- Cloud Credentials**: A section header followed by two more required fields:
 - Account Name ***: A text input field with a red asterisk.
 - Account API Key ***: A text input field with a red asterisk.
- Buttons**: A blue "Connect" button is positioned below the API key field. At the bottom right of the dialog, there are "Save" and "Cancel" buttons.

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.

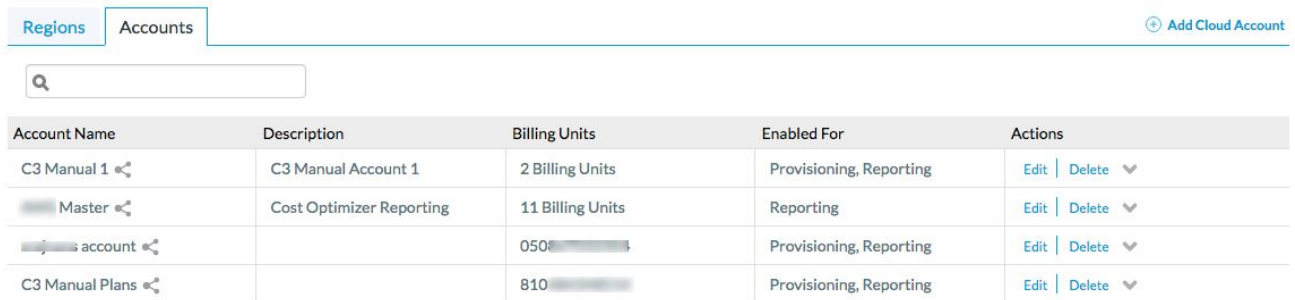
Value	Usage
Provisioning	Workload Manager can deploy jobs using this account.

Reporting	<p>Cost Optimizer and Workload Manager will track cloud costs for this account. Typical usage: master cloud accounts that are used for billing aggregation.</p> <div style="border: 1px solid #ccc; padding: 5px; margin: 5px 0;"> <p> It is recommended that you do not add a <i>Reporting</i> account to the same tenant through different cloud groups.</p> </div> <div style="border: 1px solid #ccc; padding: 5px; margin: 5px 0;"> <p> Enabling a public cloud account for <i>Reporting</i> may incur expenses to retrieve cost data. These expenses are proportional to the number of configured cloud accounts and regions.</p> </div>
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.



Account Name	Description	Billing Units	Enabled For	Actions
C3 Manual 1	C3 Manual Account 1	2 Billing Units	Provisioning, Reporting	Edit Delete
Master	Cost Optimizer Reporting	11 Billing Units	Reporting	Edit Delete
account		050	Provisioning, Reporting	Edit Delete
C3 Manual Plans		810	Provisioning, Reporting	Edit Delete

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, or [manage instance types](#) for 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](#)

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** link in 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	<p>The diagram illustrates a direct connection. On the left, a box labeled 'CloudCenter Suite' contains an 'AMQP' icon. An arrow labeled 'DIRECT' points from this box to a 'Cloud End Point' box on the right. Below the CloudCenter Suite box, another arrow labeled 'DIRECT' points to a 'Worker VMs' box at the bottom right.</p>
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	<p>The diagram illustrates indirect connectivity. On the left, a box labeled 'CloudCenter Suite' contains an 'AMQP' icon. An arrow labeled 'DIRECT' points from this box to a 'Cloud End Point' box on the right. Below the CloudCenter Suite box, a 'Worker VMs' box at the bottom right has an arrow pointing to a 'via Cloud Remote' box, which in turn has an arrow pointing to the 'CloudCenter Suite' box.</p>



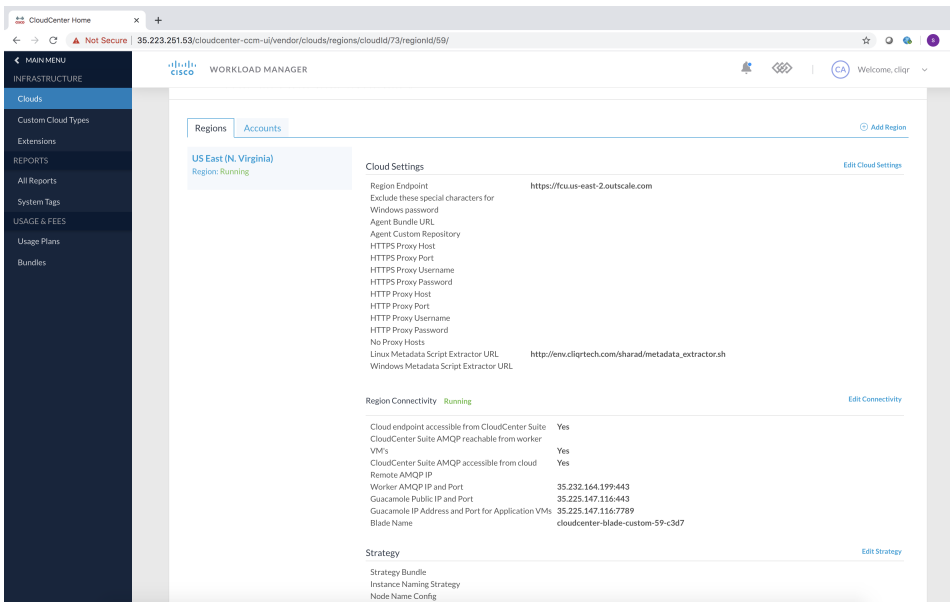
4. Click **Done** to save the configuration and close the dialog box. This brings you back to the **Clouds** page, and the cloud you just created will be added to the bottom of the list on the left side of the page.

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 the region you added shown on the right side of the page.

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** 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 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 a local 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 HTTPS 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.

Region Connectivity Running Download Configuration [Configure Region](#)

Clicking **Download Configuration** causes two things to happen:

- An encrypted zip file named artifacts.zip is 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.

Region Connectivity Enabling... Download Configuration Copy Encryption Key [Edit Connectivity](#)

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 from software.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 in [Proxy 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 with cloud (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 both Cloud 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**.

- 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 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. The exception to this guidance is when a NAT firewall or proxy server exists between the CloudCenter Suite management cluster and worker 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 the Guacamole 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 the Guacamole server running in the CloudCenter Suite cluster.
Worker VMs must use a proxy server or NAT firewall to access the Guacamole 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.

- If any of the connectivity toggles in the **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

- Obtain the Cloud Remote shared AMI from Cisco support and launch it. Follow the same guidance for obtaining and launching the [CloudCenter Suite installer appliance for AWS](#).
- 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.
- 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) > <i>Custom Port Numbers (Conditional)</i>). 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) > <i>Custom Port Numbers (Conditional)</i>).
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) > <i>Custom Port Numbers (Conditional)</i>).
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 Download Configuration 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 to Cloud Remote web UI.


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



- 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.



- Paste the encryption key that was copied to the clipboard into the **Encryption Key** field in the dialog box.
- Click **Select File** and browse to the artifacts.zip file that you downloaded through the CloudCenter Suite web UI and select it.
- Click **Confirm**.
- 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).

Region Connectivity Running		Download Configuration Configure Region
Cloud endpoint accessible from Cloud Center Manager	No	
Cloud Center Manager AMQP reachable from worker VM's	No	
Cloud Center Manager AMQP accessible from cloud	Yes	
Remote AMQP IP		
Worker AMQP IP	192.168.30.16:5671	
Blade Name	cloudcenter-blade-vmware-9-0289	
Blade Port	8443	

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 lets Workload 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 and an IP address deallocation callout script. An IP address allocation callout script may 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 and add the path in the new field that appears to the right.
	<div style="border: 1px solid #ccc; padding: 5px;"> <div style="display: flex; justify-content: space-between;"> STRATEGY BUNDLE ⓘ URL </div> <div style="display: flex; align-items: center; margin-top: 5px;"> <div style="border: 1px solid #ccc; padding: 2px 5px; flex-grow: 1;">URL</div> <div style="margin: 0 5px;">+</div> <div style="border: 1px solid #ccc; padding: 2px 5px; flex-grow: 1;">http://http.cliqrtech.com/callouts/callouts.zip</div> </div> </div>

Instance Naming Strategy	Dropdown for selecting one of three naming strategies: Default, CliQr Macro Replacement, and Hostname Callout. If you select CliQr 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 the Instance Naming Strategy field. 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 must meet the following requirements for job submission to succeed: <ul style="list-style-type: none"> Only contain the following characters: <ul style="list-style-type: none"> Lowercase a to z 0 to 9 Hyphen Start with an 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 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 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 options 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 your custom IPAM address deallocation script. See 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 for CliQr Macro Replacement:

Macro	Returned Value
%os1%	First character of OS type string in lower case
%os2%	First two characters of OS type string in lower case
%OS1%	First character of OS type string in upper case
%OS2%	First two 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.
<div style="border: 1px solid #ccc; border-radius: 5px; background-color: #fff9c4; padding: 10px; display: inline-block;"> Only macro whose value can be user-customized. </div>	

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
- Post VM 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 execute in 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) to specify external lifecycle actions at the region level.

External Lifecycle Actions	Edit External Lifecycle Actions
External Actions Bundle	
Pre VM Start	-
Pre VM Init	-
Post VM Init	-
Pre VM Stop	-
Post VM Stop	-

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

Edit External Lifecycle Actions
✕

EXTERNAL ACTIONS BUNDLE ⓘ

▼

PRE VM START

▼

PRE VM INIT

▼

POST VM INIT

▼

PRE VM STOP

▼

POST VM STOP

▼

DONE

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 figure below.

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, see [Deployment Lifecycle Scripts](#).

- 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.

Instance Types [Sync Instance Types](#) [Add Instance Type](#)

Search: Show 30 per page Page 1 of 1

Name	Instance Type	Price	Actions
m1.small	0	\$0/hr	Edit Delete
m1.large	0bbdff46-a647-4eab-827f-7367e1e6...	\$0/hr	Edit Delete
m2.small	809193e3-fb66-49a2-9c55-d318d2a...	\$0/hr	Edit Delete
m1.nano	ab1365a2-debc-44b1-b4cf-3d5dd15c...	\$0/hr	Edit Delete

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.

- 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 the package 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.

Storage Types [Add Storage Type](#)

Search: Show 30 per page Page 1 of 1

Name	Storage Type	Price	Actions
Local SSD scratch disk	local-ssd	\$0.218/GB-month	Edit Delete
Standard Persistent Disk	pd-standard	\$0.04/GB-month	Edit Delete
SSD Persistent Disk	pd-ssd	\$0.17/GB-month	Edit Delete

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.

- 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 must manually import the corresponding physical images into your region and then map the corresponding logical images to these physical images. See [Images](#) for more context.

Prerequisites

Before adding an Outscale 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 Outscale cloud account, follow this procedure.

- Locate your Outscale 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.
- Assign a cloud account **Name**.



Tip

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

- 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:
 - Outscale Account Number:** The account number from your Outscale account.
 - Outscale Access Key and Secret Key:** The security credentials to access this Outscale account.
- Click the **Connect** button. CloudCenter Suite will now attempt to validate your account credentials.
- 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. <div style="border: 1px solid #ccc; padding: 5px; margin-top: 5px;"> <p> It is recommended that you do not add a <i>Reporting</i> account to the same tenant through different cloud groups.</p> </div> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 5px;"> <p> Enabling a public cloud account for <i>Reporting</i> may incur expenses to retrieve cost data. These expenses are proportional to the number of configured cloud accounts and regions.</p> </div>
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.

Account Name	Description	Billing Units	Enabled For	Actions
C3 Manual 1	C3 Manual Account 1	2 Billing Units	Provisioning, Reporting	Edit Delete
Master	Cost Optimizer Reporting	11 Billing Units	Reporting	Edit Delete
Account		050	Provisioning, Reporting	Edit Delete
C3 Manual Plans		810	Provisioning, Reporting	Edit Delete

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, or [manage instance types](#) for the cloud account.

Cloud Maintenance

Cloud Maintenance

- [Sharing Cloud Regions and Cloud Accounts](#)
- [Deleting Clouds, 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 in [Tenant 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 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:

Clouds + Add Cloud

Could not delete region
Cloud Region "US East (Ohio)" cannot be deleted as its cloud group has associated cloud accounts.

Amazon asw2 Regions

asw2 Configure Cloud US East (Ohio) Delete Region

Cloud Accounts
• jb

Region: **Running**
1 Enabled Users

+ Add Cloud Account + Add Region

Similarly, before you can delete a single region cloud you must first delete all cloud accounts associated with that cloud. Otherwise, you will see an error message as shown below:

Clouds + Add Cloud

Could not delete cloud
Cloud Group [kubebj] has associated cloud accounts. Please delete the cloud accounts before trying to delete the cloud group.

Amazon kubernetes Configure Cloud Delete Cloud

asw2 Region: **Running**

AzureRM 0 Enabled Users

Before you can delete a cloud account you must first remove that cloud account from all [deployment environments](#) in which it is used. Otherwise, an error message as shown below is displayed:

asw2 Back to Clouds

Failed to delete cloud account.
Cloud Account is associated with the Deployment Environment.

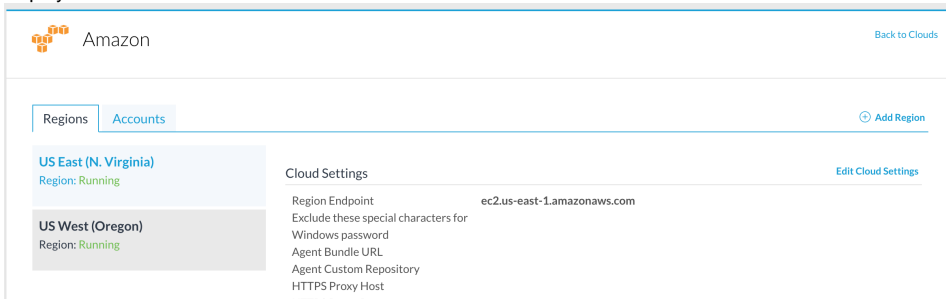
Regions Accounts + Add Cloud Account

Account Name Description Billing Units Enabled For Estimated Month to Date Cost Actions

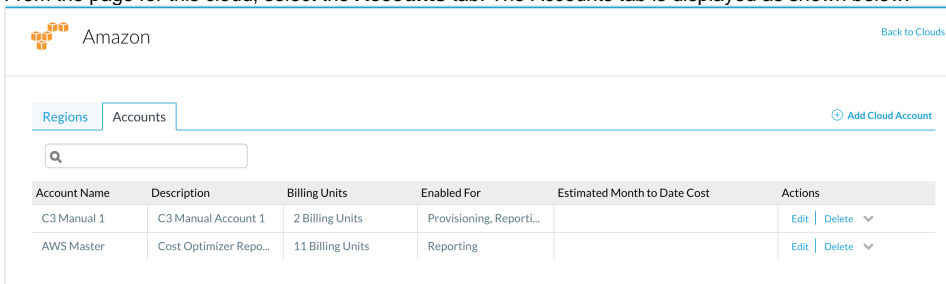
jb		068685977692	Provisioning, Reporting		Edit Delete
----	--	--------------	-------------------------	--	---------------

Therefore, to delete a cloud follow these steps:

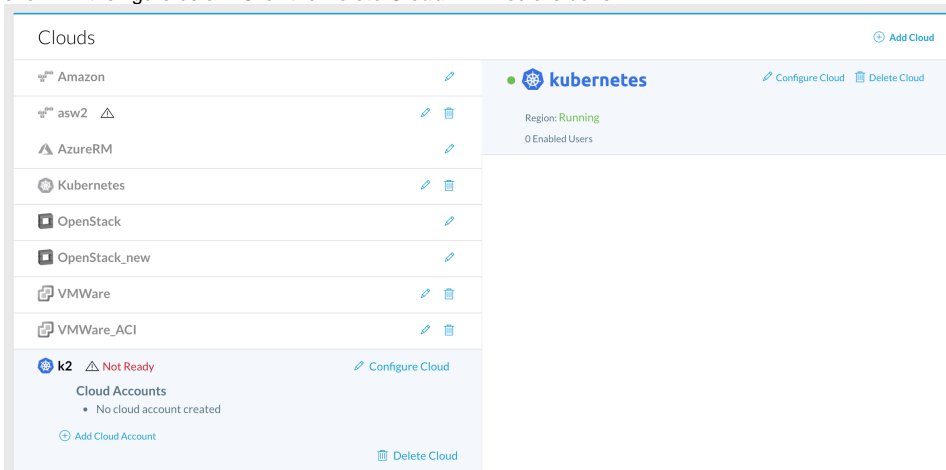
1. From the Clouds page, select the cloud and click its **Configure Cloud** link which displays the page for this cloud. The page for this cloud will be displayed as shown below.



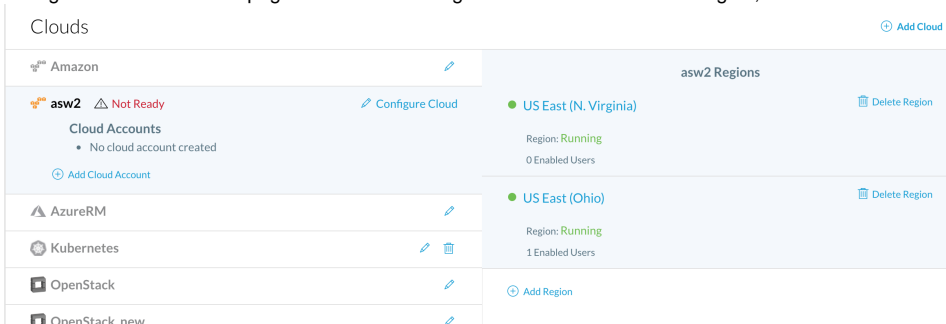
2. From the page for this cloud, select the **Accounts** tab. The Accounts tab is displayed as shown below.



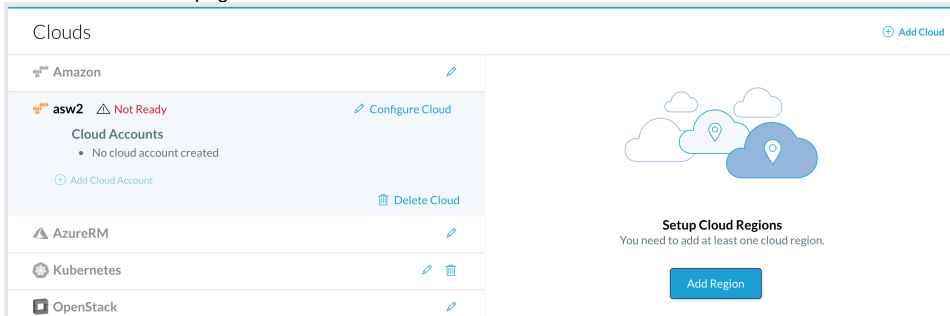
3. From the Accounts tab, delete all accounts one by one by clicking the **Delete** link in the Actions column. If an error about deployment environments appears, click on the **Main Menu > Environments** menu tab, browse the deployment environments for any references to the account, and remove the account 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 **Delete Cloud** link for that cloud will appear on the left side of the Clouds page, as shown in the figure below. Click the **Delete Cloud** 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.



The screenshot displays the 'Clouds' management interface. On the left, a list of cloud providers is shown: Amazon, asw2 (Not Ready), AzureRM, Kubernetes, and OpenStack. The 'asw2' cloud is selected, showing its 'Cloud Accounts' section with the message 'No cloud account created' and an 'Add Cloud Account' link. A 'Delete Cloud' link is also visible. On the right, a 'Setup Cloud Regions' prompt is displayed, stating 'You need to add at least one cloud region.' and featuring an 'Add Region' button.

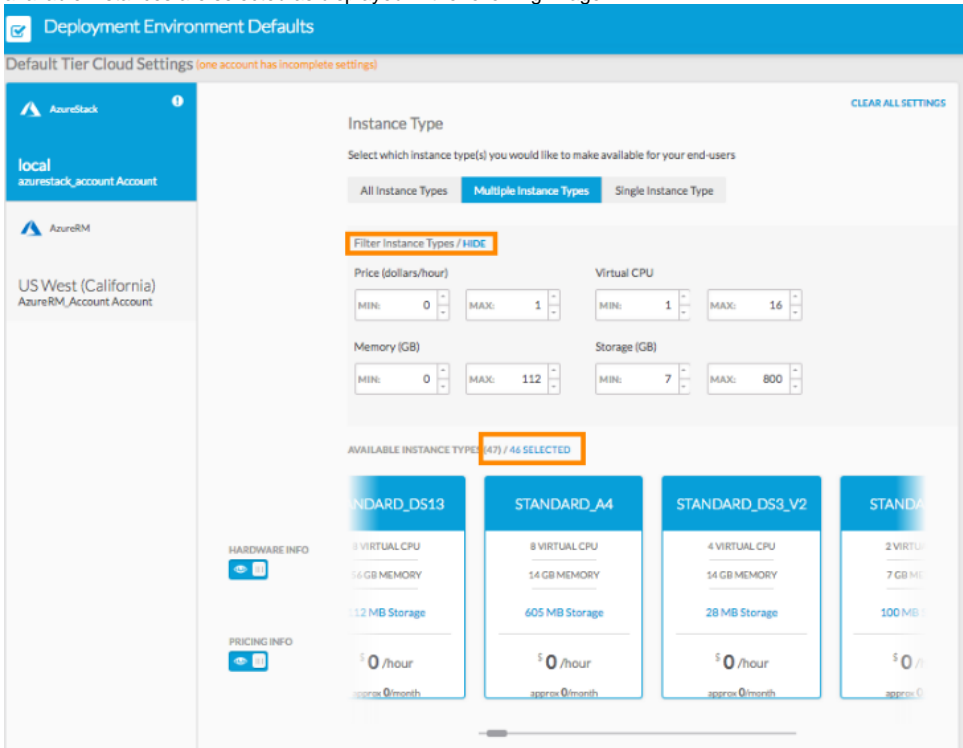
Manage Instance Types

Manage Instance Types

- [Overview](#)
- [Sample Instance Types](#)
- [Add Instance Types](#)
- [Custom Pricing](#)
- [Sync Instance Types \(Manual\)](#)
- [Auto-Sync Instance Types](#)
- [Instance Types Not Visible?](#)

The Workload Manager only displays application VM instance types that were selected by your administrator during the image mapping process:

- The Workload Manager filters the Instance Type dropdown to automatically list only those instances provided during the image mapping process. Even within the displayed list, the administrator has the option to select the instance types that should be displayed to end-users. By default, all available instances are selected as displayed in the following image.



- The Workload Manager does not display any instance type if the selected cloud does not have image mapping or if the image is not compatible with the selected cloud.
- The instance types are dynamically configured for public clouds. Once you add the instance type(s) for your deployment, you see the Configure Instance Types section display the configured instance(s).
- The administrator can edit mapped images to support additional Instance Types.
- When adding an instance type to a cloud region, users can specify a zero disk size as local storage. If 0 (zero), then the disk in the image is used and no additional disks are deployed to the VM.
- Support for instance type and image sync is available for AWS, AzureRM, and Google.
- The Instance Type reflects MilliCPUs for a Kubernetes container and (Virtual) CPUs for all other clouds.
 - The **Instance Type Storage** field is reflected in the Instance Type card cloud as applicable for each cloud.
 - The storage label for instance types differs based on cloud as listed in the following table:

Cloud	Additional Storage Label in the Instance Type Card
AWS	TEMP STORAGE (Temporary Storage)
AzureRM	For example: 2GB TEMP STORAGE
OpenStack	ROOT DISK For example: 2GB ROOT DISK

Other Clouds	This field is not displayed
--------------	-----------------------------

Enter the instance type(s) for your enterprise. The following screenshot displays some sample instance types defined in Workload Manager :

Instance Type	Name	Hardware Spec	Cost Per Hour	Actions
small	small	2CPUs, 2048GB, 1000Mbps	0.2	Edit Delete
medium	medium	2CPUs, 4096GB, 1000Mbps	0.3	Edit Delete
large	large	4CPUs, 8192GB, 1000Mbps	0.6	Edit Delete
xlarge	xlarge	4CPUs, 16384GB, 1000Mbps	1.2	Edit Delete
2xlarge	2xlarge	8CPUs, 32768GB, 1000Mbps	2.4	Edit Delete

The Add instance Type feature is only supported for vCenter and OpenStack clouds.

To add an instance type, follow this procedure:

1. Click the **Admin** link displayed in the Workload Manager UI main menu.
2. Click the **Clouds** tab in the side panel to display the Cloud list page.
3. Click the **Configure** link for the required cloud.
4. For OpenStack clouds, select the **Regions** tab and then the required region. For vCenter clouds, select the **Details** tab.
5. Scroll down to the Instance Types section and click **Add Instance Type** to add a new Instance Type.
6. Complete the required fields and click **Save**.

Administrators can perform the following actions for the selected cloud account and region:

- Set a Price Adjustment value, which will change the custom price for all the instance types. So, users don't have to change it individually for all the instance types.
- Enable/disable instance types functionality.

To set custom instance type pricing for an account/region combination, follow this procedure:

1. Navigate to **Admin > Clouds**, then click **Configure Cloud** for the required cloud.
2. Select the **Accounts** tab.
3. In the Actions column for the required cloud account, click the dropdown icon and then click **Manage Instance Types**.

4. If your cloud has only one region configured, you are brought directly the Cloud Instance Type dialog box, as shown in the screenshot below. If you have multiple regions configured, you must first select the required region from the Select Region dialog box, and then the Cloud Instance Type dialog box is displayed for that region.

NAME	SPECIFICATIONS	DEFAULT PRICE	CUSTOM PRICE	ENABLE
c1.medium	1.00GB RAM, 2 CPUs, 2 NICs, 350GB Storage, 32, 64 Bit Supported	J\$0.26/hr	J\$0.26/hr	<input checked="" type="checkbox"/>
c1.xlarge	7.00GB RAM, 8 CPUs, 4 NICs, 1680GB Storage, 64 Bit Supported	J\$1.04/hr	J\$1.04/hr	<input type="checkbox"/>
c3.2xlarge	15.00GB RAM, 8 CPUs, 4 NICs, 160GB Storage, 64 Bit Supported	J\$0.84/hr	J\$0.84/hr	<input type="checkbox"/>
c3.4xlarge	30.00GB RAM, 16 CPUs, 8 NICs, 320GB Storage, 64 Bit Supported	J\$1.68/hr	J\$1.68/hr	<input type="checkbox"/>
c3.8xlarge	60.00GB RAM, 32 CPUs, 8 NICs, 640GB Storage, 64 Bit Supported	J\$3.36/hr	J\$3.36/hr	<input type="checkbox"/>
c3.large	3.00GB RAM, 2 CPUs, 3 NICs, 32GB Storage, 32, 64 Bit Supported	J\$0.21/hr	J\$0.21/hr	<input type="checkbox"/>
c3.xlarge	7.00GB RAM, 4 CPUs, 4 NICs, 80GB Storage, 64 Bit Supported	J\$0.42/hr	J\$0.42/hr	<input type="checkbox"/>
c4.2xlarge	15.00GB RAM, 8 CPUs, 4 NICs, 64 Bit Supported	J\$0.8/hr	J\$0.8/hr	<input type="checkbox"/>
c4.4xlarge	30.00GB RAM, 16 CPUs, 8 NICs, 64 Bit Supported	J\$1.59/hr	J\$1.59/hr	<input type="checkbox"/>

5. Configure the **Price Adjustment** Value. By default, the value defaults to 100%, which indicates that the custom prices across all instance types available in the selected region will be the same as the default price. Change the Price adjustment value as required. The custom price field for all the instance types shown in the grid updates accordingly.

- The custom price of an instance = Price Adjustment value * Default price of instance/100
- Examples:
 - Price adjustment value = 10, the custom prices for all instance types will be 10% of Default price of instance type.
 - Price adjustment value = 200, the custom prices for all instance types will be 200% of Default price of instance type.
 - Valid values for Price adjustment are between (0-1000 both excluded, using up to 2 decimal places).

6. Click **OK** to save the Price Adjustment value.

7. Review the enable/disable state of each instance type and toggle the **ON/OFF** switch to change any instance type.

8. Repeat this process for other regions as applicable for your deployment.

This feature is supported for OpenStack Clouds only. Administrators have the ability to import all instance flavors defined in OpenStack as CloudCenter Suite instance types by clicking the **Sync Instance Types** link in the upper right of the instance types section of the Regions tab.

This feature is only available for public VM cloud providers.

A scheduled background task automatically syncs instance types and their corresponding costs once every 24 hours.

For public VM-based clouds, administrators:

- Cannot edit the details of the instance (except for price for AzureRM and Google) or delete it from the list.
- Cannot add any instance types manually.

This *Add Instance Type feature* (custom addition) was allowed for OpenStack environments in case users do not want to sync all instance types. In such cases, users can manually add the instance type using the same name that was defined in your OpenStack environment.



Instance types are not specific to each region. Be aware that some instances may not be available in some regions.

If the available instance types are not visible in the Instance Type dropdown for a given Base OS Image, verify the following settings:

- Base image mapping – Verify the following information when mapping the images (see [Images Page](#)):
 - The service is using the required [OOB Logical Images](#) and that you have defined your image mapping.
 - When you select the instance types for an image (either when you Add Cloud Mapping or Edit Cloud Mapping), make sure that the required instance types are mapped.
- Hardware specifications – Verify that your specifications are lower than the targeted instance type settings. See the Sample Instance Type sections above for some examples.
- Architecture settings for the instance type – Try selecting the Both option to indicate support for both 32-bit and 64-bit architecture and resubmit the deployment.
- For existing, private clouds (VMware and OpenStack) created prior to CloudCenter Suite 5.1, a user would need to click **Edit Instance Type** and save it again without modifying any fields. This step is required if you want to deploy a Windows application. For new cloud environments (VMware and OpenStack) added as part of CloudCenter Suite 5.1, this step is not needed.

Images

Images

- [Images Overview](#)
- [Images Page](#)
- [Regions Tab Image Settings Section](#)


Images Overview

Images Overview

Workload Manager includes a data structure called the logical image which is used for abstracting the physical VM image which resides in the target VM-based cloud. Workload Manager comes with several [OOB logical images](#) corresponding to various Linux and Windows releases and versions of MS SQL on Windows. In addition, users with the WM_ADMIN or WM_IMAGE_MANAGER role can create their own custom logical images corresponding to a base OS with other software installed.

All VM-based services in Workload Manager require a logical image as their starting point. When a VM-based service is deployed to a cloud region, as part of the deployment process, the logical image for each VM-based tier is translated to a corresponding physical VM image in that cloud region. This requires that the logical image be "mapped" to the corresponding physical image in the cloud region before deploy time. Workload Manager handles logical to physical image mapping as follows:

- All Workload Manager OOB logical images are automatically mapped to the corresponding physical cloud images for each [supported public cloud region](#) that you add to a cloud.
- For [private VM-based clouds](#), you must manually map the OOB logical images to physical images that you import into the private cloud.
- For custom logical images added to your tenant, regardless of whether the physical image resides in a public or private cloud, you must manually map the custom logical image to the physical images. For public clouds, the physical image could be one you imported yourself or one that was shared with your account from someone else. For private clouds the physical image is always imported.

 Since all VM-based [Services](#) are based on a logical images. If a service's logical image is not properly mapped to the correct physical cloud image for a particular cloud region, that service will not be available for deployment in that cloud region.

 A [pre-bootstrapped physical image](#) must be associated with a logical image before it can be used in a deployment.

Adding, editing and deleting logical images, and adding, editing and deleting image mappings, can be performed through the [Images Page](#).

Adding, editing and deleting image mappings, can also be performed from the [Image Settings section of the Regions tab](#).

Images Page

Images Page


- [Overview](#)
- [Add an Image Mapping](#)
- [Edit or Delete an Image Mapping](#)
- [Add a Logical Image](#)
- [Edit or Delete a Logical Image](#)
- [Share a Logical Image](#)

The Images page is the main UI screen for managing logical images and their mapping to physical cloud images. It is reached by selecting the Images tab from the main menu. The Images tab is hidden unless your login is associated the WM_ADMIN role or the WM_IMAGE_MANAGER role (see [OOB Groups, Roles, and Permissions](#)).

From the Images page you can:

- View all logical images and their mappings to physical images
- Add, edit, or delete image mappings
- Add, edit, delete, or share logical images

The Images page appears in the figure below.



The page consists of a list of rows with each row representing a logical image. If a logical image has any mappings to physical images, an expand icon appears to the left of the image name and the number of mappings defined for that image is displayed to the right of the image name.

Once you add a public cloud in Workload Manager and add regions to that cloud, for each OOB logical images, Workload Manager automatically creates the image mapping for that logical image to the corresponding physical image in that region. This is done based on data maintained by Cisco and stored in the Workload Manager public package store. No user involvement is needed to create these mappings.

All logical images are tenant owned objects. The OOB logical images are owned by the root tenant and are automatically shared with all subtenants.

You must manually add an image mapping to a logical image to support either of these two cases:

- Any physical image on a private cloud.
- Any custom physical image on a public cloud. This includes all [Pre-bootstrapped Images](#).



If your custom physical image does not correspond to one of the Workload Manager OOB logical images, you must first create a new logical image corresponding to your custom physical image. To do this, follow the instruction to add a logical image, below. If your custom image is a pre-bootstrapped image that does correspond to one of the OOB logical images, do not create a new logical image; instead, use the existing OOB logical image.

To add an image mapping to an existing logical image listed in the Images page, follow this procedure.

1. Hover over the row for the logical image. An **Add Mapping** button will appear on the right side of the row. Click it. This brings up the Add Cloud Mapping dialog box as shown in the figure below.

2. Select the cloud region from the dropdown. Only cloud regions already defined and not currently having an image mapping for this logical image will be displayed in the dropdown.
3. Enter the physical cloud image ID. To ensure you are specifying the correct image ID, use the guidance below depending on the cloud provider.

- **<VM name >|<snapshot name>**

You have two options to configure the Image ID for VMware:

- **Snapshot:** If using snapshots, add a folder in vSphere (to store your Workload Manager snapshots), name it **CliqrTemplat es**, and add this snapshot to the **CliqrTemplates** folder.
- **Template:** You can alternately use template names to configure the cloud image. In this case, specify the name of a VM or template as the image ID on the VMware console and the systems always performs a full clone to either a specified datastore or datastore cluster. Add this Template to the **CliqrTemplates** folder.

The full clone is performed on the source VM or VM template, the cloned VM can be on either datastore or datastore cluster that user specifies.

See [VMware Configurations](#) for additional context.

-
- QCOW2 Image ID (sample ID mapping highlighted in the following screenshots):


- Create, customize, and deploy a VM using the required image.
- Shut down the VM instance and clone the instance.

- In Google Cloud, launch a VM for the Workload Manager instance and click **REST** at the end of the page. You can view the **sourceImage** value in the REST output.

```
{
  "disks": [{
    "type": "PERSISTENT",
    "boot": true,
    "mode": "READ_WRITE",
    "autoDelete": true,
    "deviceName": "instance-1",
    "initializeParams": {
      "sourceImage": "https://www.googleapis.com/compute/v1/projects/centos-cloud/global/images/centos-7-v20160418",
      "diskType": "projects/x-signifier-537/zones/us-central1-f/diskTypes/pd-ssd",
      "diskSizeGb": "10"
    }
  }]
}
```

The following procedure allows you retrieve the image details using the source to map the Workload Manager image as Google Cloud Platform's dynamic bootstrapping feature allows you to temporarily access an earlier version of the image by using the REST source details. To using this method, follow this procedure.

- Access the Google Cloud Platform Compute Engine page and click the **Create Instance** link.



- b. In the Create an instance page, click **Change** in the Boot Disk field.

- c. Select one of the following options (dynamic bootstrapping is available for these options) as required for your environment and save your change: CentOS6 or 7, Ubuntu 14:04, Redhat Enterprise Linux 6 or 7, Windows 2012
- d. Back in the Create an instance page, click the **REST** link. The Equivalent REST request is displayed in the resulting popup.
- e. Scroll down to the **sourceImage** line and select the key displayed in this line.

Equivalent REST request

This is the REST request with the parameters you have selected.

```

{
  "metadata": {
    "items": []
  },
  "tags": {
    "items": []
  },
  "disks": [
    {
      "type": "PERSISTENT",
      "boot": true,
      "mode": "READ_WRITE",
      "autoDelete": true,
      "deviceName": "instance-1",
      "initializeParams": {
        "sourceImage": "https://www.googleapis.com/compute/v1/projects/windows-cloud/global/images/windows-server-2012-r2-x64-disk",
        "diskType": "projects/training4-978/zones/asia-east1-a/diskTypes/pd-standard",
        "diskSizeGb": "50"
      }
    }
  ],
  "canIpForward": false,
  "networkInterfaces": [
    {
      "network": "projects/training4-978/global/networks/default",

```

Line wrapping

[REST API reference](#)

CLOSE

f. Copy this key and paste it in the Image ID field in the Workload Manager UI's Image Mapping page.

- In AzureRM, the following command output (latest version) provides the Image ID required by the Workload Manager. Refer to <https://docs.microsoft.com/en-us/azure/virtual-machines/linux/cli-ps-findimage> for additional context.

a. **Standard AzureRM Image** – The following example queries all CentOS 7.2 images, the Image ID of the latest version is **OpenLogic:CentOS:7.2:7.2.20170105**.

```
$ az vm image list -p OpenLogic --offer CentOS --sku 7.2 --all | more
[
  {
    "offer": "CentOS",
    "publisher": "OpenLogic",
    "sku": "7.2",
    "urn": "OpenLogic:CentOS:7.2:7.2.20160303",
    "version": "7.2.20160303"
  },
  {
    "offer": "CentOS",
    "publisher": "OpenLogic",
    "sku": "7.2",
    "urn": "OpenLogic:CentOS:7.2:7.2.20160308",
    "version": "7.2.20160308"
  },
  {
    "offer": "CentOS",
    "publisher": "OpenLogic",
    "sku": "7.2",
    "urn": "OpenLogic:CentOS:7.2:7.2.20160620",
    "version": "7.2.20160620"
  },
  {
    "offer": "CentOS",
    "publisher": "OpenLogic",
    "sku": "7.2",
    "urn": "OpenLogic:CentOS:7.2:7.2.20161026",
    "version": "7.2.20161026"
  },
  {
    "offer": "CentOS",
    "publisher": "OpenLogic",
    "sku": "7.2",
    "urn": "OpenLogic:CentOS:7.2:7.2.20170105",
    "version": "7.2.20170105"
  },
  {
    "offer": "CentOS",
    "publisher": "OpenLogic",
    "sku": "7.2n",
    "urn": "OpenLogic:CentOS:7.2n:7.2.20160629",
    "version": "7.2.20160629"
  }
]
```

b. **Custom AzureRM Image:**



Image ID depends on the CloudCenter version

The Image ID differs based on the Workload Manager version – use the Resource ID of the image as Image ID (as it includes the new Azure SDK).



Managed Store Options

To launch custom AzureRM images, you must select one of the managed storage options listed (Premium or Standard).

See [Multiple Volumes](#) > *AzureRMType Nuances* for additional context.

- i. The following screenshot displays the Image ID retrieval screen via the AzureRM UI.

The screenshot shows the AzureRM UI for an image named 'testimage2'. The details are as follows:

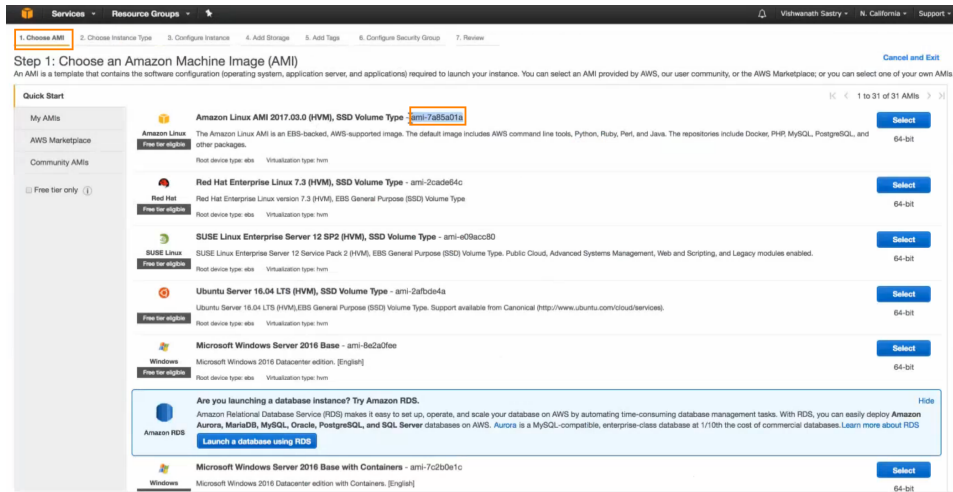
- NAME:** testimage2
- SOURCE VIRTUAL MACHINE:** gaoqi-centos73-1
- OS DISK:**

OS TYPE	SOURCE BLOB URI	ACCOUNT TYPE	CACHING
Linux	https://xuuswestpremium.blob.core.windows.net/vhds/gaoqi-centos73-120170619152212.vhd		Read/write
- DATA DISKS:** This image doesn't contain any data disks.
- RESOURCE GROUP:** gaoqi-west
- LOCATION:** West US
- SUBSCRIPTION:** Pay-As-You-Go(Converted to EA)
- RESOURCE ID:** /subscriptions/cbaba14b-e672-47d7-bb59-4a0613d6d149/resourceGroups/gaoqi-west/providers/Microsoft.Compute/images/testimage2

- ii. The following screenshot displays the Image ID retrieval via the AzureRM CLI.

```
Microsoft Azure Resource groups Search resources
Bash Azure CLI Documentation
cliqz@Azure:~$
cliqz@Azure:~$ az image list
{
  "id": "/subscriptions/0f2c89bc-0aa4-41f6-838b-3dcbcd17c166/resourceGroups/TEST-CC/providers/Microsoft.Compute/images/MyCCMVHDImage",
  "location": "westus",
  "name": "MyCCMVHDImage",
  "provisioningState": "Succeeded",
  "resourceGroup": "TEST-CC",
  "sourceVirtualMachine": null,
  "storageProfile": {
    "dataDisks": [],
    "osDisk": {
      "blobUri": "https://testab123.blob.core.windows.net/vhd/AzureCloudCCM.vhd",
      "caching": "ReadWrite",
      "diskSizeGb": null,
      "managedDisk": null,
      "osState": "Generalized",
      "osType": "Linux",
      "snapshot": null
    }
  },
  "tags": null,
  "type": "Microsoft.Compute/images"
},
{
  "id": "/subscriptions/0f2c89bc-0aa4-41f6-838b-3dcbcd17c166/resourceGroups/TEST-CC/providers/Microsoft.Compute/images/MyCCOVHDImage",
  "location": "westus",
  "name": "MyCCOVHDImage",
  "provisioningState": "Succeeded",
  "resourceGroup": "TEST-CC",
  "sourceVirtualMachine": null,
  "storageProfile": {
    "dataDisks": [],
    "osDisk": {
      "blobUri": "https://testab123.blob.core.windows.net/vhd/AzureCloudCCO.vhd",
      "caching": "ReadWrite",
      "diskSizeGb": null,
      "managedDisk": null,
      "osState": "Generalized",
      "osType": "Linux",
      "snapshot": null
    }
  },
  "tags": null,
  "type": "Microsoft.Compute/images"
}
}
```

- In AWS, the Image ID is the exact name displayed in the Machine Image (AMI) page in the AWS cloud portal.



- Expand **Advanced Instance Type Configuration** and add or remove instance types as needed.
- Click **Save** to save the changes and close the dialog box. The new mapping will be displayed in the list of cloud mappings for that logical image.

Once a mapping has been added to a logical image, if the image is owned by your tenant, you can edit it or delete it.

From the row of the logical image of interest, click the expand icon to reveal the list of current cloud image mappings for that logical image. The list of mappings will appear as shown in the figure below.

To delete a mapping, hover over the row for that mapping to cause the trash icon to appear on the right side of the row. Click the trash icon, then acknowledge the confirmation message to delete the mapping.

To edit a mapping, click on the name of the associated cloud region. This causes the Edit Cloud Mapping dialog box to appear. This dialog box is similar to the Add Cloud Mapping dialog box except the cloud region cannot be changed. Adjust the image ID and allowed instance types as needed and click Save.

If you need to deploy a custom physical cloud image that does not correspond to any of the OOB logical images, you need to add a new logical image and then map that new logical image to the custom physical image on your cloud of choice. To create a new logical image from the Images page, follow this procedure.

- Click the **Add Image** button in the upper right of the page. This causes the **Add a New Image** dialog box to be displayed as shown in the figure below.

- Enter values for image name, number of network interfaces, and OS type. Then set the toggle to enable the image.

3. Click **Save** to save the image and close the dialog box. The newly added image is now displayed as the list of logical images.
4. Add the necessary cloud mappings to the image as explained above in [Add an Image Mapping](#).

If a logical image is owned by your tenant, you can edit it or delete it.

To delete a logical image, hover over the row for that image to cause the menu dropdown icon to appear on the right side of the row. Then, click the dropdown icon and select the delete command. Acknowledge the confirmation message to delete the mapping.



Use caution when deleting logical images. Any services associated with a deleted image will stop working. And any users in subtenants that were depending on those logical images will no longer have access to them. If you are the root tenant administrator you also have the ability to delete the OOB logical images. Use extra caution when deleting an OOB logical image as that would break the OOB services that use that image.

To edit a logical image, hover over the row for that image to cause the menu dropdown icon to appear on the right side of the row. Then, click the dropdown icon and select the edit command. **This dialog box is similar to the Add a New Image dialog box except the image name cannot be changed. Adjust the number of network interfaces, OS type, and the Enable toggle as needed and click Save.**

If a logical image is owned by your tenant, you can share it with users, groups or subtenants, giving them view access. All OOB logical images are automatically shared with view access with all subtenants.

To share a logical image from the Images page, follow this procedure.

Hover over the row for that image to cause the menu dropdown icon to appear on the right side of the row. Then, click the dropdown icon and select the share command. This causes the Access Control List dialog box to be displayed.

From the dialog box, select the users, groups, and/or subtenants to share this image with. See [Permission Control](#) for additional context.

Regions Tab Image Settings Section

Regions Tab Image Settings Section

- [Overview](#)
- [Edit or Delete a Mapping](#)
- [Add a Mapping](#)
- [Sync Image Mappings](#)

In addition to the Images page, it is also possible to map logical images to physical images for a particular region through the Image Settings section of the Regions tab (see figure below).

Image Mappings

[Sync Image Mappings](#)

Show per page
 Page of 1

Name ▲	Cloud Image ID	Actions
Bare Metal Ubuntu 12.04		Add Mapping
Callout Workflow		Add Mapping
CentOS 6.x	ami-23285c35	Edit Mapping Delete Mapping
CentOS 7.x	ami-95096eef	Edit Mapping Delete Mapping
Cloud Image Helper		Add Mapping
Common AWS External Services	ami-0204ca6a	Edit Mapping Delete Mapping
External Service RDS	ami-024a2614	Edit Mapping Delete Mapping
RHEL 6.x	ami-a16eb4db	Edit Mapping Delete Mapping
RHEL 7.x	ami-c998b6b2	Edit Mapping Delete Mapping
SUSE Linux Enterprise 12	ami-62bda218	Edit Mapping Delete Mapping
Ubuntu 14.04	ami-a22323d8	Edit Mapping Delete Mapping
Ubuntu 16.04	ami-66506c1c	Edit Mapping Delete Mapping
Windows Server 2008	ami-0509af5ff9695a433	Edit Mapping Delete Mapping
Windows Server 2008 with MSS...	ami-09ffe013945d1e6bf	Edit Mapping Delete Mapping
Windows Server 2012	ami-0fba87d7f8c8744d4	Edit Mapping Delete Mapping
Windows Server 2012 with MSS...	ami-0be991d31e9d393d5	Edit Mapping Delete Mapping
Windows Server 2016	ami-050202fb72f001b47	Edit Mapping Delete Mapping

The Image Settings section automatically lists all logical images available to you as a user, and, if set, each logical image's mapping to a physical cloud image ID in the region. From this section you can edit or delete an existing mapping, or add a mapping if one does not exist. If the region is a public cloud region, Workload Manager automatically adds the appropriate physical image ID for each of the OOB logical images based on data stored in the Cisco-hosted package store when you first configure the region in Workload Manager. You can later periodically update (sync) this automatic mapping with any new image IDs stored in the package store.

If a logical image is already mapped to a physical image in that region, the physical cloud image ID is shown in the Cloud Image ID column for that logical image. That mapping can be deleted by clicking the **Delete Mapping** link in the Action column. That mapping can also be edited by clicking the **Edit Mapping** link in the Action column. This causes the Edit Cloud Mapping dialog box as described in [Images Page](#) to be displayed.

If a logical image is not yet mapped to a physical image in that region, a mapping can be added by clicking the **Add Mapping** link in the Action column. This causes the Add Cloud Mapping dialog box as described in [Images Page](#) to be displayed.

If the region is a public cloud region, then a **Sync Image Mappings** link is displayed in the upper right of the Image Mapping section. For each public cloud region supported by CloudCenter Suite, a list of OOB logical image to physical cloud image mappings for that region is stored in the Cisco-hosted CloudCenter Suite package store (see [Architecture](#)). This information is periodically updated by Cisco when newer physical cloud images are uploaded by the cloud provider. Clicking the **Sync Image Mappings** link causes any updates to this list in the package store to be downloaded to Workload Manager and the image mappings for that region to be updated.

Other Admin Functions

Other Admin Functions

- [All Reports](#)
- [System Tags](#)
- [Usage Plans and Fees](#)
- [Tenant Management](#)

All Reports

All Reports

- [Reports Overview](#)
- [Usage Summary Report](#)
- [Application Deployments Report](#)
- [Running VM History Report](#)

Reports Overview

Reports Overview

- [Access Reports](#)
- [Available Reports](#)
- [Time Period Filter](#)
- [Search Reports](#)
- [Advanced Filters](#)
 - [Save Filters](#)
 - [Delete Saved Filters](#)

To access Workload Manager Reports, follow this procedure:

1. Click **Admin > All Reports** from the Workload Manager UI. The Reports section defaults to the [Usage Summary Report](#) page.
2. Click the dropdown arrow next to the report name to view the available reports.
3. Select the required report from this list.

Workload Manager provides the following reports via the Workload Manager UI:

- [Usage Summary Report](#)
- [Application Deployments Report](#)
- [Running VM History Report](#)



The billing task runs only every calendar hour. As a result, even if a job is deploys successfully, you may see the initial cost items for this job only *after* the hourly billing task has run at least once. It may be up to an hour before costs for your running job are displayed in the Workload Manager dashboard and all Usage Reports.

The time period filter is illustrated in the following screenshot.

Spend Summary eng@cliqr.com

TENANTS: **10** CLOUDS: **9** TOTAL COMPUTE COST: **\$3,248** RUN TIME (hrs): **57,601** CURRENT RUNNING VMs: **56**

By Tenant | By Cloud Region |

MTD YTD **30D** 60D 90D Range


TENANT NAME	RUN TIME (HRS)	CURRENT RUNNING VMs	COMPUTE COST
▼ CliQr Eng	11,689.76	4	\$736.47
US East (Virginia)	1,472.62	4	\$99.87
US West (Oregon)	2,848.97	0	\$109.6
US East (Virginia)	0.37	0	\$0.02
US West (California)	0.41	0	\$0.08
US-West-MCP 2.0 NA12	5,873.8	0	\$452.35
us-central1	995.73	0	\$69.58
US	497.86	0	\$4.97

The Time Period filter option is only available for some pages (for example, the [Usage Summary Report](#) or the [Virtual Machine Management](#) page). If available, the filter options are displayed in the top right corner. The available time period filter options are explained in the following table:

Time Period Filter	Description	Notes

MTD	Month to Date	The current month
YTD	Year to date	The current year
30D (Default)	30 Days	The current 30 days ending with today <div style="border: 1px solid green; padding: 5px; margin-top: 10px;"> ✔ The data that is displayed in response to a 30-Day time period request only displays data from the 1st of the month, not for the previous 30 days. To work around this issue, use the date Range option and provide the begin and end date for the required period. </div>
60D	60 Days	The current 60 days ending with today
90D	90 Days	The current 90 days ending with today
Range	A custom range specified by the selected month and year	If using APIs, this is the only available options to display reports for a period of time based on the <i>startDate</i> and <i>endDate</i> attributes

The All Tenants dropdown list is available for all Workload Manager reports and located in the top left corner, next to the report display options (My Tenant, or All Tenants or My Data depending on your Permissions). The following screenshot illustrates the All Tenants dropdown list.



This advanced filtering options helps you directly add short cuts to filtered lists that you can quickly access at a later time. This feature is available for some pages (for example, the [Running VM History Report](#) or the [Virtual Machine Management](#) page). The following screenshots display some of the available filters.

Compare

Cloud Region/Account

- All
- AMZN US West (Oregon) AMZN
- AMZN US West (Oregon) ina
- Amazon Asia Pacific Nort... Gaurav
- Amazon Asia Pacific Nort... ina

[17 MORE](#)

Filter

OS Type

- All
- Linux
- OS Unknown
- Windows

CPU

MIN: 1 MAX: 4

Memory GB

MIN: 1 MAX: 13

Storage GB

MIN: 2 MAX: 410

All Managed VMs

All Managed VMs

Demo Filter

Favorites only

Virtual Machines

- All
- CloudCenter Deployed
- Imported VMs

Status

- All
- Running
- Starting
- Stopping
- Stopped

[3 MORE](#)

Application Profiles

- All
- OTestApp
- 5Tier_WAP
- App_Batch_shruthi
- ArtifactoryServer

[26 MORE](#)

Cloud Family

- All
- Amazon
- AzurePack
- AzureRM
- Google

[5 MORE](#)

User

- Cliq Admin

Cloud Region

- All
- Amazon Asia Pacific Nort...
- Amazon EU West (Ireland)
- Amazon EU West (Ireland)
- Amazon US East (Virginia)

[10 MORE](#)

Cloud Account

- All
- AMZN (AMZN)
- Boobalan (AMZN)
- CliQr (AzureRM)
- CliQrQA (Amazon)


[12 MORE](#)

Application Profile

- All
- OTestApp
- 5Tier_WAP
- App_Batch_shruthi
- ArtifactoryServer

[27 MORE](#)


Deployment Environment

Users 

All

User 01 CloudCenter


Vik Pary

Groups 


All

G1


G2

Cloud Region 

Amazon US East (Vir...

Cloud Account 


Vik AWS cloud ...

Application Profile 


All

Jenkins

dummyExternalServi...

Deployment Environment 


AWS only

Status 


All

Canceled


Running/Deployed

Tags 

No Tags available

CCID 

No CCID available

Project 

No Project available

i You can additionally filter CloudCenter resources using the user-based **Groups** filter (see the highlighted image above).

User **Groups** displayed in the filter lists all user groups that are configured for your tenant. Selecting any user group list filters the list of application deployments for the selected group and keeps those deployments selected in this list if they map to users in the user group.

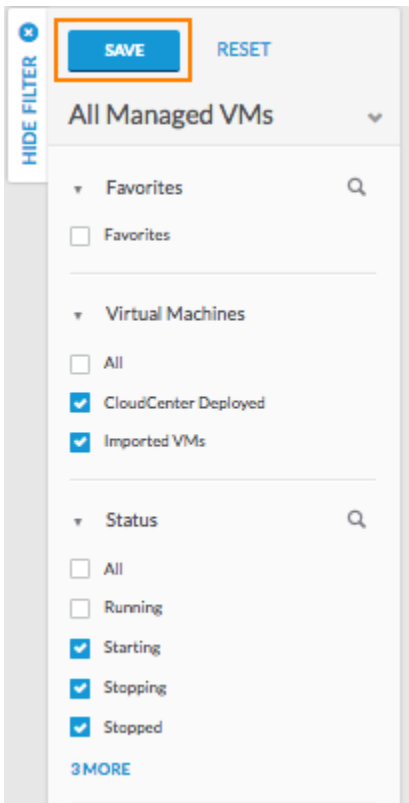
You can also combine the user and user group filters, and in this case, the report displays deployments that map to either the selected user or any user who is a member of the selected user group.

Save Filters

By saving a filter, you are directly adding short cuts to custom filtered lists that you can quickly access at a later time.

To save a custom filter, follow this procedure.

1. Select the required filters in the Filters pane and/or the Columns filter choices.
2. Click **Save**, located right above the Filters pane, as displayed in the following screenshot.



The Save Filter popup displays.

3. Enter a name for this filter and click **Save**.

Save Filters as

All fields are **required** unless otherwise specified

Name

Non-Running VMs

Selected Filters


Virtual Machines: CloudCenter Deployed, Imported VMs
Status: Starting, Stopping, Stopped, Error

Save
Cancel

4. The filter is saved and a status message displays in the page.


✓ Successfully saved selected filters with name Non-Running VMs.

5. You can access and view the saved filters from the dropdown list.



Delete Saved Filters

You can delete saved filters by clicking the Trash icon next to the saved filter live link.



The Delete Saved Filters popup confirms your intention before deleting the saved filter and displaying the status message at the [Application Deployments Report](#) page.




Advanced Filters

This advanced filtering options helps you directly add short cuts to filtered lists that you can quickly access at a later time. This feature is available for some pages (for example, the [Running VM History Report](#) or the [Virtual Machine Management](#) page). The following screenshots display some of the available filters.

The screenshot displays the 'All Managed VMs' page with a central filter sidebar and several side panels:


- Central Filter Sidebar:**
 - All Managed VMs** (selected)
 - Demo Filter**
 - Favorites** (only)
 - Virtual Machines**
 - All
 - CloudCenter Deployed
 - Imported VMs
 - Status**
 - All
 - Running
 - Starting
 - Stopping
 - Stopped
 - [3 MORE](#)
 - Application Profiles**
 - All
 - 0TestApp
 - 5Tier_WAP
 - App_Batch_shruthi
 - ArtifactoryServer
 - [26 MORE](#)
 - Cloud Family**
 - All
 - Amazon
 - AzurePack
 - AzureRM
 - Google
 - [5 MORE](#)
- Left Panel: Compare**
 - Cloud Region/Account**
 - All
 - AMZN US West (Oregon) AMZN
 - AMZN US West (Oregon) ina
 - Amazon Asia Pacific Nort... Gaurav
 - Amazon Asia Pacific Nort... ina
 - [17 MORE](#)
 - OS Type**
 - All
 - Linux
 - OS Unknown
 - Windows
 - CPU**
 - MIN: 1
 - MAX: 4
 - Memory GB**
 - MIN: 1
 - MAX: 13
 - Storage GB**
 - MIN: 2
 - MAX: 410
- Right Panel:**
 - User**
 - Cliqr Admin
 - Cloud Region**
 - All
 - Amazon Asia Pacific Nort...
 - Amazon EU West (Ireland)
 - Amazon EU West (Ireland)
 - Amazon US East (Virginia)
 - [10 MORE](#)
 - Cloud Account**
 - All
 - AMZN (AMZN)
 - Boobalan (AMZN)
 - CliQr (AzureRM)
 - CliQrQA (Amazon)
 - [12 MORE](#)
 - Application Profile**
 - All
 - 0TestApp
 - 5Tier_WAP
 - App_Batch_shruthi
 - ArtifactoryServer
 - [27 MORE](#)
 - Deployment Environment**

Users 

All

User 01 CloudCenter


Vik Pary

Groups 


All

G1


G2

Cloud Region 

Amazon US East (Vir...

Cloud Account 


Vik AWS cloud ...

Application Profile 


All

Jenkins

dummyExternalServi...

Deployment Environment 


AWS only

Status 


All

Canceled


Running/Deployed

Tags 

No Tags available

CCID 

No CCID available

Project 

No Project available

i You can additionally filter CloudCenter resources using the user-based **Groups** filter (see the highlighted image above).

User **Groups** displayed in the filter lists all user groups that are configured for your tenant. Selecting any user group list filters the list of application deployments for the selected group and keeps those deployments selected in this list if they map to users in the user group.


You can also combine the user and user group filters, and in this case, the report displays deployments that map to either the selected user or any user who is a member of the selected user group.

Save Filters

By saving a filter, you are directly adding short cuts to custom filtered lists that you can quickly access at a later time.

To save a custom filter, follow this procedure.

1. Select the required filters in the Filters pane and/or the Columns filter choices.
2. Click **Save**, located right above the Filters pane, as displayed in the following screenshot.



The Save Filter popup displays.

3. Enter a name for this filter and click **Save**.

Save Filters as

All fields are **required** unless otherwise specified

Name

Non-Running VMs

Selected Filters


Virtual Machines: CloudCenter Deployed, Imported VMs
Status: Starting, Stopping, Stopped, Error

Save
Cancel

4. The filter is saved and a status message displays in the page.


✓ Successfully saved selected filters with name Non-Running VMs.

5. You can access and view the saved filters from the dropdown list.



Delete Saved Filters

You can delete saved filters by clicking the Trash icon next to the saved filter live link.



The Delete Saved Filters popup confirms your intention before deleting the saved filter and displaying the status message at the [Application Deployments Report](#) page.



Time Period Filter

The time period filter is illustrated in the following screenshot.

The screenshot shows the 'Spend Summary' dashboard. At the top, there are summary statistics: TENANTS: 10, CLOUDS: 9, TOTAL COMPUTE COST: \$3,248, RUN TIME (hrs): 57,601, and CURRENT RUNNING VMs: 56. Below these are filter options: 'By Tenant' (selected), 'By Cloud Region', and a search icon. A time period filter dropdown is open, showing options: MTD, YTD, 30D (selected), 60D, 90D, and Range. Below the filter is a table with the following data:

TENANT NAME	RUN TIME (HRS)	CURRENT RUNNING VMs	COMPUTE COST
CliQr Eng	11,689.76	4	\$736.47
US East (Virginia)	1,472.62	4	\$99.87
US West (Oregon)	2,848.97	0	\$109.6
US East (Virginia)	0.37	0	\$0.02
US West (California)	0.41	0	\$0.08
US-West-MCP 2.0 NA12	5,873.8	0	\$452.35
us-central1	995.73	0	\$69.58
US	497.86	0	\$4.97

The Time Period filter option is only available for some pages (for example, the [Usage Summary Report](#) or the [Virtual Machine Management](#) page). If available, the filter options are displayed in the top right corner. The available time period filter options are explained in the following table:

Time Period Filter	Description	Notes
MTD	Month to Date	The current month
YTD	Year to date	The current year
30D (Default)	30 Days	The current 30 days ending with today <div style="border: 1px solid green; padding: 5px; margin-top: 10px;"> The data that is displayed in response to a 30-Day time period request only displays data from the 1st of the month, not for the previous 30 days. To work around this issue, use the date Range option and provide the begin and end date for the required period. </div>
60D	60 Days	The current 60 days ending with today
90D	90 Days	The current 90 days ending with today
Range	A custom range specified by the selected month and year	If using APIs, this is the only available options to display reports for a period of time based on the <i>startDate</i> and <i>endDate</i> attributes

Usage Summary Report

Usage Summary Report

- [Overview](#)
- [Report Details](#)
- [Additional Options](#)

To view the Usage Summary Report, click **Admin > All Reports**. The **Usage Summary Report** displays by default.



The billing task runs only every calendar hour. As a result, even if a job is deployed successfully, you may see the initial cost items for this job only *after* the hourly billing task has run at least once. It may be up to an hour before costs for your running job are displayed in The Dashboard and [All Reports](#).

The top section of the Usage Summary report displays usage summary information for your tenant hierarchy (includes the sub-tenant details):

- Tenants (see [Manage Tenants](#))
- Clouds (see configured [Clouds](#))
- Total Compute Cost (see [Track Cloud Costs](#))
- Run Time (the operational VM hours)
- Current Running VMs (includes NodeStarting, NodeStarted, NodeReady, NodeReachable, NodeResumed, NodeRebooted, and NodeError states. See [Deployment, VM, and Container States](#) for additional context)

This report provides aggregate numbers by Tenant or by Cloud Region.



The aggregated information is only displayed for your direct sub-tenant. It does not include the numbers for your sub-tenant's sub-tenant(s).

Each Tenant or Cloud Region is aggregated as a group as well as individually when you click the dropdown arrow for each:

- **Tenant:** By cloud region for every sub-tenant. The following screenshot displays the aggregated summary **By Tenant**.

NAME	RUN TIME (hrs)	CURRENT RUNNING VMs	CLOUD COST
[-] Cisco Systems	354,742.38	352	\$13,294.69
[-] AMZN EU West (Ireland)	7,816.14	0	\$258.53
[-] AMZN US West (Oregon)	2,937.59	1	\$76.27
[-] Amazon Asia Pacific North East (Tokyo)	8,688.88	6	\$358.44
[-] Amazon EU West (Ireland)	2,611.1	4	\$110.86
[-] Amazon US East (Virginia)	150,757.72	126	\$4,098.56
[-] AzureRM US East (Virginia)	35,489.77	73	\$1,252.87
[-] Didata US-West-MCP 2.0 NA12	1,716.93	0	\$182.2
[-] Google us-central1	1,346.89	0	\$94.73

- **Cloud Region:** By sub-tenant for every cloud region. To view the cloud region breakdown within this CloudCenter instance, click **By Cloud Region**, as shown in the following screenshot.

NAME	RUN TIME (hrs)	CURRENT RUNNING VMs	CLOUD COST
[-] AMZN EU West (Ireland)	7,881.08	0	\$260.35
[-] Cisco Systems	7,816.14	0	\$258.53
[-] ina	64.93	0	\$1.82
[-] AMZN US West (Oregon)	17,026.52	18	\$415.57
[-] Cisco Systems	2,937.59	1	\$76.27


You can filter this report using additional date options and search terms as well as download this report. See [Reports Overview](#) for additional details.

Application Deployments Report

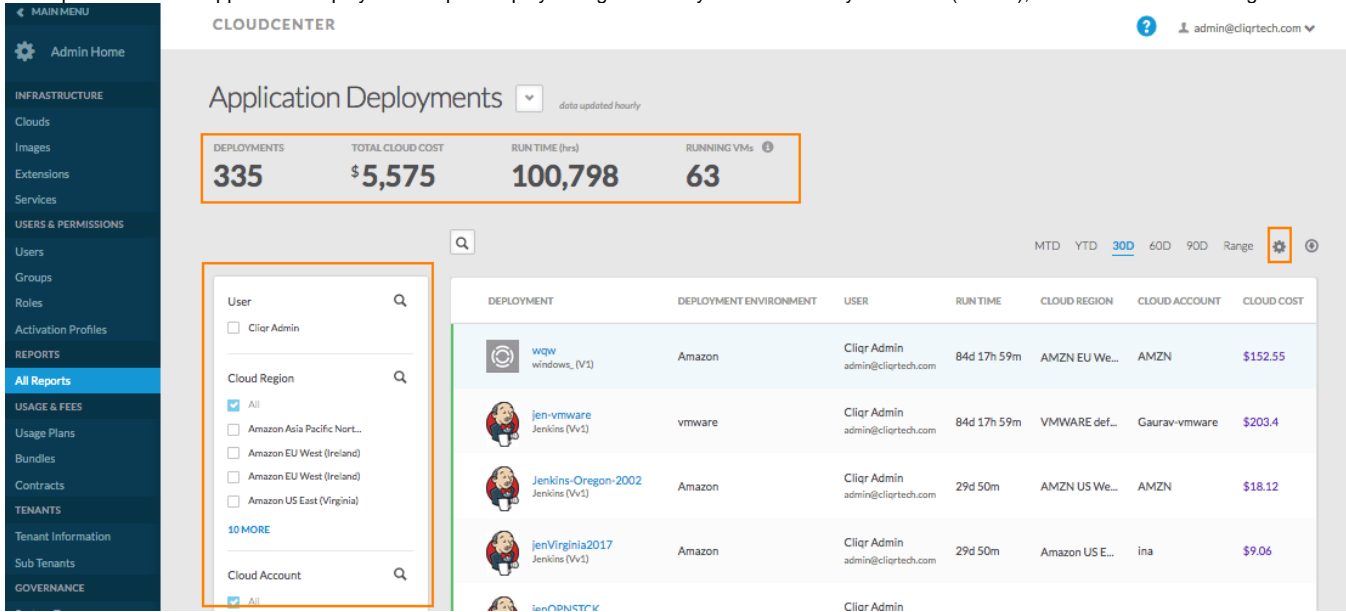
Application Deployments Report

- [Overview](#)
- [Report Details](#)
- [Other Notes](#)
- [Additional Options](#)

To view the Application Deployments Report, click **Admin > All Reports** and select Application Deployments Report from the dropdown arrow list next to the report name.

 The billing task runs only every calendar hour. As a result, even if a job is deployed successfully, you may see the initial cost items for this job only *after* the hourly billing task has run at least once. It may be up to an hour before costs for your running job are displayed in the Dashboard and all [Usage Reports](#).

The top section of the Application Deployments report displays usage summary information for your tenant (default), as shown in the following screenshot.

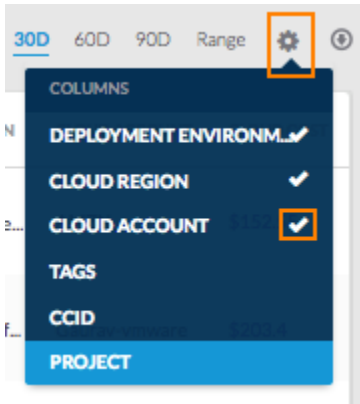


- [Deployment Details](#)
- Total Compute Cost (see [Track Cloud Costs](#))
- Run Time (hrs): The operational VM hours
- Current Running VMs (see [Financial Overview](#))

The Run Time column in the Application Deployments Report includes the usage hours for the [External Service](#).

You can filter this report using additional date options and search terms as well as download this report. See [Reports Overview](#) for additional details.

Click the *cog* icon to display additional columns or fewer columns based on the selected (check) resources. The following screenshot shows the cog icon.



Running VM History Report

Running VM History Report

- [Overview](#)
- [Report Details](#)
- [Comparison Options](#)
- [Filtering Options](#)

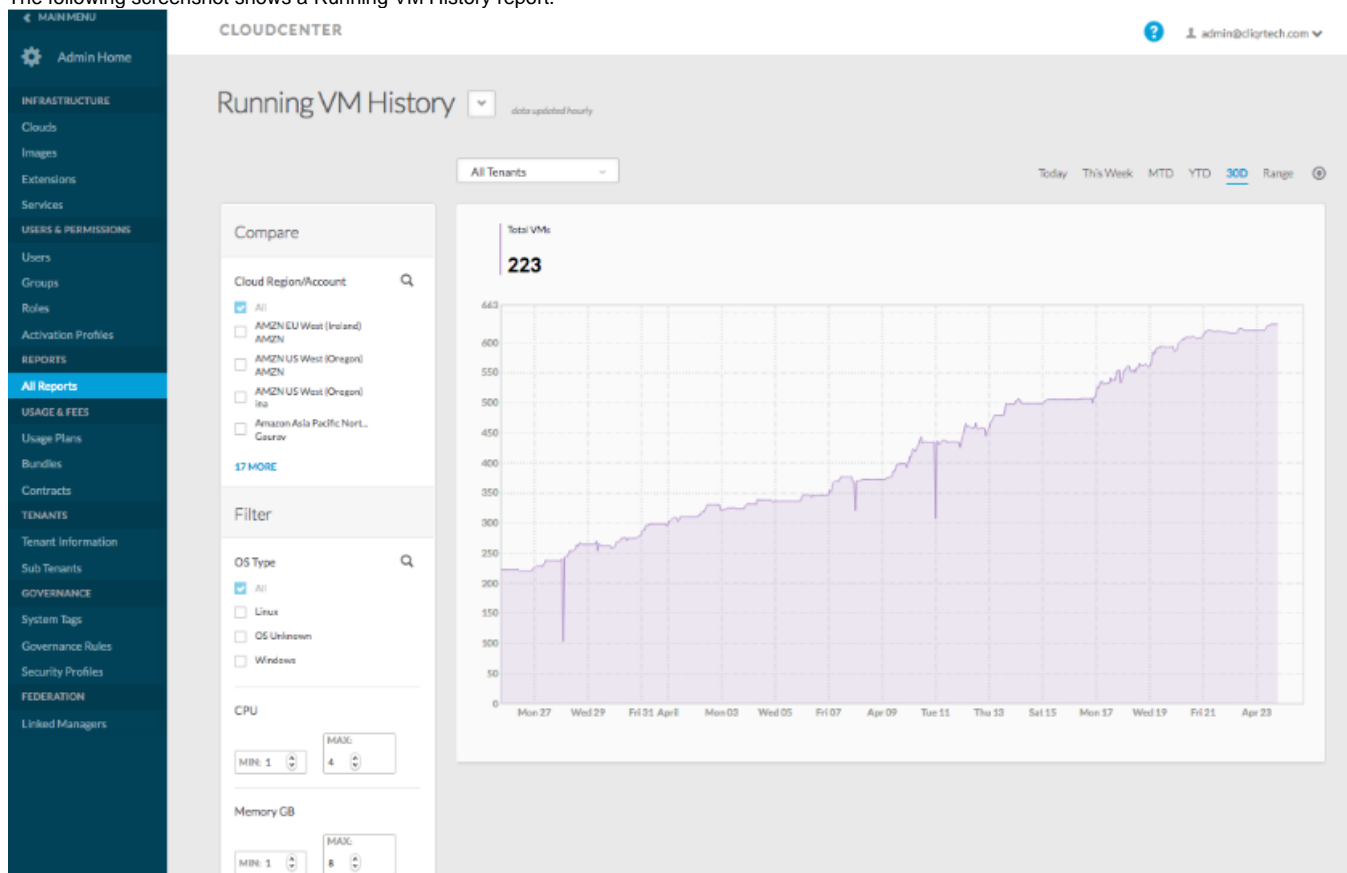
To view the Running VM History, click **Admin > All Reports** and select **Running VM History** from the arrow list next to the report name.

The Running VM History report provides a snapshot and usage pattern for VM usage for your tenant hierarchy (includes the sub-tenant details) on a per cloud account basis and provides an aggregate number of VMs being run:

- By users and sub-tenants
- Within your tenant hierarchy
- Over a period of time
- For a given cloud account

This report also allows you to view the history based on all Tenants (provides the report for the entire tenant hierarchy rooted at the current tenant level).

The following screenshot shows a Running VM History report.



- Tenants (see [sub-tenants](#) when you view your own tenant details, users and groups are also available as filtering options.)
- Users (the configured [users](#) are only displayed if you are viewing your own tenant details)
- Groups (the configured [groups](#) are only displayed if you are viewing your own tenant details)
- Cloud Regions (see configured [Configure a Region](#))

You can filter the VMs being used across cloud accounts based on users, groups, OS type, and VMs (CPU, size, memory, and storage).

This data is based on an hourly task that runs on Workload Manager and collects the snapshot of VMs running at that point of time. As the time span increases, the data increases exponentially. A purging task runs once a day to purge the data based on the following criteria:

- Hourly data for the last 30 days
- Max VMs in a day for the last 30 days – 6 months
- Max VMs in a month for anything beyond the last 6 months

This report also allows you to compare the VM usage against the limits imposed by the VM Subscription plan, if any.

You can filter this report using the OS Type, CPU, Memory, and Storage options and search terms as well as download this report. See [Reports Overview](#) for additional details.

For the same time interval setting, you can view the report for different filter options. If you choose a different time interval setting, all filters are reset.

System Tags

System Tags

- [About System Tags](#)
- [Managing System Tags](#)
- [Adding a System Tag](#)
- [Share with Sub-tenants](#)
- [Isolation Tags](#)
- [Applying Tags as Annotations](#)

A tag is a tenant-owned resource that is automatically shared with all users within the tenant. It can also be shared with read access to all users in all sub-tenants within this tenant.

The System Tags list page (**Admin > Reports > System Tags**) sorts the list by listing the most recently created first (by default), but they can also be sorted by name.

The Systems Tags page lists tags even when a tag is shared from a parent tenant as shown in the following screenshot.



Tags created by a parent tenant and shared with a sub tenant retain their name. By default, all tags are sorted with most recently updated tags at the bottom of the list.

To manage system tags, click **Admin > System Tags** in the Workload Manager UI to display the System Tags page.

The System Tags page lists lets you perform the tasks that the following table describes.

Task	Description
Add a new system tag.	Click the Add System Tag link.
View or update an existing system tag.	Click the Edit link in the Actions column for the system tag. The Edit System Tag page for the system tag displays.
Delete a system tag.	Click the Delete link in the Actions column for the system tag.

When you add a system tag, you create a new tag based on configuration settings that you make. To add a system tag, follow these steps:

1. On the System Tags page, click the **Add System Tag** link.
The Add System Tag page displays.
2. In the **Name** field, enter a brief and unique descriptive name for the system tag.
The name can include letters, numbers, and underscores (_).
3. (Optional) In the **Description** field, enter a brief description of the system tag.
4. Click the **Save** button.

A tag is a tenant-owned resource that can be shared with users/groups within the tenant as well as user/groups in other sub-tenants in the hierarchy:

- If one user/group in a tenant has manage access to a tag, then all users/groups in the tenant will also have manage access to this tags.
- If you share a tag with any user/group in a sub-tenant, that user/group will have read-only access for tags in the tenant.

The System Tags have a **Share with Sub-tenants** toggle switch column:

- **On:** Users in tenants that are further down the hierarchy can only *View* (read only) these tags.
- **Off:** Default. No user in any sub-tenant can view or use a tag in this state. The tenant admin can grant permissions to sub-tenants by enabling the **(ON)** toggle switch for the required tag.

When a deployment environment or application profile is shared with a sub-tenant or users in a sub-tenant, all policies (and system tags) associated with that environment or app profile can still be used in the shared environment or application profile.

Isolation tags are different from system tags. Isolation tags are based on the string provided by a user when launching a job. See [Security and Firewall Rules > Isolation Tags](#) and the [Submit Job API](#) pages for additional context.

You can apply tags as annotations to associate custom information with resources in the CloudCenter platform. After you add a tag, you can then use that tag to filter VMs or containers. The following table identifies the annotation terminology used for each cloud.

Cloud	Annotation Terminology
AWS	Tag
vCenter	<ul style="list-style-type: none"> • Annotation attributes • Tag
OpenStack	Metadata
Azure RM	Tag
Alibaba	Tag
Kubernetes	Labels

You can apply tags in one of two ways from the UI:

- From the application profile, on the Basic Information tab:
 - Specify the tags for in the application profile so you can use the tags as filters to search for this application. See [New Application Profile](#) to specify the tag in an application profile.
 - Associate metadata relevant to your environment in the application profile. See [Understand Application Tier Properties](#) for additional context.
 - You can also override the name-value pair in the Metadata at the time of deployment. See [Topology Modeler](#) for explicit details on the Basic Information tab.
- Predefined System tags as specified earlier in this page.
 - You can update the tag for each deployment in the **Associate Tags** field in the Deployment Details page.
 - To do this, the tag should already be added as specified in the [Adding a System Tag](#) section earlier in this page. See [Deployment Details Page](#) for additional context on updating a tag for a specific deployment.

Usage Plans and Fees

Usage Plans and Fees

- [Financial Overview](#)
- [Plan Configuration](#)
- [Bundle Configuration](#)
- [Workload Manager Cost and Fees](#)

Financial Overview

Financial Overview

- [Terminology](#)
- [Bundle and Plan Management](#)
- [Bundle and Plan Creation](#)
- [Assigning Plans to Subtenants and Users](#)
- [Plan Types](#)
- [Common Options for All Plan Types](#)
- [Meter Applications](#)

Term	Description
<i>Bundle</i>	A <i>bundle</i> is a plan type that gives the user a fixed dollar amount of usage or a fixed VM-hour limit.
<i>Concurrent VMs</i>	The number of simultaneously running VMs permitted to each user under a monthly usage plan
<i>Discount rate</i>	Contracts sometimes have a discount rate that provides users an incentive to sign a longer contract. Typically, a longer contract length has a deeper discount rate. This discount is applied to a plan's base one-time fee, annual fee, bundle price, and monthly subscription fee (Cisco does not discount overage and storage fees).
<i>Minimum charge</i>	The minimum number of minutes deducted from a plan when a user runs an application.
<i>Monthly VM hours</i>	The number of VM hours credited to the user each month.
<i>Maximum Running VMs</i>	The number of VMs permitted to the user each month.
<i>Overage limit</i>	If users exceed the limit stipulated in the subscription, Admins can decide if they want to charge a fee when users exceed the limit: <ul style="list-style-type: none"> • Limited: Stops deployment after reaching limit (specified fee is charged) • Unlimited (default): Continues deployment after overage rate kicks in (no additional charge)
<i>Plan</i>	A <i>plan</i> is an agreement determined by the admin and assigned to a user to determine the capacity or allowed usage for that user.
<i>Runtime</i>	<i>Runtime</i> refers to the number of hours: <ul style="list-style-type: none"> • VM: The number of hours that the VM has been running as part of a deployment • External Service: The number of hours for which the service has been active as part of a deployment

The root tenant automatically has a default unlimited subscription plan. All users in a tenant automatically inherit the plan assigned to the tenant.

A tenant admin can create more restrictive plans and assign them to users in the tenant or to a subtenant.

If the tenant admin does not assign a plan to a newly created subtenant, then all users in the subtenant will be unable to deploy applications, and the subtenant admin will be unable to create any bundles and plans.

Admins have the following flexibility when managing bundles and plans:

- Restrict certain bundles, and plans to just admins so it is not available to all users.
- Edit or delete bundles, and plans if no users have signed up. If users have signed up, the admin must create a *new* bundle or plan and revise the required settings.
- Delete a bundle if that bundle is not assigned to any user.

See [Bundle Configuration](#) and [Plan Configuration](#).

After creating bundles and plans they must be assigned to subtenants or users in order to be effective.

To assign a plan (and any associated bundle) to a subtenant follow these steps:

1. From the Suite Admin UI > Tenants, find the subtenant in the subtenant list and click the dropdown icon in the actions column.
2. From the dropdown menu, select Manage Plan or Manage Bundle, as appropriate. The corresponding dialog box appears.
 - a. For the Manage Plan dialog box, select the appropriate **Usage Plan**, **Bundle**, and **Plan Adjustment**. Then click **Done**.
 - b. For the Manage Bundle dialog box, select the appropriate **Bundle**. Then click **Done**.

To assign a plan (and any associated bundle) to a user follow these steps:

1. From the Suite Admin UI > Users, find the user in the users list and click the dropdown icon in the actions column.
2. From the dropdown menu, select Manage Plan or Manage Bundle, as appropriate. The corresponding dialog box appears.
 - a. For the Manage Plan dialog box, select the appropriate **Usage Plan, Bundle, and Plan Adjustment**. Then click **Done**.
 - b. For the Manage Bundle dialog box, select the appropriate **Bundle**. Then click **Done**.

When creating a plan in Workload Manager, the administrator must select one of five possible plan types that are summarized in the following table.

Category	Plan Types	
Subscription	VM Hour Subscription (Default) <ul style="list-style-type: none"> • Specifies a limit for the number of VM hours per month • Enable Rollover: If checked, any unused VM hour balance will rollover to the next month. • Overage Limit and Overage Rate 	
	VM Subscription <ul style="list-style-type: none"> • Allows a specified number of concurrent VMs across all supported clouds • Overage Limit and Overage Rate • The VM count begins as soon as the VM is deployed 	
	Unlimited Subscription <ul style="list-style-type: none"> • Does not have any limitations on usage • Users in this plan are not limited to VM hours, number of VMs, or a budget limit 	
Bundle	VM-Hour Bundle <ul style="list-style-type: none"> • Specifies a limit for the the total number of VM-hours used. You can start new deployments until this VM-hour amount is consumed. • Usage Increment Units: If checked, units above the limit are rounded up to the next level (for example, 11 minutes of actual usage with a 10-minute increment unit will be metered as 20 minutes). • Overage Limit and Overage Rate 	
	Budget Bundle <ul style="list-style-type: none"> • Specifies a limit for the total currency spent. You can start new deployments until this currency amount is consumed. • Overage Limit and Overage Rate 	

The parameters and check box selections that the following table describes are available to all plan types.

Parameter/Check Box	Description
Usage Details	
Only Visible to Tenant Admin	<ul style="list-style-type: none"> • If checked, this plan is only visible to the tenant admin. • Default = unchecked.

Integrated metering, reporting, and billing ensure that users are always aware of their consumption and are charged appropriately for what they use.

- Workload Manager monitors and meters deployed applications through the [Management Agent \(Worker\)](#). Monitoring and metering data is used for reporting purposes.
- Workload Manager meters end user activity and enforces custom plans.

Plan Configuration

Plan Configuration

- [Overview](#)
- [Create a Plan](#)
- [Discontinue Plans](#)
- [Enable Plans](#)
- [Edit or Delete Plans](#)
- [Manage Plan](#)
- [View Users or Projects Assigned to a Plan](#)

Configuring a usage plan is a multi-step process:

1. Create the plan (this page).
2. Assign the plan to a tenant or user.

See [Financial Overview](#) for definitions, types of subscription plans, definition.

Workload Manager's plans and bundles provide enterprises with two benefits:

- Enables admins to limit and restrict sub-tenants and users.
- Allows enterprises or organizations to charge users in their tenants or sub-tenants for associated cloud costs for management or cloud overhead costs. These costs are generally in addition to actual cloud costs.

To create a plan, follow this procedure:

1. Login to the Workload Manager UI and click **Admin > Plans**.
2. Click the **Create Usage Plan** link to add a new plan. The Create Usage Plan popup displays.
3. Enter the information pertaining to your enterprise based on the selected plan or bundle. See [Financial Overview > Workload Manager Subscription Types](#) for additional context.
4. When you click **Save**, the newly added plan is displayed in the Usage Plans page:

Once created, you can edit, delete, and discontinue.

If a plan is already assigned to a user, you can only **Discontinue** that plan. You cannot Edit or Delete an assigned plan.

To view or discontinue a plan, follow this procedure.

1. Login to the Workload Manager UI and click **Admin > Plans**.
2. Click **Discontinue** to stop an assigned plan.
3. Click **OK** in the popup to proceed with this action. The updated status is displayed at the top of the Usage Plans page. Once discontinued, you can Enable plans.

If a plan is discontinued, you must **Enable** this plan to make it active again.

To enable a plan, follow this procedure.

1. Login to the Workload Manager UI and click **Admin > Plans**.
2. Click **Enable** to restart the plan.
3. Click **OK** in the popup to proceed with this action. The updated status is displayed at the top of the Usage Plans page. Once enabled, you can Discontinue plans.

You can **Edit** or **Delete** any plan listed in the Plans page if the plan is not assigned to any user. However, you cannot change the following values after you assign the plan to a user:

- Plan Type
- Enable Rollover

See [Financial Overview](#) for more details on these values

To edit or delete a plan, follow this procedure:

1. Login to the Workload Manager UI and click **Admin > Plans**.
2. Click **Edit** to modify an unassigned plan or **Delete** to remove an unassigned plan.
3. Click **OK** in the confirmation popup to proceed with this action. The updated status is displayed at the top of the Usage Plans page and the plan is removed from the database.

See [Usage Plans and Fees](#).

You can view the users or projects assigned to a plan. To view the users or projects assigned to a plan, follow this procedure:

1. Login to the Workload Manager UI and click **Admin > Plans**.
2. In the Usage Plans page, click the non-zero *number* link displayed in one of the columns that the following screenshot shows.

Usage Plans + Create Usage Plan

Search plans Show 30 per page Page 1 of 1

Plan Name	Plan Type	Users Assigned	Projects Assigned	Actions
100 Dollars	Prepaid Budget Bundle	2	0	Enable
100 Hour	Prepaid VM-hour Bundles	7	0	Discontinue
1000 VM Plan	VM Subscription	1	0	Discontinue
10000 Dollars	Prepaid Budget Bundle	0	3	Discontinue
100hourpermonth	VM-hour Subscription	0	0	Edit Delete
85-hour VM Sub	VM-hour Subscription	0	0	Edit Delete
Conception VM Plan	VM Subscription	0	0	Edit Delete
Deep Plan	VM-hour Subscription	1	0	Discontinue
deusageplan2	VM-hour Subscription	1	0	Discontinue
Doc VM-hour Plan	VM-hour Subscription	0	0	Edit Delete
Free Plan	Prepaid VM-hour Bundles	3	0	Discontinue
Large Project	Prepaid Budget Bundle	0	1	Discontinue
Project 01 - Phase Dev - Budget	Prepaid Budget Bundle	0	1	Discontinue
Project 01 - Phase Dev - Budge...	Prepaid Budget Bundle	0	0	Edit Delete
Testing 123	Prepaid Budget Bundle	0	0	Edit Delete
TestingVmproject	VM Subscription	0	1	Discontinue
Unlimited Plan	Unlimited Subscription	9	0	Discontinue
Unlimited Plan 01	Unlimited Subscription	1	0	Discontinue
VF PO 57258	Prepaid Budget Bundle	0	0	Edit Delete
VMHourtest_api_18	VM-hour Subscription	5	0	Discontinue

- a. **Users Assigned:** To view the users assigned to the corresponding plan
 - b. **Projects Assigned:** To view the projects assigned to the corresponding plan
3. A popup displays the following details for each resource:
- a. Users for plan_name: Email, Name, Company, Start Date, Status, or Payment Status.

Users for VMHourtest_api_18 Close

Search users Show 20 per page Page 1 of 1

Email	Name	Company	Start Date	Status	Payment Status
Luserapi118_F...	Fuserapi118 Lu...	cliqr	05/10/2015	Enabled	
Luserapi218_F...	Fuserapi218 Lu...	cliqr	05/10/2015	Enabled	
Luserapi118_F...	Fuserapi118 Lu...	cliqr	05/10/2015	Enabled	
Luserapi218_F...	Fuserapi218 Lu...	cliqr	05/10/2015	Enabled	
coadmin1@cliq...	coadmin1 coa1	cliqr	02/11/2015	Enabled	

- b. Projects for plan_name: Name (of the project).

Projects for 10000 Dollars Close

Search projects Show 20 per page Page 1 of 1

Name
SalesIT
ATT_Test
testQAProject

4. Click **Close** when you have the required user/project information for the plan.

Bundle Configuration

Bundle Configuration

- [Create a Bundle](#)
- [Discontinue Bundles](#)
- [Enable Bundles](#)
- [Edit or Delete Bundles](#)

Workload Manager's plans and bundles provide enterprises with two benefits:

- Enables admins to limit and restrict sub-tenants and users.
- Allows enterprises or organizations to charge users in their tenants or sub-tenants for associated cloud costs for management or cloud overhead costs. These costs are generally in addition to actual cloud costs.

To create a bundle, follow this procedure:

1. Login to the Workload Manager UI and click **Admin > Bundles**.

Bundle Name	Bundle Type	Users Assigned	Actions
100 Dollars	Budget Based	2	Discontinue
100 Hours	VM-hour Based	10	Discontinue
15-Dollar Bundle	Budget Based	0	Edit Delete
Free 1000 Hours	VM-hour Based	2	Discontinue
VF PO 57258	Budget Based	0	Edit Delete

2. Click the **Create Bundle** link to add a new bundle. The Create Bundle popup displays.
3. Enter the information pertaining to your enterprise based on the bundle. See [Financial Overview > CloudCenter Subscription Types > Bundle](#) for additional context.
 - VM-hour Based
 - Budget Based
4. When you click **Save**, the newly added bundle is displayed in the Bundles page: Once created, you can Edit, Delete, or Discontinue bundles.

If a bundle is already assigned to a user, you can only **Discontinue** that bundle. You cannot Edit or Delete an assigned bundle.

To discontinue a bundle, follow this procedure.

1. Login to the Workload Manager UI and click **Admin > Bundles**.
2. Click **Discontinue** to stop an assigned bundle.
3. Click **OK** in the popup to proceed with this action. The updated status is displayed at the top of the Bundles page. Once discontinued, you can **Enable** bundles.

If a bundle is discontinued, you must **Enable** this bundle to make it active again.

To enable a bundle, follow this procedure.

1. Login to the Workload Manager UI and click **Admin > Bundles**.
2. Click **Enable** to restart the bundle.
3. Click **OK** in the popup to proceed with this action. The updated status is displayed at the top of the Bundles page. Once enabled, you can Discontinue bundles.

You can **Edit** or **Delete** any bundle listed in the Bundles page if the bundle is not assigned to any user.

To edit or delete a bundle, follow this procedure:

1. Login to the Workload Manager UI and click **Admin > Bundles**.
2. Click **Edit** to modify a bundle or **Delete** to remove a bundle.
3. Click **OK** in the confirmation popup to proceed with this action. The updated status is displayed at the top of the Bundles page and the bundle is removed from the database.

Workload Manager Cost and Fees

Workload Manager Cost and Fees

- [Overview](#)
- [What are the CloudCenter Costs?](#)

This section provides details on the costs and fees references within the Workload Manager.

When you use the Workload Manager, you may be charged for the structural costs and fees that the following table describes.

Cost	Description	Configuring this Cost
Cloud Cost	Cloud cost refers to the cloud infrastructure costs charged by the cloud providers (Public Clouds or Datacenters and Private Clouds). This cost can be broken down per deployment, per run, per user, per VM, per instance type, and so forth.	See Track Cloud Costs .
Management Cost	Management cost refers to the cost associated with the plan, bundle, or contract used for each CloudCenter user. The Financial Overview section provides the breakdown of each plan or bundle and the cost associated with related management charges.	See Configure Plan Configuration or Bundle Configuration .
Image Cost	CloudCenter billing defaults to the base cost of the image. If you launch an application using a custom image, the CloudCenter platform does not override the cost for this image. Enterprise administrators must explicitly override the instance price in the image cost so they account for accurate billing requirements.	Custom Service Definition
Service Fee	When you use one of the OOB Services , you can add a cost for the service and provide the capability to charge separately for this service. This allows admins the flexibility to charge individual users or sub-tenants for any service that is built and added to the Marketplace or the Topology Modeler Services tab. The VM-Hours column in the Application Deployments Report includes the usage hours for the External Service .	Custom Service Definition

Tenant Management

Tenant Management

- [Overview](#)
- [Manage Tenant-wide Firewall Rules](#)
- [Manage Default Usage Plan](#)
- [Manage Clouds](#)

CloudCenter Suite supports a multi-tenant model where each tenant has their own users, resources, permissions and policies. All configuration of tenants and subtenants is performed through the Suite Admin Tenant List page as described in [Manage Tenants](#). This section discusses actions in the Suite Admin Tenants page that are specific to Workload Manager

The Suite Admin allows tenant administrators to add, update and delete firewall rules that are applied to all VMs launched by all users in the tenant. Each firewall rule may be applied to one, several, or all VM-based cloud types configured for the tenant. To manage firewall rules, go to the Suite Admin home page and click on the Tenants tab to get to the Tenant List page. For your own tenant, hover over the Actions column and click on the dropdown icon to reveal the dropdown menu as shown in the figure below.

The screenshot shows the 'My Tenant' section of the Suite Admin interface. It features a table with the following columns: TENANT NAME, TENANT ID FOR LOGIN, LAST UPDATED, and ACTIONS. A single row is visible with the tenant name 'My Tenant 1', ID 'cisco', and '12 days ago' for the last updated date. The ACTIONS column contains a blue 'Edit' button. A dropdown menu is open below the 'Edit' button, showing two options: 'WORKLOAD MANAGEMENT' and 'Manage Security Policies'.

TENANT NAME	TENANT ID FOR LOGIN	LAST UPDATED	ACTIONS
My Tenant 1	cisco	12 days ago	Edit WORKLOAD MANAGEMENT Manage Security Policies

Select the Manage Security Policies menu choice. This displays the Manage Firewall Settings dialog box as shown in the figure below.

The screenshot shows the 'Manage Firewall setting for Tenant Cisco' dialog box. It has a title bar with a close button (X). Below the title bar is a table with columns: PROTOCOL, FROM PORT, TO PORT, SOURCE CIDR, and CLOUDS. Below the table is a blue '+ ADD RULE' button. Underneath, there are two toggle switches: 'CREATE DEFAULT SECURITY GROUPS FOR USERS IN THIS TENANT' and 'ALLOW LAUNCHED VMS TO COMMUNICATE WITH EACH OTHER', both currently set to 'NO'. A blue 'DONE' button is located at the bottom right of the dialog box.

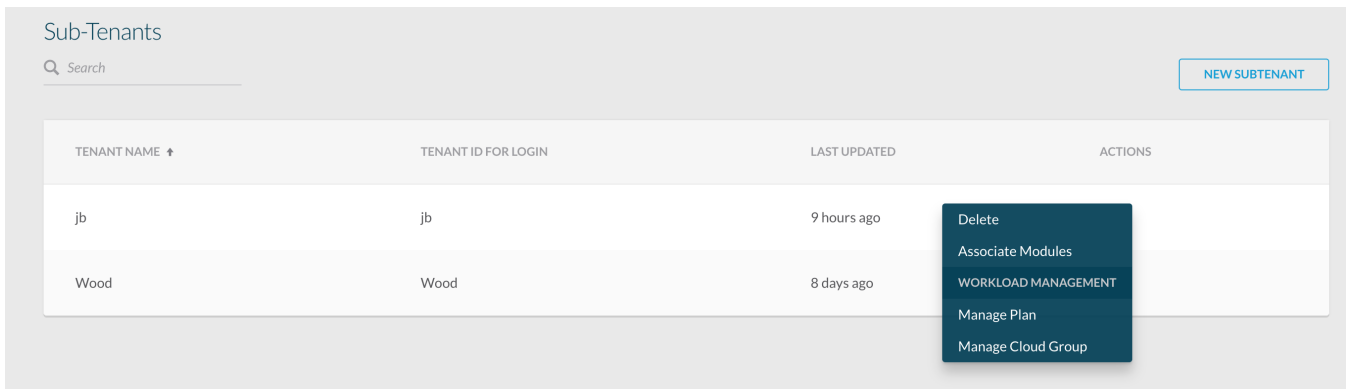
The top section of the dialog box lets you add a new firewall rule by clicking the **Add Rule** button. When you click the Add Rule button, a new line is created in the list of rules as shown in the following figure. You must then complete all of the fields for that rule. You can create more rules by again clicking the Add Rule button.

Important: In order for your newly created firewall rules to be saved when you click the Done button, you must ensure that the **Create default security groups for users in tenant** toggle is turned on before you click Done. If not, when you click Done, your newly created firewall rules will be lost.

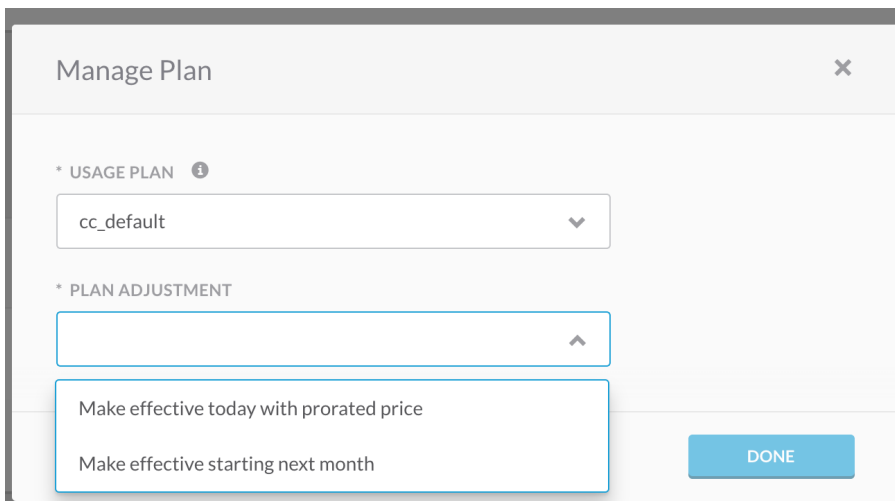
Note that additional firewall rules may be defined for each tier of an application in the [Application Tier Properties](#) section of the Topology Modeler tab of the Application Profiles form. And additional firewall rules can be defined in [Security Profiles](#) which may be selected by the user when [Deploying an Application](#). See [Security and Firewall Rules](#) for additional context.

The **Allow launched VMs to communicate with each other** toggle allows all VMs launched by a user to communicate with each other on all ports across all deployments for that user. This feature is only supported on Amazon, OpenStack, and Google clouds. Unchecking this check box puts the onus on users to set up inter-node communication for their respective deployments.

The tenant administrator can set or update the default usage plan for a subtenant. To do this, click the action dropdown icon for that subtenant as shown in the figure below, then select **Manage Plan**.



This displays the Manage Plan dialog box as shown in the figure below.



Select a usage plan from the first dropdown menu, then select a plan adjustment from the second dropdown menu. Click Done to save.

You can share any cloud region or cloud account visible to your tenant with any of your subtenants. From the subtenant actions dropdown menu select Manage Cloud Groups. This displays the Manage Cloud dialog box as shown in the figure below.

Manage Clouds for subtenant1

AWS1
1 Region(s) & 0 Account(s) assigned

AWS2
No Regions or Accounts assigned

OpenStack
No Regions or Accounts assigned

Regions Accounts

Regions (1)

- All Regions (3)
- US West (Oregon)
- Canada (Central)
- EU (Ireland)

ALLOW SUB-TENANT TO ADD CLOUDS NO

DONE

For each cloud tab on the left, select the **Accounts** tab on top to choose the cloud accounts to share, then select the **Regions** tab on top to choose the cloud regions to share. You also have the option to **Allow Sub-Tenants to Add Clouds** by toggling the corresponding switch for this option. Click **Done** after repeating this for all clouds you want to share.

Services

Services

- [Service Administration](#)
- [Custom Service Definition](#)
- [External Service](#)
- [Container Service](#)
- [Service Lifecycle Actions](#)
- [Service Versions and Usage](#)
- [Service Export/Import](#)

Service Administration

Service Administration

- [Overview](#)
- [The Services Framework](#)
- [Service Clustering](#)
- [Handling Information between Tiers](#)
- [Guidelines to Configure Service Scripts](#)
- [Export Application Scripts](#)

A *service* is a mid-tier building block for an application. The concept of services is natural to the Workload Manager. The concept of services like Tomcat, NginX, MongoDB, and other [Supported OOB Services](#) is exposed via the [Topology Modeler](#). This section discusses the services framework that is available at the administrative level.

The Services Framework enables enterprises to create and add their own services to the Service Catalog and make them available in the Topology Modeler. Once the enterprise defines the service, the same definition can be used across applications.

Workload Manager allows you deploy multiple instances of a service within a tier when that service is stateless, such as a base OS service or a web service.

The `minClusterSize` and `maxClusterSize` are system-defined parameters for services (see [Pre-Defined Parameters](#) for additional context). The *Web Server* or *OS service* groups already inherit these parameters and you do not need to explicitly define the parameters.

Using Workload Manager, you can add or reduce the number of nodes within a clustered service.

For a multi tier deployment, VMs for all tiers are launched in parallel. During the node initialization phase, every node is passed the IP address of all other nodes in the topology as environment variables by Workload Manager. For example, when Workload Manager passes the IP address of MySQL service to Apache, the scripts on the Apache service can access this IP address using a convention `%(TierName)_TIER_IP%`.

To share information (that may not be known prior to a deployment) between tiers, specify the information in a static parameter. See [Parameter Substitution](#) and [Using Parameters](#) for additional context.

Scripts is the most common method used to configure a lifecycle action is to use scripts. Ensure to adhere to these guidelines when configuring service scripts:

- You can configure scripts in any language or format.
- Add all agent lifecycle action scripts to the .zip file along with any required configuration files and use them together
- Structure the .zip file to have a top-level folder containing all the files required for the service. You can also use sub-folders to organize the scripts and supporting files.
- The top-level folder name must match the name provided for the Service ID when defining the service (for example, `tomcat6.zip`, where `tomcat6` is the Service ID).
- Define the Lifecycle Action script path is defined relative to its location in the .zip file.
- Set the Access Link URL from within a script, be sure to configure the information when you [Model Applications Profiles](#).

In the simplest case, all lifecycle actions are contained in a single script. Service parameters are defined as part of the [Custom Service Definition](#) process and accessed in the scripts through an environment variables. See [Deployment Lifecycle Scripts > Utility Files](#) and [Pre-Defined Parameters > Environment Variables for N-Tier Deployments](#) for additional details.

The following sample service script:

1. Installs Tomcat web server
2. If the environmental variable `$cliqWARFile` is null, the script terminates; if not, the script moves the file referenced by the variable to a folder known by Tomcat web server.
3. Provides the commands to execute at the time of VM start and VM stop.

Sample Service Script

```
#!/bin/bash
exec
> >(tee -a /usr/local/osmosix/logs/service.log) 2>&1
echo
"Executing service script.."
.
/usr/local/cliqr/etc/userenv
#
main entry
case
$1 in
  install)
      yum
install -y tomcat tomcat-webapps tomcat-admin-webapps
      ;;
  deploy)
      if [ -z $cliqrWARFile ]; then
          exit 0
      fi
      cp $cliqrWARFile
/usr/share/tomcat/webapps
      ;;
  configure)
      ;;
  start)
      systemctl start tomcat
      ;;
  stop)
      systemctl stop tomcat
      ;;
  restart)
      ;;
  cleanup)
      ;;
  reload)
      ;;
  upgrade)
      ;;
  *)
      exit 127
      ;;
esac
```

To use this script as the service script for a custom service named **tomcatCentOS7**:

1. Make the script file executable (chmod 755).
2. Zip the file and give it the same name as the service it will be used in; in this case, **tomcatCentOS7.zip**.
3. Upload the zip file to a repository accessible to Workload Manager.
4. In the Workload Manager Edit Service page for this service, for **Agent Actions Bundle** field, specify the repository location and the path to the zip file.
5. Ensure that any environmental variables that require user input (in this case, \$cliqrWARFile) are included as service parameters in the Edit Service page.



The **Service ID** of your custom service must match the name of the service script bundle zip file, in this case, **tomcatCentOS7**.

See [New Application Profile](#) for additional context.

Custom Service Definition

Custom Service Version Definition

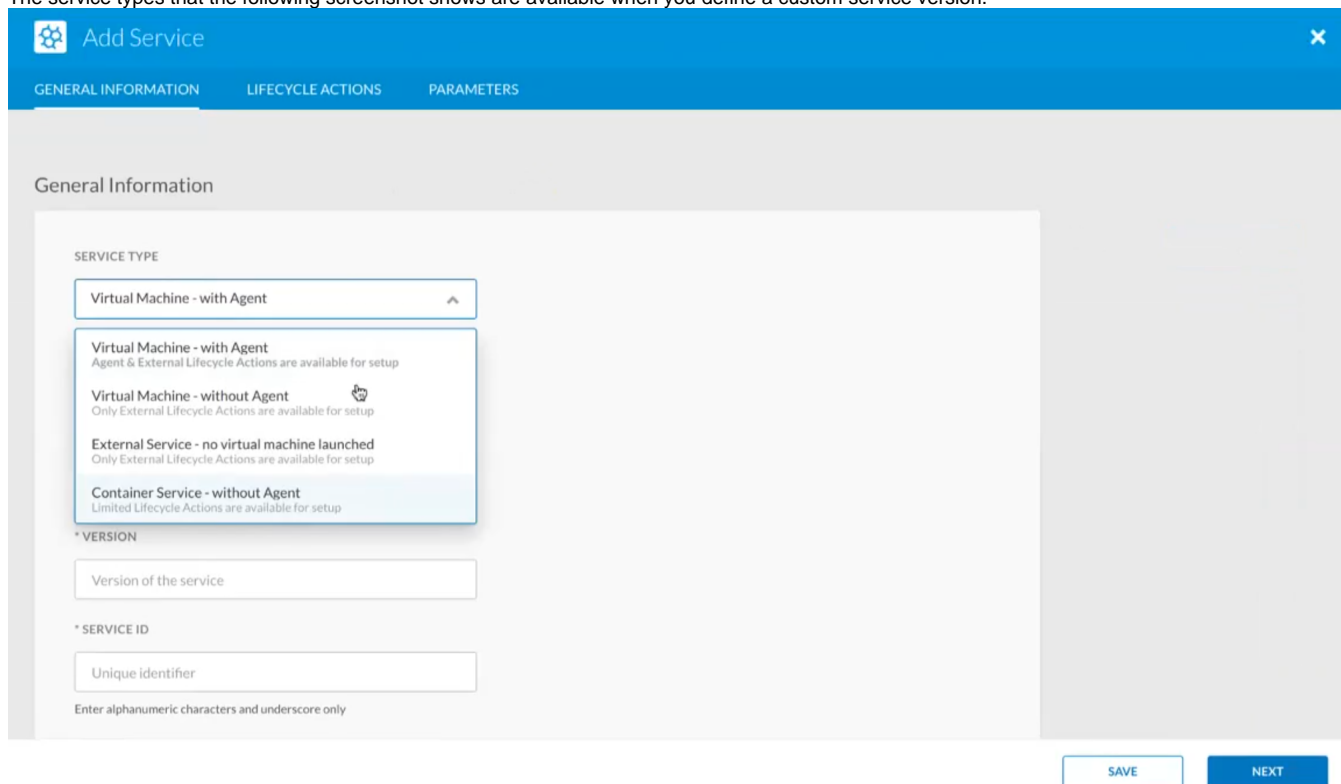
- [Overview](#)
- [Pagination for Listed Services](#)
- [Search by Name Filter](#)
- [Service Types](#)
- [Base OS Image Versions](#)
- [Prerequisites to Define Custom Services](#)
- [Guidelines to Define Custom Services](#)
- [Process to Define a Custom Service](#)
- [Using a Custom Service](#)

Workload Manager users have the flexibility to define their own service version on supported [OOB Logical Images](#). For instance, a user creates a Tomcat or Oracle WebLogic service version(s) on a specific hardened base OS image such as RHEL, CentOS, or Oracle Enterprise Linux. Accordingly, this user can define the scripts for different actions on each service version such as starting, stopping, restarting the service and so forth.

The Services list page displays 20 services by default. You can display additional services at will (clicking the **More** option) instead of all the services being loaded by default in this page.


As a limited number of services are displayed in this page (until you click the More option), the Search function allows you to search for a service by name. The search feature is listed below the Services page title.

The service types that the following screenshot shows are available when you define a custom service version.



The following table describes the service types.

Service Type	Description	Scripts /Hooks	Links
--------------	-------------	----------------	-------

Virtual Machine with a Management Agent (default)	<p>About 90% of the Workload Manager services use this type. In this case, it is nothing more than an application VM being launched with a service tied to it using an apt-get command based on the bundle script specification to install this service, and then run the required actions.</p> <p>If configured, the Pre VM Start script is the first script to be executed before the VM is launched.</p>	<ol style="list-style-type: none"> 1. Pre VM Start 2. Pre VM Init 3. Post VM Init 4. Pre VM Stop 5. Post VM Stop 	<ul style="list-style-type: none"> • Service Lifecycle Actions • OOB Services • External Initialization section in External Service
Virtual Machine without a Management Agent	<p>Use this service type to launch VMs in agent-less mode.</p> <p>While you cannot specify Agent Lifecycle Actions (because of the lack of a Management Agent), you can still specify External Initialization actions before the actual service is started.</p> <p>If there is no IP address before the actual service is started, the post/pre-init scripts inject the IP addresses and other configured information for each phase.</p>	<ol style="list-style-type: none"> 1. Pre VM Start 2. Pre VM Init 3. Post VM Init 4. Pre VM Stop 5. Post VM Stop 	<ul style="list-style-type: none"> • Supported OOB Logical Images • External Initialization section in External Service
External Service	<p>Use this service type to configure an external service that is not launching a VM.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  External service costs are not included in the cloud costs as it is considered a component of management costs. This cost is visible in the Account Details > Usage Details page. </div>	<ol style="list-style-type: none"> 1. Update 2. Start 3. Stop 	External Initialization section in External Service
Container Service	<p>Use this service type to configure container services.</p>	<ol style="list-style-type: none"> 1. Post start 2. Post stop 	Lifecycle Hooks section in Container Service

The out-of-box Base OS services are not tied to a specific version. Workload Manager provides an Ubuntu base OS service that maps to the supported versions for this service. You will find similar configurations for all other OS services.

One of the supported versions is specified as the default for each base OS image and users can change to any other supported version. See [OOB Logical Images](#) for a list of supported versions for each base OS image.

Verify these prerequisites before you define a custom service:

- Launch the Bundle Files as described in [Local Bundle Store \(Conditional\)](#).
- Review the list of [OOB Services](#) and [Understand Application Tier Properties](#) for the out-of-box services.

Define a custom service by following these simple guidelines:

- Define a skeleton service definition. The scripts can be in any language or format. Ensure to adhere to the [Service Administration > Guidelines to Configure Service Scripts](#).
- Workload Manager merely executes the command specified in the field for each phase of the [Service Lifecycle Actions](#).
- Read and understand the Service-Tier Scripts Defined in App Models explained in [Deployment Lifecycle Scripts](#).
- Upload the logo image **before** creating a service.

To add a custom service, (admins) follow this procedure:

1. Log into the Workload Manager UI as an Admin.
2. Access the Services tab: **Main Workload Manager Dashboard > Services**.
3. Click **Add Service** to add a new service.

4. In the Add Service, General Information page, enter the details of the new service.

5. Depending on the Management Agent for External availability, select one of the service types explained above from the dropdown list. The mandatory fields are validated: for example, provide a Service Version number ([Service Versions and Usage](#)), add at least one image, specify a default image, add a logo file, and so forth. This validation ensures that you add meaningful services in the Workload Manager.



While some information is specific to a version, other fields are common across all versions.

Examples of fields that are common to all versions:

- Service Logo
- Name
- Service ID (not editable)
- Category
- Description

Fields that are specific to a version:

- Version
- Protocol
- Cost per hour
- Supported Images
- Default Image

The following image displays the requirement to add a Service Logo as a missing field.


6. Add a Service Logo for this image. Click **Choose File** to upload the file to the local file system:

- The supported image formats are PNG and JPG.
- The image must be a square (W x H), not larger than 1MB.
- If not included, the logo space remains empty.



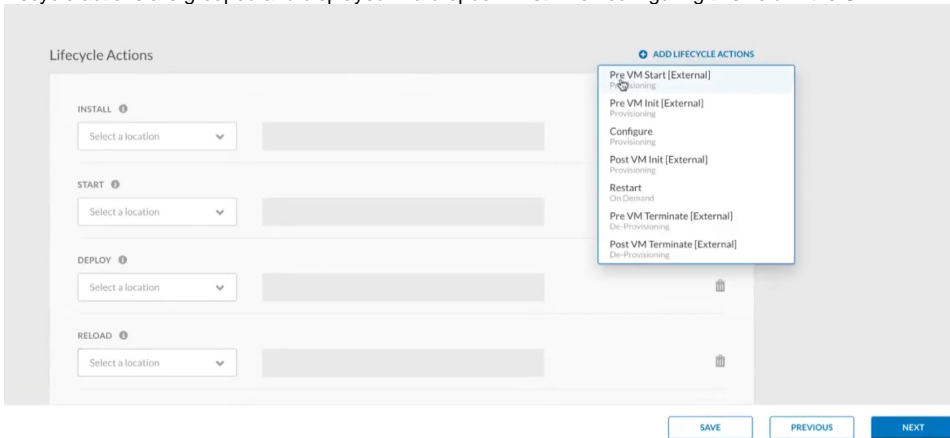
Logo file names for a Unison file synchronization process have a size limitation of 140 characters. If the file name is longer than this limit, then the Unison synchronization process fails recursively and other image files, including logo images, cannot be synchronized.

7. Provide the **Name** for this service.
8. Provide a Service ID. This must be a unique ID that uses only alphanumeric characters and/or underscore. If you are using scripts and created a .zip file, then use the same name as the .zip file. See [Guidelines to Configure Service Scripts](#) for additional details. For example, **tomcatCentOS7** is the name used in that example when creating the service lifecycle action script.
9. Add an optional description for this service.
10. Select *one* relevant Category.
 - a. After you define the service, this service will be displayed in the [Topology Modeler](#) Services tab.
 - b. The available categories are listed in the **Category** dropdown list.



- c. For example, if you add a Tomcat service, select **Web Server**, if you are adding a SQL database server, select **Database** and so forth. If you find that your new service straddles two services or does not fit into any other category, add it to the Custom Service group.
11. Select the Supported Images (VM-based services only). *At least one image is required.*
 - a. See [Services](#) for a list of supported versions for each service.
 - b. You can select multiple images for each service. By selecting multiple images in this field, you are allowing users in your enterprise the flexibility of selecting from multiple images when modeling or deploying applications. If you select multiple images, your lifecycle script must account for possible command differences between each operating system (for example, yum in CentOS vs. apt-get in Ubuntu).
 - c. Determine the underlying image for each service by selecting the applicable image for your deployment from the dropdown list. The list displays all [Operating Systems](#) supported by Workload Manager and other images that enterprises may have privately uploaded:
 - d. If you want to use an image that is different from the supported Base OS Images, you must create a logical image (see [Map Images](#)) and map it to its respective physical image for each cloud.
12. Select one Default Image from the selected supported images for this service (VM-based services only).
13. Assign a Default inbound [firewall](#) rule(s) that should be used by VMs running this service, if required (VM-based services only).
 - a. Select if the protocol should be TCP or UDP.
 - b. Add a firewall rule for each default port as applicable. For example, Port 8080 is the default port for Tomcat service.
 - c. Assign the ingress and egress port information.
14. Assign the Service Cost to allow enterprises to optionally charge an hourly cost for the service.
 - a. When modeling applications, you will not be charged when you include a charged service in your topology.
 - b. You are charged at the hourly service rate only if you deploy an application that has a chargeable service.
 - c. This cost adds up to the other costs incurred when deploying an application.
 - d. Allows admins the flexibility to charge individual users or sub-tenants for any service that is built and added to the Marketplace or Catalog. See [Cost and Fees](#) for additional context.

- Provide the applicable [Lifecycle Action\(s\)](#) in the Lifecycle Action tab by using scripts or command to execute the service on different actions. All lifecycle actions are grouped and displayed in a dropdown list when configuring this field in the UI.

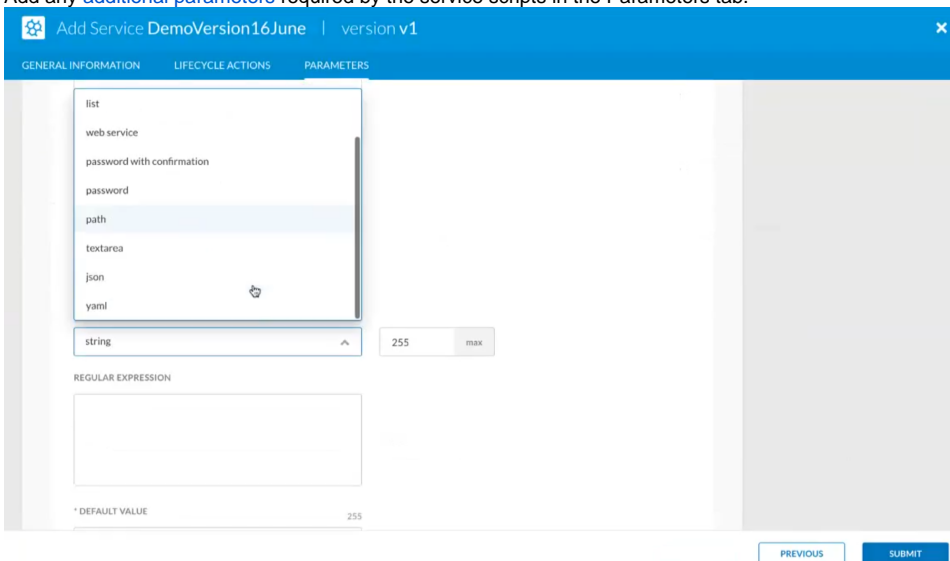


- Lifecycle actions can be one of the following:
 - Scripts: This is the most common option as scripts can be in any language or format. See [Guidelines to Configure Service Scripts](#) for additional details.
 - Commands: Workload Manager executes the command(s) specified in this field for the phase of the lifecycle
 - URLs: Downloadable by a get request
 - Workload Manager provides [multiple input locations](#) (Script Source) to execute scripts, commands, or URLs that can be executed at various phases of a service.
 - Services can be present in a repository that is already modeled in Workload Manager (*Repositories*) or hosted on any other server (*Other input*). If the bundle is hosted on any other server, then provide the URL of the script.
- In cases where you need to run external actions, define this section – Add the External Lifecycle Actions to manage the external service lifecycle.

Script Properties	The specified script is executed...
Pre-init Script	Before the service is launched
Post-init Script	After the service is launched
Pre-start Script	Before the service is started
Post-start Script	After the service is started
Pre-stop Script	Before the service is terminated
Post-stop Script	After the service is terminated

The External Actions Bundle file contains the scripts for external service lifecycle management. If you are configuring this file for a service, name this file as ServiceID.zip. For example, tomcat6.zip, where tomcat6 is the Service ID.

- Add any [additional parameters](#) required by the service scripts in the Parameters tab.





If you define any user-editable parameters, you can use one of two types of parameters when creating a service:

- **Service parameters:** Displayed in the *General Settings* section of the [Properties](#) pane in the Topology builder.
- **Deployment parameters:** Displayed in the *Deployment Parameters* section of the [Properties](#) pane in the Topology builder.

See [Using Parameters](#) for additional details.

18. Click **Save** to save this new service.

Now that the service is defined, it will appear on the Topology Modeler's Services pane. Once created, you can share the service across your sub-tenants. By default, the service is available to all users within the tenant. Users can access the service from the [Topology Modeler](#) when modeling applications.

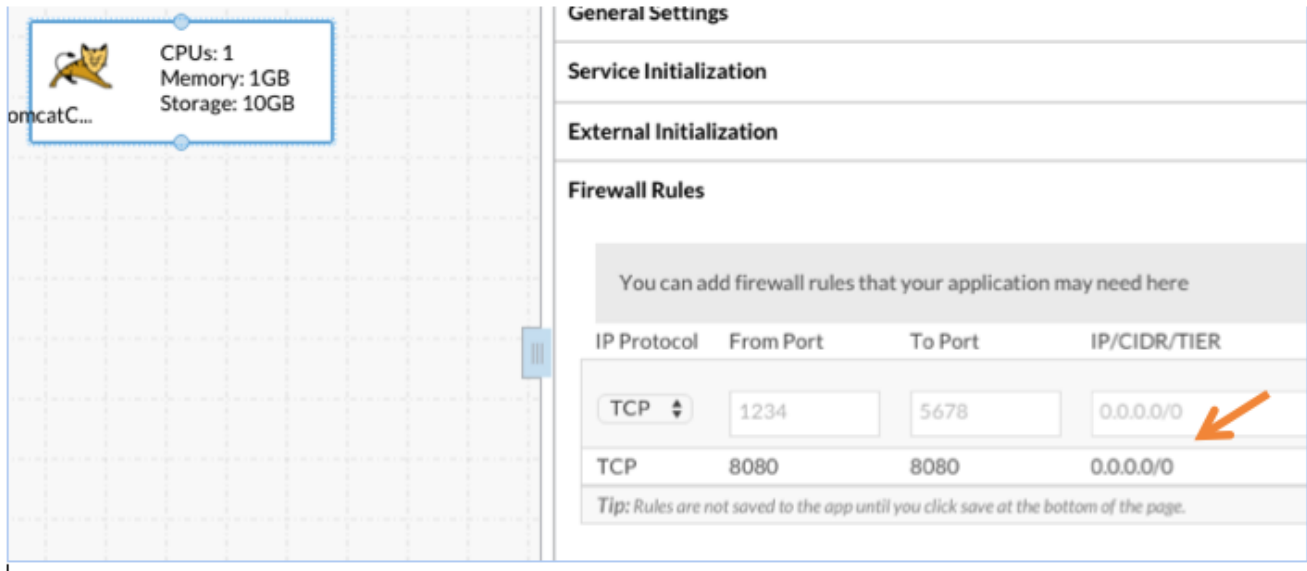
To use the newly-defined service users can drag it into the Topology Modeler as explained in [New Application Profile](#).

The new service is visible the Topology Modeler's Services Palette in the selected Category along with the specified Base Image as displayed in the following image. The service version defaults to v1 for pre-existing services that are upgraded to Workload Manager 5.4.0. If you have already configured multiple versions of a services, they are displayed in the Service Version dropdown list.

If you created additional parameter(s) specific to this service, you will see it as part of the General Settings in the [Application Tier Properties](#) pane. You can enter an optional war file from a repository you have configured within Workload Manager.

Firewall rules configured as part of this service are automatically set when you configure this service.

Firewall rules are set based on the service settings.



The screenshot displays the configuration interface for a service. On the left, a service card for 'omcatC...' is shown with a cat icon and specifications: CPUs: 1, Memory: 1GB, and Storage: 10GB. The right-hand panel is titled 'General Settings' and contains sections for 'Service Initialization', 'External Initialization', and 'Firewall Rules'. The 'Firewall Rules' section includes a text prompt: 'You can add firewall rules that your application may need here'. Below this is a table with columns for 'IP Protocol', 'From Port', 'To Port', and 'IP/CIDR/TIER'. The table contains two rows: one with 'TCP', '1234', '5678', and '0.0.0.0/0' (with an orange arrow pointing to the IP address), and another with 'TCP', '8080', '8080', and '0.0.0.0/0'. A tip at the bottom of the table states: 'Tip: Rules are not saved to the app until you click save at the bottom of the page.'

IP Protocol	From Port	To Port	IP/CIDR/TIER
TCP	1234	5678	0.0.0.0/0
TCP	8080	8080	0.0.0.0/0

Tip: Rules are not saved to the app until you click save at the bottom of the page.

External Service

External Service

- [Overview](#)
- [What Is an External Service?](#)
- [What Are External Lifecycle Actions?](#)
- [Custom Scripts Sources](#)
- [External Script Specification](#)
- [Writing an External Service Script](#)
 - [Script Language](#)
 - [Utility Bash Functions](#)
 - [Log Messages](#)
 - [Printing Results](#)
 - [Error Handling](#)
 - [Best Practice](#)
 - [Passing Information from External Service Scripts](#)
 - [Passing Information from External Initialization Scripts \(Application Level\)](#)
- [Script Timer](#)
- [Add an External Service](#)
- [Deploying an External Service](#)
- [Cost and Reports](#)

You can create an application using multiple tiers (see [Understand Application Tier Properties](#)) where you can stipulate each tier to use a different [OOB Services](#) or externally-provided service (third-party services). Additionally, you can define scripts for the each phase for each type of service *when* you add /edit a service.

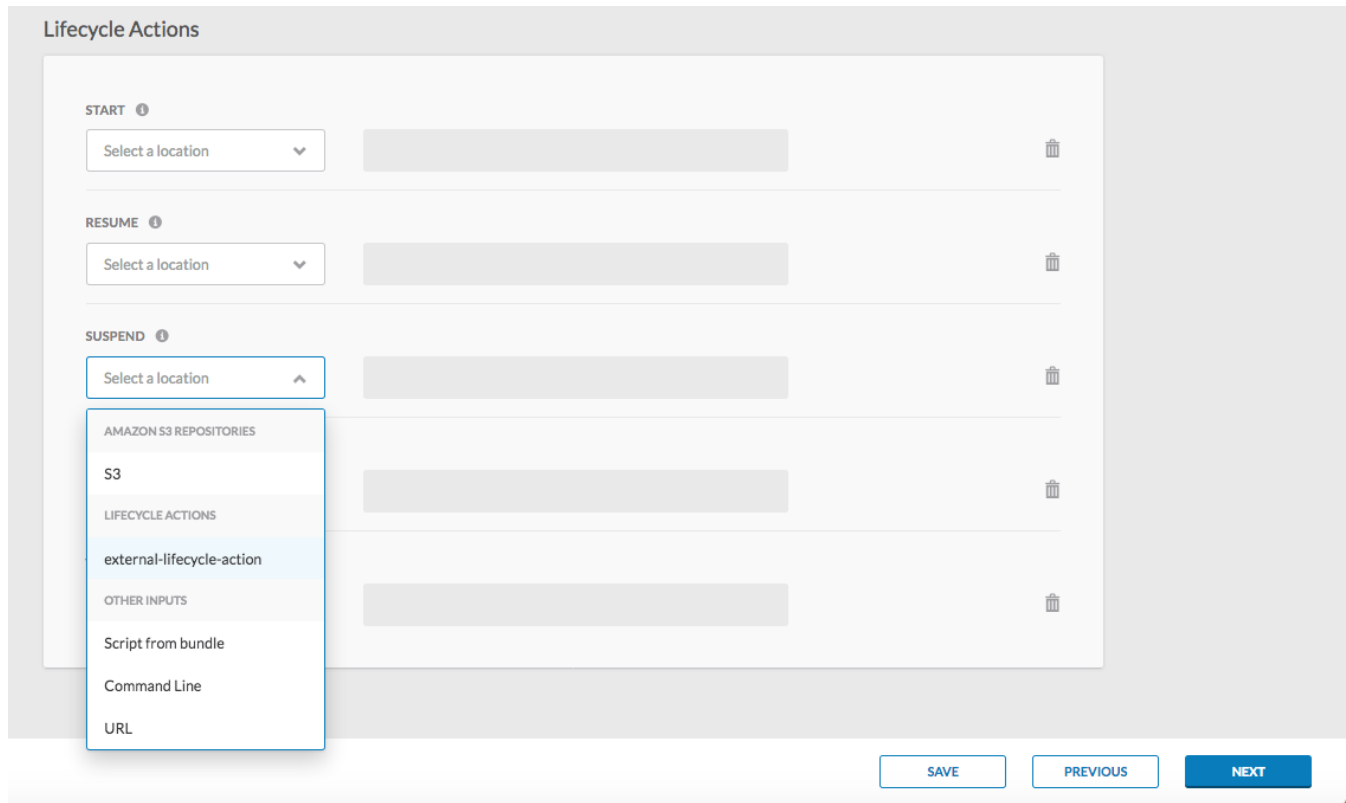
An **external service** is a service that you can define as a Workload Manager administrator. An external service does not have an associated VM – it is just a service that contains custom scripts provided by the administrator. The following image shows adding an external service.

Any script that is supported by the *External Service* type can now be added when you define a new service.

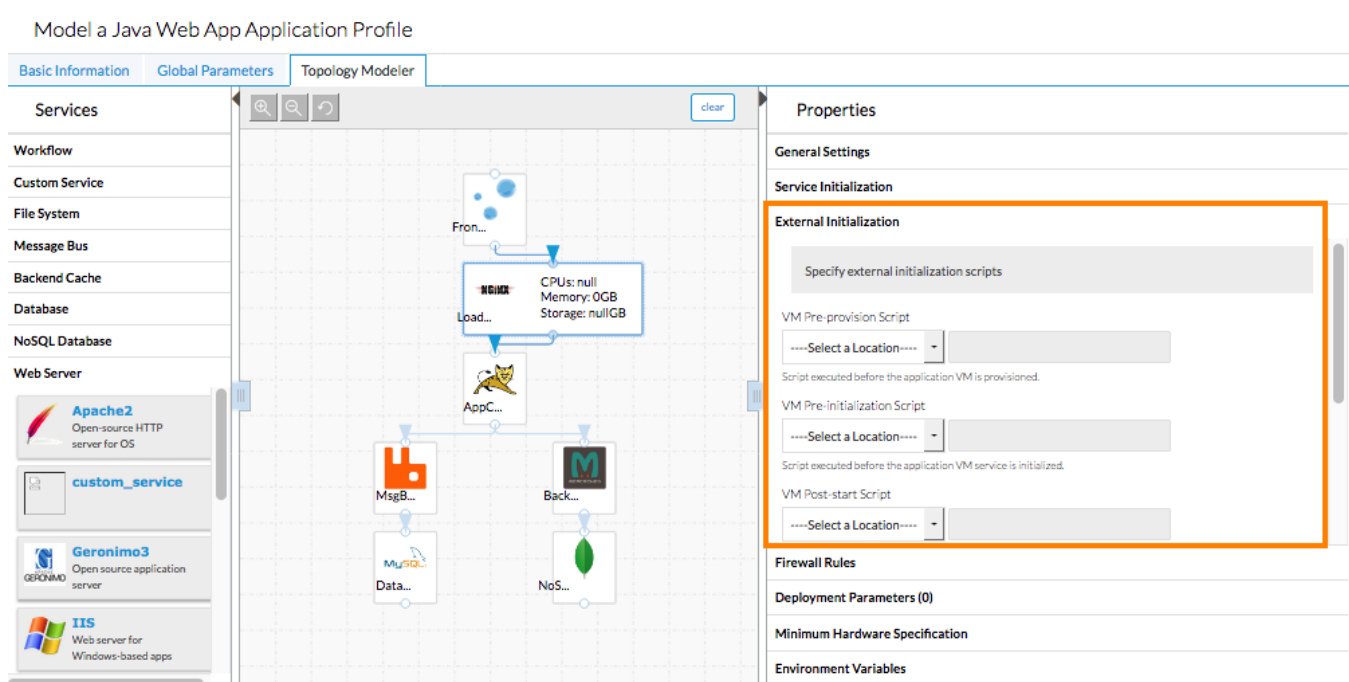
- To support services and applications, external callout scripts are executed in an isolated Docker container. While Cisco provides an [OOB Docker Service](#) image to execute callouts on any Workload Manager-supported cloud, you can also customize this container as specified in [Custom Docker Image for Scripts](#).
- When adding an External Service, you can assign any [OOB Group](#) to this service and the service (in the application service palette) is displayed in that category in the [Topology Modeler](#).

External Lifecycle Actions are actions that you can define for VM-based services or external services during a [Custom Service Definition](#) or during Application Definition.

The following image provides an example of a service-level definition.



The following image provides an example of an application-level definition.



This section allows you to augment the service by adding additional scripts that can be executed during various VM lifecycle events.

A custom script can reside in a Docker container (see next section) or a repository as identified in the [Deployment Lifecycle Scripts > Script Source Details](#) section.

The **Script from Bundle** option is not available for this service type.

External Service, External Initialization, and External Lifecycle Actions have some common areas like defining how scripts are initialized and executed when using these functions.


The following table identifies the **External Lifecycle Actions** that are specific to an *External Service*.

External Service Lifecycle Actions	The specified script is executed...
Update	When an updated IP address or a scaling operation dependency is specified when scaling up or scaling down.
Start	When a service initializes for a specific cluster or tier.
Stop	When the application terminates.
Suspend	When the application is powered off or shut down (not terminated).
Resume	When a suspended deployment resumes.


The following table identifies the **External Lifecycle Actions** that are specific to a *VM-based Service (with or without agent)*.

VM-Based Service External Lifecycle Actions	The specified script is executed...
Pre-VM Start	Before the VM is launched/provisioned
Pre-VM Init	After the IP address is returned and the application VM service is initialized
Post-VM Init	After the application VM service is started.
Pre-VM Terminate	Before the application VM is terminated.
Post-VM Terminate	The application VM is terminated.

The following table identifies the **External Initialization scripts** that are specific to a *VM-based service tier at the Application level*.

 For External Initialization Scripts, use `print_log` or `print_error` functions.

External Initialization Scripts at the Application Level	The specified script is executed...
VM Pre-Provision Script	Before the VM is launched/provisioned
VM Pre-Initialization Script	After the IP address is returned and the application VM service is initialized
VM Post-Start Script	After the application VM service is started.
VM Pre-Terminate Script	Before the application VM is terminated
VM Post-Terminate Script	application VM is terminated

 All service scripts are executed under its parent directory. To execute another script inside the current script, use the relative path of the working directory, where the parent directory is `/script` – the custom script is executed from the `/script` directory.

Script Language

The external service script is executed as a Linux bash script inside a Docker container. If you write the script in other languages, for example, Python, add the following line to the first line of your script:

```
#!/usr/bin/env python
```

Utility Bash Functions

Right before your script is executed, Workload Manager supplies all the application-specific parameters as environment variables. Some environment variables (`$CloudFamily` and `$region`) allow you to maintain the same script (that may be used for multiple cloud configurations) in a Docker container.

A `utils.sh` script in the root directory provides the utility bash functions. To include these functions, add the following:

```
. /utils.sh
```

Log Messages

All output is logged at the DEBUG level. Any output using **stdout** in the Docker container is caught by the CCO.

If you want to send log messages to the Workload Manager UI (task message list in the Job Details page), you must add delimiters around your log message so Workload Manager can execute it accordingly. The following is a Bash example:

```
echo "CloudCenter_EXTERNAL_SERVICE_LOG_MSG_START"
echo "log message here"
echo "CloudCenter_EXTERNAL_SERVICE_LOG_MSG_END"
```



Use the **print_log()** utility function in the utils.sh file to wrap your log message with delimiters.

Printing Results

If you want to send results to the Workload Manager UI, you must add delimiters around your result message so Workload Manager can execute it accordingly. The following is a Bash example:

```
echo "CloudCenter_EXTERNAL_SERVICE_RESULT_START"
echo "<JSON or YAML string>"
echo "CloudCenter_EXTERNAL_SERVICE_RESULT_END"
```

Be aware of the following Workload Manager requirements when writing your service scripts:

- The result data is case sensitive.
- The result data must use either JSON or YAML format.
- The result data must be wrapped with delimiters.



Use the **print_ext_service_result()** function in utils.sh to wrap your result with delimiters. See the *Passing Information from External Services* section for an example.

Error Handling

If your external service script encounters errors, be sure to provide a meaningful error message and exit the script with status code > 0.

To show the error message in the UI, add delimiters around your error message:

```
echo "CLIQR_EXTERNAL_SERVICE_ERR_MSG_START"
echo "error message here"
echo "CLIQR_EXTERNAL_SERVICE_ERR_MSG_END"
```



Use the **print_error()** function in utils.sh to wrap your error message with delimiters.

Best Practice

As a best practice, be sure to redirect any command execution output to a log file. You can even suppress this information – if not required.



Verify that the custom external service script does NOT dump excessive information.

When the external service script dumps excessive information, the deployments can sometimes result in an error.

Passing Information from External Service Scripts

External services may return parameters which in turn can be used by dependent tiers (any tier above the current tier).



If user scripts print a password for an external service, this password may be visible in the Deployment logs. Ensure that you are not printing sensitive passwords as part of a script output.

The Workload Manager external service can return the following parameters. These variables are injected as environment variables into dependent tiers.

- The **ipAddress** parameter is the IP address of the external service. For example, the IP address of an Amazon RDS instance.
- The **hostname** parameter is the DNS name of the external service. For example, the DNS name of an Amazon RDS instance.

Sample Script in YAML Format

```
#!/bin/bash

. /utils.sh

print_log "This is a basic log message"

result="hostName: testsite      #hostname parameter
ipAddress: 10.1.1.5           #ipAddress parameter"

print_ext_service_result "$result"
```

Sample Script in JSON Format

```
#!/bin/bash

. /utils.sh

print_log "This is a basic log message"

result="{
  \"hostName\": \"testsite\",
  \"ipAddress\": \"10.1.1.5\",
}"

print_ext_service_result "$result"
```

... returns the parameters to the dependent tier as displayed in the following *Sample userenv File* (available at `/usr/local/osmosix/etc`):

Sample userenv File

```
#passed parameters and variables
export CliqrTier_extService_1_IP="10.1.1.5"
export CliqrTier_extService_1_HOSTNAME="testsite"

#Inherited parameters and variables
export CliqrTier_extService_1_Cloud_Setting_networkName="sha-net01"
export CliqrTier_extService_1_Cloud_Setting_numNICs="1"
export CliqrTier_extService_1_Cloud_Setting_cloud="OpenstackDev-regionOne"
export CliqrTier_extService_1_Cloud_Setting_publicIpAllocate="true"
export CliqrTier_extService_1_Cloud_Setting_privateIPAllocationMode="DHCP"
export CliqrTier_extService_1_Cloud_Setting_TenantId="c8fe2db7a7cb490ba6dd1913f4e5c9c8"
export CliqrTier_extService_1_Cloud_Setting_account="3"
export CliqrTier_extService_1_Cloud_Setting_networkId="efdb81c9-eb02-4199-a97b-27aab5ec58df"
export CliqrTier_extService_1_Cloud_Setting_attachPublicIP="false"
export CliqrTier_extService_1_Cloud_Setting_TenantName="sha"
```

The injected properties are prefixed with:

```
CliqrTier_<tier_name>_<property name>
```

The **ipAddress** and **hostname** (if present in the JSON or YAML string) are propagated to the Workload Manager and displayed in the Job Details page. The injected properties are visible in the UI as highlighted in the following image:



The Workload Manager external service adds on the script details to any other information that already existed in the userenv file as displayed in the example above.

Passing Information from External Initialization Scripts (Application Level)

External Initialization scripts may return parameters which in turn can be used by associated tier.



You can pass parameters to an associated tier ONLY from the **VM Start Script** field.

The Initialization scripts can return variables that are injected as environment variables into the userenv file of the associated tier.

- The **environment** parameter is a key-value map that contains custom-defined environment variables. For example, the sample script below ...

Sample Script in YAML Format

```
#!/bin/bash

. /utils.sh

print_log "This is a basic log message"

result="environment:
  hello: world           #key-value environment parameter
  instanceName: test_instance #key-value environment parameter
  instanceType: dummy     #key-value environment parameter
  serviceType: custom     #key-value environment parameter
"

print_ext_service_result "$result
```

Sample Script in JSON Format

```
#!/bin/bash

. /utils.sh

print_log "This is a basic log message"

result="{
  \"environment\":{
    \"hello\": \"world\",
    \"instanceName\": \"test_instance\",
    \"instanceType\": \"dummy\",
    \"serviceType\": \"custom\"
  }
}"

print_ext_service_result "$result"
```

... returns the parameters to the associated tier as displayed in the following *Sample userenv File* (available at /usr/local/osmosix/etc):

Sample userenv File

```
#passed parameters and variables

export hello="world"
export instanceType="dummy"
export instanceName="test_instance"
export serviceType="custom"
```



The Workload Manager external service adds on the script details to any other information that already existed in the userenv file as displayed in the example above.

It is possible for an external service to fall into an infinite loop if using the stop/start external initialization scripts. This situation may cause the Docker container to run forever. In these cases, use the Workload Manager's `cliqrContainerExecuteScriptTimeout` property as a global parameter when modelling applications.

When you model an application in the topology builder, you can add a global parameter called **cliqrContainerExecuteScriptTimeout**.

You can specify this parameter as a floating point number that accepts the following options:

- s = seconds (default)
- m = minutes
- h = hours
- d = days

For example:

`docker.container.scriptTimeoutDuration=10m` will be overridden by:

cliqrContainerExecuteScriptTimeout=10m

This global parameter overrides the `docker.container.scriptTimeoutDuration` property in the `gateway.properties` and restricts the Docker container from running beyond 10 minutes in case the external service falls into an infinite loop.

To add an External Service, follow this process.

1. Access the Workload Manager UI > **Main Menu Dashboard** > **Services** > **Add Service** page.
2. Click **External Service** to select this service type.
3. Proceed as you would for a [Custom Service Definition](#).
4. Configure the scripts for each Service-Based **External Initialization** described in the table above.
5. Click **Save**.

To deploy an external service and ensure that it passes information from an external service to any dependent tier (any tier above the current tier), follow this procedure.

1. Define input parameters in a script and save the script in an accessible location as explained in the *Passing Information from External Services* section.

2. Add an [External Service](#) and provide the script location in one of the External Lifecycle Actions (for example, Start or Update). See the *External Initialization Scripts* section above for additional context.
3. Save the external service.
4. Model an application using this external service.
5. Launch the application. At this point, Workload Manager adds the injected properties to the dependent tier(s).

The VM-Hours column in the [Application Deployments Report](#) includes the usage hours for the External Service. See [Workload Manager Cost and Fees](#) for additional context.

Container Service

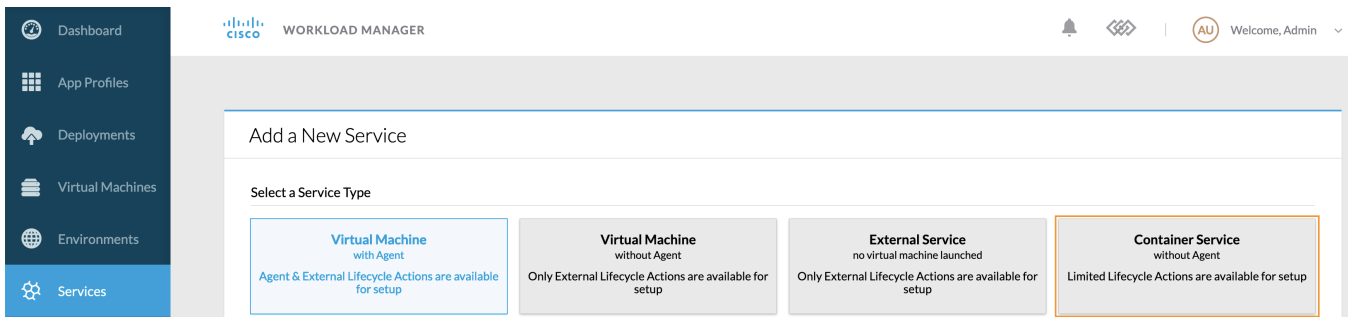
Container Service

- [Overview](#)
- [Images](#)
- [Container Ports](#)
- [Lifecycle Hooks](#)
- [Container Tier Properties in an Application Profile](#)
- [Cost and Reports](#)

Rather than being based on a logical VM image like a VM-based service, a container-based service is based on a Docker image. Other differences between VM-based services and container-based services:

- VM-based services have firewall rules, while container-based services have Container port.
- VM-based services have external lifecycle actions, while container-based services have lifecycle hooks.
- Container-based services do not have service parameters.

To add a new container-based service, from the Services page, click the Add Service link in the upper right of the screen. This brings you to the Add a New Service page as shown in the screenshot below. Select Container Service



A Container Service neither has an Agent nor does it have associated Lifecycle Actions.

When adding a Container Service, you must also complete the service definition fields not specific to VM-based services as described in [Custom Service Definition](#).



If you place the container service under the Custom Service category, you cannot select multiple replicas during deployment.

There are three sections of the service definition form that are unique to container services:

- Images
- Container Ports
- Lifecycle Hooks

Complete the input for these section as described below.

The following screenshot shows the Images fields.

Images

Image Prefix *
dockerhub.cisco.com/cloudcenter-dev-dc

Default Image Version
latest
Leave empty to default to latest version


Image Initialization Command
/bin/sh
-c
while true; do date; sleep 1; echo \$0
\$1;done

Place each command fragment on a separate line

Image Initialization Arguments
hello
world

Place each argument on a separate line

Complete the fields in the Images section as described in the following table.

Input Field	Usage
Image Prefix	Provide the full URL or relative path of the image. If you provide the relative path, the target Kubernetes cloud will use the associated image in the Docker hub.
Default Image Version	Optional. If left blank, the most recent version of the image at the location specified in the Image Prefix field is used. <div style="border: 1px solid #ffc107; padding: 5px;">  Workload Manager does not do a pre-deploy check to see if the image version you enter is valid. If you enter an invalid version name, the associated deployment will fail. </div>
Image Initialization Command	Optional. Overrides the image's entrypoint command with the specified command. The command is entered as an array of strings, each string on its own line, where each string represents the corresponding "fragment" of the entire command string. Split the command string into fragments as you would if you were to specify it as a <i>command</i> array of strings in a Kubernetes pod configuration YAML file. See https://kubernetes.io/docs/tasks/configure-pod-container/attach-handler-lifecycle-event/ for an example.
Image Initialization Arguments	Optional. Lets you pass specific arguments to the command specified in the Image Initialization Command field (if one is specified), or to the default entrypoint in the image. Each argument must be on a separate line.

In a Container Service, the Container Ports field refers to the exact port and protocol that the container listens on and must be exposed for external access.

The Container Service can expose more than one port. For example, a Web Server container can expose both Port 80 and Port 443.

A screenshot of this section is shown below.


Container Ports

Protocol	Name	Port	Actions
Select Protocol		1234	Add
TCP	cumuluswp	80	

When you add a container port in the service definition, it will use a Kubernetes service type of Cluster IP.

Lifecycle Hooks are actions that you can define for container-based services in a service definition. These scripts are executed inside the container when the application profile tier that has the Container Service is deployed on Kubernetes.

Any script in a HTTP repository that is supported by the new *Container Service* type can be added when you add a new service.

 The Container service does not support IPAM and VM naming callout scripts.

The following image shows the lifecycle hooks input fields.

Lifecycle Hooks

Container Post Start

----Select a Location----

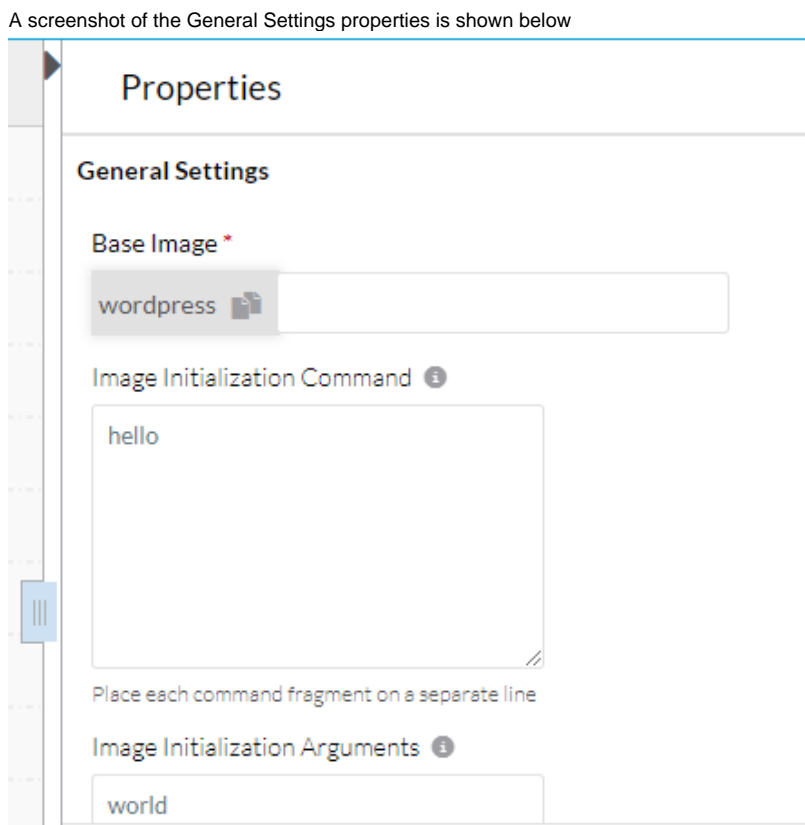
Container Pre Stop

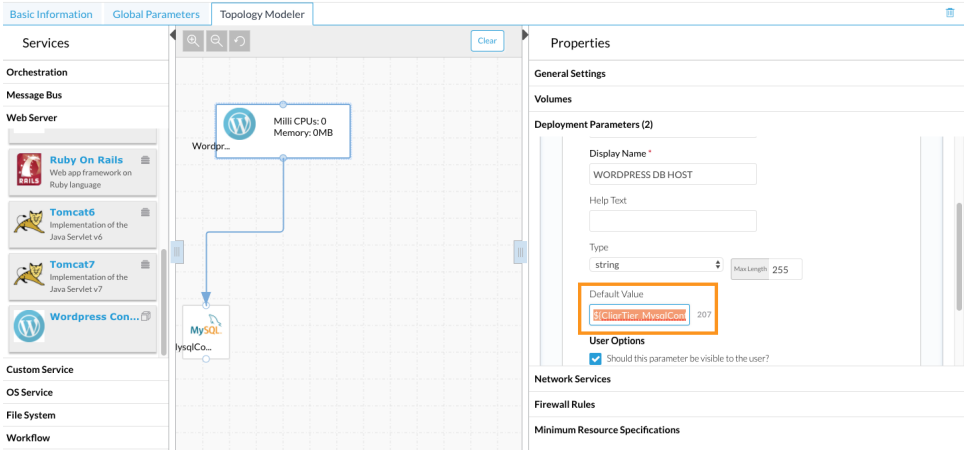
----Select a Location----

You must first select the type of lifecycle hook from the dropdown: either one of the listed HTTP repositories in the dropdown, URL or Command.

- If you select one of the HTTP repositories in the dropdown, the adjacent field to the right is used to enter the path of the command relative to that repository.
- If you select URL from the dropdown, the adjacent field to the right is used to enter the URL.
- If you select Command from the dropdown, the adjacent data entry area to the right is used to enter an array of strings, each string on its own line, where each string represents each space delimited fragment of the entire command string. Split the command string into fragments as you would if you were to specify it as a command array of strings in a Kubernetes pod configuration YAML file. See <https://kubernetes.io/docs/tasks/configure-pod-container/attach-handler-lifecycle-event/> for an example.

The following table describes the container-specific tier properties. The other tier properties are discussed in [Understand Application Tier Properties](#).

Properties	Fields	Description
General Settings	Base Image Version Tag Image Initialization Command Image Initialization Arguments	<p>A screenshot of the General Settings properties is shown below</p>  <p>The base image version tag, image initialization command, and image initialization arguments are all inherited from the service definition (see <i>Images</i> above) but can be overwritten in the General Settings section. If the base image version tag is left blank, the latest version is assumed. If you specify an image version tag, that tag is carried over to Page 1 of Deploy form in the per tier settings for that tier. In the deploy form you can delete that version tag or overwrite it with another one.</p>

Volumes	<ul style="list-style-type: none"> • Mount Path • Default Size 	Provide the mount path to the application that will be using this image along with the default size being used for this volume.
Deployment Parameters	Add a Parameter	<p>If deployment parameters have been defined at the service-level, those parameters are inherited and displayed here.</p> <p>Various Kubernetes network-related parameters can be referenced from within the Deployment Parameters section of the topology modeler. See Pre-Defined Parameters > Kubernetes Container Service Parameters for a list of these parameters.</p> <p>As an example, to pass the Internal Endpoint of a container-based MySQL service named "MysqlContainer_2" to the Wordpress tier of a Wordpress application, enter</p> <pre> \${CliqTier_MysqlContainer_2_ClusterIP_Endpoint} </pre> <p>as the default value for the WORDPRESS DB HOST parameter in the Topology Modeler, as shown in the following screenshot.</p> <p>Edit "Wordpress Container (LB)" Application Profile Version: 1 (Revision: 29)</p>  <p>You can also add parameters specific to the deployment. See Using Parameters for additional context on adding deployment-level parameters.</p>

Network Services

- Service Type
- Service Port Number
- Service Name
- Generate Unique ID

The following screenshot shows the Network Services parameters.

Network Services

	Service Port	Container Port / Protocol	Service Name	Actions
Cluster IP ▾	<input type="text"/>	80 / TCP ▾	<input type="text"/>	<input type="button" value="Add"/>
Cluster IP	80	80 / TCP	cumuluswp	<input type="button" value="Delete"/>

Network Services are not saved to the app until you click save at the bottom of the page.

Generate Unique ID ⓘ








The Container Ports you defined in the service definition are inherited and displayed here. If your tier contains multiple containers in a single pod, the container ports listed will be the union of all container ports for all containers in the pod. You can add additional container ports here by selecting the service type from the dropdown (Cluster IP, Node Port, or Load Balancer), port number, service name, and then click the **Add** button to add the port to the tier. The additional ports you create here apply to all containers in the pod.

By default, the service name created by Workload Manager and passed to the Kubernetes endpoint to deploy the tier will be the service name entered by the user appended by a hyphen, the job ID, another hyphen, and a 6 character random string. If the Generate Unique ID toggle is switched off in the application profile, at deploy time, the application will be deployed with the service name identical to name specified in the application profile without any characters appended. If the user-specified service name is not unique relative to all other existing service port names in the Kubernetes namespace associated with the Kubernetes account used for the deployment, the deployment will fail.

By default, Workload Manager will append a hyphen, the job ID, another hyphen, and a 6 character random string to the end of each service name before deploying the pod. This ensures that the service name is unique in the namespace. If you feel confident that the service names you assigned in the service definition and application profile will be unique in the namespace where you deploy the containers, you may turn the Generate Unique ID toggle to *OFF*.



If your user-specified service name is not unique relative to all other existing service port names in the Kubernetes namespace associated with the Kubernetes account used for the deployment, and you turn the Generate Unique ID toggle off in the application profile, the deployment will fail.

<p>Firewall Rules</p>	<p>Container Port / Protocol</p>	<p>Firewall Rules define who can access the container.</p> <p>By default, when you add Container Service to an application profile, a default firewall rule is added for each container port in the service to be accessed from any IP on the Internet.</p> <p>However, if you enabled Inter-Tier Communication (Firewall Rules) in the <i>Basic Parameters</i> section for an application profile, then the topmost tier has a default firewall rule added for each container port in the service to be accessed from any IP on the internet. See Security and Firewall Rules > Inter-Tier Communication (Firewall Rules) for additional context.</p> <p>For other dependent tiers, the default firewall rule is added for each container port in the service that is accessed from the dependent tier.</p> <p>You can choose to keep the pre-configured firewall definition or add/edit any firewall rules as required. All firewall rules are optional in this field.</p> <div style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;"> <p> In the Workload Manager UI, the Column Name changes to Container Port/Protocol if you are deploying a container service.</p> </div> <p>The following screenshot shows firewall rules.</p> <p>Firewall Rules</p> <p>You can add firewall rules that your application may need here</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="border: 2px solid orange;">Container Port / Protocol</th> <th>IP/CIDR/TIER</th> <th>Actions</th> </tr> </thead> <tbody> <tr> <td>TCP / 80</td> <td>0.0.0.0/0</td> <td style="text-align: right;">Add</td> </tr> <tr> <td>80</td> <td>TCP</td> <td style="text-align: right;">0.0.0.0/0 ▼ </td> </tr> <tr> <td>99</td> <td>TCP</td> <td style="text-align: right;">0.0.0.0/0 ▲ </td> </tr> </tbody> </table> <p><i>Rules are not saved to the app until you click Save at the bottom of the page.</i></p>	Container Port / Protocol	IP/CIDR/TIER	Actions	TCP / 80	0.0.0.0/0	Add	80	TCP	0.0.0.0/0 ▼ 	99	TCP	0.0.0.0/0 ▲ 
Container Port / Protocol	IP/CIDR/TIER	Actions												
TCP / 80	0.0.0.0/0	Add												
80	TCP	0.0.0.0/0 ▼ 												
99	TCP	0.0.0.0/0 ▲ 												
<p>Minimum Resource Specification</p>	<ul style="list-style-type: none"> • Mill iCP Us • Me mo ry 	<p>One thousand MilliCPUs is equal to 1 CPU and Memory is measured in bytes. This configuration depends on the container capabilities. See Managing Compute Resources for Containers for additional context.</p>												

Billing and Reporting are currently not supported for Container Services.

Service Lifecycle Actions

Service Lifecycle Actions

- [Overview](#)
- [Lifecycle Actions Based on Groups](#)
- [Lifecycle Actions Based on Service Type](#)
- [Agent Lifecycle Actions](#)
- [Container Lifecycle Hooks](#)
- [External Lifecycle Actions](#)

A service goes through various phases as it becomes operational. Service Lifecycle Actions allow enterprises to define a script or command that must be executed during different service lifecycle phases.

The root or tenant administrator can import and add services that are specific to their enterprise using the **Services > Add Service** function. The following screenshot shows this function.

The screenshot displays the 'Services' page in the Cisco Workload Manager interface. At the top, there is a search bar and buttons for 'IMPORT' and 'ADD SERVICE'. Below this is a table with the following columns: NAME, NO. OF VERSIONS, TYPE, CATEGORY, DESCRIPTION, COST PER HOUR, and ACTIONS. The table contains five entries:

NAME	NO. OF VERSIONS	TYPE	CATEGORY	DESCRIPTION	COST PER HOUR	ACTIONS
01 Agent Linux	3	Virtual Machine with Agent	Custom Service		0\$/hour	▼
01 Linux Agent CRUD	1	Virtual Machine with Agent	Custom Service	Sample Linux Custom Agent Ser...	0.01\$/hour	▼
01 Linux Agent Multiple Versions	3	Virtual Machine with Agent	Custom Service	v3 version	0.5\$/hour	▼
02 Agent Windows	1	Virtual Machine with Agent	Custom Service		0\$/hour	▼
02 Windows Agent Mod	1	Virtual Machine with Agent	Custom Service		0\$/hour	▼

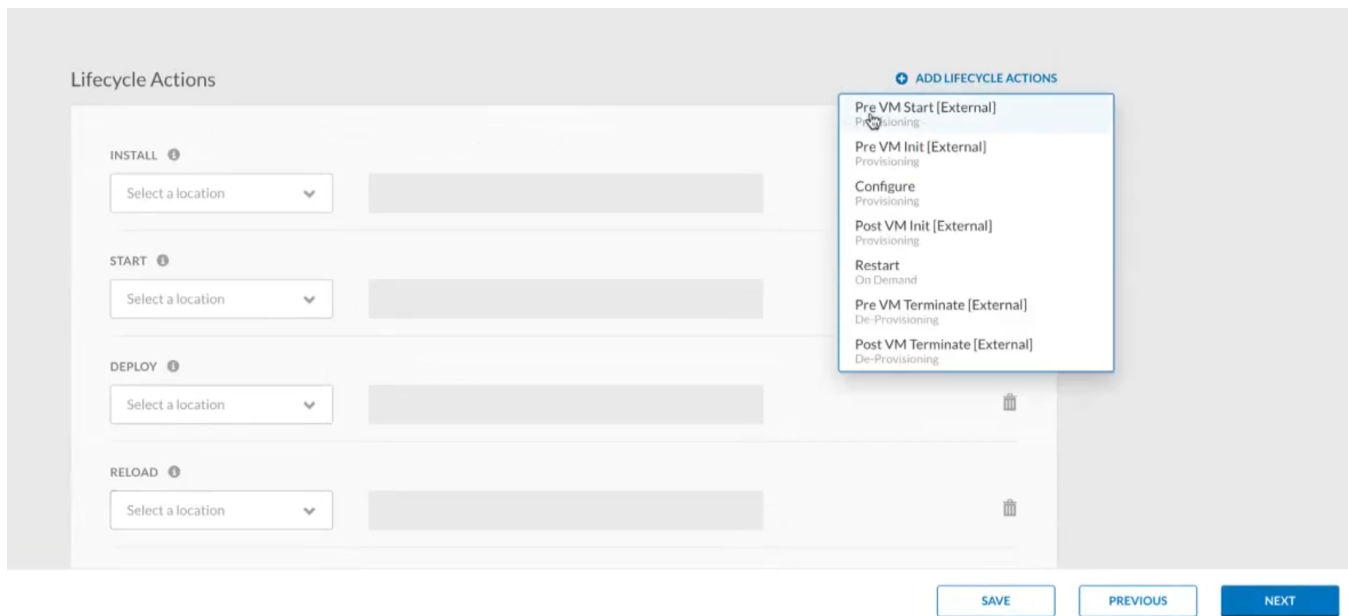
See [Custom Service Definition](#) for additional context.

The type of lifecycle actions that may be defined for a service depend on the type of service as shown in the following table.

Service Type	VM with Agent	VM w/o Agent	External (no VM)	Container
Lifecycle Action Types	Agent lifecycle actions, External lifecycle actions	External lifecycle actions	External lifecycle actions	Container Lifecycle hooks

Links to the individual lifecycle actions are specified in corresponding fields in the Edit Service and Add a Service pages.

Lifecycle actions are now displayed in a dropdown list (displayed in the following image) with the groups where each action can be used when configuring a service.



The order and grouping for lifecycle actions for each service type is relevant when importing and merging services.

The lifecycle actions order for each group are shown in the following table.


Group	Associated Lifecycle Actions
Provisioning	<ol style="list-style-type: none"> 1. Pre-VM Start 2. Install 3. Pre-VM Init 4. Configure 5. Post-VM Init 6. Start 7. Deploy
On-demand	<ul style="list-style-type: none"> • Reload • Restart • Upgrade • Only for Extension Service <ul style="list-style-type: none"> • Resume • Suspend • Update
De-provisioning	<ol style="list-style-type: none"> 1. Pre-VM Terminate 2. Terminate 3. Post-VM Terminate 4. Clean up
Container	<ol style="list-style-type: none"> 1. Container Post-Start 2. Container Pre-Terminate

The lifecycle actions order for each service type are shown in the following table.

Service Type	Associated Lifecycle Actions
--------------	------------------------------

VM with Agent	<ol style="list-style-type: none"> 1. Install 2. Start 3. Deploy 4. Reload 5. Upgrade 6. Terminate 7. Clean up
VM without Agent	<ol style="list-style-type: none"> 1. Pre-VM Start 2. Pre-VM Init 3. Post-VM Init 4. Pre-VM Terminate 5. Post-VM Terminate
External Service	<ol style="list-style-type: none"> 1. Start 2. Resume 3. Suspend 4. Update 5. Terminate
Container Service	<ol style="list-style-type: none"> 1. Container Post-Start 2. Container Pre-Terminate

The Agent lifecycle actions are only available for services based on VMs with the Workload Manager agent installed. The following table describes the agent lifecycle actions.

UI Service Action	<i>actionName</i> API Enumeration	Description
Install	INSTALL	<p>The install service phase is implemented when the node first comes up.</p> <p>Similar to the image installation process, Workload Manager also abstracts all supported OOB Services into small, individual files made available through The CloudCenter Worker1 image. When you drag and drop a service (any supported or customer-defined service) from the Topology Modeler's Services tab into the Topology Modeler graphical interface, you are instructing Workload Manager to use this service to model or deploy your application. At this point, all service files (ZIP format), regardless of each being an OOB or Custom Service Definition, are automatically extracted (from their respective locations — either Package Store (repo.cliqr.com) or your own Artifact Repository to the /usr/local/osmosix/ directory. The ZIP file for each service contains the root folder. For example, when you use the Tomcat7 service, the extracted file contains a root folder called Tomcat7. This /usr/local/osmosix/ Tomcat7 folder contains the scripts related to this service. Use this perspective accurately when you call additional scripts and parameters (see Parameters and Macros) in the Topology Modeler Properties tab.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  The install phase does not have any environment variables. </div>
Deploy	DEPLOY	Once all nodes are initialized and started, you must specify this action for services that need these files deployed into a particular service at the time of deployment.
Configure	CONFIGURE	Modify a configuration based on a requirement – If an application like NginX requires all the IP addresses for each application tier, you can configure this service to perform the related actions.
Start	START	<p>A service start action only occurs when all the nodes in the deployment are up and running.</p> <p>Some enterprises may have a start.sh script specifically for starting the services. If you have the start.sh file in the /usr/local/osmosix/<i>service</i> directory, this file executes each time the service starts.</p> <p>When you reboot a service, the IP address may change. This action may need you to reconfigure all connected IP addresses. So you need to call the Configure action and the Start action to complete the lifecycle process.</p>
Stop	STOP	When shutting down the node, you can perform related cleanup actions.
Restart	RESTART	When you restart the system from the UI after shutting it down, you can specify what actions need to be performed.
Reload	RELOAD	The IP address changes for connected nodes. For example, if any connect node restarts for any reason in a three-tier deployment, you can configure a reload command to internally reload the affected services.
Upgrade	UPGRADE	When upgrading the deployment, you can identify the dependent factors for each service.

Clean Up	CLEANUP	When a node is being terminated, you can specify the related clean up command, script, or URL.
-----------------	---------	--

Container lifecycle hooks are only available for container-based services. The following table describes the container lifecycle hooks.

UI Service Action	<i>actionName</i> API Enumeration	Description
Post Start	CONTAINER_POST_START	After the container is started – see https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/ for additional context.
Pre Stop	CONTAINER_PRE_STOP	Before the container is terminated. – see https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/ for additional context.

External lifecycle actions are available for services based on VMs (with or without agent) and external services. They are the only lifecycle actions available for agentless VMs and external services. See [External Service](#) for details on creating creating scripts for and linking to external lifecycle actions.

Service Versions and Usage

Service Versions and Usage

- [Overview](#)
- [Create a New Service Version](#)
- [Number of Service Versions](#)
- [Details for Each Service Version](#)
- [Deleting a Service Version](#)
- [Version Usage](#)

Similar to application profiles, you can create multiple versions for each service. The Workload Manager automatically converts an existing service to Version 1 (v1) and you do not need to explicitly make any configuration changes. When you add a new service, you must assign a unique service version for each new service configuration.

The *Base Version* can be one of any existing version versions for a service. The Base Version is used to create a new service version. In the following image, the new Version (v5) is created using the base version of the Agent Linux service.

Version for 01 Agent Linux

* BASE VERSION ⓘ

v1

* VERSION

v5

DESCRIPTION

This is a demo version

CANCEL SAVE

To create a new version for a service, click the **Actions** dropdown icon and select **Create Version**.

WORKLOAD MANAGER

Services

Search

IMPORT ADD SERVICE

NAME	NO. OF VERSIONS	TYPE	CATEGORY	DESCRIPTION	COST PER HOUR	ACTIONS
01 Agent Linux	3	Virtual Machine with Agent	Custom Service		0\$/hour	Export, View Usage, Edit, Delete, Share, Create Version
01 Linux Agent CRUD	1	Virtual Machine with Agent	Custom Service	Sample Linux Custom Agent Ser...	0.01\$/hour	
01 Linux Agent Multiple Versions	3	Virtual Machine with Agent	Custom Service	v3 version	0.5\$/hour	

In the Version of *ServiceName* popup, select the Base Version for the assign a version number for the new service version that you will be creating.

Once you click **Save** to create the new version, you are immediately placed in the Edit Service page for the new version so you can make the required changes for this new version. Follow the [Custom Service Definition](#) page to create a new Service Version.

The Services list page lists the number of versions that were created in the No of Versions column as displayed in the following image.

Services

Q Search IMPORT ADD SERVICE

<input type="checkbox"/>	NAME ↑	NO. OF VERSIONS	TYPE	CATEGORY	DESCRIPTION	COST PER HOUR	ACTIONS
<input type="checkbox"/>	01 Agent Linux	3	Virtual Machine with Agent	Custom Service		0\$/hour	<input type="button" value="v"/>
<input type="checkbox"/>	01 Linux Agent CRUD	1	Virtual Machine with Agent	Custom Service	Sample Linux Custom Agent Ser...	0.01\$/hour	<input type="button" value="v"/>
<input type="checkbox"/>	01 Linux Agent Multiple Versions	3	Virtual Machine with Agent	Custom Service	v3 version	0.5\$/hour	<input type="button" value="v"/>

Access the Edit Service page for a particular service and select the version to be edited from the dropdown list at the top of the page. The Service details are visible when you edit each version.

You can switch between the available service versions (by selecting the required version) in this page

You can also compare versions to detect the changes that you want to make to this new version or to verify settings and differences for each version.

You can delete a version by clicking the version dropdown displayed in the above image, selecting the version, and then clicking the **Delete** icon.

If a version is being used by an application profile/deployment, the **Delete** icon is disabled as displayed in the following image.

The last-updated version is the version that is displayed for each service. To view the resources using a particular service version, click the **Actions** dropdown icon and select **View Usage**.

Usage for 01 Agent Linux

VERSION

v1

My Tenant Sub-Tenants

APPLICATION PROFILES (1)

CustomLinux_...

DEPLOYMENTS (2)

SUSPENDING CustomLinux DEPLOYED edse

OK

The Usage for the *ServiceName* popup has a dropdown to select the required version in this popup. Each service version lists the Application Profile(s) and /or Deployment(s) where this service version is in use.

- The usage is listed on a per-tenant basis, for your tenant or any sub-tenants as applicable.
- If a version is not in use, no information is displayed as shown in the following image.

Usage for 01 Agent Linux

VERSION

v5

My Tenant Sub-Tenants

APPLICATION PROFILES (0)

There is no data to show at the moment.

DEPLOYMENTS (0)

There is no data to show at the moment

OK

Service Export/Import

Exporting and Importing Services

- [Overview](#)
- [Exporting Services](#)
 - [Select Multiple Services when Exporting](#)
 - [What Is Exported?](#)
 - [What Is Not Exported?](#)
- [Importing Services](#)
 - [Merge Duplicate Services during a Service Import](#)
 - [Requirements to Merge Services](#)
- [Deleting Services](#)

Similar to application profiles, you can import and export services and service versions.

Users can selectively export one or more services to easily share them with other CloudCenter Suite users.

- When you export a service, you are including all service versions in the export task.
- You need View (read only) service permissions to use this function.
- To export a service locate **Export** option in the Actions dropdown menu on the Services list page
- The Export function allows you to export a selected service as a .zip file that is saved to your local device.
- Click the checkbox to the left of each service to multi-select services.

Select Multiple Services when Exporting

You can select up to 5 services and export them to one .zip file. When you unzip the file, each service resides in its own folder.

If a service has multiple versions, all the versions are also exported. Each version is listed as a separate JSON file within the folder.

What Is Exported?

All aspects of the service configuration are exported to a ZIP file.

- The service metadata details for the exported file is also included in the exported folder.
- The repositories that are used by the service are listed in this file.

What Is Not Exported?

If an image is referenced by a service, then only the Image ID is exported — the image or the metadata is not exported.

Users can selectively import one or more services to easily share them with other CloudCenter Suite users.

- When you export a service, you are including all service versions in the export task.
- You need View and Modify (read and write) service permissions to use this function. With read-only permissions, service-related information that is shared from your parent tenant are not imported and details in relationship to your own tenant are imported as you do not have write permission for this service.
- To import a service, click the **Import** option that is located next to the Add Service button on the Services list page
- The Import function allows you to export a selected service as a .zip file to any other environment setup or a different tenant.

Merge Duplicate Services during a Service Import

When importing a service you also have the option to merge services by switching the toggle button to **On**. The default is Off.

For example, you have several basic Linux service instances in various tenants and another service where you have configured enhanced functions like an additional lifecycle action to invoke a specific script when VM starts in a particular tenant along with some additional parameters. In this case, you can import and merge the enhanced Linux service into the required tenants instead of creating these enhanced functions in each tenant individually.

Requirements to Merge Services

Be aware of the following requirements when you merge services:

- For services to be merged, the Names of both services must match.
- The Service ID and Type must be unique. Depending on your import location, it is possible that the ID and type details may differ. You must re-identify (or re-map) the new location details for each repository location in the Repository Mapping section for your merged setup.
- The location where the script or information reside must be one of the repository types listed in the [Artifact Repository](#).
- The lifecycle actions listed in the [Service Lifecycle Actions](#) page. Depending on lifecycle actions used in the merged service, you can re-identify (or re-map) each action to the same or different action in the Actions Mapping section for your merged setup.

Once you address these requirements, you can merge and import the service. Once successfully imported, the status banner displays at the top of the page and the merged/imported service is listed in the paginated Services page.

You can also select multiple services in preparation to delete them.

You require View and Modify (read and write) service permissions at your respective tenant levels to delete a service.

You can delete the service provided no version is being used by any application profile or deployment environment.

Deployment Environments

Deployment Environments

- [Deployment Environments Overview](#)
- [Environments Page](#)
- [Create a Deployment Environment](#)
- [Define Resource Placement](#)
- [Define Resource Validation](#)
- [Cloud-Specific Configurations](#)
- [Availability Sets and Zones](#)
- [Multi-Site, Multi-Account Deployment Considerations](#)

Deployment Environments Overview

Deployment Environments Overview


A deployment environment is a resource that consists of one or more associated cloud regions and cloud accounts and that has been set aside for specific deployment needs. Users deploy applications to deployment environments, and deployment environments can be shared with multiple users. For example, a development environment could be associated with a development cloud and a Production deployment environment could be associated with a production grade high-performance cloud. Users on a development team would have the ability to deploy only to the Development deployment environment and users on an operations team would have the ability to deploy only to the Production deployment environment.

Environments Page


Environments Page

You can add, view, manage, and delete deployment environments from the **Environments** page. To access this page, click the **Environments** tab from the main menu.

A screenshot of a sample **Environments** page is shown below.



To create a new deployment environment see [Create a Deployment Environment](#). All other actions available from the Environments page are summarized in the following table.

Deployment Environment Action	How Invoked	Notes
Edit	Click anywhere in the row for the deployment environment. OR Click on the Actions dropdown menu in the Actions column and select Edit .	Opens the edit form for the selected deployment environment. See Add a Deployment Environment, below, for details on each tab and each field in the deployment environment form.
Share	Click on the Actions dropdown menu in the Actions column and select Share .	Share a deployment environment. See Permission Control for details. <div style="border: 1px solid #ffc107; padding: 5px; background-color: #fff3cd;">  When you create a deployment environment and share it with a user without checking the Promote from option, be aware that the Migrate action <i>will not be available</i> when this user deploys an application that uses this deployment environment. </div>
Delete	Click on the Actions dropdown menu in the Actions column and select Delete .	Delete a deployment environment. If you choose to delete a configured deployment environment, the Delete Deployment Environment dialog box is displayed to confirm your intention. Your confirmation deletes the environment and displays a status message at the top of the Environment page.
Delete Deployments	Click on the Actions dropdown menu in the Actions column and select Delete Deployments	Deletes all deployments associated with the selected deployment environment.
Set Display Order of Environments	Use Up and Down arrows to the right of the deployment environment name	This determines the order of the deployment environment tiles in Page 1 of the Deploy form.

Create a Deployment Environment

Create a Deployment Environment

- [Overview](#)
- [General Settings Tab](#)
- [Cloud Settings Tab](#)
 - [Simplified Networks](#)
 - [Default Tier Cloud Settings](#)
 - [Available Instance Types](#)
 - [Resource Placement \(AWS, OpenStack, vCenter Only\)](#)
 - [Cloud Settings](#)
 - [Resource Validation](#)
 - [SSH Options](#)
- [Policy Settings Tab](#)

To create a new deployment environment, click the **New Environment** button in the upper right of the Environments page. This displays the Add Deployment Environment form which contains three tabs for setting the following parameters:

- **General Settings:** Name, Description, Tags, Enable ServiceNow integration, Require approval to deploy, Available cloud region-account combinations.
- **Cloud Settings:** For each region-account combination: Available instance types, Available deployment resources including per NIC network resources where applicable, Optional resource allocation and resource validation scripts, Optional alternate SSH key behavior for CloudCenter Suite to application VM communication.
- **Policy Settings:** Allowed aging, suspension and security policies, Require a suspension policy, Allow terminate/suspend protection.

Follow the steps for each tab described below.

To configure details in the **General Settings** tab, follow this procedure.

1. Provide the deployment environment **Name**
2. (Optional) Provide a **Description**.
3. (Optional) Configure a tag association for this deployment. See [System Tags](#) for additional details.




Google Cloud Nuance

Google Cloud does not support the attachment of tags to VMs. Although the Workload Manager UI will allow tags to be specified and shows success, tags are not added.

4. (Optional) Specify a ServiceNow extension from the dropdown menu. See [Extensions](#) for additional details
5. (Optional) Toggle the **Approval required to deploy to this environment** switch.
 - **ON:** Approval of an authorized user is required for the deployment of any application to this deployment environment.
 - a. If you want tenant users or user groups to request approval before deploying the environment.
 - b. If an environment requires approval, users and groups with this right can approve or deny deployments. If a job is submitted but pending approval, it displays *Pending* in the Job Status column for this deployment. The approving user or admin can **Approve** or **Reject** the deployment by clicking the corresponding action in the Action List. Either way, a confirmation popup confirms the action. The Job Status changes from Pending to Submitted.
 - c. Only the creator of the deployment needs to be granted access directly to the cloud or clouds associated with the deployment environment. This allows you to restrict other users to only deploy to approve deployment groups.
 - **OFF:** (Default) Approval is not required.
6. In the **Cloud Selection** section, use the checkboxes to select the desired VM-based cloud regions and/or container clouds, and for each selected region select the available cloud accounts you want to make available at deploy time from the Cloud Account dropdown list. Click the pin icon in the dropdown list to select a default cloud account as shown in the screenshot below.

Private Clouds



7. Address errors, if identified by Workload Manager, and then click **Next** to go to the Cloud Settings tab.

The **Cloud Settings** tab contains two sections:

- [Simplified Networks](#)
- [Default Tier Cloud Settings](#)

Simplified Networks

Simplified Networks allows you to create multiple network maps, where each network map contains all of the details needed in specifying the Cloud Settings within the Default Tier Cloud Settings (see below). When simplified networks is enabled for the deployment environment, when a user deploys an application to that environment, Page 2 of the Deploy form will display a dropdown of the available network maps instead of showing the detailed cloud settings.

To enable simplified networks in the deployment environment, follow this procedure.

1. Turn on the **Use Simplified Networks** toggle. This causes the simplified networks section to expand and the **New Network Mapping** link to be displayed.
2. Click the **New Network Mapping** link to cause the New Network Mapping form to be displayed.
3. Enter the required network mapping name and optional description.
4. Enter the network settings for at least one permutation of cloud region and cloud account represented by the tabs on the left side of the Network Settings section. A sample Network Settings section for a network mapping involving AWS regions is shown in the screenshot below.

These network settings fields *correspond exactly* to the Cloud Setting fields in the Default Tier Cloud Settings section when the simplified networks toggle is turned off and will vary based on the cloud provider for the region. See Default Tier Cloud Settings > Cloud Settings, below, for details.

5. Save the network mapping when done. This returns you to the deployment environment screen and your network mapping is displayed in the list of defined network mappings as shown in the screenshot below.

6. Repeat steps 2 through 5 to add additional network mappings as needed. After adding your network mappings, you can later delete or edit them.



If you enable simplified networks and do not define at least one network mapping for a region-account combination in the deployment environment, users will not be able to deploy applications using that region-account combination if they select this deployment environment in Page 1 of the Deploy form.

Default Tier Cloud Settings

You must specify the default tier cloud settings for each permutation of cloud region and cloud account represented by the tabs on the left side of this section.

The Default Tier Cloud Settings section contains the following subsections:

- Available instance types, including:
 - Visibility of instance type virtual hardware configuration
 - Visibility of instance type hourly cost
- Resource Placement
- Cloud Settings
- Resource Validation
- SSH Options

Available Instance Types

In the Instance Type subsection, select the instance type(s) that you would like to make available at deploy time.

The initial Instance type subsection display shows the **All** tab highlighted, followed by the Instance Type Filters, and below that, tiles representing instance types that may be made available at deploy time. However, no instance types are initially selected.

Use the Instance Type Filters to limit which instance type tiles are displayed below. You can select all instance types by setting the **Select All** toggle to ON (on the right side of the subsection below the instance type filers). Otherwise, you can select individual instance types by clicking on the corresponding instance type tiles. You must select at least one instance type.

CLEAR ALL SETTINGS

Instance Type

All (7)

Selected (0)

▼ Instance Type Filters

Price (dollars/hour)

MIN: 0

MAX: 0

Virtual CPU

MIN: 1

MAX: 4

Memory (GB)

MIN: 0.5

MAX: 16

Storage (GB)

MIN: 2

MAX: 30

🔍 Search
SELECT ALL

☰ NO

AVAILABLE INSTANCE TYPES (7) At least one instance type need to be selected

	M1.SMALL	M1.LARGE	M2.SMALL
HARDWARE INFO 	1 VIRTUAL CPU	4 VIRTUAL CPU	1 VIRTUAL CPU
	2 GB MEMORY	16 GB MEMORY	4 GB MEMORY
	20 GB ROOT DISK	30 GB ROOT DISK	30 GB ROOT DISK
PRICING INFO 	\$ 0 /hour approx 0/month	\$ 0 /hour approx 0/month	\$ 0 /hour approx 0/month

✔ If you do not see the required instance type listed in this subsection, make sure that this instance type appears in the Instance Type Section of the Regions tab for the cloud region. See [Manage Instance Types](#) for additional context.

When done selecting instance types, clicking the **Selected** tab at the top of the subsection. This causes a preview of the selected instance types to be displayed. This preview is comparable to what the user would see in Page 2 of the Deploy form. The following is a screenshot showing this preview after three instance types were selected.

CLEAR ALL SETTINGS

Instance Type

All (7)

Selected (3)

AVAILABLE INSTANCE TYPES (3)

HARDWARE INFO



PRICING INFO



M1.SMALL

1 VIRTUAL CPU

2 GB MEMORY

20 GB ROOT DISK

\$ 0 /hour

approx 0/month

M1.LARGE

4 VIRTUAL CPU

16 GB MEMORY

30 GB ROOT DISK

\$ 0 /hour

approx 0/month

M1.NANO

1 VIRTUAL CPU

512 MB MEMORY

2 GB ROOT DISK

\$ 0 /hour

approx 0/month

By default, the hardware information (vCPU, memory, root disk) and pricing information (cost per hour) are displayed in the instance type tiles on Page 2 of the Deploy form. Turning the corresponding **Hardware Info** and **Pricing Info** toggles off hides this information from the Deploy form.

Resource Placement (AWS, OpenStack, vCenter Only)

If you are configuring default tier cloud settings for an AWS, OpenStack or vCenter region, you may invoke a resource placement script during deploy time by turning on the **Resource Placement** toggle. This will allow you to specify the source and name of the resource placement script to be executed. The script must allocate all of the resources that are required in the Cloud Settings subsection for that cloud family. The script is executed as part of the deployment of each VM. When you specify a resource placement script for a region-account combination, the Cloud Settings subsection and the Resource Validation script option are disabled and hidden. See [Define Resource Placement](#) for more details.

Cloud Settings

The Cloud Settings subsection lets you specify default values for cloud resources and network settings used by the VMs in deployed in that region

It contains two toggles for controlling editability and visibility of the cloud settings in Page 2 of the deploy form:

- **Visibility:**
 - **ON** (default): Users can see the cloud settings in the corresponding section of Page 2 of the Deploy form when this environment is selected.
 - **OFF:** The cloud settings are not visible in the Deploy form when this environment is selected.
- **Editability:**
 - **Unlocked** (default): Users can change the cloud settings in the corresponding section of Page 2 of the Deploy form when this environment is selected.
 - **Locked:** Users cannot change the cloud settings displayed in the Deploy form when this environment is selected.



The visibility and editability toggles are disabled until you set all required fields in the Cloud Settings subsection.

The format of the rest of the Cloud Settings subsection depends on whether Simplified Networks is enabled for this environment.

If Simplified Networks *is enabled* for this environment, the rest of the Cloud Settings subsection displays a single Network Mapping dropdown field to select one of the network mappings created in the Simplified Networks section as shown in the screenshot below.



Cloud Settings

* NETWORK MAPPING

Select a network mapping

When you select a network mapping from the dropdown field, this network mapping will become the default network mapping for this region-account combination. If you have defined other network mappings, and the the Cloud Settings subsection is set as visible and editable, the user will be able to choose from any of the mappings in the corresponding dropdown field on Page 2 of the Deploy form.

If Simplified Networks *is not enabled* for this environment, all of the remaining fields in the Cloud Settings subsection depend on the cloud provider for the selected region as described below.

- The following are the Cloud Settings fields for an IBM Cloud region. All fields are optional. The Public VLAN and Public Subnet fields are only visible after you select a private VLAN and the Assign Public IP toggle is left ON.

Field	Notes
Private VLAN	Lists private VLANs that may be associated with this region-account combination.
Private Subnet	Lists private subnets that may be associated with this region-account combination.
Assign Public IP	ON (default): A public VLAN and public subnet is associated with VMs deployed using this region-account combination. OFF: Only a private VLAN and subnet is associated with VMs deployed using this region-account combination.
Public VLAN	Lists public VLANs that may be associated with this region-account combination.
Public Subnet	Lists public subnets that may be associated with this region-account combination.

- The following are the Cloud Settings fields for a vCenter region. Only the Datacenter and Cluster fields are required.

Field	Notes						
Enable Full Clone	<p>OFF (default): Directs vCenter to create a thin clone which is faster but relies on the original VM disk being available in its original location. The format is the same as the source template/snapshot disk format (the default when you create a VM from the vCenter UI).</p> <p>ON: Directs vCenter to create a full disk clone of the VM.</p> <p>If you use VM template when configuring images for vCenter cloud environments, be aware of the following considerations.</p> <table border="1"> <thead> <tr> <th>Cloning Method</th> <th>Considerations</th> </tr> </thead> <tbody> <tr> <td>Full Clone</td> <td> <ul style="list-style-type: none"> Use if deploying to a different VMware cluster from the worker image. Use if you select an image that is mapped to a Template. Add this Template to the <i>CliqrTemplates</i> folder The full clone operation is performed on the source VM or VM template, the cloned VM can be on either datastore or datastore cluster that you specify. You can use the Full clone option for both Snapshots and Templates. </td> </tr> <tr> <td>Linked Clone</td> <td> <ul style="list-style-type: none"> Use if the image is mapped to a snapshot. Add a folder in vSphere (to store your CloudCenter snapshots), name it <i>CliqrTemplates</i>, and add this snapshot to the <i>CliqrTemplates</i> folder. At the time of deployment, deploy to the datastore where the snapshot is present. </td> </tr> </tbody> </table> <p>When you use a Snapshot, both the Linked Clone and Full Clone options are possible settings.</p>	Cloning Method	Considerations	Full Clone	<ul style="list-style-type: none"> Use if deploying to a different VMware cluster from the worker image. Use if you select an image that is mapped to a Template. Add this Template to the <i>CliqrTemplates</i> folder The full clone operation is performed on the source VM or VM template, the cloned VM can be on either datastore or datastore cluster that you specify. You can use the Full clone option for both Snapshots and Templates. 	Linked Clone	<ul style="list-style-type: none"> Use if the image is mapped to a snapshot. Add a folder in vSphere (to store your CloudCenter snapshots), name it <i>CliqrTemplates</i>, and add this snapshot to the <i>CliqrTemplates</i> folder. At the time of deployment, deploy to the datastore where the snapshot is present.
Cloning Method	Considerations						
Full Clone	<ul style="list-style-type: none"> Use if deploying to a different VMware cluster from the worker image. Use if you select an image that is mapped to a Template. Add this Template to the <i>CliqrTemplates</i> folder The full clone operation is performed on the source VM or VM template, the cloned VM can be on either datastore or datastore cluster that you specify. You can use the Full clone option for both Snapshots and Templates. 						
Linked Clone	<ul style="list-style-type: none"> Use if the image is mapped to a snapshot. Add a folder in vSphere (to store your CloudCenter snapshots), name it <i>CliqrTemplates</i>, and add this snapshot to the <i>CliqrTemplates</i> folder. At the time of deployment, deploy to the datastore where the snapshot is present. 						
Datacenter	You must select the name of the datacenter where you want your VMs deployed.						
Cluster	After you select the datacenter, you must select a host cluster within that datacenter.						
VM Group	Optionally select one of the VM DRS Groups defined in vCenter.						
Datastore Cluster	Optionally select one of the datastore clusters defined in vCenter. If the Enable Full Clone toggle is turned OFF (linked clones are used), then this field is disabled. If the Enable Full Clone toggle is turned ON, and you select a datastore cluster, the Datastore data entry field appears below.						
Datastore	<p>This field is only displayed after you select a datastore cluster from the field above. If DRS is disabled on the datastore cluster selected in the Datastore Cluster field above, this field displays a list of datastores associated with that datastore cluster and you may select a datastore from the list. If DRS is enabled on the datastore cluster selected in the Datastore Cluster field above, vCenter automatically chooses the datastore and this field is disabled.</p> <p>The CloudCenter Suite supports clustered DS and manages this setting automatically.</p>						
Resource Pool	Optionally select the resource pool you want your VMs to use.						
Target Deployment Folder	Optionally select the folder where you want your VMs deployed.						

Network	You must select a network from the dropdown for NIC 1. You may optionally add more NICs by clicking the Add Network Interface Controller link below the Network dropdown field.
----------------	--

- The following are the Cloud Settings fields for a vCD region. All fields are required.

Field	Notes
vCloud Org VDC	The name of the Virtual Data Center (VDC) in vCloud Director.
vCloud Storage Profiles	The storage profiles to deploy the VMs.
vCloud Org VDC Network	Select a network for NIC 1 and for any additional NICs you add.

- The following are the Cloud Settings fields for an OpenStack region. All fields are required except Availability Zones.

Field	Notes
Availability Zones	You may select multiple default availability zones.
Cloud Tenant	
Network	Select for NIC 1 and for any additional NICs you add.
Private IP Allocation	Select for NIC 1 and for any additional NICs you add.
Assign Public IP	Select for NIC 1 and for any additional NICs you add. Default: ON
Assign IPv6 Address	Select for NIC 1 and for any additional NICs you add. Default: OFF

- The following are the Cloud Settings fields for a GCP region. All fields are required.

Field	Notes
Project	
Zone	
Network	
Subnetwork	
Assign Public IP	Default: ON
IP Forwarding	Default: OFF

- A Kubernetes cloud has only one cloud setting field and it is mandatory: **Namespace**.
- The following are the Cloud Settings fields for an AWS region. All fields are required.

Field	Notes
VPC	
Network	Select a subnet for NIC 1 and for any additional NICs you add. If you select ALL for NIC 1, you will not be able to add additional NICs.
Assign Public IP	Select for NIC 1 and for any additional NICs you add. Default: ON
Source / Destination Check	Select for NIC 1 and for any additional NICs you add. Default: ON

- The following are the Cloud Settings fields for an AzureRM region. All dropdown fields are required except Storage Account and Diagnostics.

Field	Notes
Subscription	
Resource Group	
Storage Account	
Diagnostics	
Enable Availability Set checkbox	Default: unchecked

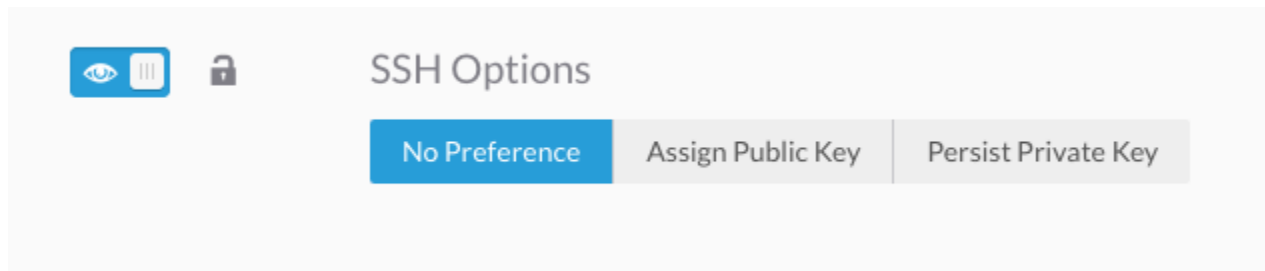
Virtual Network	
Subnet	Select a subnet for NIC 1 and for any additional NICs you add.
Assign Public IP checkbox	Defined per NIC. Default: checked.

Resource Validation

You may invoke a resource validation script during deploy time by turning on the **Resource Validation** toggle and specifying the source and name of the script. A resource validation script is designed to prevent the deployment of an application to an environment if testing reveals that less than a certain amount of one or more resources is available for supporting the deployment. When this condition occurs, the validation script can output a message that specifies which resources were low and by how much. See [Define Resource Validation](#) for more details.


SSH Options

SSH options are on the bottom of the Per Tier Default Settings section and are preceded by a visibility toggle and a lock icon, as shown in the following image.



The visibility toggle is on by default and the lock icon is unlocked by default. This means that the SSH Options subsection will be visible on page 2 of the Deploy form and can be modified at deploy time. If the visibility toggle is on but the lock icon is locked, the pre-selected choice will be visible but cannot be changed at deploy time. If the visibility toggle is off, the SSH Options section is not shown in the Deploy form and the selection made in the Deployment Environment form is automatically applied at deploy time.

Set the SSH Option by selecting one of the three radio buttons as detailed in the following table.

Button	Effect
No Preference (default)	<p>Workload Manager will generate a private SSH key to allow secure communications between the CloudCenter Suite cluster (or Cloud Remote, if deployed) and the worker VMs, but this private key is not stored on the worker VMs.</p> <div style="border: 1px solid #ffc107; padding: 5px;"> <p> If you use a custom key for deployments, you will not see the SSH or RDP connect buttons in your environment.</p> <p>The SSH and RDP options are only visible when the deployments are submitted using the default No Preference option, in which case the Workload Manager uses the default keys to establish a secure connection with the VM instance.</p> </div>
Persist Private Key	The Workload Manager generated private key is stored on all worker VMs in this deployment. Use this option to allowing SSH communication between worker VMs.

Assign Public Key

Workload Manager will use a public key specified by you for Workload Manager-to-VM communications. This key is not stored on the worker VM and therefore cannot be used for secure VM-to-VM communication. When this option is selected, new data entry fields appear under the radio buttons as shown in the following screenshot.

Give the SSH key a name, then either load the key from a file on your PC or copy and paste the key form a text file.

The private or public key is not used to create the key pair on the cloud provider. Instead, it is used by the Workload Manager agent to configure the *cliqruser* and make the VM accessible through the **cliqruser-private key** combination.



The Workload Manager has no way of knowing the private key that is held by the user – Cisco only supports SSH keys that are implicitly injected by the Workload Manager

To configure details in the **Policy Settings** tab, follow this procedure

1. In the **Policy Settings** section, identify the following options for each policy setting, which also has the info icon displaying additional details that you may need for the next step. See [Policy Management](#) for additional details on each policy type.



You cannot configure policies settings in a purely container-based deployment environment – the **Policy Settings** tab is disabled and greyed out for deployment environments configured with only container clouds.

However, you can access and configure the Policy Settings information for hybrid cloud deployment environments that include containers.



New multi-select fields enable admins to restrict what policies and tags can be selected by end-users at deploy time. Admins can also set default values for policies or completely hide policy fields from users.

2. Identify if the following settings apply to the selected policy:
 - a. Should this be **visible** to users in your tenant (toggle switch)?
 - b. Will this be the default policy (pin icon)?
 - c. Do you want to make this policy **Mandatory** (toggle switch)?
 - d. Should the **Terminate Protection** configuration for this policy be visible in the Deploy form (toggle switch)?
 - e. In the **Deploy Time Preview** section, select the required time-based option for the required policies. Refer to the info icon for the required policy in the previous step to select the corresponding option. The **Deploy Time Preview** section is a read-only preview and only displays how the policy selection would appear during the deploy flow with the currently selected options. Users cannot perform any action using this section.
3. Click **Done** to save your new deployment environment.

Define Resource Placement

Define Resource Placement

- [Overview](#)
- [Resource Placement Flow](#)

The Workload Manager has the ability to deploy enterprise applications over public, private, or hybrid clouds by configuring user-specified cloud settings in the Workload Manager UI > **Environments** > **Edit Deployment Environment** > **Define Default Cloud Settings** page.

The *Resource Placement* integration features extend the Workload Manager capabilities by allowing users to define cloud settings based on third-party infrastructure tools or quota management tools using automated scripts instead of manually-selected settings.

You can configure these integrations using an automation callout script.



This page does not list ALL available environment variables! Instead, the commonly used variables are listed in the table and in examples.

To use the resource placement script, specify a URL to the script that you want to run. The resource placement script runs inside a Docker container in the CloudCenter Suite cluster (or Cloud Remote, if installed). You cannot pass in any custom values. Your script must output the values specified in this section.




The Resource Placement feature is only supported for AWS, VMware, and OpenStack clouds.

This script is executed for each Node launch (called for each VM). For example, if you have a single-tier application with the minimum number of nodes set to 2, then this script is executed twice – 1 tier x 2 nodes = 2 executions. However, the Workload Manager passes variables such as service types which allows you to decide where to place it in the VM.

1. Define Resource Callout by enabling the Resource Callout feature on the Workload Manager UI > **Environments** tab > Add a new or edit an existing deployment environment > **Define Default Cloud Settings** (see [Deployment Environments](#) for additional context).
2. Toggle the switch to YES in the Resource Placement section, as shown in the following screenshot.

Resource Placement

ENABLE RESOURCE PLACEMENT



RESOURCE PLACEMENT CONFIG

Demo ▼



If this feature is enabled, the Cloud Settings form in the [Deployment Environments](#) > Cloud Defaults page will be disabled.

3. Identify the script location and the specific script for the Resource Placement Configuration. Based on the cloud account details received through the environment variables, the script can query the target cloud and return the cloud settings to be used for a deployment.
4. The following table describes available (input) environment variables for the Resource Placement script.

Environment Variable	Description	Cloud
env_JOB_ID	The Job ID as an environment variable can be appended to the job name.	All clouds, except Container Clouds and Cisco UCSD.
env_parentJobId	The parent Job ID as an environment variable can be appended to the job name.	
eNV_cliqrAppTierName	The tier name.	
CliqrTier_<tierName>_instanceType	The Instance Type of the tier.	
eNV_imageName	The image Name (for example: CentOS 6.x).	

serviceName	The service name to identify settings like private subnet for a database service.	
eNV_parentJobName	The unique Job Name for the deployment.	
CliqCloudAccountId	The cloud account ID.	
CliqCloudAccountPwd	The cloud account password (for AWS, access key).	
CliqCloudAccountName	The cloud account username (for AWS, account email ID).	
Cloud_Setting_CloudFamily	The cloud family of the region in the Workload Manager.	
CliqCloud_AccessSecretKey	The AWS account secret key	AWS
CliqCloud_ServiceUrl	The SDK URL for VMware.	VMware
CliqCloud_DomainId	The default OpenStack domain ID.	OpenStack
CliqCloud_Endpoint	The OpenStack Keystone authentication endpoint.	
CliqCloud_TenantName	The OpenStack default tenant name.	
CliqCloud_DomainName	The OpenStack default domain name.	
CliqCloudAccountPwd	The OpenStack password.	
CliqCloud_Region	The OpenStack region.	
CliqCloudAccountName	The OpenStack user name for this account.	
CliqCloud_TenantId	The OpenStack default tenant ID.	

5. The cloud settings listed in the following tables are returned for AWS, OpenStack, and VMware.

- The following table describes AWS-specific cloud settings for the resource placement callout script.

AWS Setting	Description
vpcId	The VPC for the node to be deployed.
subnetId	The subnet where the node should be deployed in the above VPC.
securityGroupList	The security groups where the node should be associated in the above VPC.
vmTagsList	The AWS tags to associate with the node.
assignPublicIp	Identifies if the node should be assigned with a public IP.
nodeInfo	Customizable node Information detail that is displayed in the Workload Manager UI Job Details Page for each node. If not provided, the Workload Manager generates the default nodeInfo based on the provided values.

• **Sample Amazon Resource Placement Callout Script**

```
#!/bin/bash

. /utils.sh

content="{\"vpcId\": \"vpc-1234abcd\",
\"subnetId\": \"subnet-1234abcd\", \"securityGroupList\": \"sg-1234abcd\",
\"vmTagsList\": \"Name:MyVm, PayProfile:Dev, BU:Engineering, User:DemoUser\",
\"assignPublicIp\": \"true\", \"nodeInfo\": \"VpcID:
vpc-1234abcd, subnetId: subnet-1234abcd, securityGroupList: sg-1234abcd \"}"

print_ext_service_result "$content"
```


- The following table describes OpenStack-specific cloud settings for the resource placement callout script.

OpenStack Setting	Description
TenantName	The name of OpenStack tenant.
zone	The availability zone as described in Availability Sets and Zones > OpenStack.
nicInfo	The OpenStack network interface information.
privateIPAllocationMode	The private IP allocation strategy (DHCP or PREALLOCATE_IP).
networkID	The OpenStack Network ID.
publicIpAllocate	A flag to allocate the public IP address (Boolean: true/false).
nodeInfo	Customizable node Information detail that is displayed in the Workload Manager UI <i>Job Details</i> Page for each node. If not provided, the Workload Manager generates the default nodeInfo.

- **Sample OpenStack Resource Placement Callout Script**

```
#!/bin/bash
. /utils.sh
content="{\"TenantName\": \"sample\", \"zone\": \"nova\",
\"nicInfo\": \"
[ { \"privateIPAllocationMode\": \"DHCP\", \"networkID\": \"19f56fa2-babc-4q22-9q9b-20c4e3243b85\", \"publicIpAllocate\": true },
{ \"privateIPAllocationMode\": \"PREALLOCATE_IP\", \"networkID\": \"19f56fa2-babc-4q22-9q9b-20c4e3243b85\", \"publicIpAllocate\": true },
{ \"privateIPAllocationMode\": \"PREALLOCATE_IP\", \"networkID\": \"19f56fa2-babc-4q22-9q9b-20c4e3243b85\", \"publicIpAllocate\": false } ]
\",
\"nodeInfo\": \"zone:nova, TenantName:sample\"}"
print_ext_service_result "$content"
```

- The following table describes vCenter-specific cloud settings for the resource placement callout script:

VMware Setting	Description
vmTagsList	The AWS tags to associate with the node.
UserDataCenterName	The datacenter to deploy the node.
UserClusterName	The cluster to deploy the node in the above datacenter.
UserResourcePoolName	The resource pool used to deploy the node.
UserDatastoreCluster	The datastore cluster or datastore associated with the node. Datastore can be specified here, only if the datastore is not part of any datastore cluster.
UserDatastore	Specific datastore within the datastore cluster associated with the node. If this value is populated, it is mandatory to specify the datastore cluster that the datastore belongs to, in UserDatastoreCluster.
UserFolderName	The user folder used for the node deployment.
DisableSDRS	The value set to enable or disable the storage DRS.
RootDiskResizable	Identifies if the root disk is resizable (Boolean: true/false). <div style="border: 1px solid #ccc; padding: 5px; background-color: #fff9c4;">  The Root Disk setting is available as a separate field and you can only select the Root Disk size if the Cloning Mode is set to Full Clone (Enable Full Clone = Selected). See VMware Network Settings for additional context. </div>
FullClone	Identifies if the node to be launched is with full clone (Boolean: true/false).

SystemFolderName	The folder from which the template is selected.						
networkList	The list of networks to attach to the node. The following table describes the format that various network types require. See vCenter Configurations > <i>Resource Placement</i> for additional context. <table border="1" data-bbox="537 270 1252 554"> <thead> <tr> <th>VMware Network</th> <th>Format Required</th> </tr> </thead> <tbody> <tr> <td>DVS and ACI networks</td> <td><i>DistributedPortGroupName (DistributedSwitchName)</i> For Example: DistributedPortGroupName101 (DistributedSwitch2)</td> </tr> <tr> <td>Standard Network</td> <td><i>NetworkName</i> For Example: NetworkName101</td> </tr> </tbody> </table>	VMware Network	Format Required	DVS and ACI networks	<i>DistributedPortGroupName (DistributedSwitchName)</i> For Example: DistributedPortGroupName101 (DistributedSwitch2)	Standard Network	<i>NetworkName</i> For Example: NetworkName101
VMware Network	Format Required						
DVS and ACI networks	<i>DistributedPortGroupName (DistributedSwitchName)</i> For Example: DistributedPortGroupName101 (DistributedSwitch2)						
Standard Network	<i>NetworkName</i> For Example: NetworkName101						
UserHost	The ESX host to which the node is launched.						
nodeInfo	Customizable node Information detail that is displayed in the Workload Manager UI Job Details Page for each node. If not provided, the Workload Manager generates the default nodeInfo based on the provided values.						
DisableStorageDRS	If you enable the Disable storage DRS for this virtual machine flag in the script, the values for the following variables are returned: UserDataCenterName , UserClusterName , and UserDatastoreCluster . Note that the UserDatastoreCluster value must be set to the exact datastore used from a DRS-enabled DS cluster, for example, <i>datastore-151</i> .						

- **Sample VMware Resource Placement Callout Script**

```
#!/bin/bash

. /utils.sh

content="{\"UserDataCenterName\": \"SDC-01\", \"UserClusterName\": \"wm-Cluster\", \"UserResourcePoolName\": \"\", \"vmTagsList\": \"cloud:center\", \"UserDatastoreCluster\": \"wm-DS-Cluster\", \"RootFolderName\": \"vm\", \"UserFolderName\": \"CliqrUser-2\", \"RootDiskResizable\": \"false\", \"FullClone\": \"true\", \"SystemFolderName\": \"CliqrTemplates\", \"networkList\": \"VLAN-ENG-NET (CC-DSwitch)\", \"UserHost\": \"sc2-esx02.abc.private\", \"nodeInfo\": \"UserDataCenterName: SDC-01, UserClusterName: wm-Cluster, UserDatastoreCluster: wm-DS-Cluster, networkList VLAN-ENG-NET (CC-DSwitch)\"}"

print_ext_service_result "$content"
```

- **Sample CWOM/Turbonomic Integration Script for Resource Placement for VMware**

```
#!/bin/bash

. /utils.sh

export CWOM_URL="<<VM_TURBO_HOST>"
export CWOM_USER="<<User>"
export CWOM_PASSWORD="<<pwd>"
export CWOM_RESOURCE="http://$CWOM_USER:$CWOM_PASSWORD@$CWOM_URL"
. /utils.sh

#pre fixing a datacenter in the sample
if [ -z $dcName ];
then
    export dcName="SCL2"
fi

export vmTagsList="Name:myVm"
```



```

export UserDataCenterName="$dcName"
export UserClusterName="CliQr"
export UserResourcePoolName="Eng"
export RootFolderName="vm"
export UserFolderName="CliqrUser-id"
export RootDiskResizable="false"
export FullClone="true"
export SystemFolderName="CliqrTemplates"
export networkList="10-DEV (DSwitch)"
export instanceNameVar=`echo CliqrTier_"$eNV_cliqrAppTierName"_instanceType`
eval instanceName='$'$instanceNameVar

getProfileId() {
    result=`curl -s -X GET $CWOM_RESOURCE/cwom/api/templates
| grep $instanceName`
    export profileId=`echo $result | awk -F uuid=`" '{printf $2}' | awk -F \" '{printf
$1}'`
}
getProfileId

getDatacenterId() {
    DATACENTER=$1
    result=`curl -s -X GET
$CWOM_RESOURCE/cwom/api/markets/Market/entities | grep
\"DataCenter\" | grep $DATACENTER`
    export dcId=`echo $result | awk -F
uuid=`" '{printf $2}' | awk -F \" '{printf $1}'`
}
getDatacenterId "$dcName"

#echo "dcId=$dcId"

getReservation() {
    reservationName="reserve-$RANDOM"
    export reserveId=`curl -s -X POST
$CWOM_RESOURCE/cwom/api/reservations -d
"reservationName=$reservationName&templateName=$profileId&count=1&segmentationUuid[
]=$dcId`
}
getReservation
sleep 3

getHostAndDS() {
    result=`curl -s -X GET
$CWOM_RESOURCE/cwom/api/deployitems/$reserveId`
    export datastore=`echo $result | awk -F datastore=`" '{printf $2}' | awk -F \"
'{printf $1}'`-cluster
    export host=`echo $result | awk -F host=`" '{printf $2}' | awk
-F
\" '{printf $1}'`
}
getHostAndDS
export UserDatastoreCluster="$datastore"
export UserHost="$host"

content="{\"UserDataCenterName\": \"$dcName\", \"UserClusterName\": \"$UserClusterName\", \"
UserResourcePoolName\": \"$UserResourcePoolName\", \"vmTagsList\": \"$vmTagsList\", \"
UserDatastoreCluster\": \"$UserDatastoreCluster\",
\"RootFolderName\": \"$RootFolderName\",
\"UserFolderName\": \"$UserFolderName\", \"RootDiskResizable\": \"$RootDiskResizable\",
\"FullClone\": \"$FullClone\",
\"SystemFolderName\": \"$SystemFolderName\",
\"networkList\": \"$networkList\", \"UserHost\": \"$UserHost\", \"nodeInfo\": \"
UserDataCenterName:
$dcName, UserClusterName: $UserClusterName, UserDatastoreCluster:
$UserDatastoreCluster, networkList: $networkList \"}"

print_ext_service_result "$content"

```

- You can specify media attributes for deployed VMs in vCenter for the following applications:

- Single Root I/O Virtualization (SR-IOV)
- PCI passthrough
- Shared PCI passthrough(vGPU)
- Available USB license dongles
- CPU pinning(CPU affinity)

The following are sample callout scripts.

SR-IOV example

```
#!/bin/bash

. /utils.sh

content="{\"UserDataCenterName\": \"MediaPod\", \"UserClusterName\": \"GPU-Enabled\", \"
UserResourcePoolName\": \"\", \"vmTagsList\": \"\", \"UserDatastoreCluster\": \"MP-NFS-1_Shared\", \"
RootFolderName\": \"\", \"UserFolderName\": \"/CloudCenter/CC4.10.0/Deployments\", \"
RootDiskResizable\": \"false\", \"FullClone\": \"true\", \"VmRelocationEnabled\": \"true\", \"
LocalDataStoreEnabled\": \"false\",
\"SystemFolderName\": \"\", \"networkList\": \"VM Network\", \"UserHost\": \"\",
\"mediaInfo\": [{\"networkId\": \"dv_VL041_CC-Deploy-1 (dvSwitch)\", \"type\": \"SRIOV\", \"
pciAddress\": \"0000:d8:00.0\", \"model\": \"\"}],
\"nodeInfo\": \"UserDataCenterName:MediaPod, UserClusterName: GPU-Enabled,
UserDatastoreCluster:, networkList VM Network\"}"

print_ext_service_result "$content"
```

vGPU and CPU pinning example

```
#!/bin/bash

. /utils.sh

content="{\"UserDataCenterName\": \"MediaPod\", \"UserClusterName\": \"GPU-Enabled\", \"
UserResourcePoolName\": \"\", \"vmTagsList\": \"\", \"UserDatastoreCluster\": \"MP-NFS-1_Shared\", \"
RootFolderName\": \"\", \"UserFolderName\": \"/CloudCenter/CC4.10.0/Deployments\", \"
RootDiskResizable\": \"false\", \"FullClone\": \"true\", \"VmRelocationEnabled\": \"true\", \"
LocalDataStoreEnabled\": \"false\",
\"SystemFolderName\": \"\",
\"networkList\": \"VM Network\", \"UserHost\": \"\", \"CpuAffinitySet\": \"2, 4- 6 \",
\"mediaInfo\": [{\"type\": \"PCI_SHARED\", \"vgpu\": \"\"}],
\"nodeInfo\": \"UserDataCenterName:MediaPod, UserClusterName: GPU-Enabled,
UserDatastoreCluster:, networkList VM Network\"}"

print_ext_service_result "$content"
```

xxx example

```
#!/bin/bash

. /utils.sh

content="{\"UserDataCenterName\": \"MediaPod\", \"UserClusterName\": \"GPU-Enabled\", \"
UserResourcePoolName\": \"\", \"vmTagsList\": \"\", \"UserDatastoreCluster\": \"MP-NFS-1_Shared\", \"
RootFolderName\": \"\", \"UserFolderName\": \"/CloudCenter/CC4.10.0/Deployments\", \"
RootDiskResizable\": \"false\", \"FullClone\": \"true\", \"VmRelocationEnabled\": \"true\", \"
LocalDataStoreEnabled\": \"false\",
\"SystemFolderName\": \"\",
\"networkList\": \"VM Network\", \"UserHost\": \"\", \"CpuAffinitySet\": \"2, 4- 6 \",
\"mediaInfo\": [{\"type\": \"PCI\", \"pciAddress\": \"0000:d8:00.3\", \"model\": \"MT28800
Family [ConnectX-5 Ex Virtual Function)\", \"systemId\": \"5c50c62d-def1-5e0a-8f84-
0025b50201af\"}],
\"nodeInfo\": \"UserDataCenterName:MediaPod, UserClusterName: GPU-Enabled,
UserDatastoreCluster:, networkList VM Network\"}"

print_ext_service_result "$content"
```

USB device and CPU pinning example

```
#!/bin/bash

. /utils.sh

content="{\"UserDataCenterName\": \"MediaPod\", \"UserClusterName\": \"GPU-Enabled\", \"
UserResourcePoolName\": \"\", \"vmTagsList\": \"\", \"UserDatastoreCluster\": \"MP-NFS-1_Shared\", \"
RootFolderName\": \"\", \"UserFolderName\": \"/CloudCenter/CC4.10.0/Deployments\", \"
RootDiskResizable\": \"false\", \"FullClone\": \"true\", \"VmRelocationEnabled\": \"true\", \"
LocalDataStoreEnabled\": \"false\",
\"SystemFolderName\": \"\",
\"networkList\": \"VM Network\", \"UserHost\": \"\", \"CpuAffinitySet\": \"2, 4- 6 \",
\"mediaInfo\": [{\"type\": \"USB\", \"model\": \"path:0/1/1\"}],
\"nodeInfo\": \"UserDataCenterName: MediaPod, UserClusterName: GPU-Enabled,
UserDatastoreCluster:, networkList VM Network\"}"

print_ext_service_result "$content"
```

Define Resource Validation


Define Resource Validation

- [Overview](#)
- [Resource Validation Flow](#)

The CloudCenter platform has ability to deploy enterprise applications over public, private, or hybrid clouds by configuring user-specified cloud settings in the Workload Manager UI > **Environments** > **Edit Deployment Environment** > **Cloud Settings** page.

The *Resource Validation* integration feature extends the Workload Manager platform capabilities by blocking new deployments – if users reach a configured threshold limit when using cloud resources (for example, restricting VMs being launched only if cloud resources consume < 75% of your maximum capacity).


You can configure these integrations using an automation callout script.

 The Resource Validation feature is supported for all clouds supported by the Workload Manager.

The **validation callout script is executed** on a per-deployment basis, with environment variables containing details for all the tier-level hardware requirements along with summed up values for the hardware requirements required for the deployment.

To configure the Resource Validation feature, follow this procedure.

1. Define Resource Callout by enabling the Resource Validation feature on the Workload Manager UI > **Environments** > **Edit Deployment Environment** > **Cloud Settings** page (see [Deployment Environments](#) for additional context).
2. Toggle the switch to YES in the Resource Validation section, as shown in the following screenshot.



3. Identify the script location and the specific script for the Validation Configuration.

- **Sample Resource Validation Callout Script**

```
#!/bin/bash
. /utils.sh

content="{\"validated\": \"false\", \"comment\": \"Not Enough Resources to Launch the nodes\"}"
print_ext_service_result "$content"
```

- The following table describes available environment variables for the Resource Validation script:

Environment Variable	Description	Cloud
CliqCloudAccountId	The cloud account ID.	All supported clouds
CliqCloudAccountPwd	The cloud account password (for AWS, access key).	
CliqCloudAccountName	The cloud account username (for AWS, account email ID).	
CliqTier_NameList	The comma-separated list of all tiers in the application – loop this variable for each tier in the script.	
CliqTier_Total_NumCpus	The total vCpus required to launch the complete App.	
CliqTier_Total_Memory	The total memory required to launch the complete App.	
CliqTier_Total_Local_Storage	The total local storage required to launch the complete App.	
CliqTier_<tierName>_instanceType	The Instance Type for the tier.	

CliqTier_<tierName>_instanceName	The Instance Type name (logical name) for the tier.	
CliqTier_<tierName>_cloudType	The Cloud Type for the tier.	
CliqTier_<tierName>_numOfCPUs	The number of CPU's required for the tier.	
CliqTier_<tierName>_memorySize	The memory required for the tier.	
CliqTier_<tierName>_localStorageSize	The local storage required for the tier.	
CliqTier_<tierName>_minClusterSize	The cluster size of the tier that is launched – Total vCPUs required for a tier would be minClusterSize x numOfCPUs.	
CliqCloud_AccessSecretKey	The AWS account secret key	AWS
CliqCloud_RegionEndpoint	The SDK URL for VMware.	VMware
CliqCloud_DomainId	The default OpenStack domain ID.	OpenStack
CliqCloud_Endpoint	The OpenStack Keystone authentication endpoint.	
CliqCloud_TenantName	The OpenStack default tenant name.	
CliqCloud_DomainName	The OpenStack default domain name.	
CliqCloudAccountPwd	The OpenStack password.	
CliqCloud_Region	The OpenStack region.	
CliqCloudAccountName	The OpenStack user name for this account.	
CliqCloud_TenantId	The OpenStack default tenant ID.	

Cloud-Specific Configurations

Cloud-Specific Configurations

- [AWS Configurations](#)
- [AzureRM Configurations](#)
- [Google Cloud Configurations](#)
- [OpenStack Configurations](#)
- [vCenter Configurations](#)

AWS Configurations

AWS Configurations

- [AWS ID Format](#)
- [Workload Manager AMI Details](#)
- [On-Demand Instance](#)
- [Ephemeral Disks](#)
- [Root Volume Size](#)
- [Instance Profile](#)
- [VPC](#)
- [Workload Manager ELB Representation](#)
- [Availability Zones and Sets](#)

The AWS ID is transparent to Workload Manager. If AWS returns a longer instance ID, Workload Manager accepts this AWS ID as is. While the Java string does not have a length limit the database schema is limited to 255 characters.

If you need to share Workload Manager AMIs, contact [CloudCenter Suite Support](#) with the following information:

- AWS account number
- CloudCenter version
- Contact email
- Customer name
- Customer ID (CID)

With [Multiple Volumes](#) configured when deploying the application on AWS, users have the option to select pricing by using the On-Demand Instance.

When you configure 100 GB of disk space, you may only get 20GB VM. This is because Workload Manager only used the *root disk size* in earlier Workload Manager releases. You can attach one ephemeral disk if you configure a larger size in the instance type (see [Manage Instance Types](#) for additional context).

See [Multiple Volumes](#) and the *Submit Job (v2)* API for additional context.

An optional **Instance Profile** field is available when you Setup Deployment Environments or set the Deployment Environment Defaults. If you configure this field, provide the Amazon Resource Name (ARN) used for the *Instance Profile* configured in your AWS Cloud account.

Default Tier Cloud Settings

The screenshot displays the 'Instance Type' configuration page. On the left, a sidebar shows the AWS account details: 'aws Amazon_API' and 'US West (N. California) Amazon_API Account'. The main content area is titled 'Instance Type' and includes a 'CLEAR ALL SETTINGS' link. Below the title, there are two tabs: 'All (224)' and 'Selected (198)'. A section titled 'AVAILABLE INSTANCE TYPES (198)' shows three instance type cards:

Instance Type	Virtual CPU	Memory	Temp Storage	Hourly Price	Monthly Price (approx)
T3.NANO	2	512 MB	0 GB	\$ 0.0062 /hour	approx 4.53/month
T2.NANO	1	512 MB	0 GB	\$ 0.0069 /hour	approx 5.04/month
T3.MICRO	2	1 GB	0 GB	\$ 0.0124 /hour	approx 9.05/month

Below the instance types, there are sections for 'Resource Placement' (with an 'ENABLE RESOURCE PLACEMENT' toggle set to 'NO') and 'Cloud Settings' (with a dropdown menu for VPC set to 'vpc-47230120 | CIDR 172.31.0.0/16').

If you specify the **Instance Profile** name, Workload Manager launches VMs within the IAM role that is associated with the corresponding instance profile.

To successfully launch the AWS cloud account (either using as IAM role or the account secret key) you must have the required permission to pass the IAM role associated with the specified instance profile.

If the application VMs run in isolated networks (like Amazon's VPC), be sure to set up proper NAT rule (only outgoing needed) to allow application VMs to connect to RabbitMQ.

The CCM instance that interacts with the CloudHSM server must reside inside the same VPC as the CCM.

Refer to <https://aws.amazon.com/articles/0639686206802544> for additional context.

AWS allows either *internal* or *internet-facing* ELBs and they are associated with subnets that the instances will be on. Workload Manager uses this information by allowing you to select *internal* or *external* within each ELB tier of the Workload Manager application profile. From there, the subnet for the ELB is determined by where the application tier instances are instantiated.

Refer to the [Amazon Documentation](#) for additional context.

- **UI:** See the [Availability Sets and Zones](#).
- **API:**

nics

- **Description:** Details about the AWS Network Interface Cards (NICs) configuration. See *IP Allocation Mode* for additional details. The concept of *Availability Sets and Zones* in AWS is mapped to the subnet as you can have multiple subnets for a each zone. So you must input the list of subnets as the input for an availability set. During an API job deployment, the availability set input is provided as part of the NIC information. To be more specific the **first NIC** will contain the information as the comma separated subnet list as shown in the example.
- **Type:** Object

networkId	
<ul style="list-style-type: none"> • Description: The network identifier for each required tier. • Type: String 	
Required (if configured in your application profile)	
privateIPAllocationMode	
<ul style="list-style-type: none"> • Description: Identifies the allocation strategy used to configure the NIC for an AWS cloud • Type: Enumeration 	
Enumeration	Description
DHCP (default)	This strategy allows the IP to be allocated by the DHCP server to the instance on server boot up. This IP address is not known prior to server boot up.
Pre-allocate IP	This strategy allows the cloud infrastructure IP allocation to be dynamically provided before the server boots up. This strategy is specific to the following OpenStack applications: <ul style="list-style-type: none"> • CISCO CSR1000: Configuration drive file IP populated with the pre-allocated IPs known before server boot up. • CISCO F5 Load Balancer: Multiple NIC support.
Static IP (only CloudCenter 3.x)	This strategy allows the customer to provide the IP address. As this IP address may or may not be available to the server (based on the availability), you must perform adequate checks to ensure IP availability before using this strategy.
order	
<ul style="list-style-type: none"> • Description: The number at which a resource is to be attached. When updating a phase, use this order to re-order the resource to a different position in the array of resources. • Type: Long 	
Required (if configured in your application profile)	
nicNetworkType	
<ul style="list-style-type: none"> • Description: The type of network for this NIC. A corresponding list of domains are attached to each option. • Type: Enumeration 	
Enumeration	Description
NETWORK	A private network that supports IP ranges which overlap with another private network
BRIDGE_DOMAIN	A set of logical ports that share the same flooding or broadcast characteristics. Used for ACI environments.

- **Example 1:** Using DHCP allocation mode


```
"nics": [
  {
    "order": 1,
    "allocationMode": "DHCP",
    "allocatePublicIp": "true",
    "id": "subnet-bf7c40cb,subnet-dc3c45f4"
  }
]
```

- **Example2:** Using Static IP allocation mode

- ✔ Attach the static private address to the NIC when you [Create Cloud Instance Type](#) instance create then use that address instead of using the DHCP [allocationMode](#).

```
nics":[
  {
    "order":1,
    "allocationMode":"STATIC_IP",
    "allocatePublicIp":"false",
    "id":"subnet-7dc30d25",
    "properties":
      [
        [
          {
            "key":"PRIVATE_IP_LIST",
            "values":[
              "12.3.6.1"
            ]
          }
        ]
      ]
  }
]
```

- **Example 3:** Using IPv6 Address

- ✔ When allocating firewall rules, CloudCenter supports IPv6, in addition to IPv4, addresses in the source for app profile, tenant, and security profiles. When you assign IPv6 addresses, the CloudCenter platform validates the security rule source before accepting the IPv6 address. This support is restricted to AWS and OpenStack clouds. If you provide an invalid IPv4/6 IP address, then the CloudCenter platform rejects the deployment as invalid. See [IP Allocation Mode](#) for additional details.

```
nics":[{
  "order":1,
  "allocationMode":"DHCP",
  "allocatePublicIp":"false",
  "assignIpv6Address": "true",
  "id":"subnet-ab4afff0",
  }
]
```

- **Example 4:** Using Multiple NICs

- ✔ AWS support multiple NICs across subnets in the same availability zone

```
[
  {
    "order":1,
    "allocationMode":"DHCP",
    "allocatePublicIp":"true",
    "id":"subnet-74752b32"
  },
  {
    "order":2,
    "allocationMode":"STATIC_IP",
    "allocatePublicIp":"false",
    "id":"subnet-7dc30d25",
    "properties":[
      {
        "key":"PRIVATE_IP_LIST",
        "values":[
          "12.3.6.1"
        ]
      }
    ]
  }
]
```

AzureRM Configurations

AzureRM Configurations

- [Overview](#)
- [AzureRM Access Endpoints to VMs](#)
- [AzureRM Resource Preparation](#)
- [AzureRM Diagnostics](#)
- [Availability Zones and Sets](#)

Workload Manager supports Azure Resource Manager (AzureRM). When you [Configure a Cloud End-to-End](#), use the *cloudtype* parameter to specify your choice.


Workload Manager allows you to access/restrict AzureRM deployments via the Internet. Every tier has a value to Add Access End Point:

- **Yes:** The SSH port and any other port that is opened as part of the application profile has Internet access.
- **No:** All ports are restricted and will not have Internet access.


Be sure to configure the following resources in preparation to launch a job:

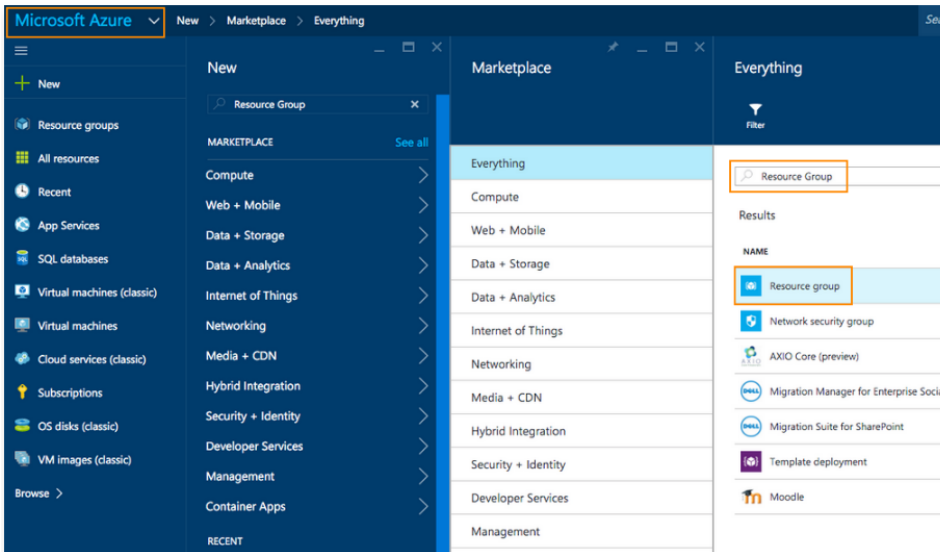
- Resource Group
- Storage Account
- Virtual Network
- Additional Resource Groups

To configure the AzureRM resources in preparation to launch a job, follow this procedure.


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


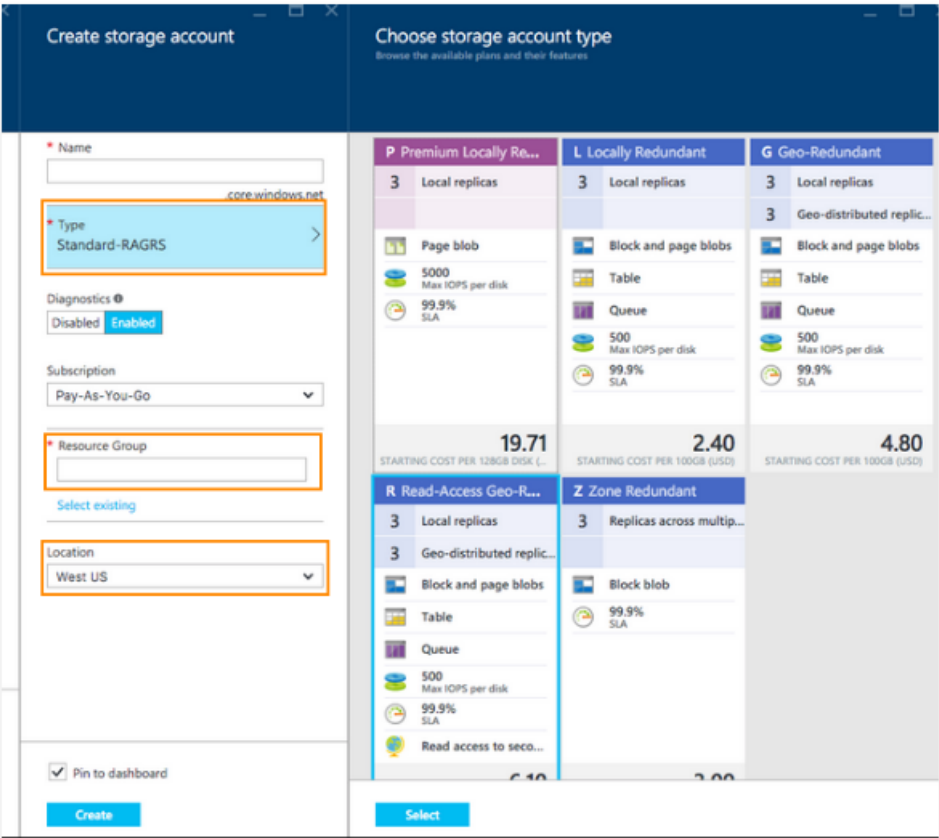
1. Using a valid [Windows Azure Resource Manager account](#), access the new Microsoft Azure Portal.

 You may need to click Microsoft Azure dropdown at the top left corner to toggle between these interfaces.



In the new Microsoft Azure Portal, search for each of the following items and add the corresponding information to each portal.


Search for	Click ...	Configure ...
<i>Resource group</i>	Resource group to access this portal	Select the same Cloud region that you will be adding as the Resource Group location. <div style="border: 1px solid gray; padding: 5px; margin-top: 10px;">  Additional resource groups You can create additional resource groups to place different VM types. At deployment time, you can select the resource group to deploy the VM. </div>

<p><i>Storage account</i></p>	<p>Storage Account to access this portal</p>	<p>Configure the following fields:</p> <ol style="list-style-type: none"> 1. Resource group – select the group created in the first row. 2. Location – provide your cloud region location. 3. Type – create two storage accounts: <div data-bbox="448 275 1484 432" style="border: 1px solid green; padding: 10px; margin: 10px 0;"> <p> Recommendation</p> <p>The reason to create two storage accounts is some instance types (for example, Standard_DS1, Standard_GS1) can use the premium storage account to enhance performance and use standard storage account. The other instance type can use the standard storage account only.</p> </div> <ol style="list-style-type: none"> a. One with Premium Locally Re... (if the premium is not available for some regions, one account is enough) b. One with Standard-RAGS (the default). 
<p><i>Virtual network</i></p>	<p>Virtual Network to access this portal</p>	<p>Configure the following fields:</p> <ol style="list-style-type: none"> 1. Resource group – select the group created in the first row. 2. Location – provide your cloud region location.

Workload Manager users can view diagnostics provided by Azure Resource Manager from multiple places in the Azure console.

1. Access the Azure console from the target VM and open the **Monitoring** page.
 - If you see a graph displayed on the Monitoring page, then diagnostics are enabled.

- If you do not see any data displayed on the Monitoring page, then diagnostics are disabled.

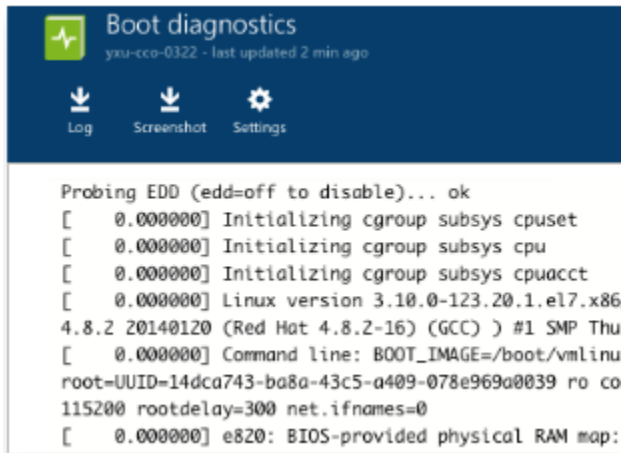


2. Go to the resource, click the **Settings** command, and select **Diagnostics**. Verify the following settings:

- The status is **On**.
- The Basic Metric and Boot Diagnostics are **checked**.




3. Navigate to the Boot diagnostics tab (at the same menu level as Diagnostics, but on the top). Verify that boot diagnostics are available.



4. Enable AzureRM diagnostics in Workload Manager as this feature is disabled by default.

- Deploy the application that uses the AzureRM diagnostics feature.
- Scroll down to the Advanced section for this deployment version.
- Click the Diagnostics dropdown list and select the applicable file for this deployment.

✓ The metrics and logs are stored in the related storage account.
Be aware that this change in settings may incur an extra billing amount.



See [Availability Sets and Zones](#).

Google Cloud Configurations

Google Cloud Configurations

See [Add a Google Cloud Account](#).

- **UI:** See the [Availability Sets and Zones](#).
- **API:**

LaunchZone

- **Description:** While launching a job from the CloudCenter platform, you have the option to create a zone at the tier level. All VMs within a cluster are placed in the same subnet. So all VMs inside the same zone are placed in the same subnet.
- **Type:** Name-Value Pair
 - Name = **LaunchZone**
 - Value = us-central1-b
- **Example:**

```
"cloudProperties": [  
  {  
    "name": "LaunchZone",  
    "value": "us-central1-b"  
  }  
]
```

OpenStack Configurations

OpenStack Configurations


- [Overview](#)
- [OpenStack Settings](#)
- [Workload Manager Tenant Name or ID Deployment](#)
- [Key Pair Settings](#)
- [Submitting Jobs](#)
- The OpenStack Cloud Tenant selector (visible for supported cloud families) displays the list of tenants obtained via the OpenStack identity management service (not tied to any tier).
- [Create New Boot Volume](#)

Multiple OpenStack tenants share cloud accounts in Workload Manager. At deployment time, Workload Manager allows you to select the required OpenStack tenant. You can create access key pairs in the OpenStack console so these key pairs are visible when submitting jobs using Workload Manager.

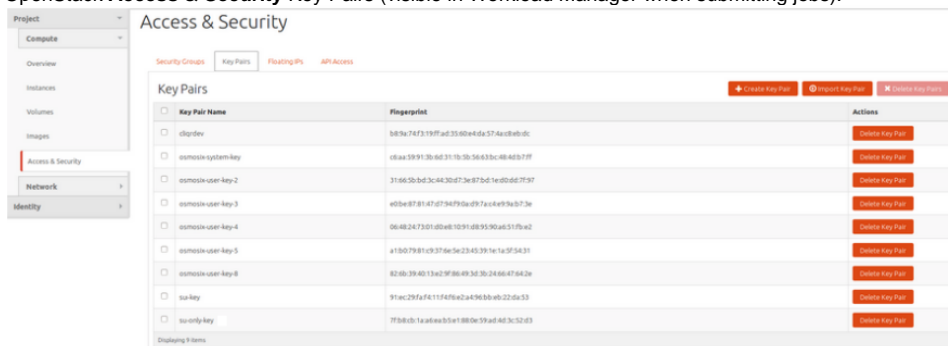
In this case, the concept of **Tenant Name** and **Tenant ID** is specific to the OpenStack cloud, *not Workload Manager*. These two fields specify the OpenStack project information (*TenantName* and *TenantId* in the APIs).

This feature uses the following OpenStack Settings:

- OpenStack Admin sharing multiple **Projects** (tenants) for deployment. In the following image, *ven* is the cloud account user and the *ven* (project) is paired up with the *su* (project) by the OpenStack admin. Consequently, *ven* can deploy applications to both projects as can *su*.



- OpenStack **Access & Security** Key Pairs (visible in Workload Manager when submitting jobs):



When configuring cloud accounts in the Workload Manager UI, you can provide the OpenStack project name as the Tenant Name. Alternately, you can provide the Tenant ID as well. If set, the Default Tenant ID (OpenStack setting in Workload Manager) has precedence over Default Tenant Name (OpenStack setting in Workload Manager). Both the Default Tenant ID and Default Tenant Name fields are optional in the Workload Manager UI. See [Add an OpenStack Cloud Account](#).

Add Cloud Account

Cloud Credentials

OpenStack User Name *

User Name associated with your OpenStack account

OpenStack Account Password ***Default Domain Name (V3)****Default Domain Id (V3)**

Either Default Domain Id or Default Domain Name is needed for V3 API

Default Tenant Name (V3 Project Name)**Default Tenant Id (V3 Project Id)**

Either Default Tenant Id or Default Tenant Name is needed

Workload Manager users can select keys defined in OpenStack when launching deployments. The SSH key pairs are displayed in the Workload Manager UI based on properties that Workload Manager retrieves from the cloud.

Advanced

All fields are **required** unless otherwise specified.

Attach Network Interfaces
Select cloud network(s) and IP allocation strategies to attach network interface(s)

Network Interface Controller	Network	Assign Public IP	Private IP Allocation
NIC 1	rak-net	<input checked="" type="checkbox"/>	DHCP

SSH Options

Assign SSH Key
Check to assign own public/private key for this tier to access the instance.

Cloud provider key User defined custom key

dd
dd
cliqrdev
test
osmosix-system-key
garv

Storage -----No Storage-----
Please select a running storage app.

apache2_0

To view a list of keys, [View Cloud Properties API](#).

The OpenStack CloudTenant is a universal property for the selected cloud/deployment environment (not attached to any tier).

Web application

General Information

We need to know a little about the application you would like to run. This will help us figure out the best solution for your needs.

All fields are **required** unless otherwise specified.

Deployment Name Deployment Name

Description (optional) What does this app do?

App Version 2

Associate Tags Enter tag name

Deployment Environment Venkat

Cloud OpenStack OpenStack-Ver
Cloud Account: venkat@12

Cloud Tenant su
Please select a cloud tenant to deploy application.

See [Multiple Volumes](#) > *OpenStack Type Nuances* and *UI Configuration Options* for additional details

Availability Zones and Sets

- **UI:** See the [Availability Sets and Zones](#).
- **API:**

availabilitySet

- **Description:** While launching a job from the CloudCenter platform, you have the option to create an Availability Set at the tier level. All VMs within a cluster are placed in the same subnet. So all VMs inside the same Availability Set are placed in the same subnet. If you do not configure the Availability Set, an availability set is NOT created. See [Availability Sets and Zones](#) for additional context.
- **Type:** Name-Value Pair
 - Name = **availabilitySet**
 - Value = nova, zone1
- **Example:**

```
"cloudProperties": [  
  {  
    "name": "availabilitySet",  
    "value": "nova,zone1"  
  }]  
]
```

vCenter Configurations

vCenter Configurations

- [Dynamic Cloud Resource Enumeration for VMware](#)
- [VMware Folders](#)
- [Resource Placement](#)
- [Configure Cloud and Network Settings](#)

Workload Manager dynamically pulls cloud resources during the VMware cloud configuration in the vCenter cloud Details tab. As a result, fields like Datacenter name, Resource Pool, and Cluster information are automatically populated during the cloud configuration phase. See [VMware Network Settings](#) for additional context.

Workload Manager enables a user to specify the Target Deployment Folder when provisioning a VM.

The screenshot shows the 'Advanced' configuration tab for a VMware cloud. It includes several dropdown menus for configuration: Data Center (Home), Cluster (cluster1), Datastore (DatastoreCluster), and Resource Pool (CliQr). The 'Target Deployment Folder' dropdown is highlighted with an orange border and contains the value '/test1'. Below these are sections for 'Attach Network Interfaces' and 'SSH Options'.

Network Interface Controller	Network
NIC 1	VM Network

SSH Options:

- Assign SSH Key
Check to assign own public/private key for this tier to access the instance
- Persist the Private Key
Persist the private key in the instance ssh config

This field allows you to define the default folder from which to deploy the VM. The following macros are supported when configuring the folder name:

- %USER_ID%
- %USER_NAME%
- %VENDOR_ID%

A sample configuration example can be `/CliQr/dev/%VENDOR_ID%/%USER_ID%`. If the folder does not exist, it will be created.

Workload Manager supports Distributed Virtual Switch (DVS) networks, ACI networks, and Standard network.

When defining cloud settings using [Define Resource Placement](#), Workload Manager requires you to specify the parsing value for the network by providing the switch and port information in the following format for the DVS and ACI networks in the resource placement script.

```
<distributed port group name> (<distributed switch name>)

# Example for DVS and ACI networks
DistributedPortGroupName1 (DistributedSwitch1)

# Example for standard networks
NetworkName1
```

See the vCenter details in [Create a Deployment Environment](#) for additional context.

Availability Sets and Zones

Availability Sets and Zones

- [Terminology](#)
- [Clustered VMs in an Application Tier](#)
- [Configuration Details](#)
- *Availability Zone* is a common cloud concept that refers to the logical grouping of resources. Each cloud can have one or more availability zones.
- *Availability Set* is a Workload Manager-specific concept to denote a group or list of availability zones.



Cloud Nuances

Be aware of the nuances that the following table describes when configuring Availability Sets using the Workload Manager UI or the Workload Manager REST API calls:

Supported Cloud	Cloud Reference	Workload Manager Concept
AzureRM	Concept of Availability Set (does not include Availability Zones)	Workload Manager provides this input as a boolean value to enable/disable Azure's Availability Set concept.
AWS	Concept of Availability Zone (includes Availability Sets)	Workload Manager provides this input via the VPC NIC(s) field.
OpenStack	Concept of Availability Zone (includes Availability Sets)	Workload Manager provides this input via the Availability Zone(s) field.
Google	Concept of Availability Zones (includes Launch Zones)	Workload Manager provides this input via the Launch Zone field.

This feature is only supported for these clouds.

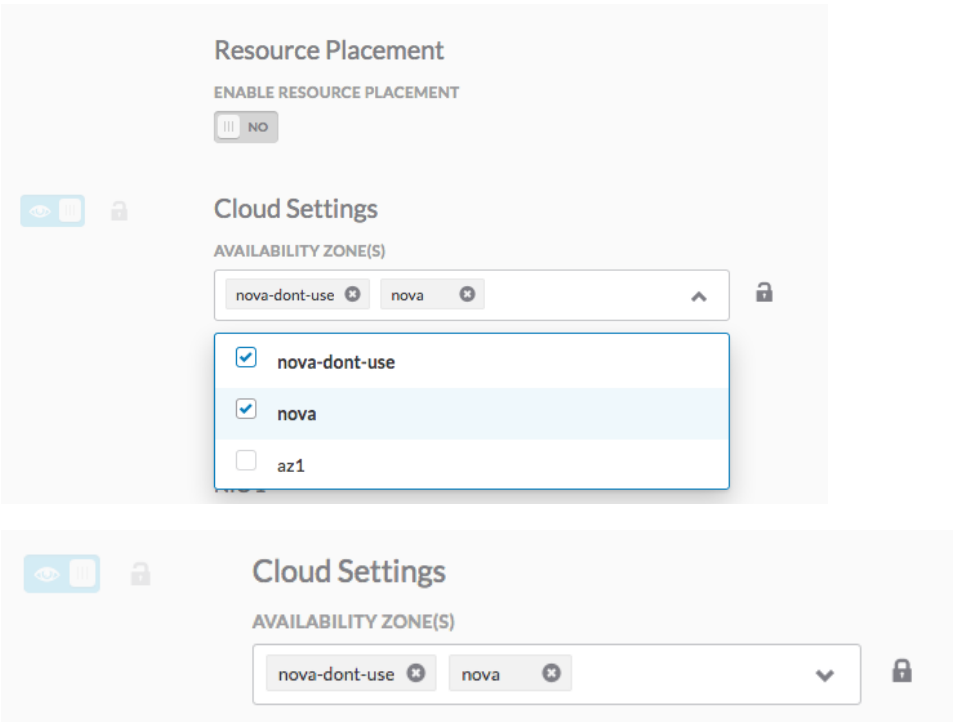
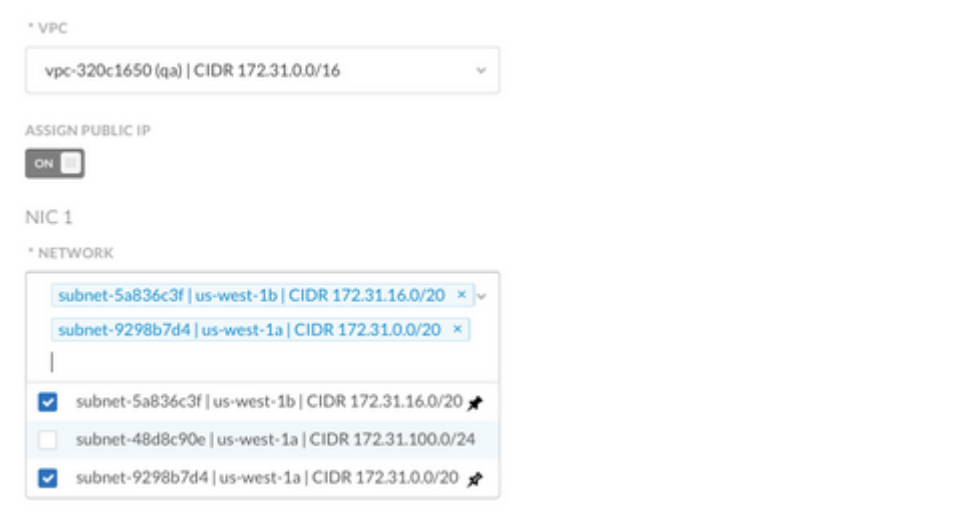

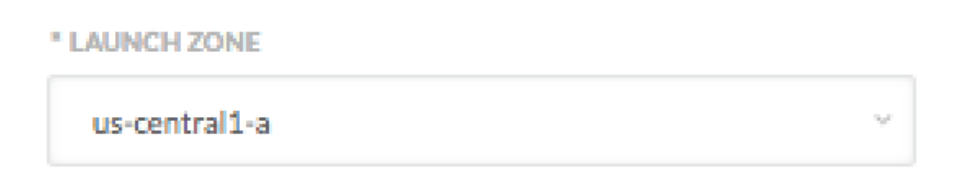
When an application tier is in a clustered environment (with more than one VM), then all the VMs in the cluster may need to be launched into the same availability zone. An availability set refers to a group or list of availability zones and ensures that the cluster VMs are evenly distributed across the specified zones using a Round-Robin implementation. Consequently, even during scaling operations (up or down), the VMs continue to be scaled across the zones. This feature is specific to each tier within an application.

For example, if a tier is a cluster of 3 VMs (being the minimum) and you have the availability set as zone1 and zone2, then when the job is launched the first node is launched into zone1, the second node into zone2, and the third node again into the zone1. After a successful job deployment, if you scale up the tier, then a new node is launched to zone2.

This section provides the related images if configuring using the UI. See the [Create a Deployment Environment](#) section to deploy to and set defaults for each supported cloud.

See the following table for cloud-specific configuration details.

Cloud	UI Image	<i>Additional References</i>
-------	----------	------------------------------

<p>OpenStack</p>		<p>See OpenStack Configurations.</p> <p>The lock icon (only available for OpenStack environments) next to the Availability Zone field is only enabled when Availability Zones are selected using the multi-select component. If you lock the Availability Zone field after selecting some of all of the Availability Zones in the Environment configuration, then you can only select within this selection (any/all /none) of Availability Zones in the Deployment page. For example, the image displayed on the left displays two Availability Zones as being selected. If you lock the selection with just these two, then the Deployments page will only display these two zones as being available – at that point, you can choose to select one or both or none of these Availability Zones.</p>
<p>AWS</p>		<p>See AWS Configurations.</p>
<p>AzureRM</p>		<p>See AzureRM Configurations.</p>
<p>Google</p>		<p>See Google Cloud Configurations.</p>

Multi-Site, Multi-Account Deployment Considerations

Multi-Site, Multi-Account Deployment Considerations

Be aware that multi-site, multi-account deployments are NOT supported in the following scenarios:

1. When you migrate or promote deployments. See [Deployment, VM, and Container States](#) for additional context.
2. For single-tier application (for example, interactive applications).
3. When you use the [Benchmark an Application](#) feature.

You can configure a multi-site deployment to enable users to deploy N-Tier applications with each tier being configure in a different cloud or in the same cloud in different segmented networks while ensuring SLA guaranty and data sovereignty. This feature allows you to use different clouds for different tiers.



While the multi-site feature is supported for all supported clouds (see [What Is Supported?](#) for a complete list), the Set Defaults functionality is only supported for AWS, OpenStack, and vCenter clouds.

For example, in a 2-Tier application, the load balancer and app cluster can be in AWS or any other public cloud offering and the database can be in the private datacenter like VMware with/without [ACI Extensions](#). This example is also applicable in cases where the two different datacenters (regions) of an enterprise is maintained as two different cloud families and the user wants the different tiers in the application to be deployed to these two datacenters or clouds.

If a deployment environment has more than one cloud selected in the Deployment page, you see a new option in the Cloud dropdown list called **Hybrid**. When you select this option, you can choose different instance types and provide Advanced cloud/network settings specific to the selected cloud for *each* tier. The following screenshot shows an example.

Deploy "DocDemo"

General Settings

* DEPLOYMENT NAME

DocDemo

* APPLICATION VERSION

3.7

* TAGS ⓘ

Enter tag name

Default_Dep_Env

Amazon_Dep_Env

Bug_Verify

AGING POLICY

Tier Settings

WebServer

TAGS ⓘ

Enter tag name

* NUMBER OF NODES

1 max: 2

SCALING POLICY

STORAGE SYSTEM

None available

Deploy "DocDemo"
✕

General Settings

* DEPLOYMENT NAME

* APPLICATION VERSION

* TAGS

TERMINATE PROTECTION

AGING POLICY

METADATA

NAME	VALUE	ACTIONS
Name	%JOB_NAME%	

[+ ADD METADATA](#)

Tier Settings

tomcat7_1

TAGS

Default_Dep_Env
 Amazon_Dep_Env
 Bug_Verify

SCALING POLICY


STORAGE SYSTEM
None available

Deploy "DocDemo"

General Cloud Settings


* DEPLOYMENT ENVIRONMENT
Amazon

Select a Cloud



Amazon

US West (Oregon) ▾




Hybrid

Deploy individual Tiers on different clouds

Tier Settings

tomcat7_1 \$ 0.013 /hour

Select a Cloud



Amazon

US East (Virginia) ▾

* CLOUD ACCOUNT
Fayas-IAM-CCO ▾

Filter Instance Types / SHOW

available instance types (49)

T2.MICRO	T2.SMALL	M1.SMALL	T2.MEDIUM
1 VIRTUAL CPU	1 VIRTUAL CPU	1 VIRTUAL CPU	1 VIRTUAL CPU
1 GB MEMORY	2 GB MEMORY	2 GB MEMORY	4 GB MEMORY
0 GB STORAGE	0 GB STORAGE	160 GB STORAGE	0 GB STORAGE
\$ 0.013 /hour approx: \$9.49/month	\$ 0.026 /hour approx: \$18.98/month	\$ 0.044 /hour approx: \$32.12/month	\$ 0.052 /hour approx: \$37.96/month

* VPC
vpc-b0e989d5 (vpc-qa-cases) | CIDR 10.0.0.0/16 ▾

ASSIGN PUBLIC IP
 OFF

NIC 1

* NETWORK
subnet-c97ad590 | us-east-1d | CIDR 10.0.1.0/24 ○ ▾

PRIVATE IP ALLOCATION

You can also configure each application tier to be deployed in different cloud accounts (multi-account) within the same datacenter. The **Hybrid** option allows you to choose the *same cloud* and *different cloud accounts* for each tier.

For example, a datacenter admin maintains a single cloud (for example, VMware) for the entire datacenter but maintains different cloud accounts for different segments that are managed by different Cisco Application Policy Infrastructure Controllers (APICs). In this case, each tier for this application can be deployed to these different segments. The database can be deployed to the pod or segment that has stricter security policies enforced by APIC1 and the AppCluster tier can go into different pod that is managed by a different APIC2.

Application Profiles

- [Application Profile Overview](#)
- [Application Profiles Page](#)
- [Application Tasks](#)
- [Model Applications](#)
- [Deploy Applications](#)
- [Benchmark an Application](#)

Application Profile Overview

Application Profile Overview

An application profile is a cloud-agnostic data structure used by Workload Manager for storing the definition of your multi-tier application and all of its dependencies. The [Application Profiles page](#) is the starting point where you can create a new application profile and where any application profiles created by you or shared with you are listed.

You can [create a new application profile](#) using the Model Application Profile form. This form lets the you specify dependencies between tiers, the service associated with each tier, and the service's associated packages, files, parameters, settings, and lifecycle actions or lifecycle workflows.

Workload Manager does not come with any OOB application profiles, but it does include [OOB Application Templates](#) which serve as a starting point for building your application profiles.

Once you create an application profile you can:

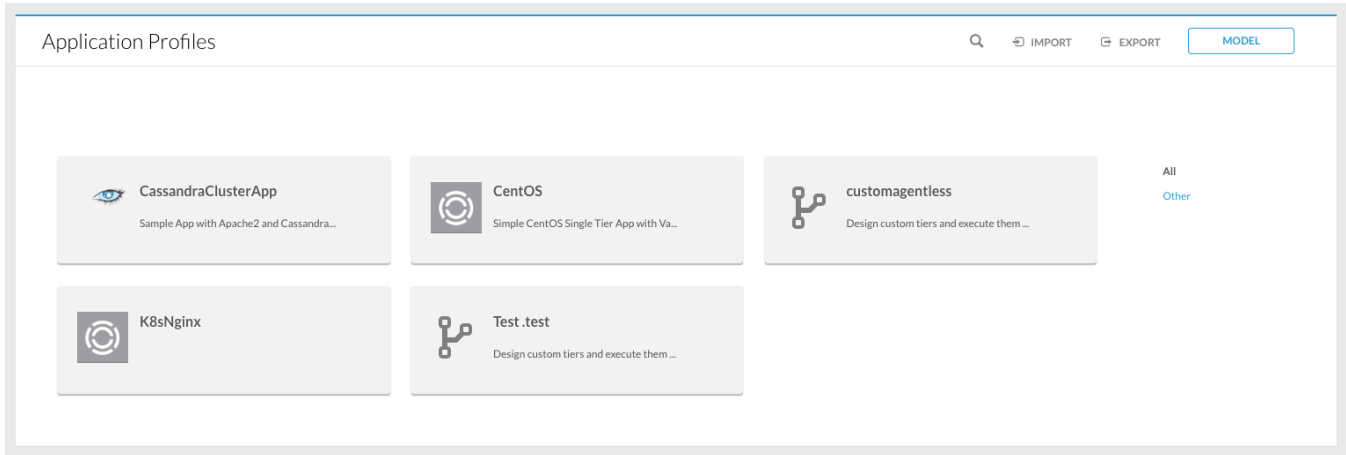
- [deploy the associated application](#) to a deployment environment,
- edit or delete it,
- share it with users, groups, or subtenants,
- export it in a format that can be imported to another instance of Workload Manager, or
- [benchmark](#) the associated application.

Once an application is deployed, you can manage it from the Deployments page.

Application Profiles Page

Application Profiles Page

The Application Profiles page is where you can create a new application profile and where any application profiles created by you or shared with you are listed. A sample Application Profiles page with some application profiles is shown below.




The right side of the page header contains three icons and one button:

- **Search.** Clicking on this icon replaces the icon with a text entry space for the search string. Entering text in this data entry field cause the list of application profiles in the body of the page to be limited to those that contain the entered search string somewhere in the application profile name or description.
- **Import.** Clicking on this icon opens up a file selection window of your desktop OS. You can then search for a compatible zip file containing the exported version of the application profile you want to import.
- **Export.** Clicking on this icon causes a dialog box to be displayed where you can select some or all of the application profiles you have access to. After selecting the profiles you want to export, click **Export**. This causes a zip file for each of the application profiles to be saved in the directory of your choice.
- **Model.** Clicking on this button begins the process of creating a new application profile. See [Create an Application Profile](#) for details.

The main body for the Application Profile page displays all application profiles created by you and shared with you. Clicking on an application profile card displays the Deploy form which lets you deploy the application to a deployment environment. See [Deploy an Application](#) for details.

When you hover over an application profile card, a drop-down menu icon appears in the lower right of the card. Click the drop-down menu icon to display the application profile action menu as shown in the figure below.



The actions listed in the dropdown depend on your permissions relative to the application profile. The table below lists all actions available for an application profile and the required user permissions. If you created the application profile, you implicitly have manage and deploy permission for it and therefore can perform all actions on it. If the application profile is created by another user and shared with you, the actions you can perform are based on the permissions you were granted when that profile was shared with you.

Action	Description	Minimum Permission
Deploy	Causes the Deploy form to be displayed where you can specify the deployment environment and enter parameters for controlling the deployment. See Deploy an Application for details.	view and deploy
Edit /Update	Causes the Edit Application Profile form to be displayed. This form is essentially the same as the Model Application Profile form. See Create an Application Profile for details.	modify
Clone	Creates a copy of the application profile and opens the Clone Application Profile form which is essentially the same as the Edit Application Profile form. The clone operation updates the name of the application profile to the original profile name prepended with "Copy of". You can edit the profile further before saving it.	view
Benchmark	Launches a benchmark job as described in Benchmark an Application .	view and deploy
Export	Lets you export the application profile as a zip file which can later be imported to another instance of Workload Manager.	view
Share	Lets you share the application profile with another user, group, or subtenant with view, modify, or manage permission	manage
Show Deployments	Shows all of your deployments associated with the application profile	view
Delete	Deletes the application profile. You are prompted with a confirmation as this action is not reversible.	manage
Download JSON	Lets you download a JSON file which defines the properties of the application profile.	view
More Info	Lets you view the owner's tenant name, the application profile category, and any tags associated with the profile.	view

Application Tasks

Application Tasks

- [Overview](#)
- [The Application Profiles Page](#)
- [Application Tasks](#)
 - [Deploy Applications](#)
 - [Edit/Update](#)
 - [Clone](#)
 - [Schedule Deployment](#)
 - [Benchmark](#)
 - [Schedule Benchmark](#)
 - [Export](#)
 - [Share](#)
 - [Show Deployments](#)
 - [Delete](#)
 - [More Info](#)

Application management refers to the runtime management activities of applications on the cloud from Workload Manager. This section provides details on where you can manage these application runtime activities.

When you save an application, it becomes available in the Workload Manager > **Apps Profile** list page. The applications displayed in this tab and the tasks you can perform for each application differ based on your [Permission Control](#) level.

Note

When you first start using the %ui, you may not find any applications displayed in the Applications page. Your administrator must first add applications to this page and then permit you to view these applications (see [Permission Control](#)).

Any permitted user can perform the tasks described in the following sections when modeling an application listed in the Applications page.

Deploy Applications

The Deploy task allows you to launch an application at a scheduled time.

To deploy for an application, follow this procedure.

1. From the **Apps Profile** page, click the dropdown arrow for the desired application.
2. Select **Deploy** from the dropdown list. The Deploy *Application Name* page displays.
3. Enter information in the **General Information**, **Select Cloud Configuration**, **Parameter Values**, and **Schedule Options** sections as required for your deployment.
4. Click **Submit** to deploy the application.

Edit/Update

The Edit/Update task allows you update the parameters and topology an application.

To edit or update an application, follow this procedure.

1. From the **Apps Profile** page, click the dropdown **arrow** for the application.
2. Select **Edit/Update** from the dropdown list. The Edit App *Application Name* page displays.
3. Enter information in the **Basic Information**, **Application Configuration**, and **Advanced Onboarding Configuration** sections. See the [Topology Modeler](#) for additional context.
4. Click **Save App**. The application information is updated.

Clone

The Clone task allows you to clone an application.

To clone an application, follow this procedure.

1. From the **Apps Profile** page, click the dropdown **arrow** for the application.
2. Select **Clone** from the dropdown list. The Clone Application Profile popup displays.
3. Select the version of the application to clone from the dropdown list.
4. Click **OK**. The **Clone Name Application Profile** page displays.
5. Enter the information in the **Basic Information**, **Application Configuration**, and **Advanced Onboarding Configuration** sections.

Schedule Deployment

The Schedule Run task allows you to launch an application at a scheduled time.

To schedule a run for an application, follow this procedure.

1. From the **Apps Profile** page, click the dropdown **arrow** for the application.
2. Select **Schedule Deployment** from the dropdown list. The **Scheduling app: "Application Name"** page displays.
3. Enter information in the **General Information**, **Select Cloud Configuration**, **Parameter Values** sections as needed.
4. Enter the scheduling information in the in the **Schedule Options** section.

Schedule Starting

02/03/2016 3:30 PM

Repeat schedule Daily

Every 3 day(s)

Every Week Day

Schedule Ending

Never

After 2 occurrences

On 02/03/2016 at 3:30 PM



If you select the **Repeat schedule**, then you can further identify how often (**Daily**, **Weekly**, or **Monthly**) you want the schedule repeated.

For example, if you check **Daily**, you can select either **Every # day(s)**, or **Every Week Day** in the *Schedule Starting* section. If you select **Daily** and schedule a deployment by specifying a Schedule Starting time and stipulate an end after 5 occurrences, then a deployment will be scheduled at the starting time. After that a new deployment will be scheduled on a "daily" (weekly/monthly depending what is chosen) basis for next 5 days (number of occurrences). The schedule begins and is inclusive of the first occurrence which happens at the schedule starting time.

By default, if you do not configure any ending period in the *Schedule Ending* section, this schedule will continue endlessly (default = **Never**).

If you determine that you want the deployment to end **After # occurrences** in the *Schedule Ending* section, it will end after those number of deployments have been launched.

In the image above, the scheduling ends after 6 days. The setting of **Daily** + every 3 days + 2 occurrences = 6 days. In this case, even if **Daily** is selected, **Every 3 days** is also specified along with a configuration to end **After 2 occurrences**. So this schedule will conclude on Feb 5, 2015 (if you configured this schedule before 3:30 PM) or on Feb 8, 2015 (if you configured this schedule after 3:30 PM).

5. Click **Submit**. The application runs at the scheduled time.

Benchmark

The Benchmark task allows you to run a benchmark for an application. The benchmark process evaluates a variety of performance and cost metrics of the application.



Use large [Manage Instance Types](#) when issuing benchmark requests.

Concurrent benchmark requests are limited to 10,000 instance types. If you use small instance types, for example m1.small, the benchmark request fails.

To run a benchmark for an application, follow this procedure.

1. From the **Apps Profile** page, click the dropdown **arrow** for the application.
2. Select **Benchmark** from the dropdown list. The **Benchmarking app: "Application Name"** page displays.
3. Enter information in the **General Information, Select Cloud Configuration, Parameter Values** sections, if you have not done so already.
4. Click **Submit** to run the benchmark.

Schedule Benchmark

The Schedule Benchmark task allows you to run a Benchmark for an application at a scheduled time.

To schedule a benchmark for an application, follow this procedure.

1. From the **Apps Profile** page, click the dropdown **arrow** for the application.
2. Select **Schedule Benchmark** from the dropdown list. The **Scheduling app: "Application Name"** page displays.
3. Enter information in the **General Information, Select Cloud Configuration, Parameter Values** sections as needed.
4. Enter the scheduling information in the **Schedule Options** section.
5. Click **Submit** to run the benchmark at the scheduled time.

Export

The Export task allows you to export an application to a .zip file. You can then share this file with peers so they can **import** the application to their tenants.

To export an application, follow this procedure.

1. From the **Apps Profile** page, click the dropdown **arrow** for the application.
2. Select **Export** from the dropdown list. The system creates an apps.zip file of the application. Depending on your browser, the .zip file popup is automatically triggered and you are prompted to select the action..
3. If your browser prompts you for an action to take, choose to **save the file**.

The saved file is ready to be imported.

Share

The Share With task allows you to share an application with designated users and groups.

To share an application, follow this procedure.

1. From the **Apps Profile** page, click the dropdown **arrow** for the application.
2. Select **Share** from the dropdown list. The **Share** popup displays.
3. Add a user name in the **Users** section or a group name in the **Groups** section of this popup.
4. Each time you add a name or group, you can set the corresponding permissions. See [Permission Control](#) for additional context.
5. Click **OK** to share the application with the designated users and groups.

Show Deployments

This task displays a list of deployments for an application along with other pertinent details.

To display a list of runs for an application, follow this procedure.

1. From the **Apps Profile** page, click the dropdown **arrow** for the application.
2. Select **Show Deployments** from the dropdown list. The **Your Application Name Runs** page displays.

Delete

The Delete task allows you to delete an application from the Apps page. Deleting an application profile does not affect the deployments for this application profile.

To delete an application, follow this procedure.

1. From the **Apps Profile** page, click the dropdown **arrow** for the application.
2. Select **Delete** from the dropdown list. The **Delete Application Profile popup** displays.
3. Click OK to delete the application profile.

More Info

The More Info task allows you to view additional information about an application.

To view additional information about an application, follow this procedure.

1. From the **Apps Profile** page, click the dropdown **arrow** for the application.
2. Select **More Info** from the dropdown list. The resulting popup displays the following information:
 - Tenant name
 - Application Category
 - Tags
3. Click **Cancel** to close the popup and return to the Applications page.

Model Applications

Model Application Profiles and Applications

- [Application Profile Creation Considerations](#)
- [Create an Application Profile](#)
- [New Application Profile](#)
- [Topology Modeler](#)
- [Container Placement Groups](#)
- [Understand Application Tier Properties](#)
- [Deployment Parameters](#)
- [Deployment Lifecycle Scripts](#)
- [Multiple Volumes](#)
- [Security and Firewall Rules](#)
- [IP Address Allocation](#)
- [Parameters and Macros](#)
- [Application Using Imported Profile](#)
- [Application Using App Package](#)
- [Application Using Puppet or Chef](#)

Application Profile Creation Considerations

Application Profile Creation Considerations

- [Architectural Considerations](#)
- [Application Discovery Considerations](#)
- [Other Considerations](#)

Before you create an Application Profile, review your architectural requirements to determine the type of profile required for your application. Use the list in the following table to consider the different architectural aspects required for your application.

No.	Consideration	Links	Action Items and Details
1	What services are required for your application and how are these services dependent on each other?	See Supported OOB Services	Identify and provide a top-down hierarchy for the required services
2	Does your application contain multiple tiers?	See Understand Application Tier Properties	For each tier, provide the list of dependent services. The services listed in this section also identify the dependency from the service deployment perspective.
3	Are you using Workload Manager-supported services?	See Local Bundle Store (Conditional)	The bundle store contains the agent and service bundles. This repository is hosted by Workload Manager CDN (cdn.cliqr.com) or can be hosted locally in a private cloud environment using a standard Apache server.
4	Did you determine the hosting location for your service(s)?	See Local Package Store (Conditional) – if using Workload Manager-supported services and worker VMs cannot access the internet	Every Workload Manager deployment requires a Workload Manager Package Store and a hosting location for all Workload Manager-supported (<i>out-of-box</i>) Services. When you properly configure a region in Workload Manager, all services are automatically detected by the Package Store and displayed via the Topology Modeler Properties tab.
5	Are there external services that need to be connected to and communicated with (e.g., a new schema on an existing SQL server cluster, a new VIP on an existing load balancer)?	See External Service	You can create an application using multiple tiers where you can stipulate each tier to use a different Workload Manager-supported Service or externally-provided service (third-party services).
6	Are you using custom/private services?	See Custom Service Definition	Where are the required artifacts for these services to be successfully deployed?
7	Do you have all the dependent application data, files, and packages?	See Artifact Repository	Typically, enterprises maintain their application packages, data, and scripts in repositories. Use the Artifact Repository to attach your own external repository to store and access your files. Workload Manager provides a Repositories tab in the Workload Manager UI for this purpose.
8	Have you identified the infrastructure requirements for each service?	See Manage Instance Types	For private clouds, use the number of VM instances that you manage for your enterprise, and the instance type, storage and network each one requires, to calculate your hardware requirements.
9	What are the networking requirements for each service to talk to the others?	See Security and Firewall Rules	Use firewall rules and security groups to ensure proper inter-tier communication and external communication.

Workload Manager profiles are generally based on core configuration details, elaborate workflows that describe the sequence of steps required to deploy the application, or run on big-data processing workloads. Use the list in the following table to consider the different discovery aspects required for your application.

No.	Consideration	Links	Additional Details
1	Application Profile Type	See OOB Application Templates .	For each application profile, identify the category.
2	Application (or tier) requirements	See the following links: <ul style="list-style-type: none"> • Parameters and Macros • Deployment Lifecycle Scripts 	For each application (or tier), provide the dependent parameters and scripts.
3	Port and firewall accessibility	See Security and Firewall Rules .	Identify the scripts or dependencies at the time of deployment.
4	Deployment requirements	See Multiple Volumes .	You can attach multiple volumes to all tier types in N-tier applications. For each volume, you must specify the size and can optionally configure the volume type.
5	Testing and verification	Refer to your enterprise policies and requirements.	Be sure to identify your dependent test cases for each profile.

Besides the information required to model your application ensure to address the following external dependencies:

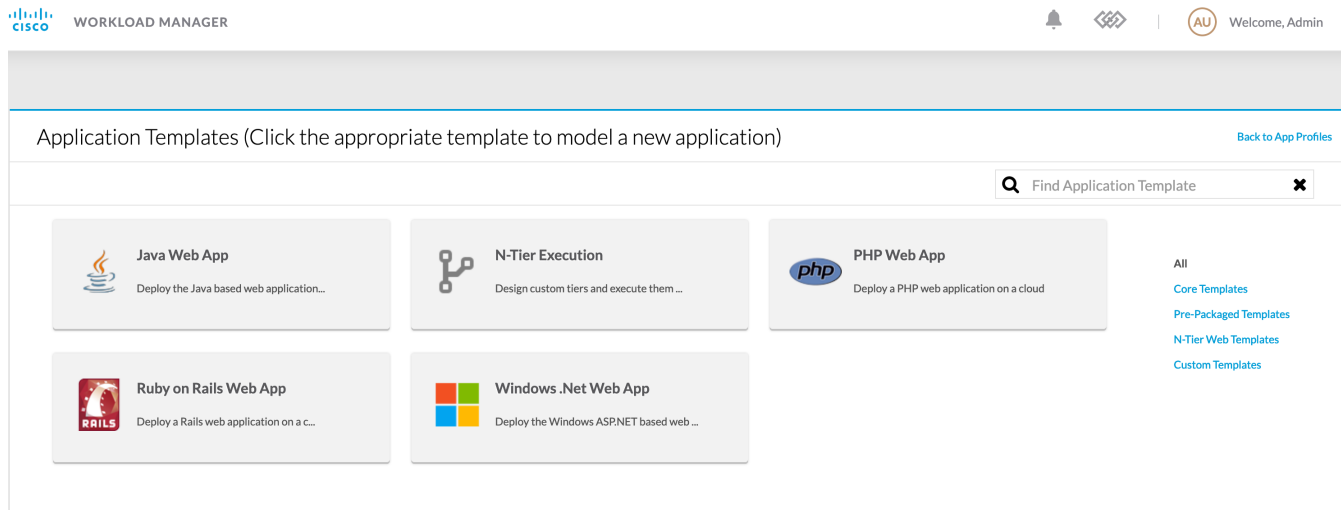
- Application packages and associated licenses, if applicable.
- Application credentials to access and verify each application, if applicable.
- Infrastructure requirements such as minimum CPUs, memory, storage, and other requirements specific to your application(s).

Create an Application Profile

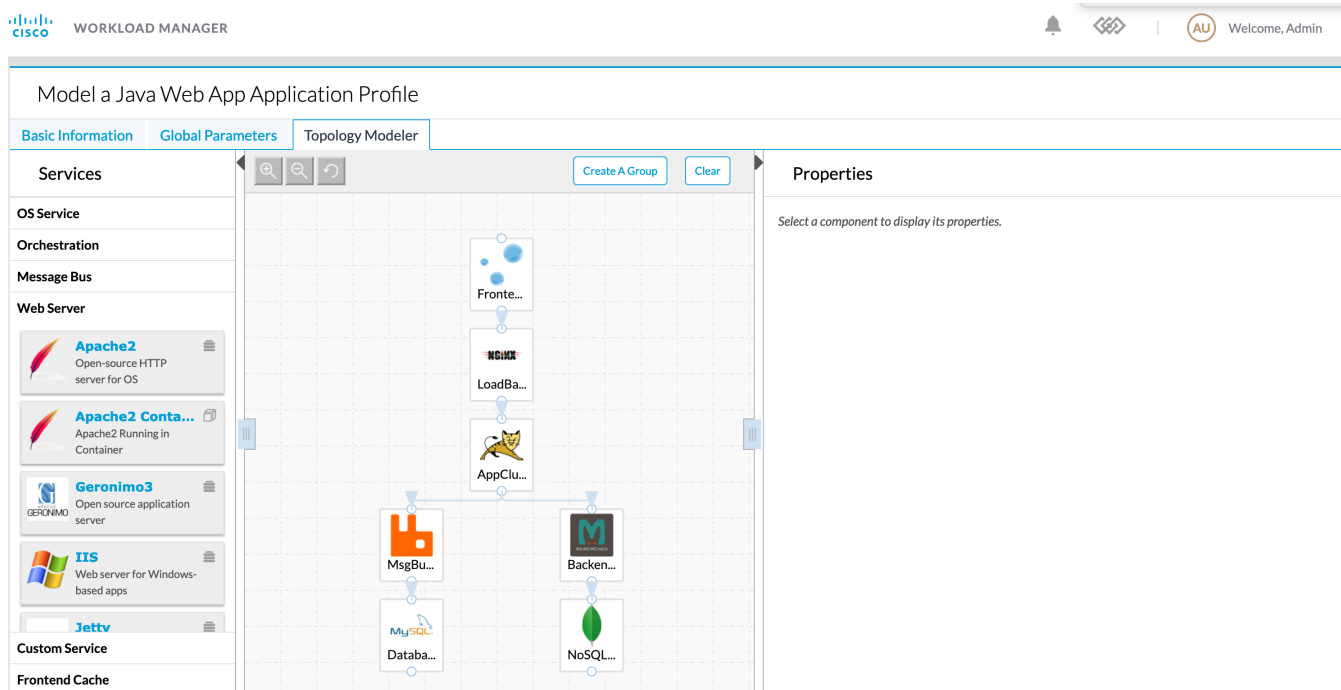
Create an Application Profile

Before creating a new application profile make sure you have reviewed the [Application Profile Creation Considerations](#) and have all of the needed components ready such as service definitions, repositories, and lifecycle actions. Once this is done, can create the application profile following this procedure.

Navigate to the Application Profiles page and click the Model button in the upper right of the page. This causes the Application Templates page to be displayed with all of the [OOB Application Templates](#) as shown in the figure below.



Click on the application template most suited for your application profile. This causes the Model Application Profile page with the Topology Modeler tab selected to be displayed. The initial Model Application Profile page associated with the N-Tier application template is shown below.



The Model Application Profile form has three tabs:

- Basic Information
- Global Parameters
- Topology Modeler

Start by defining the application topology and per tier properties in the Topology Modeler tab, then complete the required information in the remaining tabs.

The Topology Modeler tab consists of three panels from left to right (recall figure above):

- Services Palette
- Topology Canvas
- Properties Panel

The Services Palette lists all of the [Supported OOB Services](#) and any custom services created by you or created by other user and shared with you. The service details are only fetched on demand (when you drag and drop a service category option). The last updated version is used by default for the Service Version. If you have multiple versions, you can switch between versions for each application profile.

New Application Profile

Model a New Application Profile

- [Overview](#)
- [Guidelines](#)
- [Modeling Process](#)
 - [Create Real Images and Map to Logical Images](#)
 - [Services](#)
 - [Application](#)
- [Export Application Scripts](#)
 - [Sample JSON Format – N-Tier](#)
 - [Sample JSON Format – Single Tier](#)

Application Modeling is the process of capturing all images, scripts, and other dependencies required to fully deploy an existing, working application and building them into a model that you can configure using the guidelines provided in the next section.

Adhere to the following guidelines when modeling a new application profile:

- **Application Profiles:** To understand what an application profile is, see [Application Profile Overview](#).
- **Application Discovery Guidelines:** Begin your application modeling using a top-down approach focused on the application requirements. In this phase, nothing is configured in Workload Manager.
 1. Consider all of the services that make up the application (for example, Apache, Tomcat, JBoss, MySQL, Cassandra, SQL Server, and so forth).
 2. How are those services dependent on each other?
 3. Where are the artifacts that are required for the services to be successfully deployed?
 4. Are the services better deployed in a modular fashion, on small, discrete VMs or containers? Or is it better to lump some of all of them together onto a single VM?
 5. What are the networking requirements for each service to talk to others?
 6. Are there external services that need to be connected to and communicated with (for example, a new schema on existing SQL Server Cluster, New VIP on existing Load Balancer, and so forth).

The outcome from this step in the design process should be a topology diagram showing all the required VMs and external services, the ports required between each service, and which services are deployed on which VMs. See [Application Profile Creation Considerations](#) for additional context.

- **Required Service Guidelines:** Services should be developed to be usable in as many applications as possible without needing to re-modeling the same service each time. This will ensure future efficiency. For each service in the application, consider:
 1. Is there already a Workload Manager-provided *Service* that can be used? Or can an existing service be slightly modified to support this application?
 2. How can this Service be modeled to be generic?
 3. How is that service going to be deployed on the VM? Using scripts or perhaps Chef or Docker?
 4. If this is an external service, can you write a script to connect to that service and carry out whatever work is required? Are there existing libraries available that make this easier? Is one language or another easier for this? For example Python/requests VS Bash/curl VS Python library and so similar? Did you check out the corresponding libraries, for example, the Python library?
 5. What images do you need to support for this service? Is one image sufficient? Can you make this image cross-compatible with minimum effort? For example, use the `[-z /etc/redhat-release]` command to verify if you're on a CentOS/RedHat type system.
 6. Which inbound and outbound ports do you need to open for this service to function?

The outcome of this step should be a list of all services that your application requires, with a reference to the existing service that will be used if one is suitable. See [Supported OOB Services](#) to view a list of supported services in Workload Manager.

- **Supported Image Guidelines:** The VMs deployed by Workload Manager must be clones of existing [OOB Logical Images](#). These are stored on each cloud as AMIs, QCOWs, VM Snapshots or VM Template Names, and so forth. These images are the foundation to build higher-level services. In an ideal case they are very simple and generic, but often also have required security tools, monitoring agents, etc on them that will be included in every VM.

Keep in mind that services reference LOGICAL images, not real ones. The logical image references one REAL image per cloud region. This is an important part of how Workload Manager achieves cloud portability. The application and associated services are not dependent on any specific piece of infrastructure, not even the cloned images.

For each Logical Image required for your services, consider:

1. Is there an existing Logical Image that is suitable. If not, is there one that's close enough that can be made to work, or can your service be modified to fit?

IMPORTANT Considerations

- a. Don't make a new instance of a Logical Image type unless *absolutely necessary*.
- b. Ideally, you will have only one Logical Image for each type of OS that you require - and be sure to use as *few* OS types as possible.

2. If not, and if you need to create a new Logical Image:

- a. On what cloud regions do you need to run this image?
 - b. What OS do you need to use for this image?
 - c. Is there a standardized image build script that you can use?
3. What tools do you want to build into the image, if any?
 4. Do not add an application or service-specific configuration as it will be less useful for the next service.
 5. If you have a suitable Logical Image to use, does it have Real Image mappings for each region where you will want to deploy your application?

The outcome from this step should be a list of any Logical and Real Images that need to be created, how to create them, and where to create them for each services. After you have been through this a few times, the ideal outcome is generally, an empty list!

- **Modeling Process Guidelines:** Transition to a bottom-up approach when you model an application profile.
 - Workload Manager application models are composed of Services.
 - If the VM deployment is required for the service, then it can be mapped to a logical VM image. The logical VM images are in turn mapped to real images on a per-cloud basis (for example, AMI, VM-snapshot or template name, QCOW, and so forth depending on the cloud you use). See [Images Overview](#) for additional context.
 - An administrator can grant role-based permissions for application profiles. See [Permission Control > Role-Based Permissions](#) for additional context. If the user is not a member of the **Application Architects** or **Workload Manager admins** groups, the user cannot create application profiles and the model button in the app view will be disabled.

When you begin the modeling process in Workload Manager, start with configuring your images, followed by services, and finally, the application model.

Create Real Images and Map to Logical Images

If you need to create any images, create the real image and map to a logical image.

1. If you do not already have a logical Image that you can use, contact your Workload Admin to create the logical Image in Workload Manager. See [Images Page](#) (concepts and UI) for additional context.
2. Create a real Image in each cloud region.



Supported Images

You might save some time on this if you use one of the supported out-of-box [OOB Logical Images](#) for your cloud, if available and suitable to your enterprise requirements.

- a. Use your internal build process to complete this step. In some cases, you may only need to deploy the OS.
 - b. Install the Management Agent bundle using the installer. See [Worker \(Conditional\)](#) for additional context.
3. Apply the real image mapping to the logical image for the appropriate cloud region. See [Images Overview](#) for additional context.
 4. Test the images by creating blank services and applications – just for the purpose of launching and testing each image.

With your required images in place, you can start setting up services. Services have an associated lifecycle framework that calls different commands at different points in the service lifecycle. See [Service Lifecycle Actions](#) for additional context.

Services

1. Start with a dummy or other service script. This can be in any language, with any content that meets your needs, but a handy starting point written in bash is provided in the following Sample Script Code: This example contains a single script, called service, that is used to handle all VM lifecycle actions, with an argument (\$1) used to control which behavior is activated with each step.

Sample Script Code

```
#!/bin/bash
exec > >(tee -a /usr/local/osmosix/logs/service.log) 2>&1
OSSVC_HOME=/usr/local/osmosix/service
. /usr/local/osmosix/etc/.osmosix.sh
. /usr/local/osmosix/etc/userenv
. $OSSVC_HOME/utils/cfgutil.sh
. $OSSVC_HOME/utils/install_util.sh
. $OSSVC_HOME/utils/os_info_util.sh
cmd=$1
SVCNAME="dummy"
SVCHOME="$OSSVC_HOME/$SVCNAME"
USER_ENV="/usr/local/osmosix/etc/userenv"

case $cmd in
  install)
    log "[INSTALL] Installing $SVCNAME"
    ;;
  deploy)
    log "[DEPLOY] Deploying $SVCNAME"
    ;;
  configure)
```

```

    log "[CONFIGURE] Configuring $SVCNAME"
    ;;
start)
    if [ ! -z "$cliqrUserScript" -a -f "$cliqrUserScript" ]; then
        log "[START] Invoking pre-start user script"
        $cliqrUserScript 1 $cliqrUserScriptParams
    fi
    log "[START] Starting $SVCNAME"
    if [ ! -z $cliqrUserScript -a -f $cliqrUserScript ]; then
        log "[START] Invoking post-start user script"
        $cliqrUserScript 2 $cliqrUserScriptParams
    fi
    # Run restore script in case of migration
    if [ "$appMigrating" == "true" ]; then
        runMigrationRestoreScript
    fi
    log "[START] $SVCNAME successfully started."
    ;;
stop)
    log "[STOP] Invoking pre-stop user script"
    if [ ! -z $cliqrUserScript -a -f $cliqrUserScript ]; then
        $cliqrUserScript 3 $cliqrUserScriptParams
    fi
    log "[STOP] Stopping $SVCNAME"
    log "[STOP] Invoking post-stop user script"
    if [ ! -z $cliqrUserScript -a -f $cliqrUserScript ]; then
        $cliqrUserScript 4 $cliqrUserScriptParams
    fi
    log "[STOP] $SVCNAME successfully stopped."
    ;;
restart)
    log "[RESTART] Invoking pre-restart user script"
    if [ ! -z $cliqrUserScript -a -f $cliqrUserScript ]; then
        $cliqrUserScript 5 $cliqrUserScriptParams
    fi
    log "[RESTART] Restarting $SVCNAME"
    log "[RESTART] Invoking post-restart user script"
    if [ ! -z $cliqrUserScript -a -f $cliqrUserScript ]; then
        $cliqrUserScript 6 $cliqrUserScriptParams
    fi
    ;;
reload)
    log "[RELOAD] Invoking pre-reload user script"
    if [ ! -z $cliqrUserScript -a -f $cliqrUserScript ]; then
        $cliqrUserScript 7 $cliqrUserScriptParams
    fi
    log "[RELOAD] Reloading $SVCNAME settings"
    log "[RELOAD] Invoking post-reload user script"
    if [ ! -z $cliqrUserScript -a -f $cliqrUserScript ]; then
        $cliqrUserScript 8 $cliqrUserScriptParams
    fi
    log "[RELOAD] $SVCNAME successfully reloaded."
    ;;
cleanup)
    ;;
upgrade)
    log "[UPGRADE] Upgrading."
    ;;
*)
    log "[ERROR] unknown command"
    exit 127
    ;;
esac

```

The lifecycle actions would look like:

a. Install:

```
service install
```

b. **Configure:**

```
service configure
```

and so forth.

**Important Notes**

- It doesn't have to be done this way. You can use different scripts or commands or a different language like Python. This is your choice.
- See logging details on Line 2.
- See the sourced files on Lines 4-8. These are important to pick up helpful environment variables and utility functions.

2. In the Workload Manager UI, under **Main Menu Dashboard > Services**, create a new service with a descriptive name and type – use a dummy script as a starting point.
3. Complete the required fields as identified by the system validator.
4. Create a dummy application to test this service.
5. Continue building your service by launching a new test app deployment for this service often to test progress as you go along.
6. Refer to the above service tutorial for additional detail and best-practices.

Application

With the individual services built and tested using dummy applications, you can now pull the pieces together into the final working application.

1. Access the Workload Manager UI and click **Applications**.
2. Click the **Model** link in the top right corner. The Model a New Application Profile page displays with a list of application profiles.
3. Select a core, pre-packaged, or custom profile from the Workload Manager UI. For example, if you are modeling a Java web application, use the Java Web App profile.
4. Use the Workload Manager UI's [Topology Modeler](#) to define the application architecture and other components for each tier.
 - a. Drag and drop the required service from the Services pane to the Graphical pane in the Topology Modeler.
 - b. Connect the services with connectors that correspond to the order in which each service must be configured.



- These services ARE NOT related to network dependencies in any way.
- All VMs are created simultaneously (barring cloud-specific constraints)
- The lifecycle actions are executed from the bottom up according to the arrangement in the Topology Modeler.

- c. Click the service in the Graphical pane and configure the Properties for this service. See [Application Tier Properties](#) for additional details required to configure this pane.
 - d. Use debugging settings to troubleshoot issues – if required:
 - i. Add a global troubleshooting parameter (cliqIgnoreAppFailure) to your application and assign a default value of true. Be sure to come back and delete this parameter once the application, services, and images are working (see [Troubleshooting Parameters](#) for additional context).
 - ii. For each service tier, under Node Initialization, set sudo 'ALL' and make a note to come back and change this later.
 - iii. These two things will be very important during the setup and debugging phase. REMEMBER TO TAKE THEM OUT LATER.
 - e. Add firewalls rules as required for each Service Tier. See [Security and Firewall Rules](#) for additional context.
 - f. Add additional services as required for your deployment. The configuration may differ based on the service. See the following sample services for additional context:

Unable to render {children}.
5. Enter the **basic information** (component, logo, name, and so forth) according to the requirements for your application. See [Topology Modeler > Basic Information](#) tab for additional context.
 - a. Enter the general settings values (HTTP/HTTPS/Both/None based on the invoked protocol).



The Protocol field provides a None option (in addition to HTTP, HTTPS, and Both) when modeling N-tier applications. If you select **None**, the Workload Manager does not add any access link URL in the application deployment detail page.

- b. If it is a non-standard port, check the corresponding box.



For example, the default Tomcat server is started on port 8080. By setting the protocol to this port, you can access the server directly from the Deployment Details page.

- c. Enter the required values to access applications (port number in the URL, for example, if the non-standard port is 3309 and the URL will be `http://IpAddress:3309`).
- d. Specify the categories and tags for this application, if any. These fields are used as filters to search for this application.
- e. Provide the Access Link (path) for the application's launch page (landing page).
- f. Add metadata as relevant for your cloud. This is a useful way to flag the VMs.

AWS displays metadata as tags on the VMs. For example, if you use

```
Name / %USER_NAME%-%JOB_NAME%
```

When you log into the AWS console, you can see a list of app deployments and the user who launched the VM each deployment.

6. Define overriding parameters, if any, in the Topology Modeler's Global Parameters tab (for example, administrator, username, password, and so forth). See the Global Parameters section for additional context.
7. Once you enter all the definitions, you have multiple choices:
 - a. Click **Save as App** to save it to the [Application Tasks](#) page. If you save the modified profile as an Application, the definition is saved as a metadata file that can be exported in JSON format.
 - b. Click **Save as Template** to save it in the Model a New Application page.
8. Test your application or application profile in each cloud as applicable to your environment.

If you save the modified profile as an Application, the definition is saved as a metadata file that can be exported in JSON format. This section provides some sample application profile formats.

Sample JSON Format – N-Tier

This section provides an N-tier Jenkins application profile JSON Data Transfer Object (DTO) from the UI.

```
{
  "actionType": "saveApp",
  "sourceTemplateName": "N-Tier Execution",
  "appName": "Jenkins",
  "appDesc": "Leading Open Source continuous integration server built in Java to support building and
testing.",
  "appVersion": "1.54",
  "owner": "vik@cliqr.com",
  "helpLink": "",
  "cloneTemplateId": "",
  "storageClouds": [],
  "sourceAppId": null,
  "uiSystemTags": [],
  "microSegmentation": false,
  "parentVersion": "",
  "appId": 959,
  "appParamSpecs": [{
    "paramName": "cliqrWebappAccessLink",
    "type": "string",
    "defaultValue": "",
    "userVisible": true,
    "userEditable": true,
    "systemParam": true
  }, {
    "paramName": "cliqrExternalHttpsEnabled",
    "type": "string",
    "defaultValue": "0",
    "userVisible": true,
    "userEditable": true,
    "systemParam": true
  }... , {
    "paramName": "appPackage",
    "type": "path",
    "defaultValue": "",
    "userVisible": true,
    "userEditable": true,
    "systemParam": true
  }, {
    "paramName": "cliqrIgnoreAppFailure",
    "displayName": "cliqrIgnoreAppFailure",
    "helpText": "cliqrIgnoreAppFailure",
    "type": "string",
    "valueConstraint": {
      "maxLength": 255,
      "allowSpaces": true
    },
    "valueList": null,
    "collectionList": null,
    "defaultValue": "false",
    "userVisible": true,
    "userEditable": true,
    "optional": false,
    "webserviceListParams": null,
  }
}
```

```

        "systemParam": false
    }
],
"categoryIds": [],
"newCategories": [],
"categoriesArePublicVisible": false,
"storageChoice": "NONE",
"childSteps": [{
    "appName": "tomcat6_0",
    "internalImageName": "CloudWorker-CentOSx.x",
    "appVersion": "1.0",
    "sourceTemplateId": "",
    "cmdLines": [
        "nop"
    ]
},
"nodeReusable": false,
"serviceId": 7,
"appId": 961,
"appParamSpecs": [{
    "paramName": "cliqrWebServerType",
    "type": "string",
    "defaultValue": "tomcat6",
    "userVisible": false,
    "userEditable": false,
    "systemParam": true
}, {
    "paramName": "tierOrder",
    "type": "integer",
    "defaultValue": 1,
    "userVisible": false,
    "userEditable": false,
    "systemParam": true
}, {
    "paramName": "defaultService",
    "displayName": null,
    "helpText": null,
    "type": null,
    "valueList": null,
    "defaultValue": "",
    "userVisible": false,
    "userEditable": false,
    "systemParam": true,
    "exampleValue": null,
    "dataUnit": null,
    "optional": false,
    "valueConstraint": null,
    "scope": null,
    "webserviceListParams": null,
    "collectionList": []
}, {
    "paramName": "minClusterSize",
    "displayName": "Minimum number of nodes",
    "helpText": null,
    "type": "number",
    "valueList": null,
    "defaultValue": "1",
    "userVisible": true,
    "userEditable": true,
    "systemParam": true,
    "exampleValue": null,
    "dataUnit": null,
    "optional": false,
    "valueConstraint": {
        "minValue": 1,
        "maxValue": 1000,
        "maxLength": 0,
        "regex": null,
        "allowSpaces": false,
        "sizeValue": 0,
        "step": 0,
        "calloutWorkflowName": null
    }
}

```

```

    },
    "scope": null,
    "webserviceListParams": null,
    "collectionList": []
  }, {
    "paramName": "maxClusterSize",
    "displayName": "Maximum number of nodes",
    "helpText": null,
    "type": "number",
    "valueList": null,
    "defaultValue": "2",
    "userVisible": true,
    "userEditable": true,
    "systemParam": true,
    "exampleValue": null,
    "dataUnit": null,
    "optional": false,
    "valueConstraint": {
      "minValue": 1,
      "maxValue": 1000,
      "maxLength": 0,
      "regex": null,
      "allowSpaces": false,
      "sizeValue": 0,
      "step": 0,
      "calloutWorkflowName": null
    }
  },
  "scope": null,
  "webserviceListParams": null,
  "collectionList": []
}, {
  "paramName": "cliqrNoOfVolumes",
  "displayName": "Number of Volumes",
  "helpText": "Number of persistent data storage volumes for Web Server.",
  "type": "number",
  "valueList": null,
  "defaultValue": "0",
  "userVisible": true,
  "userEditable": true,
  "systemParam": true,
  "exampleValue": null,
  "dataUnit": null,
  "optional": true,
  "valueConstraint": {
    "minValue": 0,
    "maxValue": 65536,
    "maxLength": 0,
    "regex": null,
    "allowSpaces": false,
    "sizeValue": 5,
    "step": 1,
    "calloutWorkflowName": null
  }
},
"scope": null,
"webserviceListParams": null,
"collectionList": []
}, {
  "paramName": "cliqrDBDataStorageSize",
  "displayName": "Default Volume Size",
  "helpText": "Persistent data storage for Web Server.",
  "type": "number",
  "valueList": null,
  "defaultValue": "0",
  "userVisible": true,
  "userEditable": true,
  "systemParam": true,
  "exampleValue": null,
  "dataUnit": "GB",
  "optional": true,
  "valueConstraint": {
    "minValue": 0,

```

```

        "maxValue": 65536,
        "maxLength": 0,
        "regex": null,
        "allowSpaces": false,
        "sizeValue": 5,
        "step": 5,
        "calloutWorkflowName": null
    },
    "scope": null,
    "webserviceListParams": null,
    "collectionList": []
}, {
    "paramName": "cliqrJDKVersion",
    "displayName": "App Run-time",
    "helpText": "",
    "type": "list",
    "valueList": "JDK 6:JDK6,JDK 7:JDK7",
    "defaultValue": "JDK6",
    "userVisible": true,
    "userEditable": true,
    "systemParam": true,
    "exampleValue": null,
    "dataUnit": null,
    "optional": false,
    "valueConstraint": {
        "minValue": 0,
        "maxValue": 0,
        "maxLength": 0,
        "regex": null,
        "allowSpaces": true,
        "sizeValue": 0,
        "step": 0,
        "calloutWorkflowName": null
    },
    "scope": null,
    "webserviceListParams": null,
    "collectionList": []
}, {
    "paramName": "cliqrWARFile",
    "displayName": "App Package",
    "helpText": "Application package file. The file is in relative path from %rootPath%. ",
    "type": "path",
    "valueList": null,
    "defaultValue": "%REPO_ID_21%apps/jenkins/jenkins.war",
    "userVisible": true,
    "userEditable": true,
    "systemParam": true,
    "exampleValue": null,
    "dataUnit": null,
    "optional": false,
    "valueConstraint": {
        "minValue": 0,
        "maxValue": 0,
        "maxLength": 255,
        "regex": null,
        "allowSpaces": true,
        "sizeValue": 0,
        "step": 0,
        "calloutWorkflowName": null
    },
    "scope": null,
    "webserviceListParams": null,
    "collectionList": []
}, {
    "paramName": "cliqrWebappConfigFiles",
    "displayName": "App Config files",
    "helpText": "Application config files that contain CliQr system tokens and will be modified at deployment time. The config file is a relative path from the webapp context folder, e.g., WEB-INF/classes/db.conf. If there are multiple files, separate with semicolon.",
    "type": "string",
    "valueList": null,

```



```

    "defaultValue": "",
    "userVisible": true,
    "userEditable": true,
    "systemParam": true,
    "exampleValue": null,
    "dataUnit": null,
    "optional": true,
    "valueConstraint": {
      "minValue": 0,
      "maxValue": 0,
      "maxLength": 255,
      "regex": null,
      "allowSpaces": true,
      "sizeValue": 0,
      "step": 0,
      "calloutWorkflowName": null
    },
    "scope": null,
    "webserviceListParams": null,
    "collectionList": []
  }, {
    "paramName": "cliqrWebappContext",
    "displayName": "Deploy Context",
    "helpText": "",
    "type": "string",
    "valueList": null,
    "defaultValue": "ROOT",
    "userVisible": true,
    "userEditable": true,
    "systemParam": true,
    "exampleValue": null,
    "dataUnit": null,
    "optional": false,
    "valueConstraint": {
      "minValue": 0,
      "maxValue": 0,
      "maxLength": 255,
      "regex": null,
      "allowSpaces": true,
      "sizeValue": 0,
      "step": 0,
      "calloutWorkflowName": null
    },
    "scope": null,
    "webserviceListParams": null,
    "collectionList": []
  }, {
    "paramName": "cliqrEARPath",
    "displayName": "EAR file",
    "helpText": "",
    "type": "string",
    "valueList": null,
    "defaultValue": "",
    "userVisible": false,
    "userEditable": false,
    "systemParam": true,
    "exampleValue": null,
    "dataUnit": null,
    "optional": false,
    "valueConstraint": {
      "minValue": 0,
      "maxValue": 0,
      "maxLength": 255,
      "regex": null,
      "allowSpaces": true,
      "sizeValue": 0,
      "step": 0,
      "calloutWorkflowName": null
    },
    "scope": null,
    "webserviceListParams": null,

```

```

    "collectionList": []
  }, {
    "paramName": "cliqrPlanPath",
    "displayName": "Plan file",
    "helpText": "",
    "type": "string",
    "valueList": null,
    "defaultValue": "",
    "userVisible": false,
    "userEditable": false,
    "systemParam": true,
    "exampleValue": null,
    "dataUnit": null,
    "optional": false,
    "valueConstraint": {
      "minValue": 0,
      "maxValue": 0,
      "maxLength": 255,
      "regex": null,
      "allowSpaces": true,
      "sizeValue": 0,
      "step": 0,
      "calloutWorkflowName": null
    },
    "scope": null,
    "webserviceListParams": null,
    "collectionList": []
  }, {
    "paramName": "appMigrationFiles",
    "displayName": "Application Migration Files",
    "helpText": "",
    "type": "string",
    "valueList": null,
    "defaultValue": "cliqrWARFile,cliqrEARPath,cliqrPlanPath",
    "userVisible": false,
    "userEditable": false,
    "systemParam": true,
    "exampleValue": null,
    "dataUnit": null,
    "optional": false,
    "valueConstraint": {
      "minValue": 0,
      "maxValue": 0,
      "maxLength": 255,
      "regex": null,
      "allowSpaces": true,
      "sizeValue": 0,
      "step": 0,
      "calloutWorkflowName": null
    },
    "scope": null,
    "webserviceListParams": null,
    "collectionList": []
  }, {
    "paramName": "cliqrTomcat6PreStartAction",
    "type": "path",
    "defaultValue": "",
    "userVisible": true,
    "userEditable": true,
    "systemParam": true
  }, ...
  {
    "paramName": "cliqrExternalPostStopAction",
    "type": "path",
    "defaultValue": "",
    "userVisible": true,
    "userEditable": true,
    "systemParam": true
  }, {
    "id": null,
    "paramName": "GREET",

```

```

        "displayName": "GREET",
        "helpText": "GREET",
        "type": "string",
        "valueConstraint": {
            "minValue": 0,
            "maxValue": 0,
            "maxLength": 255,
            "regex": null,
            "allowSpaces": true,
            "sizeValue": 0,
            "step": 0,
            "calloutWorkflowName": null
        },
        "valueList": null,
        "collectionList": [],
        "linkedParams": null,
        "defaultValue": "HELLO",
        "userVisible": true,
        "userEditable": true,
        "systemParam": false,
        "exampleValue": null,
        "optional": true,
        "dataUnit": null,
        "scope": null,
        "webserviceListParams": null,
        "order": 1
    }, {
        "paramName": "resumeScript",
        "type": "path",
        "defaultValue": "",
        "userVisible": true,
        "userEditable": true,
        "systemParam": true
    }, {
        "paramName": "topTier",
        "displayName": "topTier",
        "type": "boolean",
        "defaultValue": true,
        "userVisible": false,
        "userEditable": false,
        "systemParam": true
    }
    ],
    "dependentStepNames": [],
    "topologyParams": {
        "x": 160,
        "y": 160
    },
    "envvarSpecs": [],
    "firewallRules": [{
        "protocol": "tcp",
        "fromPort": 80,
        "toPort": 80,
        "sourceIPRanges": ["0.0.0.0/0"]
    }
    ],
    "hwProfile": {
        "numOfCPUs": 1,
        "memorySize": 1024,
        "localStorageSize": 0,
        "numOfNICs": 1,
        "cudaSupport": "false",
        "supportHardwareProvision": false
    },
    "swProfile": {
        "support32bit": "false",
        "support64bit": "true",
        "support32On64": "false"
    },
    "nodeInitSpec": {
        "initScript": "",

```

```

        "cleanupScript": "wget http://<HOST>:<PORT>/cleanup.sh && chmod 0755 cleanup.sh && ./cleanup.sh",
        "depPkgs": "",
        "sudoCmdList": "ALL"
    },
    "migrationSpec": {
        "preMigrateScript": "",
        "backupScript": "",
        "backupLocation": "",
        "restoreScript": "",
        "postMigrateScript": ""
    },
    "upgradeSpec": {
        "upgradeType": "0",
        "preUpgradeScript": "",
        "stop": false,
        "upgradeScript": "",
        "start": false,
        "postUpgradeScript": "",
        "rollbackScript": ""
    },
    "sourceAppId": null,
    "uiSystemTags": [],
    "repositories": [{
        "id": "21"
    }
    ],
    "storageChoice": "NONE"
}],
"repositories": [],
"jobRuntimeTags": [{
    "tagCollectionId": 401,
    "tagName": "Username",
    "value": "%USER_NAME%",
    "required": true,
    "editable": false
}, {
    "tagCollectionId": 401,
    "tagName": "DeploymentEnv",
    "value": "%DEPLOYMENT_ENV%",
    "required": false,
    "editable": true
}, {
    "tagCollectionId": 401,
    "tagName": "Jobname",
    "value": "%JOB_NAME%",
    "required": true,
    "editable": true
}, {
    "tagCollectionId": 401,
    "tagName": "Whatever",
    "value": "what",
    "required": false,
    "editable": false
}
]
}

```

Sample JSON Format – Single Tier

This section provides the single-tier application profile JSON DTO from the UI.

```

{
  "id": "285",
  "resource": null,
  "perms": [],
  "name": "App_Parallel_Ex_1",
  "description": "A Parallel Execute Command",
  "version": "PRL-v1.0-U10",

```

```

"revisionId": 0,
"sourceTemplateId": 1,
"sourceTemplateName": "Parallel Execute",
"sourceAppId": 0,
"parentVersion": "",
"executorBeanName": "cmdLineExecutor",
"owner": null,
"hwProfile": {
  "memorySize": 512,
  "numOfCPUs": 1,
  "networkSpeed": null,
  "numOfNICs": 1,
  "localStorageCount": 0,
  "localStorageSize": 10,
  "cudaSupport": false,
  "ssdSupport": false,
  "supportHardwareProvision": false
},
"swProfile": {
  "osName": "Linux",
  "support32Bit": true,
  "support64Bit": true,
  "support32On64": true
},
"nodeInitSpec": {
  "depPkgs": "",
  "initScript": "%REPO_ID_2%checkExecNew.sh lapp node-init \"NodeInit Executed\"",
  "cleanupScript": "",
  "licServer": "",
  "licPort": 0,
  "sudoCmdList": "ALL"
},
"firewall": null,
"migrationSpec": null,
"upgradeSpec": null,
"nodeReusable": true,
"leafLevel": true,
"interactiveApp": false,
"defaultApp": true,
"applicationType": "defaultApp",
"templateType": "APPLICATION",
"policyId": 0,
"statusDefId": 0,
"tags": "",
"nodeCombination": null,
"internalImageName": "CloudWorker-CentOSx.x",
"helpLink": null,
"numNodes": 0,
"dynamicScalable": false,
"serviceName": null,
"topologyParamText": null,
"templateDependencies": [],
"appParamSpecs": [{
  "paramName": "NumNodes",
  "displayName": "NumNodes",
  "helpText": "",
  "type": "number",
  "valueList": null,
  "defaultValue": "2",
  "userVisible": true,
  "userEditable": true,
  "systemParam": true,
  "exampleValue": null,
  "dataUnit": null,
  "optional": false,
  "valueConstraint": {
    "minValue": 0,
    "maxValue": 255,
    "maxLength": 0,
    "regex": null,
    "allowSpaces": true,

```

```

        "sizeValue": 0,
        "step": 0,
        "calloutWorkflowName": null
    },
    "scope": null,
    "webserviceListParams": null,
    "collectionList": []
}, {
    "paramName": "OutputDir",
    "displayName": "Output Directory",
    "helpText": "Output Directory",
    "type": "string",
    "valueList": null,
    "defaultValue": "%OUTPUT_DIR%/%DATE%/%JOB_NAME%",
    "userVisible": true,
    "userEditable": true,
    "systemParam": true,
    "exampleValue": null,
    "dataUnit": null,
    "optional": false,
    "valueConstraint": {
        "minValue": 0,
        "maxValue": 0,
        "maxLength": 255,
        "regex": null,
        "allowSpaces": true,
        "sizeValue": 0,
        "step": 0,
        "calloutWorkflowName": null
    },
    "scope": null,
    "webserviceListParams": null,
    "collectionList": []
}, {
    "paramName": "NUM_PARAM",
    "displayName": "NumParam",
    "helpText": "Sample Number Param",
    "type": "number",
    "valueList": null,
    "defaultValue": "5",
    "userVisible": true,
    "userEditable": true,
    "systemParam": false,
    "exampleValue": null,
    "dataUnit": null,
    "optional": false,
    "valueConstraint": {
        "minValue": 0,
        "maxValue": 255,
        "maxLength": 0,
        "regex": null,
        "allowSpaces": true,
        "sizeValue": 0,
        "step": 0,
        "calloutWorkflowName": null
    },
    "scope": null,
    "webserviceListParams": null,
    "collectionList": []
}, {
    "paramName": "TEXT_PARAM",
    "displayName": "TextParam",
    "helpText": "Sample Text Param",
    "type": "textarea",
    "valueList": null,
    "defaultValue": " Sample Test Area Content      ",
    "userVisible": true,
    "userEditable": true,
    "systemParam": false,
    "exampleValue": null,
    "dataUnit": null,

```

```

    "optional": false,
    "valueConstraint": {
      "minValue": 0,
      "maxValue": 0,
      "maxLength": 0,
      "regex": null,
      "allowSpaces": true,
      "sizeValue": 0,
      "step": 0,
      "calloutWorkflowName": null
    },
    "scope": null,
    "webserviceListParams": null,
    "collectionList": []
  }],
  "envvarSpecs": [{
    "name": "test_macro_spaces",
    "value": "%APP_DIR% test value 1",
    "userVisible": false,
    "userEditable": false
  }],
  "agentTasks": [{
    "commandType": "cmdExec",
    "params": [{
      "name": "cmdLine1",
      "value": "touch %OutputDir%/execfile"
    }, {
      "name": "cmdLine2",
      "value": "echo "command executed" >>> %OutputDir%/execfile"
    }, {
      "name": "cmdLine3",
      "value": "echo %NUM_PARAM% >>> %OutputDir%/execfile"
    }, {
      "name": "cmdLine4",
      "value": "echo %NODE_INDEX% >>> %OutputDir%/execfile"
    }, {
      "name": "cmdLine5",
      "value": "echo %TEXT_PARAM% >>> %OutputDir%/execfile"
    }, {
      "name": "numCmdLines",
      "value": "5"
    }
  ]
}],
  "childAppTemplates": [],
  "categories": [],
  "metadataTags": [],
  "systemTags": []
}

```

Topology Modeler

Understand the Topology Modeler

- [Overview](#)
- [Basic Information Tab](#)
- [Global Parameters Tab](#)
- [The Topology Modeler Tab](#)
 - [Graphical Workflow](#)
 - [Properties](#)

The Workload Manager metadata descriptors are categorized in the following tabs:

- Basic Information
- Global Parameters
- Topology Modeler

This section provides details on each tab.

Workload Manager *models* N-tier applications (or application profiles) based on basic information, parameters, topology-specific services and properties that you define for each application. This information provides the metadata descriptors (JSON objects) for each application and Workload Manager stores the metadata information to determine cloud-compatibility and application benchmarks when deploying your application.

This tab collect basic information about the application such as Name, Version, Description, Category, Tags, Metadata, and other details when creating the profile. The following screenshot shows this tab.

Edit "Centos7Auto" Application Profile

Version: [1.1](#) (Revision: 0)

Basic Information

Global Parameters

Topology Modeler

Web App Name *

Centos7Auto

Version *

1.1

Revision

0

Specify Deployment Environment?

OFF

Description

Sample App with Basic
Centos7

Categories

Other
Advertising
Analytics
Bioinformatics

Note: To select multiple categories, CTRL/CMD+Click on the options.

Add New



example: category1, category2, category3. Click "Add" to add it to the list.

Help Link

Access Link

Protocol

HTTP HTTPS Both None

Non standard port

App Content Package

----Select a Location----

Archive of your application content, including application binaries and scripts. Currently support .zip format.

The **Specify Deployment Environment** toggle switch, if enabled (**ON**), allows you to specify select deployment environments for an application.

- **OFF** (Default): You can deploy this application to any deployment environment to which you have *Deploy To* access as specified in [Permission Control](#) > *Deployment Environment Permissions* > **Deploy To** row.
- **ON**: The **Allowed Deployment Environments** multi-select field is visible and can further restrict the deployment environments available at deploy time to those entered in the list. You must select at least one choice from the dropdown list. The first environment that you select becomes the pinned choice as indicated by the pin icon. When you select multiple environments, you can pin any one of those environments. Pinning an environment indicates that it will be the default environment at deploy time.

Specify Deployment Environment?



Allowed Deployment Environments *

CCQA_deploy_env sidenv

- CCQA_deploy_env
- kubernetes
- kub
- sidenv

The **Protocol** field provides an additional None option (in addition to HTTP, HTTPS, None, and Both) when modeling N-tier applications. If you select **None**, Workload Manager does not add any access link URL in the application deployment detail page.

i The HTTP, HTTPS, or Both option selection using a standard or non-standard port is only for the proper application access link URL setting. Be sure to verify that the top tier application firewall rules is set accordingly. The top application tier is set to use firewall rules specified in the service definition by default (for example, the Apache service opens port 80 and 443). Based on the application's requirement, however, users can modify this rule in the application profile.

In this tab, you can also add metadata tags that are carried over with the job/deployment when submitted.

Metadata | You can add metadata here for this job

Name	Value	Mandatory	Editable	Actions
<input type="text" value="Name"/>	<input type="text" value="Value"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="button" value="Add"/>
Name	%JOB_NAME%	Yes	No	<input type="button" value="⊗"/>

Tip: Metadata will not be saved until you click submit

The parameters defined in the Global Parameters tab apply to all tiers and steps in the application or application profile.

Edit "Jenkins_vol" Application Profile Close

Version: 1.0 (Revision: 16)

Basic Information | **Global Parameters** | Topology Modeler

If you would like to add additional parameters, this is the place.

Global Parameters | add a parameter +

Parameter

Parameter Name

Display Name

Help Text

Type: string Max Length: 255

Default Value

User Options:

- Should this parameter be visible to the user?
- Should this parameter be editable by the user?
- Should this parameter be optional?

Continue to Topology +

The [Using Parameters](#) page explains the differences between defining parameters at different levels.

The [Troubleshooting Parameters](#) page explains the process and a use case for global parameters.

The Topology tab contains three main sections:

- Graphical Workflow
- Services
- Topology Modeler Properties

Graphical Workflow

The graphical workflow for each application. When you click on a specific tier in the graphical workflow for each application, you see the configurable parameters and fields for that tier.

WORKLOAD MANAGER

Edit CustomLinux Application Template

Version: R7 (Revision: 0)

Basic Information | **Global Parameters** | **Topology Modeler**

Services

OS Service

Orchestration

Message Bus

Web Server

Apache2

Apache2 C

DDTESTS

Geromimo

IIS

LinuxAgent

CPU: 1
Memory: 1GB
Storage: 0GB

Properties

General Settings

v3
v2
v1

Name *

LinuxAgent_1

Base Image *

RHEL 7.x

SERVICE_PARAM_LIST

!@#

SERVICE_PARAM_LIST

SERVICE_PARAM_NUMBER *

Service Initialization

External Initialization

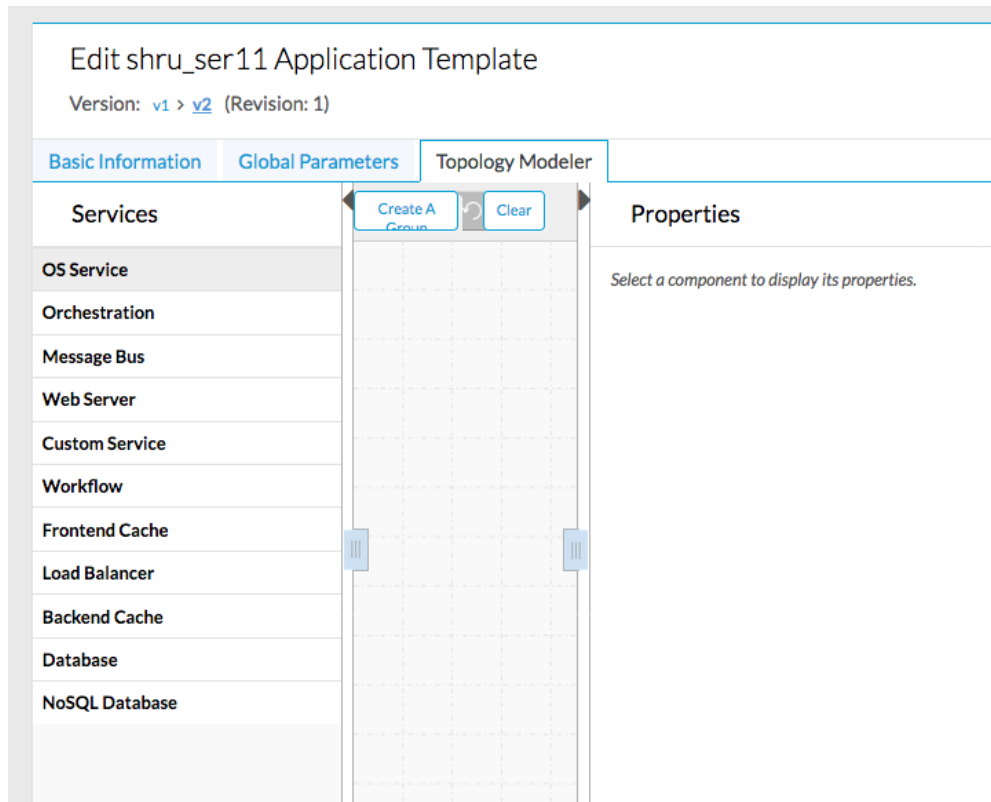
Firewall Rules

Use the Workload Manager UI's Topology Modeler to define the application architecture and components. The Topology Modeler allows users to model complex application topologies rapidly in a simple, cloud-agnostic fashion and maintains intricate service dependencies. This Topology Modeler also allows users to plug in custom scripts (application configuration scripts) and services. See [Parameters and Macros](#) for additional information.

The newly created service is displayed on the Services palette in the Topology Modeler. When you select a service from the palette to add to your application, you can further qualify this instance of the service by defining additional parameters in the Properties pane. When modeling an application, the dependencies between tiers are indicated using arrows going from one tier to the next tier. The arrowhead represents the VM instance that must function with all service scripts fully executed before you configure the VM instance at the tail of the arrow.

Services

Use the Services palette to identify the required services for your application. Workload Manager allows users to select CloudCenter-supported services or define (admins only) custom services via the Services framework. Admins can choose to make the custom services available to all users within a tenant. The Service categories are loaded default.



When you (admin) create a new service, you need to select *one intuitive* category within which to display this service in the Topology Modeler's **Services** tab.

When you add a service, the newly-created service is displayed on the Services palette in the Topology Modeler.

- To view a list of supported out-of-box services, see [OOB Services](#).
- To create/modify a new/existing service, see [Custom Service Definition](#).
- To understand the nuances related to adding/modifying services, see [Service Administration](#).
- To understand activities that users can perform with custom services, see [Permission Control](#).
- To understand the properties for each service at the application tier level, see [Understand Application Tier Properties](#).

Workload Manager also allows you to define a custom service to launch a VM on the cloud, install component(s), configure component(s), start the script, scale component(s), and shutdown/delete component(s). You can further qualify each service by defining property parameters specific to each instance of this service.

Properties

The Properties palette allows you to specify scripts or binaries to install, customize, configure, start, stop, upgrade, or run the application for any given tier.

The parameters specified in the Global Parameters tab apply to the entire application. The parameters specified in this section only apply to a specific tier within the application.

When you click on a specific tier in the Topology Modeler graphical workflow, you see the configurable parameters and fields for that tier. See [Deploy an Application](#) for additional details.

The configurable fields for each tier are explained in the [Understand Application Tier Properties](#) section.

Container Placement Groups

Container Placement Groups

- [Overview](#)
- [Process Details](#)
- [Container Service](#)

Through the topology modeler tab of the Add Application Profile form, Workload Manager lets you create a single container tier with multiple containers in the same pod. The graphical tool that lets you do this is called a container placement group.

A placement group represents a container pod and a tier in the application profile. A placement group is represented by a rectangle that you add to the topology modeler canvas by clicking the Create A Group button as indicated in the following screenshot. This support is available for some services like Apache, Nginx, MongoDB, and MySQL.

You can only drag and drop supported containers into the Create a Group rectangle.

Click and drag container-based services from the services palate on the left side of the topology modeler tab to the placement group rectangle, similar to the Apache2 profile and MongoDB profile that are visible in the canvas displayed in the following screenshot.

The properties panel on the right displays the properties for the currently selected object, which in this case is the last container dragged into the placement group, MongoDB. Notice that the minimum and maximum replicas fields in the General Settings section are not available, and the Network Services and Firewall Rules sections are also not available for this container. This is because these parameters are properties of the tier/pod. To see these parameters, click on the border of the placement group to select the placement group as a whole.

After selecting the placement group, notice the following in the Properties panel:

The name of the placement group contains the name of the first container dragged to the group (but you can edit this field).

Expanding the sections for Network Services and then Firewall Rules reveals that these sections automatically contain the union of the corresponding network services and firewall rules of the constituent containers.

The sections for Volumes, Deployment Parameters, and Minimum Resource Specifications are gone as these are specific to each container in the pod.

When deploying an application with a container tier that contains multiple containers, each container is listed within the corresponding per tier section on Page 1 of the deploy form.

If the container has visible parameters (defined in the application profile), an expand triangular icon appears to the left of the container name. Click on the expand icon to expand that section and display the deployment parameters.

On Page 2 of the deploy form you have the option to specify the instance type for each container separately.

After the application successfully deploys, the individual containers within each replica of the pod are shown in the deployment Details page as described in the Deployment Details section.

See [Container Service](#) for details on the *Generate unique service name* toggle switch.

Understand Application Tier Properties

Understand Application Tier Properties


- [Overview](#)
- [Terminology](#)
- [General Settings](#)
- [External Initialization](#)
- [Node Initialization & Clean Up](#)
- [Service Initialization](#)
- [Firewall Rules](#)
- [Deployment Parameters](#)
- [Hardware Specification](#)
- [Environment Variables](#)
- [Migration](#)
- [Upgrade](#)
- [Other References](#)

The [Topology](#) tab contains three main sections (or panes in the Workload Manager UI):

- Topology Modeler **Graphical Workflow**: The graphical workflow for each application. When you click on a specific tier in the graphical workflow for each application, you see the configurable parameters and fields for that tier.
- Topology Modeler **Services**: Type and group of the defined or selected service. Service properties differ based on the selected service.

 Depending on the service that you select in this section, the corresponding parameters and settings may differ for each service in the **P** **roperties** section.

- Topology Modeler **Properties**: Application properties defined using service parameters to further qualify application actions. This section explains the Properties that you can configure for an application tier.

 These properties include a list of property groups that are used in one or more Out-of-box (OOB) services offered by Workload Manager.

Each property group explained below will differ based on the service that you model or deploy.

This section provides a list of all the parameters in each property group. If the properties differ, they are explained in the respective **OOB Services**.

Term	Definition
<i>Service</i>	A reusable component that operates on the supported OOB Logical Images . You can define a new service to capture framework dependencies (called when you onboard or migrate one or more applications).
<i>Service library</i>	A collection of supported OOB Services and custom images that are included in Workload Manager and supports each application stack. An OOB service is sometimes referred to as a <i>system service</i> or a <i>supported service</i> .
<i>Custom service</i>	A user-defined service that resides in the service library.
<i>External service</i>	Any service that resides in another location. See External Service for additional context.
<i>Application stack</i>	A composite topology that includes a variety of formats of services, images, and containers.
<i>Topology Modeler</i>	A Workload Manager UI tool used to model an application profile using various profile templates and out of box services, images, and containers. See Topology Modeler for additional context.
<i>Service Lifecycle Actions</i>	The actions to per-determine granular steps of the service deployment process for each application. See Service Lifecycle Actions for additional context.

This section is unique based on how the service was defined.

The following table lists the configurable properties to ensure that the application is functioning.

Properties	Description
Name	A descriptive name for this service.

Service Version	Provide a service version for this service, if it is service derived from an earlier Workload Manager release, v1 is automatically assigned to it. Each service can have multiple versions. You must map the application profile to the required service version, this configuration is required and validated.
Base Image	The OOB Logical Images in which you are installing this service.
Minimum number of nodes	The minimum number of nodes (VMs) in this cluster. If the number of active VMs (nodes) is lower than the number specified in this field, then the application may not function as designed.
Maximum number of nodes	The maximum number of nodes in this cluster. Once the number specified in this field is reached, then the application cannot be scaled.
Number of Volumes	Number of persistent data storage volumes for Web Server. See Multiple Volumes for additional context.
Default Volume Size	Persistent data storage for Web Server. Persistent Storage is defined by default for database tiers but can be assigned manually to any tier See Multiple Volumes for additional context.
Allowed Scaling Policies	<p>A list of available policies for this application with each policy having an info icon that displays additional details on time duration, CPU utilization, and other details. See Policy Management for additional details on each type of policy.</p> <ul style="list-style-type: none"> • If you do not select an allowed scaling policy for a tier, the scaling policy field for that tier is hidden in the Deploy form and no scaling policies can be selected by users at deploy time. • If only one allowed scaling policy is selected, the corresponding visibility toggle is enabled and set to on by default. If one policy is selected and the visible toggle is turned off, the scaling policy cannot be disabled at deploy time but is still applied to the deployment behind the scenes. • If at least one policy is selected, a <i>mandatory</i> toggle is displayed to the right of the visibility toggle and a deploy time preview link appears under the allowed scaling policies field. <p>Visible: Identifies if the selected policy should be visible (ON) to users in your tenant and makes it visible in the deploy flow.</p> <p>Mandatory: This setting is only available if you switch on visibility and allows you to determine if this application should require this policy (ON) or it is an optional setting (OFF)</p> <p>Deploy Time Preview: Identifies the required time-based option for the selected policies. Refer to the info icon for the required policy in the previous step to select the corresponding option.</p> <p>Click the Deploy Time Preview link for a popup window to display the tier scaling policy dropdown field's appearance for that tier in the per-tier section on Page 1 of the Deploy form.</p>
Associate Tags	A label that consists of a name and an optional description for metadata association and for tracking purposes. See System Tags for additional details.
App Config files	This field is specific to the web server. Application config files that contain system Parameters and Macros that may need to be modified at deployment time. The config file uses a relative path from the web app root folder, for example, webinfo.conf. If you are configuring multiple files, separate each file in the list using a semicolon. See Configuration Files and Parameter Substitution for additional context.
Deploy Folder	This field is specific to the web server. When you deploy a web server application, use this field to specify the <code>/service/var/www/html</code> root path where the folder must be created. For example, if your folder is called Doc, then this folder is created as <code>/apache/var/www/html/Doc</code> . The app binaries are unzipped in this folder before deployment.
Associate Tags	You can associate system tags with application profiles and application deployments, either at the tier level or globally. See System Tag for additional context.

Defines the scripts to be run before (pre) and after (post) the start of the Service or pre/post stopping of the service when the application is rebooted, suspended, or terminated.

External Initialization scripts are run at the same time as their equivalent service initialization scripts, but are run from an isolated container on Docker.

See [External Service](#) for additional details on this section.




You must configure the information explained in this section:

- If your deployment uses external services, to ensure that the application is functioning.
- To ensure that the node is functioning.


The following table lists the configurable properties that define each script execution.

Properties	Description
VM Pre-Provision Script	Script executed before the application VM is provisioned.
VM Pre-Initialization Script	Script executed before the application VM service is initialized.
VM Post-Start Script	Script executed after the application VM service is started.
VM Pre-Terminate Script	Script executed before the application VM is terminated.
VM Post-Terminate Script	Script executed after the application VM is terminated.

 You must configure the information explained in this section to ensure that the application is functioning.

See [Deployment Lifecycle Scripts](#) for additional details on this section.

See [Service Administration](#) for additional details on configuring the service location.


 Each node initialization and clean up script is run once for each VM function.


The following table lists the configurable properties for node initialization.

Properties	Description
Initialization script	Downloaded and executed at deployment time. See Deployment Lifecycle Scripts > Lifecycle Action Script Definition
Cleanup script	Downloaded and executed at the time of deployment termination. See Deployment Lifecycle Scripts > Lifecycle Action Script Definition
Resume Script	Downloaded and executed when the VM is suspended and restarted. See Deployment Lifecycle Scripts > Lifecycle Action Script Definition
Deploy Packages	External packages required for each VM are bundled and referenced as an application content package in the application profile. For example, <i>php5 sendmail</i> . See the following pages for additional details: <ul style="list-style-type: none"> Deployment Lifecycle Scripts > Application Content Package Create Package Store on Premises Service Lifecycle Actions
Sudo command list	Specify a list of semicolon-separated commands to allow additional SUDO access to certain scripts if using the Linux OS. For example, <i>SUDO ALL</i> (to allow the user to run at the root level).

This section defines the scripts to be run before (pre) and after (post) the start of the Service or pre/post stopping of the service when the application is rebooted, suspended, or terminated.

Once the node initialization is complete, CloudCenter downloads the Bundle Store to the VM and unzips the file in the specified location. See [Deployment Lifecycle Scripts](#) for additional details.

 You must configure the information explained in this section to ensure that the application is functioning.


 *Service* initialization actions at the application tier level are not called automatically. You must call them explicitly from within the service scripts. This behavior is different from other initialization actions (like external initialization and node initialization actions) that are called automatically.

The following table lists the configurable properties after node initialization is complete.


Properties	Description
Pre-Start Script	Script executed before the service is started.
Post-Start Script	Script executed after the service is started.
Pre-Stop Script	Script executed before the service is stopped.

Post-Stop Script	Script executed after the service is stopped.
------------------	---

The firewall rules that are specified as part of the application profile are used to configure the cloud provider firewall. These firewall rules can be applied to Cloud-based Security Groups or features like ACI contracts.

-  Although it is not possible to explicitly specify security policies for each tier in the application profile, any allowed security policy specified in the deployment environment form can be selectively applied to each tier at deploy time. If you want to define per tier security policies that are unique to each tier, consider defining firewall rules for each tier in the application profile instead.

If microsegmentation is enabled:

-  If you disable microsegmentation, after enabling it and configuring any change, all firewall rules revert to the initial state – any changes associated with the microsegmentation configuration are lost.


- The default firewall rule for a service are automatically displayed in this section if it is the top tier for the app or at any tier level if .
- If service-level microsegmentation is enabled, you can restrict any firewall rule to any tier by specifying the tier name or IP for the source.
- See [Security and Firewall Rules](#) for additional context.

As soon as you connect the service in the [Topology Modeler](#), the IP column in the Firewall Rules section automatically updates to display the Apache service tier name. You can choose to keep the preconfigured firewall definition or add additional Firewall rules as required.

The following table lists the configurable properties for service-level firewall rules.

Properties	Description
IP Protocol	Defines the protocol to be used by VMs running this service. <ul style="list-style-type: none"> • TCP: Transmission Control Protocol • UDP: User Datagram Protocol
From Port	The initial port number of the port range to use for the inbound firewall rule (or security rule).
To Port	The final port number of the port range to use for the inbound firewall rule (or security rule).
IP/CIDR/TIER	The default CIDR for the application tier defaults to 0.0.0.0/0.

Once the Node initialization is complete, the application is capable of functioning as configured. Once the node is up, the utility files can be sourced in the script. See [Deployment Parameters](#) for additional context.

-  Custom parameters are part of the userenv file (/usr/local/osmosix/etc/userenv). This file contains the dynamically-generated information that is used for advanced scripting and orchestration purposes. See [Deployment Lifecycle Scripts > Utility Files](#) and [Configuration Files](#) for additional context.

To add additional parameters that are specific to each step or tier, provide the parameter information listed in the following table.

Properties	Description
Parameter Name	Identifies the name for this custom parameter.
Display Name	Identifies the display details for this custom parameter.
Help Text	Identifies the help text required when using this custom parameter.
Type	See Using Parameters > Parameter Type .
Default Value	Identifies a default value, if provided, for this custom parameter.
User Option Checks	See Using Parameters > Granular Control for User-Defined Parameters .

Identifies the hardware specification for each application tier. The following table lists the properties used to filter the [instance type](#) based on your minimum hardware provisioned for your deployment.

Properties	Description
CPUs Needed	See Manage Instance Types .
Memory	See Manage Instance Types .
Network Interfaces	See IP Address Allocation .

Scratch Disk Storage (Local Storage)	See Manage Instance Types .
---------------------------------------	---

Environment variables help servers find the directory to install files, the location to store temporary files, and the settings to find the user profile(s).

This section is similar to Custom Parameters for N-Tier applications. The only difference is that you can define the type for Custom Parameters (type is not configurable for Environment Variables).

The following table lists the configurable properties for environment variables.

Properties	Description
Name	Identifies the name for this environment variable.
Value	Identifies a value for this environment variable.
User Options	See the following pages for additional context: <ul style="list-style-type: none"> Visible Editable Optional <ul style="list-style-type: none"> Using Parameters Pre-Defined Parameters

You can set migration parameters that your application may need here. The migration for the Database tier will be done automatically if the scripts are not specified. See [Deployment Lifecycle Scripts](#) for additional context.

The following table lists the configurable properties for migration scripts.

Properties	Description
Pre Migrate Script	Script is executed before the backup.
Backup Script	Backup Script writes the backup files to a Backup Folder.
Backup Location	Location (path) where the application files are backed up/restored. This is a required field for any application requiring data backup.
Restore Script	Read from the backup files in the Backup folder. The Restore Script runs after the service is started and the Post-Start Script has run, except in case of the Database tier, where the Post-Start Script runs after the Restore Script.
Post Migrate Script	Script executed after the restore is run.

You can upgrade an existing deployment by specifying the upgrade scripts for each tier. Specify the upgrade parameters in the sequence to be executed. See [Deployment Lifecycle Scripts](#) for additional context.

The following table lists the configurable properties for deployment upgrades.

Properties	Description
Upgrade Type	Define upgrade scripts for the node: <ul style="list-style-type: none"> Auto: Upgrade the tier with latest package and any backup/restore content (if specified). Advance: Allow additional scripts and steps during the upgrade process. You can specify scripts for each step separately. None: Exclude the Node/Tier from upgrading.
Pre Upgrade Script	Script is executed before the upgrade.
Post Upgrade Script	Script is executed after the upgrade.

The following table lists additional references related to application tier configuration.

Task	Reference Link
Model a new application profile	See New Application Profile .
Model an application	See New Application .
Understand the Topology Modeler	See Topology Modeler .
Understand application modeling	See Application Lifecycle Management .

Understand script lifecycle flow	See Deployment Lifecycle Scripts .
List of supported services	See OOB Services .
Add a New Service	See Custom Service Definition .
Manage Services and Scripts	See Service Administration .
Parameters	See Using Parameters .

Deployment Parameters

Deployment Parameters

- [Overview](#)
- [Understanding Service and Deployment Parameters](#)
 - [Service Parameters](#)
 - [Deployment Parameters](#)
 - [When to Use Deployment Parameters](#)
- [Adding Deployment Parameters](#)
- [Using Deployment Parameters](#)
- [Inherited Deployment Parameters](#)
- [Granular Control for User-Defined Parameters](#)
- [Deployment Parameters Icons](#)
- [Deleting Deployment Parameters](#)
- [Editing Deployment Parameters](#)
- [Defining Parameters](#)
- [Inherited Parameters](#)
- [Where Do You See Defined Parameters?](#)

Workload Manager includes an improvement to allow service parameters to be displayed to end users during deployments. Users can define parameters as deployment parameters and set preferences for them. These preferences govern if the parameters should be visible, editable, and optional during a deployment.

Service Profiles and Application Profiles use two types of parameters:

- Service Parameters
- Deployment Parameters

Service Parameters

Service Parameters are initially defined in a *Service* along with optional values. Service Parameters are not available in a deploy flow.

When you add a Service to a tier in an Application Profile, all Service Parameters and values are copied over from the defined service to the Application Profile tier.

- You can optionally change the value(s) of a service parameters during the application modeling process.
- Once Service Parameters are saved in an Application Profile, they decouple themselves from the Service Parameters in the Service.
 - Any change in the values of Service Parameter in the Service will not be reflected in the Application Profile.
 - For String, List or any other parameter type, the value will be stored in the Application Profile.
- New Service Parameters that are added in the Service are not automatically added to all Application Profiles that have this Service as a tier.
- When modifying Service Parameters in Application Profiles:
 - New Service Parameter = you must edit and save the Application Profile with the correct values.
 - Existing Service parameters:
 - Are *NOT* automatically updated based on the values in the Service Parameter (during an edit process).
 - *These parameters need to be explicitly overridden.*

Example 1:

- If the MySQL Service in your environment has a Service Parameter, called *mysql_password* with the value set to *mysql_service_password*.
- When modeling an Application Profile, you saved the Application Profile without editing the password. Now the value *mysql_service_password* is saved in the Application Profile for this parameter.
- Later, you edit the *mysql_password* Service Parameter in the Service to a new password (for example, *new_mysql_service_password*). Now, the older password *mysql_service_password*, **is still retained and used** in the Application Profile.

Example 2:

- This example uses the List type for a Service Parameter, called *number_of_instances* defined in a Service. The initial values for this List is defined as **1** and **2**, with **2** being the default value.
- When modeling an Application Profile, you saved the Application Profile with that Service using the value as **2** for this Service Parameter.
- Later, the Service was modified to add Values **3** and **4** for the *number_of_instances* Service Parameter. However, the value of **2 will continue to be used in newer deployments** of the existing Application Profile, until you manually edit the Application Profile and select one of the newly added values (**3** or **4**) from the dropdown and save it again.

Deployment Parameters

Deployment Parameters can be defined in either a Service or an Application Profile. Deployment Parameters defined in a Service are *inherited* by the Application Profile that uses the Service as the tier.

- The value for a Deployment Parameter:
 - Is *inherited* from a Service or *created* in an Application Profile.
 - Can be modified during the Application Modeling process and at Deployment time, based on the visibility preference set in the parameter.
- If you define a Deployment Parameter in a Service and use the Service in an Application Profile, and later edit the Service with the parameter value, then:
 - The deployment parameter is automatically updated for the existing saved applications.
 - The default values for the parameter type of either string, type, or number is picked from the service until it is modified and saved in the application.
- The **exception** is when a user explicitly overrides the value of the *inherited* Deployment Parameter in the Application Profile. In this case, manually overridden values are not updated when their value is changed in the Service.
- New Deployment Parameters added in a Service are **automatically** added to all Application Profiles that use the Service as a tier.

Example 1:

- If the MySQL Service in your environment has a Deployment Parameter called, *mysql_password* with a value of *mysql_service_password*.
- Two scenarios:
 - When modeling the application that uses the MySQL Service, you saved the Application Profile without editing the value for the *mysql_password* parameter. At a later time, the value of this parameter in the Service was changed to *new_mysql_service_password*. Immediately, the new value is **automatically** updated in all Application Profiles that use this Service as a tier. You do not need to explicitly edit the Application Profile and change it.
 - When modeling the application that uses the MySQL Service, a user changed the value for the *mysql_password* in the Application Profile to *new_app_level_mysql_password*. At a later time, the value of the parameter is changed to *new_mysql_service_password* in the Service. The specific Application Profile and deployments initiated from that Application Profile **will continue to use the overridden parameter** *new_app_level_mysql_password*.

Example 2:

- This example uses the List type for a Deployment Parameter, called *number_of_instances* defined in a Service. The initial values for this List is defined as 1 and 2, with 2 being the default value.
- When modeling an Application Profile, you saved the Application Profile with that Service using the value as 2 for this Deployment Parameter.
- Later, the Service was modified to add Values 3 (default) and 4 for the *number_of_instances* Service Parameter. Existing applications will pass the value of 3 to new deployments.
- However if the value was initially overridden to 1 in the Application Profile and saved, then after changing the value in the Service to 3 & 4, 1 is passed to the deployment.

When to Use Deployment Parameters

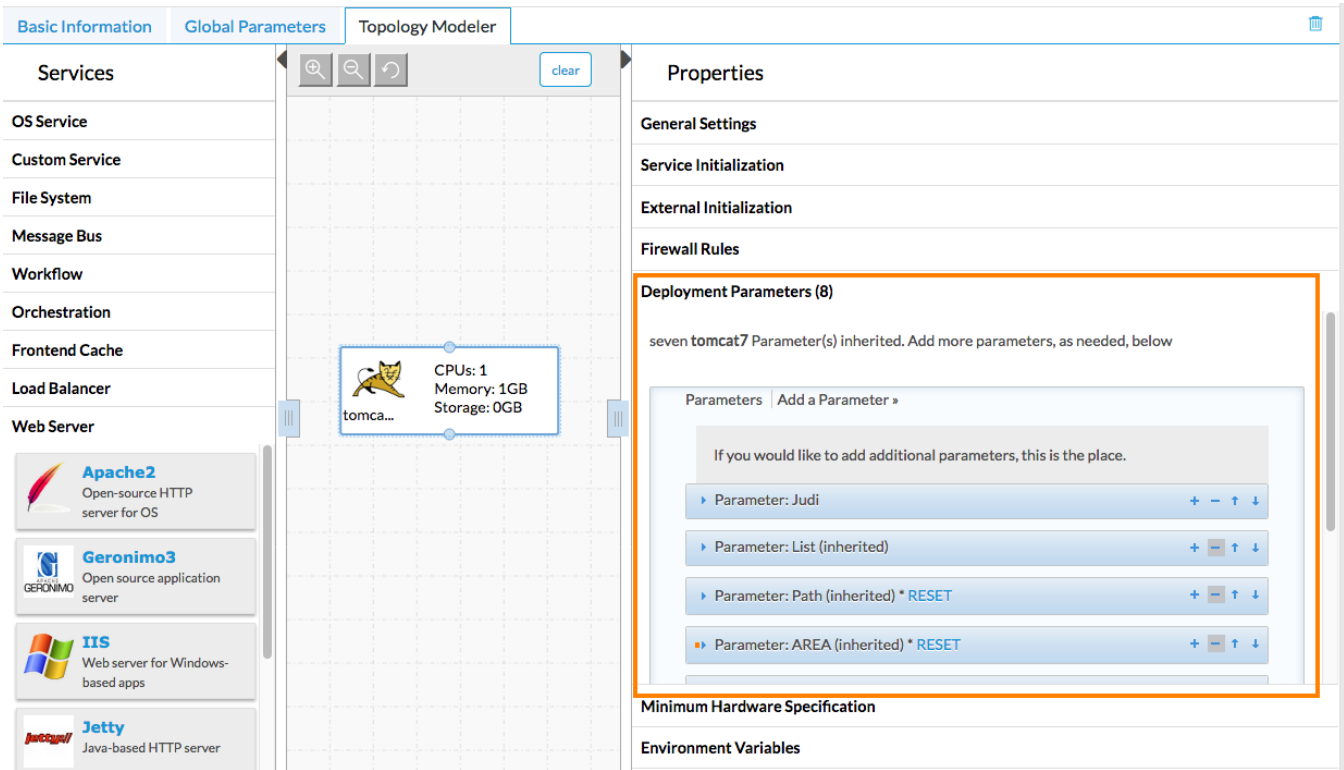
- In your environment, you may sometimes use Service Parameters instead of using Deployment Parameters.
- If you use a Service Parameter, called *SAMPLE_AB_Deployment* where you use a different set of values for each Application Profiles – (where you set **No** for Test Environments and **Yes** for Production Environments). In this case:
 - **Yes** = the old and incorrect value.
 - The Service Parameter value **does not change in any of the Application Profiles** – it continues to use the older value.
 - If you remove the Service and again add it, it will update to the new Service Parameter value.
- To avoid this issue when you require the *inherit* behavior, consider using Deployment Parameters instead of Service Parameters.
- In this example:
 - Define *SAMPLE_AB_Deployment* as a **Deployment Parameter**.
 - Do not override the value for this parameter in the Application Profile – by doing so, when you flip the value in the Service it will automatically be flipped in Application Profiles as well.

Deployment Parameters are intended to be customized by the user at deployment time and can be created at the service level and used as inherited parameters within an application or directly created when modeling an application.

- **Service-Level Parameters:** The parameter is specific to service. If you define a deployment parameter at the service level, then it is displayed as an *inherited* parameter in the tier where the service is used while modeling an application. You can customize the visibility preferences for this parameter and override the default value when modeling an application. The following screenshot shows a service-level parameter.
- **Tier-Level Parameters:** The parameter is specific to each step or tier within an application. You can define a deployment parameter at the tier level when modeling an application. You can control the visibility preferences for this parameter during the application modeling process and configure them to be visible, editable, or options during the deployment process.

See [Using Parameters](#) for additional details.


If a parameter is defined at the service level, then this parameter displays (inherited) when viewed at the tier level. The following screenshot shows inherited deployment parameters.



These parameters are made available to users that use this service when deploying applications.

When users model applications, these parameters are visible to those using this service. If used at application modeling time, you can only edit the default values for a deployment parameter.

This section explains the options that allow granular control for deployment parameters. The following image highlights the Use Defaults switch.



This switch is present in the parameters editor dialog box to the right of the value of each parameter which is visible and not locked.

- **On:** Default. Workload Manager uses the default value specified in the parameter definition.
- **Off:** If the field is not locked and you type a value for the parameter, the toggle is automatically turned off.

The **User Options** that the following table describes provide granular control over user-defined parameters:

User Option	Checked	Unchecked
Should this parameter be visible to the user?	This parameter will be visible to users during the application deployment. Users can override the visibility setting during the application modeling process.	This parameter is not visible to users during the application deployment.
Should this parameter be editable by the user?	This parameter will be editable by users during the application deployment. Users can override the edit setting during the application modeling process.	This service becomes automatically is disabled for users, and therefore not editable during the application deployment.

Should this parameter be optional ?	Marked optional – This parameter will be optional for users during the application deployment. Users can override the optional setting during the application modeling process.	Marked required – This parameter is required and the user must provide or select a value during the application deployment.
Allow selection of multiple values (only displayed if a "list" parameter is selected)	Users can select multiple values for this parameter during the application deployment.	Users can select one value for this parameter during the application deployment.

Note that the header for each parameter section has the following icons on the right side:

- + adds a new resource – at the required point
- - deletes an existing parameter
- moves the priority of the resource further up the list.
- moves the priority of the resource further down the list.

If a deployment parameter is deleted at the service level, all usage for this parameter is impacted – this deployment parameter will not be available when:

- You edit an application containing that service.
- Users deploy an application containing that service.

If the default value for a deployment parameter is changed at the service level, then this change reflects in the usage for this parameter at the application level IF the default value was not overridden during the application modeling process.

You can define deployment parameters in new and existing services:

- If an app is not modified, these parameters show on the deployment flow based on the visibility.
- At the app level, deployment parameters are shown from the service.
- You cannot remove deployment parameters at the app level.
- You can edit the default values or visibility setting; app-level values take precedence
- If you edit the default values, the values are stored only in the app.
- If you remove an updated parameter from service, it is not shown in the app or job.
- If you change parameter values to the same values of service, it is not saved at the app level and it is picked up at the service level.

You can define deployment parameters on the Service page. To do so, click **Edit** for the service and scroll down to the Deployment Parameters area. By default, parameters are collapsed. You can click any parameter to expand it and view the parameter's configuration.

When you model an app and drag a service in the Topology Modeler page, the Properties display shows how many deployment parameters exist. Click the **Deployment Parameters** field to see the name of each parameter. You can expand each parameter further and may sometimes see that a parameter has been assigned an *inherited* designation.

The *inherited* designation next to a parameter means that the parameter is derived from the Services page and is not defined within an application deployment.

Here are some general guidelines when working with inherited parameters:

- You cannot define a new parameter that has the same name as an inherited parameter.
- You cannot delete them in the Service page (the Minus icon button is disabled).
- You cannot update the Parameter Name, Display Name, Help Text, or Type fields (these fields are disabled).
- You can update the Default Value and User Options.

On the Deploy page, in the Tier Settings area, you can see each deployment parameter that is defined. If a value is defined as editable for the parameter, you can edit it here, while deploying the app.

On the Deployment Details page, you can see the parameters for a deployment under Parameters (applies to global, custom, and deployment parameters).

Deployment Lifecycle Scripts


Deployment Lifecycle Scripts

- [Script Workflow](#)
 - [Node Initialization Workflow](#)
 - [External Initialization Workflow for VM Tiers](#)
 - [External Initialization Workflow](#)
 - [Node Reboot Workflow](#)
 - [Node Update Workflow](#)
 - [Node Terminate Workflow](#)
- [Script Categories](#)
- [Using Scripts in Application Profiles](#)
- [Script Source Details](#)
- [Lifecycle Action Script Definition](#)
- [Utility Files](#)
 - [Accessing Files from a Configured Repo](#)
- [Parameters and Configuration Files](#)
- [Internal Reboot During Node Initialization](#)
- [The #!CliQrReboot: Header](#)

Scripts can be called at multiple points during a deployment. The following images depict the execution process for each service workflow:


Node Initialization Workflow

The following image shows the node initialization workflow.



External Initialization Workflow for VM Tiers

See [External Service](#) for additional context. The following image shows the external initialization workflow for VM tiers.






External Initialization Workflow

For non-VM tiers like application tiers. See [External Service](#) for additional context. The following image shows the external initialization work



Node Reboot Workflow

Triggered from either the VM Reboot or Suspend/Resume action in the Deployments page. The Suspend script is similar to the Resume script that is triggered when the node is suspended. The following image shows the node reboot workflow.



Node Update Workflow

Triggered by a VM scaling action or a reboot on other VMs. The following image shows the node update workflow.



Node Terminate Workflow

Requires all nodes to be running in order for the application to be powered off. The following image shows the node terminate workflow.

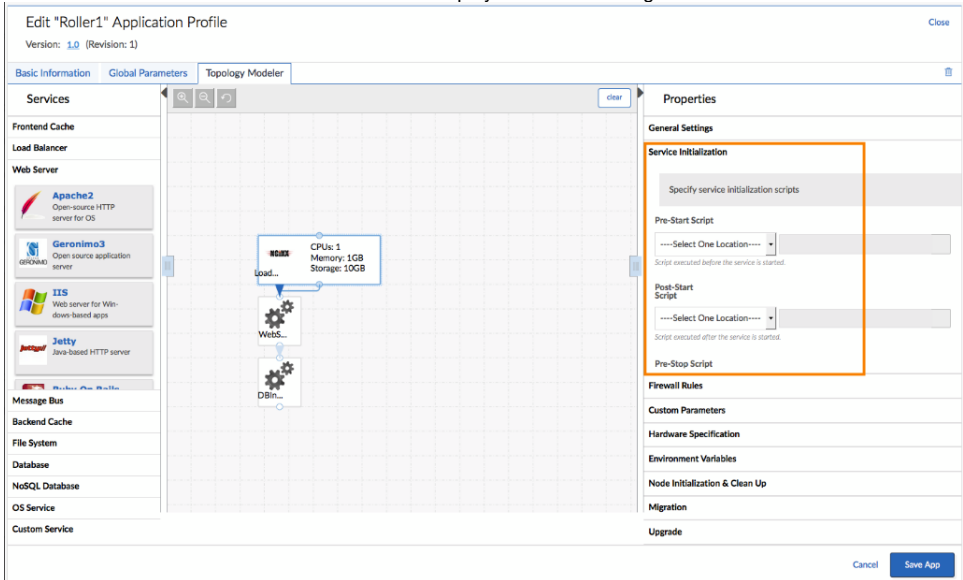


Service lifecycle actions can be divided into two categories:

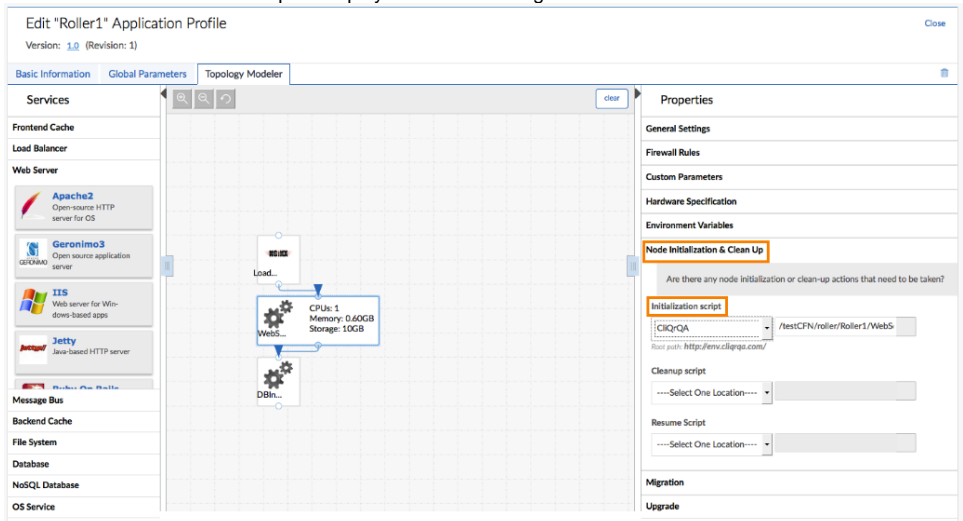
- **Service Scripts:** Each service script corresponds to a [Service Lifecycle Actions](#). See [Service Administration](#) for additional context. Once the root admin configures the services and makes it available to users, users can access the service and add on their own application scripts to each service or tier as required.
- **Application Scripts:** Users specify each application script when modeling an application or application profile. See [New Application Profile](#) for more details on these scripts.

While the [Topology Modeler](#) > [Services](#) tab (or palette) allows the root admin to add on the required services for each user or tenant, the [Topology Modeler](#) > [Properties](#) tab allows users to initialize, clean up, or resume the app tier. The following images highlight the application script configuration location in the [Topology Modeler's Properties](#) tab.

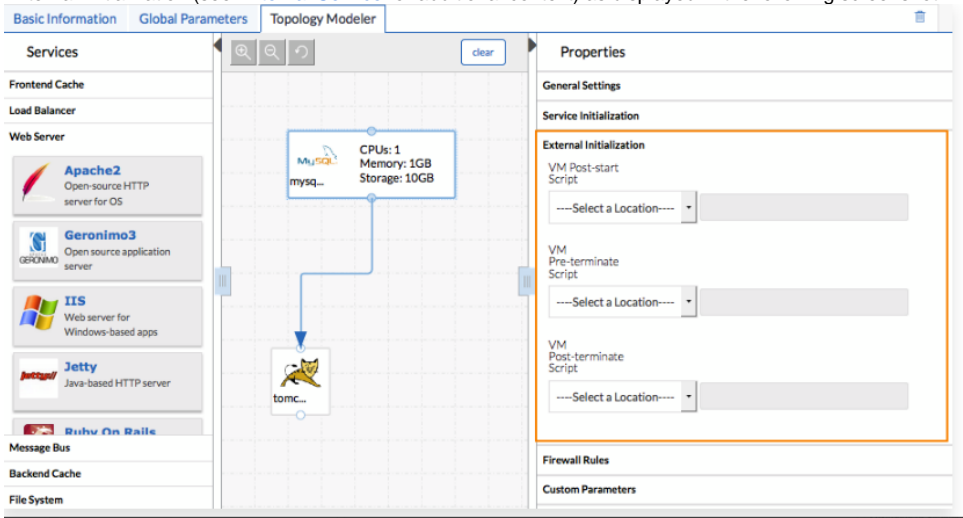
- Service Initialization Pre-Start and Post-Start as displayed in the following screenshot:



- Node Initialization and Clean Up as displayed in the following screenshot:





- External Initialization (see [External Service](#) for additional context) as displayed in the following screenshot:



- Similarly, you will also find the Migration and Upgrade in the configuration location in the Topology Modeler's Properties tab.


The different script sources are explained in the following table.

Sourcing Methods	Description	Additional Details
Repositories	Points to a repository hosted in the external site (such as HTTP or Amazon S3).	See Artifact Repository
File in package	The script resides in the Application Content Package.	See Application Using App Package
URL or Command	<p>A URL Points to a file that is downloaded and executed by the Management Agent. This URL can point to any location relevant to this service.</p> <p>A command (shell or powershell direct commands) is used to run all service-related actions. For example: <i>apt-get install foo</i>, or <i>echo hello</i>, or <i>rerun cliqr</i>, and so forth.</p> <div style="border: 1px solid #f0e68c; padding: 5px; margin-top: 10px;">  The script option does not work when using the command line. Instead, use the URL option and add a script URL. </div>	See Parameters and Macros
Script from bundle (for Service Lifecycle Actions)	<p>The script path is a relative path inside the extracted folder (<i>/usr/local/osmosix/service</i>).</p> <div style="border: 1px solid #f0e68c; padding: 5px; margin-top: 10px;">  The name of the bundle directory must match the name displayed in the Identifier field. This step is essential – the service <i>cannot</i> be created if the name or the relative path differs. </div> <p>This action is only available if the service points to a bundle Location.</p>	See Local Bundle Store (Conditional)

This section explains the level and details required to define and use each script.

See the following sections for additional information on the lifecycle action script definition at each level:

- Region Level: See [External Lifecycle Actions Settings](#).
- Service Level: See [External Service > External Initialization Scripts](#).
- Application Tier Level: See [Service Lifecycle Actions](#).

 All scripts are executed from the current script directory and allow multiple scripts to be issued at the same time.

 **Script Download Owner Permission**

The download owner permission for all scripts at all levels is **root** (-rwxr-xr-x).

The following table identifies the level and details required to define and use each script.

Script	Level of Script Definition	Script Download Location	User Running Script	Script Run Location
Pre-VM Start	Region	External script execution engine	cliqruser	Same as <i>Script Download Location</i>
Pre-VM Init				
Post-VM Init				
Pre-VM Stop				
Post-VM Stop				
Install	Service	<ul style="list-style-type: none"> • <i>/usr/local/osmosix/service/{serviceName}/</i> • <i>C:\program files\osmosix\service\{serviceName}\</i> 	root	Same as <i>Script Download Location</i>
Configure				
Deploy				
Start				
Stop				
Restart				
Reload				
Upgrade				
Clean Up				

Pre-VM Start	Service (external)	External script execution engine	cliqruser	Same as <i>Script Download Location</i>	
Pre-VM Init					
Post-VM Init					
Pre-VM Stop					
Post-VM Stop					
Initialization Script	Application	<ul style="list-style-type: none"> • /opt/remoteFiles/initScript/ • C:\temp\remoteFiles\initScript\ 	Agent user (typically cliqr user)	Same as <i>Script Download Location</i>	
Cleanup Script			<ul style="list-style-type: none"> • /opt/remoteFiles/cleanupScript/ • C:\temp\remoteFiles\cleanupScript\ 		cliqruser
<ul style="list-style-type: none"> • Resume Script • Suspend Script 			<ul style="list-style-type: none"> • /opt/remoteFiles/resumeScript/ • C:\temp\remoteFiles\resumeScript\ 		Agent user (typically cliqr user)
Pre-Start Script	Application (Service Initialization)	<ul style="list-style-type: none"> • /opt/remoteFiles/cliqr{serviceName} PreStartAction/ • C:\temp\remoteFiles\cliqr{serviceName} PreStartAction\ 	cliqruser	/usr/local/jetty/	
Post-Start Script					
Pre-Stop Script					
Post-Stop Script					
VM Pre-provision Script	Application (External Initialization)	External script execution engine	cliqruser	Same as <i>Script Download Location</i>	
VM Pre-initialization Script					
VM Post-start Script					
VM Pre-terminate Script					
VM Post-terminate Script					
Pre Migrate Script	Application (migration and upgrade)	<ul style="list-style-type: none"> • /opt/remoteFiles/preMigrateScript/ • C:\temp\remoteFiles\preMigrateScript\ 	Agent user (typically cliqr user)	/home/cliqruser/	
Backup Script					
Restore Script		<ul style="list-style-type: none"> • /opt/remoteFiles/restoreScript/ • C:\temp\remoteFiles\restoreScript\ 	cliqruser	/usr/local/jetty/	
Post Migrate Script		<ul style="list-style-type: none"> • /opt/remoteFiles/postMigrateScript/ • C:\temp\remoteFiles\postMigrateScript\ 			
Pre Upgrade Script		<ul style="list-style-type: none"> • /opt/remoteFiles/upgradeScript/ • C:\temp\remoteFiles\upgradeScript\ 			
Upgrade Script					

Post Upgrade Script			
Rollback Script			



The dynamically-generated VM password is injected as part of the lifecycle action and made available for use inside External scripts.

- **Windows deployments:** The password for **cliqr** users is available in the `cliqrWindowsPassword` environment variable in the pre-Init, post-start, pre-terminate, and post-terminate phases.
- **Linux deployments:**
 - The user name is available in the `sshUserName` environment variable
 - The SSH private key is available in the `sshKey` environment variable.
 - The SSH public key is available in the `sshPublicKey` environment variable

Users can also include utility files in scripts.

The following are some examples of utility files that can be sourced by Linux users:

- Environment Variables: `/usr/local/osmosix/etc/userenv`
- OS Information: `/usr/local/osmosix/service/utills/os_info_util.sh`
- Service Install Utility: `/usr/local/osmosix/service/utills/install_util.sh`
- Configuration Utility: `/usr/local/osmosix/service/utills/cfgutil.sh`
- Request Utility: `/usr/local/osmosix/etc/request_util.sh` (send request to agent to download a file or folder from the repo)

The following are utility files that can be sourced by Windows users:

- Environment Variables: `c:\temp\userenv.ps1`
- Utility Information: `c:\Program Files\osmosix\etc\cliqr.ps1`
- Request Utility: `c:\Program Files\osmosix\etc\request_util.ps1`



The Request Utility, for both Linux and Windows requires that the repo server you refer to use the HTTP or HTTPS protocol.

In earlier releases, users were able to use special characters in the `userenv` file. However, sourcing files with these special characters caused issues in scripts that used the `userenv` file.

- Special characters like pipe or double quotes are configured with escape sequences that use single quotes, *not* double quotes.
- An example:

```
export nameOfVar = 'ExampleABC'
```

- You can successfully source `userenv` files that contain these special characters.

Accessing Files from a Configured Repo

While you can use the **File in Package** option to model an [Application Using App Package](#), this option bundles the required files into a zip file that you can download.

Alternately, you may also opt to download something from inside a script that is stored in a configured repository. For example, you have configured a file in a repository called *MyFiles* for your lifecycle scripts using the HTTP protocol and have it password protected, as well as configured a `nodeInit` script called *myscript* as follows:

```
Repo: MyFiles, path:/myscript.sh
```

Later, you may need another user to download a file from the same *MyFiles* repo without knowing the credentials or the contents for either agent-based and external actions.

The following table provides usage details on accessing a file from a configured repository.

Usage	Syntax	Example
download \$<relative folder path on repo> <node service step>	download <path_on_repo>/<filename.txt> <initScript/ prestart/ poststart/ prestop/ poststop>	download auto/record_results.sh initScript The file is downloaded to /opt /remoteFiles
backupFile \$<relative file path on repo> \$<location of local file to be backup >	backupFile <path_backup_folder_on_repo> /<backup_filename.txt> /tmp/test_backup.txt	backupFile \$CliqrDeploymentId /dbbackup.sql /tmp/dbbackup.sql

restoreFile \$<relative file path on repo>	restoreFile <path_backup_folder_on_repo> /<backup_filename.txt>	restoreFile \$migrateFromDepld /wpbkup.zip
--	--	---

[Pre-Defined Parameters](#) are also available for use in application profiles to automate jobs without writing extensive scripts (examples include timestamp, dynamically-generated IP address, private IP address, IP address of a tier, environment variables, automation policy parameters, number of nodes in a cluster, deployment name, and so forth).

Sometimes, you may have a custom parameter already defined in a particular service that you do not want to set in the application profile. Instead, you want to leave it to your end-users to [Substitute a Parameter](#) or use [Troubleshooting Parameters](#) when they model application profiles or deploy applications.

If you use parameters (either Workload Manager-Defined Parameters or Parameter Substitution), you may need to modify the service [Configuration Files](#) to reflect the correct property defined in the relevant parameter.

The reboot referenced above is initiated externally when a user reboots the VM, or issues a Suspend/Resume command, or a direct OS reboot command. This reboot usually happens after the node is fully initialized and running.

Another reboot scenario is the internal reboot – when the reboot is part of node initialization process. This reboot is initiated by the management agent. Once the agent detects the .cliqrRebootResumelnit flag, it performs a backup and reboots itself.

You can setup multiple stage node init scripts by using appropriate flow-control logic and manipulating the following files:

- **/tmp/.cliqrRebootResumelnit**
- **\$OSMOSIX_PROD_HOME/.cliqrRebootResumelnit**



\$OSMOSIX_PROD_HOME defaults to /usr/local/osmosix



The user scripts should use SUDO to delete/create any file in OSMOSIX_PROD_HOME – in particular, the \$OSMOSIX_PROD_HOME/.cliqrRebootResumelnit file.

Workload Manager executes each pass in order, picking up where it left off each time. Consider the following sample scripts:

- Linux

```
#!/bin/bash
# Make sure to source these files to pick up all the Workload Manager environment variables.
. /usr/local/osmosix/etc/.osmosix.sh
. /usr/local/osmosix/etc/userenv

# After triggered reboot, the agent generates file $OSMOSIX_PROD_HOME/.cliqrRebootResumeInit
# with prior state from /tmp/.cliqrRebootResumeInit

# If this file DOES NOT exist, it means that this is the first pass through the script.
# The agent has not previously triggered a reboot and created this file.
if [ ! -e $OSMOSIX_PROD_HOME/.cliqrRebootResumeInit ];
then
    # Node was not rebooted already by nodeInit. First pass through this script.

    # Triggers a reboot after the script exits, then this file is deleted.
    # "Step 2" will now appear in $OSMOSIX_PROD_HOME/.cliqrRebootResumeInit
    echo "Part 2" > /tmp/.cliqrRebootResumeInit
else
    # Node was rebooted on the last pass through the script.
    # Read in prior state from file generated by Workload Manager
    step=`cat $OSMOSIX_PROD_HOME/.cliqrRebootResumeInit`

    # Check the prior state to continue with the next part of the script.
    case $step in
        "Part 2")
            # Triggers a reboot after the script exits, then this file is deleted.
            # "Step 3" will now appear in $OSMOSIX_PROD_HOME/.cliqrRebootResumeInit
            echo "Part 3" > /tmp/.cliqrRebootResumeInit
            ;;
        "Part 3")
            # Don't write to /tmp/.cliqrRebootResumeInit this time. The Workload Manager platform will
            not reboot.
            # Script exists, and node init is complete.
            ;;
    esac
fi
```

- Windows

```
#Set the "numOfReboots" variable to reflect the number of reboots required. In #this example the VM will
be rebooted 5 times.

$startTime = Get-Date
$VerbosePreference = "Continue"
$numOfReboots = 5
$logPath = "C:\nodeinit"
$logFile = "cliqr-nodeinitlog.txt"
$envscript = ' c:\temp\userenv.ps1'
$agent_utils = ". 'C:\Program Files\osmosix\service\utils\agent_util.ps1'"
$tmpfolder = "C:\tmp"
$logfolderFlag = 1
$tmpFolderFlag = 1
$logFileCreationFlag = 1

Invoke-Expression $envscript
Invoke-Expression $agent_utils

if (!(Test-Path $logPath)){
    $logDir = New-Item -ItemType Dir $logPath
} else {
    $logfolderFlag = 0
}
```

```

if (!(Test-Path $tmpfolder)){
    $tmpDir = New-Item -ItemType Dir $tmpfolder
} else {
    $tmpDir = $tmpfolder
    $tmpFolderFlag = 0
}

$logFilePath = Join-Path $logPath $logFile
if (!(Test-Path $logFilePath)){
    $logFile = New-Item -ItemType File (Join-Path $logPath $logFile)
} else {
    $logFile = $logFilePath
    $logFileCreationFlag = 0
}

function print($line,$logFileHandle) {
    $timestamp = get-date
    $line = [string]$timestamp + " : INFO " + $line
    Write-Output $line >> $logFileHandle
    Write-Verbose $line
    agentSendLogMessage $line
}

#####MAIN #####

$resumeinitfile1 = Join-Path (Get-ChildItem env:OSMOSIX_HOME).Value ".cliqrRebootResumeInit"
$resumeinitfile2 = Join-Path $tmpDir ".cliqrRebootResumeInit"

if (!(Test-Path $resumeinitfile1)) {
    #Checking for existence of file for first reboot. If not then create the flag file
    print "File .cliqrRebootResumeInit not created yet" $logFile
    Write-Output "RC=1" > $resumeinitfile2
    print "Restart Count = 1" $logFile
} else {
    #For subsequent reboots flag file will exist in a different location
    print "File '$resumeinitfile1' Found" $logFile
    $data = gc $resumeinitfile1
    $srcCount = $data.trim().split("=")[-1]
    print "Data from $resumeinitfile1 file => $data" $logFile

    if ([int]$srcCount -lt $numOfReboots) {
        $srcCount = [int]$srcCount + 1
        $strdata = "RC=" + $srcCount
        print "Writing data '$strdata' to file '$resumeinitfile2'" $logFile
        Write-Output $strdata > $resumeinitfile2
    }
    print "Restart Count = $srcCount" $logFile
}

$VerbosePreference = "SilentlyContinue"

```

The **#!CliQrReboot:** header to the /tmp/.cliqrRebootResumeInit file – Once the agent detects the **#!CliQrReboot:** header, it resumes after a reboot in the service lifecycle flow depending upon the context set in the header. You can control the resume flow after reboot by adding the resume context header to the /tmp/.cliqrRebootResumeInit file.

For example, **#!CliQrReboot:Current** as a header allows you to resume the current lifecycle action. Similarly, **#!CliQrReboot:Next** to resume to next step in the lifecycle actions and **#!CliQrReboot:Deploy** to resume from deploy service lifecycle actions. Consider the following sample scripts:

- The following sample Linux script reboots twice and resumes the same lifecycle action after the reboot. The resume context is only supported in a deployment's deploy flow. It is not supported during the deployment suspend, resume or terminate flows.

Linux Sample Script

```
#!/bin/bash

. /usr/local/osmosix/service/utils/agent_util.sh
. /usr/local/osmosix/etc/userenv
. /usr/local/osmosix/etc/.osmosix.sh

action=$1

if [ $2 == "reboot" ]; then
  if [ ! -e $OSMOSIX_PROD_HOME/.cliqrRebootResumeInit ]; then
    action="$action-step-1"
    agentSendLogMessage "Calling First Reboot... $action"
    echo "#!CliQrReboot:Current" > /tmp/.cliqrRebootResumeInit
    echo "2" > ~/step
  else
    step=`cat ~/step`
    case $step in
      "2")
        action="$action-step-2"
        agentSendLogMessage "Calling Second Reboot... $action"
        echo "#!CliQrReboot:Current" > /tmp/.cliqrRebootResumeInit
        echo "3" > ~/step
        ;;
      *)
        action="$action-step-3"
        agentSendLogMessage "Resuming after Final Reboot... $action"
        rm -fr ~/step
        rm -fr $OSMOSIX_PROD_HOME/.cliqrRebootResumeInit
        ;;
    esac
  fi
fi
```

- The following Windows sample script takes two parameters (*action* and *reboot*):



To use the **#!CliQrReboot:** header, you must use ASCII encoding – otherwise, the agent will not be able to pickup the required directives.

However, if you are not using the **#!CliQrReboot:** header, you can use other encoding as well.

Windows Sample Script

```

param (
    [string]$action,
    [string]$reboot
)

# Source userenv to get all cloud/app/job/tier data
. 'c:\temp\userenv.ps1'
. 'C:\Program Files\osmosix\service\utils\agent_util.ps1'

# Verify reboots and capture results with every reboot
$numOfReboots = 2
$tmpfolder = "C:\tmp"

if (!(Test-Path $tmpfolder)){
    $tmpDir = New-Item -ItemType Dir $tmpfolder
} else {
    $tmpDir = $tmpfolder
    $tmpFolderFlag = 0
}

$resumeinitfile1 = Join-Path (Get-ChildItem env:OSMOSIX_HOME).Value ".cliqrRebootResumeInit"
$resumeinitfile2 = Join-Path $tmpDir ".cliqrRebootResumeInit"
$resumeinitfile3 = Join-Path $tmpDir ".step"

if ( $reboot -match "reboot") {
    if (!(Test-Path $resumeinitfile1)) {
        #Checking for existence of file for first reboot. If not then create the flag file
        Set-Content -Path $resumeinitfile2 -Value "#!CliQrReboot:Current"
        Set-Content -Path $resumeinitfile3 -Value "RC=1"
        agentSendLogMessage "Calling 1 Reboot... $action"
    } else {
        #For subsequent reboots flag file will exist in a different location
        $data = gc $resumeinitfile3
        $rcCount = $data.trim().split("=")[-1]

        if ([int]$rcCount -lt $numOfReboots) {
            $rcCount = [int]$rcCount + 1
            $strdata = "RC=" + $rcCount
            agentSendLogMessage "Calling $rcCount Reboot... $action"
            Set-Content -Path $resumeinitfile2 -Value "#!CliQrReboot:Current"
            Set-Content -Path $resumeinitfile3 -Value $strdata
        } else {
            agentSendLogMessage "Resuming after Final Reboot... $action"
            Remove-Item $resumeinitfile1
            Remove-Item $resumeinitfile3
        }
    }
}
}

```

Multiple Volumes

Attach Multiple Volumes to Tiers

- [Overview](#)
- [Volume Attachment Behavior](#)
- [Configurable Volume Attributes](#)
 - [Size](#)
 - [Type](#)
 - [AzureRM Type Nuances](#)
 - [AWS Type Nuances](#)
 - [Google Type Nuances](#)
 - [OpenStack Type Nuances](#)
 - [IOPS](#)
- [UI Configuration Options](#)
- [Container-Specific Configuration Options](#)
- [API Configuration Options](#)

You can attach multiple volumes to all tier types in N-tier applications. For each volume, you must specify the size and can optionally configure the storage type for each root disk. While the storage type is optional, Workload Manager uses the default storage type if not configured.

There is no specific out-of-box storage type for clouds that are not listed in this section as most clouds support additional volumes. You can specify these details as follows:

- The number of volumes and default volume size details in the [Understand Application Tier Properties](#) > General Settings section.
- The storage type selection and the size of root/additional volume in the [Deployment Environments](#) form.


When you restart a node or deploy an application, the appropriate volumes (defined in the [Topology Modeler](#) > Properties tab) are attached. For example, if Node 1 has Volumes V1 and V2, and Node 2 has volumes V3 and V4, then on restart, the same volume combination (number of volumes, size, and type) are attached to the nodes.

The following table describes cloud-specific details for Volume attributes.

- All clouds support the ability to specify the storage type for additional volumes.
 - Public Clouds: Use out-of-box storage types.
 - Private Clouds and Datacenters: If you configure a storage type, then Workload Manager displays the configured storage type in the dropdown list at the time of Storage Type selection during a deployment.
- Only AWS and VMware vCenter support the ability to specify the Root Disk Size – The **Root Volume Size** setting is part of the *Volumes* section – You can optionally provide the *root disk size* in this field for both these clouds.
- Most [Public Clouds](#) supports specifying root disk storage type.


Size

Size is a required attribute for Additional Volumes.

 For clouds that don't support root disk resizing, this field is greyed out.


Type

Type is an optional volume information attribute that only provides values for public cloud instances identified in this section. For other cloud instances, the default volume type is used.

 For private cloud instances, you can add a type but the addition will only apply to calculate the storage cost.

AzureRM Type Nuances

Using Workload Manager's [Managed Disks in Azure](#), you can select one of three pre-defined disk sizes during deployment.

 The **Number of Volumes** field pre-populates the default volume size as specified in the *UI Configuration Options* section below. Ensure that the **Number of Volumes** field does not display 0 (zero) if you need to assign a predefined type to an AzureRM-based deployment.

Azure cloud instances display one of the following values for the *Type* attribute:

- Premium (Managed)
- Standard (Managed) = Default for non-government regions
- Unmanaged = Default for government regions

AWS Type Nuances

AWS cloud instances display one of the following values for the *Type* attribute:

- Magnetic
- General Purpose (SSD) = Default
- Provisioned IOPS (SSD)

Google Type Nuances

Google cloud instances display one of the following values for the *Type* attribute:

- Standard Persistent Disk = Default
- SSD Persistent Disk
- Local SSD Scratch disk (the root disk cannot be used for temporary storage)

OpenStack Type Nuances

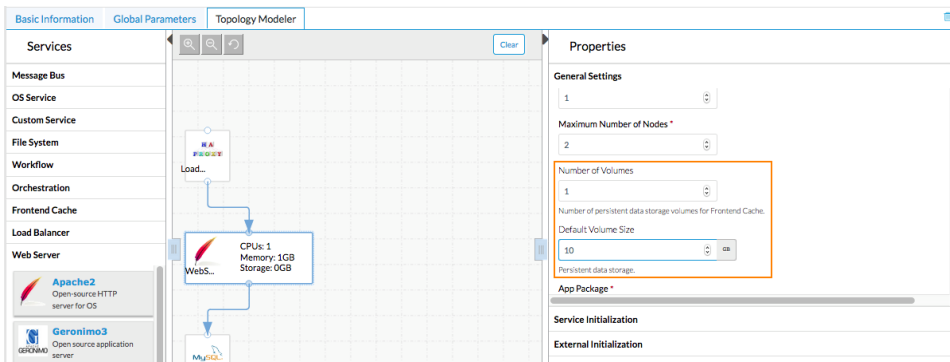
OpenStack cloud instances display the volume types configured by users in the OpenStack console.

IOPS

Input/Output Operations Per Second (*IOPS*) is an optional volume information attribute. This attribute is only applicable if the *Type* attribute displays the Provisioned IOPS (SSD) value (see the *Submit Job API*). The ratio of provisioned IOPS and the requested volume size can be a maximum of 30. That is, a volume with 3000 IOPS must be a minimum of 100 GB in size.

A Provisioned IOPS (SSD) volume can range in size from 4 GB to 16 TB and in IOPS from 100 to 20000.

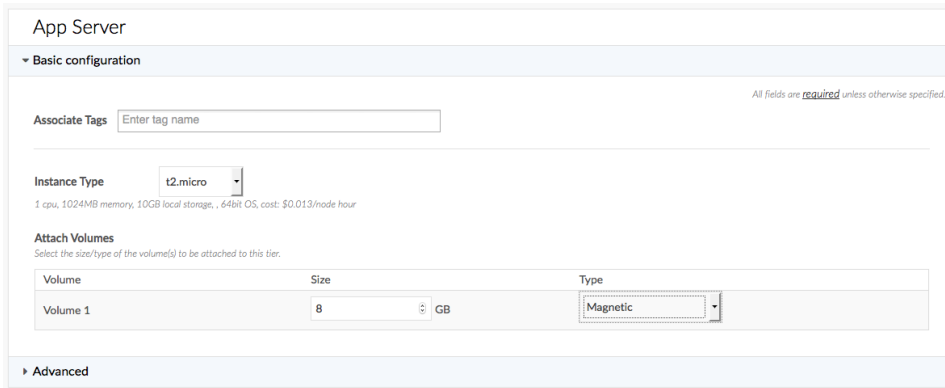
- When modeling N-tier applications, specify the Number of Volumes to be attached to each tier and the Default Volume Size for each volume in the **Topology Modeler > Properties > General Settings** field for each application tier. If the selected cloud doesn't have volume types, the default volume size is pre-populated when running this application, as shown in the following screenshot.



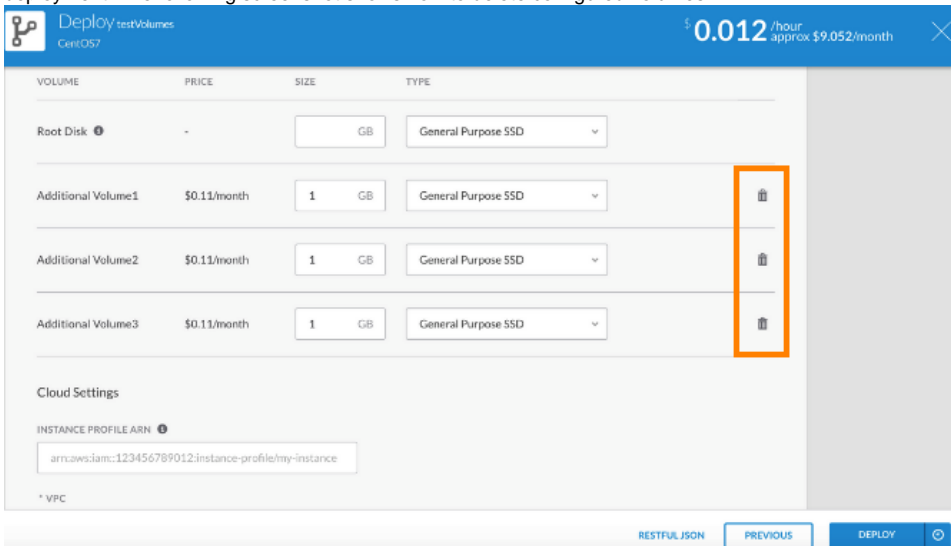
- When modeling an application, the default volume size is pre-populated in the *Size* field if the minimum size of the storage type selected is less than the default volume size for the modeled application. By default, all services do not have any persistent volume disk attached out-of-the-box. If required, you can change the default volume size at this point

The following table describes how the default volume size is pre-populated.

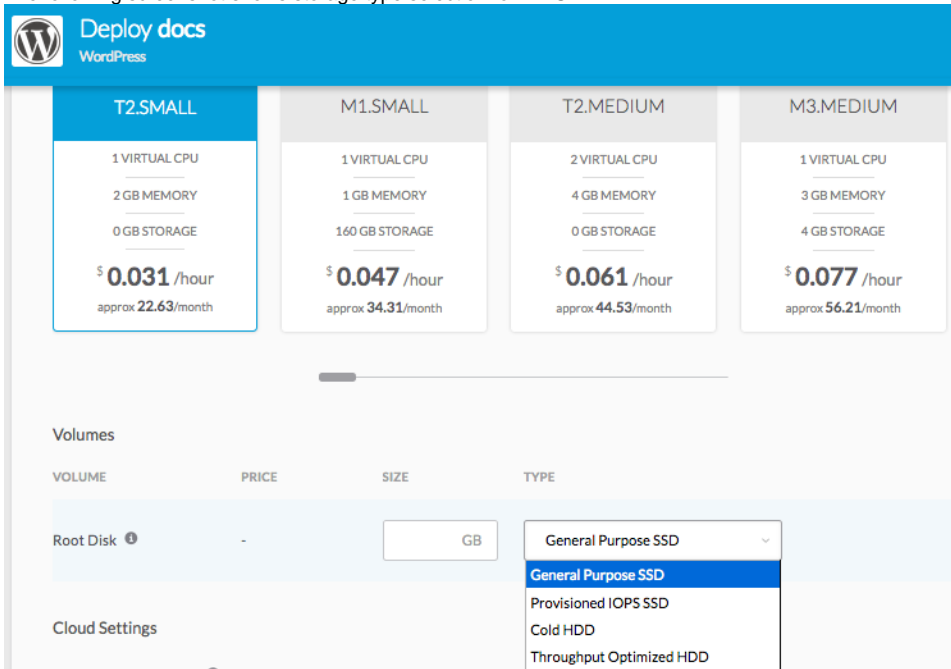
Default Volume Size from App Being Modeled	Value Pre-populated in the <i>Size</i> field
\geq minimum size of the selected storage type	The default volume size
$<$ minimum size of the selected storage type	The minimum size for the selected storage type




- You can delete the configured volumes at the time of deployment – if the preconfigured number of additional volumes are not required for your deployment. The following screenshot shows how to delete configured volumes.



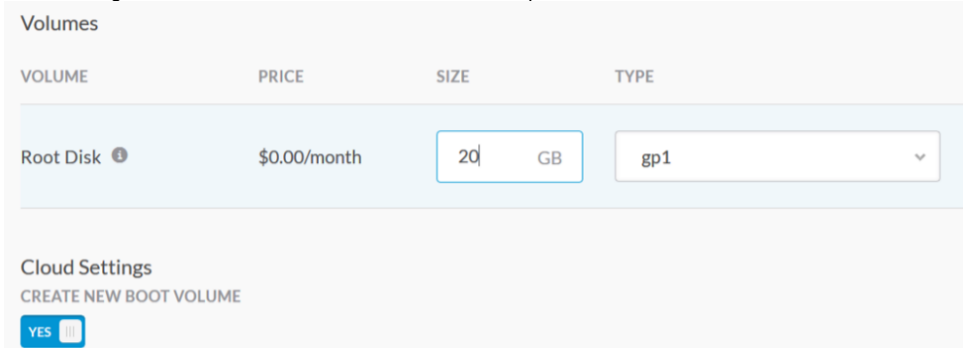
- The following screenshot shows storage type selection for AWS.



- You can configure the root volume size for AWS deployments using either the UI or the [Submit Job \(v2\) API](#).

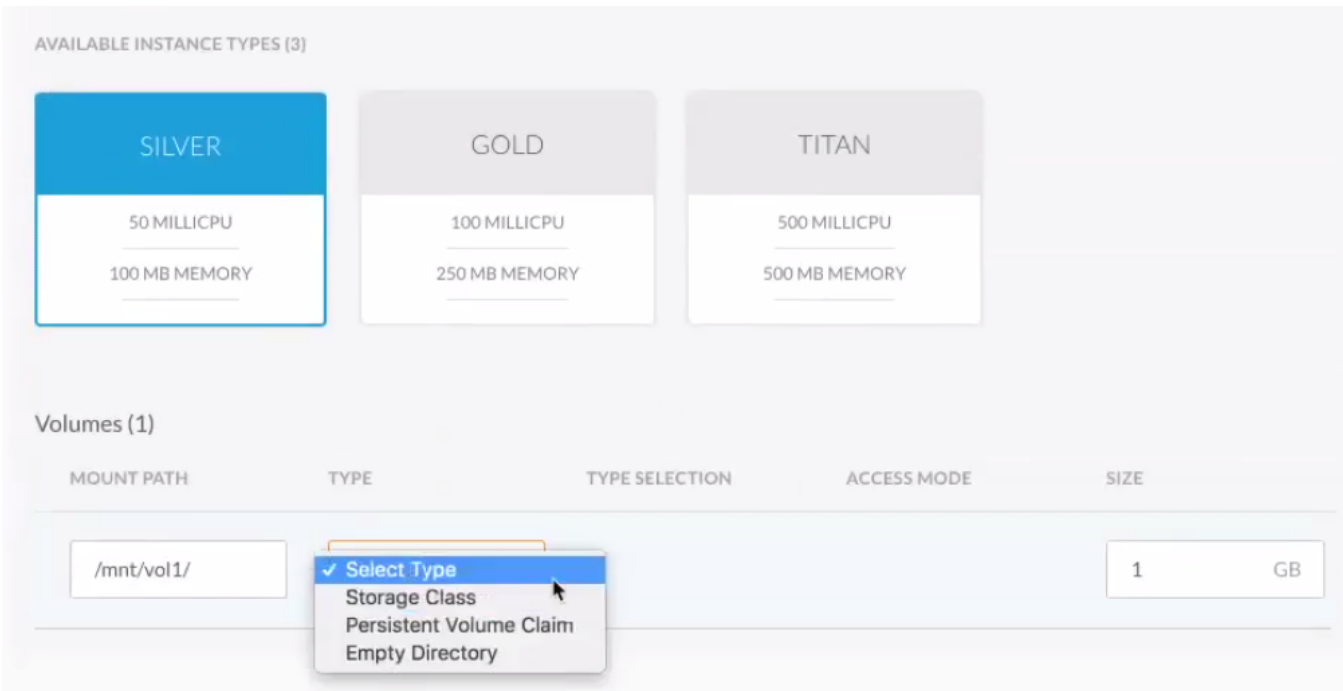
 The previous generation AWS instance types like the t1 , m1 series do not support the resizing of root volume.

- To provide a larger size for the root volume for an AWS deployment, use the Workload Manager instance type. The storage is provided in addition to the instance store and must be larger than the Root Volume size in the AMI. If set to zero, then the Root Volume size in the AMI is used.
- Root volume size = 0 (zero) value: Configure the root volume size as 0 for the older generation of instance types. This zero-configuration, by default, uses the AMI's root volume size.
- Root volume size = non-zero value:
 - The General Purpose SSD type option is automatically selected for each root volume.
 - Be sure to configure the root volume size to be higher than the size specified in the AMI but less than the size supported by AWS for General Purpose SSDs.
- The following screenshot shows the Volume selection for OpenStack.



- The *Create New Boot Volume* cloud toggle (default = No, off) in both the Environment and Deployment pages allow you to enter Root Disk Volume size. On Deploying the instance new bootable volume is created with specified size and source (instead of an image).
- The Root Disk Size setting is mandatory for OpenStack environments – this size should be greater than or equal to the minimum disk size for the image.

The following screenshot depicts the location where you can select the volume type for container deployments.



The container **Type** is a static list with the following options:

- Persistent Volume Claim – see <https://kubernetes.io/docs/concepts/storage/persistent-volumes/#persistentvolumeclaims> for additional context.
- Storage Class – see <https://kubernetes.io/docs/concepts/storage/storage-classes/> for additional context.
- Empty Dir – see <https://kubernetes.io/docs/concepts/storage/volumes/#emptydir> for additional context.

In a deployment flow, you must choose one these three options and Kubernetes mounts the volume for this container application.

The **Type Selection** is a list of available storage classes as provided by cloud administrator:

MOUNT PATH	TYPE	TYPE SELECTION	ACCESS MODE	SIZE
<input type="text" value="/mnt/vol1/"/>	<input type="text" value="Storage Class"/>	<div style="border: 1px solid #ccc; padding: 5px;"><input checked="" type="checkbox"/> Select default gp2</div>		<input type="text" value="1"/> <input type="text" value="GB"/>

The **Access Mode** is a static list as determined by the provider. See <https://kubernetes.io/docs/concepts/storage/persistent-volumes/#access-modes> for additional context.



To scale a Kubernetes deployment with volumes, be sure to configure the volume Access Mode as **ReadWriteMany** for deployments that use either scaling or multiple replicas.

See the *volumeInfos* and the *rootVolumeSize* attributes in the *Submit Job* API calls.

Security and Firewall Rules

Application Security and Firewall Rules

- [Overview](#)
- [Deleting Security Groups - Caution](#)
- [Multiple Security Groups](#)
- [Isolation Tags](#)
- [Port Behavior in Service Tiers](#)
- [Inter-Tier Communication \(Firewall Rules\)](#)
 - [Service-Level Microsegmentation](#)
 - [Define a Service with a Dynamic Firewall Rule](#)
 - [Configuring Application Tier Firewall Rules](#)

Workload Manager provisions application VMs in various cloud environments. These cloud environments provide communication security (both between nodes and external access) via Security Groups/Firewall rules. Workload Manager dynamically creates these Security Groups/Firewall rules based on your application topology to allow inter-communication between nodes.

Numerous security groups may be listed under your Workload Manager account. The following table identifies when a security group is created and deleted.



If you delete a security group via the cloud provider console, be aware that you *may* no longer be able to access the application VMs.

Application deployment VMs will not be able to connect to VMs of another application deployment without explicitly allowing connectivity through firewall rules. The following table provides information about firewall rules.

Security /Firewall Type	Security Group Creation	Notes	Security Group Name for Supported Cloud	Security Group Deletion
Tenant or User-Level Firewall Rules	At the time of user initialization. Also created during app launch process.	<ul style="list-style-type: none"> • Created per user only if the <i>Create default security groups for users in this Tenant</i> option is enabled for the tenant • Also used for tenant-level isolation (on AWS, OpenStack, and Google) to allow unrestricted access across VMs created by the user, when <i>Allow launched VMs to communicate with each other</i> is enabled for the tenant. See Tenant Management > <i>Default Security Group</i> for additional details. 	<ul style="list-style-type: none"> • AWS and OpenStack = <code>cliqr-user-security-group_userid</code> • Google = <code>net-workName-c3-user-userId-ruleId</code> 	The user-related security groups are automatically deleted when you delete this user from Workload Manager.
Application Tier-Level Firewall Rules	Created during app launch process.	Created for each application tier for unique set of firewall rules and supported.	<ul style="list-style-type: none"> • AWS and OpenStack = <code>cliqr-firewall-uniqueId</code> • Google = <code>net-workName-c3f-uniqueId-ruleId</code> 	This security group is deleted when you terminate the application deployment .
Deployment -Level Firewall Rules	The application tier /deployment security policy feature which also creates security group	See Policy Management > Security Policies for additional details	<ul style="list-style-type: none"> • AWS and OpenStack use the same format: <code>securityProfileName-uniqueId</code> 	You can only delete a Security Policy if it is not attached to any running job in <i>any</i> cloud.



Azure RM and Azure Stack

A security group is created for each AzureRM and Azure Stack VM with the name as `CliQrSecurityGroup-VMName`

- **OpenStack:** Workload Manager accepts multiple security group configurations with the same name from OpenStack but uses the first security group from the list of security groups returned by OpenStack.
- **AWS:** Workload Manager ensures that an initial job submission is verified and reused if you submit multiple security group configurations with the same name.



If deploying an application using the multi-site/multi-account feature, be sure to address all firewall requirements when you [Model Applications](#). See [Deployment Environments](#) for additional context.

To isolate a deployment to a user on OpenStack or AWS clouds, Workload Manager provides the isolation tag feature. When you launch a VM, isolation tags allow you to restrict the VM communication to a particular deployment (only VMs within this deployment can communicate with each other).

Alternately, by using the **cliqr-user-security-group_self-reference** job security group, all deployments launched by this user can communicate with each other. See [Tenant Information](#) > *Default Security Group* for additional details



Isolation tags are only supported by the *Submit Job* API for AWS and OpenStack deployments.

- When you do not specify a default port for a service tier (for example, an OS service like CentOS) – the Cloud Remote platform will not create a firewall group. In prior releases, Port 80 was assigned (based on the selected protocol) to this service when you saved the application.
- When you add a service tier, ports defined for the associated service (for example, Apache2) are added as the default firewall rule open to CIDR 0.0.0.0/0. In this example, Apache2 will have two firewall rules, Port 80 and Port 443, added by default.
- When a service tier is added below another tier, new firewalls rules are added to the bottom tier. For example, if your application has two tiers, Apache2 and MySQL, with MySQL as the bottom tier that is connected to Apache2 which is the top tier, (top), you will see an entry in the firewall rules for Port 3306 with CIDR as the name of the Apache2 tier within the MySQL tier.

When modeling application profiles or applications for clouds powered by Cisco ACI, use the *Restrict to one-way south-bound communication between connected tiers* option (previously referred to as Enable Microsegmentation) to configure firewall rules between tiers.

- **Checked:** Only a tier that is directly dependent on another tier can communicate with each other.



Be aware that you can add your own firewall rules and customize the rules as required by your application. To implement this, go the applicable tier and add the tier name for the target port in the application. In addition to the tier name, you can specify "any" in the firewall rule for this application so you can expose that port to any endpoint in the ACI context (VRF). See [ACI Extensions](#) for additional details.

- **Unchecked (default):** All tiers within the application can communicate with each other.



If unchecked (disabled), nodes within the application can communicate with each other. However, other than the **top** application tier, firewall rules relating to a service-specific port are ignored and cannot be saved. This is regardless of the firewall rule being set to open a service-specific port for public internet access or a specific CIDR/subnet.

Service-Level Microsegmentation

The following behavior applies to the service-level microsegmentation feature:

- If required, you must explicitly configure Ports 22 (SSH) and 3389 (RDP) as part of the application tier firewalls or the Vendor (tenant) Firewall list.
- If you specify any tenant firewall rules, a shared security group (*cliqr-job-worker**) with the rules is applied to all VMs deployed within the tenant.
- Services now contain a set of internal firewall rules. These rules are applied as defaults to application tiers when you model applications. The default CIDR for application tiers default to 0.0.0.0/0. The IP/CIDR/Tier automatically changes to the name of the dependent tier as soon as you make the connection in the [Topology Modeler](#).
 - Nodes within the same tier can communicate with each other.
 - Nodes in different deployments cannot communicate with each other.
 - If you do not specify a default port for a service tier (for example, an OS service like CentOS) – the Cloud Remote platform will not create a firewall group.
 - If you add a service tier, ports defined for the associated service (for example, Apache2) are added as the default firewall rule open to CIDR 0.0.0.0/0. In this example, Apache2 will have two firewall rules, Port 80 and Port 443, added by default.
 - When a service tier is added below another tier, new firewalls rules are added to the bottom tier. For example, if your application has two tiers, Apache2 and MySQL, with MySQL as the bottom tier that is connected to Apache2 which is the top tier, (top), you will see an entry in the firewall rules for Port 3306 with CIDR as the name of the Apache2 tier within the MySQL tier.

Define a Service with a Dynamic Firewall Rule

To define a dynamic firewall rule (metadata) at the service level, see [Custom Service Definition](#).

Configuring Application Tier Firewall Rules

See [Topology Modeler](#) > *Basic Information* for additional context on the **Enable microsegmentation** checkbox to automatically set firewall rule for ACI environments.

To configure firewall rules for an application tier, follow this procedure.

1. [Model an Application](#).

2. When you model the topology configure the Firewall Rules for each tier in the Topology Modeler. For example in the image below, the Apache Firewall Rules are being configured for the MySQL tier.
3. Connect the Apache service to the MySQL service. As soon as you connect the service in the Topology Modeler, the IP column in the Firewall Rules section automatically updates to display the Apache application tier name.
4. You can choose to keep the pre-configured firewall definition or add additional Firewall rules as required.

IP Address Allocation

IP Address Allocation

- [IP allocation mode \(OpenStack only\)](#)
- [Validation of IP Usage for Imported VMs](#)
- [Managing IPv6 \(AWS and OpenStack only\)](#)
- [Assign IPv4 Public IP \(All VM-based clouds except vCenter\)](#)

This feature allows you to add multiple Network Interface Cards (NICs) to a VM, and for each NIC, discover the private IP beforehand (pre-allocate IP) and pass it through the config drive file to the OpenStack cloud provider.



This feature is only available for OpenStack clouds.

OpenStack supports the association of the public IP to any private IP and hence to any NIC as mentioned by the user.

At deploy time you can also choose the IP allocation strategy for each NIC's Private IP Address: DHCP or Pre-allocate IP:

Enumeration	Description
DHCP (default)	This strategy allows the IP to be allocated by the DHCP server to the instance on server boot up. This IP address is not known prior to server boot up.
Pre-allocate IP	This strategy allows the cloud infrastructure IP allocation to be dynamically provided before the server boots up. This strategy is specific to the following OpenStack applications: <ul style="list-style-type: none"> • CISCO CSR1000: Configuration drive file IP populated with the pre-allocated IPs known before server boot up. • CISCO F5 Load Balancer: Multiple NIC support.



Use IPAM callout scripts to configure IP allocation strategies. See [Guidance for Callout Scripts](#) for additional context.

Configuration File Attributes

Specify the configuration drive file and configuration drive file contents in the Global Parameters tab in the Topology Modeler. When you launch this job, the parameters in the file contents will be replaced by the relevant IP addresses. You can only perform IP substitutions in the configuration file contents by using the pre-allocate IP mode.

- Add a **cloudConfigFilePath** parameter with a value as file path and type as **string**.

▼ Parameter

Parameter Name: 21

Display Name:

Help Text:

Type: Max Length:

Default Value: 239

User Options:

Should this parameter be visible to the user?

Should this parameter be editable by the user?

Should this parameter be optional?

- Add a **cloudConfigFileContents** parameter with value as file contents and type as **textarea**

Parameter

Parameter Name 17

Display Name

Help Text

Type

Default Value

```
%NIC1_IP%
%NIC2_IP%
%NIC3_IP%
```

The %NIC#_IP% Parameter Substitution (OpenStack only)

The number of NICs (NIC1, NIC2...) corresponds to the number of network interfaces in the application profile. The IP address substitution parameters used in configuration files (see [Configuration Files](#)) will be in the same order as %NIC1_IP%, %NIC2_IP%, ... configuration. If the number of NICs = 1 *and* you have a preferred/default network defined, the input will default to that network with the DHCP option. Otherwise, you must select the required network for any job submission.

The NIC order is important – the order defines the parameters loaded with IP address. For example, if the order is 1 and strategy is PREALLOCATE_IP, the parameter generated is %NIC1_IP% (first NIC), %NIC2_IP% (second NIC) and this parameter is replaced in the configuration drive file provided to VM with the corresponding IP address.

When installing the CloudCenter agent on an imported VM the Workload Manager verifies if the imported VM uses IPv4 or IPv6 IPs when installing the agent and issues a corresponding error.

For more information on importing VMs, see [Virtual Machine Management](#).

i IPv6 Addresses

By default, all networks default to using IPv4 and no additional configuration is required when using IPv4 addresses.

When allocating firewall rules, CloudCenter Suite supports IPv6, in addition to IPv4, addresses in the source for app profile, tenant, and security policies. When you assign IPv6 addresses, Workload Manager validates the security rule source before accepting the IPv6 address. This support is **restricted to AWS and OpenStack** clouds. If you provide an invalid IPv4/6 IP address, then Workload Manager rejects the deployment as invalid.

To use IPv6 addresses, follow this process.

1. Configure Firewall Settings in [Tenant Management](#) and check the following boxes:
 - *Allow launched VMs to communicate with each other*
 - *Create default security groups for users in this Tenant*
 - See [Security and Firewall Rules](#) for additional context
2. Add the required firewall rule using the IPv4/6 address and update the Tenant.
3. Add a [Policies](#) > [Security Policy](#) and configure it to use the firewall rule(s) using the IPv6 address.
4. Configure the firewall rule in the application profile (Application > [Topology Modeler](#) > [Firewall Rule](#) > Add the IP address and the validation works here as well.
5. Deploy the application using the IPv6 address.

The [Virtual Machine Management](#) > [Managed VMs](#) > [VM Details](#) section displays a configured IPv6 IP address.

- If an IPv6 IP address is not configured, this field displays a dash (-).
- If multiple IP addresses are configured, this field displays each configured address in its own line.

Note these cloud nuances for managing IPv6 addresses:

- AWS

- The **Assign IPv6 Address** feature is also available on a per-NIC basis as Workload Manager supports IPv6 addresses as a source for application profiles, tenants, and security policies.
 - Deploy the *Application Profile* by selecting the configured *Security Policy* in the *General Settings* and *Tier Settings* sections and the *Network* with IPv4/6 along with the IP allocation requirements.
 - The CloudCenter Deployments page displays the new deployment and your Cloud Console displays the spawned instance. When you view the details for this deployment in the cloud console, you see the corresponding security group and firewall settings deploying the IPv6 address configuration if you have checked the *Create default security groups for users in this Tenant* and *Allow launched VMs to communicate with each other boxes*.
- OpenStack
 - The **Assign IPv6 Address** feature is also available on a per-NIC basis as Workload Manager supports IPv6 addresses as a source for application profiles, tenants, and security policies.
 - Deploy the *Application Profile* by selecting the configured *Cloud Tenant* and *Network* with IPv4/6 and providing the IP allocation requirements in the *Cloud Settings* section.
 - The CloudCenter Deployments page displays the new deployment and your Cloud Console displays the spawned instance. When you view the details for this deployment in the cloud console, you see the corresponding security group and firewall settings deploying the IPv6 address configuration if you have checked the *Create default security groups for users in this Tenant* and *Allow launched VMs to communicate with each other boxes*.
 - The IPv6 field merely identifies if IPv6 is enabled or disabled for the subnet in the cloud-level settings. This is an information field and cannot be configured.

Note these cloud nuances for managing IPv4 addresses:

- AWS
 - The **Assign IPv4 Public IP** feature indicates that the eth0 interface can be associated with an IPv4 public IP.



During application deployment or when setting deployment environment defaults for AWS configurations, the **Assign Public IP** setting reflects the status of the **Enable auto-assign Public IP** setting for the selected AWS subnet. If **Assign Public IP** = Unchecked + **Enable auto-assign Public IP** = Unchecked, the Public IP is not assigned and the deployment fails. For this reason, Workload Manager automatically toggles the **Assign Public IP** setting based on the pinned subnet. You can override this setting during the CloudCenter application deployment. Be aware that if you do not have access to the internet, the subnet routing table will not be connected to an internet gateway.

- Google
 - The **Assign IPv4 Public IP** feature is available on a per-NIC basis as Google supports private-public IP pair.
 - Multiple NIC support is not currently available.
- OpenStack
 - The **Assign IPv4 Public IP** feature is available on a per-NIC basis as OpenStack supports private-public IP pair.
- vCenter
 - *vCenter does not support public IP addresses.*

Parameters and Macros

Parameters and Macros

- [Using Parameters](#)
- [Pre-Defined Parameters](#)
- [Parameter Substitution](#)
- [Configuration Files](#)

Using Parameters

Using Parameters

- [About Parameters and Macros](#)
- [Defining Parameters](#)
- [Parameter Type](#)
- [Default Value Usage](#)
- [Deployment-Specific Parameters](#)
- [Granular Control for User-Defined Parameters](#)

Parameters (also referred to as *macros*) are global or tier/step-specific system- or user-defined variables for which values can be provided at deployment time. These parameters can be used as tokens or macros inside the application's configuration file or configuration scripts. At deployment time, the macros are replaced by the actual values specified in the parameter.

Workload Manager allows you to define your own parameters or use the Workload Manager-supported parameters.

See [Deployment Parameters](#) for additional details on when to use Deployment Parameters.

When adding a new service, you can use global parameters, tier/step-specific system parameters, or user-defined parameters and add values at deployment time.

- **Service Parameters:** These parameters are specific to a newly-added/modified service. The parameter is called when a user calls this service. See [Custom Service Definition](#) and [Service Administration](#) for additional context.
- **Global Parameters:** These parameters apply to all tiers and steps within this application as displayed in the following screenshot.

- **Deployment (Custom) Parameters:** These parameters are specified for each step or tier within an application as displayed in the following screenshot. See [Understand Application Tier Properties](#) for additional context.



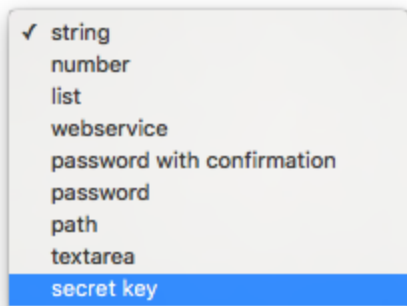
When you save a custom **Parameter Type** as a *string*, the **Deployment Environments > General Settings** section displays it as a *textarea* field.

When changing existing applications, edit and save the application if you prefer the *textarea* field to be displayed as a *text* field.

Regardless of the parameter being defined at the service level, the global application level, or at a tier level, you can specify the Type for each parameter. The Type dropdown enables you to determine your own value for a custom field:

by selecting one of the following options from the **Type** dropdown list.

Type



Type	Description
string	<p>Provide an alpha-numeric value (limited to 255 characters).</p> <p>You can optionally add a Regular Expression for <i>this</i> parameter type, and the Workload Manager validates (when saving an application using this value) the provided value against this regular expression.</p> <p>API Enumeration = string (valueConstraint)</p>
number	<p>Provide a range of numbers (limited to 255 characters).</p> <p>API Enumeration = number</p>
list	<p>Provide a list of comma separated text values (limited to 255 characters).</p> <p>API Enumeration = list (valueList)</p>
webservice	<p>WebService is a parameter type introduced in CloudCenter Legacy 4.4.</p> <p>You can provide list of dynamic webservice parameters while deploying a job. From this list, users can select one parameter. The webservice parameter type is available in custom parameters, global parameters, and services. The output should be in the following format:</p> <pre>[{"name": "p1", "displayName": "Param 1"}, {"name": "p2", "displayName": "Param 2"}]</pre> <p>If you configure this parameter type, you must provide the Protocol (HTTP or HTTPS), Web Service URL, and the credentials (Username and Password) for this webservice.</p> <div style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;"> To use this parameter type, you must set the Content-Type to be returned by webservice as application/html. </div> <p>API Enumeration = webservice (webserviceListParams)</p>
password with confirmation	<p>Provide an alpha-numeric value (limited to 255 characters).</p> <p>The UI provides a confirmation textbox.</p> <p>API Enumeration = password_input (valueConstraint)</p>
password	<p>Enter text with no constraints – for example, use this field for a password that just needs to be entered without validation or confirmation.</p> <p>The UI does not provide a confirmation textbox.</p> <p>API Enumeration = password</p>

path	A URL for the download location at the time of orchestration.	
	Options for path	Description
	Repository	Different repositories supported by Workload Manager. See Artifact Repository > Workload Manager Repository Types for a list of options. Format: %REPO_ID_{id}%xyz.war Example: %REPO_ID_2%xyz.war where 2 is the ID of the repository in Workload Manager
	File in Package	Path of the ZIP file defined in the application package of the application profile parent tier. Example: %PACKAGE_DIR%script.sh
	Storage	Path for the mounted storage location. Linux VM Example: /shared
	URL	URLs of type HTTP, HTTPS, and FTP
API Enumeration = path (text that encodes a URL)		
textarea	You can add any free form text in this area – for example , use this field for to define a parameter with values as file contents, or to define a private key, or to define a script, and so forth. API Enumeration = textarea (text of any length (long))	
secret key	This type is useful when configuring a Container Service . When you select this parameter, you can configure another layer of abstraction when specifying the Key Name (group-level coordinate) and the Key (the location) to retrieve the actual value of the secretkey for the deployment parameter that uses this type. The secret value is not visible to users nor is it maintained by Workload Manager – Workload Manager retrieves the secret from the specified location and uses that secret as configured.	
json	You can add any JSON value in the Default Value field if you select this type. The value added in this field is validated for accurate JSON syntax. This value can then be called as a parameter at the required point in the service lifecycle process.	
yaml	You can add any YAML value in the Default Value field if you select this type. The value added in this field is validated for accurate YAML syntax. This value can then be called as a parameter at the required point in the service lifecycle process.	

The **cliqIgnoreAppFailure** parameter does not terminate nodes during a failure and continues to keep all failed application deployment nodes running. If you configure this parameter, then you can enter true or false in the Default Value field to ensure that a VM is or is not terminated if an error is encountered at start up. See [Troubleshooting Parameters](#) for additional context.

Workload Manager includes an improvement to allow service parameters to be displayed to end users when they complete a Deployment form. Users are prompted to confirm if a parameter is a deployment parameter. Based on the [Permission Control](#) settings for this service a user can also be assigned permissions to change the parameter value at deployment time.

The **User Options** that the following table describes provide granular control over a user-defined parameters:

User Option	Checked	Unchecked
Should this parameter be visible to the user?	This parameter is visible to users in the Topology Properties sections.	This parameter is not visible to users. This allows admins to configure hidden parameters for licensing or tracking purposes (for example, to track how many users are using this service).
Should this parameter be editable by the user?	Users can override this parameter in the Topology Properties settings.	This service becomes automatically invisible to the user, and therefore not editable.
Should this parameter be optional ?	Marked optional – The Cloud Center platform accepts the provided default values, if any (default = None), and the user must select a value to change the default.	Marked required – If a default value is not provided, then the user must select one of the values returned by the web service.

Back to: [Parameters and Macros](#)

Related Page: [Deployment Parameters](#)

Pre-Defined Parameters

Parameters Defined by Workload Manager

- [Overview](#)
- [Using Workload Manager-Defined Parameters](#)
- [Standard IP Address Replacement Parameters](#)
- [Input Parameters for Scripts](#)
- [SSH Keys Parameters](#)
- [Automation Policy Parameters](#)
- [RDS Service Parameters](#)
- [Kubernetes Container Service Parameters](#)
- [Docker Container Environment Variables](#)
- [Docker Custom Service Variables](#)
- [VM-Based Environment Variables](#)
- [N-Tier Environment Variables](#)
- [Cluster-Based Environment Variables](#)
- [Lifecycle Actions Variables](#)
 - [Job Variables for Lifecycle Actions](#)
 - [Application Profile Variables for Lifecycle Actions](#)
 - [Service Variables for Lifecycle Actions](#)
 - [Cloud Region Variables for Lifecycle Actions](#)
- [Parameters as Environment Variables](#)
- [The userenv Utility File](#)

Automated parameters are defined in Workload Manager and contain the following characteristics:

- Values provided by Workload Manager.
- Designed to cover unknown, dynamically generated values or system-dependent values.
- Reusable in application profiles to automate jobs without writing extensive scripts (examples include timestamp, dynamically-generated IP address, private IP address, IP address of a tier, environment variables, automation policy parameters, number of nodes in a cluster, deployment name, and so forth).

To use a system-defined parameter, follow this procedure:

1. Use the parameter name to reference this parameter inside a configuration script or configuration file as a macro using Workload Manager convention. For example:
 - Parameter name: `XYZ_TIER_IP`
 - Macro reference for this parameter: `%XYZ_TIER_IP%`
2. Reference the macro inside the Application Profile.

The following sections list all parameters, macros, and environment variables defined for Workload Manager for use in Workload Manager deployments.

These parameters are available at the application level. Some parameters are specific to some application profiles as identified in the Description column.



Use the same case and format (UPPER_CASE).

Workload Manager-Defined Parameters	Description
<code>%XYZ_TIER_IP%</code>	<ul style="list-style-type: none"> • Private IP address of a custom tier. XYZ is custom tier name • If the tier name is Custom1, then macro name will be <code>%Custom1_TIER_IP%</code>
<code>%DB_TIER_IP%</code>	<ul style="list-style-type: none"> • Private IP address of the database tier. • The tier name must be <i>Database</i>
<code>%NOSQL_TIER_IP%</code>	Private IP address of NoSQL Database tier
<code>%MB_TIER_IP%</code>	Private IP address of Message BUS
<code>\$(CliqTier_(Tiername)_IP</code>	Private IP of tier named XYZ
<code>%BC_TIER_IP%</code>	Private IP address of Backend Cache

These parameters are available at the application level. Some parameters are specific to some application profiles as identified in the Description column.



Use the same case and format (UPPER_CASE)



See VM Name Config for the pre-defined macros that you can use for the CliQr Macro Replacement Option.

Workload Manager-Defined Parameters	Description
%OUTPUT_DIR%	Path for output directory: /shared/output
%JOB_NAME%	Name of the deployment
%USER_NAME%	Name of the user deploying the application
%FIRST_NAME%	First name of the user deploying the application
%LAST_NAME%	Last name of the user deploying the application
%EMAIL_ADDRESSES%	Email address of the user deploying the application
%USER_EXTERNAL_ID% or %userExternalId%	External ID of the user deploying the application (if using SSO) Both versions of this macro are the same – the <i>%userExternalId%</i> macro is replaced by the user's external ID in the URL for <i>webservice</i> (see Using Parameters > Parameter Type for additional details on <i>webservice</i>)
%DEPLOYMENT_ENV%	Name of the deployment environment
%DEP_ENV_NAME%	Name of the Puppet or Chef environment
%TIME%	Time stamp string in yyyyMMddHHmmss format
%TASK_INDEX%	<ul style="list-style-type: none"> Task index starting from 1 Applies to Batch Compute Application Profiles
%NODE_INDEX%	<ul style="list-style-type: none"> Node index starting from 1 Applies to: <ul style="list-style-type: none"> Cluster Compute Application Profiles Parallel Execution Application Profile
%HOST_IPS%	<ul style="list-style-type: none"> Host IP address list delimit by comma Applies to Cluster Compute Application Profiles
%NUM_CPU%	Number of CPUs on the application VM
\$VM_NODE_INDEX	<p>This is an additional macro to create VMs using a sequential number for each VM that is newly created. For example, if you use this macro to launch an N-Tier app with 3 nodes:</p> <ul style="list-style-type: none"> The VM_NODE_INDEX is set as 1, 2, and 3. If you remove Node 2 and add one more node, the VM_NODE_INDEX is set as 1, 3, 4, and 5. For Tiers with only one node the VM_NODE_INDEX is set to 1. <p>This macro is exported as part of userenv file. To consume this macro, add \$VM_NODE_INDEX to your scripts and source the userenv(/usr/local/osmosix/etc/userenv) file use the variables as part of your scripts.</p>

The parameters listed in the following table are available when using SSH keys.

Workload Manager-Defined Parameters	Description
-------------------------------------	-------------

sshUserName	The SSH user name for the current cloud VM
sshKey	The SSH private key for the current cloud VM
sshPublicKey	The SSH public key for the current cloud VM

These parameters are *only* applicable when defining custom action policies and are not available at the application level.



Use the same case and format (camelCase)

Workload Manager-Defined Parameters	Description
%myEmail%	User's email address
%firstName%	User first name
%vendorName%	Tenant name to send notifications, if any
%jobName%	Name of the deployment job
%jobUrl%	The URL for the deployment
%jobType%	Type of deployment job (for example, benchmark)
%appName%	Name of the application
%owner%	Owner of the deployment job
%status%	Current status of the job (for example, Running)
%cloudName%	Name of the cloud (for example, Amazon US East (Virginia))
%paramCloudType%	Identifies the cloud in which the node is launched

When defining the RDS OOB service, be sure to pass the values for the VPC ID and DB Subnet Group at deployment time. Workload Manager requires these values to provision the RDS instance. See [External Service](#) for additional context. The following table identifies a list of parameters that are required when configuring an RDS service. These following RDS-specific deployment parameters are also available on Workload Manager.

Parameter	Description
vpclid	Required. The VPC ID with which the RDS Instance is associated, its the same VPC ID of the DBSubnetGroup – a Security Group will be created using this VPC ID.
dbSubnetGroup	Required. The DB Subnet Group with which the RDS Instance is associated.
cliqrIsPublicAccessible	Required. The RDS service to be toggled to true (default) or false so you can control public accessibility.
cliqrRdsEngineVersion	You can customize the RDS service using this parameter. For the RDS MySQL OOB service, the default version is listed as 5.6.35. You can customize this version by adding a custom, deploy parameter with this name and pass the value during deployment.
cliqrStorageType	You can customize the RDS service using this parameter. For RDS MySQL OOB service, the default storage type used by Workload Manager is Magnetic (standard storage). You can customize this storage type by adding a custom, deploy parameter with this name and pass the value as “gp2” for general purpose SSD or “io1” for provisioned IOPS. See Multiple Volumes for additional context.
port	You can customize the RDS service using this parameter. For RDS MySQL OOB service, the default port used by Workload Manager is 3306. You can customize this port by adding a custom, deploy parameter with this name and pass the value as desired.
cidr	You can customize the RDS service using this parameter. When creating an RDS security group, the default port used by Workload Manager opens to CIDR 0.0.0.0/0. You can customize the CIDR by adding a custom, deploy parameter with this name and pass the value as desired.
sgid	You can customize the RDS service using this parameter. For RDS MySQL OOB service, Workload Manager creates a new security group with security rules mentioned as <i>port</i> and/or <i>cidr</i> parameters, You can customize the sgid by adding a custom, deploy parameter with this name and pass the value of security group so that RDS Instance will be launched using the specific security group. See Security and Firewall Rules for additional context.

The Cloud Remote platform supports variables to reference various Kubernetes network parameters as listed in the following table. Users can reference the namespace of the tier, and various network parameters for a specific service name in the namespace.

In the following table:

- *<ServiceName>* is the string entered by the user in the Port Name field in the Network Services section of the topology modeler for a container tier.
- *ServiceName* is the string generated by Workload Manager at deploy time and injected into the Kubernetes cluster using Port Name, Job ID, and a 6-character hexadecimal random number generator.

The following network related parameters are available from a Kubernetes [Container Service](#):

Output value	Output format	Variable syntax	Applicable service types
Cluster IP Service Name List	<ServiceName1>, <ServiceName1>, . . .	CliqrTier_<tier_name>_ClusterIP_ServiceName	ClusterIP
NodePort Service Name List	<ServiceName1>, <ServiceName1>, . . .	CliqrTier_<tier_name>_NodePort_ServiceName	NodePort
LoadBalancer Service Name List	<ServiceName1>, <ServiceName1>, . . .	CliqrTier_<tier_name>_LoadBalancer_ServiceName	LoadBalancer
ClusterIP Internal Endpoint List	<ServiceName1>:<port_no>, <ServiceName2>:<port_no>, . . .	CliqrTier_<tier_name>_ClusterIP_Endpoint	ClusterIP
NodePort Internal Endpoint List	<ServiceName1>:<service_port_no>, <ServiceName1>:<node_port_no>, <ServiceName2>:<service_port_no>, <ServiceName2>:<node_port_no>, . . .	CliqrTier_<tier_name>_NodePort_Endpoint	NodePort
LoadBalancer Internal Endpoint List	<ServiceName1>:<port_no>, <ServiceName2>:<port_no>, . . .	CliqrTier_<tier_name>_LoadBalancer_InternalEndpoint	LoadBalancer
External Endpoint List	<ip_address1>:<port_no>, <ip_address2>:<port_no>, . . .	CliqrTier_<tier_name>_LoadBalancer_ExternalEndpoint	LoadBalancer
Public IP List	<ip_address1>, <ip_address2>, . . .	CliqrTier_<tier_name>_PUBLIC_IP	LoadBalancer
Cluster IP List	<ip_address1>, <ip_address2>, . . .	Any of these: CliqrTier_<tier_name>_ClusterIP CliqrTier_<tier_name>_PRIVATE_IP CliqrTier_<tier_name>_IP	All
Pod IP List	<ip_address1>, <ip_address2>, . . .	CliqrTier_<tier_name>_PodIP	All
Namespace	string	CliqrTier_<tier_name>_Namespace	All
Cluster IP	<ip_address>	Any of the following: CliqrTier_<tier_name>_<service_name> CliqrTier_<tier_name>_<service_name>_ClusterIP CliqrTier_<tier_name>_<service_name>_PRIVATE_IP CliqrTier_<tier_name>_<service_name>_IP	All
Internal Endpoint	For ClusterIP and LoadBalancer service types: <ServiceName>:<port_no> For NodePort service type: <ServiceName>:<service_port_no>, <ServiceName>:<node_port_no>	CliqrTier_<tier_name>_<service_name>_Endpoint	All
External Endpoint	<ip_address>:<port_no>	CliqrTier_<tier_name>_<service_name>_ExternalEndpoint	LoadBalancer only
Public IP	<ip_address>	CliqrTier_<tier_name>_<service_name>_PUBLIC_IP	LoadBalancer only
Service Name	<service_name>-<job_id>-<rnd_no>	CliqrTier_<tier_name>_<service_name>_ServiceName	All

These parameters may be referenced within the Deployment Parameters section of the Topology Modeler when the application profile contains a container service tier. See [Container Service](#) for an example.



If only one network service for a particular network service type is defined for a container-based tier in the Topology Modeler, then the corresponding service name and IP parameters will contain single values. If multiple network services of a particular network service type are defined for a tier in the Topology Modeler, then the corresponding service name and IP parameters will each contain a comma separated string representing the values for each of the network services of that type.

These parameters may be referenced within the Deployment Parameters section of the Topology Modeler when the application profile contains a container service tier. See [Container Service](#) for an example.



If only one network service for a particular network service type is defined for a container-based tier in the Topology Modeler, then the corresponding service name and IP parameters will contain single values. If multiple network services of a particular network service type are defined for a tier in the Topology Modeler, then the corresponding service name and IP parameters will each contain a comma separated string representing the values for each of the network services of that type.

The variable listed in the following table are specific to Docker containers.

Cloud Remote-Defined Variables	Description
Cloud_Setting_CloudFamily	The cloud family of the region in Workload Manager.
Cloud_Setting_cloud	The region name, for example: <i>Amazon-us-west-2</i> .
CliqrTier_NameList	The names list of all tiers in the application.
CliqrCloud_RegionEndPoint	The region endpoint, example: <i>ec2.us-west-2.amazonaws.com</i> .
CliqrDepEnvId	The Deployment environment ID created by Workload Manager.
cliqrAppName	The application Name in Workload Manager.
launchUserName	The launch username in Workload Manager, for example: <i>cliqradmin</i>
cliqrAppType	The type of application in Workload Manager, for example: <i>n-tier</i>
CliqrDepEnvName	The deployment environment name n Workload Manager

Right before your script is executed in a custom Docker service, Cloud Remote supplies all the following application-specific parameters as environment variables.

Cloud Remote-Defined Variables	Description
\$region	Identifies the region information in the environment variables for an External Service .
\$CloudFamily	Identifies the cloud family in the environment variables for an External Service .
cliqrContainerExecuteScriptTimeout	Limits the Docker container from running forever if an external service falls into an infinite loop. See External Service > Script Timer .

These variables are also part of the application and password parameters list.

Cloud Remote-Defined Variables	Description
cliqrNodeid	<ul style="list-style-type: none"> The Node ID of the VM for which the external initialization script is executed. Used for the external, pre-init, Deployment Lifecycle Scripts.
cliqrNodeHostname	<ul style="list-style-type: none"> The hostname of the VM for which the external initialization script is executed. Used for the external, pre-init, Deployment Lifecycle Scripts.
cliqrNodePublicIp	The Public IP of the VM for which the external initialization script is executed.
cliqrNodePrivateIp	The Private IP of the VM for which the external initialization script is executed.
\$cliqrWindowsPassword	<ul style="list-style-type: none"> Read only Windows password variable. Used for the external, pre-init, Deployment Lifecycle Scripts.

The parameters listed in the following table are available at the application level. Some parameters are specific to some application profiles as identified in the Description column.



These values cannot be overridden!

Use the same case and format (InitCaps).

Workload Manager-Defined Variables	Description
\$CliqrDependents	Comma separated list of connected top tiers name of the tier
\$CliqrDependencies	Comma separated list of top tier names connected to bottom tier names
\$CliqrDepEnvName	Deployment environment name

\$CliqrDepEnvId	<ul style="list-style-type: none"> Helpful on the <i>restore</i> side of the migration as it is available from the source (where the backup is being run). For example: To backup, use <code><path>/\$CliqrDepEnvId</code>
\$CliqrDeploymentId	<ul style="list-style-type: none"> Current deployment ID that ties the pre- and post-migration operations together for scripting purposes (backups and so forth) For example: To backup, use <code><path>/\$CliqrDeploymentId</code>
\$migrateFromDepld	<ul style="list-style-type: none"> Only available for post-migrate operations and points to the previous deployment ID. For example: To restore, <code><path>/\$migrateFromDepld</code>
CliqrTier_<job_name>_PUBLIC_IP	Public IP address for the tier. This parameter is also available as a Cluster-Based Environment Variable (Cluster Compute Application Profiles).
\$CliqrTag_{Tagname}	Job tags defined during run time with value
ScaleType	<ul style="list-style-type: none"> The value can be either ScaleUp or ScaleDown. Displays the type of scaling on the tier with the name \$CliqrTier_scalingTierName)
CliqrTier_scalingTierName	<ul style="list-style-type: none"> The tier name on which the scaling operation occurred.
CliqrTier_{\$CliqrTier_scalingTierName}_scaleUpPublicIps	<ul style="list-style-type: none"> ScaleUp parameter. The public IPs of the VMs that are scaled up on the tier with the name \$CliqrTier_scalingTierName.
CliqrTier_{\$CliqrTier_scalingTierName}_scaleUpNodeIds	<ul style="list-style-type: none"> ScaleUp parameter. The nodeIds of the nodes that are scaled up on the tier with name \$CliqrTier_scalingTierName.
CliqrTier_{\$CliqrTier_scalingTierName}_scaleUpPrivateIps	<ul style="list-style-type: none"> ScaleUp parameter. The private IPs of the nodes that are scaled up on the tier with name \$CliqrTier_scalingTierName.
CliqrTier_{\$CliqrTier_scalingTierName}_scaleUpHostnames	<ul style="list-style-type: none"> ScaleUp parameter. The hostnames of the nodes that are scaled up on the tier with name \$CliqrTier_scalingTierName.
CliqrTier_{\$CliqrTier_scalingTierName}_scaleDownPublicIps	<ul style="list-style-type: none"> ScaleDown parameter. The public IPs of the nodes that are scaled down on the tier with name \$CliqrTier_scalingTierName.
CliqrTier_{\$CliqrTier_scalingTierName}_scaleDownNodeIds	<ul style="list-style-type: none"> ScaleDown parameter. The nodeIds of the nodes that are scaled down on the tier with name \$CliqrTier_scalingTierName.
CliqrTier_{\$CliqrTier_scalingTierName}_scaleDownPrivateIps	<ul style="list-style-type: none"> ScaleDown parameter. The private IPs of the nodes that are scaled down on the tier with name \$CliqrTier_scalingTierName.
CliqrTier_{\$CliqrTier_scalingTierName}_scaleDownHostnames	<ul style="list-style-type: none"> ScaleDown parameter. The hostnames of the nodes that are scaled down on the tier with name \$CliqrTier_scalingTierName.
CLIQR_EXECUTE_SCRIPT	The execution command/script for the current Docker container run.
externalServiceBundle	The path/URL to the external service bundle (specified in the Agent Bundle URL field in the Cloud Settings section of the Regions/ Details tab for a cloud).
actionType	The action type of the current external service action, the possible values: SCRIPT or URL or CMD
serviceName	The name of the External Service .

CUSTOM_REPO_URL	Workload Manager custom bundle store URL (specified in the Agent Custom Repository field in the Cloud Settings section of the Regions/ Details tab for a cloud).
-----------------	--


The variables listed in the following table are available in Cluster Compute Application Profiles.

Workload Manager Variables	Description
CliqrTier_<job_name>_IP	Private IP address for the tier
CliqrTier_<job_name>_PUBLIC_IP	Public IP address for the tier. This parameter is also available for N-tier deployments
CliqrTier_<job_name>_HOSTNAME	Hostname for the tier

Lifecycle actions have access to a separate list of environmental variables that are associated with the resources to which they are mapped. You can use these variables in a script, default value, or body fields modeled as part of the action, and are passed on as environment variables to the VM. See [Actions Library > Lifecycle Actions](#) for additional context. The following table lists Workload Manager-defined environment variables for Lifecycle Actions.

Workload Manager Variables	Description
Job Variables for Lifecycle Actions	
JOB ID: %jobId%	The internal ID for a deployed run
JOB NAME %jobName%	The name of the deployed run
JOB URL: %jobUrl%	The deployment URL
JOB TYPE: %jobType%	The deployment type
JOB OWNER: %jobOwner%	The owner of the Deployment
Application Profile Variables for Lifecycle Actions	
APP ID: %appId%	The Application Profile ID
APP NAME: %appName%	The Application Profile Name
APP OWNER: %appOwner%	The owner of the Application
APP OWNER ID: %appOwnerId%	The ID of the Application Owner
Service Variables for Lifecycle Actions	
SERVICE ID: %serviceId%	The Service ID
SERVICE NAME: %serviceName%	The Service Name
SERVICE OWNER: %serviceOwner%	The Service Owner
SERVICE OWNER ID: %serviceOwnerId%	The ID of the Service Owner
Cloud Region Variables for Lifecycle Actions	
CLOUD ID: %cloudId%	The Cloud ID
CLOUD NAME: %cloudName%	The Cloud Name
CLOUD OWNER: %cloudOwnerId%	The ID of the Cloud Owner

The parameters listed in the following table are available as environment variables. Some parameters are specific to some application profiles as identified in the Description column.

 Use the same case and format (InitCaps or camelCase)

Workload Manager Variables	Description
\$UseBatchTaskList	<ul style="list-style-type: none"> • Applies to Batch Execution Application Profile. • When set to: <ul style="list-style-type: none"> • 1 = command is read from a task list file • 0 = command is an execution command line

\$NumTasks	<ul style="list-style-type: none"> • Only applicable when UseBatchTaskList=0. • Represents the number of tasks to be performed (number of times to repeat a series of execution commands). • Applies to Batch Compute Application Profiles
\$numClusterNodes	Number of nodes in the tier
\$minClusterSize	Minimum number of nodes in a cluster
\$maxClusterSize	Maximum number of nodes in a cluster
\$minAppClusterSize	Maximum number of nodes for java container-based tiers

Tier-based firewall information is available in a VM's userenv file as well as the environment for external initialization scripts as evident in the following examples:

• **Example 1**

```

firewall={
  "id":null,
  "name":"cliqr-firewall-58E00D70D552698B4355B1F2135F0229",
  "description":"Cliqr user firewall",
  "rules":[
    {
      "protocol":"tcp",
      "fromPort":22,
      "toPort":22,
      "sourceIPRanges":["0.0.0.0/0"],
      "sourceGroups":null,
      "cloudId":null,
      "serviceName":null
    }
  ]
}

```

• Example 2

```
firewall={
  "id":null,
  "name":"cliqr-firewall_C1_VMW2_CentOS_1_6438",
  "description":"Cliqr user firewall",
  "rules":[
    {
      "protocol":"tcp",
      "fromPort":22,
      "toPort":22,
      "sourceIPRanges":["0.0.0.0/0"],
      "sourceGroups":null,
      "cloudId":null,
      "serviceName":null
    },
    {
      "protocol":"tcp",
      "fromPort":80,
      "toPort":80,
      "sourceIPRanges":["1.1.1.1/0"],
      "sourceGroups":null,
      "cloudId":null,
      "serviceName":null
    },
    {
      "protocol":"tcp",
      "fromPort":1,
      "toPort":65535,
      "sourceIPRanges":[
        "TIER_WebServer"
      ],
      "sourceGroups":null,
      "cloudId":null,
      "serviceName":null
    }
  ]
}
```

Back to: [Parameters and Macros](#)

Parameter Substitution

Parameter Substitution

- [Overview](#)
- [Substitute Parameters at Deployment Time](#)
- [Substitute Parameters with Macros](#)
- [Macro Substitution Examples](#)
- [Parameter and Marco Usage Examples](#)

You need to modify the configuration file (or script) or the property file for your application to include Workload Manager-defined macros so these macros automatically plugin the appropriate values for each parameter inside the configuration file/script. These parameters cover cases where the users may not know settings and value in advance and rely on the underlying infrastructure or environment or system to generate the value once the application is deployed. Different profiles may need different system macros.

Sometimes, you may have a custom parameter already defined in a particular service that you do not want to set in the application profile. Instead, you want to leave it to your end users to set this parameter when they deploy the application.

Workload Manager allows you to substitute such parameters at deployment time.

To substitute custom parameters at deployment time, follow this process.

1. Define a custom parameter in the Application Profile.
2. Verify that this parameter has the same name as the parameter defined in the service of choice.
3. When end users deploy the application, they can enter the parameter value required for this deployment and override the default value defined in the service.

To substitute a user-defined parameter with a macro, follow this process:

1. Add a new custom parameter (provide a parameter name, display name, description, type, and an optional default value for this parameter).
2. Use the parameter name to reference this parameter inside a script or configuration file as a macro using Workload Manager convention. For example:
 - Parameter name: *myParameterName*
 - Macro reference for this parameter: *%myParameterName%*
3. Reference the macro inside the Application Profile.
4. If the parameter is editable, the person executing the deployment for this application is queried for the parameter value.

Some valid use cases of how the macro is replaced by the value specified in the parameter are provided in the following examples.

See [Deployment Lifecycle Scripts](#) for additional details on using scripts.

Pass the Macro

Pass the macro as an Argument in a Specified Script

Properties

General Settings

Service Initialization

Specify service initialization scripts

Pre-Start Script

Script or Command

-

cfg.sh %MyCustomParameter%

229

Script executed before the service is started. Please type in command(s) or path(s) to your script(s) (semicolon separated)

Post-Start Script

----Select One Location----

Script executed after the service is started.

Pre-Stop Script

Insert Macro in Configuration File

Insert the macro (in this case, the %DB_TIER_IP% macro) in a configuration file, if used. Here the private IP address of the database tier is replaced by the parameter value. Refer to a custom tier.

```

installation.type=auto
mediafiles.storage.dir=/usr/local/rollerdata/mediafiles
search.index.dir=/usr/local/rollerdata/searchindex
log4j.appender.roller.File=/usr/local/rollerdata/roller.log

database.configurationType=jdbc
database.jdbc.driverClass=com.mysql.jdbc.Driver
database.jdbc.connectionURL=jdbc:mysql://%DB_TIER_IP%:3306/rollerdb?
autoReconnect=true&useUnicode=true&characterEncoding=utf-8&mysqlEncoding=utf8
database.jdbc.username=scott
database.jdbc.password=tiger
mail.configurationType=properties
mail.hostname=smtp-server.example.com
mail.username=scott
mail.password=tiger

```

Refer to a Custom Tier from the Script

Refer to a custom tier (DB, in the following example) directly from the script by referencing a parameter (private IP address of the DB tier) that is passed as an environment variable (\$CliqrTier_DB_IP) to the cloud during orchestration:

```

#create config file put in DB IP and change permissions
sudo cp $APP_DIR/openerp/openerp-server.conf /etc/
sudo sed -i "s/DB_IP/$CliqrTier_DB_IP/g" /etc/openerp-server.conf

```

Next, call the configuration file (with the defined variable) from the Application Profile similar to what displayed in the following screenshot:



AppCluster

Minimum Number Of Nodes *

1

Maximum Number Of Nodes *

2

App Run-Time *

JDK 6

App Config Files

classes/roller-custom.properties 215

Application config files that contain Cliqr system tokens and will be modified at deployment time. The config file is a relative path from the webapp context folder, e.g., WEB-INF/classes/db.conf. If there are multiple files, separate with semicolon.

The following examples provide use cases of how the Workload Manager-defined parameters replace the value specified in the parameter.

Automate IP Address Configuration

Use the IP address parameters to automate IP address configuration inside configuration files. See [Configuration Files](#) for further details.

Automate Job Picking

Use the TASK_INDEX variable to automate the task of picking the right jobs for batch computing tasks. For example:

```
query%TASK_INDEX%_result.txt
```

picks the following files for each task execution:

- query1_result.txt
- query2_result.txt
- And so forth

Define Path in Initialization Script

Use the %APP_DIR% macro in the initialization script to define the right path (user-defined or marketplace path) for each application:

```
sudo %APP_DIR%/openerp/setup-openerp.sh -APP_DIR %APP_DIR%
```

The runtime application binaries pick the applicable path automatically. See [Deployment Lifecycle Scripts](#) > *Script Source Details* for additional context.

Replace Deployment Environment

Use the %CliqrDepEnvName% macro to replace the deployment environment name during runtime. If you define the path for application binaries, script, or data files to be:

```
/global/collaborate/%CliqrDepEnvName%
```

If the deployment environment is dev, Workload Manager automatically selects the files from the following location at runtime:

```
/global/collaborate/dev at runtime.
```

Define Output Path Files

Use the %OUTPUT_DIR%/DATE%JOB_NAME% macros to define the path of output files. If you deploy a JMeter application with jmeter_001 job name and launch it on March 23rd, 2014, you can obtain the output/log files at:

```
/shared/output/20140322/jmeter_001
```

Back to: [Parameters and Macros](#)

Configuration Files

Configuration Files

- [Overview](#)
- [Profiles Using Varied Parameters](#)
- [Configuration File Using the Macro](#)

When modeling applications using either application packages or images, you may need to call application-specific install and configuration scripts (config files) when deploying the application. Specifically, these files or scripts may have hard-wired values for configuration parameters that may need to be changed at deployment time. These parameters represent settings that must be reconfigured when your application is deployed on the target cloud.

If you use parameters (either [Pre-Defined Parameters](#) or [Parameter Substitution](#)), you may need to modify the configuration files to reflect the correct property defined in the relevant parameter.

[Pre-Defined Parameters](#) are designed to cover cases where deployment owners may not know settings values in advance as those values are dynamically generated by the underlying infrastructure or is specific to the deployment environment or is dependent on the system. Based on the application deployment context, *cf* automatically replaces appropriate values at runtime for each parameter referenced in your application configuration files. Alternately, you can pass the parameters as arguments to install/configuration scripts. See [Parameter Substitution](#) for additional context.

For example, an N-tier Web Application may require the following parameters or macros:

- `%DB_TIER_IP%`: private IP address of the Database node (tier name = *Database*)
- `%NoSQL_TIER_IP%`: private IP address of the NoSQL Database node (tier name = *NoSQLDatabase*)
- `%MB_TIER_IP%`: private IP Address of the MessageBus node (tier name = *MessageBus*)
- `%BC_TIER_IP%`: private IP address of the BackendCache node (tier name = *BackendCache*)



These parameters are for used for legacy compatibility purposes. Be sure to use the exact name for each underlying tier: *Database*, *NoSQLDatabase*, *MessageBus*, and *BackendCache*.

To replace tokens in the configuration file, use the `%CliqrTier_TierName_IP%` token, where *TierName* is the name of the dependent tier.

In the above example, you can use the `%DB_TIER_IP%` macro as a placeholder for the IP address of the database tier, within your configuration file. If you do so, then Workload Manager (at deployment time) will replace the `DB_TIER_IP` parameter with the actual private IP address on which the database is accessible to this application.

The following example displays a configuration file using this macro:

```
../WEB-INF/classes/myconfig-file
installation.type=auto
mediafiles.storage.dir=/usr/local/rollerdata/mediafiles
search.index.dir=/usr/local/rollerdata/searchindex
log4j.appender.roller.File=/usr/local/rollerdata/roller.log
database.configurationType=jdbc
database.jdbc.driverClass=com.mysql.jdbc.Driver
database.jdbc.connectionURL=jdbc:mysql://%DB_TIER_IP:3306/rollerdb?
autoReconnect=true&useUnicode=true&characterEncoding=utf-8&mysqlEncoding=utf8
database.jdbc.username=scott
database.jdbc.password=tiger
mail.configurationType=properties
mail.hostname=smtplib-server.example.com
mail.username=scott
mail.password=tiger
```

Prior to macro substitution, the configuration file for the above application had the following description in the highlighted area:

```
database.jdbc.connectionURL=jdbc:mysql://localhost:3306/rollerdb?
```

Effectively, the dynamically generated private IP address (of database tier) replaces the *localhost* IP address during runtime by using the `%DB_TIER_IP%` macro value.



For supported services, if you modify the configuration file, save the file in the same package format and extension as the original file. For example, save as a .war package (for Tomcat), .zip package (for Apache), .sql package (for MySQL), and so forth.

Back to: [Parameters and Macros](#)

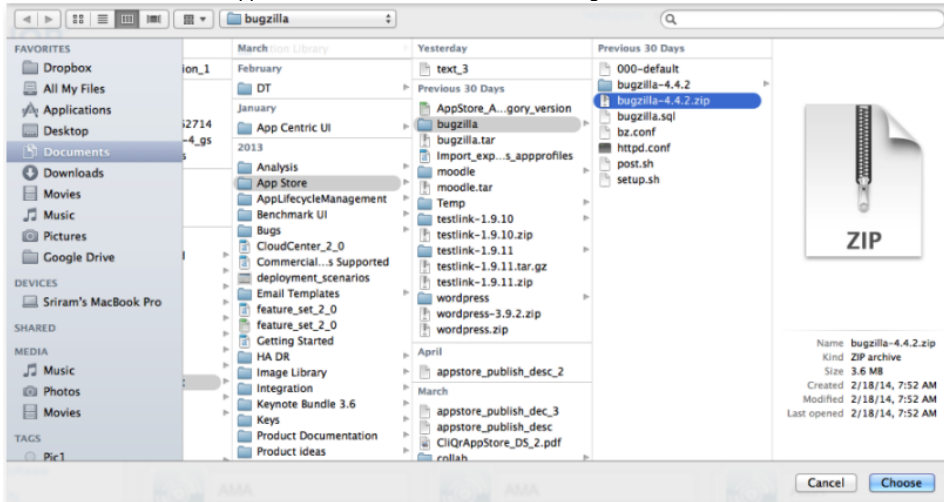
Application Using Imported Profile

Model an Application by Importing a Profile

You can model applications by importing profiles for applications that use the Workload Manager format. Workload Manager supports this process if the metadata information for this application conforms to the Workload Manager application definition format.

To import application profile that already uses the Workload Manager format, follow this process:

1. From the **Applications** page, click **Import App**. The browser window displays.
2. Select the ZIP file for the application that uses the Workload Manager format.



You may need to upload the application binaries and data files before deploying the application.

3. Workload Manager validates the format and displays the application in the Apps tab.
4. The imported profile is now available for deployment.

Application Using App Package

Model an Application Using a File in Package

- [Overview](#)
- [Use Case](#)
- [Script Reference](#)
- [Best Practices](#)

All Artifacts (scripts, binaries, and so forth) for a profile are now bundled into a single .zip file. This allows scripts to be available in the application VMs. When you model or deploy an application, select the **File in Package** option from the Service Initialization scripts (see [Understand Application Tier Properties](#) for additional context) dropdown menu and specify the file name.

Scripts for each VM and its associated scripts can be bundled into a single ZIP file and referenced as an application content package in the application profile. This packaging allows for easier management of the scripts as well as the ability to enable scripts to refer to each other using relative paths without having to contain explicit download logic.

The Application package file field in the Topology Modeler's Properties General Setting field refers to the path for the following:

- The binaries for the web service. The file is in a relative path from `http://env.ciqrtech.com/`.
- The scripts for the database service.

Scripts for each VM and its associated scripts can be bundled into a single ZIP file and referenced as an application content package in the application profile. This packaging allows for easier management of the scripts as well as the ability to enable scripts to refer to each other using relative paths without having to contain explicit download logic.

The application package extracts to a temporary folder, but any scripts inside are executed from the current user's home directory (**/home/cliqruser** by default).

Follow these best practices when using application packages:

- Before you package the scripts in a ZIP file, ensure the script has set proper execute permission.
- When terminating scripts, ensure to terminate the script with an exit statement and the appropriate return code.
- To reference one script from another, look for the directory of the currently executing script and reference the other script from there:
 - To store this script's directory:
SCRIPT_DIR=`dirname "\$0`
 - To point to the other script relative to this script's path
bash \$SCRIPT_DIR/script2.sh

Application Using Puppet or Chef

Model An Application Using Puppet or Chef

- [Overview](#)
- [Out-of-Box Service Additions](#)
- [Associate a Puppet Service](#)
- [Associate a Chef Service](#)
- [Puppet Deployment Process](#)
- [Chef Deployment Process](#)

CloudCenter enables enterprises to integrate with Puppet and Chef, as a client agent that is installed on every application VM using CloudCenter, and automates the configuration management process. This automation helps enterprises by not requiring admins to manually configure the environment each time a new VM is launched by combining the VM provisioning and the [Service Lifecycle Actions](#) to provide seamless end-to-end integration at deployment time.

Puppet and Chef are available as [OOB Services](#) on the CloudCenter platform.

When modeling an application, drag and drop the Puppet and Chef services in the [Topology Modeler](#) and add app-specific parameters in the **Properties > General Settings** tab. See the following sections for service-specific details

- [Puppet Service](#)
- [Chef Service](#)

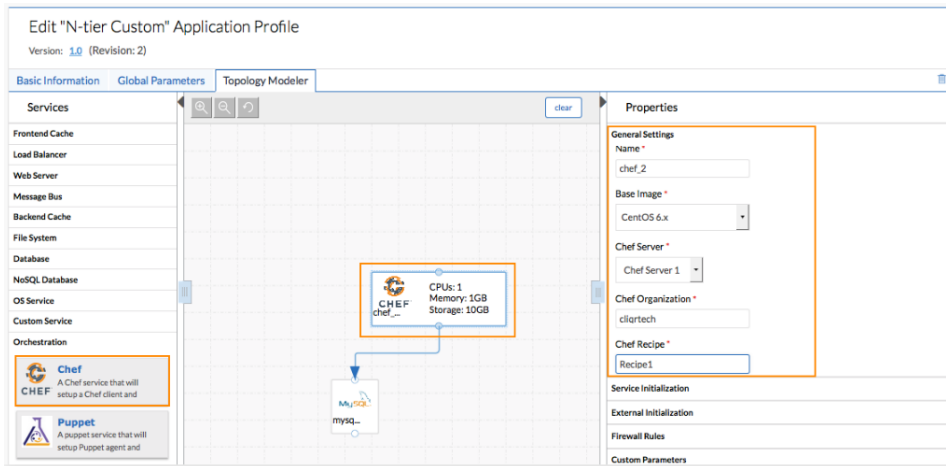
To associate a Puppet service with the application tier by following this process:

1. Select a Puppet primary server that is already registered as a repository (see share [Artifact Repository](#) for additional context).
2. Provide the required information to setup the Puppet Agent – the role and repository path.
3. Deploy the Puppet app – the service script leverages the function of Puppet Agent to configure the environment and deploy the application. Once deployed, the VM communicates with the selected Puppet primary server to retrieve data. The Puppet primary server compiles the manifests (see table above) into system-specific catalogs per the agent requests, and sends the catalogs to the agent to configure the environment.

The screenshot displays the 'Edit "N-tier Custom" Application Profile' interface. The 'Topology Modeler' tab is active, showing a diagram with a 'Puppet' service node connected to a 'MySQL' service node. The 'Puppet' node is highlighted with an orange border. On the right, the 'Properties' panel is open to the 'General Settings' section, which includes fields for 'Name' (puppet_3), 'Base image' (CentOS 6.x), 'Puppet Servers' (Puppet Master), 'Puppet Role' (Wordpress), and 'Puppet Environment'. Below this, there are sections for 'Service Initialization', 'External Initialization', 'Firewall Rules', and 'Custom Parameters'.

To associate a Chef service with the application tier by following this process:

1. Select a Chef Server to communicate with the Chef Client.
2. Provide the required information to set up the Chef Server – the *recipe* (see table above) and environment name.
3. Deploy the Chef app – the service script invokes the Chef Client and registers itself to the selected Chef Server. After receiving the recipe, the client configures the environment.



1. Set up a Puppet primary server, install and configure the Puppet primary server open source version, and host the Puppet modules (for example, Apache, MySQL, Tomcat). Refer to the Puppet website for detailed information.
 2. Create Puppet profiles and roles and set up the configuration. Refer to the Puppet website for detailed information.
 3. Specify the Certname pattern on the Puppet primary server to allow certified access from the Puppet Agent. Refer to the Puppet website for detailed information.
 - a. Hostname: The hostname of the Puppet primary server.
 - b. Certname Suffix: The suffix used by the CloudCenter Puppet service to generate the Certname for authentication to the Puppet primary server.
 4. Create and Share an [Artifact Repository](#) for Puppet and include the metadata of the Puppet primary server.
 5. [Model an Application](#) using a Puppet tier. Configure the Puppet primary server parameters that will be used on the Puppet Agent.
 6. [Deploy an Application](#) to invoke the Puppet Agent and set up the app.
-
1. Set up a Chef Server and host Chef cookbooks and recipes (for example, Apache or Tomcat). Refer to the Chef website for detailed information.
 2. Create and Share an [Artifact Repository](#) for Chef and include the metadata of the Chef Server. You will need to download the validation key from the Chef Server Web page. If the Chef server is configured with SSL, specify those credentials in the Chef repository.
 - a. Chef User Key: The private key used to manage the Chef client.
 - b. Chef Validation key: Used to authenticate the Chef Client when it communicates with the Chef Server for the first time.
 - c. Trusted Certificate: The SSL certificate of the Chef Server used to authenticate the Chef Client to the Chef Server.
 3. [Model an Application](#) using a Chef tier. Configure the Chef Server and recipe that will be used on the Chef Client.
 4. [Deploy an Application](#) to invoke the Chef Client and set up the app.

Deploy Applications

Deploy Applications

- [Deploy an Application](#)
- [JSON for the Deploy API Call](#)

Deploy an Application

Deploy an Application

- [Overview](#)
- [Prerequisites](#)
- [Process](#)
- [Cloud-Specific Configuration](#)
- [Deployment Process](#)

Once you model the Application Profile, you can deploy it to a single cloud, or if a multi-tier application, you can deploy each tier to a different cloud (hybrid deployment).

Workload Manager orchestrates the following deployment tasks:

- Provisions and configures cloud infrastructure and services as determined in the Application Profile (for example, compute, storage, networking)
- Launches VMs and mounts them to the storage repository on the target cloud so that the VMs can access the specified application packages, data, and scripts
- Deploy each application component (applies to different tiers in a multi-tier application)
- Start application services in the correct order based on service dependencies

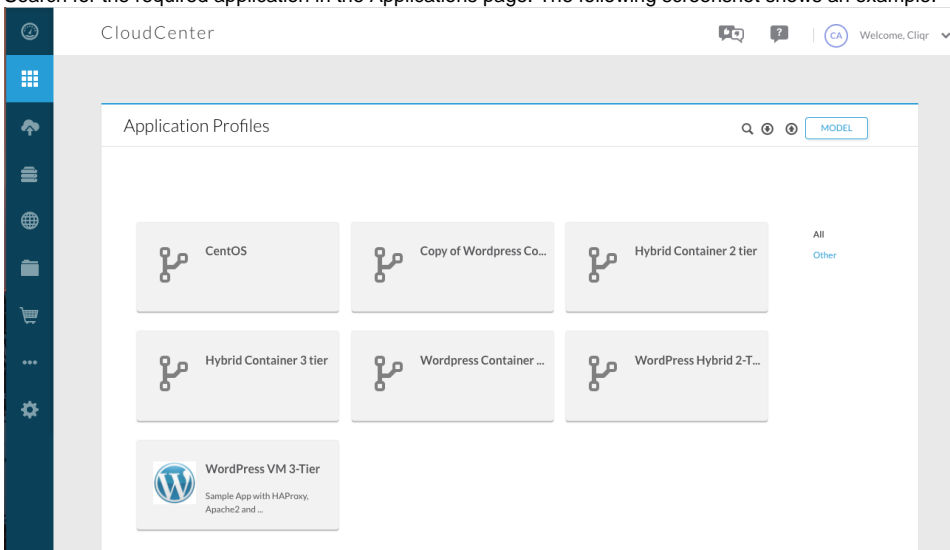
During orchestration, Workload Manager sends status messages to the UI so that you can view and track the application deployment status.

To deploy an application, you should have already performed the following tasks:

- Created an application: See [Create an Application Profile](#) for additional context.
- Created a deployment environment: See [Deployment Environments](#) for additional context.

To deploy an application, follow this process:

1. Access the application from the Workload Manager UI and click **Applications**.
2. Search for the required application in the Applications page. The following screenshot shows an example.



3. To deploy an application, left click on it. This brings up Page 1 of the Deploy form, as shown in the following screenshot.

General Settings

* DEPLOYMENT NAME
DocDeploy

* APPLICATION VERSION
6

DEPLOYMENT ENVIRONMENT

CCQA_deploy_env kubernetes sidenv

AGING POLICY
aging_policy_time

SUSPENSION POLICY
suspension_policy_BA_11P

* TIMEZONE
(GMT-07:00) Mountain Standard Time/Pacific Daylight Time

TAGS

Metadata

NAME	VALUE	ACTIONS
META_EDIT	Editable	
META_MAND	Mandatory	
META_MAND_EDIT	Mandatory and Editable	
Name	%JOB_NAME%	

ADD METADATA

In the **General Settings** section, you must give the deployment a name. You can also optionally select a deployment environment, policy or policies, timezone, tags, terminate/suspend protection, and add metadata as required. The deployment environment boards display the name, description, and list of clouds available in this environment.

Tags do not influence the application of policies and users can enter custom tags in the deploy form in addition to selecting from the pre-populated tags for both the application profile and the deployment environment. Additionally, users can delete tags if they are not locked.

i While the Aging, Suspension, and Security policies dropdown options are populated based on the deployment environment, the Scaling policies dropdown options are populated based on the application profile. See [Policy Management](#) for additional context.

The tags are populated based on both the deployment environment and the application profile. You can also add on custom tags for the entire deployment and for each tier in the deploy form at deploy time. See [System Tags](#) for additional context.

⚠ Google Cloud Nuance

Google Cloud does not support attachment of tags to VMs. Although the Workload Manager UI will allow tags to be specified, and shows success, tags are not added.

4. Scroll down to the enter any global parameters, if defined in the [New Application Profile](#), and any per tier settings. The following screenshot

shows an example.

Global Parameters

GLOBAL_PARAM_STRING ⓘ

* GLOBAL_PARAM_NUMBER ⓘ

* GLOBAL_PARAM_LIST ⓘ

* GLOBAL_PARAM_PWDWDC ⓘ

* CONFIRM PASSWORD

* GLOBAL_PARAM_WS ⓘ

GLOBAL_PARAM_PATH ⓘ

* GLOBAL_PARAM_TEXT ⓘ

Tier Settings

CentOS_1

* NUMBER OF NODES
 max:10

STORAGE SYSTEM
 None available

TAGS ⓘ

You can optionally provide per-tier settings for each tier in your application. Keep scrolling down to see all tiers. Some tiers will have settings for other parameters based on the Application Profile requirements.

Tier Settings

CentOS_1

* NUMBER OF NODES
1 max:10

STORAGE SYSTEM
None available

TAGS ⓘ
[]

PARAMETERS

* TIER_PARAM_NUMBER ⓘ
999999

TIER_PARAM_LIST ⓘ
ZXC

* TIER_PARAM_PWDWC ⓘ
[]

* CONFIRM PASSWORD
[]

TIER_PARAM_PWD ⓘ
[]

* TIER_PARAM_TEXT ⓘ
I@m a Tier T3xt Area with "-!@#%&*_+=[]|:~<>/?";

TIER_PARAM_WS ⓘ
None

ENVIRONMENT VARIABLES

ENV_PARAM_VISIBLE
Visible Only

* ENV_PARAM_EDITABLE
Visible and Editable

* ENV_PARAM_OPTIONAL
Optional

NEXT

5. Click **Next** to go to Page 2 of the Deploy form and select the cloud provider and then the cloud account, as shown in the following screenshot.

The screenshot displays two sections of the deployment form:

- General Cloud Settings:** This section is titled "SELECT A CLOUD" and features a large blue button with the AWS logo and the text "CCQA_AWS" and "US East (Virginia)". Below this is a "CLOUD ACCOUNT" dropdown menu with "CCQA_AWS_cloudaccount" selected.
- Tier Settings:** This section is titled "CentOS_1" and shows a price of "\$ 0.0116 /hour". It includes a "Filter Instance Types / SHOW" link and a list of "AVAILABLE INSTANCE TYPES (101)". Four instance types are visible:
 - T2.MICRO:** 1 VIRTUAL CPU, 1 GB MEMORY, 0 GB TEMP STORAGE. Price: \$ 0.0116 /hour (approx 8.468/month).
 - T2.SMALL:** 1 VIRTUAL CPU, 2 GB MEMORY, 0 GB TEMP STORAGE. Price: \$ 0.0230 /hour (approx 16.79/month).
 - M1.SMALL:** 1 VIRTUAL CPU, 1 GB MEMORY, 160 GB TEMP STORAGE. Price: \$ 0.0440 /hour (approx 32.12/month).
 - T2.MEDIUM:** 2 VIRTUAL CPU, 4 GB MEMORY, 0 GB TEMP STORAGE. Price: \$ 0.0464 /hour (approx 33.872/month).

6. Scroll down further to specify the tier settings for each tier. The available instance types, cloud settings, and SSH options shown in the tier settings of the deploy form is determined by the corresponding Default Tier Cloud Settings in the Deployment Environment form as shown in the screenshot below.

Tier Settings

CentOS_1 \$ 0.0116 /hour

Filter Instance Types / SHOW

AVAILABLE INSTANCE TYPES (101)

T2.MICRO	T2.SMALL	M1.SMALL	T2.MEDIUM
1 VIRTUAL CPU 1 GB MEMORY 0 GB TEMP STORAGE	1 VIRTUAL CPU 2 GB MEMORY 0 GB TEMP STORAGE	1 VIRTUAL CPU 1 GB MEMORY 160 GB TEMP STORAGE	2 VIRTUAL CPU 4 GB MEMORY 0 GB TEMP STORAGE
\$ 0.0116 /hour approx 8.468/month	\$ 0.0230 /hour approx 16.79/month	\$ 0.0440 /hour approx 32.12/month	\$ 0.0464 /hour approx 33.872/month

Volumes

VOLUME	PRICE	SIZE	TYPE
Root Disk	\$0.10/month	GB	General Purpose SSD
Additional Volume 1	\$2.00/month	20 GB	General Purpose SSD

Cloud Settings

INSTANCE PROFILE ARN
arnaws:iam::123456789012:instance-profile/my-instance

* VPC
vpc-3e41595b (vpc-default-us-east-1) | CIDR 172.31.0.0/16

NIC 1
* NETWORK
subnet-1fafee46 | us-east-1a | CIDR 172.31.16.0/20

* PRIVATE IP ALLOCATION
DHCP

ASSIGN PUBLIC IP
 ON

NETWORK INTERFACE CONTROLLER

SSH Options
 No Preference Assign Public Key Persist Private Key

You may use the default values or modify them when alternate choices are available. In the above example, detailed fields for the cloud settings are displayed because the simplified network option was not selected in the deployment environment, and the cloud settings were left unlocked and visible. If the deployment environment had simplified networking enabled, the various fields in the Cloud Settings section are replaced by a single Network Mapping dropdown field where each choice in the dropdown corresponds to one of the network mappings defined in the deployment environment. See [Deployment Environments](#) for more details. The Volume settings shown in the Tier Settings depend on the cloud family and the application profile

7. Click the Deploy button to deploy this application immediately, or click the clock icon to the right of the Deploy button to schedule deploying this application in the future.

If you choose to schedule the deployment, you also have the option to automatically terminate the deployment at a certain time by turning on the **End Date** toggle in the Schedule Deployment dialog box as shown in the screenshot below.



If you specify an end date in the Schedule Deployment dialog box, the deployment will terminate at that time even if terminate protection is turned on for that deployment.

This section provides details on configuring the **Cloud Settings** section and the **Network Settings** section for each cloud.



If you do not see the required instance type listed in this section, be sure to add the instance type first. See [Manage Instance Types](#) for additional context.

Once you configure the cloud settings, you have two control options to identify if the **Cloud Settings** should be:

- **Visible** (Default) to your end users by toggling the control switch.
- **Configurable** (Default) by your end users by toggling the control switch.

Toggle the required settings as required for your deployment environment defaults.

- **VMware Network Settings**

- a. Toggle the **Visibility** switch to determine if you want to allow your end users to use pre-configured settings.

- **OFF:** (Default) End users are not allowed to use preconfigured ACI extensions.
 - i. Select the Network in the NIC section. See [IP Address Allocation](#) for additional context on NIC configuration.
 - ii. Add additional NICs, if required.
- **ON:** End users are allowed to use preconfigured ACI extensions.
 - i. Select the required extension, the corresponding options are displayed in the dropdown list for the remaining fields (see [Extensions](#) for additional details):
 - ii. Select the APIC Extension from the dropdown list (see [ACI Extensions](#) for additional details).
 - iii. Select the APIC **Virtual Machine Manager** (VMM) associated with this APIC Extension from the filtered dropdown list.
 - iv. Select the **APIC Tenant** associated with this APIC Extension from the filtered dropdown list.

- b. Select the Network in the NIC section.

- If you select **VMware**, select the Network in the NIC section. See [IP Address Allocation](#) for additional context on NIC configuration.
- If you select **Cisco ACI**, select the type in the **End Point Group (EPG) Type** field.
 - i. **Existing EPG:** If you select this type, you must further select a pre-existing EPG (that is already connected to one of the Bridge Domains) from the **Existing EPG** dropdown, which appears if you select this type.
 - ii. **New EPG:** If you select this type, you must further select a pre-existing Bridge Domain (to which this EPG must connect) from the **Bridge Domain** dropdown list.

- iii. **Bridge Domain Template:** See [Extensions](#) for additional context.
- c. Add additional NICs, if required.

Back to:

- [Deploy an Application](#)

• **OpenStack Network Settings**

- a. Toggle the **Visibility** switch to determine if you want to allow your end users to use pre-configured settings.
 - **OFF:** (Default) End users are not allowed to associate the public IP with the NIC.
 - **ON:** End users are allowed to associate the public IP with the NIC.
- b. Select the required **Network** in the NIC section.
- c. Select the **Private IP Allocation** mode in the NIC section. See [IP Address Allocation](#) for additional context on NIC configuration.
 - i. **DHCP:** (Default) This strategy allows the IP to be allocated by the DHCP server to the instance on server boot up. This IP address is not known prior to server boot up.
 - ii. **Preallocate IP:** This strategy allows the cloud infrastructure IP allocation to be dynamically provided before the server boots up.
- d. Add additional NICs, if required.

Back to:

- [Deploy an Application](#)
- [OpenStack Configurations](#)

• **Google Network Settings**

The NIC configuration and Simplified Networks configuration is available for both networks and sub-networks. The CloudCenter platform only supports a single NIC configuration.

- a. Select one of the configured networks retrieved from Google cloud:
 - **Non-legacy Networks** – Select the required network and then select the sub-network-based Google project (or the CloudCenter-specific Shared VPC host project) contained within that network.

The screenshot shows the 'Network Settings' interface for Google cloud. On the left, there's a sidebar with 'Google' and 'us-central1 GoogleAccount Account'. The main area has a 'PROJECT' dropdown set to 'clqimages'. Below it, a box highlights the 'NIC 1' configuration: '* NETWORK' is set to 'test' and '* SUB NETWORK' is set to 'test | CIDR 10.1.0.0/20'. At the bottom, there's an 'ASSIGN PUBLIC IP' section with a 'ON' button.

- **Legacy Networks** – Select the legacy network and the sub-network selection is no longer available.

The screenshot shows the 'Network Settings' interface for Google cloud. On the left, there's a sidebar with 'Google' and 'us-central1 GoogleAccount Account'. The main area has a 'PROJECT' dropdown set to 'clqimages'. Below it, a box highlights the 'NIC 1' configuration: '* NETWORK' is set to 'legacy | CIDR 10.1.0.0/16'. Below this, it says '* SUB NETWORK: The selected network is a legacy network.' At the bottom, there's an 'ASSIGN PUBLIC IP' section with a 'ON' button.

- b. Select the checkbox to indicate if a public IP in the NIC section. See [IP Address Allocation](#) > *Cloud-Specific Nuances* > *Google* for additional context.

Back to:

- [Deploy an Application](#)

- **Kubernetes Cloud Settings**

Select the **Kubernetes Settings** from the **Namespace** dropdown list, if applicable. The list is automatically populated and only requires a selection to be made. Go back to the following pages for related details.

- [Container Clouds](#) for additional cloud-specific details.
- [Deploy an Application](#) for the deployment process

Back to:

- [Deploy an Application](#)
- [Setup Deployment Environments](#)

- **AzureRM Network Settings**

Configure the fields described in the table below for AzureRM cloud regions. These fields are configurable for the following features:

- [Deploy an Application](#)

To configure network settings for AzureRM environments, follow this procedure.

- Toggle the **Visibility** switch to determine if you want to allow your end users to use pre-configured settings.
 - **OFF:** (Default) End users are not allowed to associate the public IP with the NIC.
 - **ON:** End users are allowed to associate the public IP with the NIC.
- Select the required **Subnet** in the NIC section.
- Add additional NICs, if required.

Back to:

- [Deploy an Application](#)
- [AzureRM Configurations](#)
- [Availability Sets and Zones](#)

**AWS Subnet and Deployment Nuances**

Pinning Behavior in AWS Network Settings

When configuring the Deployment Environments defaults for AWS, be aware of the following nuances:

- You have the option to select multiple subnets for the first NIC (NIC1) and among those subnets you can pin one subnet. The pinned subnet becomes the default network for this VM
- Subsequent NICs (NIC2, NIC3, ...) only list subnets that belong to the availability zone of the pinned subnet of the first NIC. Example, in NIC1 the pinned subnet belongs to the availability zone *us-west-1b*, then subsequent NICs only list subnets belonging to the first NIC's *us-west-1b* zone.



Multiple NICs do not span across different zones.

AWS Availability Sets Behavior

During a job deployment - If you launch

- A clustered VM setup, the subnet set is passed as the job payload.
- A single VM setup, the VM is launched as the pinned subnet.

AWS Network Settings

Configure the fields described in this section for AWS cloud regions. These fields are configurable for the following features:

- [Deploy an Application](#)

To configure network settings for AWS environments, follow this procedure.

- Select the required option from the dropdown list for the **VPC** field. See [AWS Configurations](#) for additional context.
- Toggle the **Visibility** switch to determine if you want to allow your end users to use pre-configured settings.
 - **OFF:** (Default) End users are not allowed to associate the public IP with the NIC.

- **ON:** End users are allowed to associate the public IP with the NIC.
- c. Select the required **Network** in the NIC section.
- d. The **Private IP Allocation** mode in the NIC section defaults to **DHCP**. The DHCP strategy allows the IP to be allocated by the DHCP server to the instance on server boot up. This IP address is not known prior to server boot up. See [IP Address Allocation](#) for additional context on NIC configuration.
- e. Add additional NICs, if required.

Back to:

- [Deploy an Application](#)
- [AWS Configurations](#)
- [Availability Sets and Zones](#)

When you submit the application for deployment, Workload Manager reads the application profile and orchestrates the following tasks with one click:

- Provisions and configures cloud infrastructure and services based on the Application Profile requirements (compute, storage, and network).
- Launches VMs to access the application packages, data, and scripts referenced in the profile and mounts them to the storage repository on the selected cloud.
- Deploys the application components for the various tiers or steps as determined in the application profile.
- Starts each application service in the right order based on service dependencies.

During each orchestration stage, Workload Manager sends status messages to the UI for users to track the application deployment phase and status.

After the application is deployed in the cloud, it is listed. Click **Access** <app name> to open the IP address where the application is hosted.

i Note

Depending on the context path of the application in the application server, you may need to specify the full path in the browser (for example, http(s)://<IP address>/app-context-path).



Configure the Launch URL field with this path when creating the application profile.

Once the application is running (in general), Workload Manager provides the ability to securely access cloud VMs using SSH or RDP from any browser. You can use this method to troubleshoot or run one-off custom commands or scripts.



This secure access is not available if you use a custom key for deployments. See [Create a Deployment Environment](#) > *SSH Options* section for additional details.

Simplified Network Settings

Simplified Network Settings

- [AWS Network Settings](#)
- [AzureRM Network Settings](#)
- [Google Network Settings](#)
- [Kubernetes Network Settings](#)
- [OpenStack Network Settings](#)
- [VMware Network Settings](#)

AWS Network Settings

AWS Network Settings

Configure the fields described in this section for AWS cloud regions. These fields are configurable for the following features:

- [Deploy an Application](#)

To configure network settings for AWS environments, follow this procedure.

1. Select the required option from the dropdown list for the **VPC** field. See [AWS Configurations](#) for additional context.
2. Toggle the **Visibility** switch to determine if you want to allow your end users to use pre-configured settings.
 - **OFF**: (Default) End users are not allowed to associate the public IP with the NIC.
 - **ON**: End users are allowed to associate the public IP with the NIC.
3. Select the required **Network** in the NIC section.
4. The **Private IP Allocation** mode in the NIC section defaults to **DHCP**. The DHCP strategy allows the IP to be allocated by the DHCP server to the instance on server boot up. This IP address is not known prior to server boot up. See [IP Address Allocation](#) for additional context on NIC configuration.
5. Add additional NICs, if required.

Back to:

- [Deploy an Application](#)
- [AWS Configurations](#)
- [Availability Sets and Zones](#)

AzureRM Network Settings

AzureRM Network Settings

Configure the fields described in the table below for AzureRM cloud regions. These fields are configurable for the following features:

- [Deploy an Application](#)

To configure network settings for AzureRM environments, follow this procedure.

1. Toggle the **Visibility** switch to determine if you want to allow your end users to use pre-configured settings.
 - **OFF**: (Default) End users are not allowed to associate the public IP with the NIC.
 - **ON**: End users are allowed to associate the public IP with the NIC.
2. Select the required **Subnet** in the NIC section.
3. Add additional NICs, if required.

Back to:

- [Deploy an Application](#)
- [AzureRM Configurations](#)
- [Availability Sets and Zones](#)

Google Network Settings

Google Network Settings

The NIC configuration and Simplified Networks configuration is available for both networks and sub-networks. The CloudCenter platform only supports a single NIC configuration.

1. Select one of the configured networks retrieved from Google cloud:

- **Non-legacy Networks** – Select the required network and then select the sub-network-based Google project (or the CloudCenter-specific Shared VPC host project) contained within that network.

Network Settings

Google

us-central1
GoogleAccount Account

* PROJECT
cliqrimages

* NETWORK
test

* SUB NETWORK
test | CIDR 10.1.0.0/20

ASSIGN PUBLIC IP

CLEAR ALL SETTINGS

- **Legacy Networks** – Select the legacy network and the sub-network selection is no longer available.

Network Settings

Google

us-central1
GoogleAccount Account

* PROJECT
cliqrimages

* NETWORK
legacy | CIDR 10.1.0.0/16

* SUB NETWORK
The selected network is a legacy network.

ASSIGN PUBLIC IP

CLEAR ALL SETTINGS

2. Select the checkbox to indicate if a public IP in the NIC section. See [IP Address Allocation > Cloud-Specific Nuances > Google](#) for additional context.

Back to:

- [Deploy an Application](#)

Kubernetes Network Settings

Kubernetes Network Settings

Kubernetes container configurations do not require additional network settings.

Back to:

- [Deploy an Application](#)

OpenStack Network Settings

OpenStack Network Settings

1. Toggle the **Visibility** switch to determine if you want to allow your end users to use pre-configured settings.
 - **OFF**: (Default) End users are not allowed to associate the public IP with the NIC.
 - **ON**: End users are allowed to associate the public IP with the NIC.
2. Select the required **Network** in the NIC section.
3. Select the **Private IP Allocation** mode in the NIC section. See [IP Address Allocation](#) for additional context on NIC configuration.
 - a. **DHCP**: (Default) This strategy allows the IP to be allocated by the DHCP server to the instance on server boot up. This IP address is not known prior to server boot up.
 - b. **Preallocate IP**: This strategy allows the cloud infrastructure IP allocation to be dynamically provided before the server boots up.
4. Add additional NICs, if required.

Back to:

- [Deploy an Application](#)
- [OpenStack Configurations](#)

VMware Network Settings

VMware Network Settings

1. Toggle the **Visibility** switch to determine if you want to allow your end users to use pre-configured settings.
 - **OFF:** (Default) End users are not allowed to use preconfigured ACI extensions.
 - a. Select the Network in the NIC section. See [IP Address Allocation](#) for additional context on NIC configuration.
 - b. Add additional NICs, if required.
 - **ON:** End users are allowed to use preconfigured ACI extensions.
 - a. Select the required extension, the corresponding options are displayed in the dropdown list for the remaining fields (see [Extensions](#) for additional details):
 - b. Select the APIC Extension from the dropdown list (see [ACI Extensions](#) for additional details).
 - c. Select the APIC **Virtual Machine Manager** (VMM) associated with this APIC Extension from the filtered dropdown list .
 - d. Select the **APIC Tenant** associated with this APIC Extension from the filtered dropdown list.
2. Select the Network in the NIC section.
 - If you select **VMware**, select the Network in the NIC section. See [IP Address Allocation](#) for additional context on NIC configuration.
 - If you select **Cisco ACI**, select the type in the **End Point Group (EPG) Type** field.
 - a. **Existing EPG:** If you select this type, you must further select a pre-existing EPG (that is already connected to one of the Bridge Domains) from the **Existing EPG** dropdown, which appears if you select this type.
 - b. **New EPG:** If you select this type, you must further select a pre-existing Bridge Domain (to which this EPG must connect) from the **Bridge Domain** dropdown list.
 - c. **Bridge Domain Template:** See [Extensions](#) for additional context.
3. Add additional NICs, if required.

Back to:

- [Deploy an Application](#)

JSON for the Deploy API Call

JSON for the Deploy API Call

If using the *Submit Job (v2) API*, you have the added advantage of completing the forms associated with the **New Deployment** workflow and retrieving the corresponding JSON request body for use with the API(s). The following screenshot shows an example.

Deploy "DocTest"
\$ **0.013** /hour
approx \$9.49/month
✕

General Cloud Settings

*** DEPLOYMENT ENVIRONMENT**

Amazon

Select a Cloud

Amazon

US West (Oregon) ▾

Hybrid
Deploy Individual Tiers on different clouds

*** CLOUD ACCOUNT**

Fays-IAM-CCO ▾

Tier Settings

WebServer
\$ **0.013** /hour

Filter Instance Types / SHOW

available instance types (48)

T2.MICRO	T2.SMALL	M1.SMALL	T2.MEDIUM
1 VIRTUAL CPU 1 GB MEMORY 0 GB STORAGE	1 VIRTUAL CPU 2 GB MEMORY 0 GB STORAGE	1 VIRTUAL CPU 2 GB MEMORY 160 GB STORAGE	1 VIRTUAL CPU 4 GB MEMORY 0 GB STORAGE
\$ 0.013 /hour <small>approx \$9.49/month</small>	\$ 0.026 /hour <small>approx \$18.98/month</small>	\$ 0.044 /hour <small>approx \$32.12/month</small>	\$ 0.052 /hour <small>approx \$37.96/month</small>

*** VPC**

vpc-b51cb8d0 (vpc-qa-cases) | CIDR 10.0.0.0/16 ▾

ASSIGN PUBLIC IP

OFF

NIC 1

*** NETWORK**

subnet-1148b548 | us-west-2c | CIDR 10.0.1.0/24 ▾

PRIVATE IP ALLOCATION

DHCP ▾

NETWORK INTERFACE CONTROLLER

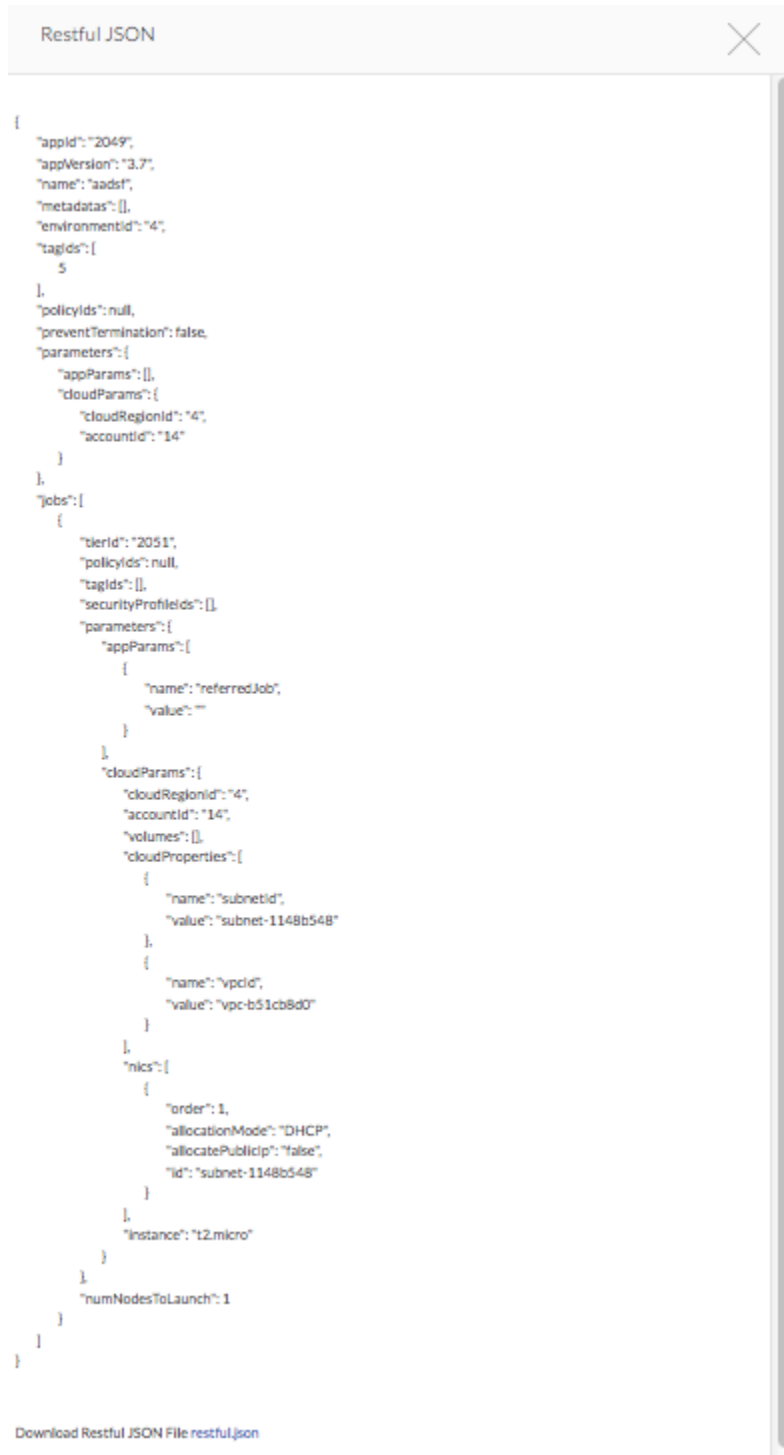
Restful JSON

PREVIOUS

DEPLOY

The **Restful JSON** button (or the **restful.json** link) becomes available when you complete the General Cloud Settings details (with the required selection for the clouds and network settings) and you are able to proceed without any errors or missing fields in this workflow. This button generates the contents of the job deploy page before a submit operation.

This button currently displays for N-tier deployments. When you click the **Restful JSON** button (or the **restful.json** link), a popup responds with the corresponding JSON request body. You can copy the entire REST payload and paste it in your RESTClient application to issue the API call. The following image shows an example.



```

{
  "appId": "2049",
  "appVersion": "3.7",
  "name": "aadsf",
  "metadatas": [],
  "environmentId": "4",
  "tags": [
    5
  ],
  "policyIds": null,
  "preventTermination": false,
  "parameters": {
    "appParams": [],
    "cloudParams": {
      "cloudRegionId": "4",
      "accountId": "14"
    }
  }
},
{
  "jobs": [
    {
      "jobId": "2051",
      "policyIds": null,
      "tags": [],
      "securityProfiles": [],
      "parameters": {
        "appParams": [
          {
            "name": "referredJob",
            "value": ""
          }
        ],
        "cloudParams": {
          "cloudRegionId": "4",
          "accountId": "14",
          "volumes": [],
          "cloudProperties": [
            {
              "name": "subnetId",
              "value": "subnet-1148b548"
            },
            {
              "name": "vpId",
              "value": "vpc-b51cb8d0"
            }
          ],
          "nics": [
            {
              "order": 1,
              "allocationMode": "DHCP",
              "allocatePublicIp": "false",
              "id": "subnet-1148b548"
            }
          ],
          "instance": "t2.micro"
        }
      },
      "numNodesToLaunch": 1
    }
  ]
}

```

Download Restful JSON File restful.json

Benchmark an Application

Benchmark an Application

- [Overview](#)
- [Benchmark Process](#)
- [Load Generator](#)

The Benchmark feature lets you deploy an application profile across multiple cloud providers or cloud regions from a single cloud provider and produces a report that helps you optimize cost and performance.

On the cloud, price and performance are inextricably linked. Knowing one without the other is of little use given inherent trade offs and changing cloud landscape.

Workload Manager automates the benchmark tasks for the price, performance, and price-performance index for each application, across any cloud, instance type or provider, so you can pick your target before you migrate.

The Schedule Benchmark task allows you to run a benchmark for an application at a scheduled time. See [Application Profiles Page](#) > *Schedule Benchmarks* for additional context.

When running benchmarks for legacy applications on VMware or OpenStack clouds, you must wait until the data is loaded. If you enter any information before the data is loaded then the **Submit** button is disabled.

The last page of the Benchmark form is similar to Page 2 of the Deploy form, where you can set the instance type, volumes, and cloud settings for each tier of your application. However, there is one important difference: before the per tier settings for each of the tier of your application is an extra unlabeled tier. This is the load generator tier. Select the instance type, volumes and cloud settings for this tier as you would for other application tiers.



The Load Generator tier is based on the CentOS 6.x logical image. Therefore, for the benchmark deployment to succeed, a CentOS 6.x physical image must exist in each region where the benchmark will be run and the CentOS 6.x logical image must be mapped to the corresponding CentOS 6.x physical image in each of these cloud regions. Furthermore, each mapping must include at least one instance type. See [Images Page](#) for more context.


Policy Management

Policy Management

- [About Policies](#)
- [Share with Sub-tenants](#)
- [Event Policies](#)
 - [Configuring an Event Policy](#)
 - [Event Policy Guidelines](#)
- [Scaling Policies](#)
 - [Configuring a Scaling Policy](#)
- [Aging Policies](#)
 - [Configuring an Aging Policy](#)
- [Suspension Policies](#)
 - [The Savings Feature](#)
 - [Configuring a Suspension Policy](#)
 - [Suspension Policy Guidelines](#)
- [Security Policies](#)
 - [Configuring a Security Policy](#)
 - [Deleting a Security Policy](#)


A policy causes Workload Manager to perform configured activities when certain events or conditions occur. For example, a policy could cause Workload Manager to send an email alert message to a designated administrator if a cloud goes down. You use the Policies window to configure the following types of policies:

- **Event Policy**– Causes Workload Manager to send an email message, invoke a web service, execute a command or script, or perform any number and combination of these activities when a designated event occurs.
- **Scaling Policy** – Causes Workload Manager to increase or decrease VM resources for each application deployment tier that is associated with the policy when one or more designated conditions occur.
- **Aging Policy** – Causes Workload Manager to suspend and optionally terminate each application deployment that is associated with the policy after the deployment has been running for a designated period of time term.
- **Suspension Policy** – Specifies a schedule for when a deployment should be in Running state.


 **Custom Actions** are now defined in a new [Actions Library](#) tab.

A policy is a tenant-owned resource that can be shared with other sub-tenants and users in the hierarchy.

- The users directly under the owner tenant have *Manage* access to the policy as long as they have the Policy permission inherited through a role. However, subtenants and users further down the hierarchy can only have *View* access to the policy in read-only mode, if the policy is shared with subtenants.
- If you do not have the Policy permissions inherited through a role, the Policy tab is hidden from your view.
- All Policy pages have a **Share with Sub-tenants** toggle switch column:
 - **On:** Users in tenants that are further down the hierarchy can only *View* (read-only) these policies, if shared.

 This setting provides visibility to having a role-based permissions for the required policy (see [Permission Control > Role-Based Permissions](#) for details).

- **Off:** Default. No user in any sub-tenant can view or use a policy in this state. The tenant admin can grant permissions to sub-tenants by enabling the **(ON)** toggle switch for the required policy.

 When a deployment environment or application profile is shared with a sub-tenant or users in a sub-tenant, all policies (and system tags) associated with that environment or application profile can still be used in the shared environment or application profile.

An event policy causes Workload Manager to perform one or more configured activities when a designated event occurs for the designated resource. These activities can be:

- **Email**–Sends an email message with the designated subject and body text to the designated recipient or recipients
- **Invoke a web service**–Executes the designated web service request
- **Execute a command**–Executes the designated command or script on an application VM

Check the **Auto Enable for Sub-Tenants** checkbox (only applicable to Event policies) to automatically enable a policy for all shared users and groups. This option determines if the policy should be auto enabled for the sub tenants and users with whom the policy has been shared.

If you select the **Auto Enable for Sub-Tenants** checkbox, you also have the option to select the **Restrict users from disabling this Policy** checkbox to determine if the sub-tenants and users with whom the policy has been shared would be able to disable this policy.

To manage event policies, click **Policies** in the Workload Manager UI main menu and then click the **Event** link to display the Events Policies page.

The Events Policies page lists configured event policies and lets you perform the tasks that the following table describes.

Task	Description
Add a new event policy.	Click the New Event Policy link. See <i>Configuring an Event Policy</i> for details.
View configurations for an existing event policy.	Click the policy name in the Name column. The View page for the policy displays. See <i>Configuring an Event Policy</i> for details. <ul style="list-style-type: none"> Description – Optional description that was configured for the policy Last Updated – Date and time that the policy was last updated Share with Sub-tenants – Described in the <i>Share with Sub-tenants</i> section above.
Enable or disable an event policy.	Click the ON/OFF toggle button in the Enable Column for the policy. When a policy is enabled, it executes when the configured event occurs.
Set permissions for an event policy.	Choose Share from the dropdown list in the Actions column for the policy. See Permission Control > Policy Permissions for details.
Update configurations for an existing event policy.	Choose Edit from the dropdown list in the Actions column for the policy. The Edit page for the policy displays. See <i>Configuring an Event Policy</i> for a description of the fields that you can update. If you make updates, click the Save button on the Edit page to save your changes.
Delete an event policy.	Choose Delete from the dropdown list in the Actions column for the policy.
See a history of event policy executions.	Click the Execution History link to the right of the list of policies. The User Defined Policy Executions page displays. This page lists in reverse chronological order each policy that has executed. For each policy, the page displays the local date and time of execution, the name of the policy, and the name of the resource against which the policy was enforced. You can see details about a policy or resource by clicking the name of the policy or entity.

Configuring an Event Policy

When you add an event policy, you create a new policy based on configuration settings that you make. To add an event policy, click **New Event Policy** on the Event Policy page and then configure the settings that display. Click the **Save** button after you complete the configuration.

The policy must be enabled before it can execute.

The following table describes the settings for configuring an event policy.

Setting	Description
Execute For drop down list	Choose the resource against which the policy is enforced with the configured event occurs
On Event dropdown list	Choose the event that, when it affects the configured resource, causes the policy to execute. See <i>eventName</i> for additional details.

Action Type dropdown list	<p>Enter information in the fields that appear for the action type that you choose as follows. When entering information, you can include variables that are described in the Available Variables section at the top right of the page to customize an email message, web service request, command, or script with dynamic information.</p> <ul style="list-style-type: none"> • Email action type – Enter information in the To, Bcc, Subject, and Body fields that appear, as appropriate. • Invoke a web service action type – Enter information in the Web Service URL, Http Request Type, Content Type, Command Params, and Body fields that appear, as appropriate. See Using Parameters > <i>Parameter Type</i> for additional context. • Execute a Command action type – Enter the desired command or script in the command/script field. <p>When you specify the script or command action type globally, these global policies are executed on all jobs/deployments. You can also configure these global policies to be executed on specific state changes – for example, if a job/deployment goes to a <i>Resumed</i> or <i>Deployed</i> state. On reaching those states, all the VMs for this job/deployment are executed.</p> <ul style="list-style-type: none"> • Launch a new deployment action type – Set the cloud-specific default in the destination deployment environment for the cloud burst configuration. <p>Be aware that the instance types and network defaults are used from the Deployment Environments in the destination job when using a <i>cloud bursting</i> policy. This action type allows other choices as well, however, a cloud bursting policy is only applicable if you use the following settings:</p> <ul style="list-style-type: none"> • Execute for = <i>Application Deployment</i>: • On Event = <i>Maximum cluster size limit</i> • Action Type = <i>Launch a new deployment</i> <p>If you configure multiple instance types in the destination Deployment Environment Defaults, Workload Manager selects the first instance type.</p>
----------------------------------	--

Event Policy Guidelines



Do not use your system administrator credentials to change the user's event policy. This change adversely affects Workload Manager and does not notify the user.

Adhere to the following event policy guidelines:

- If you use a group email and you don't want to disable all email notifications for this policy, **all users** who enabled that policy in Workload Manager must disable it for the group to no longer receive emails.
- If you use a token (for example, %myemail%) and this token is replaced with a user email, then the user with this email must disable the policy in Workload Manager.
- If you create custom email policies, only the default email policy is triggered.

A scaling policy causes Workload Manager to increase or decrease VMs for each application deployment tier that is associated with the policy. You can configure the VM scaling to occur once or daily, or when one or more designated conditions relating to system metrics such as CPU, memory, or network usage occur at a specified polling interval.



VM Metrics Frequency

The frequency of updates for VM metrics depends on the polling interval configured in your scaling policy. Two schedulers are used for each VM:

- The first scheduler fetches the metrics from the operating system (for example, CPU, memory, and so forth at a polling interval frequency defined by user. If the user defines a polling interval at 15-second intervals Workload Manager provides the metrics from the OS at 5-second intervals.
- The second scheduler calculates the average of all values provided by the first scheduler and uploads the metrics to Workload Manager. This upload to Workload Manager occurs at the user-defined, polling interval frequency. Workload Manager does not poll each VM.



CPU Calculation

The CPU threshold is calculated based on the first VM – assuming that all balanced servers have a similar usage. Once the scaling (up or down) is executed, Workload Manager waits for the *breach* period to re-evaluate the policy.

To manage scaling policies, click **Policies** in the Workload Manager UI main menu and then click the **Scaling** link to display the Scaling Policies page.

The Scaling Policies page lists configured scaling policies and lets you perform the tasks that the following table describes.



Task	Description
Add a new scaling policy.	Click the New Scaling Policy link. See <i>Configuring a Scaling Policy</i> for details.


View or update configurations for an existing scaling policy.	<p>Choose Edit from the dropdown list in the Actions column for the policy.</p> <p>The Edit page for the policy displays. See <i>Configuring a Scaling Policy</i> for a description of the fields that you can update. If you make updates, click the Save button on the Edit page to save your changes.</p> <ul style="list-style-type: none"> • Description – Optional description that was configured for the policy • Last Updated – Date and time that the policy was last updated • Share with Sub-tenants – Described in the <i>Share with Sub-tenants</i> section above.
Set permissions for a scaling policy.	<p>Choose Share from the dropdown list in the Actions column for the policy.</p> <p>See Permission Control > Policy Permissions for details.</p>
Delete a scaling policy.	Choose Delete from the dropdown list in the Actions column for the policy.

Configuring a Scaling Policy

When you add a scaling policy, you create a new policy based on configuration settings that you make. To add a scaling policy, click **New Scaling Policy** on the Scaling Policy page and then the settings that display. Click the **Save** button after you complete the configuration.

The following table describes the settings for configuring a scaling policy.

Setting	Description
Name field	Enter a brief and unique descriptive name for the policy
Description field	Optionally enter a brief description of the policy.
Type of scaling radio button	<p>Choose Scheduled Scaling or Elastic Scaling.</p> <ul style="list-style-type: none"> • Scheduled scaling causes the policy to increase and then decrease the number of VMs to the values that you designate according to the schedule that you designate. The schedule can cause the scaling cycle to occur once at the designated dates and times, or it can cause scaling cycle to occur every day at designated times. • Elastic scaling causes the policy to increase or decrease the number of VMs as needed based on configured conditions that the system detects at a specified polling interval and breach period. <div style="border: 1px solid green; padding: 10px; margin-top: 10px;"> <p> The <i>breach</i> period refers to the period of time that Workload Manager waits after one scale up or scale down action, before stabilizing the load.</p> <p>Ideally, the breach period should be greater than (or equal to) the time taken by the scaling operation to launch a VM.</p> <p><i>For example, if the time taken to:</i></p> <ul style="list-style-type: none"> • Launch and configure a node = 5 to 7 minutes • Polling interval = 30 seconds • Breach period should be = 8 to 10 minutes </div>
Add Scaling Schedule fields	Appears if you have selected Scheduled scaling . Configure how often this policy runs, when it runs, and the number of VMs to which to increase and decrease.
Poling Interval fields	<p>Appears if you have selected Elastic scaling. Choose a time length in the left field and choose a time unit in the right field. The system polls system metrics at this interval to determine if the conditions that this policy requires to execute are met.</p> <div style="border: 1px solid green; padding: 10px; margin-top: 10px;"> <p> Poling Interval</p> <p>For example, if you configure a 6-second polling interval, your metrics collected in this time is divided by 3 to obtain the average value. The average value is sent to Workload Manager at the polling interval frequency. So in this example, the CPU data is collected every 2 seconds, and automatically sent to Workload Manager at 6-second intervals.</p> </div>

Breach Period fields	<p>Appears if you choose Elastic scaling. Choose a time length in the left field and choose a time unit in the right field. If the policy executes, it will not execute again for this period of time. In this way, the system can stabilize after a condition occurs without the policy continually adjusting the system.</p> <div style="border: 1px solid green; padding: 10px; margin: 10px 0;"> <p> Breach Period</p> <p>If any VM meets the specified criteria, Workload Manager executes the policy. In a load balanced cluster if one VM's metrics crosses the threshold, it is most likely that the other VM will reach that threshold. If Workload Manager detects any VM crossing the threshold, the policy is executed and the breach period is set. Once the breach period is set, the other VMs cannot execute the policy.</p> </div>
Auto Scale Percentage	<p>Optional. Allows Workload Manager to determine when auto scaling must be triggered, as described in the following table.</p> <p>Auto scaling is triggered when metric results from a <i>minimum number of nodes</i> have crossed the defined threshold.</p> <p>This <i>minimum number of nodes</i> is calculated based on auto scale percentage that is defined as part of the policy metadata.</p> <p>If auto scale percentage is not defined as part of policy metadata, then Workload Manager's Auto Scale Percentage calculation defaults to 70%. This default is defined in the config file.</p> <p>If auto scale percentage is defined as part of policy metadata, then the user-configured Auto Scale Percentage value takes precedence. Distributed locking mechanism handles HA scenarios.</p>
Scale out condition fields	<p>Appears if you have selected Elastic scaling and causes the policy to increase the number of VMs according to the designated rule or rules. The Scale in condition fields cause the policy to decrease the number VMs per the designated rule or rules. For each condition, you configure one or more rule sets. Within each rule set, you configure one or more rules.</p> <ul style="list-style-type: none"> • From the Match dropdown list for each rule set: <ul style="list-style-type: none"> • Choose All to cause the scaling to execute when the situations that are defined by all rules in the rule set occur or • Choose Any to cause the scaling to execute when a situation that is defined by any rule in the set occurs. • To add a rule set, click the ... icon. • To add a rule to a rule set, click the + icon within the rule set. • To remove a rule, click the Trash Can icon under the rule, and then click the yes link.

An aging policy causes Workload Manager to suspend and optionally terminate each application deployment that is associated with the policy after the application deployment has been running for a designated period or reaches a designated deployment cost.



An aging policy and the prevent termination feature cannot be used simultaneously for a deployment because both items control the terminate and suspend behavior of VMs that are mapped to a deployment.

The Prevent Termination feature is only applicable to N-tier jobs. See [Terminate Protection](#) for additional context.

You can configure a grace period for an aging policy, which allows you to keep deployments in suspended state before the automated termination takes place. A grace period is particularly useful when approvals via ServiceNow are set up for extensions. You also can add extensions to a policy, which allows deployments to keep running as needed.

If an aging policy is configured to terminate a deployment when the time or cost limit is reached, you optionally can configure a grace period, which allows you to keep deployments in suspended state before the automated termination takes place. A grace period is particularly useful if approvals via ServiceNow are set up for extensions. You also can add extensions to a policy, which allows deployments to keep running by adding more time/cost to permit runtime.

An aging policy can include notifications. A notification is an e-mail message to a deployment owner that informs the owner that the deployment is going to terminate or suspend or that a grace period for a deployment is going to expire. You can configure how far in advance of a termination or grace period expiration the system sends the message, and you can send additional messages as reminders.

You cannot modify duration, extension, or grace period settings for an Aging policy if a deployment is running with the policy enforced. You can modify notification options, which apply to future policy executions. If you need to modify duration, extension, or grace period settings, you can do so only when deployments that use the policy are no longer running or if you first manually remove the policy from all deployments that are using it.

The Aging Policy field in the Deploy form differs based on the following choices:

- **Off** = A dropdown for selecting an aging policy
- **On** = Displays the aging policies associated with tags that are assigned to the deployment

The Deployment Details page contains an Aging Policy dropdown with the following choices:

- **Change Policy** lets you replace the current aging policy with another one that you pick from a list. If a policy specifies a time or cost that is less than what has accrued for the deployment, the policy is not available.
- **Add Policy** lets you add one from a list if one is not associated.
- **Remove Policy**

To manage aging policies, click **Policies** in the Workload Manager UI main menu and then click the **Aging** link to display the Aging Policies page.

The Aging Policies page lists configured aging policies and lets you perform the tasks that the following table describes.

Task	Description
Add a new aging policy.	Click the New Aging Policy link. See <i>Configuring an Aging Policy</i> for details.
Set permissions for an aging policy.	Choose Share from the dropdown list in the Actions column for the policy. See Permission Control > Policy Permissions for details.
View information about an aging policy.	You can view the following information for each policy: <ul style="list-style-type: none"> • Description – Optional description that was configured for the policy • Last Updated – Date and time that the policy was last updated • Share with Sub-tenants – Described in the <i>Share with Sub-tenants</i> section above. • Extensions – If one or more extensions are configured for the policy, the number of extensions appears, followed by the time length or cost limit of each extension, in parentheses • Notify – Indicates whether email notification is configured for the policy • Age By – With the choice of how this policy should be aged, Time Duration or Cost Limit.
View or update configurations for an existing aging policy.	Choose Edit from the dropdown list in the Actions column for the policy. The Edit page for the policy displays. See Configuring an Aging Policy for a description of the fields that you can update. If you make updates, click the Save button on the Edit page to save your changes.
Delete an aging policy.	Choose Delete from the dropdown list in the Actions column for the policy
Enable or disable an aging policy.	Turn on or off the Enable switch for the policy. When a policy is disabled, future deployments cannot use it. The policy remains in effect for existing deployments.

Configuring an Aging Policy

When you add an aging policy, you create a new policy based on configuration settings that you make. To add an aging policy, click **New Aging Policy** on the Aging Policy page and then configure the settings that display. Click the **Save** button after you complete the configuration.

The following table describes the settings for configuring an aging policy.

Setting	Description
Age By buttons	Click Time Duration if you want to suspend and optionally terminate an application deployment after it has been running for a designated period, or click Cost Limit if you want to suspend or terminate a deployment when it reaches a designated deployment cost.
Policy Name field	Enter a brief and unique descriptive name for the policy
Policy Description field	Optionally enter a brief description of the policy.
Policy Duration field	Appears if you choose a time duration policy type. Enter the duration as a number of units, such as 15 days.
Policy Cost Limit field	Appears if you choose a cost limit policy type. Enter the cost amount.
Terminate Deployment after Policy Duration switch	Turn on if you want the deployments that are associated with the policy to terminate when the configured time duration or cost limit is reached. By default, this switch is turned off and policies suspend but do not terminate when a duration or limit is reached.
Allow a Grace Period before Terminating switch	Appears if you choose to terminate deployments after the configured time duration or cost limit. Turn on this switch if you want to provide a grace period before a deployment is terminated. A grace period is designated amount of months, days, or hours after configured time duration or cost limit is reached that deployments remain suspended before terminating. Enter the grace period as a number of units, such as 5 days, in the Grace Period fields. For example, if you choose to terminate deployments after 10 days and configure a grace period of 5 days, deployments become suspended when the 10 day duration is reached, remain suspended for 5 days, and then terminate.

Allow Extensions to this Policy switch	Turn on if you want to provide the option of extending the time that the policy is in effect beyond the originally configured time duration or cost limit or a grace period
No. of Extensions field	Appears if you allow extensions for the policy. For a policy that is to terminate a deployment after a designated period, enter duration of each extension as a number of units in the Length of Each Extension field. For a policy that is to terminate a deployment after a designated deployment cost, enter the additional cost that is allowed for each extension in the Cost Limit of Each Extension field. See Extensions for additional context.
Notify Deployment Owner of Policy Expiry switch	Turn on if you want the system to send an email message to a deployment owner notifying the owner that the deployment is going to terminate. In the Notify fields that appear, enter, as a number of units, how long before a deployment terminates the emails should be sent. For example, you could enter 3 days. If you want to configure multiple messages regarding deployment termination, click Notification Alert and enter the number of units for each additional message. For example, in this way, you can configure the system to send messages 3 days, 2 days, and 1 day before termination. Edit text in the email body field, if needed, and optionally click View Sample E-Mail to see how the email message will appear to recipients.
Notify before Grace Period Ends switch	Turn on if you want the system to send an email message to a deployment owner notifying the owner that a grace period for a deployment is going to expire. In the Notify fields that appear, enter, as a number of units, how long before the grace period expires the emails should be sent. The period that you configure cannot be longer than the configured grace period so that the system does not send messages after the grace period expires. For example, you could enter 3 days. Edit text in the email body field, if needed, and optionally click View Sample E-Mail to see how the email message will appear to recipients.

Suspension policies are designed to reduce cloud cost by suspending all tiers in a deployment at specific times – you can use a suspension policy to specify a schedule for when a deployment should be in the Running state. At other times, the deployment remains suspended. You can use a suspension policy to put a deployment in the Running state every day during a certain time period, or on specific days during a certain time period.

For example, you could configure the policy to put a deployment in a Running state every day from 8:00 a.m. through 9:00 p.m., or on Mondays, Wednesdays, and Fridays from 10:00 a.m. through 6:00 p.m. You also can specify one or more *blockout* dates, during which the deployment is suspended all day.

You can use this policy to conserve resources by taking a deployment out of Running state when it is not needed, or for preventing a deployment from running during times that it should not be accessed. It can also be useful to keep a deployment in a suspended state during a holiday or a holiday period.

On the Deployment Details page:

- The **Suspension Policy** field's Add button allows you to associate a suspension policy. If a suspension policy is already associated, then you see a dropdown menu with the **Change Policy** and **Remove Policy** options.
- Other new fields on the Deployment Details page include Start Time, Action History (new actions per new policies), and so forth.

To manage Suspension policies, click **Policies** in the Workload Manager UI main menu and then click the **Suspension** link to display the Suspension Policies page.

The Suspension Policies page lists configured suspension policies and lets you perform the tasks described in the following table.

Task	Description
Add a new suspension policy.	Click the New Suspension Policy link. See <i>Configuring a Suspension Policy</i> for details.
Set permissions for a suspension policy.	Choose Share from the dropdown list in the Actions column for the policy. See Permission Control > <i>Policy Permissions</i> for details.
View information about a suspension policy.	You can view the following information for each policy: <ul style="list-style-type: none"> • Description – Optional description that was configured for the policy • Last Updated – Date and time that the policy was last updated • Share with Sub-tenants – Described in the <i>Share with Sub-tenants</i> section above. • Schedule – Runtime schedule and blockout dates that are configured for the policy
View or update configurations for an existing suspension policy.	Choose Edit from the dropdown list in the Actions column for the policy. The Edit page for the policy displays. See <i>Configuring a Suspension Policy</i> for a description of the fields that you can update. If you make updates, click the Save button on the Edit page to save your changes.
Delete an suspension policy.	Choose Delete from the dropdown list in the Actions column for the policy.
Enable or disable a suspension policy.	Turn on or off the Enable switch for the policy. When a policy is disabled, future deployments cannot use it. The policy remains in effect for existing deployments.

The Savings Feature

The suspension policy savings feature displays expected compute time cost savings in the Workload Manager UI as described in the following table.

Savings Data	Displayed in the UI
<p>The Save Up to column to the right of the Schedule column in the Policies > Suspension Policies list page.</p>	<p>This column lists the savings percentage for a suspension policy provide over time.</p> <p>This indeterminate time period calculation is based on the <i>uptime</i> schedule (not <i>blockout dates</i>).</p> <p>Blockout dates are not used in this calculation as the timeframe is indeterminate.</p> <p>In the Create/Edit Suspension Policy form, the percentage savings field in the upper right of the blue header is updated when the user updates any of these four fields: start time, end time, repeat period, repeat every.</p> <p>In the second part of the Deploy form, the percentage savings field in the upper right of the blue header is updated when the user changes the cloud, cloud region, or instance type.</p>
<p>Expected monthly Savings per deployment based on a suspension policy is displayed in the currency for the logged in user</p>	<p>In the second part of the Deploy form, a Savings percentage is displayed next the Cost (per month) field. This is updated whenever the user changes the cloud, cloud region, or instance type.</p> <p>Blockout dates are not used in this calculation as the timeframe is indeterminate.</p> <p>In the Deployment list page, the monthly savings field is displayed in the Approx Savings column located next to the Cloud Cost column as displayed in the following screenshot.</p> <div data-bbox="337 772 1292 1129" data-label="Image"> </div> <p>For <i>terminated</i> deployments, this columns is blank. An Approx Cost column is also added to the right of the Run time column showing the hourly cost and estimated monthly cost based on the <i>uptime</i> schedule.</p> <p>You can also add a policy in the Savings column on the Deployments list page. If a deployment does not already have an associated suspension policy, the Add Policy link is visible in the Savings column. Click the link to open a popup displaying a dropdown list of suspension policies. Next to each policy in the list is an info icon. Hovering over the icon displays an info balloon with the schedule of the policy and expected percent cost savings. Also, hovering over any of the savings entries in the savings column brings up a similar info balloon. You can select one of these suspension policies to add them to this deployment. The following screenshot displays the info icon and associated details for the specific suspension policy.</p> <div data-bbox="337 1381 990 1950" data-label="Image"> </div>

Configuring a Suspension Policy

When you add a suspension policy, you create a new policy based on configuration settings that you make. To add a suspension policy, click **New Suspension Policy** on the Suspension Policy page and then configure the settings that display. Click the **Save** button after you complete the configuration.

The following table describes the settings for configuring a suspension policy.

Setting	Description
Policy Name field	Enter a brief and unique descriptive name for the policy
Policy Description field	Optionally enter a brief description of the policy.
Suspend Time Schedule fields	Enter the start time and end time during which a deployment is to be in Running state on the designated days. You can select half hour time slots from the dropdown.
Repeats field	Choose Daily if the policy should put the deployment in Running state every day during the designated start time through end time period. Choose Weekly and then choose specific days if the policy should put the deployment in Running state only on certain days during this time period.
Blockout Dates fields	Optionally configure blockout dates as follows, which are individual dates or date ranges on which the deployment remains suspended for 24 hours beginning at 12:00 a.m., regardless of the runtime schedule that you configured. To add an individual blockout date, click DATE and then choose the date. To add a date range, click Date Range and then choose the start date and end date of the range. You can enter as many dates and date ranges as needed. To delete a date or date range entry, click the trash can icon in its Actions column.

Suspension Policy Guidelines

- A suspension policy does not prevent you from performing manual suspend and resume operations on a deployment. For example, if a policy sets a deployment to suspend a 9:00 a.m. and resume at 5:00 p.m., you could manually resume the policy at 4:00 p.m.
- By default, a suspension policy uses the time zone of the user that deploys the application using that suspension policy. In this way, no matter which time zone the application is physically running in, the suspension policy will be enforced according to the time zone of the user that launched the application.

A security policy is a policy that can contain ingress and egress rules and can be dynamically attached to a Workload Manager deployment.

Security policies are configured at the tenant level and can be associated with [System Tags](#), the security policy is automatically selected and attached.

The **Security** tab enables you to create a security policy and add a list of firewall rules. The source and destinations of these rules could be IP CIDRs or other security policies.

The Security Policies page lists configured security policies and lets you perform the tasks described in the following table.

Task	Description
Add a new security policy.	Click the New Security Policy link. See Configuring a Security Policy for details.
Set permissions for a security policy.	Choose Share from the dropdown list in the Actions column for the policy. See Permission Control > Policy Permissions for details.
View information about a security policy.	You can view the following information for each policy: <ul style="list-style-type: none"> • Description – Optional description that was configured for the policy • Last Updated – Date and time that the policy was last updated • Share with Sub-tenants – Described in the Share with Sub-tenants section above. • Ports – Lists the open inbound and outbound ports that are configured for each security policy • Enable – Identifies if this new security policy should be enabled for the associated resources.
View or update configurations for an existing security policy.	Choose Edit from the dropdown list in the Actions column for the policy. The Edit page for the policy displays. See Configuring a Security Policy for a description of the fields that you can update. If you make updates, click the Save button on the Edit page to save your changes.

Delete an security policy.	Choose Delete from the dropdown list in the Actions column for the policy. See <i>Deleting a Security Policy</i> for a description of the fields that you can update.
Enable or disable a security policy.	Turn on or off the Enable switch for the policy. When a policy is disabled, future deployments cannot use it. The policy remains in effect for existing deployments.

Configuring a Security Policy

The Workload Manager UI (**Admin > Policies > Security**) enables you to assign firewall rule sets when adding a security policy.

The following screenshot shows the Add a New Security Policy page.

Deleting a Security Policy

You can manually delete the security policy (as long as it does not have any running job associated) from the Security policy page.



Azure Cloud Nuances

Due to the Azure limitation on the number of Security Groups, the Azure security group lifecycle is tied to an [Instance](#) – the security group is created when you create an instance is deleted when you delete the instance.

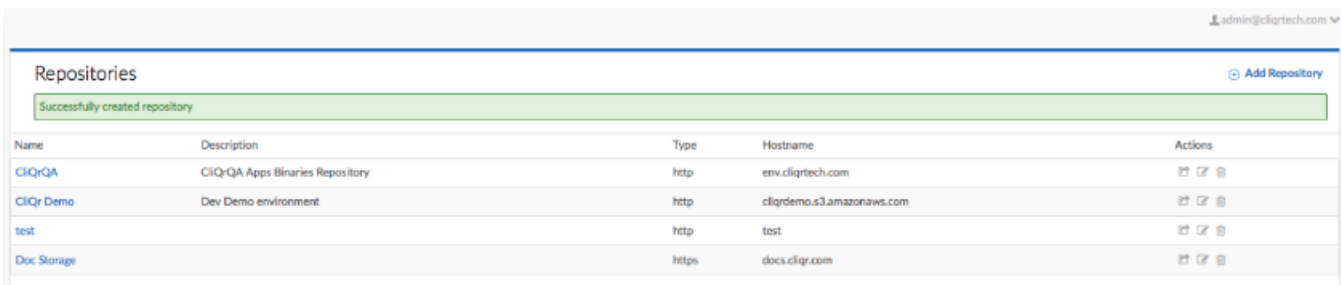
Artifact Repository

Artifact Repository

- [Overview](#)
- [Purpose](#)
- [Using the Artifact Repository in an Application Profile](#)
- [Automatic Endpoint URL Detection](#)
- [Workload Manager Repository Types](#)
- [Create a Repository](#)
- [Share the Repository](#)
- [Back Up and Restore Data for Application Migration](#)

Artifacts (application artifacts) refer to packages, images, binaries, file, scripts, and application data included in physical images that are stored in one or more repositories.

Typically, enterprises maintain application packages, data, and scripts in multiple repositories of their choice. Use the Artifact Repository to attach your own external repository to store and access your files. Workload Manager provides a **Repositories** tab in the Workload Manager UI for this purpose. The following screenshot shows the **Repositories** tab.



Name	Description	Type	Hostname	Actions
CLIQrQA	CLIQrQA Apps Binaries Repository	http	env.cliqrtech.com	[Edit] [Share] [Delete]
CLIQr Demo	Dev Demo environment	http	cliqrdemo.s3.amazonaws.com	[Edit] [Share] [Delete]
test		http	test	[Edit] [Share] [Delete]
Doc Storage		https	docs.cliqr.com	[Edit] [Share] [Delete]

Workload Manager users can define their own repositories and store the required application binaries, scripts, and shared files. At deployment time, you can download (or upload for supported write operations) the required file from this repository.

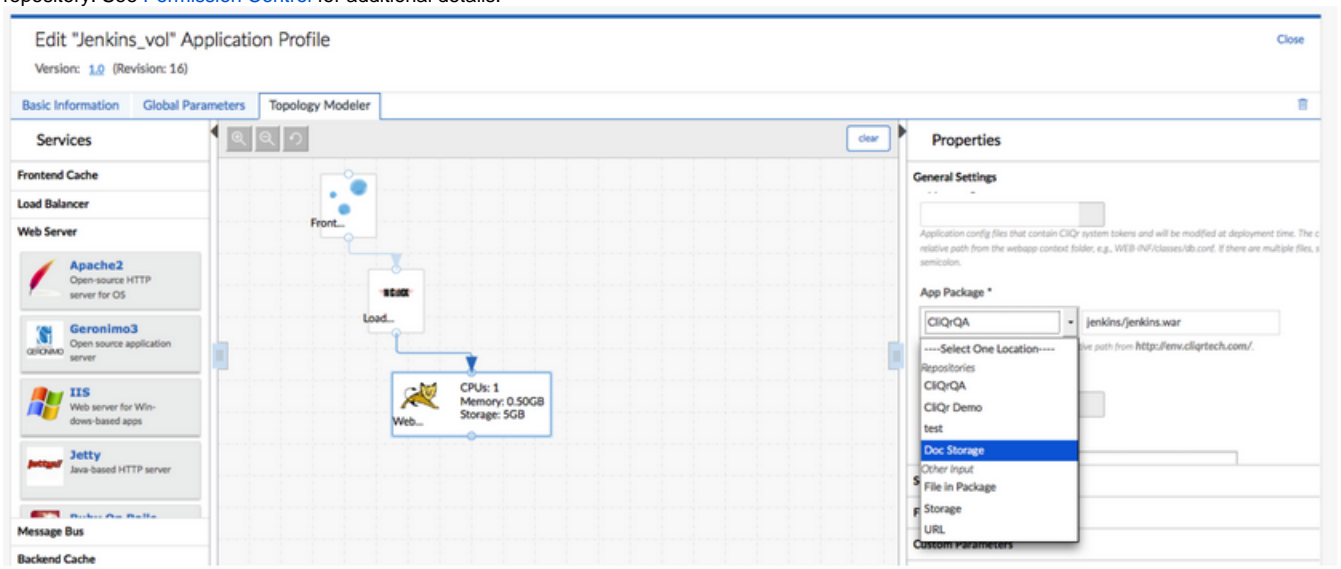
You can provide the required repository credentials to authenticate supported repositories using secure access (Private key, Client Certificate, and Trusted Certificate).

Use the Artifact Repository to:

- Point to application binaries, scripts, and shared files.
- View and maintain the list of external repositories and point to the repository name in application profiles by providing a relative path.

Enterprises may decide to make the artifact repository (or multiple artifact repositories) specific to a user, a tenant, a cloud, or any combination of these resources based on their respective deployment requirement.

Tenant and users can view repositories specific to their tenant (or as permitted by their admin). Admins can enforce access permissions for each repository. See [Permission Control](#) for additional details.



When modeling the application or application profiles, users can select the relevant repository to provide the relative path to the application packages /scripts/file path. The list of added/available repositories is displayed for user selection. When you select a repository, the endpoint URL is automatically appended and you only need to provide the name of the folder where the packages/scripts/files are located.

For each repository type, the Workload Manager platform provides additional fields to identify the configuration specific to each type.

Add a New Repository

Basic Information

Name *


Description


Type *

Username

Password

The following table lists the supported repository types and the inputs required when configuring them in Workload Manager.

Repository Type	Description
HTTP	<p>Use your own HTTP repository for application artifacts.</p> <p>Provide the Hostname, Port, and optionally, Username/Password of your HTTP repository.</p> <p>You can create HTTP repositories using Apache, Nginx or any other web server of your choice. You can then use HTTP repositories to host application artifacts.</p>
HTTPS	<p>Use your own HTTPS repository for application artifacts.</p> <p>Provide the Hostname and Port. You can choose to provide the username/password or SSL credentials that comprise of the Private Key, Client/Trusted certificates of your HTTPS repository.</p>
FTP	<p>Use your own FTP repository for application artifacts.</p> <p>Provide the Hostname, Port, Username/Password of your FTP repository.</p> <div style="border: 1px solid #ffc107; padding: 5px; margin-top: 10px;">  Workload Manager only supports PASSIVE mode when using FTP repositories. </div>

S3	Provide the Region, Access Key, Access Secret for your AWS account, and the name of the S3 Bucket, that you would like to use as your repository.
Artifactory	Provide the Hostname, Port, Username/Password or the SSL Credentials of your Artifactory repository. See repository documentation to set up your Artifact repository.
Artifactory_AUTH	?? Jagadeesh ??
Puppet Server	The central server that manages Puppet Agent(s). Refer to http://docs.puppetlabs.com (install Puppet) for additional context. Only displayed if you use the Puppet service. Provide the Hostname (of the Puppet primary server) and the Certname Suffix (authentication credentials for the Puppet primary server). See Application Using Puppet or Chef for additional details.
Chef Server	The central server that manages the Chef clients. Refer to https://docs.chef.io/ (install server) for additional context. Alternately, you can also use the free Chef server image provided in the AWS Marketplace. Only displayed if you use the Chef service. Provide the Hostname, Chef User Key, Chef Validation key and optionally Trusted Certificate of the Chef Server. See Application Using Puppet or Chef for additional details. You can download the <i>Starter Kit</i> from the Chef console to obtain the validation key that must be specified in CloudCenter UI. <div style="border: 1px solid green; padding: 5px; margin-top: 10px;"> If your Chef Server is configured with a Public DNS, add the Public DNS in the Hostname field and copy the public DNS to the Trusted Certificate field when creating or editing the repository.</div>
Other Input	Provide the entire URL for the repository – this is the only option that is directly configurable for Workload Manager resources.

To define a shared repository, you must first create the repository and then share the repository with applicable users or groups.

To create a repository, follow this procedure:

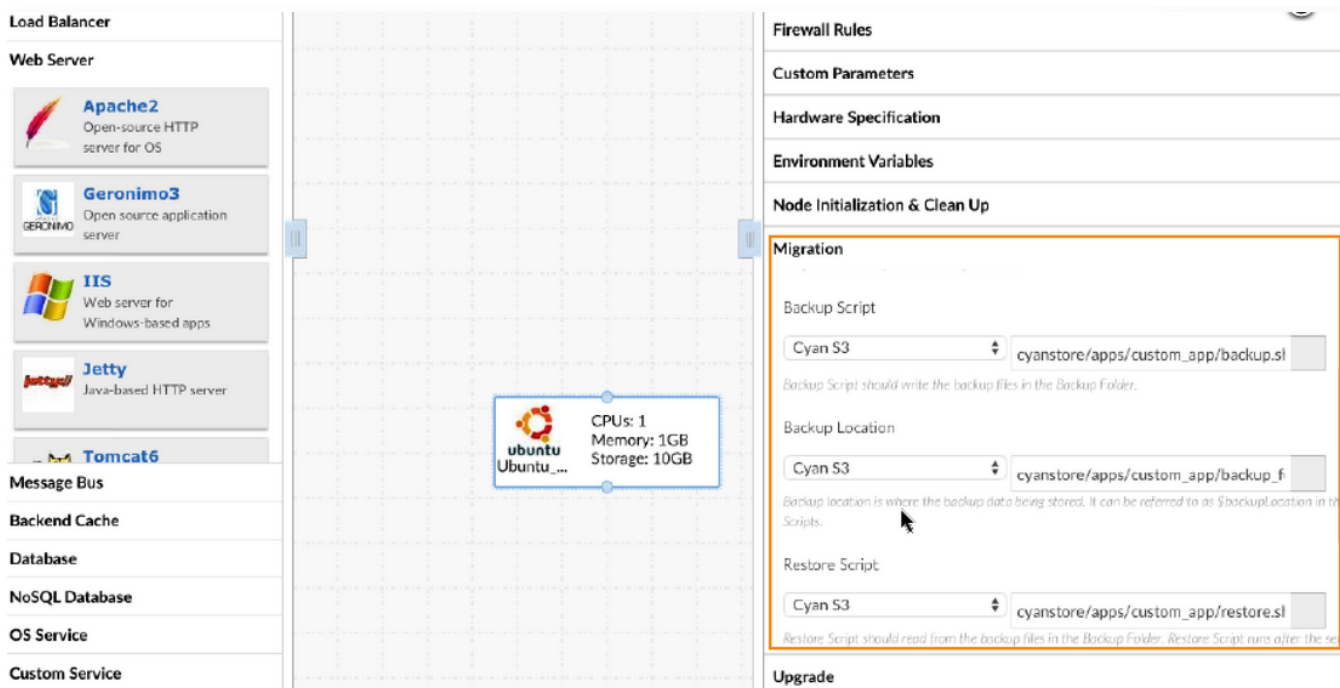
1. Access the **Repositories** tab in the Workload Manager UI.
2. Review the repositories (if any):
 - If the required repository is listed in this page, click the **Edit** icon. You can only edit a repository if you are the owner.
 - If the required repository is not listed, follow the rest of the procedure to add a new repository.
3. Click the **Add Repositories** button to add a new repository.
4. Provide the Basic Information required to model this repository.
5. Provide the authentication credentials, depending on your authentication scheme, provide username/password or certificate details for SSL.
6. Click **Save**. The newly added repository is saved and visible in the Your Repositories page.

Each tenant and users within a tenant can only view shared repositories specific to their tenant (or as permitted by their admin).

To share a repository, follow this procedure:

1. Access the the Workload Manager UI and click **Repositories**.
2. For the required repository, click the **Share Repository** icon.
3. Associate the required Users, Groups, and Tenants who are permitted to access this repository and identify the [Permission Control](#) for each association. You can also choose to share this repository with all users and/or tenants.
4. Click **Save**. The newly added repository is saved and visible in the Repositories page.

Some repositories also support Write operations can be used to backup (or restore) data to/from the repository as identified in the table above (see the Write Operations column).



To use this capability, you must provide the backup location (repository endpoint) for the backup/restore operation. Once you provide the relative path for the **Backup Script**, **Backup Location**, and/or the **Restore Script**, in the Migration Properties tab, these will be used to perform application migration between clouds.

Operation	Path Description
Backup Location	Application data is backed up to the repository specified in the Backup Location .
Backup Script	Executed as root user. Backup file syntax: backupFile <file name and backup path on Backup location> <file with path to be backed up> Use the backup shell script to call the util function. Include the following line to backup a file: backupFile \$path_on_repo \$localFilePath , replace the \$localFilePath with the path of the file that must be backed up.
Restore Script	Executed as root user. Restore file syntax: restoreFile <file name> By default, the file is restored to /opt/remoteFiles/ . Use the restore shell script to call the util function. Include the following line to restore a file: restoreFile \$path_on_repo

You can repeat calls to backupFile and restoreFile functions for each file that must be backed up or restored.

Actions Library

Actions Library

- [Actions Library Overview](#)
- [Actions Library Page](#)
- [Create a Custom Action](#)
- [Manage Actions](#)
- [Create Snapshot](#)

Actions Library Overview

Actions Library Overview

- [Overview](#)
- [On-Demand Actions](#)
- [Lifecycle Actions](#)
- [Platform Actions](#)

The Actions Library is the Workload Manager data structure for storing on-demand actions and lifecycle actions. Users with the WM_ADMIN or WM_DEV_OPS role can view, create and share custom actions using the [Actions Library page](#) which is accessible from the Actions Library tab of the main menu. The Actions Library includes a set of OOB on-demand actions which are also referred to as *platform actions*.

Characteristics of on-demand actions, lifecycle actions, and the out-of-box platform actions are summarized below.

On-Demand actions have the following characteristics:

- Are manually invoked by the user from various screens in Workload Manager UI, including the Deployments List page, Deployment Details page, Virtual Machines List page, Virtual Machine Details page, and Application Profiles page.
- Include Workload Manager OOB [platform actions](#) and custom on-demand actions.
- Platform actions:
 - Are visible to all users and subtenants by default but can be restricted to certain users and subtenants by a root tenant administrator.
- Custom on-demand actions:
 - May be one of these action types:
 - Invoke Web service
 - Command or script
 - Puppet
 - UCS Director workflow
 - Chef
 - Ansible
 - The Invoke Web service action type may be executed against a VM, a deployment, or an application profile. All other action types are only executed against a VM.
 - Action types executed against a VM can be defined to be executed on the VM if the management agent is installed on the VM (agent actions), or externally (external actions).
 - May include custom fields that are entered by the user at the time of action execution.

Lifecycle actions:

- Are automatically invoked by Workload Manager when a deployment initiated through Workload Manager passes through a specific phase of its lifecycle.
- Are always custom actions defined by the user.
- Are always of type Command or script.
- Are always executed against VMs deployed through Workload Manager.
- Can be defined to be executed on the VM if the management agent is installed on the VM (agent actions), or externally (external actions).
- Can be made available to only certain VMs based on the VM's associated service or application profile, and, for external actions, the associated cloud region.
- Agent actions:
 - Execute in the target VM.
 - May be referenced in:
 - Node Initialization and Cleanup settings of a tier in the Application Profile Topology Modeler tab.
 - Agent Lifecycle Actions section of the Service Definition form.
- External actions:
 - Execute in a container in Cloud Remote if Cloud Remote is deployed in the region where the VM is deployed.
 - Execute in a container in the management cluster if Cloud Remote is not deployed in the region where the VM is deployed.
 - May be referenced in:
 - External Initialization settings of a tier in the Application Profile Topology Modeler tab.
 - External Lifecycle Actions section of the Service Definition form.
 - External Lifecycle Actions section of the Regions tab for a cloud region.
- May include custom fields which are entered by the user when the action is referenced in an application profile, a service definition, or a regions tab.

Platform actions refer to OOB on-demand actions. These actions are always executed against a particular VM. Like custom on-demand actions executed against a VM, the platform actions are visible as entries in the VM action dropdown menu in the Virtual Machines List page and the VMs tab of the Tiers tab of the Deployment Details page. They are also visible as action buttons on the right panel of the VM Details page.

Platform actions are visible to all users and subtenants by default, but visibility and use of platform actions can be restricted to only certain users and subtenants by a root tenant administrator. The following table summarizes properties of the platform actions.



All platform actions require the VM to be a managed VM except:

- Terminate, which can be applied to both managed and unmanaged VMs, and
- Import, which can only be applied to unmanaged VMs.

Action Name	Description	Required VM State ¹	Supported Clouds
Start	Powers on the VM	Stopped	All
Stop	Powers off the VM	Running	
Reboot	Reboots the VM. While VMware and OpenStack support both soft and hard reboot, the CloudCenter Suite performs a soft reboot on OpenStack environments and a hard reboot on VMware environments. Here is how each cloud currently handles a reboot operation: <ul style="list-style-type: none"> • OpenStack has both soft and hard reboot options. • VMware has a separate option for graceful (GuestOS) reboot/shutdown and supports it only when VMware tools is installed. • Google cloud issues a warning in their UI that they implement a hard reboot. • AWS performs a hard reboot if the instance does not cleanly shut down within four minutes. • Azure RM attempts a graceful shutdown by default and if it does not complete in 5 minutes they perform a hard shutdown. 	Running	
Terminate	Stops (if running) and removes the VM from the cloud	Running or stopped	
Import	Moves a VM from the Unmanaged category into the Managed category. See Virtual Machine Management for additional context.	Running or stopped	
Install Agent	Installs the management agent on an Imported VM from the Workload Manager UI. See Virtual Machine Management > Install Agent for additional details. If any version of the agent is already installed on a VM, then this action will not be available for this VM. <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">i The Install Agent action requires the user to specify the OS type and login credentials during submission.</div>	Running	
Upgrade Agent	Upgrades the agent on a Managed VM to the latest released Workload Manager agent version from the Workload Manager UI. See Virtual Machine Management > Upgrade Agent for additional details. <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">i You can only upgrade a VM running CloudCenter Legacy 4.7.3 or later versions. See Virtual Machine Management > Install Agent for additional details.</div> If the latest version of the agent is already installed on a VM, then this action will not be available for this VM.	Running	
Create and Attach Volume	Creates and attaches a volume of a storage type defined in Workload Manager. The user specifies the storage type and size via a dialog box when the action is initiated. Cisco does not support the attachment of volumes created outside of Workload Manager. <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">i Requires at least one storage type to be defined for the region where the VM is deployed. Upon execution of this action you are prompted for volume type and size.</div>	Running or stopped	
Detach Volume	Detaches volumes from the existing VM.	Running or stopped	
Sync VM information	Updates the current VM information from Workload Manager so the latest VM metadata information is visible.	Running, stopped, or started	
Resize Instance Type	Lets you select a different instance type from among the instance types defined for	Stopped	
Create Snapshot	Creates the image snapshot for the given VM.	Running or stopped	vCenter

¹ For more information on VM states, see [Deployment, VM, and Container States](#).

Actions Library Page

Actions Library Page

- [Overview](#)
- [Page Contents](#)

The Actions Library page is where you can create and edit custom actions and view all actions created by you or shared with you. You can access the Actions Library from the Actions Library tab of the main menu. This tab is only visible to users with the WM_ADMIN or WM_DEV_OPS role. If you are a tenant administrator for the root tenant, you can also control which users can see and use the individual platform actions, or completely disable individual platform actions.

A sample Actions Library page as viewable by a root tenant administrator is shown below.

NAME	TYPE	LAST UPDATED	DESCRIPTION	WHERE USED	ENABLE	ACTIONS
jb1 ON DEMAND	Command or Script	June 4, 2019 6:32 PM		Virtual Machines	ON	
Chef ON DEMAND	Chef	May 26, 2019 6:16 PM		Virtual Machines	ON	
Connect ON DEMAND	SSH VNC		Connect	Virtual Machines	ON	
Resend Approval Request ON DEMAND	RESEND APPROVAL REQUEST		resend approval request	Virtual Machines	ON	
Cancel Approval Request ON DEMAND	CANCEL APPROVAL REQUEST		cancel approval request	Virtual Machines	ON	
Create Snapshot ON DEMAND	INSTANCE CREATE SNAPSHOT		create snapshot	Virtual Machines	ON	

The Actions Library page has a header and a main table. You can perform functions on the Actions Library from both the header and the table as described below.

From the **page header** you can:

- Click on the **Search** icon (upper left) and enter text to limit the list of actions displayed.
- Click on the **Export** link (upper right) to export all custom actions into a single JSON file. This is useful if you want to later import these actions into another instance of Workload Manager.
- Click on the **Import** link (upper right) to import custom actions from a JSON file. When you do this, a dialog box is displayed where you must map the repository locations included in the imported file to your existing repository locations.
- Click on the **New Action** button (upper right) to create a new custom on-demand action or lifecycle action. Details on creating a custom action are explained in [Create a Custom Action](#).

The main section of the Actions Library page is the **table** containing the list of actions. The meaning of the table columns is summarized in the table below. You can sort the table by clicking the header of any of these columns: Name, Type, Last Updated, and Enable.

Column	Notes
Name	Shows the name of the action on the first line, and the category of action (on-demand or lifecycle) on the second line. If the action is a custom action, the name is shown in blue. If you have modify or manage permission for the action, you can click anywhere on the row to open the action editor form for that action. See Manage Actions for more information on editing actions. All OOB on-demand actions have a closed padlock icon next to the name which means they cannot be edited and cannot be deleted.
Type	Each OOB on-demand action has its own unique type. Lifecycle actions are all of type Command or script. Custom on-demand actions may be of type Invoke Web service, Command or script, Puppet, UCS Director workflow, Chef, or Ansible.
Last Updated	Workload Manager tracks when a user updates an action through the editor. This provides a convenient way to see which actions were recently updated.
Description	Free-form text field entered by the creator of the custom action.

Where Used	<p>This field has two different meanings depending on whether the action is on-demand or lifecycle.</p> <p>On-demand actions: <i>Where Used</i> means the target of the action: the object that the action will query, modify or control. Possible values for <i>Where Used</i> for on-demand actions vary based on the action type, as summarized in the table below.</p> <table border="1" data-bbox="256 254 927 443"> <thead> <tr> <th colspan="2">Possible values for <i>Where Used</i> for on-demand actions</th> </tr> <tr> <th>Action Type</th> <th>Possible Values</th> </tr> </thead> <tbody> <tr> <td>Invoke web service</td> <td>Application Profile, Virtual Machines, Deployments</td> </tr> <tr> <td>All other action types</td> <td>Virtual Machines</td> </tr> </tbody> </table> <p>Lifecycle actions: <i>Where Used</i> means the resources where the lifecycle action may be specified. Possible values for <i>Where Used</i> for lifecycle actions vary based on whether the command or script of the lifecycle action is executed on the VM or externally, as summarized in the table below.</p> <table border="1" data-bbox="256 554 868 743"> <thead> <tr> <th colspan="2">Possible values for <i>Where Used</i> for lifecycle actions</th> </tr> <tr> <th>Execution Location</th> <th>Possible Values</th> </tr> </thead> <tbody> <tr> <td>On-VM</td> <td>Application Profile, Service</td> </tr> <tr> <td>External</td> <td>Application Profile, Service, Cloud Region</td> </tr> </tbody> </table>	Possible values for <i>Where Used</i> for on-demand actions		Action Type	Possible Values	Invoke web service	Application Profile, Virtual Machines, Deployments	All other action types	Virtual Machines	Possible values for <i>Where Used</i> for lifecycle actions		Execution Location	Possible Values	On-VM	Application Profile, Service	External	Application Profile, Service, Cloud Region
Possible values for <i>Where Used</i> for on-demand actions																	
Action Type	Possible Values																
Invoke web service	Application Profile, Virtual Machines, Deployments																
All other action types	Virtual Machines																
Possible values for <i>Where Used</i> for lifecycle actions																	
Execution Location	Possible Values																
On-VM	Application Profile, Service																
External	Application Profile, Service, Cloud Region																
Enable	<p>For custom actions: If you have modify or manage permission for the action, this field is a toggle you can switch ON or OFF. If you only have view permission for the action, this column contains a static text field (ON or OFF). When you create a new custom action, it is automatically enabled.</p> <p>For OOB on-demand actions: If you are an administrator for the root tenant, the Enable column shows toggles for the OOB on-demand actions which you can turn ON or OFF. For all other users, the Enable column for the OOB on-demand actions just shows the ON or OFF text. All OOB on-demand actions are enabled by default. When you create a new custom action, it is automatically enabled.</p>																
Actions	<p>If you have manage permission for a custom action, hovering over the Action column for the action will display a triangular dropdown menu icon. Clicking on this icon shows two menu choices: Share and Delete. If you are a root tenant administrator, and hover over this column for for an OOB on-demand action, clicking on dropdown icon initially shows one menu choice: Disable Default Share. Once you click this choice, the menu choices change to two choices: Share, Enable Default Share. See Manage Actions for more details on sharing and deleting actions.</p>																

Create a Custom Action

Create a Custom Action

- [Overview](#)
- [General Settings Section](#)
- [Resource Mapping Section](#)
 - [Resource Mapping for Actions Other Than Invoke a Web Service](#)
 - [Resource Mapping for the Invoke a Web Service Action](#)
- [Action Definition Section](#)
 - [Custom Fields](#)

To define either an On-Demand Action or a Lifecycle Action, click the **New Action** link in the [Actions Library page](#). This brings you to the New Action form. This form has three sections you must complete before you can save the action:

- **General Settings Section** – Defines action category, action type, where the action executes, and timeout value.
- **Resource Mapping Section** – Defines which resources this action may be executed against.
- **Action Definition Section** – Defines the required parameters corresponding to the selected action type, and lets you specify input parameters ([custom fields](#)) for the action.



Complete each section of the New Action form as described below.

When you first create a new action, the General Settings section will be displayed at the beginning of the New Action form as shown below.

Follow this procedure to complete the fields in this section.

1. Choose the action category: On-Demand or Lifecycle.
2. Choose the action type. The table below summarizes the possible action types for on-demand actions. For lifecycle actions, the only available action type is Command or Script.

Type	Description
Invoke a Web Service	Executes the designated web service request against a VM, deployment, or application profile.

Command or Script (Default)	<p>Executes the designated command or script against a VM.</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"> <p> If you provide scripts that need to be executed as part of Action Library, be aware that a script is considered to be successful based on OS exit codes for script execution. This behavior is consistent with scripts that are run directly on a Windows or Linux server using their respective shells. Once a script is executed on Linux or Windows, the output shown on running <code>echo \$?</code>, is the status code.</p> <p>For example:</p> <ul style="list-style-type: none"> On Linux, script execution will be marked as Failed only if the script's last line failed to execute. In all other cases it will be marked as successful. On Windows, scripts will be marked as Failed only if an explicit Exception is thrown. <p>For user-provided scripts, the onus is on the user to ensure that the script exits successfully.</p> </div> <p>This type additionally allows you to specify the following options:</p> <ul style="list-style-type: none"> Where to run this type of action: <ul style="list-style-type: none"> On the VM – Requires that management agent is installed. Externally – In a script execution container in the management cluster, if Cloud Remote is not used to connect with the region where the VM is deployed, or in a script execution container in the Cloud Remote cluster, if Cloud Remote is used for the region. Reboot the VM after action execution – If the action creator determines that a reboot is required after the action execution, then end users using this action do not need to manually perform this extra step. This switch is only available for On-Demand Actions. Sync VM information after action execution – This causes Workload Manager to update its VM information with the latest VM metadata provided by the cloud region. You can perform this operation in bulk by multi-selecting several VMs, or just for one VM. This switch is only available for On-Demand Actions. <p>The possible execution modes for this option are based on your <i>Execute from Bundle</i> setting that is explained in the <i>Action Definition Section</i> later in this page:</p> <ul style="list-style-type: none"> Download From Bundle is True <ol style="list-style-type: none"> The content in Script From Bundle is presumed to be a script inside the bundle specified. The bundle is presumed to be an archive (zip/tar/tgz/tar.gz) that is extracted and the file is searched and executed along with any other command line parameters. Download From Bundle is False <ol style="list-style-type: none"> The content in the Executable Command field is presumed to be a series of commands separated by semicolons. A special case is that the first command in this series of commands can be a Script URL. Thus the first command, if a script URL, is processed by first downloading the script and then it is executed along with the rest of the following commands.
Puppet	<p>Causes a Puppet role to be enforced against a VM.</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"> <p> To run Puppet actions on Workload Manager instance, you must first disable the <code>requiretty</code> setting for that instance in the <code>/etc/sudoers</code> file.</p> </div>
Invoke Cisco UCS Director Workflow	<p>Causes a workflow defined in a UCS Director server to be executed against a VM. Input parameters for the workflow defined in UCS Director are displayed to the Workload Manager user in a data entry dialog box when</p>
Chef	<p>Causes a Chef run list to be run against a VM.</p>
Ansible	<p>Causes an Ansible playbook to be run against a VM.</p>

3. Specify the Action Name, Description, and Action Timeout.

The function of the Resource Mapping section is to qualify where an action may be made available. The appearance of this section changes based on what type of action is specified in the General Settings section. If the action specified is Invoke a web service, the user must first specify whether the action is directed against a deployment, VM, or application profile. If the user selects any other action type in the General Settings section, the action is always executed against a VM. If the action is executed against a deployment or a VM, the *Where To Apply* options must also be selected.

The following table identifies the available resource types for each action type and their associated *Where To Apply* options.

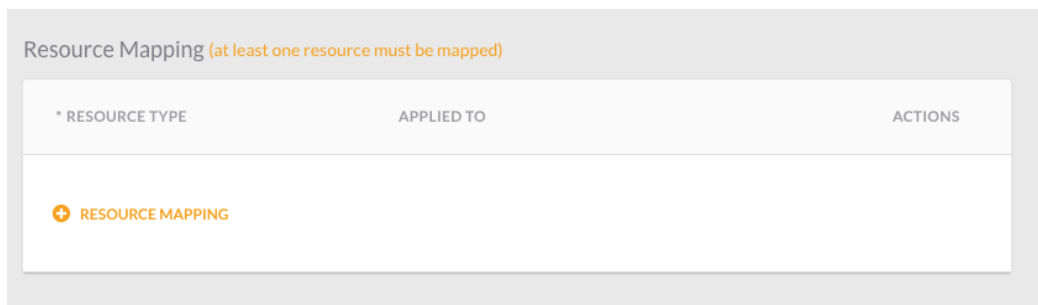
Action Type	Resource Type	Where to Apply Options
Invoke a Web Service	Deployments	<ul style="list-style-type: none"> Deployment Environment: Identifies if this action is applied to all deployment environments or specific deployment environments. Default = All deployment environments. Service: Identifies if this action is applied to all services or specific services. Default = All services.

	Workload Manager Deployed VMs	<ul style="list-style-type: none"> • Application Profile: Identifies if this action is applied to all application profiles or specific profiles. Default = All application profiles. • Service: Identifies if this action is applied to all services or specific services. Default = All services. • Cloud Region: Identifies if this action is applied to all cloud regions or specific regions. Default = All cloud regions. • Cloud Account: Identifies if this action is applied to all cloud accounts or specific accounts. Default = All clouds accounts.
	Imported VMs	<ul style="list-style-type: none"> • Cloud Region: Identifies if this action is applied to all cloud regions or specific regions. Default = All cloud regions. • Cloud Account: Identifies if this action is applied to all cloud accounts or specific accounts. Default = All clouds accounts. • OS Types: Identifies if this action is applied to all OSs or specific OSs. Default = All
	Application Profiles	<ul style="list-style-type: none"> • None: When this resource type is chosen the action is applied to all application profiles.
<ul style="list-style-type: none"> • Command or Script (Default) • Invoke Cisco UCS Director Workflow • Chef • Puppet • Ansible 	Workload Manager Deployed VMs	<ul style="list-style-type: none"> • Application Profile: Identifies if this action is applied to all application profiles or specific profiles. Default = All application profiles. • Service: Identifies if this action is applied to all services or specific services. Default = All services. • Cloud Region: Identifies if this action is applied to all cloud regions or specific regions. Default = All cloud regions. • Cloud Account: Identifies if this action is applied to all cloud accounts or specific accounts. Default = All clouds accounts.
	Imported VMs	<ul style="list-style-type: none"> • Cloud Region: Identifies if this action is applied to all cloud regions or specific regions. Default = All cloud regions. • Cloud Account: Identifies if this action is applied to all cloud accounts or specific accounts. Default = All clouds accounts. • OS Types: Identifies if this action is applied to all OSs or specific OSs. Default = All

To complete the Resource Mapping section, follow the instructions below based on the action type you selected in the General Settings section.

Resource Mapping for Actions Other Than Invoke a Web Service

If you select any action type other than Invoke a web service, the action is executed against a VM and the Resource Mapping section appears as shown below.



Clicking on the Add Resource Mapping link displays the following dialog box:

Resource Mapping
✕

RESOURCE TYPE *

Workload Manager Deployed VMs

Imported VMs (with Agent Installed)

Where To Apply

APPLICATION PROFILE

ALL
▼

CLOUD REGION

ALL
▼

CLOUD ACCOUNT

ALL
▼

SERVICE

ALL
▼

DONE

Selecting the resource type will cause the list of *Where To Apply* parameters to change as summarized in the table above.

Select the resource type (Workload Manager Deployed VMs or Imported VMs), select the *Where To Apply* options from each of the dropdown menus, and then click Done. This will cause a resource mapping to be displayed in the Resource Mapping section as shown below.

Resource Mapping
✕

* RESOURCE TYPE	APPLIED TO	ACTIONS
Workload Manager Deployed VMs	Application Profile: ALL Service: ALL Cloud Region: ALL Cloud Account: ALL	

+

RESOURCE MAPPING

Repeat this process as many times as needed to add all of your resource mappings. You can delete a resource mapping by clicking the corresponding trash can icon in the Actions column.

Resource Mapping for the Invoke a Web Service Action

If you select the Invoke a web service action type in the General Settings section, the action may be executed against a deployment, a VM, or an application profile; therefore the Resource Mapping section changes to include a dropdown menu labeled *Select Resource To Map* as shown below.

Resource Mapping (at least one resource must be mapped)

* SELECT RESOURCE TO MAP

Deployments

* RESOURCE TYPE	APPLIED TO	ACTIONS
<p>+ RESOURCE MAPPING</p>		

The default value for *Select Resource To Map* is **Deployments**. If you now click the **Add Resource Mapping** button, a resource mapping dialog box for Deployments is displayed as shown below.

Resource Mapping

Where To Apply

DEPLOYMENT ENVIRONMENT

ALL

SERVICE

ALL

DONE

Select the *Where To Apply* options from each of the dropdown menus, and then click **Done**. This will cause a new resource mapping to be displayed in the Resource Mapping section. Add and delete mappings as needed.

If you click on the *Select Resource To Map* dropdown and change it to **Virtual Machines**, clicking the Add Resource Mapping link caused the resource mapping dialog box for VMs to appear as described in the [Resource Mapping for Actions Other Than Invoke a Web Service](#) section above.

If you click on the *Select Resource To Map* dropdown and change it to **Application Profiles**, the Add Resource Mapping link disappears and is replaced with the message shown below.

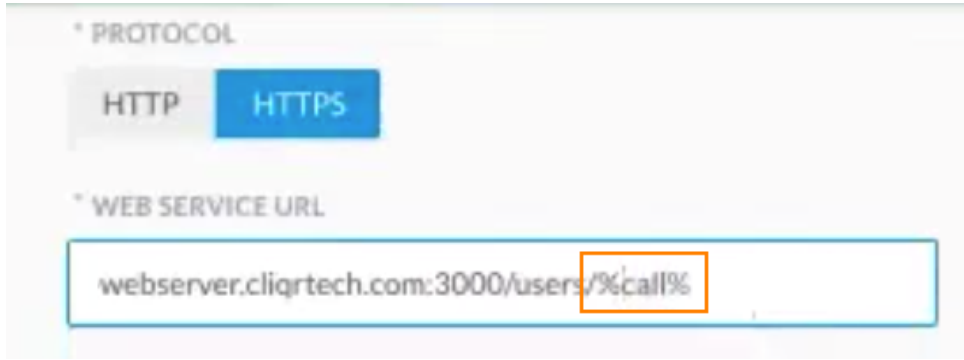
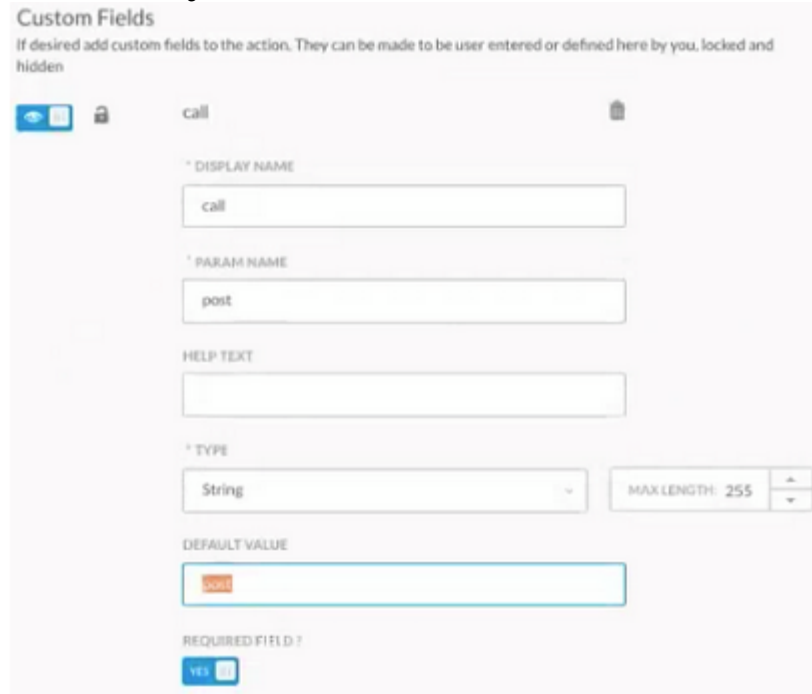
Resource Mapping

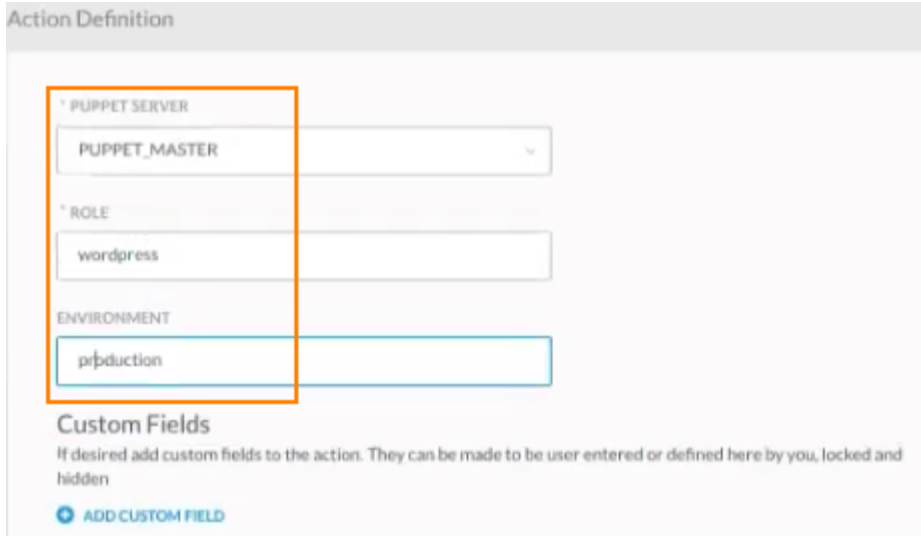
* SELECT RESOURCE TO MAP

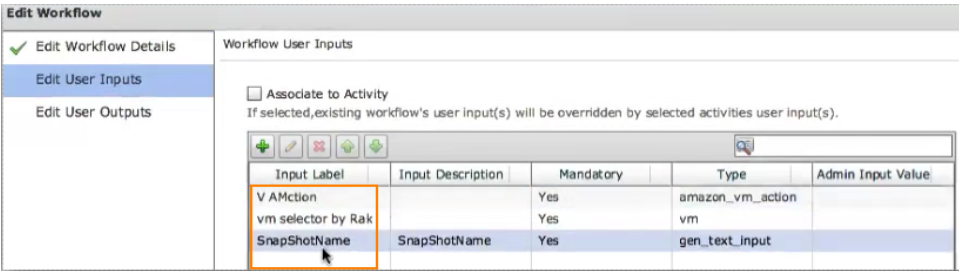

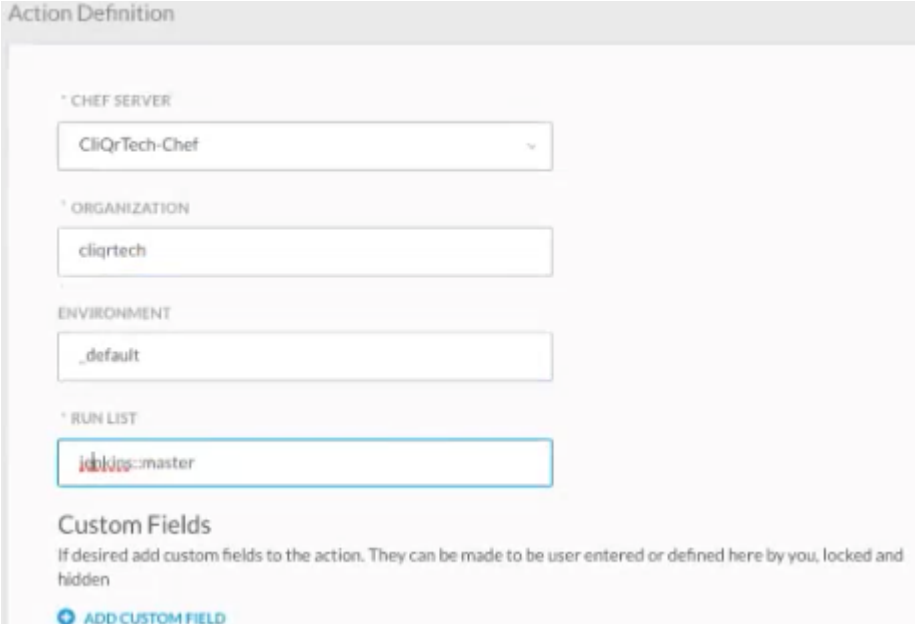
Application Profiles

This action will be applied to all Application Profiles

The input fields for the Action Definition section change. The following table identifies the available **Type** along with a brief description and identifies the permitted resources for **Resource Mapping**.

Type	Fields	Notes
Invoke a Web Service	Protocol	The transfer protocol to be used by this action when invoking the service – select HTTP or HTTPS.
	Web Service URL	<p>The URL used for the web service using the following format: <host>:<port>/<resource>/<parameter> For example: <i>webservice.cliqrtech.com:3000/users/post</i></p> <p>Use Case: You can also introduce a custom parameter in the Web Service URL field and define that parameter in the Custom Fields section (described below). For example, if you introduce a custom parameter called <i>call</i>/the URL as shown in the following screenshot.</p> 
		<p>Define the <i>call</i> parameter in the Custom Fields section to ensure that this parameter is replaced by the value defined in that section. The following screenshot shows the Custom Fields section.</p> 
	Username and Password:	Credentials required to issue the web service call
HTTP Request Type	<ul style="list-style-type: none"> • PUT • GET • POST • DELETE 	
Content Type	If the body content should be sent using the JSON or XML format	

	Body	The request body contents to be used when issuing the API call for this action.
Command or Script (Default)	Execute From Bundle	<p>If you choose to configure this setting, provide the following details:</p> <ul style="list-style-type: none"> • Location: Select from a list of previously-configured repositories as described in the Artifact Repository section. If repositories have not been configured, you must use the default URL option and provide the entire URL. • Relative Path: Specify the path to the folder where the script bundle resides. Workload Manager appends this path to the hostname defined in the repository. All compressed file formats (.tar, .zip, etc.) are acceptable in this field. • Script from Bundle: Provide the name of the script that this action should use.
	Executable Command	If you choose to configure this option, provide the command that should be executed as part of the custom action.
Puppet	Puppet Server	<p>A Puppet server that is maintained by the user can be added to the CloudCenter repository list. (See Artifact Repository to configure a repository). A Puppet server configured in Workload Manager as a shared repository is referred to as a <i>Puppet repository</i>. If multiple Puppet repositories exist, they are listed in the dropdown menu and you must pick one of the configured Puppet repositories in this field, as shown in the following screenshot.</p> 
	Role	A role that must exist in the Puppet repository configured above. For example <i>wordpress</i> .
	Environment	The Puppet environment where this action should be executed. For example, <i>production</i> .
Invoke Cisco UCS Director Workflow	API Key	The API key for the UCS Director.
	UCSD Server	The IP address of the UCS Director.

<p>Workflow Name</p>	<p>The name of the workflow created in the UCS Director.</p> <p>All of the workflow user input fields configured in the workflow in the UCS Director, as shown in the following screenshot...</p>  <p>... are displayed in a data input dialog box in Workload Manager when the action is invoked, as shown in the following screenshot.</p> 
<p>Chef Server</p>	<p>A Chef server that is maintained by the user can be added to the CloudCenter repository list. (See Artifact Repository to configure a repository). A Chef server configured in Workload Manager as a shared repository is referred to as a <i>Chef repository</i>. If multiple Chef repositories exist, they are listed in the dropdown menu and you must pick one of the configured Chef repositories in this field. The following screenshot shows an example.</p> 
<p>Organization</p>	<p>The company or department details. In this example, <i>cliqrtech</i></p>

	Environment	The Chef environment where this action should be executed. In this example, <i>default</i> .
	Run List	The combination of the cookbook name and recipe name using the <i>cookbook::recipe</i> format – in this example, <i>jenkins::master</i>
Ansible	Repository	An Ansible server that is maintained by the user can be added to the CloudCenter repository list. (See Artifact Repository to configure a repository). An Ansible server configured in Workload Manager as a shared repository is referred to as an <i>Ansible repository</i> . If multiple Ansible repositories exist, they are listed in the dropdown menu and you must pick one of the configured Ansible repositories in this field.
	Playbook	The Ansible playbook path is <i>jenkins/install.yml</i> .

Once you complete the required fields in all three section, you can preview the action to see how the action is presented to the end user by clicking the **Preview** button in the lower right, or you can save the action by clicking the **Done** button in the lower right.

Custom Fields

At bottom of the Action Definition Section you can add optional input fields and make them available for user input.

To add a custom field, click the Add Custom Field link at the bottom of the Action Definition section. This causes the section to expand showing input parameters for defining the custom field, as shown in the screenshot below.

Custom Fields

If desired add custom fields to the action. They can be made to be user entered or defined here by you, locked and hidden

* DISPLAY NAME

* PARAMETER NAME

HELP TEXT

* TYPE

String

MAX LENGTH

255

DEFAULT VALUE

REQUIRED FIELD ?

YES

+ ADD CUSTOM FIELD

Set the toggles and fields for the custom field as explained in the table below.

Field / Toggle	Description
Visibility toggle	Default is visible. Set to invisible if you do not want to display this custom field in a dialog box during execution of an on-demand action, or during selection in an application profile or service definition or region tab of a lifecycle action. When set to invisible, the Lock toggle and Required Field toggle (see below) are hidden.
Lock toggle	Default is unlocked, meaning the user can edit the value of this custom field in a dialog box during execution of an on-demand action, or during selection in an application profile or service definition or region tab of a lifecycle action. Click on the icon to lock it. This prevents the user from changing the value of the field. If this toggle is set to locked, Required Field toggle (see below) is hidden.
Display Name	Enter the name that you wish to assign for this custom field. If made visible, this label is displayed to end users when they invoke this action.

Param Name	The parameter name that is available to the script execution as environment variables. If these parameters are provided as \$parameterName or as %paramName% in the script parameters, Workload Manager replaces this field with the passed value in the script.
Help Text	Provide additional tips that you wish to provide for the end user to complete in this field. Try using a single, short sentence so your end users are not overwhelmed with too much text.
Type	See Using Parameters > Parameter Type for details
Default Value	Assign the default value to be used by the custom field – should the end user not specify any values.
Required Field toggle	Set to YES if the end user must enter information in this field in order for the action to complete.

Repeat adding custom fields as necessary until done. You can delete a custom field by clicking on the the corresponding trash icon.

Manage Actions

Manage Actions

- [Overview](#)
- [Edit an Action](#)
- [Share an Action](#)
 - [Custom Actions](#)
 - [OOB On-Demand Actions](#)
- [Delete an Action](#)
- [Disable an OOB On-Demand Action](#)

The following management functions can be applied to actions created by you or shared with you:

- **Edit** - If you have manage or modify permission for a custom action, you can modify it via the Actions Edit form.
- **Share** - If you have manage permission for a custom action, you may share it with other users, groups, and subtenants. If you are a tenant admin for the root tenant, you also have the ability to control which users, groups and subtenants can see and use the OOB on-demand actions.
- **Delete** - If you have manage permission for a custom action, you may delete it. However, if the action is a lifecycle action and is referred to in an application profile, service or cloud region setting, it cannot be deleted until all of the references to it are deleted.

To edit an action, from the list of actions on the Actions Library page, click on the row containing the action you want to edit. This displays the Edit Action form. The Edit Action form is similar to the New Action form as explained in [Create a Custom Action](#) with the following restrictions:

- The action category (on-demand or lifecycle) cannot be changed.
- The action type cannot be changed.
- For the actions using the Invoke web service action type, the *Select Resource to Map* field can only be changed after all resource mappings associated with that resource have been deleted.

When done editing an action click the **Done** button to save it.

Custom Actions

To share a custom action that you have to manage permissions for, from the list of actions on the Actions Library page, for the action you want to share, hover over the Action column until you see the dropdown menu icon appear. Once it appears, click on the dropdown icon to display the dropdown menu. Click the **Share** command from the dropdown menu. This causes the Share dialog box for custom actions to be displayed. Grant view, modify or manage access to all users in your tenant, or to any user, group or subtenant as appropriate. Click **Save** when you are done.

OOB On-Demand Actions

By default, all OOB on-demand actions are visible to all users in all tenants. However, as a tenant administrator of the root tenant, you can restrict which users, groups and tenants can see and use the OOB platform actions. To do this, from the list of actions on the Actions Library page, for the action you want to share, hover over the Action column until you see the dropdown menu icon appear. Once it appears, click on the dropdown icon to display the single command: **Disable Default Share**, then click on that command. Once the default sharing is disabled, the action is only visible to the user, groups, or tenants that you explicitly share it with. When default sharing is turned off, clicking on the Action dropdown menu icon reveals two commands: **Share** and **enable Default Share**. Select **Enable Default Share** to restore sharing with all users in all tenants. Select **Share** to bring up the Share dialog box for custom actions. This allows you to share the action with only certain, users, group and subtenants.

To delete a custom action that you have to manage permissions for, from the list of actions on the Actions Library page, for the action you want to delete, hover over the Action column until you see the dropdown menu icon appear. Once it appears, click on the dropdown icon to display the dropdown menu. Click the **Delete** command from the dropdown menu. If the action is not a lifecycle action referenced in a service definition, application profile, or region tab, Workload Manager displays a dialog box to confirm the deletion. If the action is a custom lifecycle action referenced by a resource, Workload Manager will display an information box listing all of the resources where the action is referenced and explain that all of those references to the lifecycle action must be deleted before the action can be deleted.

If you are a tenant administrator for the root tenant, you have the ability to disable some or all of the Workload Manager OOB on-demand actions via toggle switches in the **Enable** column on the [Actions Library Page](#). To disable an OOB on-demand action, click the corresponding Enable toggle which sets it to OFF. When you set an Enable toggle to OFF, the corresponding on-demand action will disappear from these areas of the Workload Manager UI for all users:

- the Virtual Machines List page,
- the Virtual Machine Details page,
- the VMs tab within a VM-based tier in the tiers tab of the Job Details page.

To re-enable a disabled OOB on-demand action, click the corresponding Enable toggle again which sets it to ON. The corresponding action will reappear in the three Workload Manager UI screens mentioned above.

Create Snapshot

Create Snapshot

- [Overview](#)
- [Limitations](#)
- [On-Demand Action](#)
- [The Create Snapshot Action](#)
- [Maximum Snapshot Limit](#)
- [Managing Snapshots](#)

A snapshot is a reproduction of the Virtual Machine (VM) captured when the snapshot was taken. The snapshot includes the state of the data on all VM disks and the VM power states (on, off, or suspended) – You can only take a snapshot when a VM is powered on, powered off, or suspended.

When you create a snapshot, the system creates a delta disk file for that snapshot in the datastore and writes any changes to that delta disk. You can later revert to the previous state of the VM.

Snapshot-related operations include creating snapshots and managing snapshots (reverting to any snapshot and removing snapshots).

This feature is only available for VMware vCenter environments.

On-Demand actions are manually invoked by the user from various screens in Workload Manager UI, including the Deployments List page, Deployment Details page, Virtual Machines List page, Virtual Machine Details page, and Application Profiles page.

The **Create Snapshot** Action is an on-demand action and allows you to take snapshot of a virtual machine (capture the settings state, disk state, and memory state at different specific times).

For more information on on-demand actions, see the [Actions Library Overview](#) section.


While Create Snapshot is an action that was available in earlier releases, the ability to display alert messages with the number of remaining snapshots and the Manage Snapshots link is being introduced in Workload Manager 5.2 for the following scenarios:

- When the retaining snapshot limit is 5 or below.
- When you reach the maximum snapshot limit.

For each region, you could and can continue to configure the Maximum Snapshot limit for VMs in the *Region Settings* page as displayed in the following image:

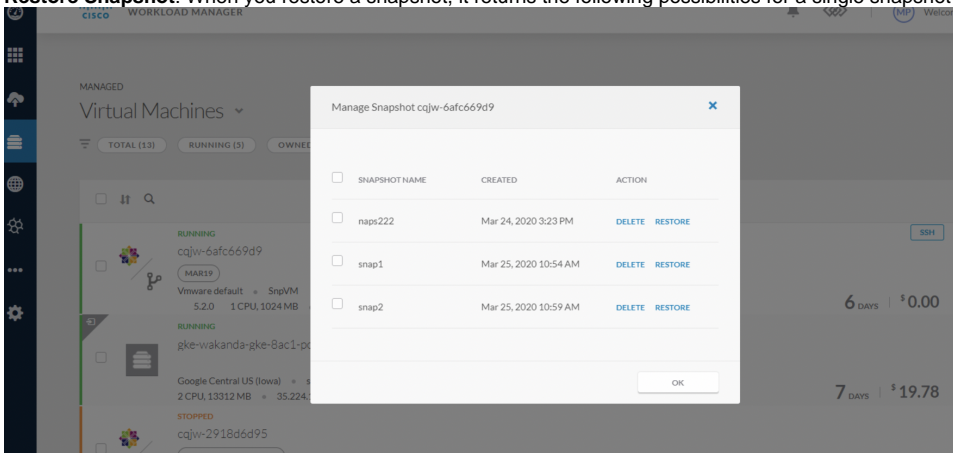
The screenshot shows the 'Cloud Settings' page in the Cisco Workload Manager interface. The page is divided into two tabs: 'Details' and 'Accounts'. The 'Details' tab is active, showing a list of configuration parameters for the cloud settings. The 'Number of snapshots allowed per VM' parameter is highlighted with an orange box and has a value of '3'. Other parameters include vCenter API Endpoint, VM Create Workflow, Exclude these special characters for Windows password, Concurrent Nodes Launches, Linux Max Sockets, Windows Max Sockets, Disable Custom Attributes, Allowed Root Disk Sizes (GB), Allowed Additional Volume Sizes (GB), Agent Bundle URL, Agent Custom Repository, HTTPS Proxy Host, HTTPS Proxy Port, HTTPS Proxy Username, HTTPS Proxy Password, HTTP Proxy Host, HTTP Proxy Port, HTTP Proxy Username, HTTP Proxy Password, and No Proxy Hosts. The 'Region Connectivity' status is shown as 'Running'. At the bottom, there are two status indicators: 'Cloud endpoint accessible from CloudCenter Suite' (Yes) and 'CloudCenter Suite AMQP reachable from worker VM's' (Yes).

The Manage Snapshots option is visible in the following screenshot:

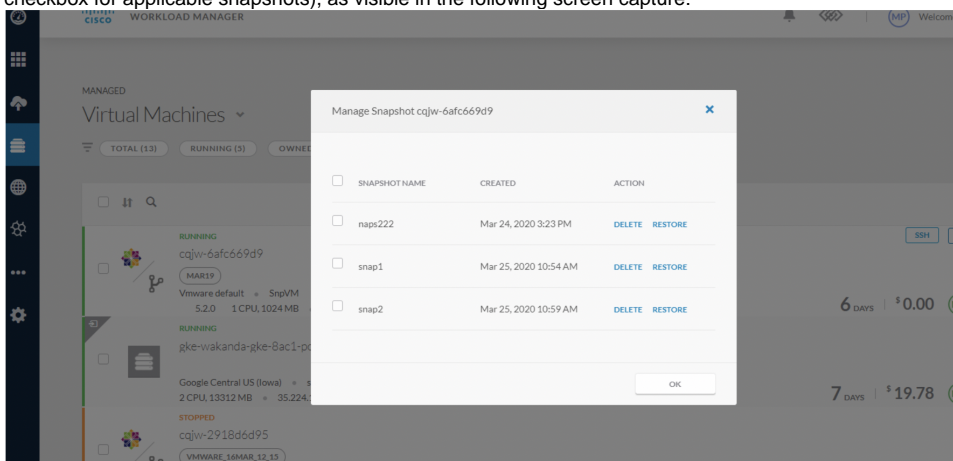


Managing Snapshots, a new feature introduced in Workload Manager 5.2.0, allows you to perform the following additional tasks:

- **Restore Snapshot:** When you restore a snapshot, it returns the following possibilities for a single snapshot (via the **Restore** link):



- The VM to its original state or to another snapshot in the snapshot hierarchy.
- The VM's memory, settings, and the state of the VM disks to the state they were in at the time of the snapshot.
- **Delete Snapshot:** When you restore a snapshot for a single snapshot or multiple snapshots (via the **Delete** link or checking the corresponding checkbox for applicable snapshots), as visible in the following screen capture:



- It removes the snapshot from the Snapshot Manager.

- The snapshot files are consolidated and written to the parent snapshot disk and merged with the VM base disk.

You can access the Manage Snapshot option from the following pages in the UI:

- Virtual Machines List page
- Virtual Machine Details page
- Job Details pPage (deployment details)

Manage Deployments and VMs

Manage Deployments and VMs

- [Deployment, VM, and Container States](#)
- [Virtual Machine Management](#)
- [Deployments List Page](#)
- [Deployment Details Page](#)
- [Terminate Protection](#)
- [Track Cloud Costs](#)
- [Project and Phase Management](#)
- [Migrating Applications](#)
- [Container Tier Rolling Updates](#)

Deployment, VM, and Container States

Deployment, VM, and Container States

- [Terminology](#)
- [Application Deployment States](#)
- [VM \(Node\) States](#)
- [Container States](#)
- [Orchestration Lifecycle Threshold Settings](#)
- [Permitted Job Operations and Actions](#)
- [ServiceNow Extension Workflow Requests](#)
- [Auto-Clean up of Resources](#)

Term	Description
<i>Job</i>	A job is a single action that contains many properties (for example, ownership, cloud cost, deployment information, and so forth). A job can contain other child jobs or is the child of a parent job.
<i>Deployment</i>	Refers to the application deployment. A deployment contains a set of jobs (for example, creating a VM). Deployments do not have the concept of child or parent, but merely refers to bundled jobs that share properties (for example, sharing properties between VMs in an application deployment).
<i>Deployed application</i>	An application running in a deployment environment.
<i>Deployment environment</i>	Provides shared access to multiple users on clouds or cloud zones for a specific use (for example, development): <ul style="list-style-type: none"> • When users with access permissions deploy an application to a deployment environment, these users can view and/or manage deployments based on their Permission Control level • See Deployment Environments for additional information.
<i>Lifecycle</i>	The set of states a deployment will go through.


The Workload Manager displays color-coded states for deployed applications in the **Deployments** page. The following table provides the description for each Deployment state.

UI Deployment State (Alphabetical Listing)	Lifecycle Order	Description	API Enumeration for <i>deploymentStatus</i>	API Enumeration for <i>jobStatus</i>
Deployment Rollback	9	There was an error during upgrade, because of which the deployment was rolled back.	DeploymentRollback	JobRollback
DeploymentRollback Error	10	There was an error in running the rollback script after an upgrade failure.	DeploymentRollbackError	JobUpgradeRollback Error
Error	22	The deployment encountered an error state.	DeploymentError	JobError
Finished	20	The deployment has finished executing or was canceled successfully. The nodes remain available for reuse until the next hour boundary, at which time they will be cleaned up if they are not reused.	DeploymentDone	JobFinished
Migrating	11	The deployment is in the process of migrating to another cloud (see Terminate Protection for additional context).	DeploymentMigrating	JobMigrating
Migration Error	13	The deployment encountered an error while migration to another cloud was in process.	DeploymentMigrationError	JobMigrationError
Migrate Pending	12	The deployment is pending approval for migration.	DeploymentMigratePending	JobMigratePending
Pending	2	If you launched a deployment in an environment that required approval and you are waiting for approval.	DeploymentPending	JobPending
Reconfiguring	6	The node is being reconfigured with the new settings.	DeploymentReconfiguring	JobReconfiguring
Rejected	4	If you launched a deployment in an environment that required approval and the phase approver rejected your request (see the previous section to set up approval).	DeploymentRejected	JobRejected
Resuming	16	The deployment was suspended and is in the process of resuming.	DeploymentResuming	JobResuming

Running/Deployed Running applies to batch, parallel, and interactive applications	3	Workload Manager has completed the orchestration steps and the job is running. The application is successfully deployed and all nodes are up and running (detected via a heartbeat within the heartbeat timeout interval window).	DeploymentDone ("Deployed")	JobRunning
Scaling	5	The deployment is in the process of increasing or decreasing the number of nodes for an app tier.	DeploymentScaling	JobScaling
Stopped	18	The deployment has been stopped and all nodes associated with the job have been cleaned up.	DeploymentStopped	JobStopped
Stopping	17	The deployment is being stopped.	DeploymentStopping	JobStopping
Stopping Error	19	An error occurred when the deployment was being stopped.	DeploymentStoppingError	JobStoppingError
Submitted	1	Workload Manager has received the request and started to process it but the orchestration steps are not yet completed.	DeploymentSubmitted	JobSubmitted
Suspended	15	The application VMs in this deployment have been powered off or shut down (not terminated).	DeploymentSuspended	JobSuspended
Suspending	14	The application is in the process of being suspended.	DeploymentSuspending	JobSuspending
Terminated	21	The application VMs and external volumes launched as part of the deployment are terminated.	DeploymentKilled	JobCanceled
Upgrading	7	The deployment is being upgraded to a more recent application profile version. Application profiles can have numerous versions. When launching an application, you may prefer to launch using the latest version of the profile. If you choose to do so, the deployment lists this state.	DeploymentUpgrading	JobUpgrading
Upgrading Error	8	The deployment encountered an error when upgrading the application profile version.	DeploymentUpgradeError	JobUpgradingError

Use the **Deployments > Application Deployments > Click Application** page to view the VM state for each job, as described in the following table.

VM State (Alphabetical Listing)	Lifecycle Order	Supported Actions	Description	API Enumeration for <i>nodeStatus</i>
Cleaned	11	None	The node is cleaned of all its configurations and dependent services based on the cleanup script provided by the user.	NodeCleaned
Error	12	None	The node encountered an error as a result of the most recent action. You can see error details in the job status message. For example: <i>Instance bootstrapping is timed out, possibly due to incorrect or missing agent bundle for node...</i>	NodeError
Not Reachable	13	None	The node is offline as a result of detecting a heartbeat loss.	NodeNotReachable
Reachable	14	None	The node is reachable after being in the not reachable state for any period of time.	NodeReachable
Rebooted	8	None	The node was restarted.	NodeRebooted
Resuming	6	None	The node was suspended and is in the process of resuming.	NodeResuming
Resumed	7	None	The node has resumed after the suspend state.	NodeResumed
Running	3	<ul style="list-style-type: none"> Power Off Reboot 	The node is started and the agent is running. This is the normal status for a running node.	NodeReady
Started	2	<ul style="list-style-type: none"> Power Off Reboot 	The node has started successfully and the initialization is complete.	NodeStarted
Starting	1	None	This is the first status message from the agent to indicate that the node initialization has started.	NodeStarting
Suspending	4	None	The node is in the process of being suspended.	NodeSuspending
Suspended	5	Power On	The node has been powered off or shut down (not terminated).	NodeSuspended
Hibernated	5	Power On	The node has been hibernated (vCenter deployments only).	NodeHibernated

Terminating	9	None	The node is in the process of being terminated	NodeTerminating
Terminated	10	None	The node is terminated in the cloud. Termination occurs as a result of the reduce action or the deployment termination action.  Terminating a job says NodeTerminated when the node is actually still in shutting-down state. The terminate job operations submits a request to the cloud to cleanup the node. The actual cleanup may happen a bit later.	NodeTerminated
Node not found	15	None	The node is not found by the Workload Manager either because the node is down for an unknown reason. This state is also displayed when an imported node is deleted from the cloud provider console instead of the Workload Manager	NodeNotFound

Container Clouds deployments provide a different set of states and actions.

Use the **Deployments > Application Deployments > Click Application** page to view the container state for each job, as described in the following table.

Container State (Alphabetical Listing)	Lifecycle Order	Supported Actions	Description	API Enumeration for <i>nodeStatus</i>
Running	1	<ul style="list-style-type: none"> • Terminate • Terminate and hide • Enable terminate protection • Share 	The pod is started and running. This is the normal status for a running pod.	RUNNING
Updating	2	None	The pod is in the process of being updated.	UPDATING
Terminated	3	Hide	The pod is deleted from the Kubernetes cluster. This state may be caused by the Terminate action or by an error.	TERMINATED
Deployed	4	None	If the pod does not start successfully after 10 tries, Kubernetes-displays the CrashLoopBackOff status in its console. However, the Workload Manager Job Details page displays the Deployed status.	DEPLOYED

A success of a lifecycle depends on the successful deployment of an application. You can set the timeout thresholds for some phases of the orchestration process. These threshold settings enable Workload Manager to proceed with the remainder of the lifecycle process when deploying an application.

This is the high-level orchestration lifecycle process when launching VMs:

1. Launch the required VMs as part of the application deployment process.
2. The VMs go through the bootstrap process.
 - a. If all the VMs are successfully bootstrapped (preinit) within their configured **bootstrap timeout**, then the VM bootstrap process is complete.
 - b. If some of the VMs reach the bootstrap timeout, then they are terminated. If the number of healthy VMs is less than the configured **minimum number of nodes**, then all VMs are terminated and the deployment ends in a failure.
 - c. If the number of healthy VMs is greater than or equal to **the minimum number of nodes**, then Workload Manager retries *once* to launch the failed VMs. Once again, Workload Manager terminates any VMs that reach the bootstrap timeout.
 - d. If the minimum number of VMs required for the deployment is met, Workload Manager proceeds to initialize the healthy VMs, thus ensuring partial success, even if all the VMs cannot be bootstrapped successfully.
3. Workload Manager runs the initialization scripts on the healthy VMs.
 - a. If the initialization scripts running on the VMs do not complete within the **node ready timeout** then Workload Manager terminates those VMs, without any retries.
 - b. Once the initialization scripts either complete or time out for all the VMs, Workload Manager checks to see if the initialization was successful for the minimum number of nodes.
 - c. If initialization was successful for the minimum number of nodes, then the deployment is completed for those VMs.
 - d. Otherwise, the deployment ends in a failure.

The orchestration lifecycle steps for pre-initialization (pre-init) and initialization in the high-level process relies on the numbers that you configure for the following settings.

The Min and Max size set in the **Application Tier Properties > General Settings** section is the minimum and maximum boundary from the scaling context.

However, the Workload Manager only allows deployments with PARTIAL SUCCESS when you additionally set at least one of the following parameters in the application profile.


- *deploymentMinClusterSize*
- *deploymentMinClusterPercentage*

The following table describes the Orchestration Lifecycle Threshold settings.

Setting	Application Parameter	Description
Bootstrap Timeout (time in seconds)	<i>cliqrNodeBootstrapTimeout</i>	<p>Default = 3600</p> <p>The <i>Bootstrap Timeout</i> setting identifies the maximum time available for VMs to bootstrap after they are launched.</p> <p>The deployment behavior when a few VMs fail to bootstrap is governed by the <i>Minimum Number of Nodes</i> setting.</p> <ul style="list-style-type: none"> • If some VMs time out and face a bootstrap failure, then those VMs are terminated. • If the number of VMs that bootstrap successfully is greater than (or equal to) the <i>Minimum Number of Nodes</i>, then Workload Manager tries a fresh launch of the failed VMs.
Minimum Number of VMs (VM count)	<i>deploymentMinClusterSize</i>	<p>You can specify the <i>Minimum Number of Nodes</i> and <i>Maximum Number of Nodes</i> in the Application Tier Properties > General Settings section determines the cluster boundary. While these settings are important in the scaling context, they are also important in the bootstrap timeout context as the Workload Manager eases the failure process for VMs that have reached the timeout limit. The deployment proceeds as long as the number of VMs that are successfully bootstrapped (despite the <i>Bootstrap Timeout</i>) is greater than or equal to the <i>Minimum number of Nodes</i> setting and ensures <i>partial success</i> even if the system encounters a bootstrap failure for a few VMs.</p> <p>To configure <i>deploymentMinClusterSize</i>, you must add this setting as a custom parameter as specified in the Deployment Parameters.</p> <p>If you add both <i>deploymentMinClusterSize</i> and <i>deploymentMinClusterPercentage</i> as custom Deployment Parameters, the <i>deploymentMinClusterSize</i> parameter takes precedence.</p> <p>If <i>deploymentMinClusterSize</i> is not configured:</p> <ul style="list-style-type: none"> • And <i>deploymentMinClusterPercentage</i> is configured, then <i>deploymentMinClusterPercentage</i> takes precedence. • And <i>deploymentMinClusterPercentage</i> is ALSO NOT configured, then the Workload Manager expects all launched nodes to complete bootstrapping and will not accept partial success scenarios.
Minimum Percentage of VMs (manually defined)	<i>deploymentMinClusterPercentage</i>	<p>Percentage which governs the minimum number of nodes that the Workload Manager should use to calculate and mark the deployment as <i>complete with partial success</i>.</p> <p>You may need to adjust this parameter until you arrive at an acceptable percentage that will suit all deployment sizes in your environment. For example, setting the value of <i>deploymentMinClusterPercentage</i> to 95 will respond differently for different deployment sizes as provided in the sample scenarios below:</p> <ul style="list-style-type: none"> • For a deployment with 100 nodes, 95 nodes will meet the partial success criteria • For a deployment with 500 nodes, 475 nodes will meet the partial success criteria • For a deployment with 1000 nodes, 950 nodes will meet the partial success criteria and so forth <p>To configure <i>deploymentMinClusterPercentage</i>, you must add this setting as a custom parameter as specified in the Deployment Parameters.</p> <p>If you add both <i>deploymentMinClusterSize</i> and <i>deploymentMinClusterPercentage</i> as custom Deployment Parameters, the <i>deploymentMinClusterSize</i> parameter takes precedence.</p> <p>If <i>deploymentMinClusterPercentage</i> is not configured:</p> <ul style="list-style-type: none"> • And <i>deploymentMinClusterSize</i> is configured, then <i>deploymentMinClusterSize</i> takes precedence. • And <i>deploymentMinClusterSize</i> is ALSO NOT configured, then the Workload Manager expects all launched nodes to complete bootstrapping and will not accept partial success scenarios.

<p>Node Ready Timeout (time in seconds)</p>	<p><i>cliqrNodeReadyTimeout</i></p>	<p>Default = 36000</p> <p>Sometimes, a VM may function but the associated application service may not have started on this VM when the initialization scripts are stuck or take too long to complete. The <i>Node Ready Timeout</i> setting ensures that the timed out VMs are terminated. In this case, the remaining VMs complete the application process as long as the requirement for the <i>Minimum Number of Nodes</i> is met. The termination of the VMs affected by the time out does not hinder the completion process for the remaining VMs and ensures partial success for the entire Orchestrator lifecycle process.</p> <p>The nodes that are terminated by this process are not given another chance – the Workload Manager does not re-launch these terminated VMs – instead, it displays an Error message for this kind of a timeout.</p> <table border="1" data-bbox="402 390 1349 842"> <thead> <tr> <th>Task Name</th> <th>Task Status</th> <th>Last Update Time</th> <th>Last Message</th> </tr> </thead> <tbody> <tr> <td>task0</td> <td>TaskRunning</td> <td>9/11/17 18:13:00.026</td> <td>Instance initialize is timed out for node i-091481d0123661615</td> </tr> <tr> <td>nodeInit</td> <td>TaskRunning</td> <td>9/11/17 18:11:01.775</td> <td>Execute node init script...</td> </tr> <tr> <td>nodeInit</td> <td>TaskRunning</td> <td>9/11/17 18:11:01.745</td> <td>Execute node setsudo: ALL</td> </tr> <tr> <td>nodeInit</td> <td>TaskRunning</td> <td>9/11/17 18:10:59.816</td> <td>Initializing cluster node...</td> </tr> <tr> <td>nodePreInit</td> <td>TaskRunning</td> <td>9/11/17 18:06:32.964</td> <td>Checking if cliqr tools are installed... Yes</td> </tr> <tr> <td>nodePreInit</td> <td>TaskRunning</td> <td>9/11/17 18:06:32.953</td> <td>Pre initializing cluster node...</td> </tr> <tr> <td>nodeBooting</td> <td>TaskRunning</td> <td>9/11/17 18:04:57.144</td> <td>Node booting up...</td> </tr> </tbody> </table>	Task Name	Task Status	Last Update Time	Last Message	task0	TaskRunning	9/11/17 18:13:00.026	Instance initialize is timed out for node i-091481d0123661615	nodeInit	TaskRunning	9/11/17 18:11:01.775	Execute node init script...	nodeInit	TaskRunning	9/11/17 18:11:01.745	Execute node setsudo: ALL	nodeInit	TaskRunning	9/11/17 18:10:59.816	Initializing cluster node...	nodePreInit	TaskRunning	9/11/17 18:06:32.964	Checking if cliqr tools are installed... Yes	nodePreInit	TaskRunning	9/11/17 18:06:32.953	Pre initializing cluster node...	nodeBooting	TaskRunning	9/11/17 18:04:57.144	Node booting up...
Task Name	Task Status	Last Update Time	Last Message																															
task0	TaskRunning	9/11/17 18:13:00.026	Instance initialize is timed out for node i-091481d0123661615																															
nodeInit	TaskRunning	9/11/17 18:11:01.775	Execute node init script...																															
nodeInit	TaskRunning	9/11/17 18:11:01.745	Execute node setsudo: ALL																															
nodeInit	TaskRunning	9/11/17 18:10:59.816	Initializing cluster node...																															
nodePreInit	TaskRunning	9/11/17 18:06:32.964	Checking if cliqr tools are installed... Yes																															
nodePreInit	TaskRunning	9/11/17 18:06:32.953	Pre initializing cluster node...																															
nodeBooting	TaskRunning	9/11/17 18:04:57.144	Node booting up...																															

For each valid state, Workload Manager displays corresponding operations and actions based on configured permissions.

 If you deploy a container-based hybrid or container-only application, some of the actions for N-tier jobs may not apply. In most cases, only the *Terminate* action is applicable.

The following table describes the actions that the Workload Manager displays for N-tier jobs.

Deployment Status	Supported Actions	Notes
Error	<ul style="list-style-type: none"> Terminate Terminate and Hide 	<ul style="list-style-type: none"> User must have Manage privilege on own/other deployments in the deployment environment used in deploying this application Deployment must not be terminated
Rejected	<ul style="list-style-type: none"> Terminate Terminate and Hide Hide 	<ul style="list-style-type: none"> User must have Manage privilege on own/other deployments in the deployment environment used in deploying this application Deployment must not be terminated
Stopped	<ul style="list-style-type: none"> Resume 	<ul style="list-style-type: none"> User must have Manage privilege on own/other deployments for the deployment environment used in deploying this application
	<ul style="list-style-type: none"> Terminate Terminate and Hide 	<ul style="list-style-type: none"> User must have Manage privilege on own/other deployments for the deployment environment used in deploying this application Deployment is not terminated
Stopping Error	<ul style="list-style-type: none"> Terminate Terminate and Hide 	<ul style="list-style-type: none"> User must have Manage privilege on own/other deployments for the deployment environment used in deploying this application Deployment is not terminated

<ul style="list-style-type: none"> • Deployed • Upgrade Rollback Error • Upgrade Rollback Error • Upgrade Error 	Suspend	<ul style="list-style-type: none"> • User must have Manage privilege on own/other deployments for the deployment environment used in deploying this application • Prevent Terminate Protection and <i>terminateProtection</i> must be disabled (off) in the deployment environment used in deploying this application • Suspend is supported
	Stop	<ul style="list-style-type: none"> • User must have Manage privilege on own/other deployments for the deployment environment used in deploying this application • Prevent Termination/<i>terminateProtection</i> must be disabled (off) in the deployment environment used in deploying this application • Suspend is not supported
	Terminate Terminate and Hide	<ul style="list-style-type: none"> • User must have Manage privilege on own/other deployments for the deployment environment used in deploying this application • Prevent Termination/<i>terminateProtection</i> must be disabled (off) in the deployment environment used in deploying this application • Deployment is not terminated
	Upgrade	<ul style="list-style-type: none"> • User must have Manage privilege on own/other deployments for the deployment environment used in deploying this application
	Promote	<ul style="list-style-type: none"> • User must have Promote From privilege for the deployment environment used in deploying this application
	Migrate	<ul style="list-style-type: none"> • User must have Promote From privilege for the deployment environment used in deploying this application • Prevent Termination/<i>terminateProtection</i> must be disabled (off) in the deployment environment used in deploying this application
	Enable/Disable Terminate Protection	<ul style="list-style-type: none"> • User must have Manage privilege on own/other deployments for the deployment environment used in deploying this application • Termination/<i>terminateProtection</i> must be supported
	Share VM Access	<ul style="list-style-type: none"> • Visible to users that have Manage privilege on own/other deployments for the deployment environment used in deploying this application • Lets the user grant SSH or VNC access to other users, groups or subtenants. • Only effective if the <i>targeted</i> user, group or subtenant also has Access or Manage privilege on own/other deployments for the deployment environment used in deploying this application
<ul style="list-style-type: none"> • Migrating • Upgrading • Reconfiguring 	Suspend	<ul style="list-style-type: none"> • User must have Manage privilege on own/other deployments for the deployment environment used in deploying this application • Prevent Termination/<i>terminateProtection</i> must be disabled (off) in the deployment environment used in deploying this application • Suspend is supported
	Stop	<ul style="list-style-type: none"> • User must have Manage privilege on own/other deployments • Prevent Termination/<i>terminateProtection</i> must be disabled (off) • Suspend is not supported
	Enable/Disable Terminate Protection	<ul style="list-style-type: none"> • User must have Manage privilege on own/other deployments for the deployment environment used in deploying this application • Prevent Termination/<i>terminateProtection</i> must be supported

<ul style="list-style-type: none"> In Progress Submitted 	Terminate Terminate and Hide	<ul style="list-style-type: none"> User must have Manage privilege on own/other deployments for the deployment environment used in deploying this application Prevent Termination/<i>terminateProtection</i> must be disabled (off) Deployment is not terminated
	Enable/Disable Terminate Protection	<ul style="list-style-type: none"> User must have Manage privilege for the deployment environment used in deploying this application
Migration Error	Terminate Terminate and Hide	<ul style="list-style-type: none"> User must have Manage privilege on own/other deployments Prevent Termination/<i>terminateProtection</i> must be disabled (off)
	Enable/Disable Terminate Protection	User must have Manage privilege for the deployment environment used in deploying this application
Suspended	Terminate Terminate and Hide	<ul style="list-style-type: none"> User must have Manage privilege on own/other deployments Prevent Termination/<i>terminateProtection</i> must be disabled (off)
	Resume	User must have Manage privilege on own/other deployments for the deployment environment used in deploying this application
	Enable/Disable Terminate Protection	User must have Manage privilege for the deployment environment used in deploying this application
<ul style="list-style-type: none"> Suspending Resuming Stopping 	Terminate Terminate and Hide	<ul style="list-style-type: none"> User must have Manage privilege on own/other deployments for the deployment environment used in deploying this application Prevent Termination/<i>terminateProtection</i> must be disabled (off) in the deployment environment used in deploying this application
	Enable/Disable Terminate Protection	User must have Manage privilege for the deployment environment used in deploying this application
<ul style="list-style-type: none"> Pending Migrate Pending 	Approve	User must have Approve privilege for the deployment environment used in deploying this application
	Reject	User must have Approve privilege for the deployment environment used in deploying this application
	Enable/Disable Terminate Protection	User must have Manage privilege for the deployment environment used in deploying this application
Scaling	Suspend Terminate Terminate and Hide	<ul style="list-style-type: none"> User must have Manage privilege on own/other deployments for the deployment environment used in deploying this application Prevent Termination/<i>terminateProtection</i> must be disabled (off) in the deployment environment used in deploying this application
	Enable/Disable Terminate Protection	User must have Manage privilege for the deployment environment used in deploying this application
Terminated	Hide	User must have Manage privilege for the deployment environment used in deploying this application

The following table describes the action and status that the Workload Manager displays for deployments/VMs.

<i>approvalRequestAction</i>	<i>approvalRequestStatus</i>	Description
DEPLOY	Deployment Pending Approval	The approval request action and status for the deployment submission
SUSPEND	Suspend Pending Approval	The approval request action and status to suspend the deployment
RESUME	Resume Pending Approval	The approval request action and status to resume the deployment
PROMOTE	Promote Pending Approval	The approval request action and status to promote the deployment

RERUN	Rerun Pending Approval	The approval request action and status to rerun the deployment
MIGRATE	Migrate Pending Approval	The approval request action and status to migrate the deployment
DEMOTE	Demote Pending Approval	The approval request action and status to demote the deployment
TERMINATE	Terminate Pending Approval	The approval request action and status to terminate the deployment
START	Start Pending Approval	The approval request action and status to start the VM
REBOOT	Reboot Pending Approval	The approval request action and status to reboot the VM
STOP	Stop Pending Approval	The approval request action and status to stop the VM
EXTEND_AGE	Extend_Age Pending Approval	The approval request action and status to extend the aging policy of a deployment

A Cron utility runs in the background at 02:00 hours (CloudCenter Suite system time) and automatically terminates and deletes jobs/deployments that meet the following conditions:

- The job/deployment is in a *JobError* state.
- The job/deployment is running on a *Managed* VM (see [Virtual Machine Management](#) for additional context).
- The instances/volumes related to this job/deployment continue to run or have not been terminated by the user, they will be terminated in the background.
- The troubleshooting parameter, `cliqrIgnoreAppFailure = false` (see [Troubleshooting Parameters](#) for additional context).



If `cliqrIgnoreAppFailure = true`, then the instances/volumes are left untouched.

The auto-clean utility:

- Retrieves the list of jobs/deployments (latest 500 ordered by time) that meet the conditions listed above.
- Deletes the related instances of the job/deployment at 02:0 hours (CloudCenter Suite system time).
- Updates the status for this job/deployment in the VM Details page – you can check if the node has been terminated by this background Cron utility by referring to the corresponding VM details page in the Workload Manager UI.

Virtual Machine Management

Virtual Machine Management

- [Overview](#)
- [Cloud Nuances](#)
- [Permissions and Access Control](#)
- [Managed VMs](#)
 - [VM Details](#)
 - [VM Errors](#)
- [Unmanaged VMs](#)
- [Import to Workload Manager](#)
- [Share VMs](#)
- [Imported VMs](#)
- [Download CSV File for Managed VMs](#)
- [Install the Management Agent](#)
 - [Prerequisites to Use the Management Agent](#)
 - [Install the Management Agent from the UI](#)
 - [Install the Management Agent Manually](#)
- [Set the Management Agent Environment Variables](#)
- [Manage the Agent](#)
- [Uninstall the Management Agent](#)
- [Upgrade Agent](#)
- [Perform VM Actions](#)
- [Billing](#)
- [Agent Time Zone Details](#)
- [Additional Functions](#)
- [Time Period Filter](#)
- [Advanced Filters](#)
 - [Save Filters](#)
 - [Delete Saved Filters](#)

Datacenters that support several thousand VMs sometimes require support across multiple clouds. Workload Manager provides a Virtual Machine (VM) management feature for such datacenters. This feature allows you to import VMs into Workload Manager and manage them directly from the UI.

From the Workload Manager perspective, VMs have two categories. Both categories are included in the CCM UI under a new tab called **Virtual Machines**. The following table describes the categories.

Category	Description	Visibility	Permitted Actions
Managed VMs	Displays VMs that are already managed by Workload Manager. This list includes Workload Manager deployed VMs and imported VMs. VMware Tools must be installed and running on VMs before users can import these VMs into Workload Manager. Refer to the VMware documentation for additional details.	All users	<ul style="list-style-type: none"> • Start • Stop • Terminate • Reboot
Unmanaged VMs	Displays VMs that are not yet managed, by Workload Manager. This list includes VMs discovered by Workload Manager. To manage a VM displayed in this list, you must first import the VM to Workload Manager.	Admin users	<ul style="list-style-type: none"> • Import to Workload Manager • Terminate

Once you import a VM from the Unmanaged list to the Managed list, VM actions are available based on the underlying cloud. Additional VM actions are available for the following clouds:

- AWS
- AzureRM
- Google Cloud
- OpenStack
- VMware vCenter

A Workload Manager user with admin permissions can import a VM listed in the **Unmanaged** into the **Managed** category.

Imported VMs do not have any default [Understand ACLs](#) permission. Permissions are derived based on the importer's invitations when importing the VM – you may have permission to log into the VM imported to Workload Manager but may not have permission (as a Workload Manager user) to upgrade the agent for this VM.

The VMs displayed in this list includes the following VMs:

- *Workload Manager Deployed VMs* display 2 logos – the Application logo (if available) and the OS logo.
- *Imported VMs* display 1 logo – the Imported VM icon.

The following screenshot displays a filtered list of *Workload Manager Deployed* VMs displayed in the **Managed** category.




The screenshot displays the 'Managed' tab of the Workload Manager interface. At the top, summary statistics are shown: 10 Total VMs, 886 Running VMs, \$239 Cloud Cost, \$288 Est Monthly Cost, and 5,611 VM Hours. Below this is a table of VMs with columns for status, name, configuration, and cost. The table includes VMs like 'windows' (Running, \$11.656), 'cqjw-cf6712cbd' (Stopped, \$0.15), 'cloudworker-b846b23e-5252-476e-8bbf-54d46a4c1c46' (Running, \$42.68), 'cqjw-b674a9174' (Running, \$16.054), 'cqjw-3b0b713a8' (Running, \$16.054), 'artifacts' (Running, \$50.881), 'i-0ce50b441d7cfadaf' (Running, \$85.808), and 'cqjw-f9db61337' (Running, \$3.000).

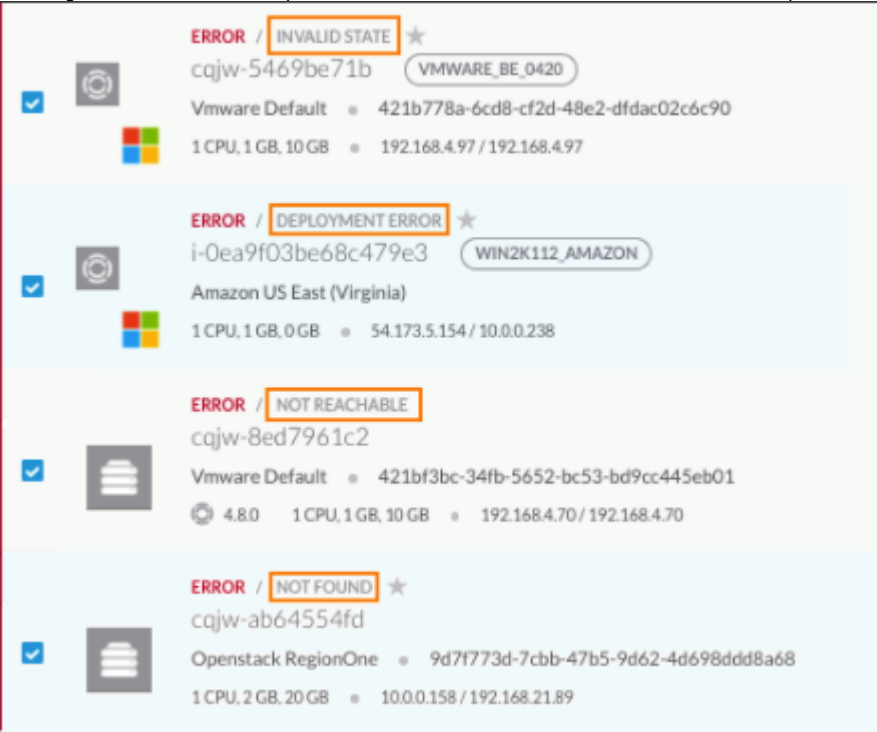
Regardless of the default filter settings, the information in the following table applies to the summary displayed at the top of the Virtual Machines page:

Summary	Description
Total VMs	<p>The total number of running VMs for the selected time period.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> This count depends on the selected filters. </div>
Running VMs	<p>The total number of running (billed) VMs for the entire deployment <i>without any time restriction</i>.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> This count includes VMs that display the ERROR / NOT REACHABLE status. </div> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> The term Running VMs in this summary differs from the term Running in the status-based filter. To co-relate the Running VMs count in the summary, check Running, Starting, and Error (NODE NOT REACHABLE) Statuses. </div>
Cloud Cost	The cost of running the VMs

Est Monthly Cost	The estimated hourly rate of running VMs (based on the VM status)
VM Hours	The total number of running VM hours during the selected time period

The following table identifies various aspects of the Virtual Machines tab:

Identity	Screenshot and Description																
Favorite	 <p>Mark any VM as a favorite by clicking the star icon next to the VM.</p>																
Parent deployment link	 <p>Displays the deployment name as a link. Click the link to view details about the parent deployment</p>																
VM details link	 <p>Displays the VM name as a link. Click the link to view details about the deployment. Depending on the cloud, the information displayed for this link differs:</p> <ul style="list-style-type: none"> • Hostname: The <i>hostName</i> for the VM – if configured. If not configured, then the <i>nodeId</i> is displayed. • Display Name: AWS VMs display the user-configured name for this VM – if configured. If not configured, then the <i>nodeId</i> is displayed. • VM ID: The <i>nodeId</i> is displayed for all other cases – the <i>vmId</i> is a unique identifier but it is not used for any VM operation, it is only used for the metadata purposes. 																
High-level status	<p>Color-coded status indicator to identify the high-level status of the VM displayed in the Virtual Machine tab, they do not indicate the status of the deployment. The following table describes the statuses.</p> <table border="1"> <thead> <tr> <th>Displayed Status</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>Error</td> <td>Identifies several types of errors as displayed by the VM status. See the <i>Error VMs Status</i> row below for additional context.</td> </tr> <tr> <td>Stopping</td> <td>Identifies a VM that is in the NodeSuspending state.</td> </tr> <tr> <td>Stopped</td> <td>Identifies a VM that is in the NodeSuspended state.</td> </tr> <tr> <td>Starting</td> <td>Identifies a VM that is in the NodeStarting and NodeResuming states.</td> </tr> <tr> <td>Running</td> <td>Identifies a VM that is in the NodeStarted, NodeReady, NodeResumed, NodeRebooted, and NodeReachable states.</td> </tr> <tr> <td>Terminating</td> <td>Identifies a VM that is in the NodeTerminating state.</td> </tr> <tr> <td>Terminated</td> <td>Identifies a VM that is in the NodeTerminated and NodeCleaned states.</td> </tr> </tbody> </table> <p>See Deployment, VM, and Container States for a complete list and additional details.</p>	Displayed Status	Description	Error	Identifies several types of errors as displayed by the VM status . See the <i>Error VMs Status</i> row below for additional context.	Stopping	Identifies a VM that is in the NodeSuspending state.	Stopped	Identifies a VM that is in the NodeSuspended state.	Starting	Identifies a VM that is in the NodeStarting and NodeResuming states.	Running	Identifies a VM that is in the NodeStarted, NodeReady, NodeResumed, NodeRebooted, and NodeReachable states.	Terminating	Identifies a VM that is in the NodeTerminating state.	Terminated	Identifies a VM that is in the NodeTerminated and NodeCleaned states.
Displayed Status	Description																
Error	Identifies several types of errors as displayed by the VM status . See the <i>Error VMs Status</i> row below for additional context.																
Stopping	Identifies a VM that is in the NodeSuspending state.																
Stopped	Identifies a VM that is in the NodeSuspended state.																
Starting	Identifies a VM that is in the NodeStarting and NodeResuming states.																
Running	Identifies a VM that is in the NodeStarted, NodeReady, NodeResumed, NodeRebooted, and NodeReachable states.																
Terminating	Identifies a VM that is in the NodeTerminating state.																
Terminated	Identifies a VM that is in the NodeTerminated and NodeCleaned states.																

<p>Error VM Status</p>	<p>This information is only provided if a Workload Manager deployed VM is in the Error state – Identifies one of the types of errors that the following screenshot shows and provides additional details on cause and correction tool tips.</p>  <p>If a VM is not deployed by Workload Manager is in the error state, then Error / Not Found state is displayed.</p>
<p>Management Agent</p>	<p>The Management Agent (also referred to as <i>agent</i>) can be installed on VMs that have already been imported into Workload Manager. This agent is an alternate option for VMs that do not require the capability to launch applications but do require some basic Workload Manager functionality like performing platform actions. If installed, the Virtual Machines page and the VM Details page displays the icon and version as identified in the following screenshot.</p>  <p>The benefit to installing this agent on an imported VM is that you have additional actions types that become available for this VM.</p> <p>You can only install this agent on Managed VMs – You can have Workload Manager Deployed VMs or Imported VMs that are managed by Workload Manager but do not have agent installed.</p> <p>You cannot install this agent on Unmanaged VMs. You must first import the VM to the Managed list and then install the agent on that VM.</p>
<p>SSH</p>	 <p>Click to SSH into the VM. See Specify SSH Options for additional details.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p>i This option is not visible to users with <i>onlyView</i> permissions. Users require Manage permission to view this option that is only available on Managed VMs. Users will not see this option for Managed VMs that are in a terminated state.</p> </div>

VM Details

When you click the link for a VM in the Virtual Machines tab, you see the server Details page, as shown in the following screenshot:

- The **Details** tab (default) provides exhaustive details for the *VM*.

If any of your VM details are missing or showing zero values, click the Sync VM Details button to refresh the data.

- The **Logs** tab provides the entire list of *VM deployments* details.
- The **History** tab provides a complete history for all *actions* (succeeded, failed, occurred) performed on this VM.
- The **Available Actions** (in the Actions panel) that you can perform for this VM. See [Actions Library](#) for additional context.
- The **IPv6 Address** field identifies validated IPv6 addresses. See [IP Address Allocation](#) for additional context.

VM Errors

VM errors, if any, are displayed at the top of the Virtual Machines page and the VM Details page, as shown in the following screenshot:

- Click the **X** to dismiss the error.
- Click **View Details** to access the reason for the error.
- Click **Dismiss** to remove it from being listed in the page.

The VMs displayed in this list include VMs discovered by Workload Manager. These VMs were launched/deployed outside of Workload Manager. When Workload Manager connects to a cloud account, all VMs in that cloud account are displayed in the Unmanaged VMs list.

This category is only visible to Admin users – the administrator *and* the cloud account owner to import a VM from this category into the **Managed** category.

The following screenshot displays a filtered list of *ALL* VMs displayed in the **Unmanaged** category.


A Workload Manager user who is the administrator *and* who is the cloud account owner can import a VM listed in the **Unmanaged** into the **Managed** category. To Import a VM, you must:

- Import a VM which into Workload Manager only if it is already running.
- Be the administrator (only Admin users can create cloud accounts).
- Be the cloud Account Owner – the following screenshot illustrates where you can locate the account owner permissions for each cloud region.

You can import unmanaged VMs in one of two ways:

- **Individually:** Click the dropdown arrow next to the VM and selecting **Import to Workload Manager**. The following screenshot illustrates importing an unmanaged VM individually.

- **Batch:** You can also multi-select VMs by clicking the corresponding check boxes next to each VM and then selecting **Import to Workload Manager** from the Actions menu which displays the number of selected VMs.


 When performing bulk operations, if two instances have the same name then Workload Manager rejects this bulk operation request.

The following screenshot illustrates a batch import configuration.



See [Actions Library](#) for additional context on types of actions and other details.

Imported VMs do not have any default [Understand ACLs](#) permission. Permissions are derived based on the importer's invitations when importing the VM.



Once imported to Workload Manager, the VM is considered to be an **Imported VM** and the following behavior applies to this VM:

- Only listed in the **Managed** category
- Workload Manager licensing and billing begins as soon as the import is successful
- Available for VM actions (see [Actions Library](#) for additional context).
- Visible in Workload Manager but still does not have an agent installed
- Eligible to have an Agent screenshot installed

The following screenshot displays a filtered list of *Imported VMs* displayed in the **Managed** category and identifies the imported VM icon as well as the icon for an unknown OS:

When you click the *Imported VM link* to view details about the imported VM, a page as shown in the following screenshot displays:


Click the **Download CSV** option to download global and deployment-specific parameters for a managed VM to a Comma-Separated Value (CSV) file. By default all global and deployment-time parameters are included in the CSV file that is downloaded to your default download location.

You can also use an optional parameter, *includeJobParameters=true*, to include deployment-specific parameters when downloading this file using the [CCM Calls 5.4.0 > virtual-machines-controller](#) API calls. By default the value for this parameters is *false*. If set to true, the API call additionally returns all the global and deployment-time parameters for the corresponding VM.


If a VM was launched as part of an application deployment using Workload Manager, then the Management Agent may already be running on it (unless the VM was installed as an Agentless node or a user manually stopped the Agent). Such VMs (with the Management agent installed as part of deployment) are referred to as **Workload Manager Deployed VMs**.

Installing the Management Agent is an alternate option for VMs that do not require the capability to launch applications but do require some basic Workload Manager functionality like performing platform actions. You can only install the agent on **imported VMs**.

The following screenshot identifies an imported VM without the agent installation:



If the Management Agent is installed, you see the Agent icon and the version displayed for each Imported VM on the Virtual Machines page, as shown in the following screenshot:



Use one of the following options to install an agent:

i If the latest version of the agent is already installed on a VM, then the **Install Workload Manager Agent** action will no longer be available for this VM.

- **From the UI:** See the *Install the Management Agent from the UI* section below for additional details.
- **Manually:** See the *Install the Management Agent Manually* section below for additional details.

Prerequisites to Use the Management Agent

To install the Management Agent, you must meet the following conditions, you must meet the following conditions:

- **Supported Clouds:** Amazon, AzureRM, OpenStack, VMware vCenter, and Google Cloud
- **Supported OS:** Ubuntu 14/6, CentOS 6/7, RHEL 6/7, and Windows 2012/2016.
- **Required Utilities:** Clouds like VMware vCenter and AzureRM requires the following utilities:
 - **WindowsVMs:** wmic
 - **LinuxVMs:** dmidecode

Install the Management Agent from the UI

The **Install Workload Manager Agent** action installs the agent and enables Custom actions on this VM. After installing the agent, the list of allowed custom actions is listed in the Actions dropdown for this VM.


Requirements

To install the Workload Manager agent from the CCM UI, you must have the following credentials:

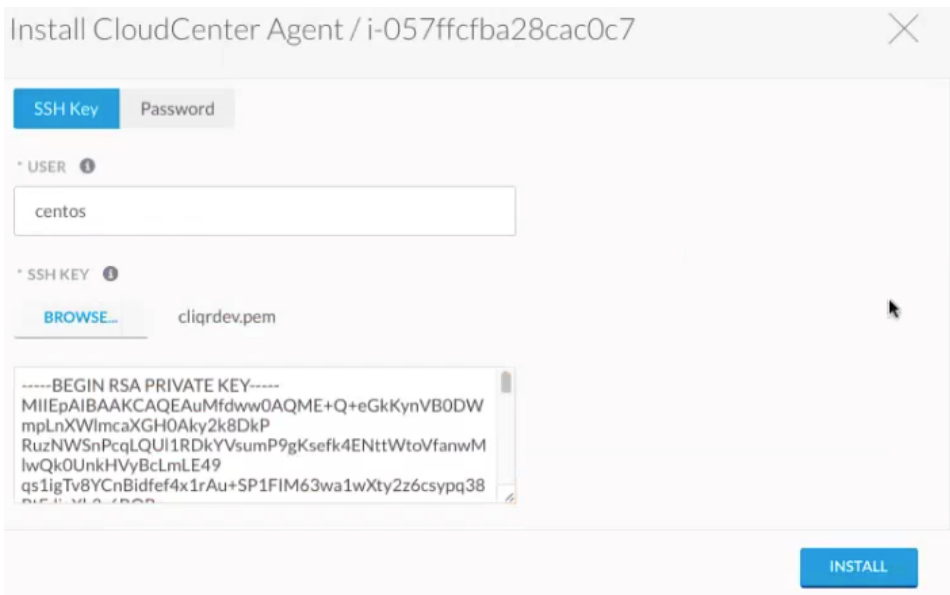
- **Linux-based VMs:**
 - SSH credentials of user with root privileges
 - Sudo access to the VM
- **Windows-based VMs** – Admin Login credentials

To install the lightweight agent on an imported VM, follow this procedure:

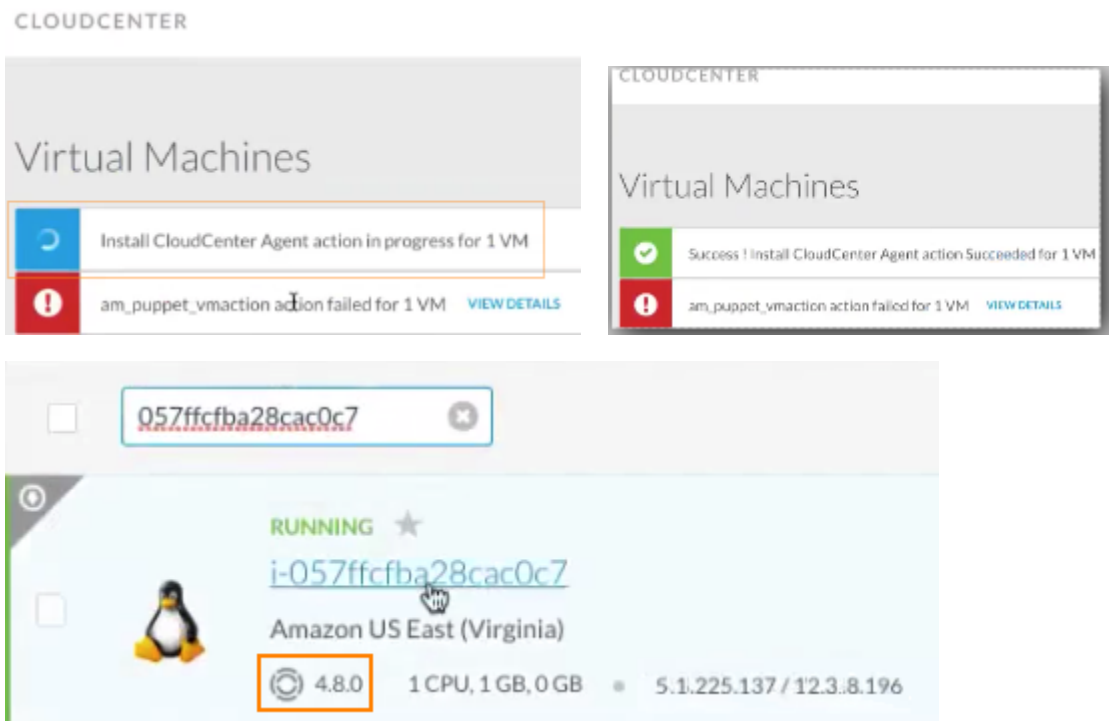
1. Identify the VM that requires the agent to be installed.
2. From the Actions dropdown, select **Install Workload Manager Agent**, as shown in the following screenshot.




3. The *Install Workload Manager Agent for this VM* displays, as shown in the following screenshot.



4. Identify the OS Type and the Authentication Type for the installation.
 - a. OS Type: The OS installed on this VM – **Linux** or **Windows**
 - b. Authentication Type:
 - i. **SSH Key** – provide key
 - ii. **Password** – provide password
 - c. User: Can be a **root** user (if you have root permissions for this VM) or a *configured* user (*centos* in this example) with sudo privileges for installation tasks.
 - d. Click **Install**. The agent installation commences and the Virtual Machines page displays the status at the top of the page. On successful completion, you see the success message and the agent icon for this VM, as shown in the following screenshots:



5. Click the Node name for this VM and click the History tab, as shown in the following screenshot.



6. Verify that the custom actions you configured are displayed in the Actions panel along with the platform actions (see [Actions Library](#) for details).

- a. Verify that previously configured custom actions (if any) are displayed in the Actions panel *along* with the platform actions (see [Actions Library](#) for details).
- b. Verify that the history reflect the right events.

7. Click **View Details** for the agent installation event and verify that the bundle file highlighted as displayed in the following screenshot.

Install CloudCenter Agent Action Details

```

+ '[' -f agent-lite-linux-bundle.tar.gz ']'
+ echo 'Downloaded http://build-rel.cliqr.com/release-4.8.0-20170425/bundle/agent-lite-linux-bundle.tar.gz
successful'
Downloaded http://build-rel.cliqr.com/release-4.8.0-20170425/bundle/agent-lite-linux-bundle.tar.gz successful
+ export AGENT_HOME=/usr/local/agentlite
+ AGENT_HOME=/usr/local/agentlite
+ '[' -d /usr/local/agentlite ']'
+ tar -xvf agent-lite-linux-bundle.tar.gz -C /usr/local/
agentlite/
agentlite/bin/
agentlite/bin/agent-pre-upgrade.sh

```

You have successfully installed the agent on an imported VM!

Install the Management Agent Manually

Install the Management Agent on an imported VM by following a manual procedure to login to the VM as a Root or Admin, install the agent manually by logging into the VM, and downloading the lightweight bundle.



Use the manual procedure in the following situations:

- **Lack of Credentials:** You may have permission to log into the VM imported to Workload Manager but may not have permission (as a Cisco Workload Manager user) to upgrade the agent for this VM.
- **Lack of Remote Access:** The OS version installed on the VM may not allow remote logins.
- **CCM UI Install Agent Failure** – If the install agent process fails for Windows-based VMs when using the **Install Workload Manager Agent** action from the CCM UI. See [Actions Library > Failure Behavior](#) for additional details.
- **VM Running CloudCenter Legacy 4.7.3** – The Management Agent is tested and available for VMs running CloudCenter Legacy 4.7.3 and later versions.

Locating the vhost Value

You must specify the vhost (brokerVirtualHost) value for the region in which the vhost instance exists. To retrieve the vhost value, you must configure a load balancer instance for RabbitMQ, log into the load balancer instance, and retrieve the value using the `/v1/tenants/tenantId/clouds/cloudId/regions?size=0` API. The following screenshot provides a sample response of where this value is located.

```

importRegion: null
lastInstanceSyncTime: null
numUsers: 0
▶ perms: ["administration"]
  regionName: "mygcp-europe-west1"
▼ regionProperties: [{name: "guacamoleHost", value: "35.225.119.255"}
▶ 0: {name: "guacamoleHost", value: "35.225.119.255"}
▶ 1: {name: "guacamoleRoutingKey", value: "mygcp-europe-west1"}
▶ 2: {name: "connectionBrokerHost", value: "35.225.119.255"}
▶ 3: {name: "connectionBrokerEgressPort", value: "7150"}
▶ 4: {name: "RegionEndPoint", value: ""}
▶ 5: {name: "guacamoleServicePort", value: "443"}
▶ 6: {name: "enableProxyForWorker", value: "false"}
▶ 7: {name: "bladeIp", value: "10.128.0.10"}
▶ 8: {name: "WorkerAMQPPort", value: "443"}
▶ 9: {name: "WorkerAMQPIpAddress", value: "35.225.119.255"}
▶ 10: {name: "connectionBrokerIngressPort", value: "7150"}
▶ 11: {name: "brokerVirtualHost", value: "127.0.0.1:5671/vhost"}
▶ 12: {name: "LocalAMQPPort", value: "443"}
▶ 13: {name: "PreferZone", value: "europe-west1-b"}
▶ 14: {name: "enableProxyForCloudEndpoint", value: "false"}
▶ 15: {name: "LocalAMQPIpAddress", value: "35.225.119.255"}
▶ 16: {name: "Region", value: "europe-west1"}
▶ 17: {name: "amqpAccessibleFromCloud", value: "true"}
▶ 18: {name: "lpPort", value: "7150"}
▶ 19: {name: "bladePort", value: "8153"}
▶ 20: {name: "bladeStatusUpdateTime", value: "2017-04-20T14:00:00Z"}
resource: "https://www.googleapis.com/compute/v1/projects/mygcp-europe-west1/zones/europe-west1-a/instances/mygcp-europe-west1-1"

```

The Node ID Input Parameter

The Node ID input parameter is optional for all supported clouds

- If you do not provide the Node ID, the cloud's metadata service retrieves the Node ID.
- If the user-provided Node ID does not match the Node ID generated by the cloud, then Workload Manager raises a WARNING but proceeds to use the user-provided Node ID for the installation.

To provide the Node ID value for Windows, issue the following command:

```

.\install.ps1 -brokerHost <IP> -brokerPort 5671 -cloudFamily vcd --nodeId 421b7140-eda1-da9e-4fa2-370b00d6bf16

```

To provide the Node ID value for Linux, issue the following command:

```

./install --brokerHost <IP> --brokerPort 5671 --cloudFamily vcd --nodeId 421b7140-eda1-da9e-4fa2-370b00d6bf16

```

Download the the Management Agent Bundle

To install the agent bundle, you must first download one of the following bundle store files:

1. SSH into the VM instance designated for this component by using the key pair that you used to launch the VM.



Along with the key pair, you may need to use your login credentials for *sudo* or *root* access based on your environment.

2. Download the following required files for this component from software.cisco.com. Be aware that the following files are contained in a file name that uses the following syntax:

```

cloudcenter-release-<release.tag>-installer-artifacts.tar

```

- agent-lite-linux-bundle.tar.gz (for Linux-based VMs)



Workload Manager installs this file on Linux servers ONLY if the WGET or CURL utilities are available in the image. If either of those utilities are not available, Workload Manager does not install this file. See [CURL and WGET Utilities](#) for additional context.

- agent-lite-windows-bundle.zip (for Windows-based VMs, use RDP access)



The downloaded bundle locations are required to configure the corresponding URL in the Cloud Settings section of the Regions / Details tab.

Install the Management Agent on a Linux VM

- **Syntax:** `http://${s3Bucket}/${agentBundlePath}/bundle/agent-lite-linux-bundle.tar.gz`
- **Example:** <http://build-rel.cliqr.com/release-4.10.0-20180701/bundle/agent-lite-linux-bundle.tar.gz>
- **Structure:** The initial structure of the downloaded Linux agent bundle is:

```
agent\
  bin\
    (scripts and executable...)
  config\
    (config files...)
```

- **Process:**

To install the Management Agent bundle on a Linux VM, follow this procedure.

1. Extract the Management Agent bundle to the `/usr/local` folder.

The `/usr/local/agent` folder is the Agent's home directory `$AGENT_HOME`

2. The `install` command's help (`-h`) option provides help on using this command and provides multiple options to run this command:

```
$AGENT_HOME/bin/install-h
```

- a. Specify environment variables:

Environment Variable Example

```
Set BROKER_HOST, BROKER_PORT and CLOUD_FAMILY env variables and run install
```

- b. Use one of the following options:

- Command line parameters:

Command Line Examples

```
./install -bh 192.168.2.4 -bp 5671 -c vmware

#or
./install -bh 192.168.2.4 -bp 5671 -c vmware -S

#or
./install --brokerHost 192.168.2.4 --brokerPort 5671 --cloudFamily vmware
```

- **Interactive** mode (user is prompted for each required parameters)
- **Silent** mode (`-S` or `--silent`):
 - If silent mode is not specified then installer prompts you for the missing values.
 - If silent mode is specified, then the installation will abort if all mandatory params are not provided.

```
./usr/local/agent/bin/install -bh 12.1.4.1 -bp 5671 -c vmware -S

#or
./usr/local/agent/bin/install -bh 12.1.4.1 -bp 5671 -c vmware --silent

#or
[root@agent]# bin/install --silent -bh 10.6.1.1 -c azurem
Broker port was not specified. Default port 5671 will be used.
Configured Agent
Installing Agent as daemon
Agent Service instantiated
Agent home is /usr/local/agent
Initiating Agent start sequence ...
Agent PID: 27891
Agent start sequence will complete within 40 seconds...
Agent started
```

3. Run the **install** command to install and start the agent in **interactive** mode :

```
$AGENT_HOME/bin/install
```

The agent is started as a background process. – use the following command to verify this process.

```
ps aux | grep agent-lite
```



During the installation process, Workload Manager sets the location for the following files:

- Agent log file: **\$AGENT_HOME/log/agent.log**
- Agent version file: **\$AGENT_HOME/version**
- Config.json file: **\$AGENT_HOME/config/config.json**

If you reboot the VM, either intentionally or unintentionally, then the agent is automatically started as soon as the VM is started.

The config.json file records the following properties:

Sample contents of the config.json file

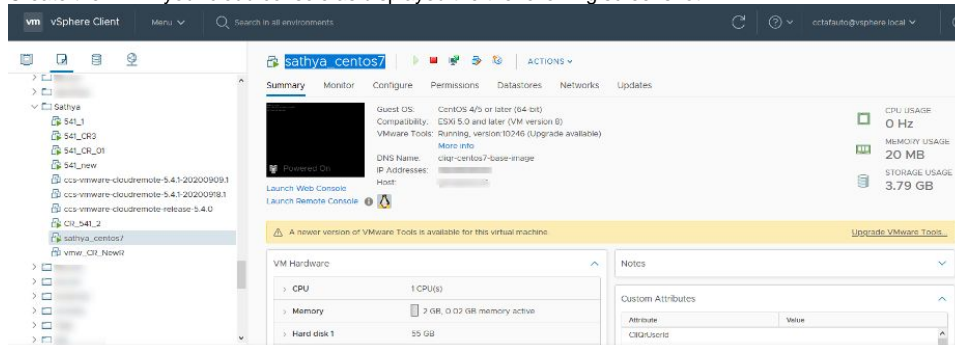
```
{
  "AmqpAddress": "%AMQP_HOST%:%AMQP_PORT%",
  "AmqpUsername": "cliqr",
  "AmqpPassword": "cliqr",
  "AmqpVHost": "/cliqr",
  "AmqpExchange": "cliqr.gateway.exchange",
  "AmqpExchangeType": "direct",
  "AmqpRoutingKey": "cliqr.gateway.key",
  "AmqpRoutingQueue": "cliqr.gateway.queue",
  "CloudFamily": "%CLOUD_FAMILY%",
  "NodeId": "%NODE_ID%",
  "AgentReleaseVersion": "4.10.0",
  "AgentBuildVersion": "1.0.0"
}
```

Installing the Management Agent on **imported VMs** is an alternate option for VMs that do not require the capability to launch applications but do require some basic Workload Manager functionality like performing platform actions.

Workload Manager 5.4.1 provides secure communication when installing the Management Agent on imported VMs. The procedure to use this improved functionality differs from the procedure provided in previous releases. To use this new functionality, you must download the ZIP file which contains the certificates required to communicate with RabbitMQ.

To install the Management Agent on imported VMs and ensure secure communication when communicating with RabbitMQ, follow this procedure.

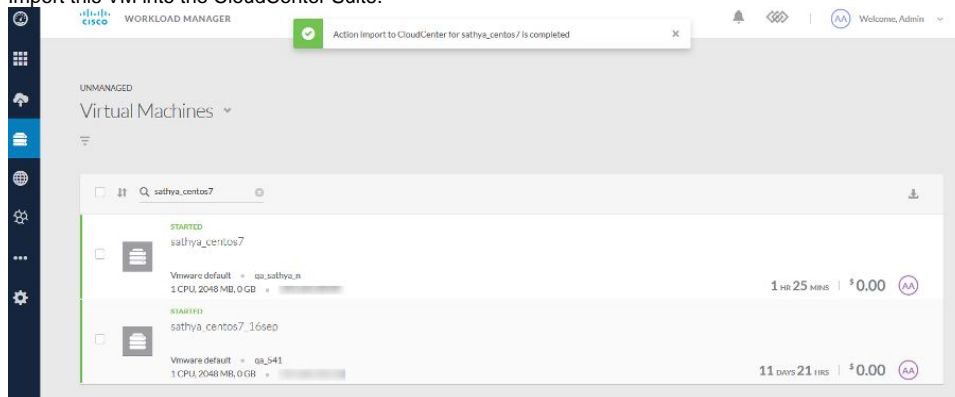
1. Create the VM in your cloud console as displayed in the following screenshot.



2. Navigate to Workload Manager Dashboard > Virtual Machines > Unmanaged VMs, and search for the VM you created in Step 1 above.



3. Import this VM into the CloudCenter Suite.



4. Login to the VM and run the following command.

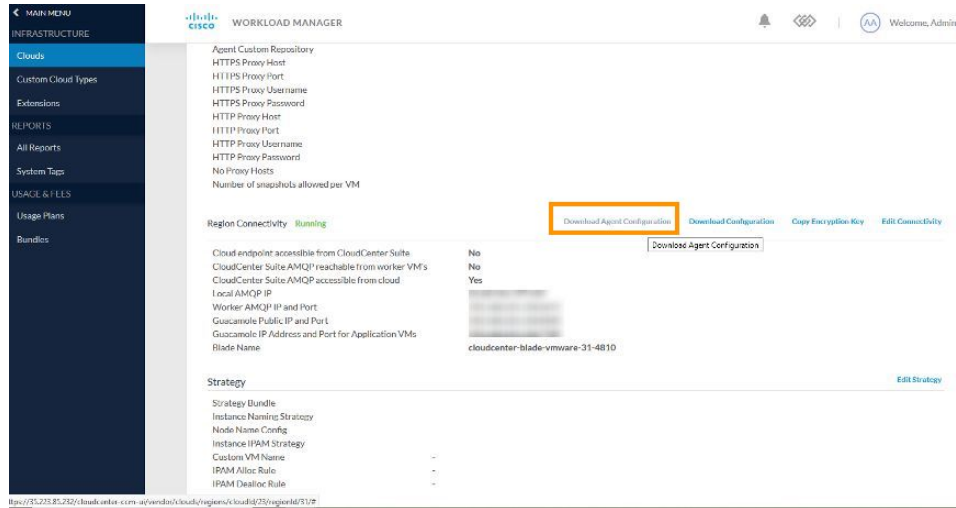
```
wget <bundle file location>
```

```
#For example: wget https://cdn.cliqr.com/release-5.4.1-RC3/bundle/agent-lite-linux-bundle.tar.gz
```

5. Create an archive file of this bundle and extract the file in verbose mode using the following command.

```
tar -C /usr/local -xvf agent-lite-linux-bundle.tar.gz
```

6. Access the Regions page in Workload Manager and download the agentartifacts.zip by clicking the **Download Agent Configuration** link.



7. Copy agentartifacts.zip file to the VM and execute the install file as displayed in the following command.

```
cd /usr/local/agentlite/bin
execute ./install
```

8. Provide the following details as input for each field.

- RabbitMQ broker host: From <https://35.223.85.232/cloudcenter-cloud-setup/api/v1/tenants/1/clouds/23/regions/31> (workerA mqplpAddress: "192.168.131.118")
- RabbitMQ broker port: 5671 for VMware
- RabbitMQ broker VHost: From <https://35.223.85.232/cloudcenter-cloud-setup/api/v1/tenants/1/clouds/23/regions/31> (broker VirtualHost: 31-cloudendpoint986647)
- Cloud Family: VMware (in this example)
- Action: Install
- Agent mode: brownfield
 - The manual install mode refers to the mode when User manually logs in to the VM and then runs the **CLI** to install brownfield agent. Hence the port , host have to be provided manually.
 - The Push mode, which is done via UI, doesn't require user to input anything apart from remote login details.
- RabbitMQ broker auth ZIP file local path: The path for the agentartifacts.zip file



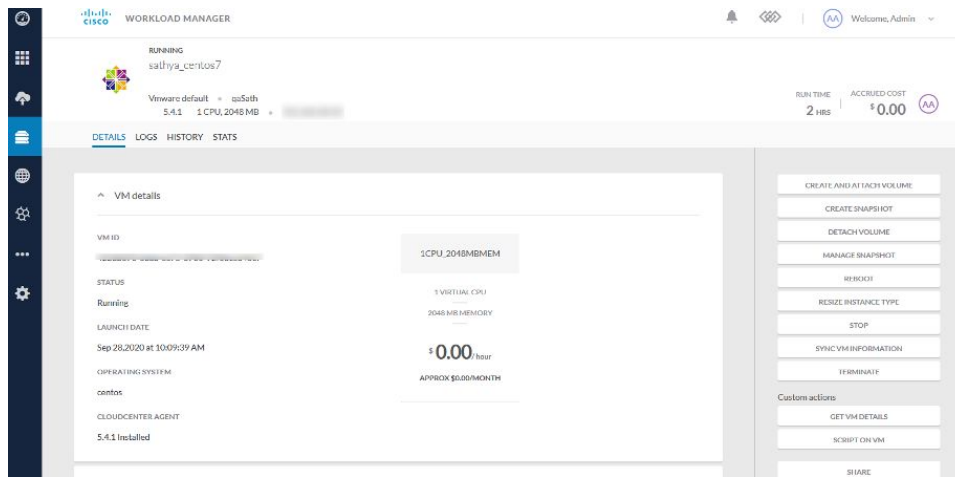
The password is automatically retrieved from the agentartifacts.zip file.

The /usr/local/agentlite/config/config.json allows you to validate the following parameters after the installation process is complete:

- The AMQP Username
- The AMQP Password
- The AMQP VHost
- The AMQP Auth, Certificate, and CA Certificate

Providing those details allows the VM to connect back to the CloudCenter Suite via the RabbitMQ server.

9. Once you save the edited file, the OS flavor and agent version will be displayed in the UI after the VM is registered with the CloudCenter Suite.



Install the Management Agent on a Windows VM

- **Syntax:** `http://${s3Bucket}/${agentBundlePath}/bundle/agent-lite-windows-bundle.zip`
- **Example:** <http://build-rel.ciqr.com/release-4.10.0-20180701/bundle/agent-lite-windows-bundle.zip>
- **Structure:** The initial structure of the Windows agent bundle is:

```
agent\
  bin\
    (scripts and executables...)
  config\
    (config files...)
  utils\
    (utils..)
```

Prerequisites:

- Have a Windows instance enabled for RDP.
- Enable the SMB2 protocol on the Windows server where you are installing the Management Agent. See <http://docs.microsoft.com/en-us/windows-server/storage/file-server/troubleshoot/detect-enable-and-disable-smbv1-v2-v3> for details on how to enable the SMB2 protocol.
- Open Ports 139 and 445 on Windows for the **Administrator** user.

```
New-NetFirewallRule -DisplayName "Winxe-port-139" -Direction Inbound -LocalPort 139 -Protocol TCP
New-NetFirewallRule -DisplayName "Winxe-port-445" -Direction Inbound -LocalPort 445 -Protocol TCP
```

- **Process:** The process to install the Management Agent on a Windows instance differs based on the Windows version.

- To install the Management Agent for Windows 2012, follow this procedure.

1. Extract the agent bundle to **<SystemDrive>:\opt** (example: **C:\opt**).

The **<SystemDrive>:\opt\agent** folder is the Agent's home directory **\$AGENT_HOME**

2. The **install.ps1** command's help (**-detailed**) option provides help on using this command and provides multiple options to run this command:

```
get-help C:\opt\agent\bin\install.ps1 -detailed
```

- a. Use command line parameters:

Command Line Example

```
C:\PS>.\install.ps1 -brokerHost 192.168.2.4 -brokerPort 5671 -cloudFamily vmware
```


Cloud Nuances when Installing AgentLite

Be aware of the following cloud nuances when you provide the cloudFamily name while installing the Management Agent :

Node cloudFamily	Cloud that the VM is Hosted	Behavioral Description
<ul style="list-style-type: none"> VMware vCenter AzureRM 	Irrespective of the node being hosted on AWS, OpenStack, Azure, or VMware vCenter	The Node ID retrieval logic uses the OS command and does not result in a cloudFamily error.
<ul style="list-style-type: none"> OpenStack AWS (Amazon) 	OpenStack or AWS respectively	The cloud provider APIs are used to retrieve the node information and these two cloud providers detect if the node exists – if the node does not exist in this cloud, it results in an error.

- b. Use **interactive** mode (user is prompted for each required parameters).
 3. Run the **install.ps1** command to install and start the agent in **interactive** mode:

```
powershell.exe -ExecutionPolicy Bypass -NoProfile -File C:\opt\agent\bin\install.ps1
```

The agent is started Agent as a windows service called **AgentService**. You can verify it by opening Windows service manager and viewing the list of running services.



During the installation process, Workload Manager sets the location for the following files:

- Agent log file: **\$AGENT_HOME/log/agent.log**
- Agent version file: **\$AGENT_HOME/version**
- Config.json file:
\$AGENT_HOME/config/config.json

If you reboot the VM, either intentionally or unintentionally, then the agent is automatically started as soon as the VM is started.

The config.json file records the following properties:

Sample contents of the config.json file

```
{
  "AmqpAddress": "%AMQP_HOST%:%AMQP_PORT%",
  "AmqpUsername": "cliqr",
  "AmqpPassword": "cliqr",
  "AmqpVHost": "/cliqr",
  "AmqpExchange": "cliqr.gateway.exchange",
  "AmqpExchangeType": "direct",
  "AmqpRoutingKey": "cliqr.gateway.key",
  "AmqpRoutingQueue": "cliqr.gateway.queue",
  "CloudFamily": "%CLOUD_FAMILY%",
  "NodeId": "%NODE_ID%",
  "AgentReleaseVersion": "4.10.0",
  "AgentBuildVersion": "1.0.0"
}
```

You can set the environment variables to provide easy access to logs and configuration files.

1. Source the environment vars files:

```
source /usr/local/agent/bin/vars.sh
```

2. Use the following variables to set the values:

```
AGENT_HOME : /usr/local/agent
CONFIG_FILE : /usr/local/agent/config/config.json
LOG_FILE : /usr/local/agent/log/agent.log
LOG_LEVEL : DEBUG
```

Managing the agent file and feature is specific to the OS in use. The following table provides information about managing the agent .

Management Task	Linux	Windows
Agent Registration by Workload Manager	Registered as a daemon program agentd .	Registered as a service, called AgentService .
Start the Management Agent	Use one of the following commands: <pre>/etc/init.d/agentd start #or service agentd start #or /bin/bash /usr/local/agent/bin/agent-start.sh</pre>	Use the service manager to start the service.
Stop the Management Agent	Use one of the following commands: <pre>/etc/init.d/agentd stop #or service agentd stop #or /bin/bash /usr/local/agent/bin/agent-stop.sh</pre>	Use the service manager to stop the service.
Restart the Management Agent	Use one of the following commands: <pre>/etc/init.d/agentd restart #or service agentd restart</pre>	Use the service manager to restart the service.
Modify the configuration	To modify the configuration, follow this process: <ol style="list-style-type: none"> 1. Modify /usr/local/agent/config/config.json as required. 2. Restart the agentd service. 	To modify the configuration, follow this process: <ol style="list-style-type: none"> 1. Modify C:\opt\agent\config\config.json as required. 2. Stop and start AgentService using service manager.

By uninstalling the agent, you are only removing the agent daemon from the Linux server and the agent service from the Windows server – you are not removing the folder/directory.

To uninstall an existing agent instance on a VM and install a new agent version, follow this procedure.

1. Uninstall the agent file using the following OS-specific command:

Linux

```
/usr/local/agent/bin/uninstall.sh
```

Windows

```
powershell.exe -ExecutionPolicy Bypass -NoProfile -File C:\opt\agent\bin\uninstall.ps1
```

2. Delete the agent Home folder/directory

Example:

```
Linux: rm -rf /usr/local/agent
Windows: rm -r C:\opt\agent
```

3. Install the Management Agent using one of the options provided in the sections above (from the UI Actions dropdown or Manually).



You can only upgrade agent for a Workload Manager VM if the agent was previously installed and if the VM is using CloudCenter Legacy 4.7.3 or later versions.

If the latest version of the agent is already installed on a VM, then the **Upgrade Workload Manager Agent** action will no longer be available for this VM.

You can upgrade the agent either from the UI or the API.

Before you can upgrade the agent on an imported Linux VM with the version 5.0.0 agent installed, you must perform these pre-upgrade steps:

1. Create a new on demand action to prepare the VM before upgrade (see [Actions Library](#) for more details on creating a custom on demand action):
 - a. Navigate to the Actions Library tab and click **New Action**.
 - b. Give the new action a memorable **Action Name**, eg, "PreUpgrade".
 - c. Add a resource mapping corresponding to all managed imported Linux VMs and click Done. See the screenshot below for guidance.

- d. Copy and paste the following string in the **Executable Command** field:

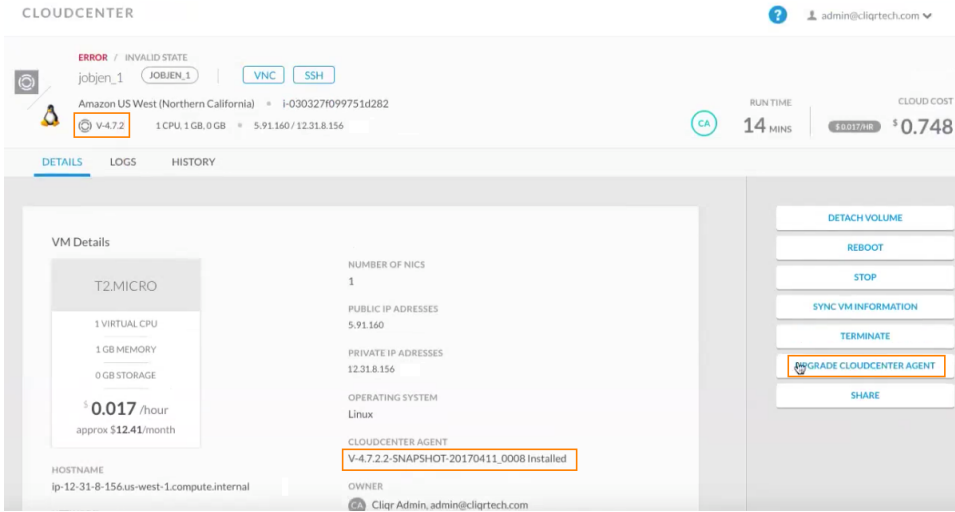
```
sed -i.bak -e ':a;N;$!ba;s/}\nfunction jsonval {/function jsonval {/g' /usr/local/agent/bin/agent-pre-upgrade.sh
```

- e. Click **Done** to save the new on demand action.
2. Execute the on demand action you created on the VM needing the agent upgrade:
 - a. Navigate to the VM details page of the VM you want to upgrade.
 - b. Click the action button on the right side of the screen corresponding to the new "PreUpgrade" action.

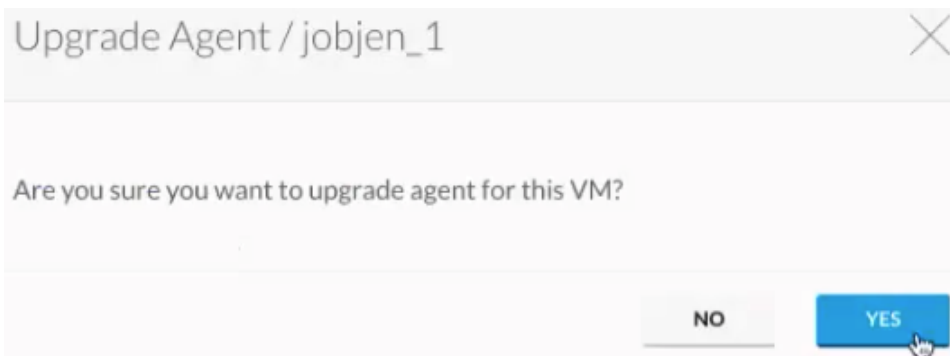
You can now proceed to the steps below to upgrade the agent on the VM.

To upgrade the agent on a Workload Manager VM, follow this procedure:


1. Identify the VM that requires the agent to be installed and access the details page or the list page for this VM. The following screenshot displays a VM deployed with CloudCenter Legacy 4.7.2 agent installed on it.



2. In the Actions panel, select **Upgrade Workload Manager Agent**. The Upgrade Agent page displays for this VM, as shown in the following screenshot.



3. Click **Yes**. The upgrade process begins as displayed in the following status screenshot.



If the upgrade does not complete successfully, the agent reverts to the previous Workload Manager version and provides reasons for the failure in the History tab.

4. Once upgraded, the status screenshot reflects the status of the upgrade. Verify that the version information displays the upgraded details.

You can perform VM operations such as start, stop, and reboot VMs from the [Deployment Details Page](#) or the Virtual Machine list page. Admins can manage Workload Manager VMs and take several actions from the Virtual Machines list page or a VM Details page. See [Actions Library](#) for additional details.

If the [Terminate Protection](#) feature is enabled, you will not be able to perform any stop or reboot actions on the corresponding VM(s). These operations are only permitted when the N-tier jobs/VMs/deployments are in the running state. See [Deployment, VM, and Container States](#) for additional context.


A new billing process calculates the run time and cost usage of imported VMs. This process is similar to existing billing process and runs once an hour.

The Workload Manager costs are accrued for each child job as well as for the total cost of the deployment.

A Workload Manager-created or dynamically bootstrapped worker image has the NTP daemon automatically started and the timezone set to UTC.

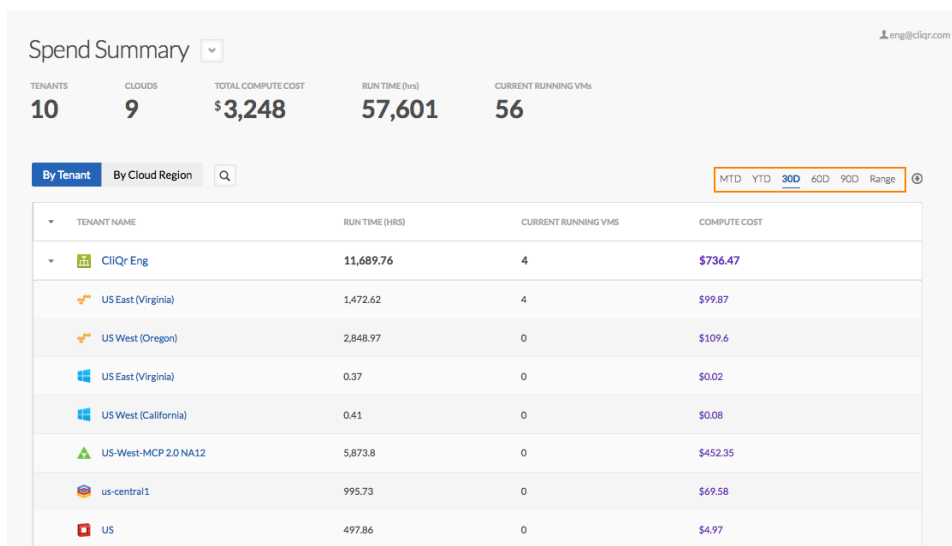
The various timestamps generated for startTime, importedTime, terminatedTime generated and stored on behalf of any VM in Workload Manager are based on the Workload Manager timestamps.

- **Filter Favorites:** Any time you mark a deployment as a favorite (see [Deployment Details Page](#) > *Favorite Deployments*) by clicking the star icon, you can also view a filtered list of favorite deployments, as shown in the following screenshot.



- **Search Strings:** Specify strings in the Search icon field based on strings that identify the following resources:
 - High-level status (status)
 - Public IP address (publicIpAddr)
 - Private IP address (privateIpAddr)
 - Node name (nodeId)
 - Hostname (hostName)
 - Cloud family (cloudFamily)
 - Region name (regionName)
 - Cloud account name (cloudAccountName)
 - Cloud group name (cloudGroupName)
 - Parent deployment name (name)
 - VM name (displayName)
- The time period filtering options displayed in the top right corner of the Virtual Machines tab enables you to filter VMs based on the VM **Run Time** (default).

The time period filter is illustrated in the following screenshot.



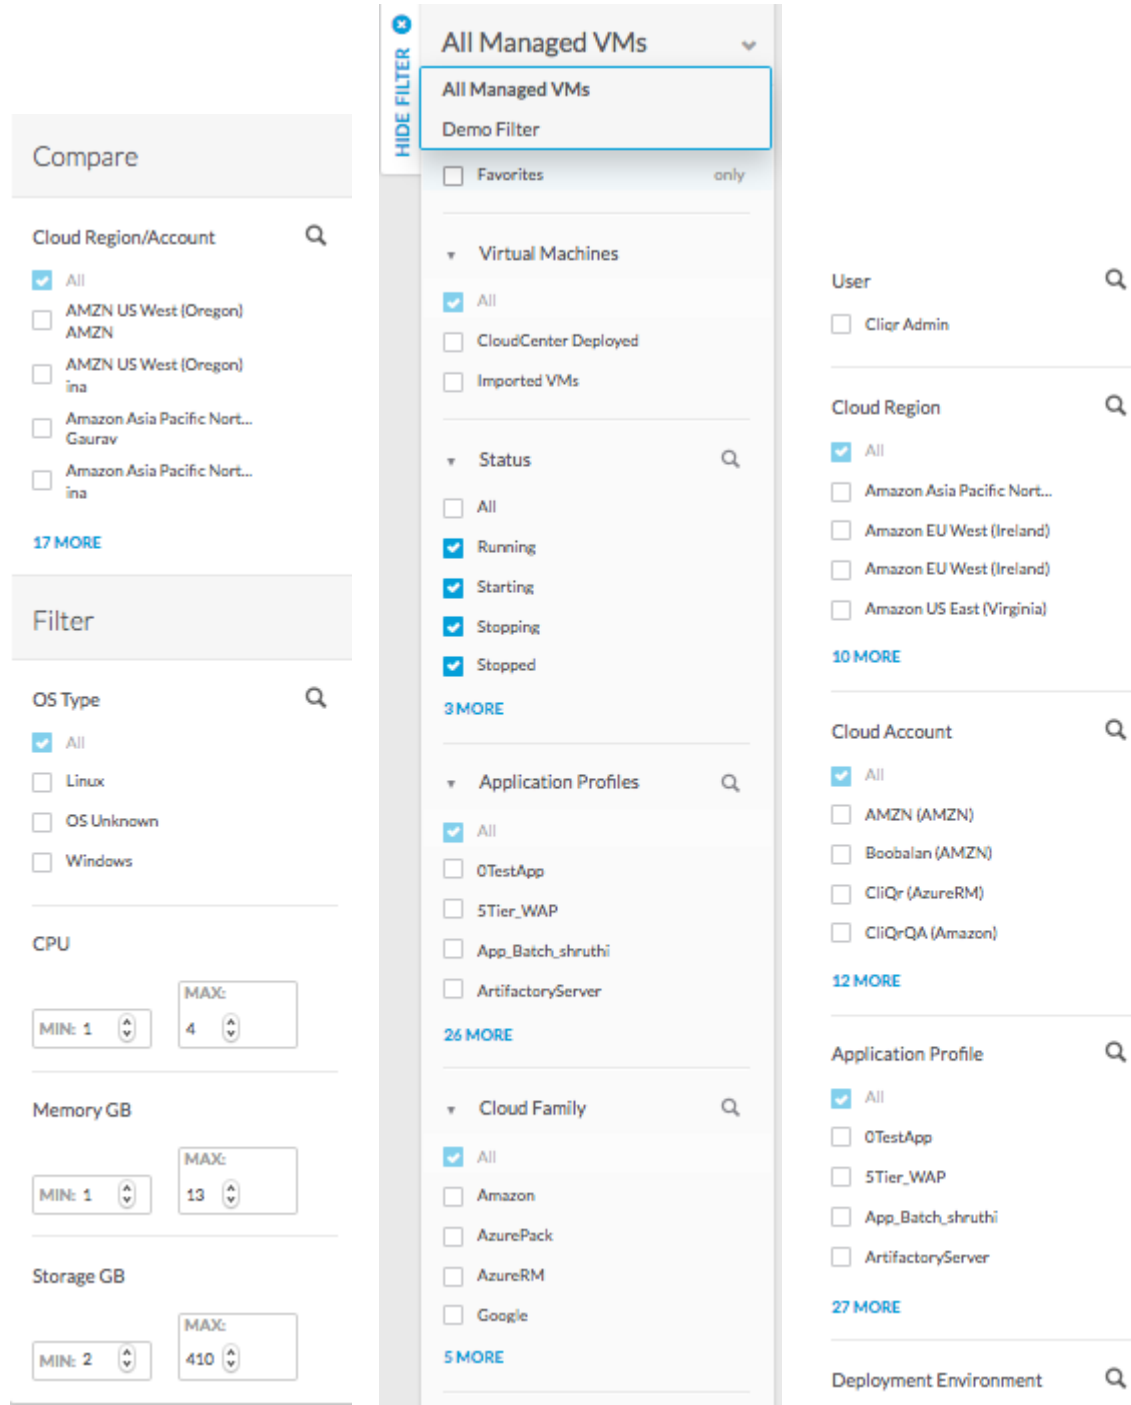
The Time Period filter option is only available for some pages (for example, the [Usage Summary Report](#) or the [Virtual Machine Management](#) page). If available, the filter options are displayed in the top right corner. The available time period filter options are explained in the following table:

Time Period Filter	Description	Notes
MTD	Month to Date	The current month
YTD	Year to date	The current year
30D (Default)	30 Days	The current 30 days ending with today <div style="border: 1px solid green; padding: 5px; margin-top: 10px;"> The data that is displayed in response to a 30-Day time period request only displays data from the 1st of the month, not for the previous 30 days. To work around this issue, use the date Range option and provide the begin and end date for the required period. </div>
60D	60 Days	The current 60 days ending with today
90D	90 Days	The current 90 days ending with today

Range	A custom range specified by the selected month and year	If using APIs, this is the only available options to display reports for a period of time based on the <i>startDate</i> and <i>endDate</i> attributes
--------------	---	---

- The **Hide Filter/Show Filter** option enables you to hide or expand advanced filtering options. You can save custom filters just as you would for Workload Manager reports

This advanced filtering options helps you directly add short cuts to filtered lists that you can quickly access at a later time. This feature is available for some pages (for example, the [Running VM History Report](#) or the [Virtual Machine Management](#) page). The following screenshots display some of the available filters.



Users Q

- All
- User 01 CloudCenter
- Vik Pary

Groups Q

- All
- G1
- G2

Cloud Region Q

- Amazon US East (Vir...

Cloud Account Q

- Vik AWS cloud ...

Application Profile Q

- All
- Jenkins
- dummyExternalServi...

Deployment Environment Q

- AWS only

Status Q

- All
- Canceled
- Running/Deployed

Tags Q

No Tags available

CCID Q

No CCID available

Project Q

No Project available

i You can additionally filter CloudCenter resources using the user-based **Groups** filter (see the highlighted image above).

User **Groups** displayed in the filter lists all user groups that are configured for your tenant. Selecting any user group list filters the list of application deployments for the selected group and keeps those deployments selected in this list if they map to users in the user group.


You can also combine the user and user group filters, and in this case, the report displays deployments that map to either the selected user or any user who is a member of the selected user group.

Save Filters

By saving a filter, you are directly adding short cuts to custom filtered lists that you can quickly access at a later time.

To save a custom filter, follow this procedure.

1. Select the required filters in the Filters pane and/or the Columns filter choices.
2. Click **Save**, located right above the Filters pane, as displayed in the following screenshot.



The Save Filter popup displays.

3. Enter a name for this filter and click **Save**.

Save Filters as

All fields are required unless otherwise specified

Name


Non-Running VMs

Selected Filters


Virtual Machines: CloudCenter Deployed, Imported VMs
Status: Starting, Stopping, Stopped, Error

Save
Cancel

4. The filter is saved and a status message displays in the page.




5. You can access and view the saved filters from the dropdown list.




Delete Saved Filters

You can delete saved filters by clicking the Trash icon next to the saved filter live link.



The Delete Saved Filters popup confirms your intention before deleting the saved filter and displaying the status message at the [Application Deployments Report](#) page.



Deployments List Page

Deployments List Page

The deployments list page is where you can see all of the deployments created by you or shared with you. You can access this page by clicking the Deployments tab from the main menu. A screenshot of a sample Deployments List page is shown below.

The screenshot shows the 'Deployments' page with the following data:

NAME	STATUS	ENVIRONMENT	START TIME	BILLABLE TIME	APPROX COST	APPROX SAVINGS	ACCRUED COST	ACTIONS
cont <small>cont (V1)</small> <small>Kubernetes default</small>	Starting	Kubernetes	Jul 29, 2019 at 07:00 PM	Not Available	\$0.00/hr \$0.00/mo		\$0	
centos <small>CentOS (V7)</small> <small>Amazon_API US West (N. Califor...</small>	Starting	user01_1	Jul 29, 2019 at 06:58 PM	Not Available	\$0.00/hr \$0.00/mo		\$0	
wcont <small>cont (V1)</small> <small>Kubernetes default</small>	Deployed	Kubernetes	Jul 27, 2019 at 08:02 PM (UTC)	Not Available	\$0.00/hr \$0.00/mo		\$0	

At the bottom right of the table, there is a pagination control: 'Show 50 per page Page 1 of 1'.

From this page you can:

- Deploy an application by clicking on the New Deployment button in the upper right. This will redirect you to the Application Profiles page.
- Go to the [Deployment Details page](#) for a deployment by clicking on the line where the deployment is listed.
- Perform an on-demand action on the deployment by hovering over the actions column, clicking the dropdown action field, and selecting an action. The list of available actions will depend on the deployment type and its [state](#), and whether any [Projects](#) were created by you or shared with you.

Deployment Details Page

Deployment Details Page

- [Overview](#)
- [Accessing Deployment Details](#)
- [The Header](#)
- [The Main Panel](#)
- [The Tiers Tab](#)
- [VM-Based Tiers](#)
 - [The VMs Tab](#)
 - [The VM Details Tab](#)
- [Container-Based Tiers](#)
 - [The Replicas Tab](#)
 - [The Container Details Tab](#)
- [The Details Tab](#)
- [The History Tab](#)
- [The Topology Panel](#)
- [Notifications](#)

The **Details** page provides deployment details for each deployment that you submit when [Deploy an Application](#).

Once deployed, you can access the *Deployment Details* page in one of two ways:

- From the **Deployments** list page, by clicking on the deployment *name* as highlighted in the following screenshot:


NAME	STATUS	ENVIRONMENT	START TIME	RUN TIME	APPROX COST	APPROX SAVINGS	ACCRUED COST	ACTIONS
si placementgroupable(N1) Kubernetes default	Error	kub	Sep 13, 2018 at 04:21 PM		\$0.00/hr \$0.00/mo		\$0.00	
capache cap1-capache(N1) Kubernetes default	Deployed	kub	Sep 13, 2018 at 09:50 PM(UTC)		\$0.00/hr \$0.00/mo		\$0.00	
wocomt_kub wocomt(N1) Kubernetes default	Terminated	kub	Sep 13, 2018 at 09:09 PM(UTC)		\$0.00/hr \$0.00/mo		\$0.00	
pub placementgroupable(N1) Kubernetes default	Deployed	kub	Sep 13, 2018 at 08:45 PM(UTC)		\$0.00/hr \$0.00/mo		\$0.00	
complextopo CCQA_AWS_US_East(Virginia)	Deployed	CCQA_deploy_env	Sep 13, 2018 at 08:41 PM(UTC)	9 hrs 13 mins	\$0.0125/hr \$9.02/mo	\$24.39/mo	\$0.11	-Actions-
ib ib(N1) CCQA_AWS_US_East(Virginia)	Deployed	CCQA_deploy_env	Sep 13, 2018 at 06:59 PM(UTC)	4 hrs	\$0.0031/hr \$2.26/mo	\$6.10/mo	\$0.05	

- From the **Virtual Machines** list page, by clicking on the *deployment button* as highlighted in the following screenshot:

TOTAL VMs	RUNNING VMs	CLOUD COST	EST MONTHLY COST	VM HOURS
141	55	\$ 521	\$ 808	16,795

SHOW FILTER	SEARCH	MTD	YTD	30D	60D	90D	RANGE	##	↓
<div style="display: flex; align-items: center;"> <div style="margin-right: 10px;"> RUNNING </div> <div> COMPLEXTOPO </div> </div>	i-0d72a0b1ce8ab54d	Amazon US East (Virginia)	4.10.0	1 CPU, 1 GB, 0 GB	35.168.17.97 / 172.31.31.2	3 hrs	\$0.03		
<div style="display: flex; align-items: center;"> <div style="margin-right: 10px;"> RUNNING </div> <div> COMPLEXTOPO </div> </div>	i-08ecc1a07722c7c13	Amazon US East (Virginia)	4.10.0	1 CPU, 1 GB, 0 GB	35.172.117.118 / 172.31.28.50	3 hrs	\$0.03		

- Either way, you see the **Details** page as shown in the following screenshot:

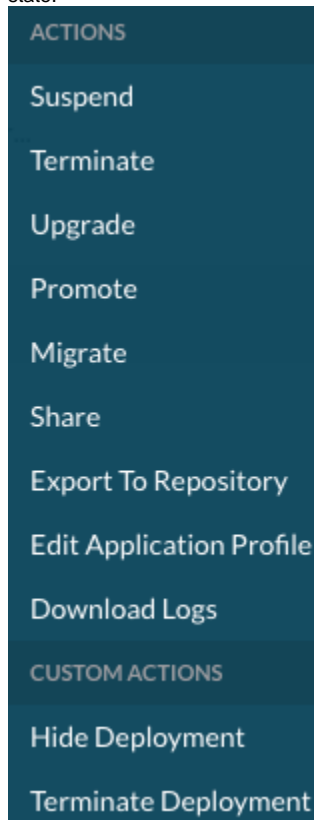


The Details page has three sections:

- The Header
- The Topology Panel (right)
- The Main Panel (left)

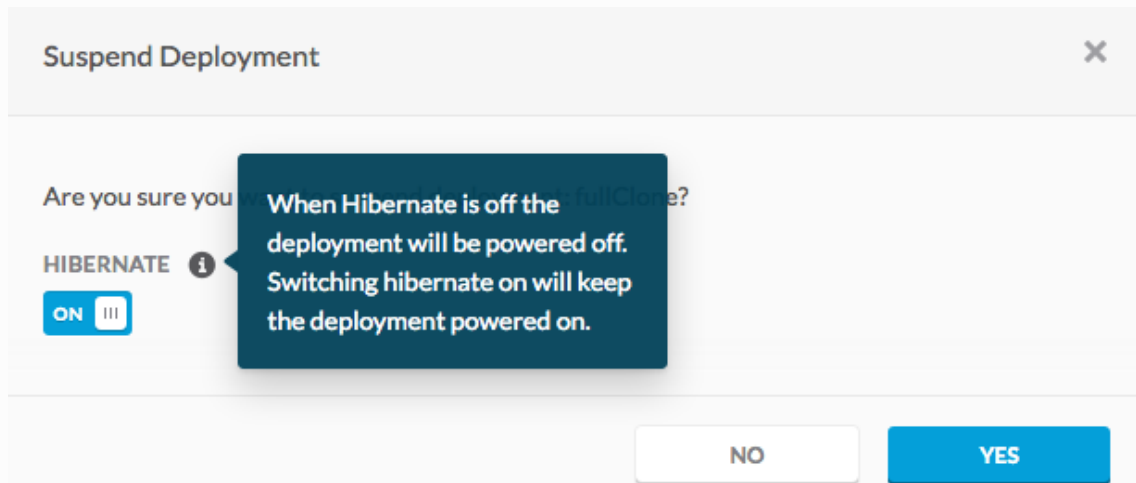
The App logo, state (stopped, running, suspended, etc.), favorites star, deployment name, access link (when applicable), and the job actions dropdown menu are included in **the top part of this section**.

Click on the job actions dropdown icon to display a menu as visible in the following screenshot – the list of available actions vary based on the current job state.



i The **Suspend** action is not available for VMs, pure container, or hybrid applications.

The **Hibernate** action (only available for VMware deployments) is available from the **Suspend** action as a toggle (default = OFF) the switch as displayed in the following screenshot.

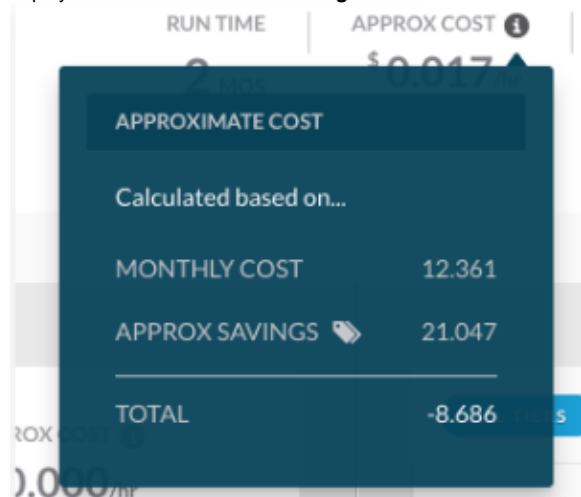


For applications where all tiers are deployed to a vCenter cloud region, an option is added to the Workload Manager deployment suspend function that hibernates the associated VMs instead of causing a full power off / shut down. When such a deployment is later resumed from within Workload Manager, the vCenter Power On function is applied to the associated VMs. This causes the VMs to resume from hibernation instead of being booted up from a cold start. See [Deployment, VM, and Container States](#) for additional details.

The application name, environment name, and cloud region are listed in **the lower left section** and the truncated names expand on hover as displayed in the following screenshot:



The run time, approx cost (hourly compute charges, projected monthly cost factoring in suspension policy savings, and on hover, the following details as well: projected monthly cost before savings from suspension policy and monthly savings - see figure below), accrued cost (since the job started), deployment owner are listed in **the right side of this section**. The following screenshot displays the cost factoring details.



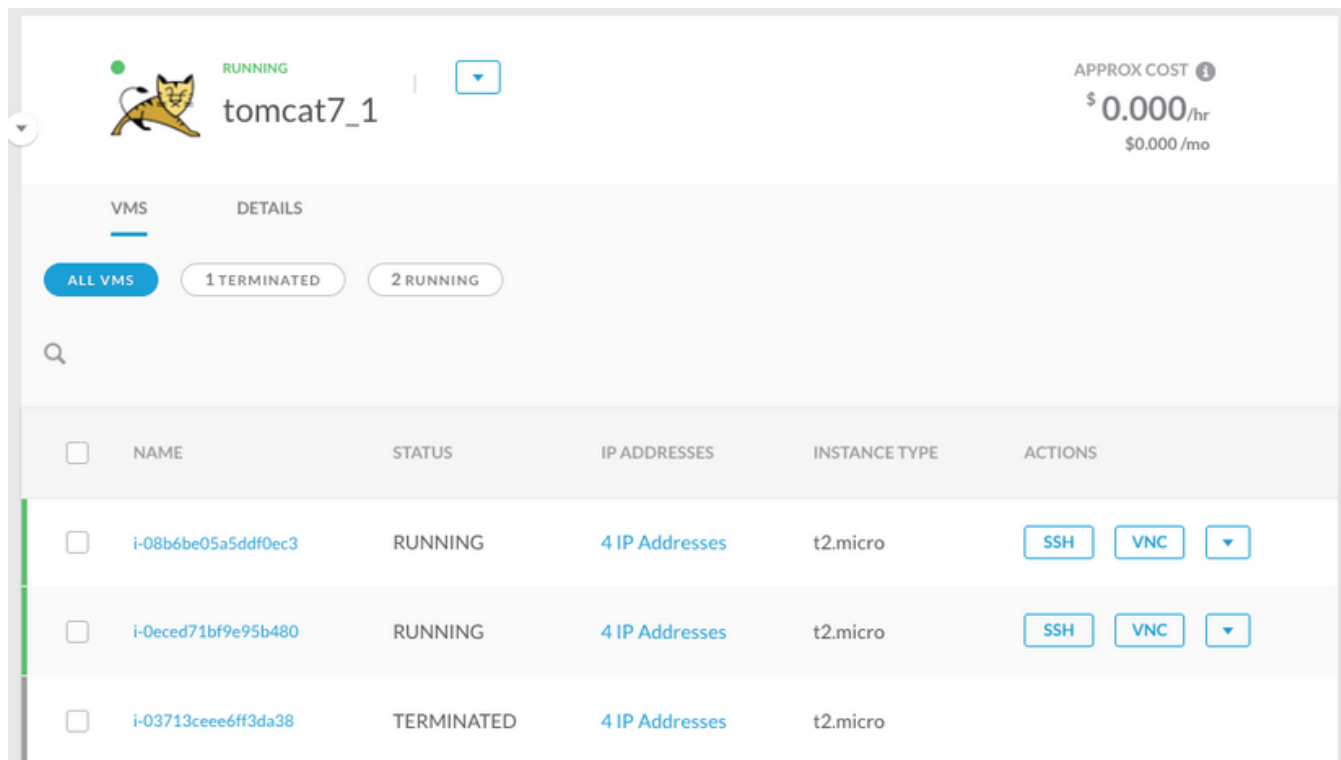
The Main panel has the following subsections:

- The Tiers Tab
- The Job Details Tab
- The History Tab

The Tiers Tab contains a list of per-tier sub-panels:

- Up to 5 tiers display initially in the expanded mode (showing the tier header and tier main section). If more tiers exist a **Load more** button allows you to load 10 more tiers at a time. You can compress/expand each tier by clicking the triangular icon.
- The displayed information differs for VM-based tiers and container-based tiers.

The following screenshot displays a VM-based tier.



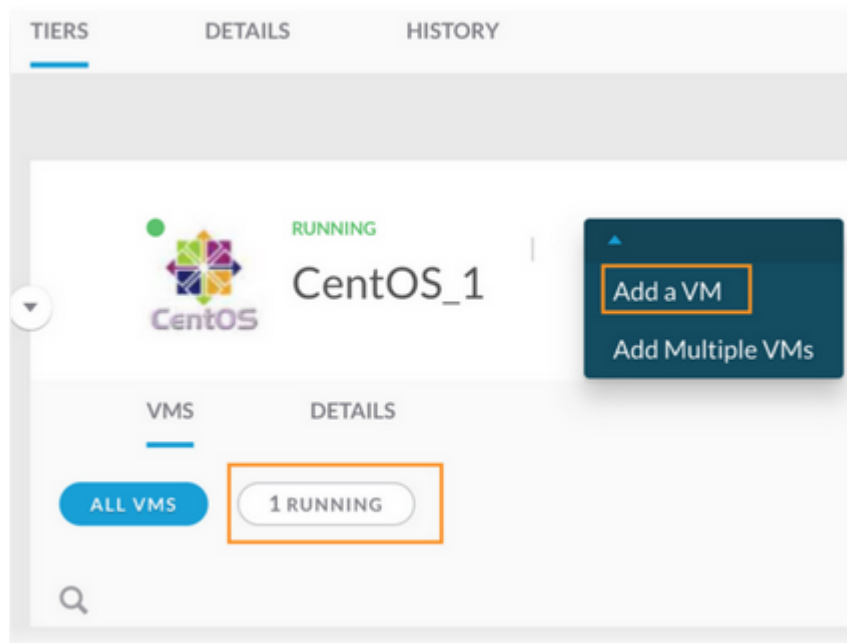
- **Tier Header:** Service logo, tier state (running, starting, etc.), tier name, approx cost (hourly compute charges, projected monthly cost factoring in suspension policy savings, on hover: projected monthly cost before savings and suspension policy).
 - Bulk actions are supported on VMs – you can select multiple VMs and perform the action from a menu that appears on the top right corner of the table.
 - Search VMs using keywords.
 - Load More: Initially 5 VMs are loaded. You can click **Load More** to load 5 more VMs or the remaining VMs if this number is less than 5.
 - The scaling actions dropdown is only present for tiers that scale when the tier is in the running state.
 - Click the dropdown icon to the right of the tier state to display and select from the available choices for each tier as described in the following table:

i Both batch actions or single actions are available to terminate VMs.

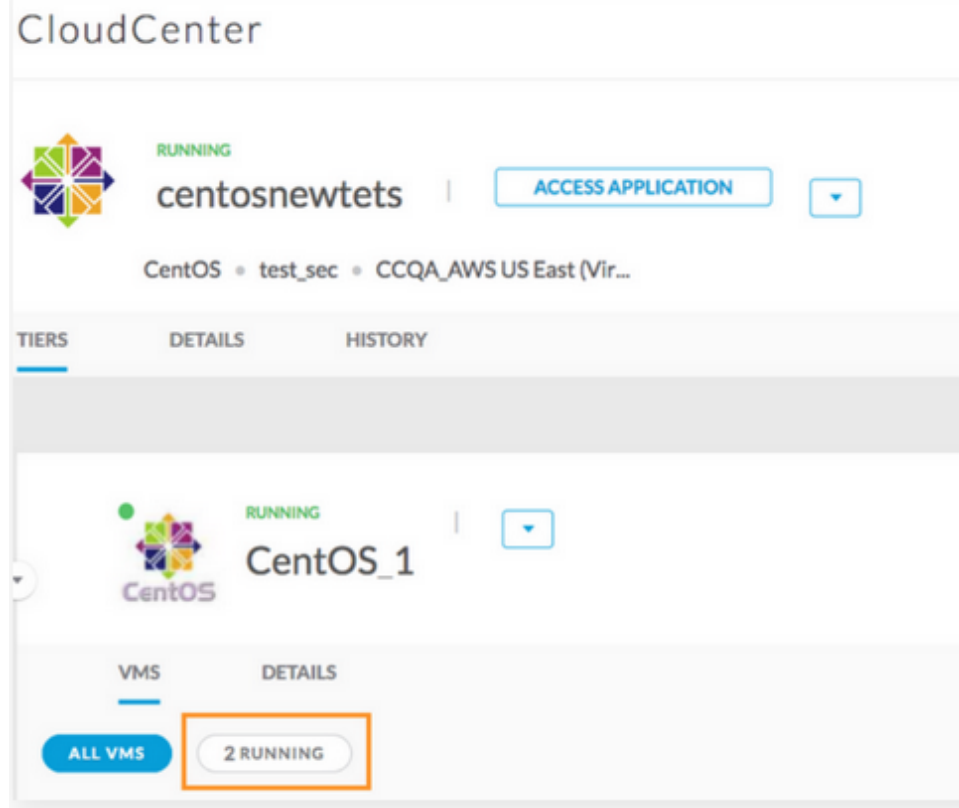
Scaling Action	Description
----------------	-------------

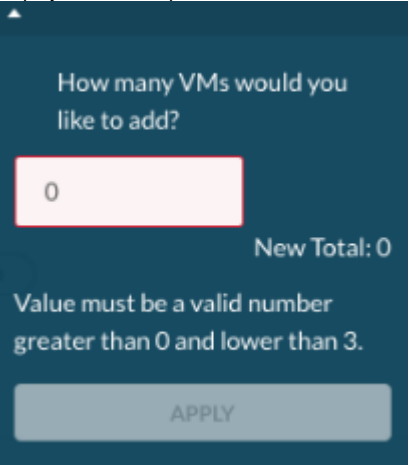
Add a VM

If you click **Add a VM** the system automatically adds one additional VM without any further dialog. The alternate choice is to add multiple VMs as explained in the next row in this table. The following screenshot displays the scaling action location and only displays 1 VM running.



The following screenshot displays the second VM added after clicking the **Add a VM** option. The status message indicates the VM being added and the running VM count increases by one after the VM is successfully added.

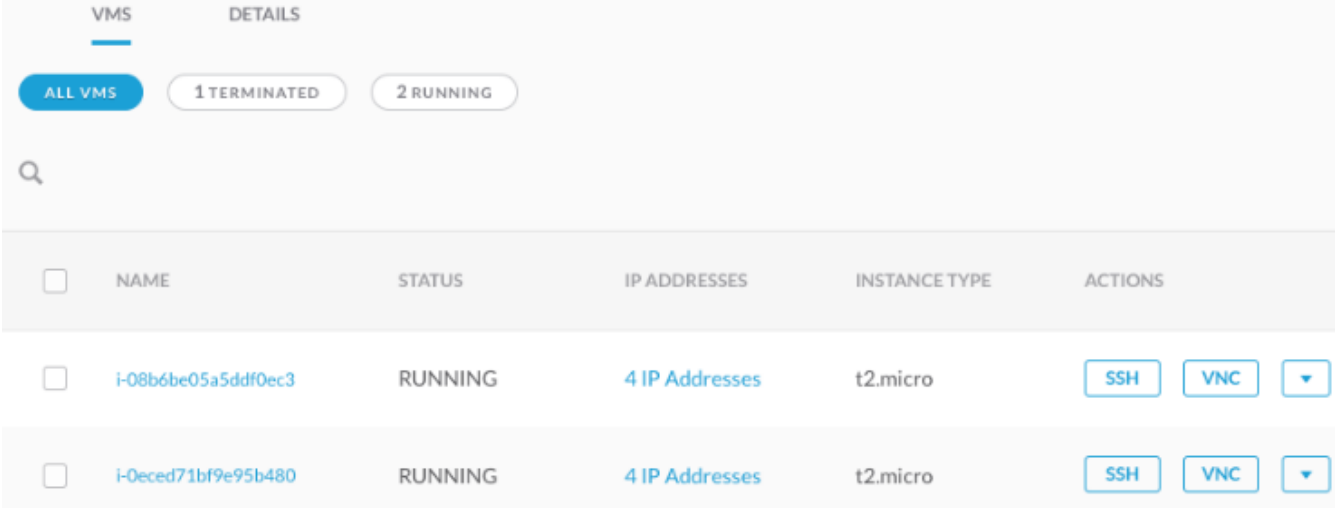


Add Multiple VMs	<p>The dialog box displays a range of VMs that you can add. Enter value and click Apply. The following screenshot displays the subsequent screenshot.</p>  <p>The maximum number of VMs allowed for scale up is based on the <i>maxNodes</i> set for a tier when you model the application.</p>
Remove Multiple VMs	<p>The dialog box displays a range of VMs that you can remove. Enter the value and click Apply.</p> <p>The minimum number of VMs allowed for scale down is based on <i>minNodes</i> set for a tier when you model the application.</p>

- **Tier Main Section:** For each tier, the main section contains tabs for **VMs** and **Details**.

The VMs Tab


The VMs page is divided into sections with self explanatory filters and status details.



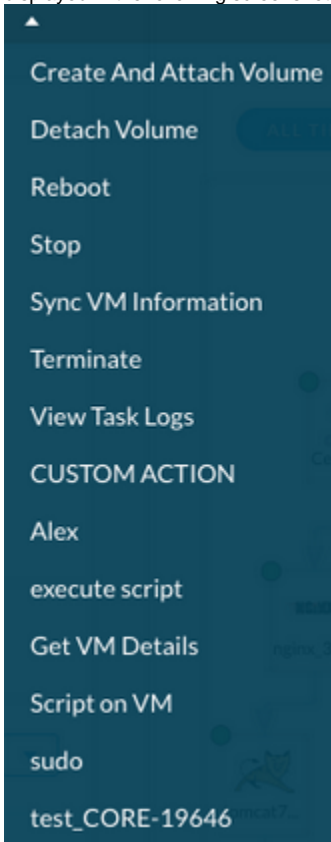
<input type="checkbox"/>	NAME	STATUS	IP ADDRESSES	INSTANCE TYPE	ACTIONS
<input type="checkbox"/>	i-08b6be05a5ddf0ec3	RUNNING	4 IP Addresses	t2.micro	SSH VNC ▾
<input type="checkbox"/>	i-0eced71bf9e95b480	RUNNING	4 IP Addresses	t2.micro	SSH VNC ▾
<input type="checkbox"/>	i-03713ceee6ff3da38	TERMINATED	4 IP Addresses	t2.micro	

- The Status buttons display the VM count in a certain state. The buttons are displayed based on the state of VMs (running, stopped, terminated, error, and so forth) in the tier.


- The list of VMs details in a tier along with additional information when you hover over most fields. For example, hovering over the IP addresses entry displays a balloon with specific IP addresses as displayed in the following screenshot.



- Click the VM name to open a new tab showing VM details.
- Click the dropdown icon next to the VNC link to display the VM actions dropdown menu. The list of actions depends on the current VM state as displayed in the following screenshot.



- Click the **View Task Logs** to view a pop up window with the task log messages for that VM in reverse chronological order as as displayed in the following screenshot.



- You can select all or some of the VMs in the list using check boxes to the left of each VM entry.

- Click on at least one checkbox to view the dropdown menu icon to appear above the list of VMs and to the right as shown in the following screenshot.

The screenshot shows a table with the following columns: NAME, STATUS, IP ADDRESSES, INSTANCE TYPE, and ACTIONS. A red vertical bar on the left indicates that one VM is selected. The selected VM is 'i-08ebc0948a711bd05' with a status of 'ERROR'. A dropdown arrow icon is visible in the ACTIONS column for this VM. The top right of the table shows '1 SELECTED'.

NAME	STATUS	IP ADDRESSES	INSTANCE TYPE	ACTIONS
<input type="checkbox"/> i-00f38c1b0351ab85	TERMINATED	4 IP Addresses	t2.micro	
<input checked="" type="checkbox"/> i-08ebc0948a711bd05	ERROR	4 IP Addresses	t2.micro	▼
<input type="checkbox"/> i-02863c362fd622cd6	TERMINATED	4 IP Addresses	t2.micro	

- Click the dropdown icon to display the VM actions menu applicable to the least common denominator of all VMs selected. The VM action menu corresponding to the one selected error state VM is shown in the following screenshot.

The screenshot shows the same table as above, but with the dropdown menu open for the selected VM. The menu contains the following options: Reboot, Sync VM Information, Terminate, CUSTOM ACTION, execute script, and sudo.

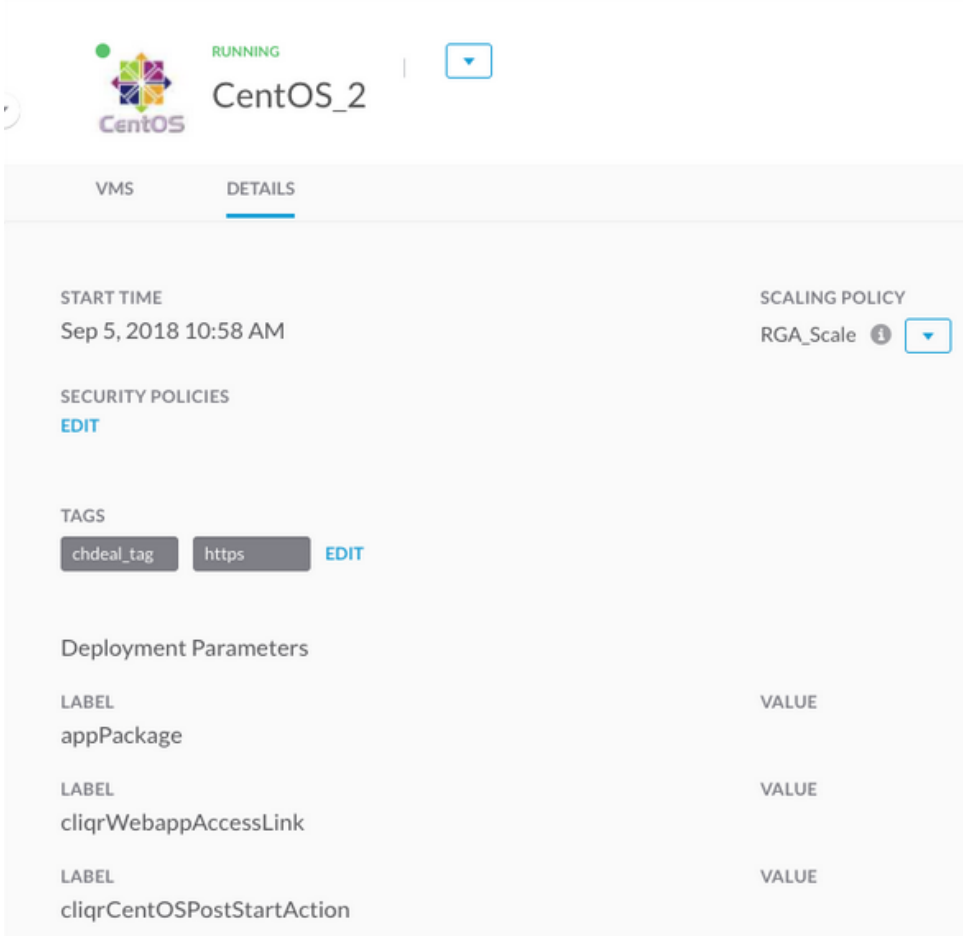
NAME	STATUS	IP ADDRESSES	INSTANCE TYPE	ACTIONS
<input type="checkbox"/> i-00f38c1b0351ab85	TERMINATED	4 IP Addresses	t2.micro	
<input checked="" type="checkbox"/> i-08ebc0948a711bd05	ERROR	4 IP Addresses	t2.micro	<ul style="list-style-type: none"> Reboot Sync VM Information Terminate CUSTOM ACTION execute script sudo
<input type="checkbox"/> i-02863c362fd622cd6	TERMINATED	4 IP Addresses	t2.micro	

- If there are no actions available for a particular VM, then there will not be an actions dropdown available.


The VM Details Tab

This tab contains basic information about the tier such as start time, stop time (if the tier has terminated), current policy settings and tags, and deployment parameters.


- If enabled in the deployment environment at the time of deployment, you have the ability to update policy settings and tags. The following screenshot is an example.



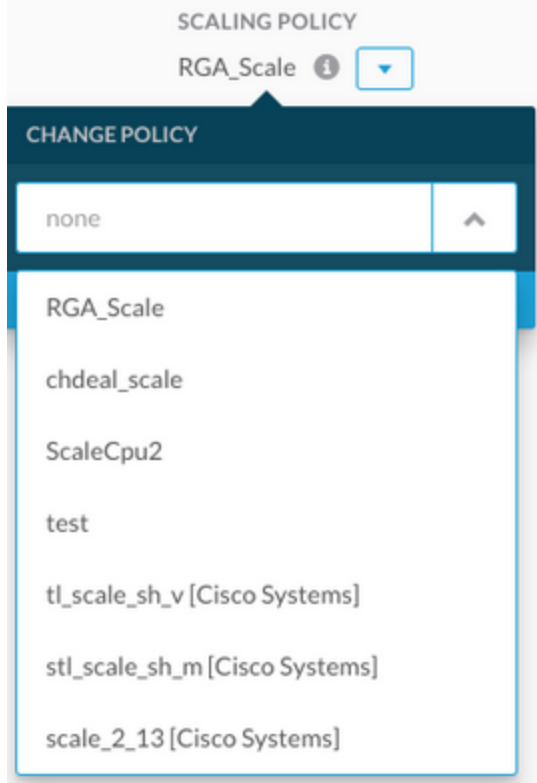
- Hover over the scaling policy's info icon to view the info balloon displayed in the following screenshot.



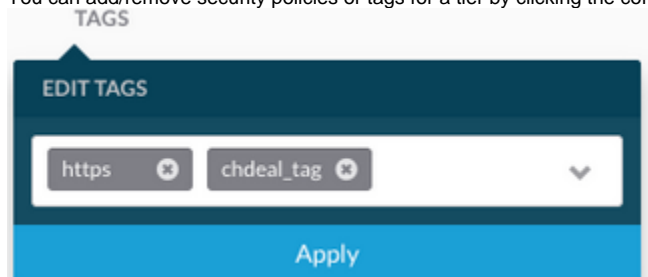
- You can change or remove an existing scaling policy or add one, if none is set, by clicking the dropdown menu icon next to the scaling policy label as displayed in the following screenshot.



- The **Change Policy** option displays a new dropdown list as displayed in the following screenshot that shows all available scaling policies for this tier based on restriction specified in the application profile's Topology Modeler tab.



- Select a policy and click **Apply** to replace the old policy with a new one.
- You can add/remove security policies or tags for a tier by clicking the corresponding **Edit** link as displayed in the following screenshot.



Google Cloud Nuance

Google Cloud does not support attachment of tags to VMs. Although the Workload Manager UI will allow tags to be specified, and shows success, tags are not added.

Container-based tiers differ based on the number of containers per pod.

- The following screenshot displays a single container per pod.

- The following screenshot displays multiple containers per pod (for more details on *Placement Groups*, see the [Define Resource Placement](#)).


- The tier header details are the same as the VM-based tier header except for the following factors:
 - In Approx cost, the projected monthly cost does not factor in suspension policy savings because suspension policies cannot be applied to containerized tiers.
 - In the Scaling actions dropdown, units are replicas, not VMs.
- The main section contains tabs for Replicas and Tier details.
- A placement group is represented by a **rectangle that you add to the topology modeler canvas** by clicking the **Create A Group** button. See [Define Resource Placement](#) > *Container-Specific Resource Placement* for additional details.

The Replicas Tab

The Replicas tab contains the following details:

- The Status buttons display the replica count in a certain state. The buttons are displayed based on the state of replica state (running, stopped, terminated, error, and so forth) in the tier.

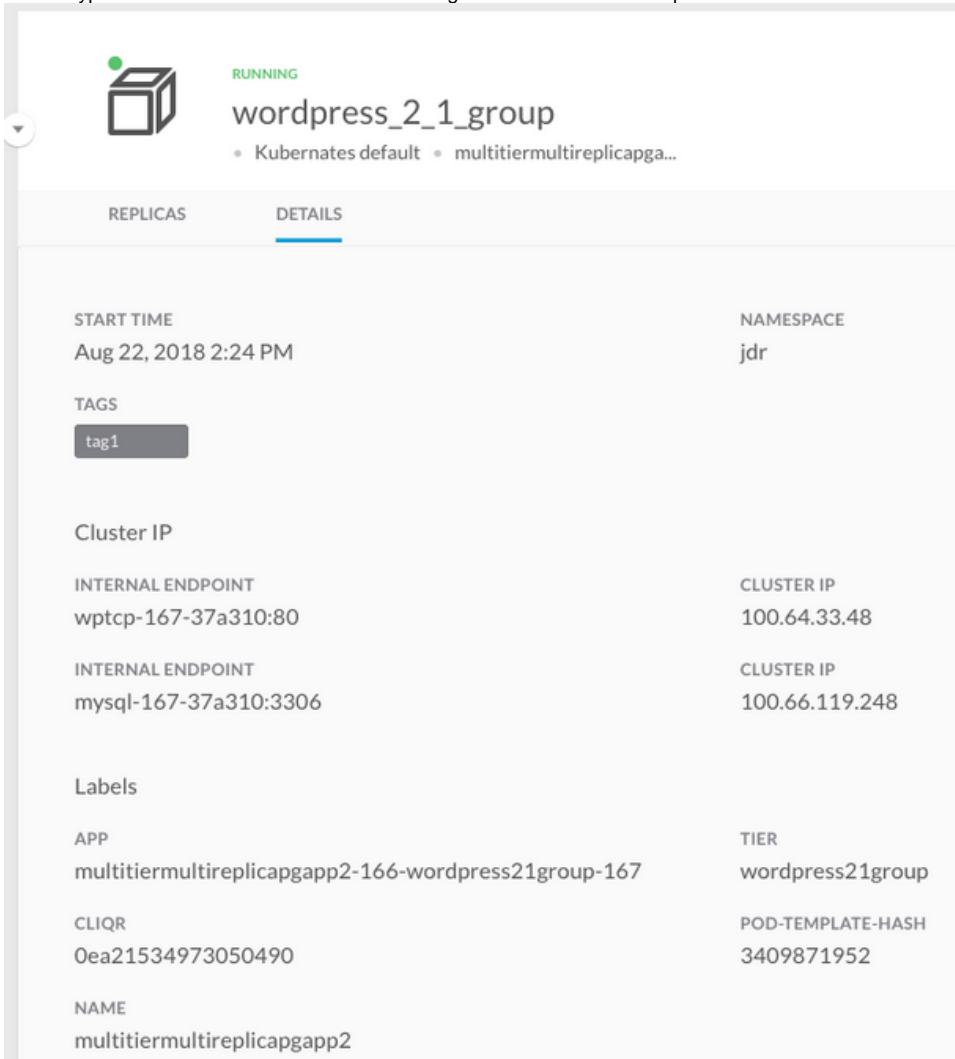
- The list of replica details in a tier along with additional information when you hover over most fields. For example, hovering over the number of parameters provides additional parameter details as displayed in the following screenshot.



The Container Details Tab

The container tier details tab is similar to the VM tier details tab with the following exceptions.

- There are no scaling policies for container tiers.
- For Kubernetes container tiers, the namespace, and IP addresses and internal endpoints related to the ClusterIP, NodePort and LoadBalancer service types in the tier are listed. See the following screenshot for an example.



General Information about deployment, for example, Approval/Start time and so forth are listed in the Details tab. This tab also contains actions related to Aging/Suspension/Security Policies (Remove/Change/Add) along with Tags, Global Parameters, and Metadata information. See the [Policy Management](#) section for additional details.

A list of all job state changes and policy changes since the job was deployed, in reverse chronological order are listed in the History tab.

TIERS	DETAILS	HISTORY																														
<div style="display: flex; justify-content: space-between; align-items: center;"> <input type="text" value=""/> Download </div> <table border="1"> <thead> <tr> <th>ACTION TAKEN</th> <th>USER</th> <th>DATE</th> <th>STATUS</th> <th>DETAILS</th> </tr> </thead> <tbody> <tr> <td>Addsuspensionpolicy</td> <td>Dev Cisco</td> <td>Sep 5, 2018 At 11:44 AM</td> <td>Completed</td> <td>Suspension policy 'tl_susp_sh_n' added to job.</td> </tr> <tr> <td>Addsuspensionpolicy</td> <td>Dev Cisco</td> <td>Sep 5, 2018 At 11:44 AM</td> <td>Submitted</td> <td>Initiating add suspension policy 'tl_susp_sh_n' to job.</td> </tr> <tr> <td>Removesuspensionpolicy</td> <td>Dev Cisco</td> <td>Sep 5, 2018 At 11:44 AM</td> <td>Submitted</td> <td>Initiating remove suspension policy 'suspension_2_14' to job.</td> </tr> <tr> <td>Removesuspensionpolicy</td> <td>Dev Cisco</td> <td>Sep 5, 2018 At 11:44 AM</td> <td>Completed</td> <td>Suspension policy 'suspension_2_14' removed from job.</td> </tr> <tr> <td>Addagingpolicy</td> <td>Dev Cisco</td> <td>Sep 5, 2018 At 11:43 AM</td> <td>Completed</td> <td>Aging policy 'aging78' added to job.</td> </tr> </tbody> </table>			ACTION TAKEN	USER	DATE	STATUS	DETAILS	Addsuspensionpolicy	Dev Cisco	Sep 5, 2018 At 11:44 AM	Completed	Suspension policy 'tl_susp_sh_n' added to job.	Addsuspensionpolicy	Dev Cisco	Sep 5, 2018 At 11:44 AM	Submitted	Initiating add suspension policy 'tl_susp_sh_n' to job.	Removesuspensionpolicy	Dev Cisco	Sep 5, 2018 At 11:44 AM	Submitted	Initiating remove suspension policy 'suspension_2_14' to job.	Removesuspensionpolicy	Dev Cisco	Sep 5, 2018 At 11:44 AM	Completed	Suspension policy 'suspension_2_14' removed from job.	Addagingpolicy	Dev Cisco	Sep 5, 2018 At 11:43 AM	Completed	Aging policy 'aging78' added to job.
ACTION TAKEN	USER	DATE	STATUS	DETAILS																												
Addsuspensionpolicy	Dev Cisco	Sep 5, 2018 At 11:44 AM	Completed	Suspension policy 'tl_susp_sh_n' added to job.																												
Addsuspensionpolicy	Dev Cisco	Sep 5, 2018 At 11:44 AM	Submitted	Initiating add suspension policy 'tl_susp_sh_n' to job.																												
Removesuspensionpolicy	Dev Cisco	Sep 5, 2018 At 11:44 AM	Submitted	Initiating remove suspension policy 'suspension_2_14' to job.																												
Removesuspensionpolicy	Dev Cisco	Sep 5, 2018 At 11:44 AM	Completed	Suspension policy 'suspension_2_14' removed from job.																												
Addagingpolicy	Dev Cisco	Sep 5, 2018 At 11:43 AM	Completed	Aging policy 'aging78' added to job.																												

Besides viewing the details on this page, you can perform the following actions:

- Click the magnifying glass icon in the upper right to filter the list to entries that contain the text in the filter.
- Click the download icon in the upper right to download the complete history as a CSV file.

The **History** tab provides details on actions taken or incidents that occurred during the life of this deployment. You can download the details on this page to a CSV file.

The Topology panel displays the application topology with status indicator lights (red, yellow, green) for each tier. Filter buttons appear above the topology diagram representing all tiers and tiers in various states: running, error. Click on a filter button includes only tiers in that state as displayed in the following screenshot.

The screenshot shows the application topology panel with two tiers:

- apache2_1** (ERROR status):
 - Buttons: ALL VMS, 1 TERMINATED
 - Table:

NAME	STATUS	IP ADDRESSES	INSTANCE TYPE	ACTIONS
cqjw-4401f61cc	TERMINATED	10.88.163.207	small	
- MySQL Database** (CANCELLED status):
 - Buttons: ALL VMS, 1 TERMINATED

The topology diagram on the right shows a red dot on the 'apache2_1' tier icon and an orange dot on the 'MySQL Database' tier icon, connected by a blue line.

A colored dot on the upper left corner of each tier icon represents the status of the tier during the deployment process as described in the following table.

Indicator Color	Application Tier Status
Blinking Yellow	The application tier is in the pre-initialization state.
Blinking Green	The application tier is initializing.
Solid Green	The application tier is up and running.
Solid Red	The application tier has an error.

For all actions, notifications are displayed at the top of the page. Three types of notifications are displayed in the details page:

- In Progress – stays on the page during the action.
- Success – stays for 7 seconds.
- Error – you must specifically acknowledge this notification.

Three or more notifications of the same type are grouped together and you can click **View Details** to see details for this group.

Terminate Protection

Terminate Protection

- [Overview](#)
- [Termination Notes](#)
- [Toggle Visibility](#)
- [Enable Terminate Protection](#)
- [Allow Termination](#)
- [Terminate Existing Deployment](#)

The Workload Manager provides a feature to prevent the termination of nodes when an application is in the process of being deployed. This feature allows you to prevent an inadvertent stoppage from the Workload Manager and thus avoid interrupting production deployments.



The Terminate Protection feature is only applicable to N-tier jobs.

If you enable terminate protection for a job (Job A1), and disable terminate protection in the deployment environment and this deployment is active, then Job A1 cannot be terminated.

If you enable both the aging policy and the prevent termination feature, the prevent termination feature takes precedence.

Be aware of the automatic termination by the Workload Manager in the following cases:

- The Workload Manager terminates VMs randomly once a scale down is executed.
- The Workload Manager does not terminate jobs when the user runs out of credits.
- Both batch actions or single actions are available to terminate VMs as listed in the [Deployment Details](#) section.
 - You can terminate a single running VM(s) using one of two methods:
 - Use the **Scale Down** option from the [Deployment Details](#) > *Tier Header*. This method allows you to terminate one VMs or even, multiple VMs at the same time.
 - Use the **Terminate** action from the Actions dropdown for the VM. See [Deployment Details](#) > *The VMs Tab*. This method allows you to terminate the specified VM.
 - You can also terminate multiple, specific VMs using batch actions as specified in the [Deployment Details](#) > *VM-Based Tiers* section.

The **Allow Terminate/Suspend Protection** toggle switch is OFF by default in the Deployment Environment form's **Policy Setting** tab. This default setting makes the **Terminate/Suspend Protection** toggle invisible to users in the *Deploy* form.

To configure the visibility of the **Terminate/Suspend Protection** toggle in the Deploy form, you can change the toggle switch as highlighted in the following screenshot.

New Deployment Environment

General Settings Cloud Settings **Policy Settings**

Aging Policy

ALLOWED AGING POLICIES

Suspension Policy

ALLOWED SUSPENSION POLICIES

All

MANDATORY

OFF

ALLOW TERMINATE/SUSPEND PROTECTION

ON

Deploy Time Preview

SUSPENSION POLICY

tl_susp_sh_n [Cisco Systems]

Security Policy


ALLOWED SECURITY POLICIES

The **Allow Terminate/Suspend Protection** toggle switch allows you to set the visibility of the **Terminate/Suspend Protection** setting:

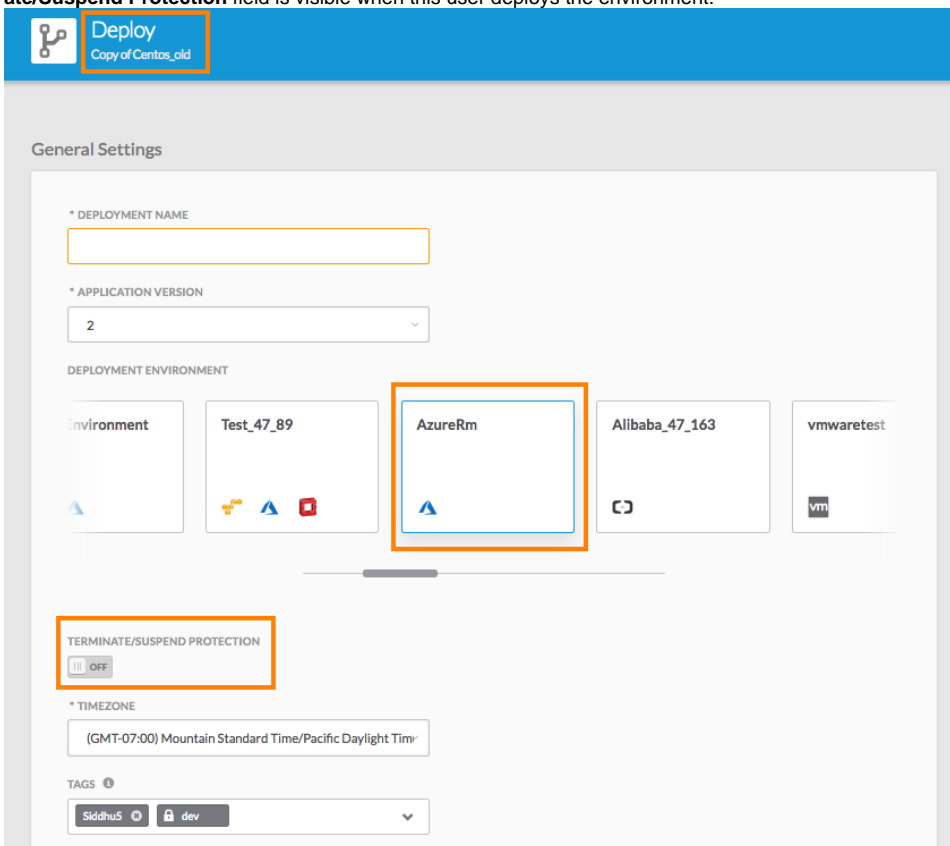
- **OFF** (Default): The **Terminate/Suspend Protection** setting is **not visible** from the Deploy form when this environment is selected at deploy time.
- **ON**: The **Terminate/Suspend Protection** setting is **visible** on the Deploy form when this environment is selected at Deploy time.

When [Deploying an Application](#) (Step 3), the General Settings section of the Deploy form will correspondingly display the **Terminate/Suspend Protection** field *ONLY* if the selected Deployment Environment was configured to ON as evident in the following screenshots:

- The following screenshot displays an environment where the default setting (OFF) was not changed and thus the **Terminate/Suspend Protection** field is not visible when this user deploys the environment:



- The following screenshot displays an environment where the **Allow Terminate/Suspend Protection** setting was toggled ON and thus the **Terminate/Suspend Protection** field is visible when this user deploys the environment:



✔ If a termination policy executes on a VM on which terminate protection is enabled, then protection overrides termination.

You can also prevent termination by selecting the **Enable Terminate Protection** option in the Deployments page from the **Actions** dropdown list and confirm your intention in the Enable Terminate Protection popup. The following screenshot shows this Actions drop-down list.

The screenshot shows the 'Deployments' page in Cisco CloudCenter. At the top, there are filters for 'Running' status, 'All Applications', and a 'Show Hidden' checkbox. A search icon, a refresh icon, and a 'New Deployment' button are also present. Below the filters is a table with the following columns: NAME, STATUS, ENVIRONMENT, START TIME, RUN TIME, CLOUD COST, and ACTIONS. Two deployments are listed:

NAME	STATUS	ENVIRONMENT	START TIME	RUN TIME	CLOUD COST	ACTIONS
WP-K8-1 Wordpress Container 2-Tier (V1) K8s-Google default	Deployed	AllClouds	Apr 13, 2018 at 01:18 PM		\$0.00	-Actio... ^
apr12-1 Copy of Wordpress Container 2-T... K8s_CCP default	Deployed	AllClouds	Apr 12, 2018 at 11:28 AM		\$0.00	-Actions-

The expanded actions menu for the second deployment shows the following options: Terminate, Terminate A..., Enable Term..., and Share. At the bottom right of the table, there is a 'Show 50 per page' and a 'Page' indicator.

When you enable the *Terminate Protection* feature, you will not be allowed to Stop, Suspend, or Migrate the deployment. However, you can promote the deployment, if required, to a different environment.

To allow termination, select the **Disable Terminate Protection** option and confirm your intention to Disable Terminate Protection popup.

You can terminate an existing deployment using the **Terminate** existing deployment option.

You can set this option when promoting or migrating an application. See [Deployment, VM, and Container States](#) for additional context.

This feature has the following options:

- **Terminate:** Terminates the deployment
- **Hide:** Just hide this job (for jobs with Errors)
- **Terminate and Hide:** Terminates and hides the deployment

Track Cloud Costs

Track Cloud Costs

- [Overview](#)
- [Instance Type Pricing](#)
- [Compute Cost and Storage Cost](#)

Workload Manager reports cloud costs of a deployed VM based on run time and per units costs of the VM's instance type and storage types. These costs are summarized in four place in the Workload Manager UI:

- [Usage Summary Report](#)
- [Applications Deployment Report](#)
- [Deployments Tab](#)
- [Virtual Machines Tab](#)

You must wait for the next billing cycle (runs at the top of every hour) to view the cost and usage details. See [Financial Overview](#) for additional context.

Administrators can customize cloud instance type pricing. Consider a scenario where a company has a volume pricing agreement for a public cloud such that consumption is below list price or a private cloud that a central governing authority wants to pass on costs to a set of constituents. Using the administrative tools provided by the the CloudCenter Suite, enterprises have the ability to control and override prices and address other similar situations.

The CloudCenter Suite offers granular compute and storage cost calculations customized for the following cloud providers:

- AWS
- Google Cloud Platform
- Microsoft Azure

The cost reporting models for compute and storage for these public cloud providers are summarized in the following two tables.

Compute charges:

Cloud Provider	Image Type	Billing Cycle	Minimum Duration	Rounding
AWS	Amazon Linux, Ubuntu, Marketplace AMIs without hourly charge	1 second	60 seconds	None
AWS	Windows, RHEL, SLES, Marketplace AMIs with hourly charge	1 hour	1 hour	Round up
AzureRM	All	1 minute	1 minute	Round down
GCP	All	1 second	60 seconds	None

Storage charges:

Cloud Provider	Billing Cycle	Minimum Duration	Rounding
AWS	1 second	60 seconds	None
AzureRM	1 hour	None	Round up
GCP	1 second	60 seconds	None

For these public clouds:

- When a VM or deployment is suspended, compute charges *cease to accumulate*.
- Storage charges *continue to accumulate* until the VM or deployment is terminated.

CloudCenter uses these billing models to calculate and update the cumulative cloud costs as follows:

- For a deployment, it is displayed in the following list pages in the Deployments list page.
- For a VM, it is displayed in the Virtual Machines list page.

The calculation considers all suspend and resume events for the VM or deployment since the last calculation. These updates occur at the top of the hour, when:

- A VM or a deployment is manually suspended.
- A deployment is automatically suspended due to a suspension policy.

The storage cost is *not* calculated for unmanaged VMs.

 **AWS Cloud Nuance**

For Unmanaged VMs imported from AWS, before the CloudCenter agent is installed, the OS is unknown; therefore the one hour billing cycle model is used in this situation.

For VMs on *all other public and private clouds*, the one hour billing cycle model is used.

Project and Phase Management

Project and Phase Management

- [Overview](#)
- [Project Participants](#)
- [High-Level Project Flow](#)
- [Create a Project](#)
- [Project View](#)
- [Status Indicators](#)
- [Draft vs. Published Mode](#)

Projects enable enterprises to create a workflow to manage their devops process in Workload Manager. The devops process may involve different stages (phases) such as Development, Test, Stage, Production and so forth. Each phase can have its own participants, environments for deployment, budget, as well as resource constraints. Applications would have to go through these phases before going live.

Workload Manager already has the concept of [Setup Deployment Environments](#) with additional [System Tags](#) and various integration points like the [Jenkins Integration](#) plugin. However, the devops process has to be indirectly modeled using [Access Control Lists](#) on Deployment Environments to move an application from one environment to another.

Using the Workload Manager *Projects* feature, enterprises can create a workflow to represent their devops process, setup participants and environments, and enforce monetary and resource limits from one central location.

A project allows you to:

- Create and configure various phases that map to the devops process.
- Identify permitted users
- Choose participant applications
- Impose budget and resource limits

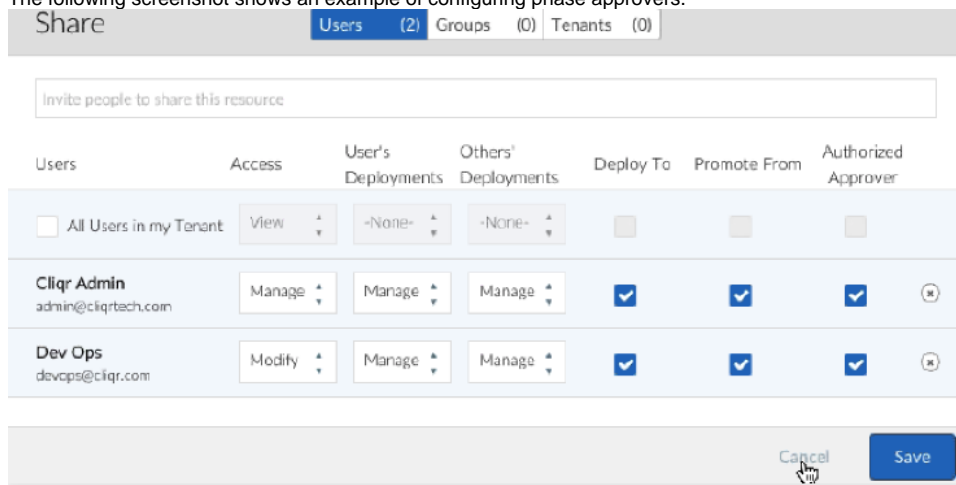
A project does not allow you to:

- Define application profiles – You *must* define an [Application Profile](#) first before associating it with a project.
- Create users – You *must* create users in the Suite Admin UI before adding them as participants.

Workload Manager project participants include:

- **Admin/Project Manager:** A project manager may be a tenant admin or any user that is assigned the Project Admin role and has the ability to perform the following functions:
 - Create a project.
 - Define lifecycle phases. Each phase must set up a unique deployment environment.
 - Identify users who can view and promote applications from one phase to another (see [Permission Control](#)).
 - Set a usage plan for each phase.
- **Authorized (Phase) Approver:** A phase approver is given *Promote From* and *Authorized Approver* privileges for the deployment environment corresponding to the phase and can perform the following functions:
 - Approve the deployment of applications into a phase.
 - Promote/demote applications from one phase to another.

The following screenshot shows an example of configuring phase approvers.



- **Project User:** A project user must have *Deploy To* permissions for the deployment environment corresponding to the phase and can perform the following functions:
 - View projects, phases, and deployments within.
 - Deploy Application Profiles within permitted phases.

- Receive notifications each time a deployment is promoted/demoted.

Projects follow this workflow:

1. The tenant admin/project manager creates a project and adds participants.
2. The project manager:
 - a. Adds the phases and associates plans with them.
 - b. Adds participants for each phase to projects.
3. Project users make phased deployments on a day-to-day basis.
4. Phase approvers:
 - a. Promote a deployment to the next phase.
 - b. Approve the promotion (of the next phase).
5. Permitted users deploy the application in the approved phase.

To create a project, follow this process:

1. Access the Workload Manager UI and click **Deployments > Projects** (default tab).
2. Click the **New Project** button to create a project. The Create Project popup displays.



The **New Project** button is only visible to users who have permission to create projects. Be sure to assign the Project Admin role to the required user(s).

The following screenshot shows the Create Project popup.

Create Project close

Name *

Description

Plan Type

Maximum Running VMs
 VM(s)
The maximum number of concurrently running VMs allowed.

Target End Date

Application Profiles | Only these Application Profiles can be deployed in this Project

- CCA
- CentOS**
- Centos6x_Parallel_Ex_1

Note: To select multiple application profiles, CTRL/CMD+Click on the options.

Continue **Cancel**

3. Provide a Name for this project.
4. Select the Usage Plan Type and budget for the entire project.



Once selected, the plan type applies to all phases of the project. The corresponding plan for the phase is used for billing purposes and tracked based on either budget or VMs, but not both.

5. Select the Application Profiles that can be deployed with this project.
6. Click **Continue** to proceed to the Share popup. The following screenshot shows this popup.

Share **Users (2)** Groups (0) Tenants (1)

Invite people to share this resource

Users	Access
<input checked="" type="checkbox"/> All Users in my Tenant	View
Cliqr Admin admin@cliqrtech.com	Manage
QA Admin qaadmin@cliqr.com	Modify

Save Cancel

7. Identify the users who can view this project, modify the project, or have the ability to share this project with others.



Project are **only** displayed in the Project owner's dashboard. Even if other users are added to it, it will show up in the user dashboard *after* the project is **published**.

Share Users (2) **Groups (1)** Tenants (1)

Invite groups to share this resource

Groups	Access
Doc Group	View <input type="button" value="v"/>

Share Users (2) Groups (1) **Tenants (1)**

This resource will be shared with all users in the shared Tenant's hierarchy.

Invite tenants to share this resource

Tenants	Access
<input type="checkbox"/> My Tenant & Sub-Tenants	View <input type="button" value="v"/>
Doc1	View <input type="button" value="v"/>

- Click **Save** to save the permissions assigned to the users for this project and proceed to the project board, as shown in the following screenshot.

Project Doc Project

[edit project detail](#) | [sharing and access](#) | [delete](#)

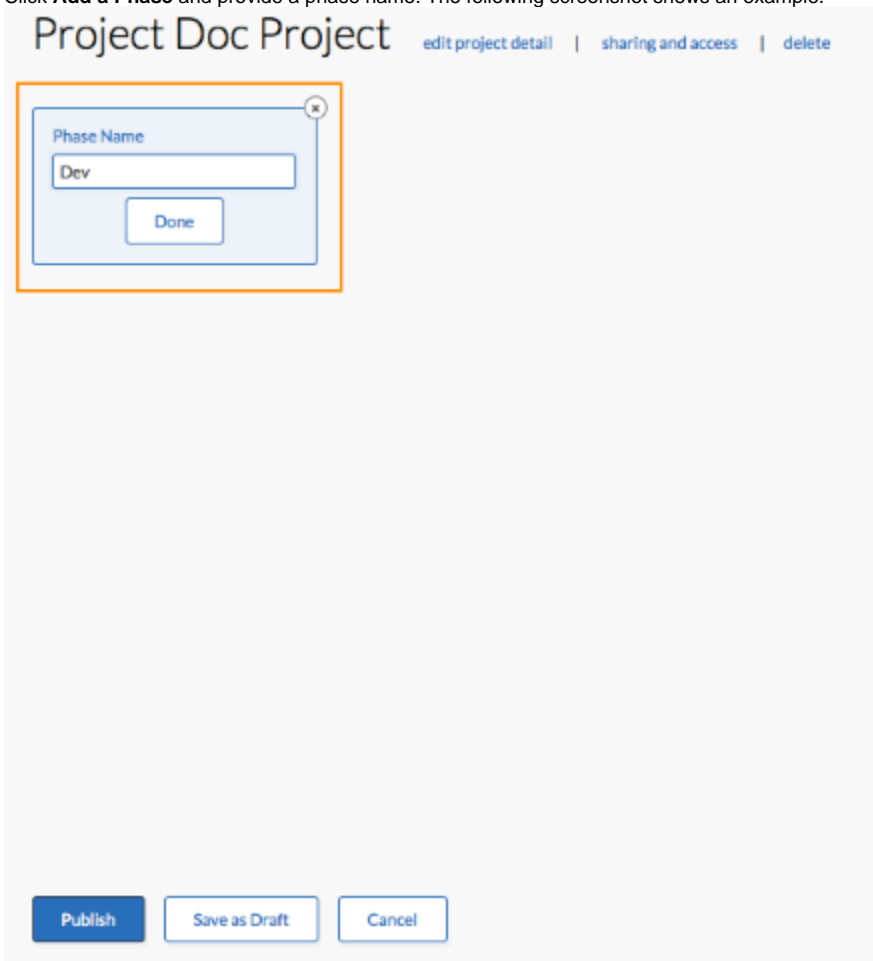
Add a Phase

Publish

Save as Draft

Cancel

- Click **Add a Phase** and provide a phase name. The following screenshot shows an example.



The screenshot shows the 'Project Doc Project' interface. At the top, there are links for 'edit project detail', 'sharing and access', and 'delete'. Below this, a modal dialog box is open with the title 'Phase Name'. Inside the dialog, there is a text input field containing the text 'Dev' and a 'Done' button below it. The dialog box is highlighted with an orange border. At the bottom of the main interface, there are three buttons: 'Publish', 'Save as Draft', and 'Cancel'.

- Choose a phase type, (the deployment environment to be associated with this phase) or click **New Deployment Environment** to create a new one.



The same deployment environment cannot be used by two phases.

The following screenshot shows an example of choosing a phase type.

Project Doc Project

[edit project detail](#) | [sharing and access](#) | [delete](#)

Dev

Select Phase Type

Dev

- azuretesting
- qa_softlayer
- testasd
- tej_hp

GCE Test


- amazon_uswest1_de
- OpenStack Only
- depenv-gj
- sudepenv
- test_mj
- test_mj2
- test_mj3
- test_mj4

DepEnv - Project - 01 - VP - Pha...


qatestsunitha

New Deployment Environment

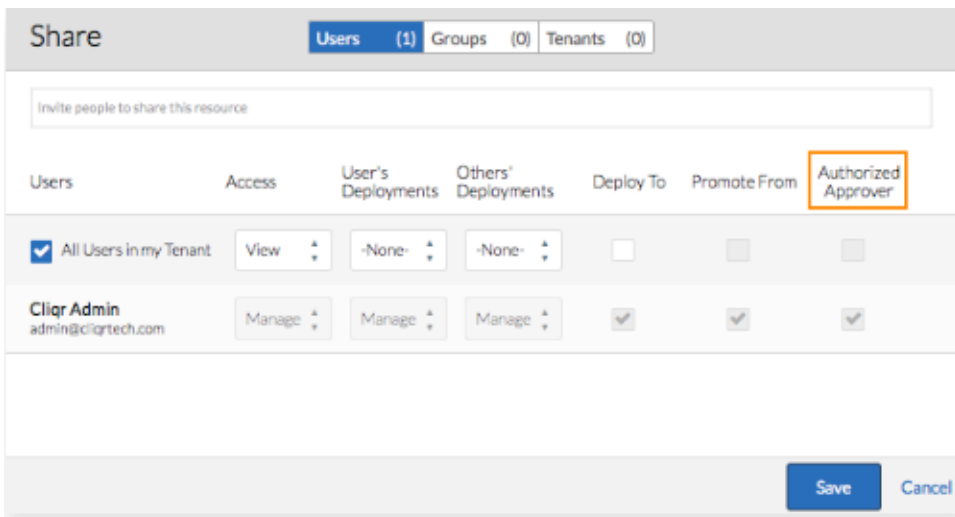
11. Choose an appropriate plan for this phase based on the project budget type, as shown in the following screenshot.



12. Add the participants of the phase.


 Any change you make in the sharing permission for each Phase is automatically saved for the underlying deployment environment's permission as well.

The following screenshot highlights the Authorized Approver options.



- 13. Review all phases for this project as required by your enterprise.
- 14. Modify permissions for each phase, if required.

15. Click **Save as Draft** until you configure all project details, as shown in the following screenshot.



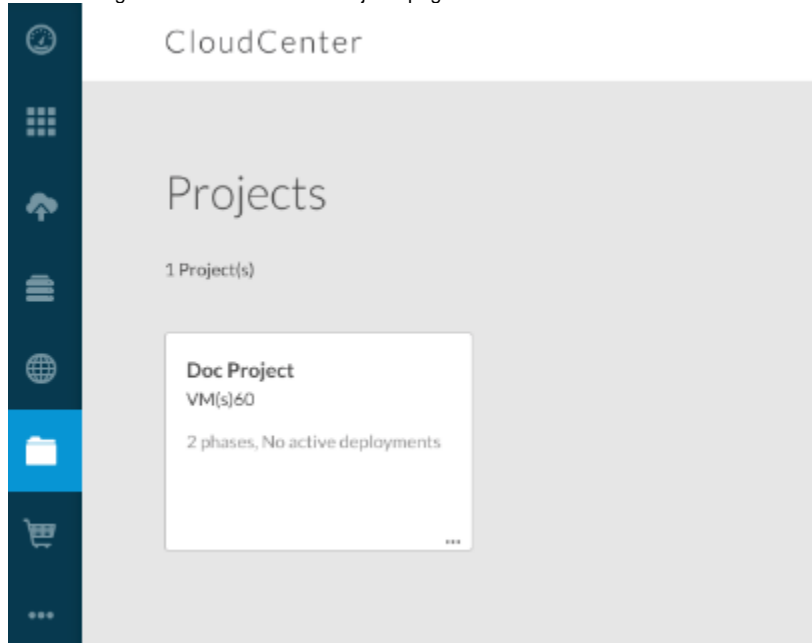
16. Click **Publish** when you have configured all the project details.

You can view each project by accessing the Workload Manager UI > **Projects** link. This end-to-end view allows you to see projects, the phases for each project, and the status indicators for your application deployments.

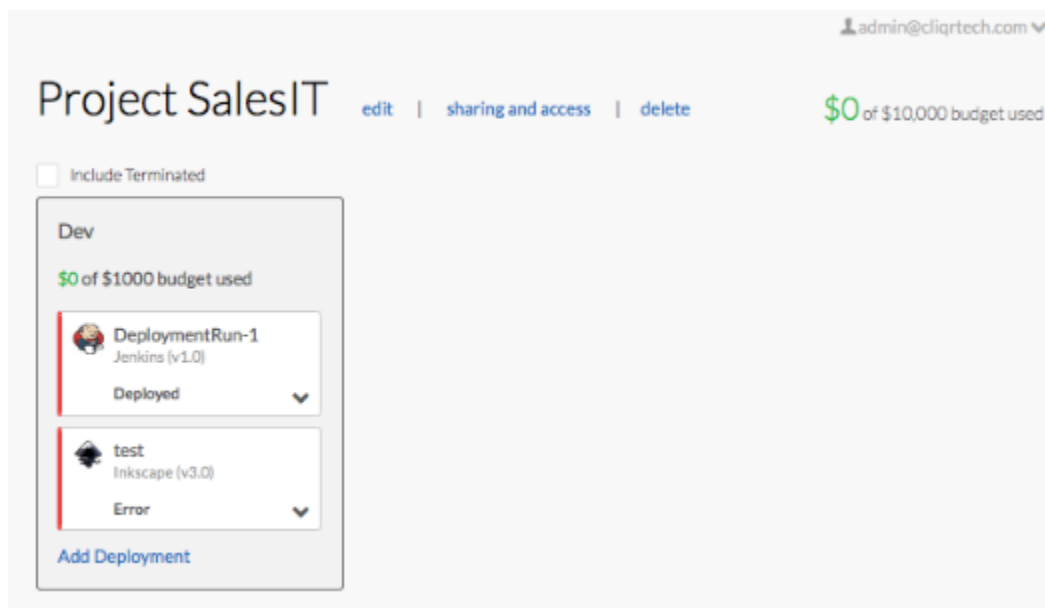


Deployments made using the **Projects** link can only be viewed and managed through the Projects dashboard.

The following screenshot shows the Projects page.



Click any project in the Project *name* page to drill down into the project, as shown in the following screenshot.



The Project *name* page provides color-coded indicators for each phase, as described in the following table.

Status Indicator	Description
Red	This deployment status is Errored or Rejected (approval is rejected by the phase approver).
Orange	This deployment phase is suspending, suspended or pending approval.
Green	This deployment is in progress or deployed.
Grey	This deployment has been terminated.

If a project does not have associated phases and application, you cannot publish a project. Once published, you cannot move the project back to the Draft mode.

When you save a project, the phase attributes (Bundle, Plan, Phase Name, and Deployment Environment) behavior differs based on the mode in which the project was saved, as described in the following table.

Operations	Project in Draft Mode	Published Project
Change name	Allowed	Allowed
Change deployment environment of a phase	Allowed	Not Allowed
Change plan	Allowed	<ul style="list-style-type: none">• Plan Type = Pre-paid Budget: Not Allowed

Migrating Applications

Migrating Applications

- [Overview](#)
- [Migrating an N-Tier Application Containing Data](#)
 - [Migration Guidelines](#)
 - [Migrate an N-Tier Application with Data](#)
- [Migrating an N-Tier Application Using Images](#)
- [Migrate without Suspending Deployments](#)

To migration applications from your existing cloud to your target cloud, be aware of the following requirements:

- Stateless application profile that can deploy a fresh/clean copy of the application without a state in either cloud
- A method to preserve and move the application state (containing data) – for example, take a database dump, copy it to cloud storage, then copy back, and restore the database.

If you have these two requirements in mind, the CloudCenter platform can help with your migration! The CloudCenter migrate feature helps you deploy a new, clean copy of the application to your target cloud and then perform the application state move based on the scripts you provide.

There is no reason to delete the VM(s) on the source side first – you can just leave them running or delete them after the migration is successful.

The migration scenarios in this section are just two examples to consider when migrating applications.

This section uses an application called RollerWeBlog to describe the CloudCenter process to migrate a three-tier Java web application with data running at the customer's location.

Migration Guidelines

- Services:
 - CloudCenter does not provide out-of-box mapping for any cloud-specific service to an alternate service in another environment.
 - If using cloud-specific services (such as ELB or RDS), then the author/owner of the Application Profile must be aware that are making a decision to lock that Application Profile to a particular cloud (for example, AWS).
 - However, you can create a custom service definition and include the scripts required to detect whether it was being deployed in AWS or in VMWare or another cloud and then connect that application to either the ELB or their F5 load balancer.
- DNS:
 - CloudCenter does not inherently update the DNS to point to a migrated deployment.
 - However, you can configure the migration process to include any scripts required to change the DNS mapping in your environment.

Migrate an N-Tier Application with Data

To migrate an N-tier application with data, follow this process:

1. Back up the data for the RollerWeBlog application. For this example:
 - The data backup file is called **dbbkup.sql**.
 - The RollerWeBlog application runs on a different cloud than CloudCenter.
2. Identify the relevant components, [Parameters and Macros](#), and [Configuration Files](#) for this application.
3. Modify the Configuration File (or properties file) to include CloudCenter-defined system macros to automatically plug in the appropriate values for parameters defined in the configuration file. Alternately, you can pass these parameters as arguments to install or configuration scripts.
 - The RollerWeBlog application configuration file changes:

```

../WEB-INF/classes/myconfig-file
installation.type=auto
mediafiles.storage.dir=/usr/local/rollerdata/mediafiles
search.index.dir=/usr/local/rollerdata/searchindex
log4j.appender.roller.File=/usr/local/rollerdata/roller.log
database.configurationType=jdbc
database.jdbc.driverClass=com.mysql.jdbc.Driver
database.jdbc.connectionURL=jdbc:mysql://%DB_TIER_IP%:3306/rollerdb?
autoReconnect=true&useUnicode=true&characterEncoding=utf-8&mysqlEncoding=utf8
database.jdbc.username=scott
database.jdbc.password=tiger
mail.configurationType=properties
mail.hostname=smtp-server.example.com
mail.username=scott
mail.password=tiger

```

- Prior to the macro substitution, the configuration file had the following description in the highlighted area: [database.jdbc.connectionURL=jdbc:mysql://localhost:3306/rollerdb?](#) Effectively, the localhost IP address is replaced by a dynamically generated private IP address (database tier) during runtime using `%DB_TIER_IP%`.
 - In Application Profiles, users can define parameters to pass as arguments to the username and password field. Optionally, users can also include default values in scripts.
4. After modifying the configuration files with system macros, create a .war application package and include the updated configuration file.
 5. Upload the application data (packages, configuration files, backup data, SQL script, and other scripts) to your shared directory in the [Artifact Repository](#).
 - a. Create an application directory under `/storage/app/<application name>`.
 - b. Upload the application data to the newly created directory.

Model (define) the Application. For the RollerWeBlog application, select the Java Web App Profile. For each tier in the application, associate System Tags and if required, configure the Environments, the target cloud, the Prevent [Terminate Protection](#) behavior, the [Aging and Scaling Policy](#) details the [Metadata for Custom Properties](#), and [instance type\(s\)](#) as applicable to your deployment.

6. If the application requires certain parameters to be defined or overridden (administrator, username, and password), include them in the [Topology Modeler's Global Parameters](#) section.
7. Define the application architecture and services using the Topology Modeler. For the RollerWeBlog example, use Tomcat as the web server, MySQL as the database, and NginX as the load balancer. For each tier, use the Properties panel to provide additional details. For example:
 - MySQL tier: Parameters, username, password, the path for the database script that has the backup data from a previous environment, and so forth.
 - RollerWeBlog tier: Define the path for the application binary file (roller.war) and the configuration file.
 - Similarly, provide the dependent details for the Tomcat and NginX tiers as well.
8. Save the Java web app profile. Access your new application from the Apps tab and verify your changes for each tier. The application is now ready for deployment.
9. Submit the application for deployment.
10. View the deployment progress in the Status field.
11. SSH or VNC to the cloud VM for additional troubleshooting access or to run additional commands and scripts.
12. Once the deployment is complete, click the Access link to open the IP address for this application. You may need to specify the full path in the context of the application in the Topology Modeler's Basic Information pane.
13. Use the full path URL to access the migrated RollerWeBlog application.

You can import existing application images running in a cloud and use them as custom images for a CloudCenter tenant.

This example describes a Siebel CRM application that uses an Oracle database with data running at the customer's location. The data is being migrated from a dedicated on-premises server to a cloud.

To migrate an N-tier application with data, using images, follow this process:



You need Workload Manager administrator privileges to perform this procedure.

1. Back up the data running on the Siebel CRM application.
2. Create an image for each tier in the application:
 - a. The database with data.
 - b. The Siebel CRM application with all dependencies.
3. Identify the relevant components, [Parameters and Macros](#), and [Configuration Files](#) for this application.
4. Log into Workload Manager UI and access **Admin > Image**:
 - a. Click the **Add New** link to add a new image.
 - b. In the Add a New Image page, fill in the fields to create a logical entry for your image in CloudCenter:
 - OS Type: The OS on which the image is based. For Siebel CRM, the OS is Linux.
 - Number of Network Interfaces: The base image dictates the number of NICs.
 - Enabled: Check this box if you want to use this image.
5. Map the logical images to your actual image:
 - a. In the Image Mappings section, select the Cloud Type for the image from the dropdown list.
 - b. Enter the [Image ID](#) associated with the Cloud Type.
 - c. Select all applicable Instance Types supported for this image. You can also override the image cost by not modifying the cost.
 - d. After adding the mapping, you can view the mapped image by clicking the image name. You can edit/convert/delete images after creating them.
 - e. Repeat the step for each image associated with each tier in this application (for example, the database image, SiebelDB).
 - f. Users will see the newly added images during their application deployment process.
6. [Configuration Files](#) (or properties file) to include CloudCenter-defined system macros to automatically plug in the appropriate values for parameters defined in the configuration file. Alternately, you can pass these as arguments to install or configuration scripts.
7. After modifying the configuration files with system macros, create a .war application package and include the updated configuration file.
8. Upload the application data (packages, configuration files, backup data, SQL script, and other scripts) to your secure, shared directory in the [Artifact Repository](#):
 - a. Create an application directory under `/storage/app/<application name>`.
 - b. Upload the components to the Application Repository to the newly created directory.
9. Define the [Application Profile](#). For the Siebel DB application, select the N-Tier Execution App Profile: For each tier in the application, associate System Tags and if required, configure the Environments, the target cloud, the Prevent [Terminate Protection](#) behavior, the [Aging and Scaling Policy](#) details, and [instance type\(s\)](#) as applicable to your deployment.
10. If the application requires certain parameters to be defined or overridden (administrator, username, and password), include them in the [Topology Modeler's Global Parameters](#) tab.

11. Define the application architecture and services using the Topology Modeler. For each tier, use the Properties panel to provide additional details. For example:
 - a. Siebel Database tier: Parameters, username, password, the path for the database script that has the backup data from a previous environment, and so forth.
 - b. Siebel CRM tier: Define the path for the application binary file and the configuration file.
 - c. Similarly, provide the dependent details for other applicable tiers as well.
12. Save the N-Tier app profile as Siebel CRM application. Access your new application from the Apps tab and verify your changes for each tier. The application is now ready for deployment.
13. [Deploy the Application](#):
 - a. Depending on the hardware requirement and application-specific requirements, select the cloud and instance types from the displayed list.
 - b. Select the required [Instance Types](#).
14. Submit the application for deployment.
15. View the deployment progress in the Status field.
16. SSH or VNC to the cloud VM for additional troubleshooting access or to run additional commands and scripts.
17. Once the deployment is complete, click the Access link (access URL) to open the IP address for this application. You may need to specify the full path in the context of the application in the Topology Builder's Basic Information pane.
18. Use the full path URL to access the migrated Siebel CRM application.

Suspension is an optional setting during the migration process. This change helps when testing your migration – without suspending your deployment. The following screenshot highlights suspending a deployment.

Migrate "Sample Tomcat App"

Migrate your cloud application by specifying application parameters and selecting cloud to run.

General Information

We need to know a little about the application you would like to run. This will help us figure out the best solution for your needs.

Deployment Name *

Description

Associate Tags

Deployment Environment

Cloud

▶ **Advanced**

tomcat7_1

▼ **Basic configuration**

Associate Tags

Instance Type *

1 cpu, 1024MB memory, 8GB local storage, 64bit OS, cost: \$0.013/node hour

▶ **Advanced**

Suspend Deployment from previous phase
Note: Uncheck to leave this application in Deployed Status in the previous phase.

If you check (unchecked by default) the **Suspend Deployment from previous phase** checkbox, the application is terminated and NOT suspended.

Container Tier Rolling Updates

Container Tier Rolling Updates

- [Overview](#)
- [Restrictions](#)
- [Update Process](#)
- [Troubleshooting](#)

When you deploy an application with at least one container-based tier, and the deployment is in the deployed state, the *Update* on-demand action becomes available for that deployment. This action is visible from the the Action dropdown field in the the Deployments List page and in the deployment's Deployment Details page.

The rolling update action is only available from the Deployments list page, not the Deployment details or Job details page.

When you select this action you are presented with a dialog box enumerating all of the container images in your deployment. The following screenshot shows an example dialog box with two container images.

The screenshot shows a dialog box titled "Update / ap-mon-c". It lists two container images with their respective version tags in input fields:

- apache2container_1_group** (parent group):
 - Child: **apache2container_1**
 - Image Version:
- mongocontainer_1**:
 - Image Version:

At the bottom right of the dialog, there are two buttons: "Update" (in blue) and "Cancel" (in grey).

Enter the version tag for the version to update for all container images in the dialog box and click **Update**. If you leave the version tag field blank, the container will be updated to the latest version of the image stored in the Docker registry you are using.

This is a live update. Your existing replica sets continues to receive traffic until the replica sets with the new images are fully functional. Once that happens, traffic is redirected to the new replica sets. During the process, the status of your deployment changes to Updating until the update is complete, at which point the deployment status returns to deployed.



Workload Manager does not do a validity check for your version tag before attempting to perform the update procedure. If you enter a version tag that does not exist in the Docker registry you are referring to (either Docker Hub or your own private Docker registry). Even if the update action fails, the Container deployments continue to be restored with the previous image – before the rolling update action is initiated.

If you see the *Unable to update container images on Deployment Details page* error, go to Deployments list page and select Update from the actions dropdown menu.

Advanced Configuration Topics

Advanced Configuration Topics

- [Guidance for Callout Scripts](#)
- [Advanced Configuration using Kubernetes ConfigMaps](#)
- [Limiting Acceptance of Self-Signed Certificates for HTTPS Web Services](#)
- [Update Self-Signed Certificates for Guacamole](#)
- [Adjusting Action Message Purge Frequency](#)
- [Password-Based Authentication](#)
- [Custom Docker Image for Scripts](#)

Guidance for Callout Scripts

Guidance for Callout Scripts

- [Overview](#)
- [VM Naming Script](#)
 - [VM Naming Script Supported Clouds and Cloud Nuances](#)
 - [VM_NODE_INDEX](#)
 - [Sample VM Naming Callout Script](#)
- [IPAM Callout Scripts](#)
 - [IPAM Callout Script Supported Clouds and Cloud Nuances](#)
 - [IPAM Address Allocation Callout Output Parameters](#)
 - [Sample IPAM Callout Script for Single-NIC Scenarios](#)
 - [Sample IPAM Callout Script for Multi-NIC Scenarios](#)
 - [Sample IPAM Script which Returns IPV6 IP Address](#)
 - [Callout Script OS-Specific Output Parameters](#)
 - [Alternate Windows Guest OS Customization \(vCenter only\)](#)

Callout scripts may be used for VM naming and for IPAM address allocation and deallocation for a particular VM-based cloud region. These as specified in the Strategy section of the cloud's Details tab or Region's tab

Callout scripts can be shell executable scripts or Python scripts; however, Python scripts must be embedded in a wrapper shell script. Scripts can refer to Workload Manager environmental variables as input. Scripts must also output their results in **JSON** format as a series of key-value pairs. Callout scripts must also include the wrapper utility named **utils.sh** and use the **print_ext_service_result** command which encloses the output in a Workload Manager accepted format, as shown in the examples below. Certain output variables are expected of VM naming scripts and IPAM address allocation scripts as explained below.

The VM naming script is called before each node is launched. It is provided (injected into the script) with all the name variables (name of application, name of tier, image selected) for each job. It must create a string of up to 15 characters and assign that string to the reserved Workload Manager output field: **vmName**. The allowed characters of the VM name are as specified the **Node Name Config** field above.

VM Naming Script Supported Clouds and Cloud Nuances

VM naming scripts are supported on all VM-based clouds. The following cloud nuances apply:

Cloud Provider	Nuances
vCenter	The <i>Hostname Callout</i> option in the Instance Naming Strategy dropdown sets the osHostname and vmName inside the guest OS. These two settings are the same: <ul style="list-style-type: none"> • As the vCenter settings • For Linux (CentOS7)
AWS	The <i>Hostname Callout</i> option in the Instance Naming Strategy dropdown sets the osHostname and vmName inside the guest OS. These two settings are the same for both Windows and Linux.
Azure	VM naming callout script support an optional osHostname as output variable in addition to the vmName . If osHostname is specified, it is used to set the hostname, or else vmName is used to set the hostname. This is applicable for both Linux and Windows. For Windows, the osHostname should not be more than 15 characters.

Typical Workload Manager environment variables accessed

Variable	Sample value or type
eNV_JOB_ID	integer (application VM only)
eNV_launchUserId	integer
eNV_launchUserName	string

VM_NODE_INDEX

Supported Clouds: AWS, VMware, Google, AzureRM, OpenStack, IBM, and VCD

- The environment variable VM_NODE_INDEX is passed to VM naming callout scripts with a unique value – when there are multiple nodes being deployed or scaled up.
- For each tier, the VM_NODE_INDEX starts from 1.
- For example, if 3 nodes are deployed in a tier, then 1, 2, and 3 are sent respectively as environment variable VM_NODE_INDEX for the corresponding VM naming callouts.
- During scale-up (if you add 2 more VMs), 4 and 5 are sent respectively as environment variable VM_NODE_INDEX for corresponding VM naming callouts.

Supported Clouds: AWS, VMware, Google, and OpenStack

- The VM_NODE_INDEX environment variable is passed to IPAM callouts (alloc(create vm) and dealloc (terminate vm)) with unique values when multiple nodes are deployed or scaled up.
- For each tier, VM_NODE_INDEX starts from 1.
- If 3 nodes are deployed in a tier, then 1, 2, and 3 are respectively sent as environment variable VM_NODE_INDEX for the corresponding IPAM callouts.
- During scale-up (if you add 2 VMs), 4 and 5 are sent respectively as environment variable VM_NODE_INDEX for corresponding for IPAM callouts.
- During scale down/terminate, the environment variable VM_NODE_INDEX is passed as the same value, which is used during scale-up /deployment, to corresponding IPAM callouts.

Sample VM Naming Callout Script

run.sh

```
#!/bin/bash

. /utils.sh

# Install uuidgen
apk add --no-cache util-linux

# Creating custom vm name/hostname
custom_tag="ccqa-"
vmName="$custom_tag$(uuidgen | fold -w 8 | head -1 | tr '[:upper:]' '[:lower:]')"
#vmName="$custom_tag$((($RANDOM + ($RANDOM % 2) * 32768)))"
osHostName="$custom_tag$(uuidgen | fold -w 8 | head -1 | tr '[:upper:]' '[:lower:]')" //Supported for Azure//
content="{ 'vmName': '$vmName', 'osHostName': '$osHostName' }";
# Print the results
print_ext_service_result "$content"

# Recording IPAM execution and env variables
curl -s -o /tmp/record_results.sh http://http.cliqrtech.com/auto/record_results.sh
chmod 755 /tmp/record_results.sh
source /tmp/record_results.sh callout-vmNaming $vmName > /dev/null 2>&1
rm -fr /tmp/record_results.sh
```

There two types of IPAM callout scripts: IPAM address allocation scripts, for assigning IP addresses to vNICs on a VM; and IPAM address deallocation scripts, for freeing up used IP addresses when a VM is terminated.

IPAM Callout Script Supported Clouds and Cloud Nuances

IPAM callout scripts are supported on all VM-based clouds except AzureRM. The following cloud nuances apply to the rest.

Cloud Provider	Nuances
vCenter	<p>The osHostname setting:</p> <ul style="list-style-type: none"> • Not mandatory for IPAM callouts. • Works for Windows only. • Linux setting is overwritten by the vmNaming setting.
AWS	Workload Manager uses the IP address , network , and mask to set the DHCP scope in the specified subnet.

IPAM Address Allocation Callout Output Parameters

All IPAM address allocation scripts must output their data in JSON formatted key:value pairs where the key name corresponds to one of these well known key names in the following table.

Key	Description	Required?
osHostname	OS hostname	Yes. <i>Not supported on AWS and OpenStack.</i>
DnsServer List	DNS server list (comma separated)	

DnsSuffixList	DNS Suffix list (comma separated).	
nicCount	The number of virtual NICs (vNICs).	No.
nicIP_0	vNIC IP address	Yes Set on a per-NIC basis. A new IP address must be set for each vNIC.
nicNetmask_0	vNIC netmask	Yes Set on a per-NIC basis. A new netmask must be set for each vNIC.
nicGateway_0	vNIC gateway IP address	No Set on a per-NIC basis.
nicDnsServerList_0	vNIC DNS server list (comma separated)	Yes
nicUseDhcp_0	Set to TRUE if DHCP should be used for assigning an IPv4 address to this vNIC	Yes, if using IPAM callout and the addressing is assigned to use DHCP.
nicIPv6_0	IPv6 IP address	Yes Set on a per-NIC basis. A new netmask must be set for each NIC.
nicGatewayIPv6_0	IPv6 gateway IP address	Yes Set on a per-NIC basis. A new netmask must be set for each NIC.
nicNetmaskIPv6_0	IPv6 netmask	Yes Set on a per-NIC basis. A new netmask must be set for each NIC.
nicUseDhcpIPv6_0	Set to TRUE if DHCP should be used for assigning an IPv6 address to this vNIC	If nicUseDhcpIPv6_0 is set to true and static IP information is also provided, DHCP takes precedence over STATIC allocation strategy.
Custom	Example key:value pair output: " {<myCustomParam>:<myValue>}	Custom IPAM Callout variables are not set in the userenv file on target deployment VMs.



If the VM configuration includes multiple NICs, Workload Manager makes one IPAM call per NIC. You can assign multiple IPs to each NIC by using keys with `_n` suffix as described earlier.

Sample IPAM Allocation Scripts

Sample IPAM Callout Script for Single-NIC Scenarios

A single-NIC script is executed once and the `nicIP_0` (for example) value is set to the first interface of the VM. The IPAM script is executed once for every NIC. If the VM has n NICs, the same IPAM script is triggered n times.

Single NIC Script

```
#!/bin/bash

. /utils.sh
content="{ 'DnsServerList' : '8.8.8.8,10.0.0.100', 'nicIP_0' :
'10.0.0.100','nicDnsServerList_0' : '1.2.3.4,5.6.7.8','nicCount': '1',
'nicGateway_0':'10.0.0.1','nicNetmask_0':'255.255.255.0','domainName':'test.
org','hwClockUTC':'true','timeZone':'Canada/Eastern','osHostname':'testhost1'}"

print_ext_service_result "$content"
```

Sample IPAM Callout Script for Multi-NIC Scenarios

A multi-NIC script is executed for each NIC in your VM. In a multi-NIC scenario, the single-NIC script is called multiple times corresponding to the number of NICs defined in your Workload Manager deployment. The IPAM script is executed once for every NIC. If the VM has n NICs, the same IPAM script is triggered n times. An OOB parameter `nicIndex`, starts from 1 and increments when the script is called for each NIC.

For each execution of this multi-NIC script, a new `nicIP_0` value is set to the corresponding interface of your VM.

Other than changing the `nicIP_0` value, you can also change the values for all other parameters – other than `nicCount` (which is always set to 1)

Multi-NIC Script

```
#!/bin/bash

. /utils.sh
content="{ 'DnsServerList' : '8.8.8.8,10.0.0.100', 'nicIP_0' :
'10.0.0.100','nicDnsServerList_0' : '1.2.3.4,5.6.7.8','nicCount': '1',
'nicGateway_0':'10.0.0.1','nicNetmask_0':'255.255.255.0','domainName':'test.
org','hwClockUTC':'true','timeZone':'Canada/Eastern','osHostname':'testhost1'}"

print_ext_service_result "$content"
```

Sample IPAM Script which Returns IPV6 IP Address

When you assign IPv6 addresses, Workload Manager validates the security rule source before accepting the IPv6 address. See [IP Allocation Mode > Cloud-Specific Nuances > IPv6](#) Note.

Working Script for IPv6 Allocation

```
#!/bin/bash


. /utils.sh
count=1
#have you logic for maintaining the count
content="{ 'DnsServerList' : '8.8.8.8,10.0.0.100', 'nicIP_0' : '###.###.###.
###.$count','nicDnsServerList_0' : '1.2.3.4,5.6.7.8','nicCount': '1', 'nicGateway_0':'###.
###.###.###','nicNetmask_0':'255.255.255.0','nicUseDhcp':'true','nicIPv6_0':'2600:1f14:5aa:
2f00:524a:fbf5:3377:a$count'}"

print_ext_service_result "$content"
```

Sample IPAM Deallocation Script	This script provides notification only, no output.
	<pre>run.sh #!/bin/bash ./delete_record_by_ip.sh \$IP</pre>

Callout Script OS-Specific Output Parameters

The following table shows the multiple key-value pair that is output for each callout script.

OS Properties	Linux	Windows	Required?
timezone	<i>Supported</i> for VMware. <i>Not supported</i> for AWS and OpenStack.	Not used	Yes
timeZoneId	Not used	The Windows Index ID for this time zone. <div style="border: 1px solid #ffc107; padding: 5px; margin: 5px 0;">  For Windows-specific VMware IPAM config scripts, you may see the changes after the deployment has completed. </div> AWS: No effect as instance timing is internally managed.	Yes
fullName	Not required	The name of the Admin user	Yes
organization	Not required	The name of the organization (string)	Yes
productKey	Not required	The Windows product key	Yes
setAdminPassword	Not required	The Admin password	Yes
changeSid	Not used	A true or false value for the Microsoft SID You must set the changeSid option to true .	Yes
deleteAccounts	Not used	A true or false value.	Yes
dynamicPropertyName	Not used	Reserved name holder for arbitrary property	Yes
dynamicPropertyValue	Not used	Reserved value holder for arbitrary property	Yes
custSpec (see the <i>VMware Customization Spec</i> section below)	The Guest Customization Specification name in VMware	The Guest Customization Specification name in VMware	No
hwClockUTC	<i>Supported</i> for VMware – Identifies if the hardware clock should follow UTC or local time. <i>True</i> = UTC <i>False</i> = Local time	<i>Not supported</i> .	Yes
domainName	Used for FQDN resolution of Linux VMs as it is visible when using <code>hostname -f</code> or <code>cat /etc/hosts</code>	Used to automatically join a domain – Only supported for VMware. Not supported on AWS and OpenStack.	No
domainAdminName	Not used	Used to automatically joining a domain	
domainAdminPassword	Not used	Used to automatically joining a domain	
workgroup	Not used	The workgroup in which to place the VM. If any of the 3 domain values are missing, the workgroup key is required. If all three domain values are present, the workgroup is not required.	

Alternate Windows Guest OS Customization (vCenter only)

SysPrep is a tool that is executed to customize Windows deployments. Windows IPAM optimization allows you to skip the *SysPrep* execution for Windows deployments.

Running *SysPrep* to customize Windows may affect performance. You can bypass the *SysPrep* by providing IPAM properties in the corresponding callout.



The *SysPrep* process will be triggered when properties in the following table are *returned* by the callout script.

The following table lists OS-Specific IPAM Properties for alternate Windows guest OS customization.

OS-Specific IPAM Properties*
changeSid
deleteAccounts
fullName
organization
timezone
setAdminPassword
domainName
domainAdminName
domainAdminPassword
workgroup

* These properties are described in the [OS-Specific Parameters](#) section.

Advanced Configuration using Kubernetes ConfigMaps

Advanced Configuration using Kubernetes ConfigMaps

- [Overview](#)
- [List of ConfigMaps](#)
- [ConfigMap Properties](#)
 - [cloudcenter-ccm-backend](#)
 - [cloudcenter-cco](#)
 - [cloudcenter-cloud-setup](#)
 - [cloudcenter-ces](#)
 - [cloudcenter-blade-{cloud-family}-{region-id}-{random-4-digit}](#)

Several of the Workload Manager services running in the CloudCenter Suite cluster can be configured by editing their corresponding ConfigMap. An example of how to do this is shown in [Limiting Acceptance of Self-Signed Certificates for HTTPS Web Services](#). This section is a reference for the various Workload Manager ConfigMaps and their corresponding properties, usage, and default values.

Kubernetes Service	ConfigMap Name
cloudcenter-ccm-backend	cloudcenter-manager
cloudcenter-cco	cloudcenter-orchestrator
cloudcenter-cloud-setup	cloud-setup
cloudcenter-ces	cloudcenter-ces
cloudcenter-blade-{cloud-family}-{region-id}-{random-4-digit}	cloudcenter-blade-{cloud-family}-{region-id}-{random-4-digit}

The lists of properties for each service's ConfigMap follows below.

cloudcenter-ccm-backend

Property Name	Type	Default Property Value	Description
external.hosts	String	host:port	To import certificate for a given list of external hosts
allow.self.signed.certs	Boolean	true	If "true" then allow all self signed certs while calling https service, if "false" then allow certs only when they are valid
cloud.packages.url	String	http://repo.cliqrtech.com/cloud-packages/version/@project.version@	
synchronizelImagesOnStartup	Boolean	false	"true" if synchronization of images is allowed on startup "false" if not allowed.
autoStorageMigration	Boolean	true	"true" if storage migration is auto enabled, "false" if not allowed.
use.summary.report.view	Boolean	false	"true" if summary report view can be used, "false" if cannot be used.
minio.defaultVendorPath	String	assets/default/vendor	default path of minio vendor.
minio.defaultVendorAssetsPath	String	assets/default/vendor/content	default path of minio vendor assets.
minio.appLogoPath	String	assets/img/appLogo	path of minio application logo
minio.serviceLogoPath	String	assets/img/serviceLogo	path of minio service logo

cloudcenter-cco

Property Name	Type	Default Property Value	Description
bootstrap.waittime	Integer	15	Time gateway waits before performing bootstrapping
bundle.store.url	String	http://build-rel.cliqr.com/	The base url which contains all the agent bundles
custom.repository.url	String	http://repo.cliqrtech.com	optional path to the repo server containing all the .deb and .rpm packages
maximum.bootstrap.wait.time	Integer	3600	maximum waiting time for bootstrapping

node.heartbeat.time	Integer	180000	heartbeat time for a node
node.cleanup.timeout	Integer	300	timeout age for a node cleanup
node.ready.timeout	Integer	3600	timeout age for a node ready
docker.container.scriptTimeoutDuration	String	10m	timeout for scripts which run inside docker container. This is to prevent any scripts from running forever
default.auto.scale.percentage	Integer	70	Default auto scale percentage for auto scaling
valid.metric.result.time.period.multify.factor	Integer	3	Time period needed to consider a metric result valid as a factor of the policy polling interval, e.g: 3 times the polling interval
lifecycle.tier.terminationInParallel	Boolean	true	By default, all tiers are terminated in parallel. When it is set to false, the termination would follow the reverse of dependency graph
allow.self.signed.certs	Boolean	true	If "true" then allow all self signed certs while calling https service, if "false" then allow certs only when they are valid

cloudcenter-cloud-setup

Property Name	Type	Default Property Value	Description
cloud.packages.url	String	http://repo.cliqrtech.com/cloud-packages/version/5.0.0/	path of cloud packages
cb.worker.image.repo	String	devhub-docker.cisco.com/cloudcenter-dev-docker/cliqrimages/cloudcenter	path of cb worker image repository.
c2healthcheck.enabled	Boolean	true	"true" if c2 health check is enabled,"false" if not enabled.

cloudcenter-ces

Property Name	Type	Default Property Value	Description
ces.global.script.timeout	String	10m	timeout age for ces global script in minutes

cloudcenter-blade-{cloud-family}-{region-id}-{random-4-digit}

Property Name	Type	Default Property Value	Description
allow.self.signed.certs	Boolean	true	If "true" then allow all self signed certs while calling https service, if "false" then allow certs only when they are valid
vmware.thread.pool.size	Integer	30	Value specified is the max concurrency limit when creating virtual machines
vmware.usePropertyCollector	Boolean	false	If "true", cloud properties collection is done by Property Collector api, If "false", cloud properties collection is done by getting all managed objects and iterating parent and child objects

Limiting Acceptance of Self-Signed Certificates for HTTPS Web Services

Limiting Acceptance of Self-Signed Certificates for HTTPS Web Services

- [Overview](#)
- [Procedure](#)

Workload Manager has the ability to initiate connections to web services via the HTTPS protocol. These web service calls are invoked in the following use cases:

- Configuring [ServiceNow extensions](#)
- [On-demand custom actions](#) that call a web service
- [Deployment parameters](#) that are of type *webservice* that are defined at the [service](#) level or application tier level.
- [Service parameters](#) that are of type *webservice* that are defined at the [service](#) level.
- [Global parameters](#) that are of type *webservice* that are defined at the application level.

Web service calls from Workload Manager to web sites that do not have a certificate authority signed certificate **will succeed by default**. As a CloudCenter Suite administrator, you can configure CloudCenter Suite to reject self-signed certificates but with exceptions for certain host:port combinations. This is done by editing the [Kubernetes configmap](#) for the cloudcenter-manager pod.

The procedure for allowing calls to only certain web services that use self-signed certificates is as follows.

1. Install kubectl on your computer, [download the CloudCenter Suite provided Kubeconfig file](#), then move the kubeconfig file to the directory specified in the \$Kubeconfig environmental variable. This allows you to connect to the CloudCenter Suite cluster with kubectl from your computer.
2. From the command prompt of your computer, use kubectl to ensure your CloudCenter Suite instance has a configmap for the cloudcenter-manager pod:

```
kubectl get configmaps -n cisco
```

3. Once confirmed, use kubectl to launch your default editor to edit the cloudcenter-manager configmap:

```
kubectl edit configmaps cloudcenter-manager -n cisco
```

Your editor should display the configmap in edit mode to allow you to modify it:

```
apiVersion: v1
data:
  external.hosts: "hostA:portA"
kind: ConfigMap
metadata:
  creationTimestamp: 2018-12-19T15:31:04Z
  labels:
    app: cloudcenter-ccm-backend-5.0.0
    chart: cloudcenter-ccm-backend-5.0.0
    heritage: Tiller
    purpose: configuration
    release: workload-manager
  name: cloudcenter-manager
  namespace: cisco
  resourceVersion: "18309514"
  selfLink: /api/v1/namespaces/cisco/configmaps/cloudcenter-manager
  uid: 18c30efe-03a3-11e9-bd86-42010a80004a
```

4. Using your editor:
 - a. Edit the **external.hosts** property: Replace the default value with a comma separated list of host:port combinations of the services you want to allow enclosed in double quotes.
 - b. Create a property called **allow.self.signed.certs** and set the value to **false**. Insert this right after the **external.hosts** property.
 - c. Save the file and exit your editor.

Update Self-Signed Certificates for Guacamole

Update Self-Signed Certificates for Guacamole

- [Overview](#)
- [Requirements](#)
- [Procedure](#)
- [Update Cloud Remote Environments](#)

This feature is useful when you receive a certificate error (similar to the error in the following screenshot) when attempting to use the SSH button in the Deployments tab to reach a VM. At this point, you may need to replace the certificate using the procedure provide later in this section.

The user performing this procedure must have `cluster.admin` privileges to the Kubernetes cluster.

To update the certificates for the Guacamole service, follow this procedure.

1. Create a Kubernetes resource using the certificate and .key files.

```
$ cat private.key | base64
LS0tLS1CRUdJTiBSU0EgUFJVVkFURSBURVktLS0tLQpNSU1FcFFJQkFBSONBUUVBc09yNjVWnmprN0pacWxZOFNYSitBZ1FONitaUzRMC
lozcldYbJz6V3J5WlNlc0drCktZRHVSQytEKzYyQ3JVWjzCTEVkNmXkcmqMmVFRXFuc2hDNjk5ODRGTnF1clJjd1k2eVRHaGFwdXhFak
lONHAKZU4vWk1oeUVEb0pNQXJUS8yZG9Fdf1ST0xju2k1VWk4K1V4bzJOM2F5dkQ0Q1pZb3NvdCtoRjg0TUtVNTZuWApQYnRuVlk0L2t
heHZPbXZmUmh5eno3K2ZkamlnTlDnMTJGcVp4NFBzMNBPd2hFUGJzakZxVUuzYXYxTjg1bk1zCmZ0b0xHRkhGdG1ac21pRVZBa1VjMVUr
ZmJqaUt5VVNrzcLDyZiza1NhoXN6Zk1qrW5YnNB2NzdQOErVYjBfaUwKbGVTR3pIbXlnMnN4c3dydlVBdmMxdHNMM0w1Y1ZWWwCvQmt6W
ndJREFRQUJb0lCQVFDdUdqQ1o5SzVXTTNnBa02OE5jVEMydzRtBw8wbnFRNm5bG1iTGkwZ1piU3cwei9NZitiPQR5WFFnM2J1TUFmak
NwM1ZzbE9HNTJ0d2tMCjQwDndNZ0tzMTZdCtLr2c2TDbY0Thhemo4WHNiOHIVmWtGZDBIdy9ra3hxc3RMmUrSysxUlpZSXczRgtIWF1
Kb3l4SQRDNUNXZ0ROR0w2a0VabWnhZmFnOXDdn2RWMHlQnnpOTzgd2dU1XbWpsdnRMSVJVMHhNwMh1VXBROetoOq0QkpEemJBUUREQ1FM
bU2WtYtYSHVzYncxQldWZUNBWEhBdlMvamNJa1NxxVWV4djlNUdM2OE5Xb3BkNFRRRk82CmVXRnZ6cmJlQWHzcUk1K2RrUVrLb3MvbmZoc
FY3clFQSG0vM1ZPcWtEaWtPTT1lcXU1bGUvQWk2cjhFM2NCclIKZZJ10W9sYkjb0dCQU9iSGt2Y2Y0aFzScdVTUHZ4UmJlSk1qaFE2ND
AxMDQ5aU1wcjFEdlpVbVZHUvQ1k4agpTYTRtTkNdcWRvR1VaOTR2Z3EYqMjQ2Q1E4RysvdVdxbmpuSHZNQ0JPL3ZISHRKvzBna2s1Wws
3U1lraJNNeTBSCKiYUEZ5Mm0YVpaZnPNYTrzWm04UENEbjZBL25HMmE3L2YrMiTRnk12YTNvdmxMNwVLaW5sa3R4QW9HQkFNuKELaWhZ
Q1RKTtVZa1JTK0p2TEdCcVpDY1o4vzNTNmlsb012bjJNWG41exZITzhveGRKSjFBZVQwek1DYm5PdUtHZApndHZNSW5iv1ZvVE5ib1ZpS
WJ4Y0FUNlLnMnc5WGxnWGRqL01nL3VZME5pYUZKQjBKL1NCRXcVdkdOSGZsvUN0CkswaU15SjBnrjdValhJtGMvNkZWOENONEdoQ2Ivb0
RjcVhsdjlwQ1hBb0dCQUtCT2VaUVlIwi9aZktNqnh5V0k0UpQdkfIcGRnTLQ4d0Yxa2sxZVJNN2FOYyt2MULSTncyN21Rnkj6WxpFRHV
xy2Y1RUxrZGM4Y9S9wNFM0bFQRMA03S5W5RTUlvQTFHOfTfucVc0SWRoamVCChdvYjAxwVwM05VWgXhNG16OGdsMndGS0M4VHR0T2dkUmtD
MmJ4N2p4C1Zwc1A2dWxrY2szZm9uV1l2YXZKM2VqQkFvR0FQMG9NL21JsnJlVhdvdFliSktydmdBeGdrYwTqR1ziYkxVVTAKM1dvn1F3O
GZaVGv0Wk9JTGtWUGp3UUDBRlhMc0tmcVE2KzgrSWHmblhmYVWLNjDVOEZpL2NnZWxlCUVuM3NMYQpPS0dpcHEzV2xEUEVjLzF4clRFN0
EOVjQvsjNkR1RtRFh1Yy8veDjhTzZsR3VWRXFbMDZpbmQycWtqYXNjY1EzCngxanJMqXNDz1lFQWlUOWRGY0NCUKRuNjfxUkg4Tke2NUd
iZEgrb1FZN1VVL3V0S21Y0XV2OE3ZXhTeW56S1YKMVBmNTFOeUFOVzlaZ2N0UjJtVWNxcno4cW5taEFMeFgzS05oeDRJaS9yRTAWdTFr
NVVpTghiUnNISzVtQ3NaOqPlbTNzS240WTVuM1dhTUI5RfLgd1NwL0tRZ1RRU1FUdmw4N05FSwX5NTdHeGdpLy9DQXg2dVMwPQotLS0tL
UVORCBSU0EgUFJVVkFURSBURVktLS0tLQo=
$ cat custom.crt | base64
LS0tLS1CRUdJTiBDRVJUSUZJQ0FURS0tLS0tck1JSUViakNDQTFhZ0F3SUJZB01KQU82dk5NZEecwG5TTUEwR0NtCudTSWIzRFFQkJRv
UfNSudbTVFzd0NRWUQkVlFRROV3S1ZVekVPTUF3R0ExVUVDQk1GVkdWNF1YTXhfREFPQmdOVkYjY1RCMhZ2ZFHOMGIyNHhFekFSQmdOVg
pCQW9UQ2sXNULFTnZiWEJoYm5reZEZEVQVNCZ05WQkFzVEMvVnVaMmx1WldWeWFXNW5U1F3SWdZRFZRUURFeHR2CmJtSnZWEprTGpFNU1
pNHhOamd1T1RrdU5pNXVhWEF1Yvc4d0hoY05NVGN4TURnd0lUTTFnekKxv2hjTkl1UZ3gkTURnd0lUTTFnekKxv2pDQmdERUxNQWtHQTFV
RjUoTUNwVkl1RGPBTUJN1lZCQWduQ1ZSbGVHRnpNukF3RjGZRAPWUVFIRXdkSWIzVnpkRz11TVJNd0VRWURWUVFLRXdWtmVTQkRiMjF3W
Vc1NUIsUXdfZ11EVLFRTEV3deZibWRwCmJtVmXjbWxlWnpFa01DSudBMVVFQXhNyMyNW1iMkZ5WkM0eE9USXVNVFk0TGprNUxqXWVibV
x3TG1sdK1JSUIKSwpBTkna3Foa21hOXcwQkFRRUZBQU9DQVE4QU1JUSUJDZ0tDQVFFQXNpcjY1VjZqZ0dKWnF5WThTWEorQWZRTgo2K1p
TNEzXWjYnV1huNnpXcnlaU2VzR2tLWURlUkMrRCs2MkNyVVo2QkxZDZS2hJnaJlRUVxnbNoQzY50Tg0CkZocXVzUmN3WTZ5VEdoYXB1
eEVqtU40cGVOLpNaHlFRG9KTUFySFEvMmRvRrXZUk9MY1NpNVVpOCTVeG8ydmkYXl2RDRCw1lvc290K2hgODRNS1U1Nm5YUGJ0b1ZZN
C9rYXh2T21Z2lJ0eXp6NytmZGptZ05XZzEYRnFaEDRQcwoyce93aEVQYnNqRnFVVTNhdjFODVUuTXNmDG9MR0ZIRnRpnNtaUVWQWpVYz
FVKZ2iaamlLeVvTa3M5Q2MyM2tTCmE5c3pmSWpfb1g2cHY3N1A4RG9IMEVpTGx1U0d6SG15ZzJzeHN3cnZVQXZjMxRzTDNMNNWV1lnL0J
relp3SUQKQVFBqM80SG9NSuhsTUIwR0ExVWREZ1FXQkJtCTJib1RaMGRjkzVFN1ZWtWwrZmJicGhsaEFUQ0J0Uv1Ev1IwagpCSud0TU1H
cWdCU3EyYm5UWjBkYys1RTZWVklSk2ZiSHBobGHbYudCaHFTQmd6Q0JnREVMtUFrR0ExVUVCaElDClZWTXhEakFNQmdOVkYjBz1RCVlJzS
UdGek1SQxdEZ11EVLFRSEV3ZELim1Z6Zec5dU1STXdfUV1EVLFRS0V3cE4KZVNCrGIyMxdZVzU1TVJRd0VnWURWUVFMRxd0RmJtZHBibV
ZsY21sdVp6RwtNQ0lHQTFVRUF4TWJ1mJvYjJGepQaQzR4T1RjDU1UWTRMams1TgPzdWjtBdMbwX2Z2drQtdXODB4MGJTBwrJd0RBWUR
WUjBUQkFv0F3RU1IvekFOckJna3Foa21hOXcwQkFRVU9DQVFFQXJtdUUVeOckH0YzZLdmpxbjJmY0Z5TGZxMWRqQTJWZXPJU2FmUhhx
WWUKd1pJenc2SkxOaUI3MVRibk1YWVp3UGpNTXVBOGJYc0ZHR0wzQzRGRz12NGIvYzd0V2hsWks3R1ZucjYyQ0Q2dgpss3pLeG5ZT2E0M
VN4UzdRdG9RNOtdYkc5K0ZPeWhaUjIzS2hPv0UrQ0p1T09MVE5XUW1QVlRkRWhuWmVeaW1BclhDWE1DNHE1a25UmNTK1ZXRWUTjrhd2
R2bUvSeSs2bXlTNksZQt11VTQvaJQ5dlRwcWNUVTh0LzI4SzhTQUMKewX1QkF3RG1pRURFU1JzUmVDWVBhanZ2Tmx2c2xvUElJUVWwXTTA
vN2V1bjzkaU96R3ZueGfKN01tZ29Ody9hNgpsQm5pYW1OaXNCEjViU25Gw1pLbXNDWV1oN25teFJsODNBZWRoN3VCNW05S1hBPT0KLS0t
LS1FTkQgQOVSVB1GSUNBEUuTLS0tLQo=
```

2. Create certificates for the Guacamole service.

```

apiVersion: v1
kind: Secret
metadata:
  name: custom-ca-tls
  namespace: default
type: kubernetes.io/tls
data:
  tls.crt:LS0tLS1CRUdJTiBDRVJUSUZJQ0FURSB0tLS0tCk1JSUViakNDQTFhZ0F3SUJBZ0
  tls.key:LS0tLS1CRUdJTiBSU0EgUFJJVkFURSBLRVktLS0tLQpNSU1FcFFJQkFBS0


```

3. Using the above secret Name, edit the deployment YAML file.

```

volumes:
  - name: guacamole-tls
    secret:
      defaultMode: 420
      items:
        - key: tls.crt
          path: certificate.pem
        - key: tls.key
          path: private_key.pem
      secretName: custom-ca-tls

```

 The cloud region, where the Cloud Remote is installed, must be in a RUNNING state.

To update self-signed certificates for Guacamole in Cloud Remote environments, following this procedure.

1. Create a keystore.jks file and add the key and both certificates (the actual certificate and the CA certificate) to this file.
2. Create a truststore.jks file and only add the CA certificate to this file.
3. Copy these .jks files to the Cloud Remote VM.
4. SSH into the Cloud Remote VM.
5. Change to the following directory.

```
cd /opt/cisco/pilot/builds/pilot_5.1.0-XXXXXXXXXX/bin/
```

6. Run the command to update the self-signed certificate.

```
./update_guac_certs.sh <path to keystore.jks> <path to truststore.jks> <password of keystore and truststore>
```

You have now updated the self-signed certificates for Guacamole in Cloud Remote environments.

Adjusting Action Message Purge Frequency

Adjusting Action Message Purge Frequency

- [Overview](#)
- [Procedure](#)

By default, Workload Manager will retain the last 90 days worth of action messages related to actions executed again a deployment or a VM. You can override this default value by adding an `action.message.purge.frequency` key-value pair to the `cloudcenter-manager` configmap.

The procedure for adjusting the action message purge frequency is as follows.

1. Install `kubectl` on your computer, [download the CloudCenter Suite provided Kubeconfig file](#), then move the kubeconfig file to the directory specified in the `$KUBECONFIG` environmental variable. This allows you to connect to the CloudCenter Suite cluster with `kubectl` from you computer.
2. From the command prompt of your computer, use `kubectl` to ensure your CloudCenter Suite instance has a configmap for the `cloudcenter-manager` pod:

```
kubectl get configmaps -n cisco
```

3. Once confirmed, use `kubectl` to launch your default editor to edit the `cloudcenter-manager` configmap:

```
kubectl edit configmaps cloudcenter-manager -n cisco
```

Your editor should display the configmap in edit mode to allow you to modify it:

```
apiVersion: v1
data:
  external.hosts: "hostA:portA"
kind: ConfigMap
metadata:
  creationTimestamp: 2019-08-19T15:31:04Z
  labels:
    app: cloudcenter-ccm-backend-5.1.0
    chart: cloudcenter-ccm-backend-5.1.0
    heritage: Tiller
    purpose: configuration
    release: workload-manager
  name: cloudcenter-manager
  namespace: cisco
  resourceVersion: "18309514"
  selfLink: /api/v1/namespaces/cisco/configmaps/cloudcenter-manager
  uid: 18c30efe-03a3-11e9-bd86-42010a80004a
```

4. Using your editor:
 - a. Insert a line after the `external.hosts` property.
 - b. In that new line, create a property called **`action.message.purge.frequency`** and set the value to the number of days that action messages should be retained. For example, to set the purge frequency to 30 days, insert this text:

```
action.message.purge.frequency: "30"
```

- c. Save the file and exit your editor.

Password-Based Authentication

Password-Based Authentication

- [The Default cliqruser Account and Authentication](#)
- [Override the Default Configuration](#)



This procedure must be performed by an Dev Ops or Enterprise Admin with an intricate understanding of your security architecture and your IT infrastructure.

Cliqruser.

- Is the default user to log into the Application VM from the Workload Manager UI – Contact the [CloudCenter Support team](#) for details on a [Custom Windows Image \(Installer\)](#).
- Refers to an OS user in the Application VM (Worker) – see [Deployment Lifecycle Scripts](#) for additional context.

By default, key-based authentication is configured using *cliqruser*.

If you do not want to use the default cliqruser account, you can alternately use password-based authentication to access Application VMs on a per-tenant basis.



You cannot configure password-based authentication from the CCM UI. You can only do so from the CCM API as described in this section.

To configure password-based authentication, follow this procedure.

1. Use the following APIs to retrieve some basic information:



You can override the default property as a ROOT Admin or a Tenant Admin. Either way, your login credentials determine if you are an admin (platform (root), tenant admin, or co-admin) or a user.

- a. Use the *View Tenants* API for your current tenant to determine your Tenant ID:
 - i. Provide the additional *vendorId* for a ROOT admin
 - ii. Just the *tenantId* is sufficient for a Tenant Admin
 - b. Use the *View All Tenant Properties* API to determine the *propertyId* for the required Property name (*propertyName*) for this password-based authentication (**enable.password.auth**) setting, which has an ID of **132**.
2. Use the *Create Tenant Properties Override* API to override the default (false) setting for the **enable.password.auth** property name with a *propertyId* of **132**.

POST API Request

```
https://<HOST>:<PORT>/v1/tenantProperties/override
```

ROOT Admin – Request Body

```
{
  "propertyId": "132",
  "vendorId": "2",
  "propertyValue": "true"
}
```

Tenant Admin – Request Body

```
{
  "propertyId": "132",
  "propertyValue": "true"
}
```

You have now overridden the default *cliqruser* account usage in order to use password-based authentication to access Application VMs on a per-tenant basis.

Custom Docker Image for Scripts

Custom Docker Image for Scripts

- [Overview](#)
- [Prerequisites](#)
- [Customize the Docker Image](#)
- [Update the configMap Image set](#)
- [Additional Considerations for Cloud Remote](#)

CloudCenter Suite provides a base [Docker Service](#) image to execute callouts on any Workload Manager-supported cloud.

This section provides details on the custom worker image for callouts/scripts if you prefer to execute external callout scripts in an isolated Docker container.

- Create a public or private docker registry to host the new Docker worker image.
- This Docker registry must be reachable from the CloudCenter Suite cluster
- If the registry requires authentication, create a Kubernetes secret.

```
kubectl create secret docker-registry custom-image-pull-secret
--docker-username=<user> --docker-password=<pwd>
--docker-email=custom@image.com --docker-server=devhub.com -n
<namespace>
```

To customize the Docker image, follow this procedure.

1. Access the Local Bundle Store (Conditional) and verify that you have access to the software.cisco.com.
2. Log into any server that has access to custom Docker image.
3. Download the file and create a worker:release-5.1.0 docker image in your Docker registry by issuing the following command.

```
docker load < worker-release-5.1.0.tar
```

4. Now create a file within this folder with the following contents and name it **Dockerfile**.

```
FROM worker:<TAG>
RUN apk update; pip install requests
# Add your customization here - for example, python requests module
ENTRYPOINT ["/worker.sh"]
```

5. Save this file and run the following commands.

```
docker build -t devhub.com/worker:custom-version .
```

6. Upload created image in Step 4 to your docker repository.

```
docker push devhub.com/worker:custom-version
```

To update the configured image set you must edit the following name, value pairs in the cloudcenter-ces config map as listed in this procedure.

1. Edit the cloudcenter-ces config map.

```
kubectl edit configmap cloudcenter-ces -o yaml
```

2. Add the following lines into the data section of the ConfigMap.

```
ces.image.pull.secret: custom-image-pull-secret
ces.container.image: devhub.com/worker:custom-version
```

3. Save the configMap changes. This should cause the cloudcenter-ces pod to restart.
4. Run callouts to pick up the new images for all callouts.

If you are using Cloud Remote in your environment to execute the callout scripts, the above process to build the Docker images remains the same.

Be sure to save the Docker image and load this on the Cloud Remote worker nodes.

1. Save the Docker image as a tar file.

```
docker save worker:custom-version > custom-version.tar
```

2. As Docker is already installed on the Cloud Remote instance(s), load the above tar file on each instance.

```
docker load --input=custom-version.tar  
docker tag worker:custom-version ces_worker:latest
```

Integrations

Integrations

- [Workload Manager - Action Orchestrator Integration](#)
- [Extensions](#)
 - [ACI Extensions](#)
 - [ServiceNow Extensions](#)
 - [ACI Multi-Site Extensions](#)
- [Jenkins Integration](#)
- [Service Libraries](#)

Workload Manager - Action Orchestrator Integration

Workload Manager - Action Orchestrator Integration

- [Overview](#)
- [Northbound Integration](#)
 - [Where Action Orchestrator Workflows Appear in Workload Manager](#)
 - [Creating a Workload Manager Lifecycle or On Demand Workflow](#)
 - [Manual Method](#)
 - [Automated Method](#)
 - [Other Useful Activities for Building Lifecycle and On Demand Workflows](#)
 - [Execute Action on Virtual Machine atomic workflow activity](#)
 - [JSONPath Query atomic workflow activity](#)
 - [Generic CCS API Request atomic workflow activity](#)
- [Southbound Integration](#)

Workload Manager 5.0 has built in northbound and southbound integration with Action Orchestrator 5.0:

- Northbound integration consists of the ability for Action Orchestrator workflows to be initiated as lifecycle workflows or on demand workflows from within Workload Manager.
- Southbound integration consists of prebuilt activities in Action Orchestrator designed for capturing data from Workload Manager or controlling actions in Workload Manager.

Where Action Orchestrator Workflows Appear in Workload Manager

In Action Orchestrator it is possible to define workflows that can be used as lifecycle workflows or on demand workflows within Workload Manager. When properly defined, these workflows appear in the same places in the Workload Manager UI as the corresponding lifecycle actions and on demand actions as summarized in the table below.

Workflow Type	Locations in Workload Manager UI
Lifecycle	The Properties panel in the application profile's Topology tab. When the workflow is mapped to application profiles, it is visible in the dropdown lists for the Service Initialization and Node Initialization and Clean Up scripts if it is a VM-based workflow, or in the dropdown lists for External Initialization scripts if it is an external workflow. Workflows are not available for use as migrate or upgrade scripts.
	The Add or Edit Service forms. When the workflow is mapped to a service, it is visible in the dropdown lists for Agent Lifecycle Actions scripts if it is a VM-based workflow, or External Lifecycle Actions scripts if it is an external workflow.
	The External Lifecycle Actions Settings section of the Regions tab. When an external workflow is mapped to a cloud region, it is visible in the dropdown list for External Lifecycle Actions scripts for that region.
On Demand	The actions dropdown menu for running VMs in the Virtual Machines list page.
	The actions dropdown menu for running VMs in the tiers tab of the Job Details page.
	The action buttons in the lower right of the VM details page.

The Action Orchestrator workflows are displayed at the bottom of each corresponding dropdown list, or at the bottom of the list of on demand action buttons for on demand workflows.

Creating a Workload Manager Lifecycle or On Demand Workflow

A workflow to be used Workload Manager lifecycle workflow or on demand workflow can be created in one of two ways:

- Automated method: The workflow is created by clicking an "add" menu selection from one of the dropdown lists in the Workload Manager UI. In this case, some key initial steps in creating the workflow are performed automatically.
- Manual method: The workflow is created from scratch starting in the Action Orchestrator module.

Both methods assume you are using Workload Manager 5.0.1 or later and Action Orchestrator 5.0.1 or later. The automated method is preferred, but the manual method is explained first to illustrate what happens behind the scene with the automated method.

Manual Method

1. From Action Orchestrator home screen, open the workflow editor as explained in [Creating a Basic Workflow](#).
2. From the workflow editor, find the **Get Workload Manager Context** activity in the Activities panel on the left, then click and drag it to the canvas. This is the key activity that drives the integration between Workload Manager and Action Orchestrator.
3. Before selecting the Get Workload Manager Context activity in your new workflow, the Properties panel on the right corresponds to the new workflow as a whole. Enter the name of the workflow in the display name field in the Properties panel.
4. Scroll the Properties panel down to the Variables section to create a new variable using the Add Variable button. Name the variable **CC_RUN_ID** and set it as type String and scope Input. This is needed in order to link the workflow with the deployment or VM associated with the workflow. The Get Workload Manager Context activity will automatically set this variable to the Job ID if it is used for a lifecycle workflow, or to the VM ID if it is used for an on demand workflow.

5. Optional. If this workflow needs to run a command or script on a VM, you also need to add a second variable of type String and scope Input, and name it **NODE_ID**. This is used by the [Execute Action on Virtual Machine](#) atomic workflow activity. You will use this atomic workflow activity when you want to run a command or script on a VM (see below). If you do not plan to use this workflow to run a script or command on a VM, the NODE_ID input field is not required.
6. Select the Get Workload Manager Context activity that you dragged to the canvas. This causes the Properties panel to correspond to this activity. In the Properties panel, set the **Job ID** field to the CC_RUN_ID variable by clicking the [Variable Reference](#) (puzzle piece) icon and then selecting Workflow > Input > CC_RUN_ID and then Save. This completes the linking between the new workflow and the VM or deployment from where the workflow is to be called.
7. Define the type and scope of the workflow. Click the add button under the Workload Manager Configuration section in the Properties panel. (See [Get Workload Manager Context](#) for details). This brings up dialog boxes comparable to those used in specifying the type and scope of Actions Library actions. Through these dialog boxes specify:
 - a. **Workflow Availability:** Lifecycle or On Demand.
 - b. **Execute Action:** On Virtual Machine OS or External.
 - c. **Resource mappings:** Select the appropriate resources in a fashion similar to mapping Actions Library actions to resources.
8. Add additional activities to the workflow appropriate to your use case. In particular, If you want to execute a command or script on a VM, you would add the [Execute Action on Virtual Machine](#) atomic workflow activity to your workflow. See *Other Useful Activities for Building Lifecycle and On Demand Workflows* below for more details.

Automated Method

From within the Workload Manager UI, where ever lifecycle actions or on demand actions are displayed (see previous table), an Add workflow menu selection or button will appear. Clicking Add will cause a new workflow to be generated in Action Orchestrator with steps 2 through 6 from the manual method (above) to be performed automatically, and a new tab to be opened in your browser that points to the newly created workflow in the workflow editor screen.

This newly created workflow includes an automatically generated workflow name per the following convention: if the workflow is created from the context of adding an on demand action workflow, or from adding a lifecycle action from a service, the name of the workflow is of the form: WM_AO_OOB_WF_<random_6_char_string>; if the workflow is created from the context of adding a lifecycle action workflow to an application profile, the name of the workflow is of the form: <app_profile_name>_<random_6_char_string>. You can change the name as you see fit.

The newly created workflow also includes the partial creation of the Workload Manager Configuration settings. If the workflow is created from the context of adding an on demand workflow, the workflow type is set to On Demand. If the workflow is created from the context of adding a lifecycle workflow, the workflow type is set to Lifecycle.



The automated method requires that you first import the WM_CREATE_WF workflow into your user account from the public Cisco Action Orchestrator repository. This is included in the instructions below.

The steps to create a new workflow from within Workload Manager are as follows:

1. If this is the first time you are attempting to create workflows from within Workload Manager, you must first import the WM_CREATE_WF workflow into your user account from the public Cisco Action Orchestrator repository as described below.
 - a. From within Action Orchestrator, create a Git repository integration with the public Cisco Action Orchestrator repository. This requires that you have the tenant admin role. If another tenant admin in your tenant has created this repository link already, get the name of that integration and skip to step b, below.
 - i. In the Action Orchestrator UI, navigate to Admin > Git Repositories and click New Git Repository.
 - ii. In the New Git Repository dialog box, ensure the following fields are set as specified:
 1. No Account Keys = TRUE
 2. Rest API repository = api.github.com/repos/cisco/ActionOrchestratorContent
 3. Branch = master
 4. Code Path = /workflow-examples
 - iii. Give the Git repository integration a memorable name, for example, CiscoPublic, and save it.
 - b. Import the WM_CREATE_WF workflow action:
 - i. Navigate to Main Menu > Workflows and click Import in the upper right.
 - ii. In the Import Workflow dialog box, ensure the following fields are set as specified:
 1. Imported From = Git
 2. Git Repository = The name the Git repository integration created in step 1a, above.
 3. Filename = CloudCenterSuite-WMCreateWF
 4. Git Version = The latest version of the workflow available in the dropdown
 - iii. Click the Import button in the lower right of the dialog box. A new WM_CREATE_WF workflow should appear in your list of workflows in the My Workflows tab.
2. Switch from from the Actions Orchestrator UI to the Workload Manager UI. From any location where it is possible to select a lifecycle action or an on demand action (see previous table), scroll down to the end of the dropdown menu, or list of buttons, and select Add. This will open a new tab and eventually display the newly created workflow in the Action Orchestrator workflow editor.
3. In the Action Orchestrator workflow editor canvas, click the Get Workload Manager Context activity to cause its properties to be displayed in the Properties panel, scroll down to the Workload Manager Configuration section, and click the edit button. Use the dialog boxes to specify:
 - a. **Workflow execution location:** externally or on the VM.
 - b. **Resource mappings:** select the appropriate resources in a fashion similar to mapping Actions Library actions to resources.

4. Add additional activities to the workflow appropriate to your use case. In particular, If you want to execute a command or script on a VM, you would add the [Execute Action on Virtual Machine](#) atomic workflow activity to your workflow. See *Other Useful Activities for Building Lifecycle and On Demand Workflows* below for more details.

Other Useful Activities for Building Lifecycle and On Demand Workflows

Execute Action on Virtual Machine atomic workflow activity

The [Execute Action on Virtual Machine](#) activity is required if you want your workflow to execute a command or script on a VM. This activity should only be used in workflows where the Get Workload Manager Context is the first activity. To use this activity in a workflow, perform these steps:

1. Check to see if this atomic workflow is visible in the activities panel on the left side of the workflow editor screen. If yes, skip to step 2. If no, it must be imported by a user in your tenant with tenant admin privileges using these steps:
 - a. If you already have a Git repository integration with the /atomic-workflows code path in the public Cisco Action Orchestrator repository, skip to step b. Otherwise, create one now using the following procedure.
 - i. In the Action Orchestrator UI, navigate to Admin > Git Repositories and click New Git Repository.
 - ii. In the New Git Repository dialog box, ensure the following fields are set as specified:
 1. No Account Keys = TRUE
 2. Rest API repository = api.github.com/repos/cisco/ActionOrchestratorContent
 3. Branch = master
 4. Code Path = /atomic-workflows
 - iii. Give the Git repository integration a memorable name, eg, CiscoPublicAtomic, and save it.
 - b. Import the Execute Action on Virtual Machine workflow action:
 - i. Navigate to Main Menu > Workflows, select the and click Import in the upper right.
 - ii. In the Import Workflow dialog box, ensure the following fields are set as specified:
 1. Imported From = Git
 2. Git Repository = The name the Git repository integration created in the step a, above.
 3. Filename = CloudCenterSuite-ExecuteActionOnVM
 4. Git Version = The latest version of the workflow available in the dropdown
 - iii. Click the Import button in the lower right of the dialog box. A new Execute Action on Virtual Machine atomic workflow should appear in the list of workflows in the Atomic Workflows tab.
2. In the Actions Orchestrator workflow editor screen, drag and drop this atomic workflow from the activities panel to the canvas at a position after the Get Workload Manager Context activity.
3. Select this newly added activity to cause the properties panel to refer to this activity. From the Properties panel, populate the following required input fields:
 - a. **Get CloudCenter Context Response:** Click the variable reference icon and select Activities > Get Workload Manager Context > Response Body and then Save.
 - b. **Action Type:** Click the variable reference icon and select Activities > Get Workload Manager Context > Action Type and then Save.
 - c. **NODE_ID:** If this workflow is a lifecycle workflow, click the variable reference icon and select Workflow > Input > NODE_ID and then Save. Otherwise, this field can be left at its default value which is null.
 - d. **Script:** Enter the path of the script or command on the VM that is to be executed. The script may be a shell script for Linux VMs or a PowerShell script for Windows VMs.

JSONPath Query atomic workflow activity

The [JSONPath Query](#) activity lets you parse the response body of a previous activity in the workflow from which it is called. The response body of the Get Workload Manager Context is either a job details API call response body (for lifecycle workflows) or a VM details API call response body (for on demand workflows); therefore, the JSONPath Query activity can be used to extract specific fields related to the job or VM so that they can be used later in the workflow. For example, to parse the owner email address from the Get Workload Manager Context response body, follow the steps below.

1. Drag and drop activity from the activities panel to the canvas at a position after the Get Workload Manager Context action.
2. Select this newly added activity to cause the Properties panel to refer to this activity. In the Properties panel, set the **Source JSON to Query** field to the Get Workload Manager Context response body by clicking the variable reference icon and then selecting Activities > Get Workload Manager Context > Response Body and then Save.
3. Under the **JSONpath Queries** label, click the Add button. This reveals three fields:
 - a. **JSONpath Query:** Enter the field name as specified in the response body output. In the case of the owner email address, enter: "\$.ownerEmailAddress".
 - b. **Property Name:** Enter a user friendly field name.
 - c. **Property Type:** Select a field type from the dropdown menu. In the case of the owner email address, select String.
4. Optional: To extract addition fields from the response body, repeat the the previous step as necessary.

Generic CCS API Request atomic workflow activity

The [Generic CCS API Request](#) activity lets you execute a CloudCenter Suite API call within your workflow.

1. Drag and drop activity from the activities panel to the canvas at the appropriate position of your workflow.
2. Select this newly added activity to cause the properties panel to refer to this activity. In the Properties panel, under the **CCS API Request** label, enter values for these three fields:
 - a. **Relative URL:** Enter the portion of the API call that comes after <address>:<port>.
 - b. **Method:** Enter GET, POST, PUT or DELETE as appropriate.

c. **Request Body:** Enter the JSON request body if required.



The request body can include any of the variables available within the workflow: position your cursor at the appropriate position of the request body, click the variable reference icon, and then use the variable reference browser to select the appropriate variable.

Action Orchestrator has a set of activities under the [CloudCenter Suite adapter](#) for performing actions on various components within the CloudCenter Suite. Some of these activities are unique to Cost Optimizer, some are unique to Workload Manager, one is shared between Workload Manager and Cost Optimizer, and one is common to all modules as summarized below.

- Workload Manager Specific Activities
 - [Get Workload Manager Context](#). See also *Creating a Workload Manager Lifecycle or On Demand Action Workflow*, above.
 - [Manage Deployment Environment](#). With this activity, you can add a cloud account to one or more regions within a deployment environment. You must specify the deployment environment and regions already associated with the deployment environment. You must specify the cloud account using its CloudCenter Suite account ID. When specifying regions, you must ensure that all regions you select are associated with the same cloud type as the cloud account you are adding.
 - [Execute Action on Virtual Machine](#). See *Other Useful Activities for Building Lifecycle and On Demand Workflows*, above.
- Workload Manager and Cost Optimizer Shared Activity
 - [Add Cloud Account](#). Use this activity to add a cloud account created within the cloud provider to the CloudCenter Suite.
- CloudCenter Suite Common Activity
 - [Generic CCS API Request](#). See also *Other Useful Activities for Building Lifecycle and On Demand Workflows*, above.

Extensions

Extensions

- [ACI Extensions](#)
- [ServiceNow Extensions](#)
- [ACI Multi-Site Extensions](#)

ACI Extensions

Configure ACI Extensions

- [Overview](#)
- [ACI Fundamentals](#)
- [Benefits](#)
- [Availability](#)
- [Requirements](#)
 - [APIC Requirements](#)
 - [Workload Manager Requirements](#)
 - [VMware vSphere Requirements](#)
- [CloudCenter Extensions](#)
- [Using ACI with CloudCenter Extensions](#)
- [Configuring Extensions](#)
- [Launch the ACI Extension](#)
- [Bridge Domain Template](#)
- [ACI Actions](#)
- [Troubleshooting](#)

Workload Manager users can use out-of-the-box application profiles to create infrastructure-independent models of any application. Once modeled, the Cisco CloudCenter platform and Cisco [Application Centric Infrastructure \(ACI\)](#) can work together to provide automated, end-to-end provisioning of compute, storage, and network configuration of the application as well as its set of required components.

See the [Cisco ACI Fundamentals Guide](#) for additional details on the ACI policy model.

The Workload Manager – ACI integration provides the following benefits:

- Use a fully automated creation of ACI policy objects.
- Gain the security and efficiency of network microsegmentation without the need to program or modify application code, write cloud-specific scripts, or have special network expertise.
- Users get self-service/on-demand deployment and management of applications with fully integrated Cisco ACI network policy and configuration.

The Workload Manager – ACI integration is available for **VMware** cloud environments.

Workload Manager supports the following [APIC](#) releases:

- Cisco APIC, Release 1.0
- Cisco APIC, Release 1.1
- Cisco APIC, Release 1.2
- Cisco APIC, Release 2.0 (only Distributed Virtual Switch – DVS mode)
- Cisco APIC, Release 2.1
- Cisco APIC, Release 2.3
- Cisco APIC, Release 3.0
- Cisco APIC, Release 3.1
- Cisco APIC, Release 4.0

The CloudCenter platform automates the end-to-end-provisioning of the overlay infrastructure and deployments of applications. On ACI, this includes the provisioning and management of the following resources:



Ensure that the APIC tenant being configured in the CloudCenter has the privileges to create these resources.

- Application Network Profiles (ANP)
- Endpoint Groups (EPG)
- Contracts
- Subjects/Filters

As a prerequisite for the CloudCenter platform to provision and configure the applications on APIC, first complete the following requirements to have a working Cisco ACI environment:

- Leaf switch profiles, Switch Selectors, Interface Profile, and Policy Groups
- VLAN Pool
- VMware's Virtual Machine Manager (VMM) Domain
- Routable IP subnet to a New Tenant and Bridge Domain(s) configured with Layer 3 out (L3 Out) for external internet connectivity.



The CloudCenter platform uses the L3 Out network to associate the *Common* tenant (or the selected tenant).

- Routing protocols
- VRF

APIC Requirements

The Cisco Application Policy Infrastructure Controller (Cisco APIC) functions over both HTTP or HTTPS.

- **HTTPS:** By default, Cisco APIC listens to HTTPS for both the UI and REST APIS.



Ensure that the APIC is configured with a valid SSL certificate that corresponds to the APIC host name.

- **HTTP:** Enable the HTTP access for APIC and ensure accessibility using either the host name or IP address

To ensure the sanity of the environment, follow this procedure.

1. Using the APIC UI, manually add a new application network profile with one EPG.
2. Verify that a new VMware Virtual Distribute Switch (vDS) port group is provisioned and displayed in the APIC UI.
3. Using the vCenter UI, provision/clone a new VM with the network pointing to the created port group.
4. If operating in *Strict* mode, you will not have SSH/RDP access to the VM:
 - a. Create a *Contract* for Port 22/3389 with its provider being the EPG from Step 1.
 - b. Create a new L3 out setting to be consumed by the *Contract* created in Step 4a.
5. SSH/RDP into the VM launched in Step 3 and verify that you can access the CloudCenter [Local Bundle Store \(Conditional\)](#) repository and the AMQP server.

Workload Manager Requirements

The cloud region being used in the ACI Extension should be able to access the corresponding APIC endpoint – activate the native APIC integration in CloudCenter by adding the endpoint URL of the APIC as an Extension in the Admin area

The CloudCenter platform assumes that you have configured the ACI extension based on the following setup requirements:

- The typical ACI constructs for the tenant (Bridge Domain, DHCP Policy/Relay Label, VRF, External Routed Network - whether tenant specific or shared from the common tenant) are preconfigured and operationally health.
- CloudCenter will create a new ACI Application Profile, new EPGs - one per tier of the application, new filters and contracts and apply them to the new EPGs according to the design of the CloudCenter application profile.
- The ACI objects created by CloudCenter are named after the original deployment name so that they can be quickly and easily traced to the CloudCenter deployment.
 - Configure the ACI Extension in the CloudCenter platform.
 - Once the extension is selected, the CloudCenter will auto-discover the objects relevant to the privileges of the user whose credentials were used to configure the ACI Extension.
 - The following resources are specified by the APIC:
 - Virtual Machine Manager – My-vCenter is specified by the APIC
 - APIC Tenant
 - L3 Out
 - Network Type = ACI
 - End Point Group = Existing EPG
 - Existing EPG:
 - Web-Servers
 - DB-Servers
- To enable an application for ACI compatibility, enable the micro-segmentation capability in the application profile. The default firewall rule for a service are automatically displayed in if micro-segmentation is enabled – you can restrict any firewall rule to any tier by specifying the tier name or IP for the source. See [Security and Firewall Rules](#) > *Inter-Tier Communication (Firewall Rules)* for additional context.

VMware vSphere Requirements

The following table describes the VMware vSphere requirements.

Requirement	Details
A working VMware vCenter 5.0/5.5/6.0 environment	The minimum VMware vSphere version is v5.0, but vSphere v5.5 U2 is optimal.
The CloudCenter platform automates the provisioning of virtual machines into the VMware private datacenter.	The CloudCenter platform requires <i>access credentials</i> to the vCenter setup.
All ESX host(s) must be physically connected to the ACI leaf switches.	The prerequisite installation requirements for the datacenter are: <ul style="list-style-type: none"> • A physical ESX host capable of running at least 10 medium sized instances • An ESX cluster (cluster could comprise of just the one host) • A datastore (or datastore cluster for DRS support), at least 100gb of free space

If the ESXi hosts are Cisco UCS based

- The VLANs for the CMM must be mapped to the vNIC template.
- The uplinks from the Fabric must interconnect trunking VLANs to the leaf switches.

The [APIC policy model](#) is available as a standalone extension on the CloudCenter platform and provides increased ease when creating ACI objects by allowing better, faster, and easier network isolation by:

- Using Extensions on the CloudCenter platform, network administrators can access CloudCenter from the UI or the API to create, update, or delete the following objects:
 - Bridge domains
 - Virtual Machine Manager (VMM) domains
- Allowing the consumption of newly-created bridge and VMM domains during the application deployment process or the deployment environment process without having to manually sync configurations.

Once you configure an extension (procedure provided later in this page), select the cloud and cloud accounts in the Network Settings section to see that a configured network, such as Cisco ACI, is available for selection when configuring this deployment environment.

You can create a CloudCenter extension to extend the capabilities of the cloud region to provision networks in an ACI environment. You can then *Launch the ACI Extension* to configure the following CloudCenter resources:

- **Deployment Environment Flow:** ACI Extensions are also integrated in the deployment environment and you can determine the extension to be used by each cloud account. CCMs do not need to make the request to the cloud provider. See Deployment Environments Defaults for additional context.
- **Application Deployment Level:** Configure tenant and VMM domains to be populated into the application profile when [Deploy an Application](#). You can configure the External Routed Network field (Layer 3 out) for your APIC setup, as shown in the following screenshot, and connect to that tenant network.

USE ACI EXTENSION

ON

* APIC EXTENSION

CiscoLabACI_ByAdmin01

* VIRTUAL MACHINE MANAGER

VMM

* APIC TENANT

CliQr-Dev-3

L3 OUT

External

- **NIC:** When you select the Cisco ACI tab, you have the option to select one of the options from the Endpoint Group (EPG) Type, as shown in the following screenshot.

NIC 1

NETWORK TYPE

VMware Cisco ACI

* END POINT GROUPT (EPG) TYPE

New EPG

* BRIDGE DOMAIN

CliQr-BD

CONTRACTS

CliQr-Dev-3/default common/maple-l3-out

CliQr-Dev-3/imported

- **Existing EPG:** Uses a preconfigured EPG from the APIC setup.
- **New EPG:** Creates a new EPG for this deployment. Optionally, you can also select contracts or interfaces that are consumed by this new EPG.
- **Bridge Domain Template:** Creates a new bridge domain using the selected template.
- **ACI as an External Service:** When you [Deploy an Application](#) that contains an [External Service](#), you can configure the ACI extension in the **Advanced** section for this service tier to use the APIC Service Graph Template. The following screenshot shows this section.

Cloud Settings

APIC EXTENSION
ACI

APIC TENANT
CliQr-Dev-2

SERVICE GRAPH TEMPLATE
LBSERVICEGRAPHCOMMON

When you have configured the cloud or datacenter resources (for example, the tasks listed in the ACI Integrations section), verify your network connectivity and launch a sample application to ensure everything is working from end-to-end. If all the requirements worked, you are ready to configure the extension from the CCM UI.

To configure an extension from the Workload Manager UI, follow this procedure.

1. Access the Workload Manager UI and navigate to **Admin > Extensions**. The Extensions page displays and you can edit an existing extension or add a new extension as required for your ACI integration. The following screenshot shows the Extensions page.

NAME	EXTENSION URL	TYPE	ACTIONS
ext1	http://12.18.11.0	ACI	-Actions-

i The TYPE column in the Extensions page currently displays ACI for all extensions as this is the only type of extension that is currently accepted by the CloudCenter platform.

2. Click **Add Extension**. The New ACI Extension page displays, as shown in the following screenshot.

New ACI Extension

Connection Settings

*** NAME**

*** APIC CONTROLLER URL**

*** USERNAME**

*** PASSWORD**

*** MANAGED ORCHESTRATOR**

Vmware
▼

3. Configure the following Cisco APIC endpoint information in the **Connection Settings** section:

- The APIC Name
- The APIC endpoint URL (HTTP or HTTPS)
- The APIC access credentials (Username and Password) – Use the ACI admin credentials for APIC access.



If you do not use Admin credentials for the ACI account

Most integrations use the ACI admin credentials for APIC access.

If you use the Admin credentials, you do not need the information in this note.

If you prefer to limit the CloudCenter platform's access to ACI, then make sure that the ACI user for the ACI–CloudCenter integration account is configured as follows:

- Create the Security Domain (SD) and associate the SD with the VMM(s) and Tenant(s) used for the ACI – CloudCenter integration.
- Create an ACI *Role* with the following privileges:
 - vmm-connectivity
 - tenant-security
 - tenant-network-profile
 - tenant-epg
- Add (associate) the SD to the ACI–CloudCenter integration account and assign the created ACI Role with *writePriv* (write privilege).

- The cloud region used to manage this APIC endpoint (select the required CCO from the dropdown list)
4. Click **Connect** to connect and save the ACI configuration information.
 - a. The CloudCenter software validates the APIC endpoint connection and displays a status message displays at the top of this page.
 - b. Once the APIC endpoint connects successfully, you also see the New ACI Extension page refresh to display the **Bridge Domain Template** section below the **Connection Settings** section. You can use this section to provide additional placement information. See the *Bridge Domain Template* section below for additional details.
 5. Click **Save** to save this new extension. The Extensions page refreshes to display the newly-configured extension to the list of configured and validated Extensions.

To launch the ACI integration in your cloud, follow this procedure.

1. Access the CCM UI and navigate to **Deployments**. The Deployments page displays
2. Click the **Environments** tab. The Deployments page refreshes to display the configured environments and you can edit an existing environment or add a new environment as required for your ACI integration.
3. Click **Add Environment**. The New Deployment Environment page displays.
4. In the **General Settings** section:
 - a. Provide the deployment environment **Name**
 - b. Optionally, provide a **Description**.
 - c. Identify if approval is required to deploy to this environment by switching **On** the button.
5. In the **Cloud Selection** section:
 - a. Select the checkbox for the required **Cloud Region**. This cloud region must be the same as the CCO cloud region (used to manage your new APIC extension in the above section).
 - b. Select the **Cloud Account** from the dropdown list.
6. Click **Define Default Cloud Settings** to define the Deployment Environment default settings for this cloud. See Deployment Environment Defaults for additional context.



You can pre-define much of your experience during the deployment submission in the Default Settings of the Deployment Environment. See Deployment Environment Defaults for additional context.

7. (Optional) Define the **Networks** Settings:
 - a. Turn **On** the Use Network Types button. The Networks section expands to display the Network Types.
 - b. Click **+Network Type** to add a new type. The New Network Type page displays.
 - i. Provide the network type **Name**.
 - ii. Optionally, provide a **Description**.
 - iii. Configure the **Network Settings**. The available networks for this cloud are displayed in the Network Settings section. The Network Settings section differs for each cloud.

VMware Network Settings

1. Toggle the **Visibility** switch to determine if you want to allow your end users to use pre-configured settings.
 - **OFF:** (Default) End users are not allowed to use preconfigured ACI extensions.
 - a. Select the Network in the NIC section. See [IP Address Allocation](#) for additional context on NIC configuration.
 - b. Add additional NICs, if required.
 - **ON:** End users are allowed to use preconfigured ACI extensions.
 - a. Select the required extension, the corresponding options are displayed in the dropdown list for the remaining fields (see [Extensions](#) for additional details):
 - b. Select the APIC Extension from the dropdown list (see [ACI Extensions](#) for additional details).
 - c. Select the APIC **Virtual Machine Manager** (VMM) associated with this APIC Extension from the filtered dropdown list.
 - d. Select the **APIC Tenant** associated with this APIC Extension from the filtered dropdown list.
2. Select the Network in the NIC section.
 - If you select **VMware**, select the Network in the NIC section. See [IP Address Allocation](#) for additional context on NIC configuration.
 - If you select **Cisco ACI**, select the type in the **End Point Group (EPG) Type** field.
 - a. **Existing EPG:** If you select this type, you must further select a pre-existing EPG (that is already connected to one of the Bridge Domains) from the **Existing EPG** dropdown, which appears if you select this type.
 - b. **New EPG:** If you select this type, you must further select a pre-existing Bridge Domain (to which this EPG must connect) from the **Bridge Domain** dropdown list.
 - c. **Bridge Domain Template:** See [Extensions](#) for additional context.
3. Add additional NICs, if required.

Back to:

- [Deploy an Application](#)

8. Click **Save** to save this new deployment environment. The Environments page refreshes to display the newly-configured deployment environment to the list of configured and validated Environments.
9. Designate a Bridge domain from the ACI environment. The list of bridge domains is pulled from ACI. See the *Bridge Domain Template* section for additional context.

A *bridge domain* represents a Layer 2 forwarding construct within the fabric. The Bridge Domain template (Layer 2 space) is linked to an ACI Virtual Routing and Forwarding (VRF) template (Layer 3 space). See the [Cisco ACI Fundamentals Guide](#) for additional details.

From the CloudCenter context, the ACI integration requires a routable IP subnet to a New Tenant that is configured with Layer 3 Out for external internet connectivity. When configuring an ACI Extension as part of the Deployment Environment Defaults, you have the option to select Bridge Domain Template in the Cisco ACI, End Point Group (EPG) Type field.

If you do, you should have already configured the Bridge Domain Template so it displays in the dropdown list for that field.

CloudCenter administrators can create a Bridge Domain template to configure ACI extensions:

- Each time CloudCenter admins configure an ACI extension, they also have the option to configure a Bridge Domain template.
- The L3 Out connection to the external world is through the CloudCenter EPG Type selection. If you are deploying this instance into an existing EPG type, you do not need to update the subnet mask each time.
- To restrict this subnet from being accessed by any other network, update the subnet mask with the database tier ID in the Bridge Domain template. This way, the subnet is exposed to the world on this external network and allows the destination to be open to the DB node.
- When connecting to the database tier, the database Layer 3 out is linked to one of the IP addresses displayed in a dropdown list — instead of allowing everyone to connect to a tier.

To add a Bridge Domain Template, follow this procedure.

1. Access an ACI Extension as outlined in the section above (**Admin > Extensions**) and edit an existing extension. You can also opt to create a new extension in the process outlined above and continue to add a Bridge Domain Template as an extension of that process.
2. In the Add ACI Extension page or Edit ACI Extension page, scroll down to the **Bridge Domain Templates** section, as shown in the following screenshot.

The screenshot shows the 'Edit ACI Extension' interface. The top section is 'Connection Settings' with the following fields:

- * NAME: ext1
- * APIC CONTROLLER URL: http://192.168.131.10
- * USERNAME: admin
- * PASSWORD: *****
- * MANAGED ORCHESTRATOR: VMware-default


The bottom section is 'Bridge Domain Templates', which contains a table with one entry:

NAME	VRF	POOL SUBNETS (# NETWORKS)	
t1	Tenant_auto332-VRF	192.168.203.0/24 (8)	

An 'ADD TEMPLATE' button is located in the top right of the table area. A 'SAVE' button is located at the bottom right of the entire form.

3. Click **Add Template**. The New Bridge Domain Template pages displays.
4. Configure the following Bridge Domain Template details in the **General Settings** section:
 - **Template Name:** A name reference by which you can refer to this Bridge Domain template.
 - **Bridge Domain Name Configuration:** The exact name variable for the Bridge Domain that is used by the ACI.
 - VRF Selection:
 - **Existing VRF:** Select the VRF from the dropdown list. Templates are listed by tenant in the dropdown list, be sure to select the VRF template for the correct ACI tenant.

✓ Shared resources are saved in the Common tenant, as shown in the following screenshot.



When you select a VRF from a *Common* tenant (highlighted in the dropdown list image), that Bridge Domain Template can be selected by any tenant and consequently deployed to any other tenant. If you select a VRF that is specific to just one tenant, you can only deploy the Bridge Domain Template to just that tenant.

- **Dynamic VRF:** Select a VRF that is provisioned for this APIC. The VRF hosts the Bridge Domain that is created using the Bridge Domain Template.
 - **Associated L3 Outs:** Optional. Depending on the tenant selected in the VRF settings, you can now associate the L3 Out networks from the *Common* tenant (or the selected tenant).
 - **L3 Out for Route Profile:** Optional. Depending on the tenant selected in the VRF settings, you can now select the desired L3 Out for route profile from the *Common* tenant (or the selected tenant).
 - **DHCP Relay Label:** Optional. Depending on the tenant selected in the VRF settings, you can now select the one or more DHCP relay labels from the *Common* tenant (or the selected tenant) that is applied to the new bridge domain.
5. Configure the following network details in the **Subnet** section.

- Scope: APIC concept – See the [Cisco ACI Fundamentals Guide](#) for additional details.
 - **Private to VRF:** An APIC setting that refers to a Private Network (context) is equivalent to a virtual routing and forwarding (VRF) instance in the networking world.
 - **Advertised Externally:** An APIC setting that refers to an EPG that provides a shared service must have its subnet configured under that EPG (not under a bridge domain), and its scope must be set to advertised externally, and shared between VRFs.
 - **Shared between VRFs:** An APIC setting that refers to shared subnets must be unique across the VRF involved in the communication. When a subnet under an EPG provides a Layer 3 external network shared service, such a subnet must be globally unique within the entire ACI fabric.
- Subnet Control: APIC concept – See the [Cisco ACI Fundamentals Guide](#) for additional details.
 - **ND RA Prefix:** An APIC setting to control Neighbor Discovery (ND) – Router Advertisement (RA) message communications between an outside public or private network and the ACI fabric.
 - **Querier IP:** An APIC setting to enable Internet Group Management Protocol (IGMP) snooping on the subnet.
- Subnet Pools: CloudCenter concept – Prevents any subnet in the pool from being wrongly reused. When you deploy a Bridge Domain Template on an application with multiple tiers, then each tier will use a different subnet from within this pool to ensure that the same subnet is not reused multiple times. If you deployment uses more subnets than are defined in this pool, the deployment will fail as all configured subnets are already used in this deployment.)
 - **Master Subnet:** The IP address of the first subnet in the tenant network.
 - **Pool Subnet:** A dropdown list to identify the last subnet in the tenant network.
 - **Networks:** This section automatically updates to reflect the number of networks in the pool based on the Master and Pool Subnet configurations.
 - You can add multiple subnet pools by clicking the **Add Subnet Pool** button.
 - **Delete Icon:** Allows you to delete a previously configured subnet pool from the CloudCenter platform.

✔ Once you add a subnet pool, you cannot update the pool. You can only delete the configured pool and add a new subnet pool.

6. Click **Save** to save this new Bridge Domain Template along with the configured ACI extension. The Extensions page displays the Success message below the header to state the the extension is saved.

Administrators can perform the actions that the following screenshot shows for each ACI extension listed in the Extensions page.

NAME	EXTENSION URL	TYPE	ACTIONS
ext1	http://192.168.131.10	ACI	-Actions- ▾
test123	http://192.168.131.10	ACI	-Actions- ▲ -Actions- Edit Share Delete

The Deployment Environment pages list configured information and allows you perform the actions that the following table describes.

Actions Dropdown	Description
Edit	<p>Change configurations for an existing extension. Once configured, you can only perform the following changes to an Extension:</p> <ul style="list-style-type: none"> • Change the name of the ACI Extension, URL, username, password as required. • Add a new Bridge Domain Templates for the extension. • Delete a configured Bridge Domain Template. <p>See the <i>Adding a Deployment Environment</i> section (below) for additional details.</p>
Share	Share an Extension. See Permission Control > <i>Extension Permissions</i> for details.
Delete	<p>Delete an Extension.</p> <p>If you choose to delete a configured Extension, the Delete Extension popup confirms your intention, deletes the configured Extension, and displays a status message at the top of the Extension page.</p>

If you set the cliqignoreAppFailure parameter (see [Troubleshooting Parameters](#)), then the APIC resources (ANP, EPGs, Contracts, and so forth) created using the CloudCenter platform are not removed if the deployment fails. The launched VMs and related APIC policies are only removed when the user terminates the deployment from the Deployments page. See [Terminate Protection](#) for additional context.

ServiceNow Extensions

Configuring ServiceNow Extensions

- [Overview](#)
- [Version Support](#)
- [Changes to Integration – CloudCenter Suite \(Application\) v4.2.0](#)
- [Request Integration Application from ServiceNow App Store](#)
- [Install the Integration App](#)
- [Integration Configuration in ServiceNow](#)
 - [Application Configuration](#)
 - [Disable the Checking of Untrusted SSL Certificates](#)
 - [Configure Connection with CloudCenter Suite](#)
 - [Automated Method](#)
 - [Manual Method](#)
 - [Create an Integration User](#)
 - [Create a Test User](#)
- [Integration Configuration in CloudCenter Suite](#)
 - [Create a ServiceNow Extension](#)
 - [Enable CMDDB](#)
 - [Associate ServiceNow Extension to a Deployment Environment](#)
- [Validate the Integration](#)
- [Optional Configuration](#)
 - [Connecting to Multiple CloudCenter Suite Instances or Tenants](#)
- [Additional Information](#)
 - [High-level Interaction](#)
 - [Sequence Diagram](#)
 - [User Roles](#)
 - [Catalog Item – Deploy Application Profile](#)
 - [Catalog Item – Manage Deployments](#)
 - [API Cache](#)
 - [Image Cache](#)
 - [Workflows](#)
- [Tables impacted](#)

Integration between CloudCenter Suite and the ServiceNow platform is provided by a Cisco developed and certified ServiceNow application called **Integration – Cisco CloudCenter Suite**. This application is available at no cost and can be requested from the [ServiceNow App Store](#). The integration application provides a mechanism to easily setup communication between CloudCenter Suite (Workload Manager) and ServiceNow in order to deploy and manage Application Profiles from the [Service Portal](#).

Based on new versions or platform changes, Cisco validates and certifies the integration application as needed. The contents of this section apply to the platform, product, or component versions identified in the following table.

Product Name	Version(s)
Integration – Cisco CloudCenter Suite (Application)	4.2.0
CloudCenter Suite (Workload Manager)	<ul style="list-style-type: none"> • 5.2.x • 5.3.x
ServiceNow	<ul style="list-style-type: none"> • Orlando • New York



You can upgrade to Integration - CloudCenter Suite (Application) v4.2.0 from the previous version (v4.1.1).

To upgrade, navigate to **System Applications > All Available Applications > All**. Search for the application. Select the version to upgrade and click **Update**.


The following table lists the versions of the platform, product, or component that were used to develop the information in this section.

Product Name	Version(s)
Integration – Cisco CloudCenter Suite (Application)	4.2.0
CloudCenter Suite (Workload Manager)	5.3.0
ServiceNow	New York

The ServiceNow app *Integration – CloudCenter Suite v4.2.0* includes the following enhancements and updates:


- Integration and tenant configuration settings have been consolidated for easier setup.
- A ServiceNow Requested Item record is now created regardless of the deployment request being initiated from the ServiceNow Service Portal or directly from CloudCenter Suite. In previous releases, a ServiceNow Request was created for deployments initiated from CloudCenter Suite and a Request with Requested Item was created for deployments initiated from the ServiceNow Service Portal.
- When VMs are imported into CloudCenter Suite, their details and status can now be added to the ServiceNow CMDB. This can be enabled or disabled from the *Tenant Configuration* module.
- If a deployment in CloudCenter Suite remains in a *Starting* state for 2 hours, the deployment workflow in ServiceNow is terminated. This prevents ServiceNow Requests from remaining in an *Open* state.
- If a new deployment request is canceled in CloudCenter Suite prior to being approved in ServiceNow, the Request record in ServiceNow is marked as *Closed*.
- The Approval request now contains a *Zone* value for OpenStack environments.
- The Virtual Machine table in CMDB now displays the *Start Date*, *Terminated On*, and *VM running hours* fields.
- The Storage Volumes table in CMDB now the *Storage Type* field.
- A new *info* variable of RITM contains additional details about requested action, job or VM, which can be used in workflows to automate approvals, display on the form, or define catalog tasks.
- Resolved an issue where validating the owner API key in ServiceNow created a new record in the Tenants Mapping module.
- Resolved an issue where certain Day 2 actions resulted in multiple ServiceNow approval requests. For example, a resize action may require a VM to be stopped, resize and then restarted. This action now creates a single approval request.


To setup the integration between CloudCenter Suite and the ServiceNow platform, you must request the **Integration – CloudCenter Suite** application from the ServiceNow App Store.

 This task requires ServiceNow HI credentials.

To request the integration app from the ServiceNow App Store, follow this procedure.


1. Go to <https://store.servicenow.com> and log in using your ServiceNow HI credentials.
2. Search and select the **Integration – Cisco CloudCenter Suite** application.
3. Confirm that you have the correct version selected.
4. Click **Contact Seller** to request the application.
5. Once approved by Cisco, you will have the ability to download the app to your ServiceNow instance. In the unlikely event that you do not receive approval, contact your Cisco account team for assistance.

 You can only install the integration app after your request has been approved by Cisco.


 Before you begin, you will need ServiceNow admin credentials and CloudCenter Suite root or tenant admin credentials for the installation.

To install the integration application, follow this procedure.

1. Within your ServiceNow instance, use the Filter Navigator and browse to **System Applications > Applications > Downloads > Integration-Cisco CloudCenter Suite**.
2. Click **All versions**.
3. If presented with multiple application versions, select the appropriate version for your installation.
4. **Preview** and then **Commit** the application installation. You may see some error messages during the preview state – this is normal.

 It is normal to see some error or warning messages during the Preview stage. Review the errors/warnings and take a screenshot if necessary for future reference.

Select all the errors/warnings and choose the option to **Skip remote update**. If required, click Commit Update Set to proceed with the installation.

 The overall configuration steps require that you have administrative access to both the Cisco CloudCenter Suite and ServiceNow instances.

Application Configuration

To setup the application, follow this procedure.

1. Using the ServiceNow Filter Navigator, navigate to **Cisco CloudCenter Suite > Configuration > Integration Configuration**.
2. Configure the parameters as described in the following table and displayed in the screenshot below.

✓ As an alternative, you can update the following parameters by navigating to **System Properties** and executing a search for the following pattern:

x_cqt_c3*

Parameter	Type	Default	Description
Tenant name field on the User record	Optional	NA	Field reference in the user record to identify which CloudCenter Suite tenant the user belongs to.
Automatically create accounts in CloudCenter	Required	Yes	Yes = A ServiceNow user visits a CloudCenter Suite catalog item in the ServiceNow Portal and if a matching user is not found, then an account is automatically created for this user in CloudCenter Suite. No = A ServiceNow user visits a CloudCenter Suite catalog item in the ServiceNow Portal and if a matching user is not found, then an account is NOT created for this user in CloudCenter Suite.
Use a single user mode	Required	No	Yes = All operations for CloudCenter Suite will take place using the account configured in the Owner API Keys section. User accounts for individual users will not be created in CloudCenter Suite. No = User accounts will be created for individual users.
Retry attempts on communication error	Optional	2	Number of retries to attempt upon any communication error with CloudCenter Suite.
Price visibility in the Service Portal	Required	Use Local Currency	Defines if and in what currency prices should be visible to the user.
Hide SSH, RDP, and PWD options from <i>Manage Deployments</i> catalog item	Required	No	Yes = The user will not be able to connect to the VM using the RDP or SSH options available on the <i>Manage Deployments</i> catalog item form. No = The user can connect to the VM using the RDP or SSH options available on the <i>Manage Deployments</i> catalog item form.
Logging level	Required	info	Defines a verbosity level of logs. Change to debug only at debugging time.

The screenshot shows the 'Integration Configuration' page in ServiceNow. The left sidebar contains navigation options for 'Integration', 'Cisco CloudCenter Suite', and 'Configuration'. The main content area is titled 'Integration Configuration' and contains the following settings:

- Tenant Name field on the user record:
- Automatically create accounts in CloudCenter Suite: Yes | No
- Use a single user mode: Yes | No
- Retry attempts on communication error:
- Price visibility in Service Portal:
- Hide SSH, RDP and PWD options from 'Manage Deployments' catalog item: Yes | No
- Logging level:

A 'Save' button is located at the bottom of the configuration area.

Disable the Checking of Untrusted SSL Certificates

This step is required if the CloudCenter Suite certificate is self-signed or if there are problems authenticating the Owner API credentials.

To disable the untrusted SSL certificate check, follow this procedure.

1. Using the ServiceNow Filter Navigator, type **sys_properties.list** and press **Enter**.
2. Update the following two properties.

The screenshot shows a search interface with a filter set to 'Name'. The search results table is as follows:

	Name	Value	Type	Application
<input type="checkbox"/>	com.glide.communications.httpclient.veri...	false	true false	Global
<input type="checkbox"/>	com.glide.communications.trustmanager_tr...	true	true false	Global

- Search for the **com.glide.communications.trustmanager_trust_all** property and set the value to **true**.
- Search for the **com.glide.communications.httpclient.verify_hostname** property and set the value to **false**.

Configure Connection with CloudCenter Suite

This step is required to establish communication from ServiceNow to Cisco CloudCenter Suite. You can use one of two methods to perform this task: an *automated* method or a *manual* method.

Automated Method

The automated method requires you to provide CloudCenter Suite admin credentials or a tenant admin credentials in order to retrieve and use that user's API key. The retrieved API key is then populated in the **Tenant Configurations** module in ServiceNow.

This method generates a new API key and replaces the old API key in the CloudCenter Suite for that specific user. If there is a risk that this API key is used by other external system, use the manual method (described below) instead.

To setup a connection with a CloudCenter Suite instance using the automated method, follow this procedure.

- Navigate to **Cisco CloudCenter Suite > Configuration > Tenant Configuration Form**.
- Provide the CloudCenter URL, Username, Password, and Tenant Name of the CloudCenter Suite user with the Tenant Administrator role (displayed in the following screenshot).

The screenshot shows the 'CloudCenter Tenant Configuration' form. It is divided into two main sections: 'CloudCenter Connection Setup' and 'Tenant Mapping'. The 'CloudCenter Connection Setup' section includes fields for 'CloudCenter URL', 'Username', 'Password', 'Tenant Name', and 'MID Server'. The 'Tenant Mapping' section includes a 'Company Name in ServiceNow' field. A 'Submit' button is located at the bottom right. A red text note at the bottom says 'Please fill in the required fields.'


The above action will create a new record in the Tenant Configuration module or override the existing record if an entry with the same CloudCenter URL for the same CloudCenter user or Tenant already exists.

- Validate the Tenant Configuration.
 - Navigate to Cisco CloudCenter Suite > **Configuration > Tenant Configurations**.
 - Check if the Tenant Configuration table records match your requirements.

Manual Method

To setup a connection with a CloudCenter Suite instance using the manual method, follow this procedure.

1. Navigate to **Cisco CloudCenter Suite > Configuration > Tenant Configurations** (displayed in the following screenshot)



2. Enter the CloudCenter URL, API key, and Tenant Short Name. Validate the credentials and then Submit your request.



- If the **Validate Credentials** step fails, confirm that there are no firewall rules blocking traffic between ServiceNow and CloudCenter Suite. In addition, you may need to install a ServiceNow MID server on the network where CloudCenter Suite is installed.
- On the **Tenant Configuration Form**, the **Tenant Name in ServiceNow** field is dynamically translated to a value that matches the current configuration of the application. To better understand this topic, see the *Connecting multiple CloudCenter Suite Instances or tenants* section of this document.
- Checking the **Periodic VMs update** option will enable a scheduled job that runs every 1 hour. It updates the CMBD, the status and cost of all active VMs and related child entities, including imported VMs.

Create an Integration User

In ServiceNow, create a user called **cloudcentersuite.integration** (or with any other appropriate User ID) using the details provided in the following table. Make a note of the User ID and password as this information is later used in CloudCenter Suite's ServiceNow Extension setup.



Make a note of the User ID and password because this information will later be used in Cisco CloudCenter Suite's ServiceNow Extension setup.

ServiceNow Field	Description
User ID	Enter cloudcentersuite.integration in this field.
First name	Enter CloudCenterSuite in this field.
Last name	Enter Integration in this field.
Password	Provide a password that is acceptable to your organization.
Email address	Not required for this user.
Web service access only	Check this box.
Internal Integration User	Check this box.
Role	Add this user to the x_cqt_c3_frame.admin role.

This completes the configuration on the ServiceNow side. However, the following steps to create users in ServiceNow are needed to validate the integration. You can delete these users once the validation is complete.

Create a Test User

In ServiceNow, create a user called **test.requester** (or with any other appropriate User ID) using the details provided in the following table.

ServiceNow Field	Description
User ID	Enter test.requester in this field.
First name	Enter Test in this field.
Last name	Enter Requester in this field.
Password	Provide a password that is acceptable to your organization.
Email address	Enter test.requester@yourdomain.com as this field is required for this user.
Role	Add this user to the x_cqt_c3_frame.ccs_service_portal



The overall configuration steps require that you have administrative access to both the Cisco CloudCenter Suite and ServiceNow instances.

Create a ServiceNow Extension

In CloudCenter Suite's Workload Manager, navigate to **Admin > Extensions**. Add a New Extension and configure the parameters as described in the table below and displayed in the following screenshot.

CloudCenter Suite Field	Description
Extension Type	Select ServiceNow from the dropdown menu.
Extension Name	Enter any unique name for the extension.
ServiceNow URL	Enter the URL of the ServiceNow instance being integrated.
Username	Enter the User ID of the integration user (cloudcentersuite.integration) created in ServiceNow.
Password	Enter the password of the cloudcentersuite.integration user created in ServiceNow.

New Extension

Successfully connected to ServiceNow.

Connection Settings

EXTENSION TYPE
ServiceNow

* EXTENSION NAME
ServiceNow

* SERVICENOW URL
https://dev100.service-now.com

* USERNAME
cloudcenter.integration

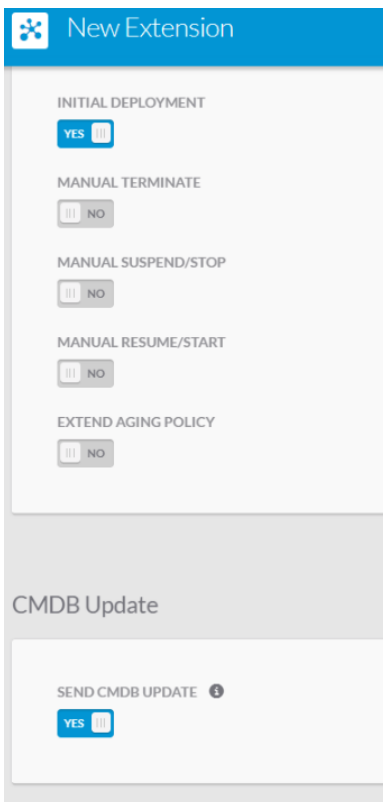
* PASSWORD
.....

[EDIT](#) ✔ CONNECTED

Enable CMDB

To enable CMDB, follow this procedure.

1. Once communication is successfully established between Cisco CloudCenter Suite and ServiceNow, scroll down to view the additional options.



2. Enable the toggle for **Send CMDB update** – enabling this toggle populates the ServiceNow CMDB with deployment information. The following CMDB tables in ServiceNow will be updated:

- Jobs (custom table part of this application)
- Virtual Machines (Out-of-box)
- Network Adapters (Out-of-box)
- Storage Volumes (Out-of-box)
- IP Addresses (Out-of-box)

A CMDB update is triggered when the following events complete:

- **Deployments:**
 - Deploy
 - Scale
 - Stop
 - Terminate
 - Suspend
 - Resume
 - Error
 - Migrate
 - Upgrade
- **Managed VMs:**
 - Start
 - Stop
 - Terminate
 - Reboot
 - Resize
 - Sync VM information
 - Attach/Detach Volumes
 - Error

Associate ServiceNow Extension to a Deployment Environment

In Workload Manager, navigate to an existing environment or create a new environment as displayed in the following screenshot.

Edit Deployment Environment DEV

General Settings

* NAME
DEV

DESCRIPTION
Development environment

SERVICENOW EXTENSION ⓘ
SNOW

APPROVAL REQUIRED TO DEPLOY TO THIS ENVIRONMENT ⓘ
NO

Within the environment's **General Settings** section, use the **ServiceNow Extension** dropdown to select the ServiceNow Extension that you created earlier.



When you associate a ServiceNow extension with a Deployment Environment, Cisco CloudCenter Suite's deployment approval toggle switch for this environment is automatically disabled.


This completes the integration configuration in Cisco CloudCenter Suite.



Before you begin, in CloudCenter Suite Workload Manager, make sure that at least one Application Profile and the Environment associated with ServiceNow Extension is shared with **All users in my tenant**.

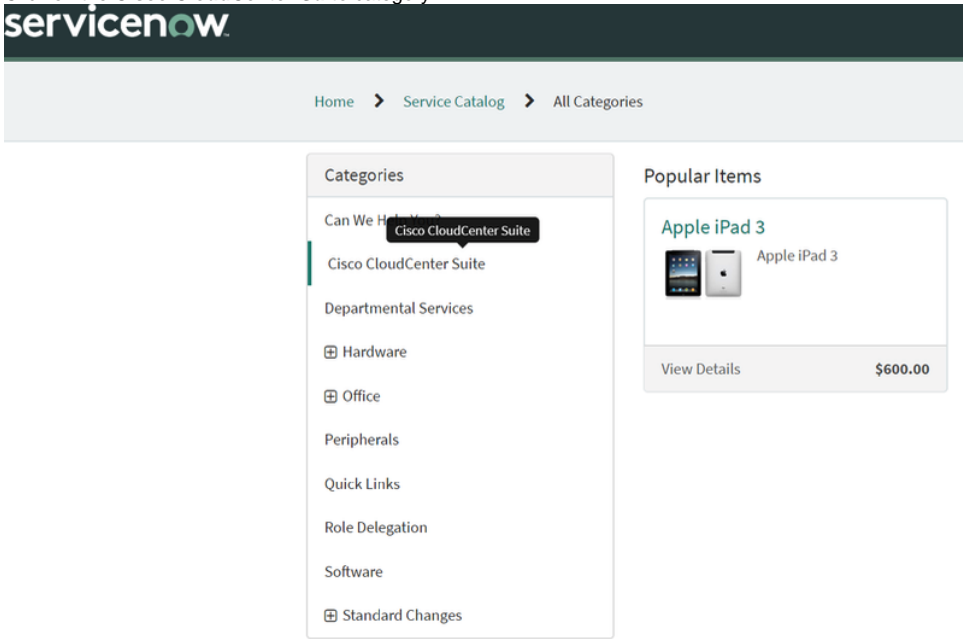
To validate the integration, follow this procedure.

1. Login to ServiceNow as *test.requester*. This is the user account that was created during the setup process.
2. Navigate to the ServiceNow Service Portal by updating the URL to <https://<hostname>.service-now.com/sp>. The out-of-the-box Service Portal page should be displayed.

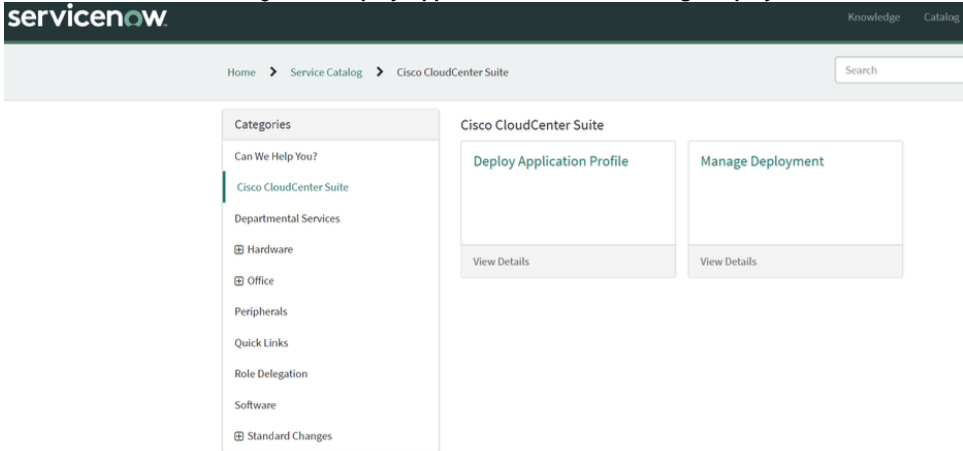


3. Click on the **Request Something** link. Alternatively, you can also type **CloudCenter** in the search bar.

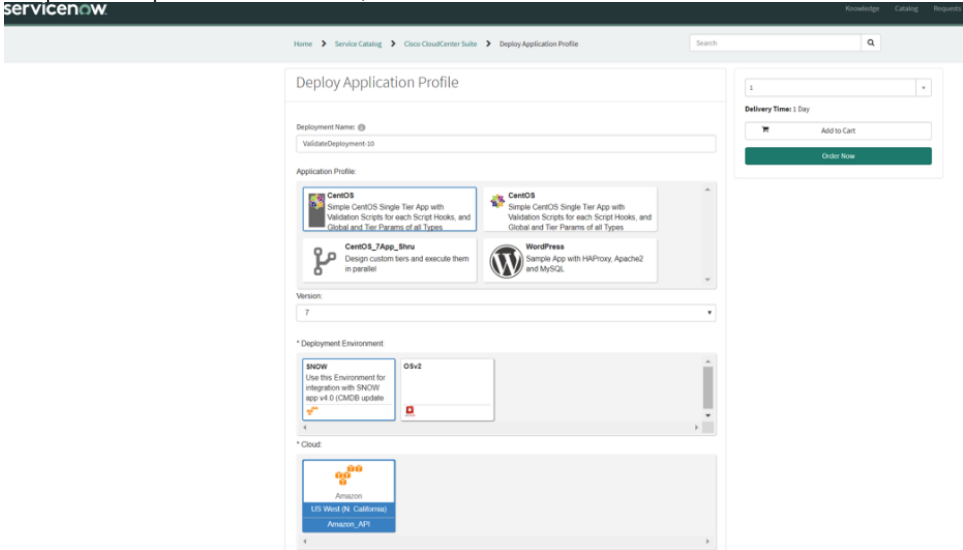
- Click on the **Cisco CloudCenter Suite** category.




- You should now see 2 catalog items: **Deploy Application Profile** and **Manage Deployments**.



- Click **Deploy Application Profile** to load the deployment order form.
- Complete the required fields on the form, and then click on **Order Now**.

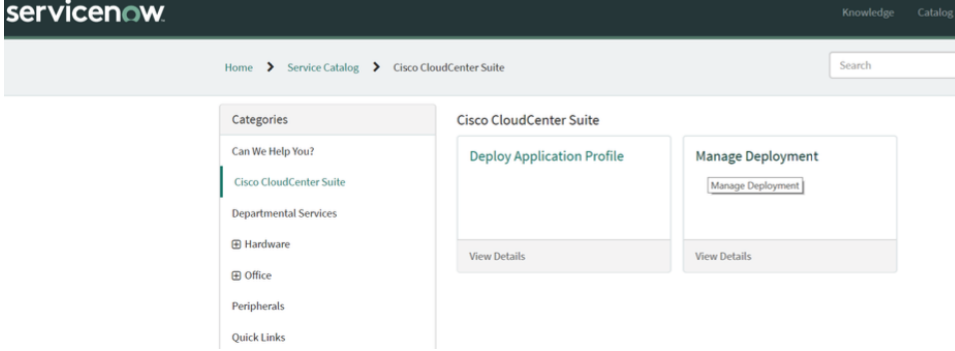


8. An order confirmation page displays.




9. Use the breadcrumb link on top of the page to return to **Home**.

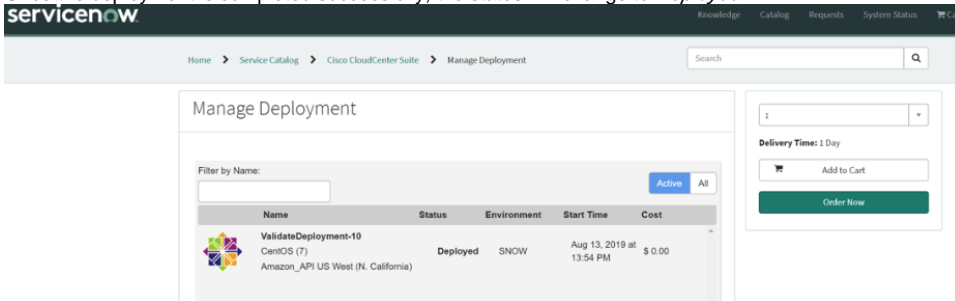
10. Now navigate back to **Manage Deployment** to view the deployment status.




11. The deployment should be listed with the current status.




12. Once the deployment is completed successfully, the status will change to *Deployed*.




- Once deployed, select the deployment by clicking on it once. Then click **Actions available for <deployment name>**. You should see the option to Suspend and Terminate.




- Click **Show VMs** and then select the VM by clicking on it once. You should see several VM options, and the ability to either SSH or RDP into the VM based on the Operating System.




- Switch out of the Service Portal.
- Using the ServiceNow Filter Navigator, navigate to **My Requests**. Then select your Request and view the details.



17. While remaining on the same Request, click on the **Requested Item** link and view the details.




18. Using the ServiceNow Filter Navigator, navigate to **CloudCenter Suite > CMDB – Virtual Machines > Jobs**. The deployed parent and child Job records should be listed.




19. Click on the parent Job Name and view the details.



20. While remaining on the same Job details page, click the **Related records** tab. The child Job should be listed.



21. View details of the child Job.






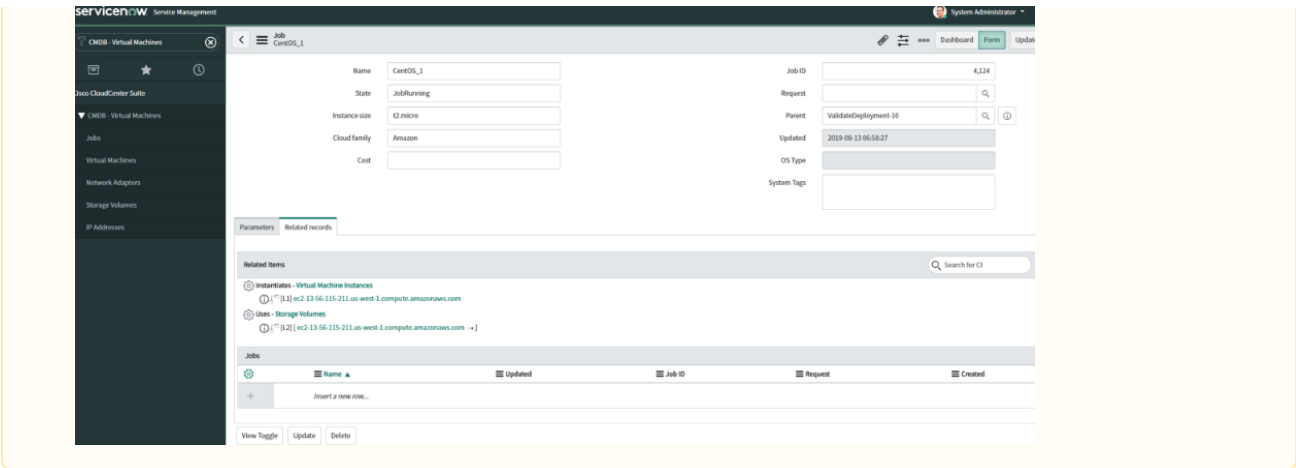
The following CMDB modules in ServiceNow should be updated with deployment details.

- a. Jobs
- b. Virtual Machines
- c. Network Adapters
- d. Storage Volumes
- e. IP Addresses



Items b - e listed above at out-of-the-box tables in ServiceNow and require elevated permission in order to view the records. Try adjusting the user permission as necessary based on your need. If the user is given appropriate permission, records in those tables are visible. The following screenshots provide an example of additional details visible to the administrator user.

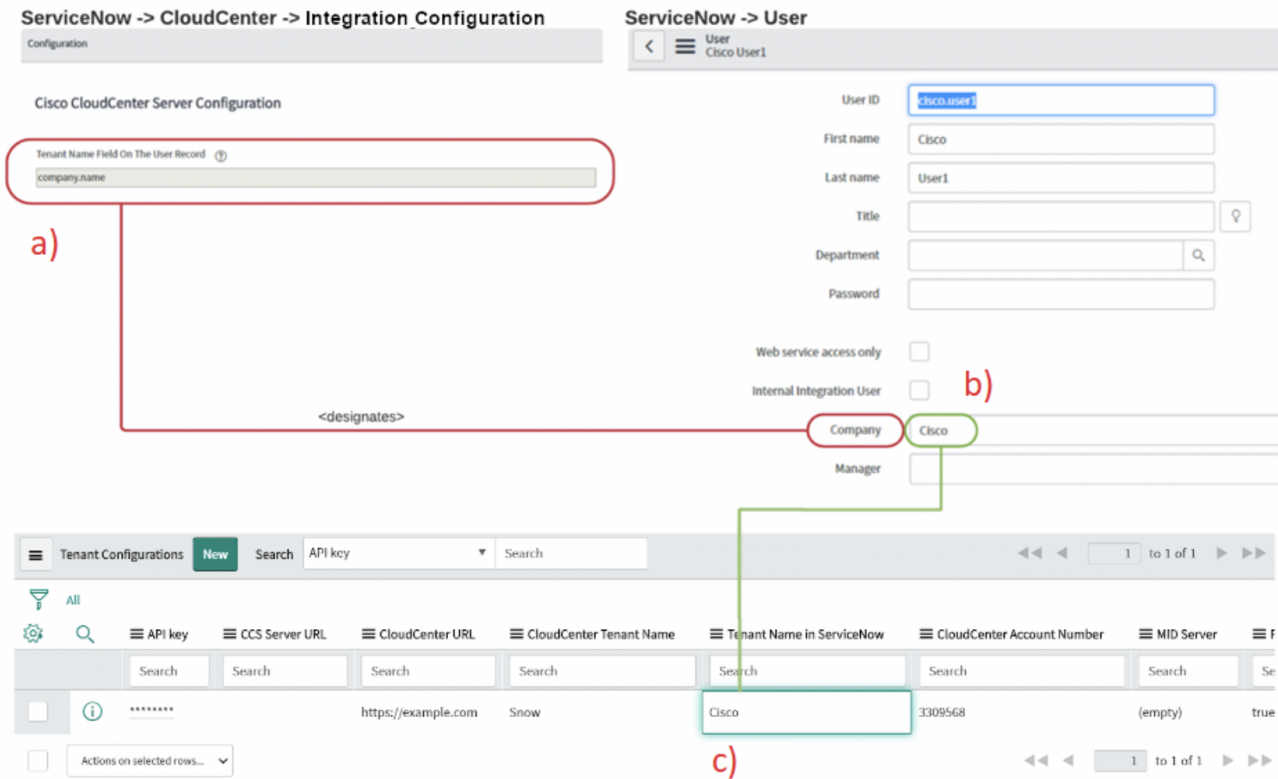




Connecting to Multiple CloudCenter Suite Instances or Tenants

It is possible to create a Tenant Configuration for more than one tenant or instance of the CloudCenter Suite. To do that, you can add another entry in the Tenant Configurations table using a **Tenant Configuration Form** module.

A requester will be matched to a specific tenant defined in Tenant Configurations table based on the value of the *Tenant Name in ServiceNow* field. The image below shows the relationships between configuration and user data.



When a user executes an action that requires connection to CloudCenter Suite (for example, visits a CloudCenter Catalog item or if an approval is required), the system (based on the sample data shown in the above image):

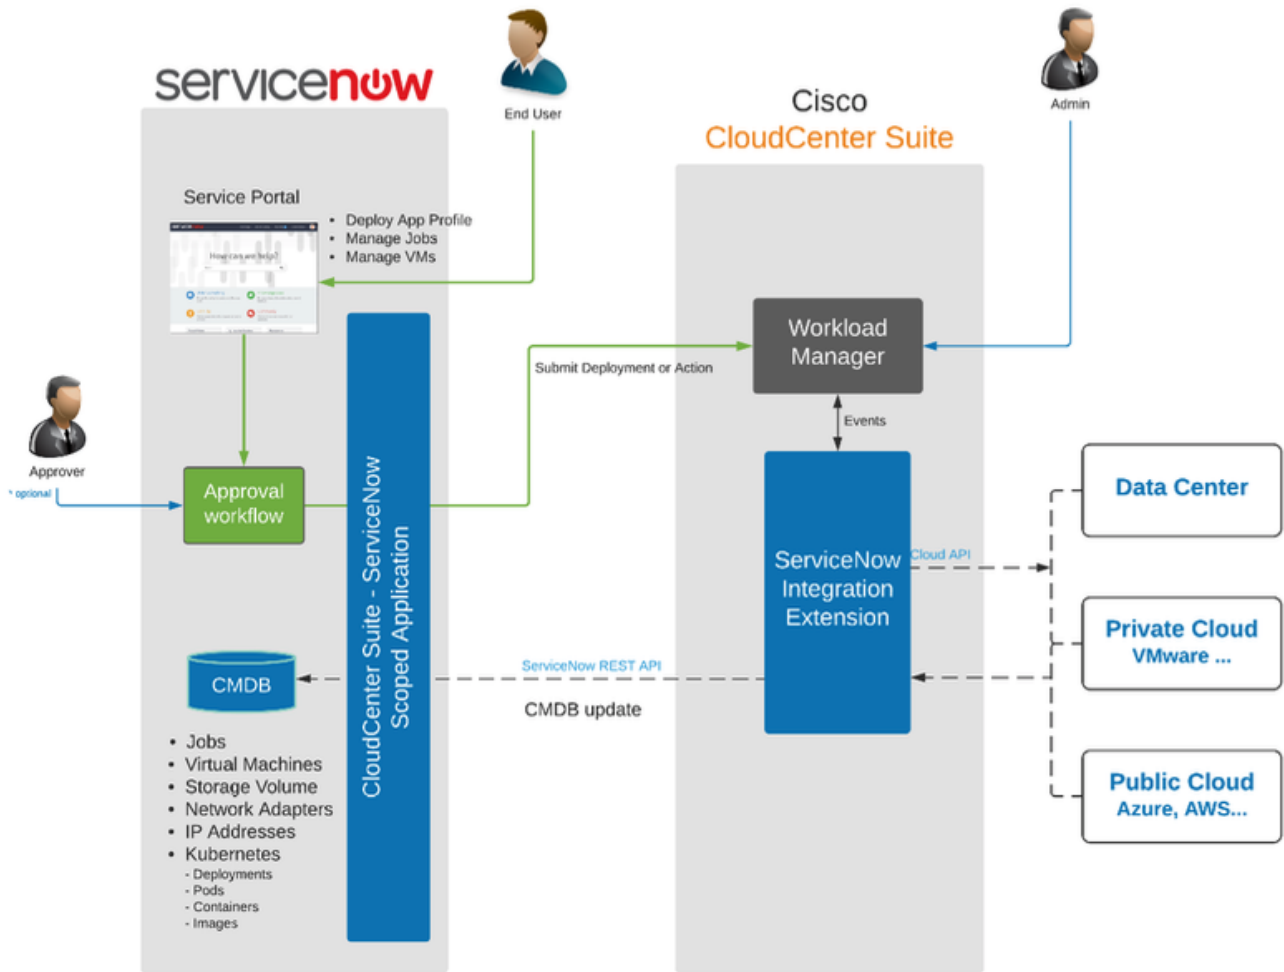
1. Gets the "Tenant Name Field On The User Record" property value in the Integration Configuration = "company.name" (identified as **a**) in the image above).
2. Reads the "Company Name" = "Cisco" (identified as **b**) in the image above).
3. Searches "Tenant Name in ServiceNow" field in "Tenant Configurations" for "Cisco" (identified as **c**) in the image above).

If the system cannot find a direct match, it will try to search a record with wildcard character *.

High-level Interaction

The following image displays the high-level interaction between Cisco CloudCenter Suite and ServiceNow.

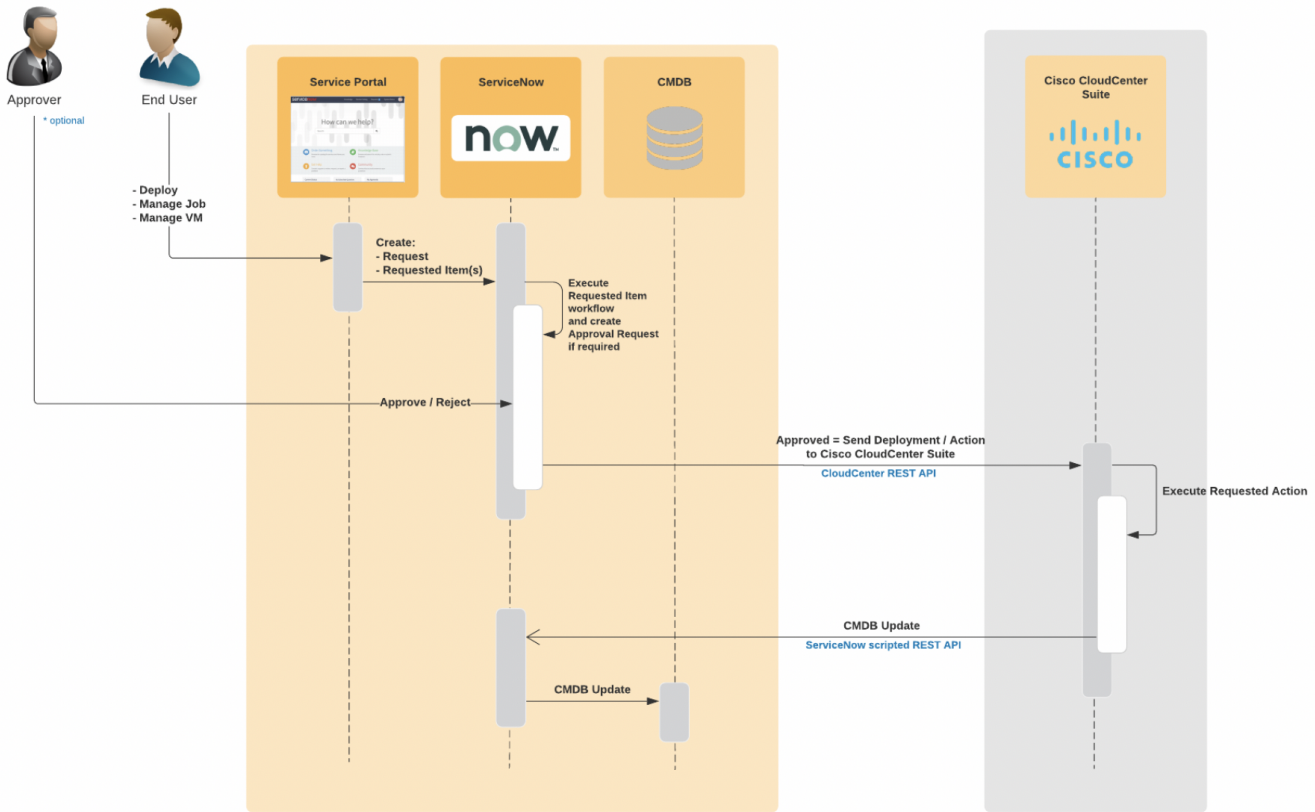
CloudCenter Suite - ServiceNow Integration



© 2016 Cisco and/or its affiliates. All rights reserved. Cisco Confidential

Sequence Diagram

The following image displays the sequence to request a new deployment or Job/VM action.



User Roles


The following table describes the user roles that are defined in the integration application.

Role	Description
x_cqt_c3_frame.consumer	Enables access to CloudCenter Suite application menu and modules: Requests, Approvals, CMDB, Deploy App, App Deployments.
x_cqt_c3_frame.ccs_service_portal	Contains x_cqt_c3_frame.consumer role and enables access to CloudCenter Suite Category and Catalog Items in Service Portal.
x_cqt_c3_frame.admin	Contains x_cqt_c3_frame.consumer role and enables access to Base Configuration menu.


Catalog Item – Deploy Application Profile

- This order form is dynamically loaded from CloudCenter Suite with respect to all user security settings in CloudCenter Suite. User can see only the environments, application profiles and clouds that are shared for a tenant or the user in CloudCenter Suite.
- The order form contains a number of parameters that lets the user configure their deployment. These parameters include:
 - Application Profile
 - Environment and Cloud
 - Metadata
 - System Tags
 - Global Parameters
 - Tier settings
 - Instance size
 - Deployment parameters
 - Number of nodes
 - Volumes
 - Assigning of Public IP
 - Tags
 - SSH Options
- Form data is cached to improve performance. Changes in CloudCenter Suite may not be reflected immediately.

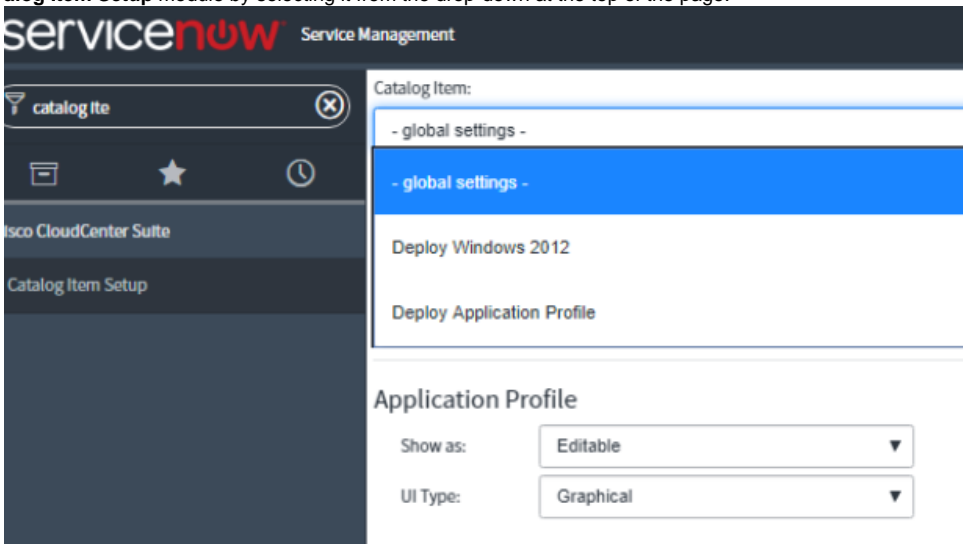
- Similar to other Service Portal catalog items, this item can be added to or edited in a Cart.
- After the order is submitted, a Request and Requested Item is created in OOTB ServiceNow tables.
- The order form can be customized by an administrator based on their preference or requirements. Visibility and appearance of the form sections can be changed by navigating to **CiscoCloudCenter Suite > Configuration > Catalog Item Setup**.



- By copying the provided **Deploy Application Profile** catalog item, the administrator can create additional catalog items in Service Portal with different appearance configuration, default values and accessible to users with a different role.

 *Never modify the original **Deploy Application Profile** catalog item. Make a copy of the item and deactivate the original one if needed.*

After copying the item, an administrator can add/remove required user roles, select a different workflow and configure it separately through the **Catalog Item Setup** module by selecting it from the drop-down at the top of the page.



- The image below shows an extreme example of simplifying the deployment of Windows Server 2012. This can be achieved by copying the original **Deploy Application Profile** catalog item, configuring it through the **Catalog Item Setup** module in ServiceNow and appropriately configuring the Application Profile in CloudCenter Suite. It takes administrative effort once, but the end-user experience can be much simplified.

Home > Service Catalog > Cisco CloudCenter Suite

Search

Categories

- Cisco CloudCenter Suite
- Departmental Services
- Role Delegation

Cisco CloudCenter Suite

- Deploy Application Profile
- Deploy Windows 2012
- Manage Deployment

View Details

View Details

View Details

Deploy Windows 2012

Deploy a Virtual Server with CloudCenter Suite

Application Profile:

Centos - Design custom tiers and execute them in parallel

Version:

6.x

Approximate Price: \$0.11 / month

1 Submit Add to Cart

Add attachments

Catalog Item – Manage Deployments

This catalog item allows a user to view and execute actions on existing CloudCenter Suite deployments.

Users can view, execute actions on Job or VMs accordingly to CloudCenter Suite security settings.

The supported actions are:

Job actions:

- Terminate
- Suspend
- Resume

VM actions:

- Stop
- Start
- Reboot
- Terminate
- Attach Volumes
- Detach Volumes
- SSH, RDP and Windows Password (Guacamole server must be running and accessible from user network)

API Cache

Data is cached in dedicated tables in ServiceNow to improve performance of ServiceNow to CloudCenter Suite communication. Two levels of cache are implemented.

- Level 1 - Server-side cache. Default TTL for this cache is 500 minutes and it can be manually deleted by an administrator from table 'x_cqt_c3_frame_api_cache' or by going to the URL https://<instance>.service-now.com/nav_to.do?uri=%2Fxcqt_c3_frame_api_cache_list.do. Default TTL for this cache can be modified by changing a value of system property "x_cqt_c3_frame.api_cache_ttl" (in minutes).
- Level 2 – Client-side cache. Data is cached in browser session (per tab) to improve performance of browser to ServiceNow communication. To clear this cache, close current tab and open a new one (refreshing a page does not clear the cache).

Image Cache

The image cache stores images associated with Application Profiles and Services that can be displayed on the catalog item forms. This cache has a very long TTL (as images are rarely changed). However, it can be cleared by an administrator by deleting entries in the table 'x_cqt_c3_frame_image_cache' or by going to the URL https://<instance>.service-now.com/nav_to.do?uri=%2Fxcqt_c3_frame_image_cache_list.do

Workflows

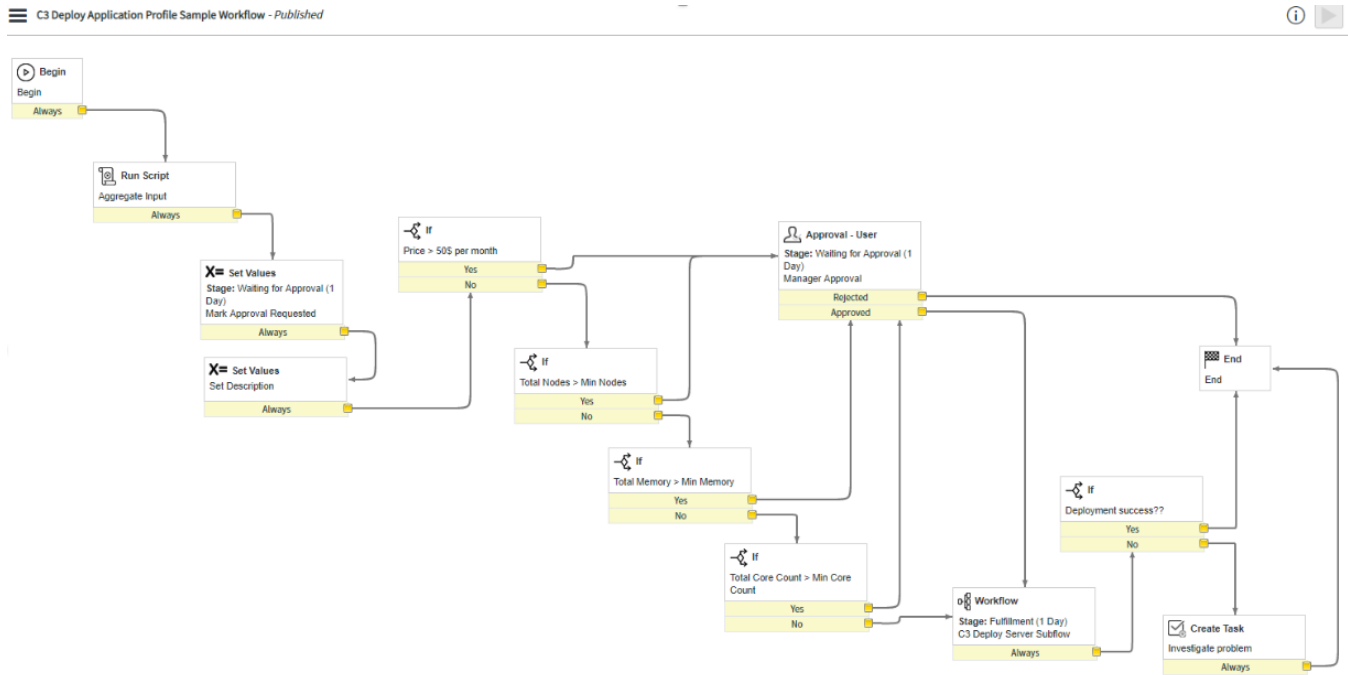
Never change the provided workflows. Copy the workflow and modify it as needed. Then select the new workflow for CloudCenter Suite catalog items under **Service Catalog > Maintain Items** module.

An important and required activity of the approval workflows is sub-flows: **C3 Deploy Server Subflow** and **C3 Manage Virtual Server Subflow**. These sub-flows return *true* or *false* to a parent flow to inform about success or failure, and also updates the requested item Activity log.

Do not customize the **C3 Deploy Server Subflow** and **C3 Manage Virtual Server Subflow** sub-flows.


The integration application provides a sample approval workflow for new deployments called 'C3 Deploy Application Profile Sample Workflow'.

The following image displays the Approval workflow associated with **Deploy Application Profile**.



The integration application provides a sample approval workflow for managing deployments called **C3 Manage Deployment Sample Workflow**.

The following image displays the approval workflow associated with **Manage Deployments**.



The following ServiceNow tables are created or updated by the integration application.

- 1. If a Manager is not assigned to the test.requester user, the deployment Request record in ServiceNow will not be created and the default approval workflow will auto approve the Request.

2. If a user logs directly into CloudCenter Suite (Workload Manager) and makes a deployment into an Environment that requires ServiceNow approval, and if this user does not exist in ServiceNow (matching email ID), the deployment will remain in Pending state.

Title	OOB SN table	Description
Owners API Keys (x_cqt_c3_frame_owner_api_keys)	No	Stores IDs and API keys of CloudCenter Suite admin users.
User API Credentials (x_cqt_c3_frame_user_creds)	No	Stores IDs and API keys of CloudCenter Suite users (requesters).
Approval workflows (x_cqt_c3_frame_approval_workflows)	No	Configuration table. Assigns workflows for specific request types.
Requested Items (sc_req_item)	Yes	Contains all Requested Items created with Service Portal catalog items related to CloudCenter Suite.
Requests (x_cqt_c3_frame_request)	No	Contains all approval requests sent directly from CloudCenter Suite.
Jobs (x_cqt_c3_frame_job)	No	Contains information about deployments in CloudCenter. <ul style="list-style-type: none"> Parent jobs (with empty "Parent" filed) stands for deployments in CloudCenter Suite. Child jobs (with "Parent" filed set) corresponds to the particular tier of deployments (Services).
Virtual Machines (cmdb_ci_vm_instance)	Yes	Corresponds directly to particular virtual machines (nodes) deployed in the cloud. If multiple nodes are selected in CloudCenter during deployment, Job may consist more than one Virtual Machines.
Network adapters (cmdb_ci_network_adapter)	Yes	Network adapters associated with VMs
IP addresses (cmdb_ci_ip_address)	Yes	IP addresses associated with Network Adapters / VMs
Storage Volumes (cmdb_ci_ip_address)	Yes	Storage volumes associated to VMs
Requested Item (x_cqt_c3_frame_requested_item)	No	Contains information about requested deployments if request was generated from "CloudCenter Integration UI" plugin.
K8s Deployment (x_cqt_c3_frame_k8s_deployment)	No	Contains information about Kubernetes deployments in CloudCenter Suite.
K8s Pod (x_cqt_c3_frame_k8s_pod)	No	Pods associated with Kubernetes deployments.
K8s Container (x_cqt_c3_frame_k8s_container)	No	Containers associated with K8s pods.
K8s Image (x_cqt_c3_frame_k8s_image)	No	Kubernetes Images.
Tenant Name Mapping (x_cqt_c3_frame_tenant_name_mapping)	No	Stores information about tenant mapping between ServiceNow and CloudCenter Suite.
CloudCenter Order (x_cqt_c3_frame_order)	No	Not used.
Ordered Deployments (x_cqt_c3_frame_ordered_deployment)	No	Not used.
Image Cache	No	Cache table for images like Application Profiles and Services icons

(x_cqt_c3_frame_image_cache)		
API Cache (x_cqt_c3_frame_api_cache)	No	Cache table for REST API "get" request to CloudCenter Suite
Catalog Item Config (x_cqt_c3_frame_item_cfg)	No	Contains information about configuration of particular Catalog Items related to CloudCenter Suite
PriceLookup (x_cqt_c3_frame_pricelookup)	No	Table with constant lookup values required to calculate and store the price of requested items. * DO NOT MODIFY RECORDS IN THIS TABLE.

ACI Multi-Site Extensions

ACI Multi-Site Extensions

- [Overview](#)
- [Requirements](#)
- [Extensions Page](#)
- [Create an ACI Multi-Site Extension](#)
- [Referencing an Extension in a Deployment Environment](#)
- [Troubleshooting](#)

Workload Manager supports integration with the Cisco ACI Multi-Site endpoint through the definition of a ACI Multi-Site extensions. These extensions may be created and viewed from the Workload Manager Extensions page. Once an extension is created and validated, it can be referenced in the Deployment Environment form in the Cloud Settings subsection for a vCenter region or in a network mapping for a vCenter region. Once the extension is referenced in a deployment environment, when the user deploys an application to that environment using the one of the vCenter regions where the Multi-Site endpoint is deployed, the ACI Multi-Site related cloud settings are applied to the deployment at deploy time. Depending on how the the environment is set up, the user may be able to see and modify the cloud settings at deploy time. Workload Manager leverages Cloud Remote when connecting to an Multi-Site endpoint that is not directly accessible to the CloudCenter Suite cluster.

The integration allows the Workload Manager user to specify the following ACI resources during a deployment:

- Application Network Profiles (ANP)
- Endpoint Groups (EPG)
- Contracts
- Subjects/Filters

Be aware of the following requirements to use Multi-Site extensions:

- ACI Multi-Site Requirements
 - ACI Multi-Site 2.0(2c) and 2.1(1i).
 - All of the sites configured within ACI Multi-Site must adhere to the APIC requirements specified in the [ACI Extensions](#) section.
- Workload Manager Requirements:
 - If Multi-Site endpoint not directly accessible from CloudCenter Suite cluster you must use the cloud region that has the applicable Cloud Remote configured.
- vCenter Requirements: The following table describes the VMware vCenter requirements.

Requirement	Details
A working VMware vCenter 5.0/5.5/6.0 environment	The minimum VMware vSphere version is v5.0, but vSphere v5.5 U2 is optimal.
Workload Manager automates the provisioning of virtual machines into the VMware private datacenter.	Workload Manager requires <i>access credentials</i> to the vCenter setup.
All ESX host(s) must be physically connected to the ACI leaf switches.	The prerequisite installation requirements for the datacenter are: <ul style="list-style-type: none"> • A physical ESX host capable of running at least 10 medium sized instances • An ESX cluster (cluster could comprise of just the one host) • A datastore (or datastore cluster for DRS support), at least 100gb of free space
If the ESXi hosts are Cisco UCS based	<ul style="list-style-type: none"> • The VLANs for the CMM must be mapped to the vNIC template. • The uplinks from the Fabric must interconnect trunking VLANs to the leaf switches.

The Extensions page can be reached directly from the Admin menu. This is where you can see all existing ACI, ACI Multi-Site, and ServiceNow extensions that were created by you or shared with you. From here you can perform actions on an existing extension, depending on your permissions, and also create new extensions. A screenshot of the Extensions page which includes ACI Multi-Site extensions is shown below.

Extensions

ADD EXTENSION

NAME	LAST UPDATED	EXTENSION URL	TYPE	BLADE STATUS	ACTIONS
testMSOExtension	Jul 1, 2019 07:09 PM	http://192.168.131.20/	ACI Multi-Site	INVALID_CREDENTIALS	
Demo_Revalidate	Jul 1, 2019 05:14 PM	https://192.168.131.119	ACI Multi-Site	SUCCESS	
InvalidEndpointDemo	Jul 1, 2019 05:10 PM	https://192.168.131.119	ACI Multi-Site	INVALID_CREDENTIALS	
MSO_Extn_1	Jul 1, 2019 04:48 AM	https://192.168.131.119/	ACI Multi-Site	SUCCESS	
APIC_02	Jun 30, 2019 04:16 PM	https://192.168.131.10	ACI	NOT_APPLICABLE	
APIC_01	Jun 30, 2019 04:15 PM	https://192.168.131.10	ACI	NOT_APPLICABLE	

The following table summarizes the actions that can be performed from the Extensions page:

Action	How Accessed	Notes
View or edit an extension	Click on a corresponding row in the list of extensions	The extension edit form is similar to the extension creation form.
Create a new extension	Click the Add Extension button in the upper right of the screen	See <i>Create an ACI Multi-Site Extension</i> section below for more details
Share an extension	Click in the Actions column for the appropriate row and select Share	See Permission Control > Extension Permissions
Delete an extension	Click in the Actions column for the appropriate row and select Delete	if an ACI Multi-Site extension is referenced in a deployment environment or deployment, it cannot be deleted.
Re-enable an extension	Click in the Actions column for the appropriate row and select Re-enable	This action is only available for extensions where the Blade Status is Blade Failure. This causes the pod associated with the failed blade to be regenerated.

The Blade Status column on the Extensions page applies to ACI Multi-Site extensions only. Possible values for Blade Status are summarized in the following table:

Blade Status	Meaning
Blade Failure	The pod required to support the extension was not successfully launched
Endpoint Inaccessible	The Multi-Site API endpoint could not be accessed
Invalid Credentials	The Multi-Site API endpoint could be accessed but the user credentials could not be validated
Success	The Multi-Site credentials were accepted and connection established

To create a new ACI Multi-Site extension, follow this procedure.

1. From the Extensions page, click **Add Extension** in the upper right. The New Extension page displays, as shown in the following screenshot.

The screenshot shows the 'New Extension' page with the following fields and values:

- EXTENSION NAME:** ACIMultiSite_Prod
- EXTENSION TYPE:** ACI Multi-Site (selected from a dropdown menu)
- APIC CONTROLLER URL:** https://192.168.131.119

2. Enter an extension name and select ACI Multi-Site from the dropdown menu. This causes the rest of the form to update and display the ACI Multi-Site specific fields as show in the following screenshot.

The screenshot shows the 'New Extension' page with the following fields and values:

- EXTENSION NAME:** ACIMultiSite_Prod
- EXTENSION TYPE:** ACI Multi-Site
- MSO CONTROLLER URL:** https://192.168.131.119
- USERNAME:** admin
- PASSWORD:** [Masked]
- DOMAIN ID:** 0000FFFF000000000000000090
- IS MSO DIRECTLY ACCESSIBLE VIA END POINT:**

3. Enter the following required fields:

- Multi-Site endpoint URL
- Multi-Site endpoint user name and password

4. The Domain ID field is always prepopulated with the domain ID of the default local domain. Update this field if you want to use the domain ID of another domain defined in the Multi-Site endpoint.
5. If the Multi-Site endpoint endpoint is directly accessible to the Workload Manager cluster, click **Save** to save the configuration. This also causes Workload Manager to connect with the Multi-Site endpoint and login using the provided credentials. The status of the connection attempt is displayed in the Blade Status column of the Extensions page. Once the connection is successful, a success message is briefly displayed at the top of the Extensions page.

6. If the Multi-Site endpoint is not directly accessible to the Workload Manager cluster, set the Multi-Site Endpoint **Directly Accessible** toggle to **No**. This causes the a new dropdown field, VMware cloud region, to be displayed as shown in the the following screenshot.

Select the vCenter region where the Multi-Site endpoint is hosted, then click **Save**. This causes Workload Manager to connect with the Multi-Site endpoint through the Cloud Remote appliance in that region.

7. After clicking **Save**, the Extensions page refreshes to display the newly added extension appears at the top of the list of extensions.


An ACI Multi-Site extension can only be used when it is referenced in an deployment environment containing a vCenter region hosting the Multi-Site endpoint referenced in the extension. To do this, create a deployment environment as instructed in [Deployment Environments > Add a Deployment Environment](#), with the following extra steps:

1. From the General Settings tab, make sure you select the vCenter region that is hosting the Multi-Site controller referenced in your extension.
2. From the Cloud Settings tab, either create a new network mapping, or specify the cloud settings for the vCenter region hosting the Multi-Site endpoint. In either case, specify the data center from the dropdown field and all subsequent deployment resources as you would with any other vCenter region.
3. If you have read access to at least one ACI or ACI Multi-Site extension, just above the per NIC settings will be a **Use Extension** toggle which if **Off** by default. Turn this toggle **On**.
4. When you turn on the Use Extension toggle, the display updates with a new section added under the toggle with new extension related fields as shown in the following screenshot.


Click the ACI Multi-Site tab.


5. Clicking the ACI Multi-Site tab displays the ACI Multi-Site related fields as shown in the following screenshot.


USE EXTENSION


ON 


ACI **ACI Multi-Site**

* MSO EXTENSION
Select extension 

* MSO SITE
Select MSO Site 

* VIRTUAL MACHINE MANAGER
Select APIC VMM 

* MSO TENANT
Select Tenant 


L3 OUT
Select L3 Out 


6. Select the *ACI Multi-Site extension* you want to use from the first dropdown field. This causes the list of Multi-Site sites associated with the Multi-Site controller to be available in the Multi-Site site dropdown.
7. Select the *Multi-Site* from the dropdown field. This causes the list of Virtual Machine Managers associated with the Multi-Site to be available in the Virtual Machine Manager dropdown.
8. Select the *Virtual Machine Manager* from the dropdown field. This causes the list of Multi-Site tenants associated with the Virtual Machine Manager to be available in the Multi-Site Tenant dropdown.
9. Select the *Multi-Site Tenant* from the dropdown field. This causes the list of L3 Out values associated with the Multi-Site tenant to be available in the L3 Out dropdown.
10. Select the *L3 Out value* from the dropdown.
11. Scroll down to the per NIC settings. When an ACI Multi-Site extension is associated with the region, a tabbed header appears above the NIC

NIC 1

NETWORK TYPE

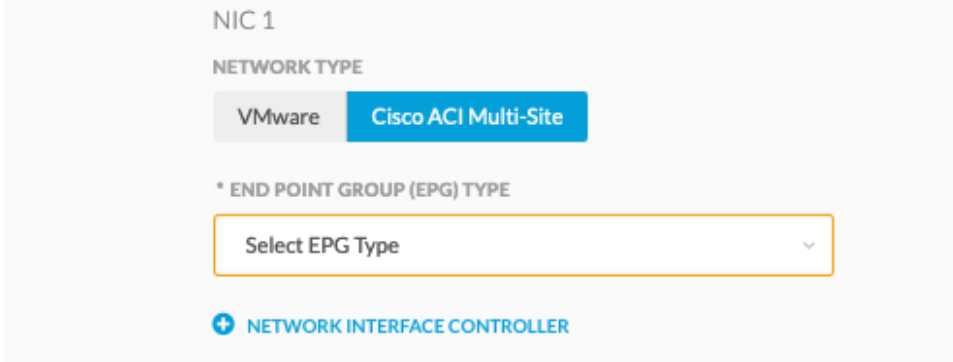
VMware Cisco ACI Multi-Site

* NETWORK
Select Network 

 NETWORK INTERFACE CONTROLLER

settings as shown in the following screenshot.
Select the **Cisco ACI Multi-Site** tab.

12. Selecting the **ACI Multi-Site** tab causes the Network dropdown to be replaced with the EPG Type dropdown as shown in the following screenshot



NIC 1

NETWORK TYPE

VMware Cisco ACI Multi-Site

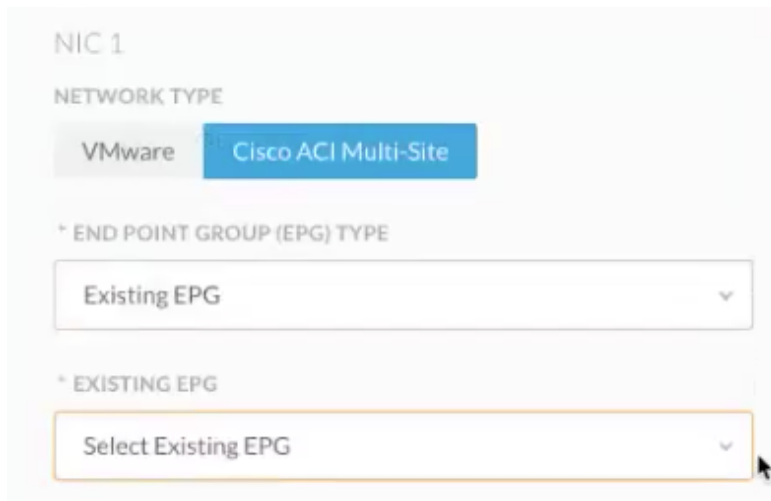
* END POINT GROUP (EPG) TYPE

Select EPG Type

+ NETWORK INTERFACE CONTROLLER

13. You can set the EPG Type dropdown to Existing EPG or New EPG.

- a. If you select Existing EPG, an Existing EPG dropdown field is displayed as shown in the following screenshot.



NIC 1

NETWORK TYPE

VMware Cisco ACI Multi-Site

* END POINT GROUP (EPG) TYPE

Existing EPG

* EXISTING EPG

Select Existing EPG

Select an existing EPG from the dropdown. You are done configuring this NIC.

- b. If you select New EPG for the EPG Type, the Bridge Domain and Contracts dropdown fields are displayed as shown in the following screenshot.

NIC 1

NETWORK TYPE

VMware Cisco ACI Multi-Site

* END POINT GROUP (EPG) TYPE

New EPG

* BRIDGE DOMAIN

CliQr-BD

CONTRACTS

Contract

- Contract
- CloudCenter-dep25-Database-to-Database-1-I3out
- CloudCenter-dep25-Database-to-WebServer
- CloudCenter-dep25-LoadBalancer-to-LoadBalancer...
- CloudCenter-dep25-WebServer-to-LoadBalancer
- CloudCenter-dep25-WebServer-to-WebServer-1-I3...

Select a Bridge Domain, and from the Contracts dropdown, select one or more contracts.

14. Add and configure additional NICs if required.

If you set the **cliqrIgnoreAppFailure** parameter (see [Troubleshooting Parameters](#)), then the APIC resources (ANP, EPGs, Contracts, and so forth) created using the Workload Manager are not removed if the deployment fails. The launched VMs and related APIC policies are only removed when the user terminates the deployment from the Deployments page. See [Terminate Protection](#) for additional context.

Jenkins Integration

Jenkins Integration

- [Overview](#)
- [The Workload Manager Jenkins Plugin](#)
- [Prerequisites](#)
- [Install the Workload Manager Jenkins Plugin](#)
- [The jenkinsBuildId Macro](#)
- [Create a New Deployment on Every Build](#)
- [Update an Existing Deployment](#)

For pre-modeled Jenkins projects (for example, Maven, to fetch the source code from Git/SVN), you can integrate with Workload Manager using the Workload Manager Jenkins plugin.

You do not need to manually copy this file, Cisco provides a download URL to make this plugin available to Jenkins users. Contact [CloudCenter Suite Support](#) to obtain the download location.

The Workload Manager Jenkins plugin provides complete integration between Jenkins and Workload Manager by allowing users to directly launch deployments on any [Supported Cloud](#) from a Jenkins server.

Additionally, users can upgrade an existing deployment by specifying upgrade scripts for each tier.

- If you are new to Jenkins, setup a maven project on Jenkins with Github as its source repository. See <http://www.youtube.com/watch?v=ITQGi5jzjvo> for additional details.
- The supported Jenkins versions value to use the Workload Manager Jenkins plugin must be **Jenkins 1.624 and later versions**.
- The required Java version for the Jenkins server must be **Java 8**.
- You must set the default cloud settings in the [Deployment Environments](#) so the Jenkins Plugin can fetch and use those cloud settings when you deploy the application.
- To attach aging/suspension/security policies to a Job deployed by Jenkins, be sure to toggle the switch for the selected policy in the **Deployment Environment > Policy Settings** tab as **Not visible to users** to ensure that the policy is automatically attached to the deployment.

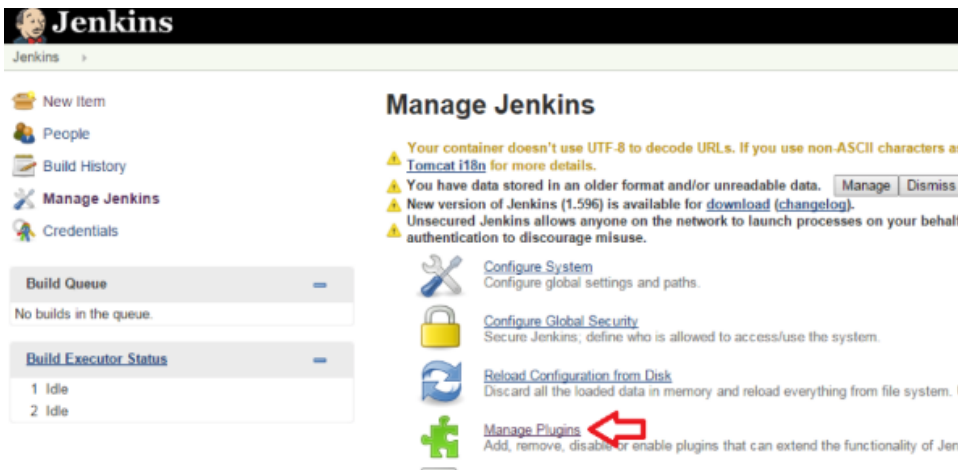
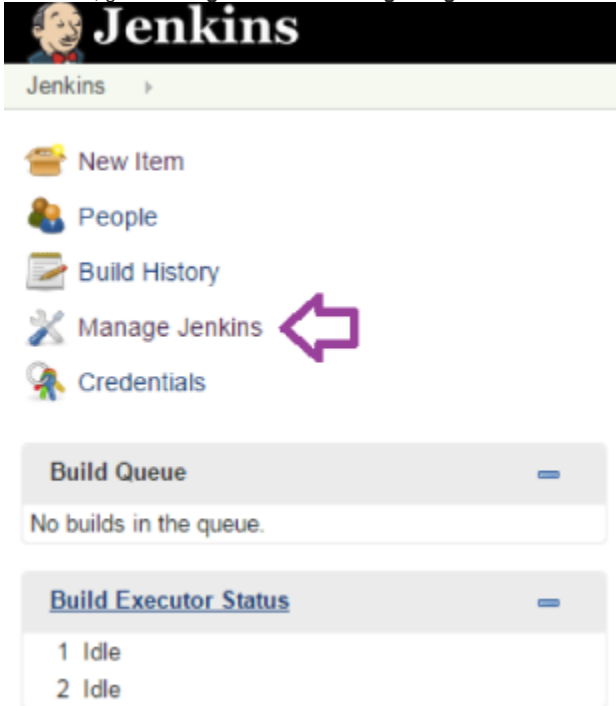
The screenshot shows the 'Edit Deployment Environment NewEnv' interface with three tabs: 'General Settings', 'Cloud Settings', and 'Policy Settings'. The 'Policy Settings' tab is active. It contains three sections: 'Aging Policy', 'Suspension Policy', and 'Security Policy'. Each section has a visibility icon (an eye) on the left. The 'Suspension Policy' section is highlighted with a red box, and a red label 'Not visible to users' with a downward arrow points to the visibility icon. The 'Aging Policy' section shows 'ALLOWED AGING POLICIES' with a dropdown menu set to '1Hour (gsharma6)'. The 'Suspension Policy' section shows 'ALLOWED SUSPENSION POLICIES' with a dropdown menu set to '4thDec (gsharma6)'. Below this is a toggle for 'ALLOW TERMINATE/SUSPEND PROTECTION' set to 'OFF'. The 'Security Policy' section shows 'ALLOWED SECURITY POLICIES' with a dropdown menu set to 'All' and 'PRESELECTED SECURITY POLICIES' with a dropdown menu set to 'sp1 (gsharma6)'.

- The Jenkins plugin for Workload Manager requires the Workload Manager APIs for job deployment).

To install the Workload Manager Jenkins plugin, follow this procedure:

1. Contact [CloudCenter Suite Support](#) to obtain the download instructions or download the installer_artifacts.tar file from software.cisco.com.
2. Log into Workload Manager using your admin credentials.

3. Generate the API Management (Access) Key for the Jenkins user. See [API Key](#) for additional details on generating the API key.
4. [Model the Application](#) so this user can access artifacts from the Jenkins build server.
5. In Jenkins, go to **Manage Jenkins > Manage Plugins > Advanced > Upload Plugin** to upload and install the Workload Manager Jenkins Plugin.



Jenkins > gg > configuration

Build History

[trend](#)

- #2 Jan 7, 2015 7:50 PM
- #1 Jan 7, 2015 7:48 PM

[RSS for all](#) [RSS for failures](#)

Disable build (no new builds will be executed until the project is re-enabled.)

Execute concurrent builds if necessary

Advanced Project Options

Source Code Management

None

CVS

CVS Projectset

Subversion

Build Triggers

Build after other projects are built

Build periodically

Poll SCM

Build

Add build step ▾

- CliQr Application Deployment Client** ←
- Execute Windows batch command
- Execute shell
- Invoke Ant
- Invoke top-level Maven targets

Jenkins

Jenkins > Update center

[Back to Dashboard](#)

[Manage Jenkins](#)

[Manage Plugins](#)

Installing Plugins/Upgrades

Preparation

CliQrJenkinsClient Success

[Go back to the top page](#)
(you can start using the installed plugins right away)

Restart Jenkins when installation is complete and no jobs are running

- After you install the CliQrJenkinsPlugin, go to your existing/new project to configure post-build step and fetch the source code from the Git/SVN using a Maven project into the Jenkins Build. Configure the Workload Manager Application Deployment Client in Jenkins for continuous integration from build system and deployment (new or upgrade on an existing node), as shown in the following screenshot.

The screenshot shows the Jenkins configuration page for the Workload Manager Application Deployment Client. The configuration includes:

- AccessKey:** A text field for the API key, with a "Validate Connection" button.
- Deploy to Project:** A checked checkbox.
- Project Name:** A dropdown menu with "4_7_cicd" selected.
- Project Phase:** A dropdown menu with "47qa" selected.
- Application Name:** A dropdown menu with "PetClinic" selected.
- Application Version:** A dropdown menu with "1.0.2" selected.
- Deployment Environment:** A dropdown menu with "de_47_test" selected.
- Cloud Type:** A dropdown menu with "Azure US West (California)" selected.
- Cloud Account:** A dropdown menu with "CICQQA" selected.
- Instance Type:** A dropdown menu with "Large" selected.
- Tags:** A text field with "stacking_tag47" entered.
- AppParameters:** A text field with "JOB_BUILD_NUMBER=\${PARENT_BUILD_NUMBER}" entered.
- Application Deployment Configuration:** A section with an "Add another Tier Deployment Configuration" button.
- Choose to copy your binaries:** Radio buttons for "Copy Binaries to External Location" (unchecked) and "Don't copy Binaries" (checked).
- Choose the Deployment Behavior:** Radio buttons for "Create a new Deployment on every Build" (checked), "Skip Old Deployment if exists" (unchecked), and "Update Existing Deployment (Create new if doesn't Exist)" (unchecked).
- Wait For Deployment And Export Details:** A checked checkbox.
- Buttons:** "Add build step", "Save", "Apply", and "Delete".

The following table lists the parameters for the Workload Manager Application Deployment Client page.

Parameter	Description
Cisco CloudCenter Suite URL	The the IP address of the CloudCenter Suite UI. Verify that trusted certificates are used.
Username	Username listed in the Manage Access Key section.
AccessKey	The API Management Key for the user listed in the he Manage Access Key section.
Deploy to Project	Use this flag to deploy to a project, instead of a general deployment environment (see Project and Phase Management for additional context).
Project Name	If you select the Deploy to Project flag, this parameter is displayed and the list of projects are fetched from CloudCenter Suite. Based on this parameter value, only applications associated with this project are filtered out in the Application Version field.
Project Phase	Based on the project name, a dropdown list of all Phases in that project are displayed. Use this value in the Deployment Environment field to filter the project.
Application Name	From the dropdown list of applications listed in the Workload Manager UI, select the required application to deploy. <div style="border: 1px solid green; padding: 5px; display: inline-block;"> After entering your credentials, be prepare to wait for some time as the <i>Application Management APIs</i> APIs may take a while to load. </div>
Application Version	Based on your application, select the application version from this dropdown list.
Deployment Environment	Select the required deployment environment to deploy your application. Be sure to verify and check all default settings like default cloud, default instance type, and so forth as Workload Manager uses these default settings for each deployment.
Cloud Type	Select one cloud type from this dropdown list of cloud types that are present in your deployment environment.

Cloud Account	The cloud account with which you deploy the application.
Instance Type	Select a single instance type as the default instance type.
Tags	A comma-separated list of tags associated with this job.
AppParameters	<p>A comma-separated list of key-value pairs to pass as global parameters. For example: abcd=wow, cdef=cliqrRocks</p> <p>CloudCenter includes the \$BUILD_ID, \$BUILD_NUMBER, \$BUILD_TAG, \$JOB_NAME, and if available, \$BUILD_TIMESTAMP from other plugins.</p> <p>You can add a variable in this section to fetch parameters that are shared by other jobs. For example: abc=\${BUILD_PARENT_NUMBER} or abc=\$BUILD_PARENT_NUMBER</p> <p>You also have the option to retrieve parameters from the \$WORKSPACE/appParams file that contains multiple lines of key-value parameter pairs. You can then use these parameters to pass passwords or other sensitive information without displaying them in the Workload Manager UI.</p>
Choose the binaries to be Copied	<p>Identifies if the files must be copied to an external location</p> <ul style="list-style-type: none"> • Copy Binaries to External Location – all your binaries will be available under /tmp/app1/latest • Don't copy Binaries – the binaries will not be copied to any other location
Choose the Deployment Behavior	<p>You can only do this once for deployments.</p> <ul style="list-style-type: none"> • Create a new Deployment on every Build: This option creates a brand new deployment for every build. • Stop Old Deployment if exists: This option creates a brand new deployment for every build and stops the older deployment if it exists. • Update Existing Deployment (Create new if doesn't Exist): <ul style="list-style-type: none"> • Updates a previous deployment that was launched from the Workload Manager Jenkins plugin during a previous build for the same project. • If the previous deployment is still in progress (job) and is not yet in the Running state, the Workload Manager Jenkins plugin waits till the deployment is in the running state before triggering an update (see Deployment, VM, and Container States for additional context). • If the previous deployments ends up as an error or if that deployment is stopped or cancelled from the Workload Manager UI, the Workload Manager Jenkins plugin launches a new deployment as part of the Update process. • If this is the first build, the Workload Manager Jenkins plugin creates a new Deployment and from the next successful build it uses the existing deployment. • UpdateScripts: A comma separated list of tierName:actionId that are executed in the order mentioned here. For example: AppCluster:1, Database:2, AppCluster:3 where 1,2,3 are the Action ID's created with the update scripts and parameters like BUILD_ID. • See Actions Library for additional details on creating a custom deployment update action. • Wait For Deployment And Export Details: Waits for Deployment to enter the <i>Running</i> state and exports the Job Details to the \$WORKSPACE/userenv file.

In Update Scripts, \$BUILD_ID or %jenkinsBuildId% can be passed as an argument to point the *Binaries to be Copied* during an update deployment.



The %jenkinsBuildId% macro is not an applications-specific macro. This Workload Manager-defined macro applies to deployments that are launched using the Jenkins plugin.

The jenkinsBuildId macro is mainly used to pass the Jenkins Build ID to the userenv of the app deployment. Any deployment triggered by the Jenkins plugin will automatically have jenkinsBuildId in the userenv and will be used to point to the right binaries in repo/storage. For example, if a web server has a previous war file path set to /shared/app/petclinic/latest/, then this war file (petclinic.war) can now use this macro to point to /shared/app/petclinic/%jenkinsBuildId%/petclinic.war.

In update deployment scenarios, the jenkinsBuildId macro changes the value that should be passed to existing deployments as userenv has old the jenkinsBuildId value during the deployment.

The folder name that Workload Manager creates in the target location will now use the jenkinsBuildId value instead of the random timestamp value.

This option creates a Brand new deployment on every build.

When you update an existing deployment:

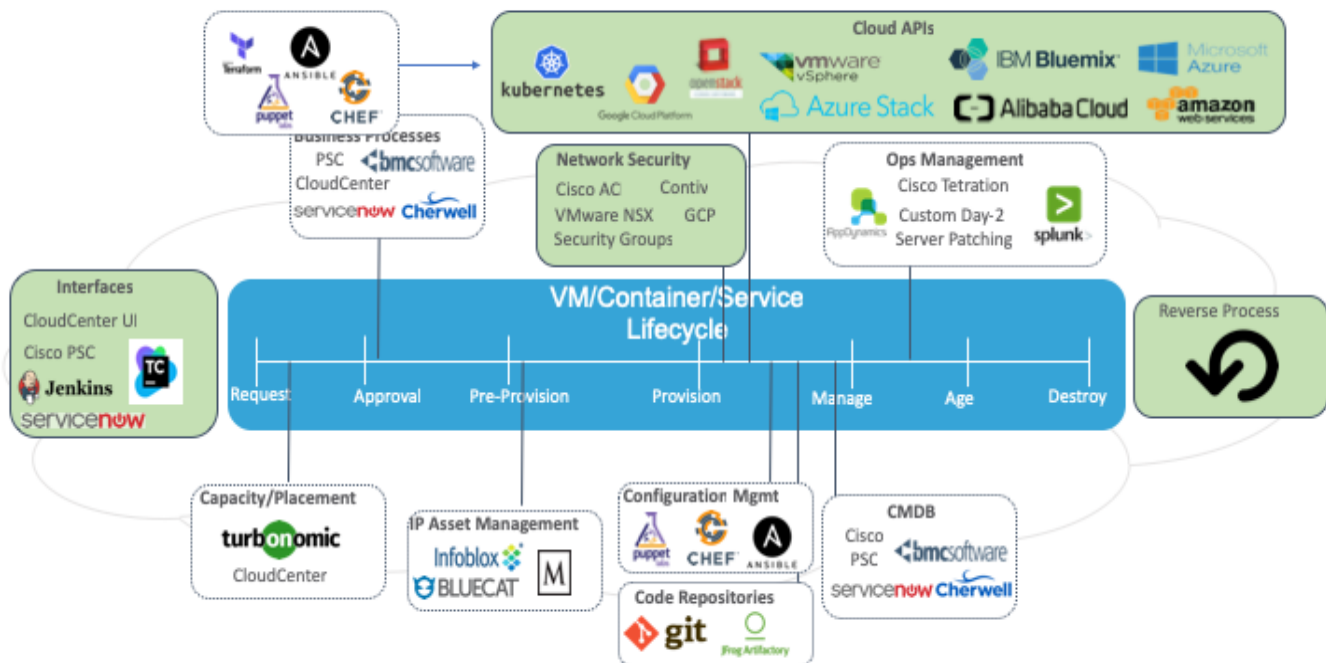
- It updates a previous deployment that is launched from the Jenkins plugin during the previous builds of the same project.
- If the previous deployment job is still in progress and is not in the Running state, the plugin waits till it enters the running state and then triggers an update.
- If the previous deployments ends in an error or if that deployment is stopped/cancelled from the Workload Manager UI, the plugin launches a fresh deployment as part of update.
- If it is the first build, this plugin creates a new deployment and for the next successful build it uses the existing deployment.
- If you receive a security group error, be sure to verify if more than one system is accessing the same account.

Service Libraries

Workload Manager Service Libraries

Cisco CloudCenter Suite is a multicloud management solution that allows enterprise IT teams to securely design, deploy, and optimize infrastructure and applications across multiple clouds from a single point of access. It delivers a consistent user experience, while controlling costs and helping you meet compliance requirements. The suite simplifies multicloud management by providing workflow automation, application lifecycle management, cost optimization, governance and policy management across various cloud types. The [CloudCenter Suite](#) consists for 3 keys modules: Workload Manager, Action Orchestrator, and Cost Optimizer.

Workload Manager allows users to integrate to various products as part of the orchestration process. This integration can occur at various points during the lifecycle of the application. We can introduce integrations during the Day 0, Day 1, or Day 2. These integrations or service libraries can include databases, PaaS services, networking solutions, CMDBs, and so forth. This section list sample service libraries Workload Manager has with various products and how they tie into the lifecycle of the application.



This is a continuously growing list. Click a category to continue.

- [Backup and Recovery](#)
- [Caching](#)
- [Compute](#)
- [Databases](#)
 - [DBaaS](#)
 - [NoSQL Databases](#)
 - [Relational Databases](#)
- [ITSM](#)
- [Logging](#)
- [Middleware](#)
- [Monitoring](#)
- [Networking](#)
 - [DNS](#)
 - [IPAM](#)
 - [Load Balancers](#)
- [Sample App Profiles](#)

Back to: [Workload Manager Integrations](#)

Troubleshoot

Troubleshoot Workload Manager

- [Obtaining Your SSH Private Key](#)
- [Bundle Download Failure](#)
- [Kubernetes Troubleshooting](#)
- [Agent Troubleshooting](#)
- [Troubleshooting Parameters](#)

Obtaining Your SSH Private Key

Obtaining Your SSH Private Key

- [Overview](#)
- [API to Use](#)
- [Process](#)
- [Using the PEM File](#)

When you launch a VM-based deployment in Workload Manager, Workload Manager generates a private key for you for each region where you deploy VMs. You will need this private key to establish a direct SSH session with any of your application VMs for the purpose of troubleshooting.

Workload Manager provides the [View Keys](#) API call which returns a list of private SSH keys associated with each cloud region where you have deployed VMs through Workload Manager. For the API call to successfully return your keys, you must first be logged in to your CloudCenter Suite account with the same browser you will issue the API call from.

To obtain your SSH private key, follow this process.

1. Login to your CloudCenter Suite account through a supported browser.
2. Using the same browser, open a new tab to enter the [View Keys](#) API call.
3. From the API response body, copy the private key associated with the cloud region containing the VM you need to access. Copy all text between the opening and closing double quotes.
4. Open a text editor, paste the private key, perform a global replace of all occurrences of "\n" with the new line character, and save the file.
5. Change the file extension to .pem
6. Set the appropriate file permission level:

```
chmod 400 <Key_file>.pem
```

You can now use the pem file in your ssh command:

```
ssh -i <Key_file>.pem <ssh_username>@<vm_ip_address>
```

Bundle Download Failure

Bundle Download Failure

- [Overview](#)
- [Diagnosis](#)
- [Workaround](#)

In some cases when you try to download a bundle after installing CloudCenter and importing the VM to CloudCenter, you may see the following error:

```
Failed to download specified bundle: BundleName
```

A Maximum Transfer Unit (MTU) value mismatch is indicative of this being an instance-related issue. In these cases, check the MTU settings on the instance:

```
ifconfig | grep MTU
```

This kind of an error is issued if the MTU settings on the instance are different from the default MTU value of **1500**.

- Change the MTU value back to the default of 1500.

Kubernetes Troubleshooting

Kubernetes Troubleshooting

- [Available Documentation](#)
- [Log Files](#)
- [Insufficient Permission](#)
- [Incorrect API Version](#)
- [Rolling Updates](#)
- [Deployment Migration](#)
- [Expired Certificates](#)
- [Container Clouds](#)
- [Configure a Kubernetes Cloud](#)
- [Container Service](#)

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	<p>Login to Kibana (Monitor Modules > View Logs in Kibana) and look for error messages in logs with the text "cloudcenter-cco".</p> <p>You may find the following warning message in the Kubernetes cloudcenter-cco logs – you can safely ignore this message as it does not impact product functionality.</p> <pre>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!</pre>
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".

Failure to Deploy a New Container

If you are unable to deploy a new container, revisit the following steps to ensure that you follow the prescribed process: See [Configure a Kubernetes Cloud](#) for additional details.

- Check your **clusterrole** assignment and ensure that it is set to **cluster-admin**: Role binding the service account to the admin is essential to access the dropdown in the Cloud Defaults page.

```
kubectl create clusterrolebinding <name> --clusterrole=cluster-admin
```

- If the details in the previous bullet did not address the issue, then create a dedicated service account for CloudCenter Suite. The following example, walks you through the required steps for this process

```
kubectl create serviceaccount cloudcenterSA

kubectl create clusterrolebinding cloudcentersabinding --clusterrole=cluster-admin --
serviceaccount=default:cloudcenterSA

#The following commands use jq. If not installed, you can install it using this command: sudo apt-get
install jq

kubectl get serviceaccount cloudcenterSA -o json | jq -Mr '.secrets[.name]'

#The cloudcenterSA-token-XXXXX name is unique and is gathered from this command -- be sure to replace
the token in the following command

kubectl get secrets cloudcenterSA-token-XXXXX -o json | jq -Mr '.data.token' | base64 -d
```

- If the above two workarounds did not address the issue, verify the Kubernetes setting for the **Default API version** or the **API version override**. The API version is optional and not required.
 - To verify the version, access the Kubernetes Region UI > Kubernetes Settings.
 - If you have configured a specific API version in your environment, try leaving it blank and retry the deployment
- **Issue:** Your deployment fails with *forbidden* (networkpolicies.extensions is forbidden) or *Code 403* (Received status: Status(apiVersion=v1, code=403) in the containerblade.log
- **Reason:** The Service Account is associated with cluster role has insufficient permissions or has a non-existing cluster role
- **Solution:** Update the Service Account to map to right cluster role. See [Configure a Kubernetes Cloud](#) for additional details.
- **Issue:** Your deployment fails with a *Code 400* (no kind Network Policy is registered for v1 version. Received status:Status(apiVersion=v1, code=400) error in the containerblade.log
- **Reason:** The API version for the object (Network Policy) in this case is not sent correctly. Either a wrong version is specified or the CloudCenter platform could not auto-detect the version.
- **Solution:** Try one of the following solutions:
 - Leave the Default and Override API versions blank – this often corrects the issue.
 - Alternately, find the right version by examining an existing object instance in the Kubernetes dashboard or using the kubectl GET API. In the CloudCenter Kubernetes region settings, set the **API Version Override** field with the identified version. For example, "NetworkPolicy: v1beta1".
- **Issue:** *Unable to update container images on Deployment Details page error*
- **Reason:** The rolling update action is available in Workload Manager 5.1.0 from the Deployments list page, not the deployment details or job details page.
- **Solution:** Go to the Deployments list page and select **Update** from the **Actions** dropdown menu.
- **Issue:** Workload Manager POST APIs fail when the request data size is more than 5MB of data.
- **Reason:** The Suite Admin validation restricts the size of each payload to be less than 5MB. You can change to a bigger payload size using the solution below.
- **Solution:** To change the payload size, follow this process:
 1. Edit the annotations using the following command:

```
kubectl edit ingress common-framework-suite-ui -n cisco and change the proxy-body-size to "0"
```

2. Run the following command for each ingress gateway:

```
kubectl edit ingress common-framework-suite-gateway -n cisco
```

3. Edit the configmap for the ingress controller and add these details to the data map – similar to the following example:

```
kubectl get cm common-framework-nginx-ingress-controller -n cisco -o yaml
apiVersion: v1
data:
  client-body-buffer-size: 32k # CHANGE the client body buffer size in this line
  enable-vts-status: "false"
  proxy-body-size: "0" # CHANGE the proxy body size in this line
  proxy-set-headers: cisco/common-framework-nginx-ingress-custom-headers
  server-tokens: "false"
kind: ConfigMap
metadata:
  creationTimestamp: "2020-06-24T18:05:15Z"
  labels:
    app: nginx-ingress
    chart: nginx-ingress-0.31.0
    component: controller
    heritage: Tiller
    release: common-framework
  name: common-framework-nginx-ingress-controller
  namespace: cisco
  resourceVersion: "378066"
  selfLink: /api/v1/namespaces/cisco/configmaps/common-framework-nginx-ingress-controller
  uid: 5b96d0a6-7680-4526-932a-e2b8f9c7f8aa
```

See [Troubleshooting > Expired Certificates](#) at the CloudCenter Suite level for details on addressing this issue.

Agent Troubleshooting

Agent Troubleshooting

- [Node Not Reachable](#)
 - [Symptom](#)
 - [Applies to](#)
 - [Solution](#)
- [Agent and RabbitMQ communication issue](#)
 - [Symptom](#)
 - [Applies to](#)
 - [Solution](#)

Symptom

While the Deployment succeeded, after sometime the Agent was not in the Running state and the deployment shows the VM in red (Unreachable).

Applies to

All Clouds and Apps

Solution

1. Verify the following files for additional details:
 - \$AGENT_HOME/log/agent.log
 - \$AGENT_HOME/log/nohup.err (Linux)
2. Reboot the VM or manually start the Agent.

Symptom

A deployment failed due to Bootstrap Timeout.

Applies to

All clouds and applications.

Solution

To troubleshoot issues caused by a communication between the agent and RabbitMQ, do the following.

1. Check if the broker host IP and broker host port were correctly specified in user-data in the worker VM.
2. Confirm that the RabbitMQ service is running:
 - a. If you deployed Cloud Remote appliance for this region, from the command line of the Cloud Remote appliance enter:


```
docker ps | grep rabbitmq
```
 - b. If you did not deployed Cloud Remote appliance for this region, ensure kubectl on your computer is configured to communicate with the Kubernetes cluster hosting the CloudCenter Suite and issue this command:


```
kubectl get pod | grep rabbitmq
```
3. Ensure there are no firewall rules on the Worker VMs blocking outgoing RabbitMQ connections.
4. Ensure there are no firewall rules blocking the RabbitMQ service, whether it is running in the CloudCenter Suite cluster or on the Cloud Remote appliance.
5. Optional. If required, open the RabbitMQ console and confirm you can login by using the following format:

```
http://<RabbitMQ server IP>:15672

#Effective Workload Manager 5.1.1, Port 15672 is disabled by default and can be enabled when required.
```

If you are not using Cloud Remote, you can find the RabbitMQ server IP from the **Worker AMQP IP Address** field in the Region Connectivity section of the Regions / Details tab. If you are using Cloud Remote, use the public IP address assigned to your Cloud Remote appliance when you launched it.

6. Verify from the agent logs that a Node Bootstrap message was sent.
7. [Login to Kibana](#) and verify from the logs that:
 - a. A Node Bootstrap message was received by the cloudcenter-cco service
 - b. The cloudcenter-ccm-backend service received messages from the cloudcenter-cco service

Troubleshooting Parameters

Troubleshooting Parameters

- [Overview](#)
- [Parameter to Ignore App Failure](#)
- [Timeout Parameters](#)

Workload Manager allows you to define your own parameters (see [Using Parameters](#)) or use [Workload Manager-Defined Parameters](#).

For development or test debugging purposes, Workload Manager provides a troubleshooting parameter so users can call the required parameter on an as-needed basis.

The **cliqrIgnoreAppFailure** parameter does not terminate nodes during a failure and continues to keep all failed application deployment nodes running.

Currently the app profile along with the security groups are deleted.



Caution

Use the **cliqrIgnoreAppFailure** global parameter to *only* troubleshoot deployment failures. This parameter is only intended for initial debugging and must be removed for security reasons once the application is working properly.

To configure all failed nodes to run for any application, follow this procedure.

1. Model an Application.
2. Edit the application (see [Application Tasks > Edit/Update](#)).
3. Configure the cliqrIgnoreAppFailure global parameter
 - a. Access the [Topology Modeler > Global Parameters](#).

- b. Click the **add a parameter** link to open the Parameter dropdown section.

Model a N-Tier Execution Application Profile

Basic Information
Global Parameters
Topology Modeler

If you would like to add additional parameters, this is the place.

Global Parameters | [add a parameter »](#)

▼ Parameter

Parameter Name 19

Display Name 22

Help Text 206

Type string Max Length

Default Value 251

User Options:

- Should this parameter be visible to the user?
- Should this parameter be editable by the user?
- Should this parameter be optional?

Save as Template
Save as App
Cancel

- c. Add the following values for each field:

- i. Parameter Name: **cliqIgnoreAppFailure**
- ii. Display Name: **Ignore App Failure**
- iii. Help Text: **The app will not terminate nodes during a failure**
- iv. Type: **string**
- v. Default Value: **true**
- vi. Check the User Options as applicable to your deployment (see [Using Parameters](#))

> Granular Control for User-Defined Parameters for additional context).

4. Save the profile as an App or a Template as applicable to your deployment.
5. [Deploy the Application](#). You see cliqIgnoreAppFailure global parameter configuration displayed Ignore App Failure as true in the Global Parameters section.

Global Parameters

Specify values for parameters that the application needs in order to run.

Ignore App Failure:

true

The app will not terminate nodes during a failure

You have now configured the **cliqrIgnoreAppFailure** global parameter for this application and all failed nodes will not be terminated.

You can use multiple parameters to configure timeouts:

Parameter	When Used?	Additional Details
cliqrContainerExecuteScriptTimeout	It is possible for an external service to fall into an infinite loop if using the stop/start external initialization scripts. This situation may cause the Docker container to run forever. In these cases, use this parameter when modelling applications.	External Service > Script Timer
cliqrNodeBootstrapTimeout	Identifies the maximum time available for VMs to bootstrap after they are launched.	Deployment and VM States
cliqrNodeReadyTimeout	Terminates timed out VMs.	> Orchestration Lifecycle Threshold Settings

Workload Manager API

Workload Manager API

- [API Overview](#)
- [API Authentication](#)
- [API Key](#)
- [Base URI Format](#)
- [HTTP Status Codes](#)
- [CSRF Token Protection](#)
- [API Permissions](#)
- [Synchronous and Asynchronous Calls](#)
- [CCM Calls 5.4.0](#)
- [Cloud Setup Calls 5.4.0](#)
-

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 Suite APIs?](#)

The payloads for the CloudCenter Suite APIs are visible in the API documentation section for each module.

CloudCenter Suite APIs provide support for the CloudCenter 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.

The CloudCenter Suite API date and time values are formatted in [Unix time](#) to the millisecond level. The APIs are agnostic to dates and time zones.

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 Suite task or creating the resource
- **DELETE**: To remove specific aspects of the CloudCenter Suite deployment

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
```

For each API request, you see two common attributes displayed in the API response:

```
{
  "resource": "https://<HOST>:<PORT>/v1/users/",
  "size": 12,
  "pageNumber": 0,
  "totalElements": 12,
  "totalPages": 1,
  "users": [
    {
      "id": "2",
    }
  ]
  ...
}
```

- 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*.

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:


```
{
  "resource": "https://<HOST>:<PORT>/v1/users/",
  "size": 12,
  "pageNumber": 0,
  "totalElements": 12,
  "totalPages": 1,
  "users": [
    {
      "id": "2",
    }
  ]
  ...
}
```

- **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 default size (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). <div style="border: 1px solid green; border-radius: 10px; padding: 5px; display: inline-block;">  if the page does not have more than 10 elements, then only those 10 elements are displayed. </div>
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 = 20 records
 - 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 <ul style="list-style-type: none"> • Description: The page number that the client wants to fetch. Page numbers start with 0 (default). • Type: Integer
totalElements <ul style="list-style-type: none"> • Description: The number resources that an API call returns • Type: Long
totalPages <ul style="list-style-type: none"> • 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 <ul style="list-style-type: none"> • Description: Unique URL to access this resource. • Type: String
size (see above)
pageNumber <ul style="list-style-type: none"> • Description: The page number that the client wants to fetch. Page numbers start with 0 (default). • Type: Integer
totalPages <ul style="list-style-type: none"> • Description: The number of pages in a response • Type: Integer
jobs <ul style="list-style-type: none"> • Description: Array of JSON objects that use jobs as the key. • Type: Array of JSON objects
previousPage <ul style="list-style-type: none"> • Description: A resource link to the previous page. • Type: URI as a string
nextPage <ul style="list-style-type: none"> • Description: A resource link to the following page. • Type: URI as a string
lastPage <ul style="list-style-type: none"> • Description: A resource link to the last page. • Type: URI as a string

- **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 YXNpdjpwcm9kei5kb3DD6F7874E6B26B
> 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	<ul style="list-style-type: none"> • Description: Unique, system-generated identifier for this <i>resource</i>. • Type: String 										
status	<ul style="list-style-type: none"> • Description: Status of the operation. See the <i>APIs for the relevant module</i> to view a list of all job operations. • Type: Enumeration <table border="1"> <thead> <tr> <th>Enumeration</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>SUBMITTED</td> <td>The operation has been submitted</td> </tr> <tr> <td>RUNNING</td> <td>The operation is currently in progress</td> </tr> <tr> <td>SUCCESS</td> <td>The operation succeeded</td> </tr> <tr> <td>FAIL</td> <td>The operation failed</td> </tr> </tbody> </table>	Enumeration	Description	SUBMITTED	The operation has been submitted	RUNNING	The operation is currently in progress	SUCCESS	The operation succeeded	FAIL	The operation failed
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	<ul style="list-style-type: none"> • Description: Start/End time for this resource. Unix epoch time in milliseconds. • Type: <ul style="list-style-type: none"> • v1 APIs = Long • v2 APIs = Epoch time as a String 										
totalCost	<ul style="list-style-type: none"> • Description: Identifies the total cost per hour of the job for billing purposes. See the Cost Optimizer API section to view additional details. • Type: Float 										
nodeHours	<ul style="list-style-type: none"> • Description: The number of VM hours for this resource. See the Cost Optimizer API section to view additional details. • Type: Float 										
name	<ul style="list-style-type: none"> • Description: The name assigned for this <i>CloudCenter Suite Resource</i>. Valid characters are letters, numbers, underscores, and spaces. • Type: String 										
deploymentEntity.name	<ul style="list-style-type: none"> • Description: Identifies evolving resource details about the deployment. The deploymentEntity attribute uses the <i>deploymentEntity.name</i> format, where .name is a search value for deploymentEntity and deploymentEntity itself is a JSON object. <div style="border: 1px solid green; padding: 10px; margin: 10px 0;">  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. </div> <ul style="list-style-type: none"> • 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. See the *Favorite Deployments* section for the relevant release for additional context.
- **Type:** Epoch time as a String

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 <i>pattern</i> %
fl	LIKE % <i>pattern</i>
eln	NOT LIKE <i>pattern</i> %
fln	NOT LIKE % <i>pattern</i>
file	LIKE % <i>pattern</i> %"
gt	> <i>searchValue</i>
lt	< <i>searchValue</i>
ge	>= <i>searchValue</i>
le	<= <i>searchValue</i>
gtlt	> <i>searchValue</i> && <i>searchValue</i>
gtelt	>= <i>searchValue</i> && < <i>searchValue</i>
gtlte	> <i>searchValue</i> && <= <i>searchValue</i>
gtelte	>= <i>searchValue</i> && <= <i>searchValue</i>
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, file, nn, emp, noemp
nodeHours	eq, gt, ge, le, gtlt, gtlte, gtelte, gtelt
name	eq, nn, eln, file, fln, el, emp, noemp, fl
description	eq, nn, eln, file, fln, el, emp, noemp, fl

deploymentEntity.name	eq, nn, eln, fle, fln, el, emp, noemp, fl
ownerEmailAddress	eq
cloudFamily	eq, nu
status	eq, nu

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 Y2xpcXJhZG1pbjpmEM0RENky3ODc0RTZCMjZC
> 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
```

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 required Permission Control level to access any *resource*, you receive the HTTP 403 status error mentioned in the [HTTP Status Codes](#) section.

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](#)

CloudCenter Suite APIs require the following authentication details for each API call:

- Username
- API access key



The authentication HTTP header is not required when making standalone REST API calls using the username/API Key credentials.

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 the username
- **BED74F4D9BFE0DA0** is the API accessKey

Your tenant administrator can retrieve the username and API access key from the UI. See [API Key](#) for additional details.

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.

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

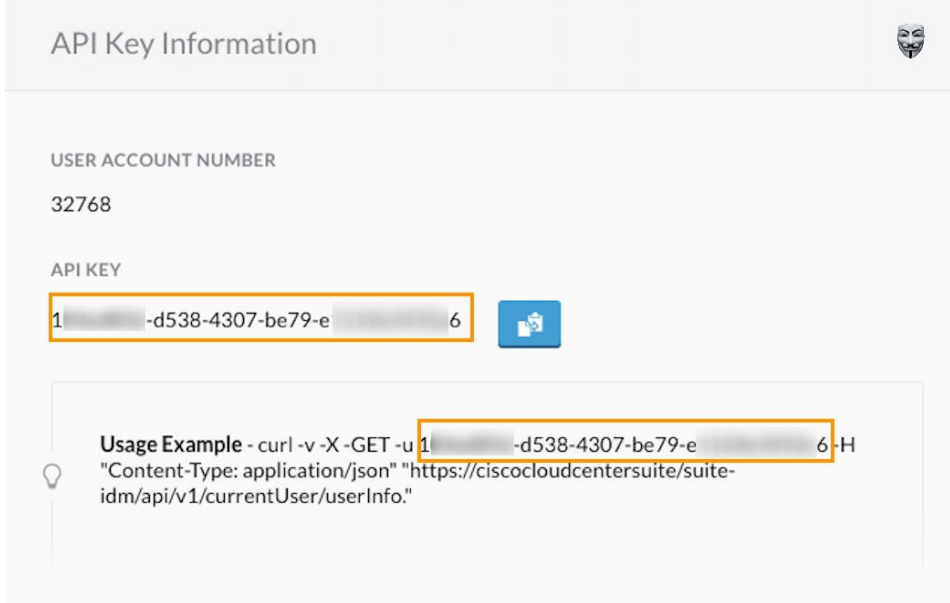
Generate API Key

- [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](#)

You need an **API key** to use CloudCenter Suite APIs. Suite administrators or tenant administrator (for their respective tenants) can generate/regenerate an API key by using the Suite Admin UI or the `user_api_key` API call.

To generate the API key from the UI for yourself, follow this procedure:

1. Navigate to the [Suite Admin Dashboard](#) and click your account 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.



API Key Information

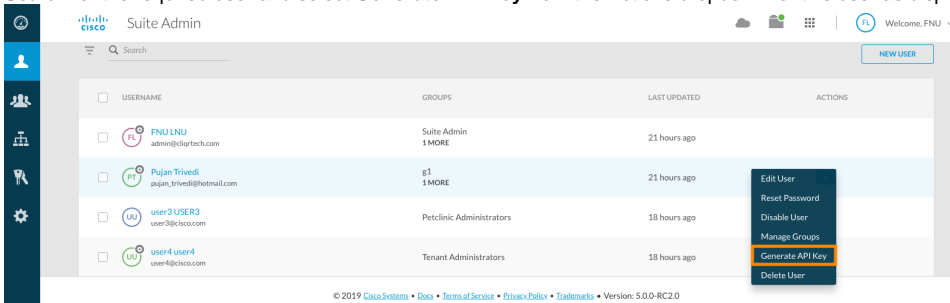
USER ACCOUNT NUMBER
32768

API KEY
1[redacted]-d538-4307-be79-e6

Usage Example - curl -v -X GET -u 1[redacted]-d538-4307-be79-e6 -H "Content-Type: application/json" "https://ciscloudcentersuite/suite-idm/api/v1/currentUser/userInfo."

To generate the API key from the UI for another user, follow this procedure:

1. Navigate to the [Suite 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.



USERNAME	GROUPS	LAST UPDATED	ACTIONS
FNU LNU admin@clirttech.com	Suite Admin 1 MORE	21 hours ago	Edit User Reset Password Disable User Manage Groups Generate API Key Delete User
Pujan Trivedi pujan_trivedi@hotmail.com	g1 1 MORE	21 hours ago	
user3 USER3 user3@cisico.com	Petclinic Administrators	18 hours ago	
user4 user4 user4@cisico.com	Tenant Administrators	18 hours ago	

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.

To generate the API key using the Suite Admin API call, follow this procedure:

1. Issue the [Password Service API Calls](#) > `/api/v1/users/{userId}/user_api_key` API POST call 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](#)

The base URI format is **https:// <host>:<port>/...**

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.

The port is generally represented at <PORT> in all CloudCenter APIs. It represents the port used to connect to the CCO 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.

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 used to make the API call are displayed after the APIs and are called out after the description.

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 where each 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

HTTP Status Codes

CloudCenter APIs return one or more of the following HTTPS status codes for all (synchronous and asynchronous) API requests:

HTTP Response Code	Status	Description
200	Success	Successful GET and PUT
201		Successful POST (when a resource is created)
202		Request accepted for a time-consuming task (asynchronous update and created requests). See Synchronous and Asynchronous APIs for more details You can issue GET calls until the request completes.
204		Successful DELETE
30x	Redirection	Only displays if a client calls an API using HTTP instead of HTTPS
400	Client failure	Validation error. This category has additional error codes in the response body for each API (as applicable).
401		Not authenticated
403		Forbidden. You do not have the required permission level to access the <i>CloudCenter Resource</i>
404		Resource not found
500	Server failure	Server error: The server failed to respond to this request due to an internal error

Back to:

- [Suite Admin API](#)
- [Workload Manager API](#)
- [Action Orchestrator API](#)
- [Cost Optimizer API](#)

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](#)

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

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.

To set a CSRF token, add **X-CSRF-TOKEN** to the header name (case sensitive, all uppercase).

To obtain the CSRF token, follow this procedure.

1. You must first pass authentication. See [API 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)

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](#)

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 in [Understand Roles](#) and user-level as explained in [Understand User Levels](#).

Users can find their permission level by executing the **GET /suite-idm/api/v1/currentUser/userInfo** API listed in the [IDM Service API Calls > User Controller](#) section.

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: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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	A SUITE_ADMIN can promote any SUITE_USER to the Suite Administrator group as described in Create and Assign Groups . This user can perform the following tasks: <ul style="list-style-type: none"> • Manage users and groups • Create, update, delete users and groups • Assign roles to users and groups • Manage passwords for users
SUITE_OUTOFBOX_USER	A SUITE_ADMIN can 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 with SUITE_ADMIN permissions and/or SUITE_TENANT_ADMIN for this tenant as described in Create and Manage Users > User Actions . This user can perform the following tasks: <ul style="list-style-type: none"> • Edit any user's profile by changing the first/middle/last name and email • Configure metadata details • Configure groups • Reset password • Disable a user

See [OOB Groups, Roles, and Permissions](#) for details.

See [Action Orchestrator Roles](#) for details.

See [Access 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

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 **Status 204 No Content** after deleting the *resource* in the background.

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 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 return HTTP Status Codes responses containing information with an HTTP 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 the HTTP 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 the resource URL for this operation.

Call States

In the following example of a [Create Cloud Account](#) API:

- 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](#)

CCM Calls 5.4.0

CCM Calls for Workload Manager 5.4.0

Refer to the 5.4_CCM_api.json file.

Cloud Setup Calls 5.4.0

Cloud Setup Calls for Workload Manager 5.4.0

Refer to the 5.4_cloud_setup_api.json file.

Manual API Documentation

Manual API Documentation

- [Swagger File Context](#)
- [Manual Documentation Availability](#)
- [List of Manually Documented CloudCenter Suite 5.0 APIs](#)

The CCM Calls section for each release provides an automated Swagger file listing the available API calls for Workload Manager. While these files do not contain the history and details that are available in the older manually-created API documentation, this section provides a list of previously-supported API calls. Going forward, the API documentation for each release will be generated automatically to ensure consistency with the rest of the CloudCenter Suite documentation.

This manual section continues to be available and provide information that was previously available in Workload Manager 5.0 and earlier releases.

- The same information is retained here for easy reference in addition to the Swagger files.
 - The [Release Notes](#) for each release will continue to provide the historical API information going forward.
 - This manual section is only intended for additional reference and *will not be* updated. Once the swagger files are in place, this section will be deprecated.
-
- [ACL Management API Calls](#)
 - [Actions Management API Calls](#)
 - [Application Management API Calls](#)
 - [Bundle Management API Calls](#)
 - [Cloud Account Management API Calls](#)
 - [Cloud Family Management API Calls](#)
 - [Cloud Instance Management API Calls](#)
 - [Cloud Image Mapping Management API Calls](#)
 - [Cloud Management API Calls](#)
 - [Cloud Properties Management API Calls](#)
 - [Cloud Provider Management API Calls](#)
 - [Cloud Region Management API Calls](#)
 - [Cloud Storage Management API Calls](#)
 - [Custom Action Management API Calls](#)
 - [Deployment Environment Management API Calls](#)
 - [Export/Import Application API Calls](#)
 - [Extensions Management API Calls](#)
 - [File Management API Calls](#)
 - [Image Management API Calls](#)
 - [Inventory Management API Calls](#)
 - [Job Management API Calls](#)
 - [Key Management API Calls](#)
 - [Operations Status API Calls](#)
 - [Ownership Management API Calls](#)
 - [Phase Management API Calls](#)
 - [Plan Management API Calls](#)
 - [Project Management API Calls](#)
 - [Policy Management API Calls](#)
 - [Report Management APIs](#)
 - [Service Management API Calls](#)
 - [VM Management API Calls](#)