



## CloudCenter Suite 5.1 Documentation

**First Published:** August 19, 2019

**Last Modified:** November 9, 2019

### **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. CloudCenter Suite 5.1	2
1.1 SaaS Access	3
1.2 Self-Hosted Access	6
1.3 Release Notes	8
1.3.1 CloudCenter Suite 5.1 Release Notes	9
1.4 Browser Compatibility	11
1.5 Support Information	12
1.5.1 Documentation Website	13
1.5.2 Documentation Accessibility	14
1.5.3 OpenSource Version Matrix	15
1.5.4 End of Support Notices	16
1.6 Security Considerations	18
1.7 Module Versions	20

# CloudCenter Suite 5.1

Welcome to CloudCenter Suite 5.1 Documentation



Suite Admin



Workload Manager



Action Orchestrator



Cost Optimizer



Training

# SaaS Access

## CloudCenter Suite SaaS Overview



On March 1st, 2021 the CloudCenter SaaS platform will be completely decommissioned. For questions or comments, please contact CloudCenter Suite Product Management.

CloudCenter Suite SaaS is a managed, cloud-based service offered by Cisco that enables you to administer and control applications, costs and workflows across multiple clouds from anywhere! Cisco provides ongoing management, maintenance and upgrades with end-to-end monitoring and our world-class 24/7 customer support. **Cisco is responsible for managing the availability, stability, and security, of both the platform and the lifecycle management of the modules and their respective toolings.**



### North America



Be sure to use the regional URL that corresponds with your original CloudCenter Suite SaaS order. CloudCenter Suite SaaS trial accounts are automatically provisioned in the North America region.



### Status Page

- [CloudCenter Suite Training](#)
- [Getting Started with CloudCenter Suite SaaS](#)
- [CloudCenter Suite SaaS General FAQ](#)
  - [What is the CloudCenter Suite SaaS Global Infrastructure?](#)
  - [How do I purchase CloudCenter Suite SaaS?](#)
  - [How do I manage private clouds with CloudCenter Suite?](#)
- [CloudCenter Suite SaaS Trial Accounts FAQ](#)
  - [What is the CloudCenter Suite SaaS 30-day trial?](#)
  - [Does the CloudCenter Suite SaaS 30-day trial include any sample content or training content to help me get started?](#)
  - [Once my CloudCenter Suite SaaS 30-day trial is complete, what are the next steps available?](#)
  - [How can I get technical assistance?](#)

## CloudCenter Suite SaaS General FAQ

What is the CloudCenter Suite SaaS Global Infrastructure?

The CloudCenter Suite SaaS platform is a highly-available, scalable service designed to meet your needs for performance and data residency. CloudCenter Suite SaaS is available in North America (US-East and US-West).



The CloudCenter Suite SaaS platform is available to customers in **ALL** geographic locations, regardless of physical presence. Expansion of CloudCenter Suite SaaS platform to additional geographic regions is a roadmap item.

High Availability

The CloudCenter Suite SaaS platform employs multiple layers of redundancy to ensure that the environment is available 24x7. Our ability to fail over locally in seconds means you are unlikely to ever notice any downtime.

#### Disaster Recovery

Our built-in processes and workflows back up data for fast recovery times in the unlikely event of a local outage. We maintain comprehensive Disaster Recovery sites worldwide in North America.

#### Support

Our experienced customer support engineering team is available 24x7 to deliver superior customer service across any geographic region, in respective time zones, following the sun.

### How do I purchase CloudCenter Suite SaaS?

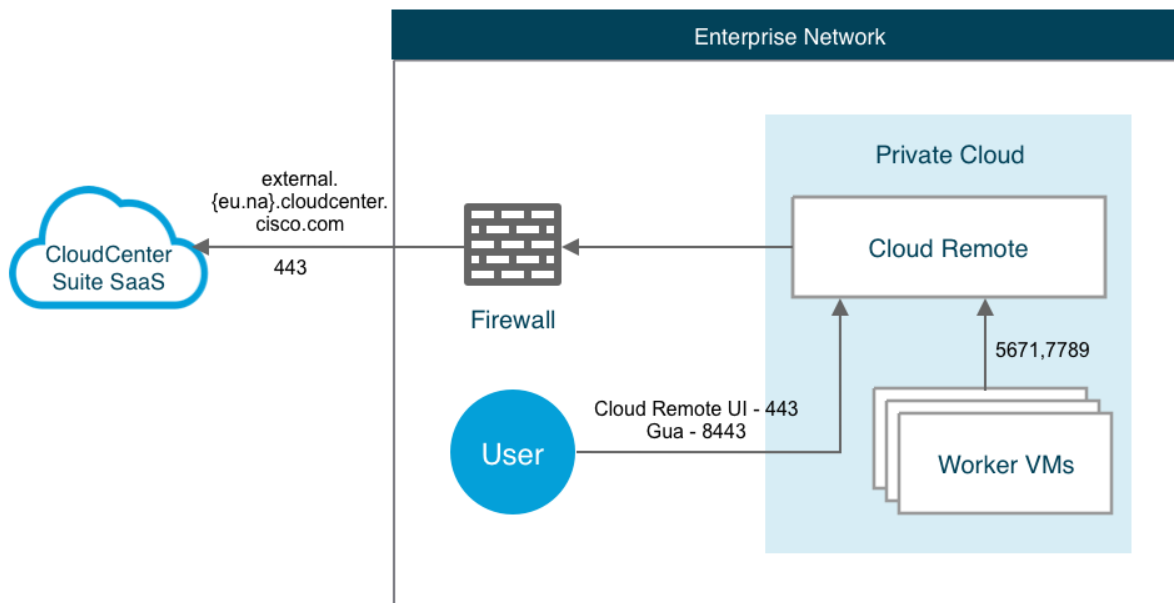
CloudCenter Suite SaaS can be purchased through your Cisco Account Manager, or by [contacting](#) a Cisco sales representative.

[Back to FAQs](#)

### How do I manage private clouds with CloudCenter Suite?

As displayed in the following image, CloudCenter Suite SaaS can manage private clouds and public clouds with established/limited connectivity using a Cloud Remote. In order to facilitate communication, each Cloud Remote only requires an outbound network connection from the cloud environment. See [Cloud Remote \(Conditional\)](#) for additional details.

#### CCS SaaS Managing Private Cloud



## CloudCenter Suite SaaS Trial Accounts FAQ

### What is the CloudCenter Suite SaaS 30-day trial?

The CloudCenter Suite SaaS 30-day trial is a "no-risk" opportunity from Cisco to explore and validate our integrated set of products for multicloud application and workflow management. Come learn how CloudCenter Suite accelerates innovation and simplifies governance and policy management across multiple clouds ... but with zero installation and maintenance!

Does the CloudCenter Suite SaaS 30-day trial include access to the entire product suite?

**YES!** Each CloudCenter Suite SaaS 30-day trial account is automatically setup with access to Suite Admin, Workload Manager, Action Orchestrator and Cost Optimizer - meaning no wasted time getting started with CloudCenter Suite!

Does the CloudCenter Suite SaaS 30-day trial include access to public/private cloud providers?

**YES!** Each CloudCenter Suite SaaS 30-day trial account is automatically setup with access to cloud regions from several popular public cloud providers - Amazon Web Services, Google Cloud Platform and Microsoft Azure. CloudCenter Suite SaaS users are expected to provide their own cloud accounts for connecting to each cloud environment. We currently **do not** support the ability to add any additional public or private clouds to CloudCenter Suite SaaS trial accounts. If you require different/additional regions, please review the section covering technical assistance.

Does the CloudCenter Suite SaaS 30-day trial include any additional Cisco or 3rd-party ecosystem integrations?

**YES!** Assuming they are publicly accessible, each CloudCenter Suite SaaS 30-day trial account can integrate Action Orchestrator with any number of Cisco or 3rd-party products. However, because CloudCenter Suite SaaS trial accounts do not support the ability to add private clouds, you will not be able to validate ACI integration.

[Back to FAQs](#)

Does the CloudCenter Suite SaaS 30-day trial include any sample content or training content to help me get started?

Our goal is to enable developers to leverage numerous integrations across many Cisco products and other ecosystem solutions to build on the strength of Cisco's ever-increasing investments in cloud technologies.

Sample Content

Each CloudCenter Suite SaaS 30-day trial account is created with a multitude of "out-of-the-box" content. We have imported several example application profiles into Workload Manager, ranging from single-OS virtual machines to multi-tier applications. Action Orchestrator has been setup with access to several Cisco-supported Git Repositories containing example workflows and atomic actions.

Training Materials

Once your CloudCenter Suite SaaS account is created, you should consider visiting our [CloudCenter Suite Training Portal](#).

The goal of this content is to help you setup and operationalize Workload Manager, Action Orchestrator, Cost Optimizer and Suite Admin. There are several types of training content offered, ranging from self-paced learning labs to instructor-led boot camps. In addition, Cisco Customer Experience (CX) / Services have begun to develop and introduce Accelerators and QuickStart Programs to enable our customers and their enterprises.

**Pro-Tip:** Get started with ... well, [Getting Started with CloudCenter Suite SaaS!](#)

[Back to FAQs](#)

Once my CloudCenter Suite SaaS 30-day trial is complete, what are the next steps available?

Each CloudCenter Suite SaaS 30-day trial account will be presented with a contact form near the end of the trial period. If you are ready to continue your multicloud journey with a POC or in-depth product demo, please complete the contact form to reach a Cisco Representative. Our sales team will engage with you as soon as possible! In the meantime, don't forget to keep your Cisco Account Team aware of your progress.

Can my CloudCenter Suite SaaS 30-day trial account be extended?

We are unable to provide extensions to CloudCenter Suite SaaS trial accounts.

Can my content be saved and/or backup?

**YES!** CloudCenter Suite SaaS trial accounts are welcome to [export application profiles](#) and [export workflows](#) from Workload Manager and Action Orchestrator.

[Back to FAQs](#)

How can I get technical assistance?

We are unable to provide one-on-one technical assistance to our CloudCenter Suite SaaS trial accounts at this time. However, we hope our rich, growing library of content can help you navigate and succeed with CloudCenter Suite SaaS. Please visit the [CloudCenter Suite Documentation](#) site and our [Training Portal](#) to get started.

Why was my CloudCenter Suite SaaS 30-day trial request rejected?

In most cases, CloudCenter Suite SaaS trial accounts are rejected for one of two reasons: the email address has previously been used to create a trial account; or the email address belongs to a "non-business" domain. We apologize for any inconvenience this may cause.

How can I add/remove cloud regions to my trial account?

CloudCenter Suite SaaS trial accounts are limited to a subset of popular cloud regions across AWS, Azure and GCP. If you would like additional regions be added to your trial account, please email [support@cloudcenter.zendesk.com](mailto:support@cloudcenter.zendesk.com) with your Trial ID and your requested regions. The same process can be used to request the removal of regions. Our team will review your request and respond as quickly as possible!

[Back to FAQs](#)

**Back to: [CloudCenter Suite Home](#)**

# Self-Hosted Access

## Suite Architecture

- [Overview](#)
- [The Suite Architecture](#)
- [Port Requirements](#)
- [The Suite Admin](#)
- [The Modules](#)

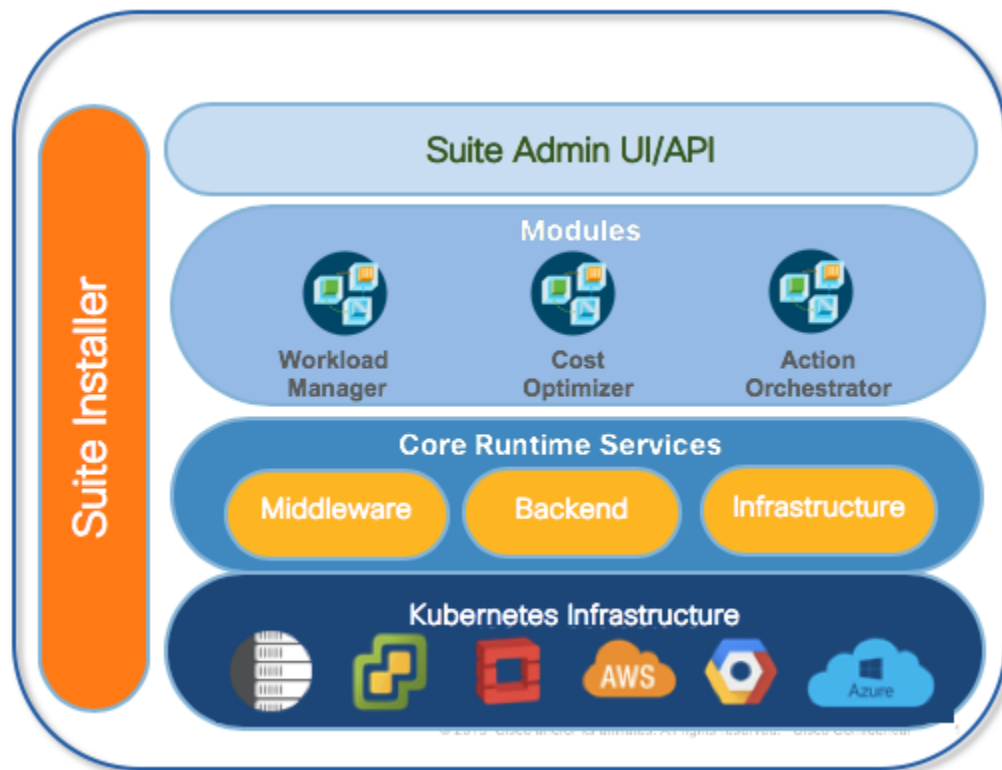
The CloudCenter Suite is Cisco's hybrid cloud deployment platform. This platform takes a unique approach to install, configure, and maintain hybrid cloud environments that are often encountered by Information Technology (IT) departments to adopt business agility and improve time-to-market solutions within an enterprise. As a cloud-based organization, your enterprise can choose from multiple cloud (*multicloud*) providers depending on your location, policies, permissions, security requirements, and governance regulations for both traditional and modern IT requirements.

The CloudCenter Suite provides a solution that is cloud agnostic, works with diverse workloads, provides cross-domain orchestration, supports cost-optimization, and integrates easily in an agile world.

The CloudCenter Suite is made up of the following components:

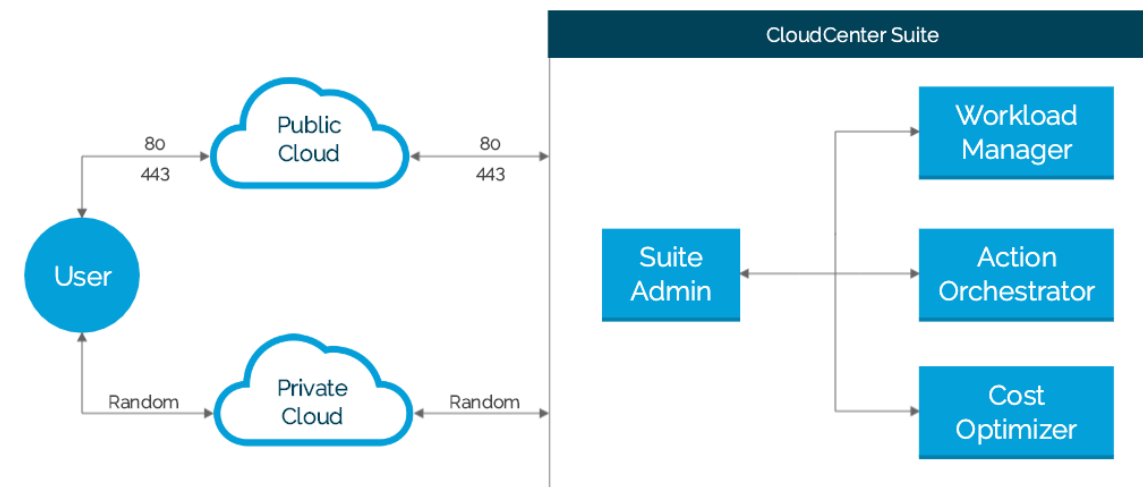
- **Suite Installer** – Installs the Suite Admin. See [Suite Installer](#) for additional details.
- **Suite Admin** – Installs and launches a suite of modules. See *The Suite Admin* section below for additional details.
- **Modules** – The Workload Manager, the Cost Optimizer, and the Action Orchestrator. See *The Modules* section below for additional details.
- **Core Runtime Platform and Kubernetes Infrastructure** – A Kubernetes-based platform that allows you to launch each module on a new or existing Kubernetes cluster.

The following image displays the Suite Admin architecture.



The following image identifies the ports that must be open for the CloudCenter Suite to function as designed.





When you download and install the Suite Installer, the **Suite Admin** is **already installed**! You have the option to use the Suite Admin UI to perform the following tasks:

- Install additional, available modules based on the list available in the Dashboard.
- Upgrade the Suite Admin or other installed modules when a new version becomes available.

The Suite Admin facilitates the installation of the following modules:

- **Workload Manager:**
  - This module allows IT organizations to provide management for clouds (public/private/container), applications, VMs/pods, governance policies with centralized visibility and permission control for enterprise environments.
  - See [Workload Manager](#) for additional details.
- **Action Orchestrator:**
  - This module allows IT organizations to use cross-domain orchestration to automate a process that has multiple, complex steps with a specific order and implemented across different technical domains.
  - See [Action Orchestrator](#) for additional details.
- **Cost Optimizer:**
  - This module allows IT organizations to use cost optimization in a pay-per-use environment to avoid consumption that does not add value.
  - See [Cost Optimizer](#) for additional details.

Each module in the CloudCenter Suite is independent and allows access to additional gateways or endpoints so you can add on module-specific components on supported clouds.

# Release Notes

## Release Notes for the CloudCenter Suite

- [CloudCenter Suite 5.1 Release Notes](#)

# CloudCenter Suite 5.1 Release Notes

## CloudCenter Suite 5.1 Release Notes

- [Release Date](#)
- [SaaS Release Cadence](#)
- [Self-Hosted Release Cadence](#)
- [Available Modules](#)
- [Security](#)
- [End of Life Notices](#)
- [Documentation](#)
- [Known Issues](#)

CloudCenter Suite 5.1.0 Release Date: August 19, 2019

Updated:

- November 5, 2019: Updated the Documentation section to list modified pages.
- November 9, 2019: Updated the Self-Hosted Release Cadence section below.

CloudCenter Suite SaaS is a managed, cloud-based service offered by Cisco. Cisco provides ongoing management, maintenance, and upgrades with end-to-end monitoring and 24/7 customer support. In addition, CloudCenter Suite SaaS offers a large number of pre-built, out-of-the-box integrations and content to ensure that your getting started experience with CloudCenter Suite is seamless. The following aspects of your CloudCenter Suite experience is handled by Cisco:

- Preparing the infrastructure based on architectural considerations
- Installing, Managing, and Upgrading CloudCenter Suite
- Installing, Managing, and Upgrading Modules
- Setting up the initial administrator
- Managing the cluster
- See [SaaS Access](#) for additional details

You can purchase CloudCenter Suite SaaS by contacting your Cisco Account Manager or a [Cisco sales representative](#).

CloudCenter Suite Self-Hosted allows you to purchase, host, and install the software from Cisco so you can access and manage the entire solution from a remote server or location that is located on your premises. You can choose to deploy CloudCenter Suite by yourself by installing it on-premises (VMware, OpenStack, and so forth) or by using a cloud provider (AWS, GCP, and so forth) as well as use the CloudCenter Suite out-of-the-box integrations to create a custom solution.

- Across modules, *only* releases within the same major/minor versions are supported with each other. For instance, Action Orchestrator 5.1.3 only works with Workload Manager 5.1.2, but Action Orchestrator 5.1.3 with Workload Manager 5.0.1 is not supported.
- The CloudCenter Suite has a common installer at the major release level to install, upgrade, and integrate all modules included in the suite.
  - The CloudCenter Suite installation corresponds directly to each major release, for instance CloudCenter Suite 5.1.
  - The minor release version is CloudCenter Suite 5.1.0, which is available as installer files for ALL components for all supported clouds.
- CloudCenter Suite 5.1 includes multiple modules that are available through the Suite Installer and initiated by the Suite Admin.
  - Each module can have access to additional gateways or endpoints that allow enterprises to add module-specific *components*.
  - Each major release can have multiple minor releases at the module level.
- The Kubernetes cluster can be upgraded to CloudCenter Suite 5.1.0 from CloudCenter Suite 5.0.x.
- The backup and restore functionality is available in CloudCenter Suite 5.1.0.
- See [Installer Overview](#) or [Suite Architecture](#) for additional details.

You can purchase CloudCenter Suite Hosted by contacting your Cisco Account Manager or a [Cisco sales representative](#) or [CloudCenter Suite Support team](#).

The following modules are part of CloudCenter Suite 5.1:

- [Suite Admin 5.1](#)
- [Workload Manager 5.1](#)
- [Action Orchestrator 5.1](#)
- [Cost Optimizer 5.1](#)

The release notes for each module is available in the links listed above.

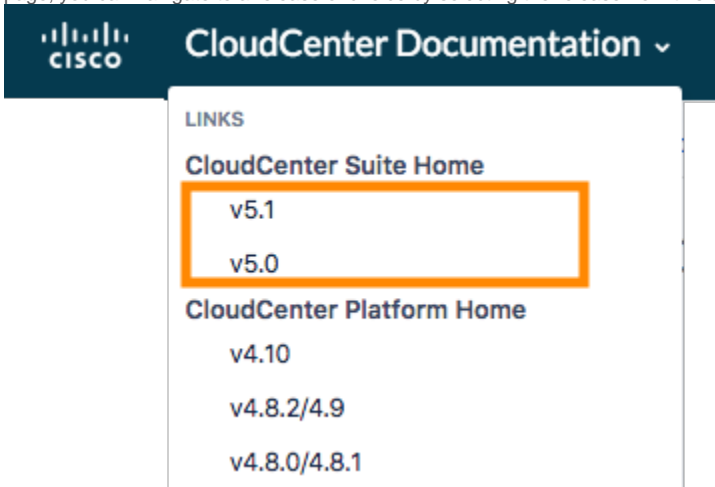
See [Security Considerations](#) for details.

See [End of Support Notices](#) for additional details.

The <https://docs.cloudcenter.cisco.com> website is the home of the following products:

- CloudCenter Suite 5.1 and later releases (includes documentation for all modules that are part of the CloudCenter Suite Suite, including the [Workload Manager](#), which is the new name for the legacy CloudCenter platform).
- The **CloudCenter Platform** 4.x releases (the legacy versions of the current Workload Manager).


- You can access one of the releases listed above from the dropdown list in the left header bar as displayed in the following screenshot. From any page, you can navigate to a release of choice by selecting the release from this dropdown list.



- The site-wide search bar in the header is now moved to the reading pane to ensure easy access and visibility.

## Welcome to CloudCenter Suite 5.1 Documentation

### Search Documentation



- The following documentation change was implemented in CloudCenter Suite 5.1.0:
  - [SaaS Access](#) (updated the overview details)

CloudCenter Suite 5.1.0 has no known issues.

**Back to:** [CloudCenter Suite Home](#)

# Browser Compatibility

## Browser and Resolution Compatibility

- [Browser Compatibility](#)
- [Resolution Requirements](#)

For CloudCenter Suite 5.1, Cisco supports the browser versions listed in the following table.

Browser	Version
Microsoft Edge	Version 42.17134.1.0 and HTML 17.17
Firefox	Version 68.0 and 67.0
Chrome	Version 75.0 and 74.0
Safari	Version 12.1.1 and 12.1.2

\* Internet Explorer is not supported.

Optimize your browser resolution by setting your monitor display to at least 1828 x 762 px to view the screen without scrolling.

**Back to:** [CloudCenter Suite Home](#)

# Support Information

## Support Information

- [Documentation Website](#)
- [Documentation Accessibility](#)
- [OpenSource Version Matrix](#)
- [End of Support Notices](#)

**Back to:** [CloudCenter Suite Home](#)

# Documentation Website

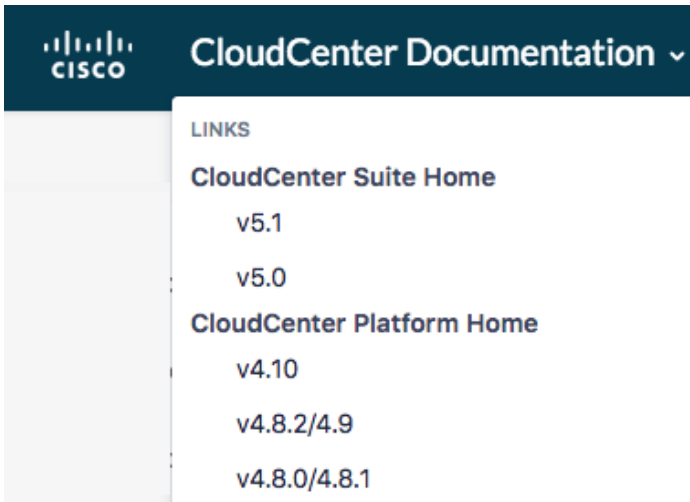
## Documentation Website

- [Website Compatibility](#)
- [Website Navigation](#)

For security and compliance reasons, CloudCenter Suite documentation (<https://docs.cloudcenter.cisco.com>) is accessible on browsers that support the Transport Layer Security (TLS) 1.2 protocol defined in RFC 5246. See your browser documentation for compliance details.

The <https://docs.cloudcenter.cisco.com> website is the home of the following products:

- The CloudCenter Suite 5.0 and later releases (includes documentation for all modules that are part of the CloudCenter Suite, including the [Workload Manager](#) module).
- The **CloudCenter Platform** 4.x releases (the older versions of the current [Workload Manager](#) module).



You can access one of the releases listed above from the dropdown list in the left header bar as displayed in the following screenshot.

From any page, you can navigate to your release of choice by selecting the release from this dropdown list!

**Back to:** [CloudCenter Suite Home](#)

# Documentation Accessibility

## Documentation Accessibility

- [Overview](#)
- [Accessibility Features](#)
- [Keyboard Shortcuts](#)

The information in this section applies to CloudCenter Suite Suite 5.1 releases.

For a list of accessibility features in CloudCenter Suite Suite 5.1, see [Voluntary Product Accessibility Template \(VPAT\)](#) on the Cisco website, or contact [accessibility@cisco.com](mailto:accessibility@cisco.com).

To expand the tree pane, follow this procedure:

1. Press the **return** key when the item is in focus.
2. Press the **tab** key to view the children for each item.

**Back to:** [CloudCenter Suite Home](#)



# OpenSource Version Matrix

## OpenSource Version Matrix

For a complete list of application versions for each module in the CloudCenter Suite, refer to the *CloudCenter OpenSource documentation* at [software.cisco.com](https://software.cisco.com) for the appropriate CloudCenter Suite *version*.

See the following image for additional details on finding this file: [software.cisco.com](https://software.cisco.com) > **CloudCenter** > *version* > **CloudCenter Suite OpenSource documentation**.

## Software Download

[Downloads Home](#) / [Cloud and Systems Management](#) / [Cloud Management](#) / [CloudCenter Suite](#) / CloudCenter- 5.0(0)

Expand All Collapse All

Latest Release ▼

5.0(0)

All Release ▼

5 ▼

5.0(0)

### CloudCenter Suite

Release 5.0(0)

[▲ Notifications](#)

Related Links and Documentation

- No related links or documentation -

File Information	Release Date	Size	↓	🛒	📄
The README describes the images required for installing CloudCenter Suite README_download_Cisco_CloudCenter-Suite-release-5.0.0	15-Feb-2019	0.00 MB	↓	🛒	📄
Jenkins Client Plugin ccs-JenkinsClient-5.0.0-20190215.hpi	15-Feb-2019	11.74 MB	↓	🛒	📄
Google CloudCenter Suite installer ccs-google-suiteinstaller-5.0.0-20190215.tar.gz	15-Feb-2019	5093.43 MB	↓	🛒	📄
Installers and artifacts needed to build Virtual Appliances using your operating system image ccs-installer-artifacts-5.0.0-20190215.tar	15-Feb-2019	121.49 MB	↓	🛒	📄
CloudCenter Suite OpenSource library documentation ccs-opensource-documentation-5.0.0-20190215.zip	15-Feb-2019	20.86 MB	↓	🛒	📄

Back to: [CloudCenter Suite Home](#)

15

Cisco Cloud Management Documentation

1

# End of Support Notices

## End-of-Sale and End-of-Life Announcements for Cisco CloudCenter Products

- [Cisco CloudCenter Suite](#)
  - [Cisco CloudCenter Suite \(On-Prem / Self-Hosted\)](#)
    - [Suite Installer](#)
    - [Suite Admin](#)
    - [Workload Manager/Cost Optimizer](#)
    - [Action Orchestrator](#)
  - [Cisco CloudCenter Suite SaaS](#)
- [Cisco CloudCenter Platform \(Legacy\)](#)
  - [Cisco CloudCenter Platform 4.10.x](#)
  - [Cisco CloudCenter Platform 4.9.x and prior](#)



*This bulletin provides the consolidated information for all Cisco CloudCenter products and replaces previously provided information.*

Cisco CloudCenter Suite releases are supported for up 18 months. However, Cisco reserves the right to change and defer support timelines as required. The Last Date of Support (LDOS) marks the last date for customers to receive applicable service and support as entitled by active service contracts for covered products. After this date, the service is no longer available.

### Cisco CloudCenter Suite (On-Prem / Self-Hosted)

#### Suite Installer

CloudCenter Release	Kubernetes Version	CCP Tenant Image	Release Date	LDOS
<a href="#">Suite Installer 5.1.1</a>	1.13.5	ccp-tenant-image-1.13.5.ova	September 26, 2019	March 26, 2021
<a href="#">Suite Installer 5.2.0</a>	1.16.3	ccp-tenant-image-1.16.3-ubuntu18-6.1.0.ova	May 9, 2020	November 9, 2021
<a href="#">Suite Installer 5.2.3</a>	1.16.3	ccp-tenant-image-1.16.3-ubuntu18-6.1.1.ova	October 13, 2020	April 13, 2022

Cisco announces the End-of-life and End-of-Support for Kubernetes clusters deployed by **5.0(x) Cisco CloudCenter Suite Installers**. No patches or maintenance releases will be provided. Support for modules running on older Kubernetes clusters will be *best effort* as determined by Cisco TAC. Customers are always encouraged to use the latest Suite Installer to backup and restore their existing CloudCenter Suite application to a supported Kubernetes cluster version.

#### Suite Admin

CloudCenter Release	Release Date	LDOS
<a href="#">Suite Admin 5.0</a>	February 16, 2019	August 16, 2020
<a href="#">Suite Admin 5.1</a>	August 19, 2019	February 19, 2021
<a href="#">Suite Admin 5.2</a>	May 9, 2020	November 9, 2021

#### Workload Manager/Cost Optimizer

CloudCenter Release	Release Date	LDOS
<a href="#">Workload Manager 5.0/Cost Optimizer 5.0</a>	February 16, 2019	August 16, 2020
<a href="#">Workload Manager 5.1/Cost Optimizer 5.1</a>	August 19, 2019	February 19, 2021
<a href="#">Workload Manager 5.2/Cost Optimizer 5.2</a>	March 31, 2020	September 30, 2021
<a href="#">Workload Manager 5.3/Cost Optimizer 5.3</a>	May 7, 2020	November 7, 2021
<a href="#">Workload Manager 5.4/Cost Optimizer 5.4</a>	July 30, 2020	January 30, 2022

#### Action Orchestrator

CloudCenter Release	Release Date	LDOS
<a href="#">Action Orchestrator 5.0</a>	February 16, 2019	August 16, 2020

<a href="#">Action Orchestrator 5.1</a>	August 19, 2019	February 19, 2021
<a href="#">Action Orchestrator 5.2</a>	May 29, 2020	November 29, 2021

### Cisco CloudCenter Suite SaaS

Cisco [announces](#) the end-of-sale and end-of-life dates for **Cisco CloudCenter Suite SaaS**. Customers with active service contracts will continue to receive support from the Cisco Technical Assistance Center (TAC) as shown below. The following table describes the end-of-life milestones, definitions, and dates for the affected product(s).

End-of-Life Milestones		
Milestone	Definition	Date
End-of-Life Announcement Date	The date the document that announces the end-of-sale and end-of-life of a product is distributed to the general public.	November 6, 2020
Last Date of Support (LDOS)	The last date to receive applicable service and support as entitled by active service contracts for covered products. After this date, the service is no longer available.	<b>TBD</b> <i>Target - Feb 2021</i>

For additional information please review [Cisco's End-of-Life Policy](#) and the [End-of-Life and End-of-Sale Notices for CloudCenter Suite](#).

### Cisco CloudCenter Platform 4.10.x

Cisco [announces](#) the end-of-sale and end-of-life dates for the **Cisco CloudCenter Platform (Legacy/4.x)**. The last day to order the affected product(s) is **May 7, 2021**. Customers with active service contracts will continue to receive support from the Cisco Technical Assistance Center (TAC) as shown below. The following table describes the end-of-life milestones, definitions, and dates for the affected product(s).

End-of-Life Milestones		
Milestone	Definition	Date
End-of-Life Announcement Date	The date the document that announces the end-of-sale and end-of-life of a product is distributed to the general public.	November 6, 2020
End-of-Sale Date	The last date to order the product through Cisco point-of-sale mechanisms. The product is no longer for sale after this date.	May 7, 2021
Last Date of Support (LDOS)	The last date to receive applicable service and support as entitled by active service contracts for covered products. After this date, the service is no longer available.	May 7, 2024

For additional information please review [Cisco's End-of-Life Policy](#) and the [End-of-Life and End-of-Sale Notices for CloudCenter Platform \(Legacy\)](#).

### Cisco CloudCenter Platform 4.9.x and prior

Cisco [announces](#) the End-of-life and End-of-Support for versions of **Cisco CloudCenter Platform 4.9.1 and earlier**. Software maintenance support for all versions listed above ended on **October 31, 2020**. No patches or maintenance releases will be provided. Customers are encouraged to migrate to Cisco CloudCenter Platform 4.10.0.10 or to Cisco CloudCenter Suite 5.2.3 or later.

**Back to:** [CloudCenter Suite Home](#)

# Security Considerations

## Security Considerations

- [Overview](#)
- [Product Overview](#)
- [CloudCenter Suite Architecture](#)
- [User Authentication](#)
- [Cloud Authentication](#)
- [REST API Calls](#)
- [UI Authentication](#)
- [Module Security](#)
- [Role-Based Access Control](#)

This section provides design specification details related to the security of the CloudCenter Suite.

This section DOES NOT provide on operational policies such as key rotation, incident management and business continuity policies are not covered in this document.

CloudCenter Suite is an enterprise-class solution that offers a secure, scalable, extendable, and multi-tenant solution that can scale to meet the needs of the most demanding IT organizations and cloud service providers.

CloudCenter Suite uses various types of metadata, authentication information (such as *customer credentials and keys*), cloud usage metrics, and users associated with cloud applications to deploy and manage applications on cloud infrastructures.

The CloudCenter Suite does not store *customer application data* (data that is created, used, or managed by the user's cloud applications).

- Customer application data is only stored on customer premises or on cloud infrastructures.
- Customer application data is not stored or accessed by CloudCenter Suite at any point.

CloudCenter Suite provides end-to-end security with:

- A comprehensive key management mechanism
- Full application and application tier network isolation (micro-segmentation)
- Data encryption for data both in transit and at rest
- User identity management and authentication control
- User, application, and object-level access control

The CloudCenter Suite architecture is deployed as a distributed architecture and is composed of several key architectural components as described in [The CloudCenter Suite Architecture](#).

CloudCenter Suite supports user password, hash-based authentication, and SAML 2.0-based Single Sign-On (SSO) authentication. CloudCenter Suite also provides authentication for REST API endpoint access.

CloudCenter Suite authenticates users through a unique username and password. The password is not stored in clear-text, but is converted using a secure one-way hash algorithm (SHA256) with a random salt. If different users use the same password, this will not result in the same password hash. This hash code is generated and stored when the user creates the password for the first time or changes the password at a later time. Upon login, the hash code is regenerated using the specified password and matched against the stored hash code to authenticate the user. Since this is a one-way hash algorithm, no Cisco employee or third-parties can discover the user password. The password is neither reverse recoverable, nor subject to brute force dictionary attack.

CloudCenter Suite leverages SAML (2.0) to integrate with customer identity platforms such as Active Directory (AD) and LDAP. For SAML-based SSO authentication, the user directory, password, and authentication mechanism are controlled by the customer. Customers may further choose to enable multi-factor authentication on their user login page through well-known identity provider platforms such as ADFS, Ping Identity, Okta, and so forth. The CloudCenter Suite only uses the user's email address as the user identity in SSO mode. Customers can configure unique SAML Identity Providers (IdP) properties on a per tenant basis. The CloudCenter Suite tenant admin can optionally set additional mapping rules to automatically sync user groups and user group membership based on custom properties provide by IdP

The CloudCenter Suite authenticates to public, private, and hybrid clouds using cloud account credentials provided to CloudCenter Suite when a user configures cloud environments. These cloud account credentials are stored securely in the CloudCenter Suite database using AES-256 encryption.

Configuring and registering clouds and cloud accounts in CloudCenter Suite is limited to CloudCenter Suite administrators. The CloudCenter Suite administrator can decide if additional tenant administrators and end-users can configure their own cloud account information. See [Initial Administrator Setup](#) for details.



Cisco provides CSRF protection for all API calls. See [CSRF Token Protection](#) for additional details.

Access to the REST API interface is limited to configured user accounts. To authenticate API requests, all CloudCenter Suite REST APIs require basic authentication using an API key as the password. For example:

```
curl -H "Accept:application/json" -H "Content-Type:application/json" -u <user_accountNumber>:<api_key> -X GET https://<HOST>:<PORT>/api/v1/suite-idm/currentUser/userInfo
```

In addition to the user's *accountNumber.apikey* combination, all CloudCenter Suite REST APIs can also accept the *JSON Web Token (JWT)*. For example:

```
curl -H "Accept:application/json" -H "Content-Type:application/json" -H "Authorization: Bearer <JWT>" -X GET https://<HOST>:<PORT>/api/v1/suite-idm/currentUser/userInfo
```

A REST API key is a 36-character, randomly generated, case-sensitive, hexadecimal UUID string. This key, combined with the user's unique Account Number (*accountNumber*), is used for REST API authentication. During authentication, the REST API key specified in the HTTPS request is matched with the REST API key stored in the CloudCenter Suite database. This prevents the user from revealing the real user password in any automation script, and also allows REST API authentication to work with either user/password hash-based or SAML SSO-based authentication.

To provide data security, all REST API requests must be issued over a secure, encrypted, HTTPS connection.

The REST API key for each user is stored securely in CloudCenter Suite database using SHA256 one-way hash. The [API Key](#) section provides additional details about secure key storage and key operations. See [Suite Admin API](#) for details on CloudCenter Suite REST APIs and how to use them.

All users can generate their own API keys – the Suite Admin has no control over this function.

The CloudCenter Suite UI requires user authentication. Each authenticated user will have a unique Session ID to track activities and a JWT to ensure API access. The JWT expires in 15 minutes and the UI auto-refreshes the JWT token if it detects the user actively using the UI. If the user is logged off or if the user is disabled or deleted, the user's active JWT is no longer valid.

The CloudCenter Suite connects to a Cisco hosted Helm repository and a Docker registry to check for available modules and updates. These repositories are fully compliant with export control and requires authentication for each user connecting to the repository. All CloudCenter Suite module are packaged as Helm Chart and Docker images. The Helm Chart refers to Docker images via the image's SHA256 hash. The Helm Chart itself is signed and verified by the CloudCenter Suite upon installation or upgrade. This way the integrity of the Helm Chart and Docker images are guaranteed.

The CloudCenter Suite offers granular control of access to each CloudCenter Suite resource through role-based, module-level access control. Access to resources like services, clouds, application profiles, deployment environments, and other CloudCenter Suite resources can be managed based on roles associated with users or user groups. See [Understand Roles](#) for details.

**Back to:** [CloudCenter Suite Home](#)

# Module Versions

## Module Versions

- [Suite Admin Versions](#)
- [Workload Manager Versions](#)
- [Action Orchestrator Versions](#)
- [Cost Optimizer Home](#)

1. Suite Installer 5.1 Home	2
1.1 Suite Architecture	3
1.2 Self-Hosted Installation	5
1.2.1 Installer Overview	6
1.2.2 Installer Virtual Appliances	7
1.2.2.1 Virtual Appliance Overview	8
1.2.2.2 Amazon Appliance Setup	10
1.2.2.3 Azure Appliance Setup	11
1.2.2.4 GCP Appliance Setup	12
1.2.2.5 OpenStack Appliance Setup	16
1.2.2.6 VMware vSphere Appliance Setup	19
1.2.3 Prepare Infrastructure	29
1.2.4 New Cluster Installation	31
1.2.4.1 Amazon EKS Installation	32
1.2.4.2 Azure AKS Installation	36
1.2.4.3 Google GKE Installation	39
1.2.4.4 OpenStack Installation	43
1.2.4.5 VMware vSphere Installation	47
1.2.5 Existing Cluster Installation	55
1.2.6 Upgrade Kubernetes Cluster	58
1.2.6.1 Upgrade Approach	59
1.2.6.2 Amazon EKS Upgrade	63
1.2.6.3 Azure AKS Upgrade	66
1.2.6.4 Google GKE Upgrade	67
1.2.6.5 OpenStack Upgrade	71
1.2.6.6 VMware vSphere Upgrade	77
1.2.7 Offline Repository	85
1.2.8 Backup and Restore	89
1.2.8.1 With Internet Access	90
1.2.8.1.1 Backup	91
1.2.8.1.2 Restore	98
1.2.8.2 Without Internet Access	127
1.2.9 Troubleshooting	133
1.3 Suite Admin Workflow	140
1.4 Initial Administrator Setup	142
1.5 Kubernetes Cluster Management	144
1.5.1 Cluster Status	145
1.5.2 Manage Clusters	147
1.6 Configure Smart Licenses	149
1.7 Module Lifecycle Management	159
1.7.1 Install Module	160
1.7.2 Update Module	164
1.7.3 Monitor Modules	167

# Suite Installer 5.1 Home

## Self-Hosted 5.1 Documentation

Cisco released Suite Admin releases as follows:

- [Suite Admin 5.1.0](#) released on August 19, 2019
- [Suite Admin 5.1.1](#) released on September 26, 2019
- [Suite Admin 5.1.2](#) released on November 25, 2019

Search

[Suite Installer 5.2 Home](#)

updated Jan 28, 2021

[view change](#)

[Backup Approach](#)

updated Jan 12, 2021

[view change](#)

[Private Cloud](#)

updated Dec 05, 2020

[view change](#)



# Suite Architecture

## Suite Architecture

- [Overview](#)
- [The Suite Architecture](#)
- [Port Requirements](#)
- [The Suite Admin](#)
- [The Modules](#)

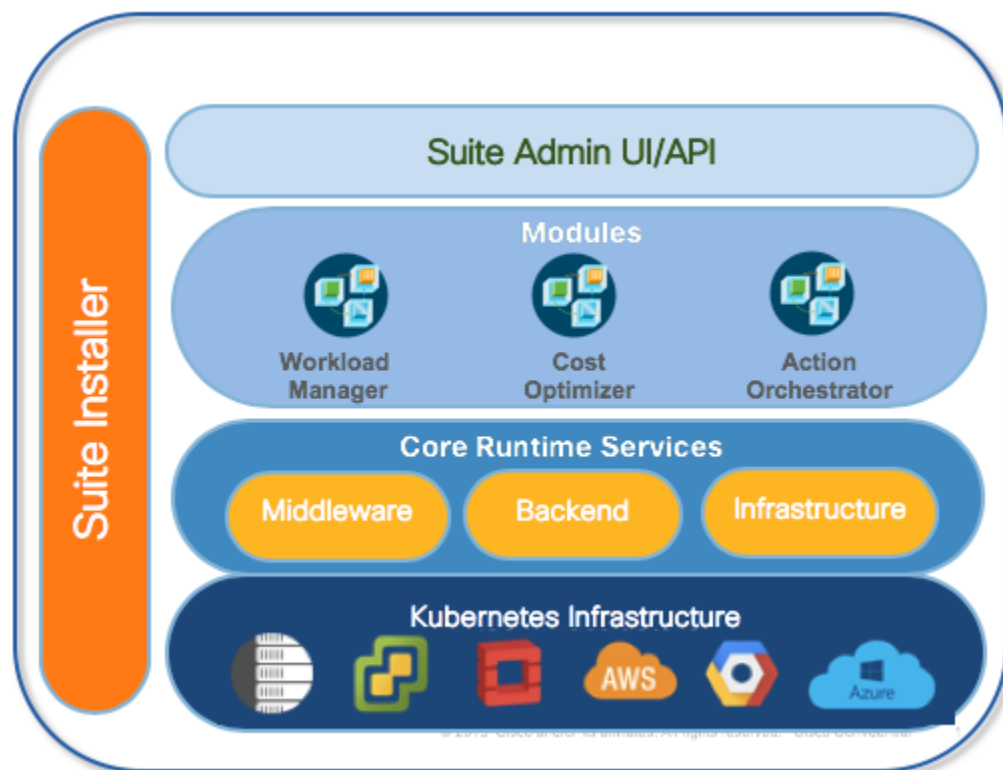
The CloudCenter Suite is Cisco's hybrid cloud deployment platform. This platform takes a unique approach to install, configure, and maintain hybrid cloud environments that are often encountered by Information Technology (IT) departments to adopt business agility and improve time-to-market solutions within an enterprise. As a cloud-based organization, your enterprise can choose from multiple cloud (*multicloud*) providers depending on your location, policies, permissions, security requirements, and governance regulations for both traditional and modern IT requirements.

The CloudCenter Suite provides a solution that is cloud agnostic, works with diverse workloads, provides cross-domain orchestration, supports cost-optimization, and integrates easily in an agile world.


The CloudCenter Suite is made up of the following components:

- **Suite Installer** – Installs the Suite Admin. See [Suite Installer](#) for additional details.
- **Suite Admin** – Installs and launches a suite of modules. See *The Suite Admin* section below for additional details.
- **Modules** – The Workload Manager, the Cost Optimizer, and the Action Orchestrator. See *The Modules* section below for additional details.
- **Core Runtime Platform and Kubernetes Infrastructure** – A Kubernetes-based platform that allows you to launch each module on a new or existing Kubernetes cluster.

The following image displays the Suite Admin architecture.



The following image identifies the ports that must be open for the CloudCenter Suite to function as designed.



When you download and install the Suite Installer, the [Suite Admin](#) is **already installed**! You have the option to use the Suite Admin UI to perform the following tasks:

- Install additional, available modules based on the list available in the Dashboard.
- Upgrade the Suite Admin or other installed modules when a new version becomes available.

The Suite Admin facilitates the installation of the following modules:

- **Workload Manager:**
  - This module allows IT organizations to provide management for clouds (public/private/container), applications, VMs/pods, governance policies with centralized visibility and permission control for enterprise environments.
  - See [Workload Manager](#) for additional details.
- **Action Orchestrator:**
  - This module allows IT organizations to use cross-domain orchestration to automate a process that has multiple, complex steps with a specific order and implemented across different technical domains.
  - See [Action Orchestrator](#) for additional details.
- **Cost Optimizer:**
  - This module allows IT organizations to use cost optimization in a pay-per-use environment to avoid consumption that does not add value.
  - See [Cost Optimizer](#) for additional details.

Each module in the CloudCenter Suite is independent and allows access to additional gateways or endpoints so you can add on module-specific components on supported clouds.

# Self-Hosted Installation

## Self-Hosted Installation

- [Installer Overview](#)
- [Installer Virtual Appliances](#)
- [Prepare Infrastructure](#)
- [New Cluster Installation](#)
- [Existing Cluster Installation](#)
- [Upgrade Kubernetes Cluster](#)
- [Offline Repository](#)
- [Backup and Restore](#)
- [Troubleshooting](#)


# Installer Overview

## Installer Overview

- [Overview](#)
- [Supported Clouds](#)
- [Installer Appliance Download Location](#)

The CloudCenter Suite provides a new way to install, configure, and maintain multiple modules that jointly make up the suite. The CloudCenter Suite has a common installer to install, upgrade, and integrate all modules included in the suite.

You can install the CloudCenter Suite by using installer appliance images provided by Cisco. As part of the installation process, the CloudCenter Suite installs the Suite Admin. Once authenticated, each user can access the CloudCenter Suite using valid credentials created by the Suite Administrator.


 Installers are already incorporated in the CloudCenter Suite SaaS offer, see [SaaS Access](#) for additional details.

Cisco supports the corresponding Kubernetes engine (or managed services) for the following public clouds for the CloudCenter Suite:

- [Amazon Elastic Container Service for Kubernetes \(Amazon EKS\)](#)
- [Google Kubernetes Engine \(GKE\)](#)
- [Azure Kubernetes Service \(AKS\)](#)

Cisco supports the following private clouds for the CloudCenter Suite:

- [VMware vSphere](#)
- [OpenStack](#)

 All supported clouds are visible and enabled for private cloud installers.  
Only public clouds are visible and enabled for public cloud installers.  
This includes both the functionality and the CloudCenter Suite UI.

Major releases include installer appliances for the following components and cloud providers.

You can download these files from [software.cisco.com](https://software.cisco.com).

The [Virtual Appliance Overview](#) section provides more details on these files.

# Installer Virtual Appliances

## Installer Virtual Appliances

- [Virtual Appliance Overview](#)
- [Amazon Appliance Setup](#)
- [Azure Appliance Setup](#)
- [GCP Appliance Setup](#)
- [OpenStack Appliance Setup](#)
- [VMware vSphere Appliance Setup](#)

# Virtual Appliance Overview

## Virtual Appliance Overview

- [Virtual Appliance Overview](#)
- [General Virtual Appliance Approach](#)
- [Cloud-Specific Setup](#)

The only way to install the CloudCenter Suite is to use the virtual appliance Installer method. Cisco builds these appliances on CentOS 7.x base images.





Installers are already incorporated in the CloudCenter Suite SaaS offer, see [SaaS Access](#) for additional details.

To prepare infrastructure for the appliance approach, follow this process.

1. Review and ensure that you have met the requirements to [Prepare Infrastructure](#) before installing the CloudCenter Suite.
2. Review the list of [Supported Suite Installers](#) to verify the supported Virtual Appliances.
3. Navigate to [software.cisco.com](https://software.cisco.com) to download virtual appliances for each supported cloud.
4. Follow directions as specified in the table below to obtain and import each image.

Cloud	Image Type	Description
AWS	Shared image (AMI)	<p>Obtain launch permissions for the AWS account. Refer to the <a href="#">AWS documentation</a> for additional context.</p> <p>Request image sharing for the AWS account by opening a CloudCenter Support case (<a href="https://mycase.cloudapps.cisco.com/case">https://mycase.cloudapps.cisco.com/case</a> or <a href="http://www.cisco.com/c/en/us/support/index.html">http://www.cisco.com/c/en/us/support/index.html</a>). In your request, specify the following details:</p> <ol style="list-style-type: none"> <li>a. Your AWS account number</li> <li>b. Your CloudCenter Suite version</li> <li>c. Your Customer ID (CID)</li> <li>d. Your customer name</li> <li>e. Specify if your setup is in production or for a POC</li> <li>f. Your Contact Email address</li> </ol>
Azure	Downloaded Virtual Appliance (VHD from the ZIP folder)	Create a new Azure image using the provided VHD file provided by Cisco and launch a VM using that image. Refer to the <a href="#">Azure documentation</a> for additional context.
GCP	Shared image	Create a new GCP image using the provided VHD provided by Cisco and launch a VM using that image. Refer to the <a href="#">GCP documentation</a> for additional context
OpenStack	Downloaded Virtual Appliance (QCOW2)	Import the QCOW2 image file using the OpenStack client. Refer to the <a href="#">OpenStack Documentation</a> for additional context.

VMware vSphere	Downloaded Virtual Appliance (OVA)	<p>Follow this procedure:</p> <ol style="list-style-type: none"> <li>a. Download the OVA image.</li> <li>b. Import the OVA to your vSphere environment by using the vSphere client             <ol style="list-style-type: none"> <li>i. When you import the OVA as a VM, ensure that it is powered <b>off</b> on vSphere.</li> <li>ii. If your environment requires a static IP, use a <a href="#">VMware Customization Spec</a> to manually configure the static IP for the installer VM.</li> </ol> </li> <li>c. A default password is required to ensure access to the VM using the console (in case the SSH has issues).             <div style="border: 1px solid green; padding: 5px; margin: 5px 0;"> <p> If you provide a default password or public-key, be aware of the following requirements:</p> <ul style="list-style-type: none"> <li>• The login user is the cloud-user.</li> <li>• If you configure a default password or public key in the VM, you must also configure the default instance ID and hostname fields as they are dependent and required fields.</li> <li>• Use this password to access the VM via vSphere console.</li> <li>• You cannot use this password to SSH into the launched VMs.</li> </ul> </div> </li> <li>d. Select the required Network for the interface to be connected.</li> <li>e. Convert the VM to a template.             <div style="border: 1px solid orange; padding: 5px; margin: 5px 0;"> <p> You <i>must</i> convert the VM to template and then create a VM from this template, so that the template can be used when installing a VMware data center. If you do not provide the template name when installing a VMware data center, your installation will fail.</p> </div> </li> <li>f. Select the template created in the previous step and <i>clone to Virtual Machine</i>, to launch the installer VM. This template will also be used as the value for the <i>vSphere Template Name</i> cloud setting, in the installer UI.</li> <li>g. After the VM is created from the template, power it on. To access the UI, go to the newly created VM IP using HTTPS protocol in a supported browser (see <a href="#">Browser Compatibility</a>).</li> </ol>
----------------	------------------------------------	--

5. Launch the installer instance using the image.



The per-cloud setup procedures are only listed below to serve as sample setup scenarios.

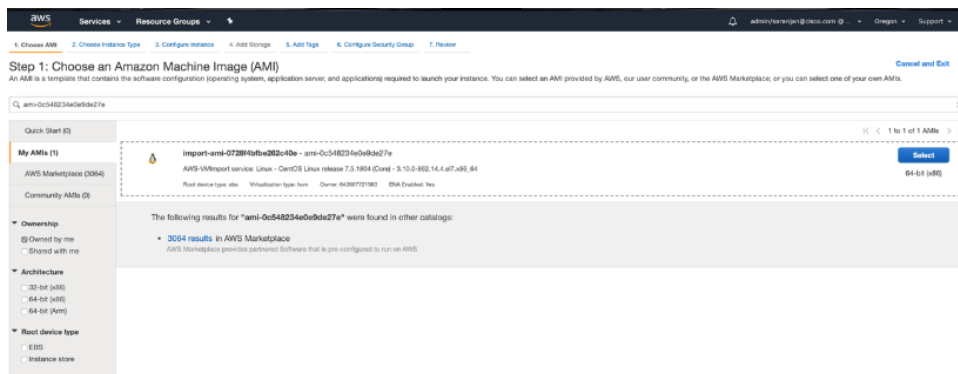
- [Amazon Appliance Setup](#)
- [GCP Appliance Setup](#)
- [Azure Appliance Setup](#)
- [OpenStack Appliance Setup](#)
- [VMware vSphere Appliance Setup](#)

# Amazon Appliance Setup

## Amazon Appliance Setup

To setup infrastructure for Amazon, follow this process.

1. Request image sharing for the AWS account by opening a [CloudCenter Support case](#). In your request, specify the following details:
  - a. Your AWS account number
  - b. Your CloudCenter Suite version
  - c. Your Customer ID (CID)
  - d. Your customer name
  - e. Specify if your setup is in production or for a POC
  - f. Your Contact Email
2. After you open a case, your support case is updated with the share AMI IDs. **Proceed to the next step only after your support case is updated with the AMI IDs.**
3. Navigate to the EC2 dashboard and search for the AMI ID name provided in the [CloudCenter Support case](#) (from Step 2 above)
4. Launch the EC2 instance using the AMI.
  - a. Navigate to the EC2 dashboard (the following screenshot displays a sample EC2 dashboard).



- b. Create EC2 instance in desired Region, VPC, subnet.
  - i. Choose an Instance Type.
  - ii. Configure the instance details for your environment.
  - iii. Review the instance launch details.
  - iv. Select an existing key-pair or create a new pair as required.
  - v. Create a security group with Ports 443, 80 (and optionally, 22) to be open.
  - vi. Launch the instance with the security group and key pair created in the previous two steps.
  - vii. Access the installer using the IP of the launched instance via HTTPS from your favorite browser.




# Azure Appliance Setup

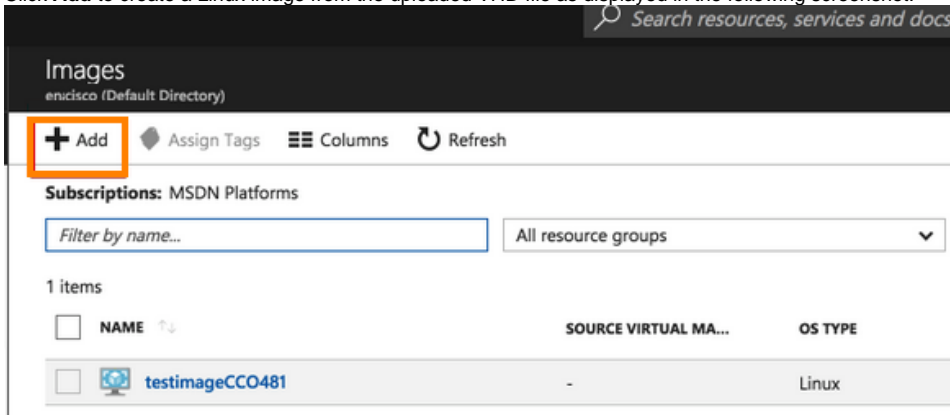
## Azure Appliance Setup

To setup infrastructure for Azure clouds, follow this process.

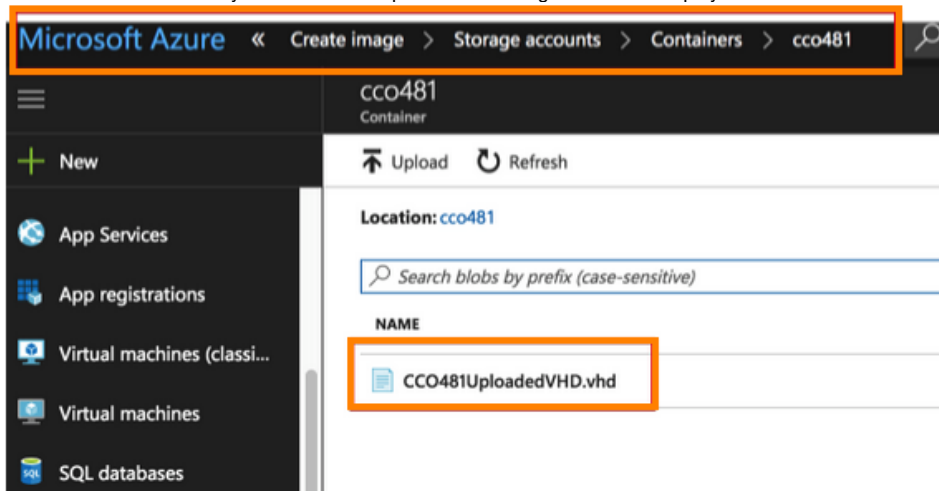
1. Upload the Cisco-provided VHD to the desired Azure Storage account/Region. See [https://www.ibm.com/support/knowledgecenter/en/SSPREK\\_9.0.6/com.ibm.isam.doc/admin/task/tsk\\_upload\\_vhd\\_azure.html](https://www.ibm.com/support/knowledgecenter/en/SSPREK_9.0.6/com.ibm.isam.doc/admin/task/tsk_upload_vhd_azure.html) for detailed instructions.

 You must use the Azure CLI to perform this upload.

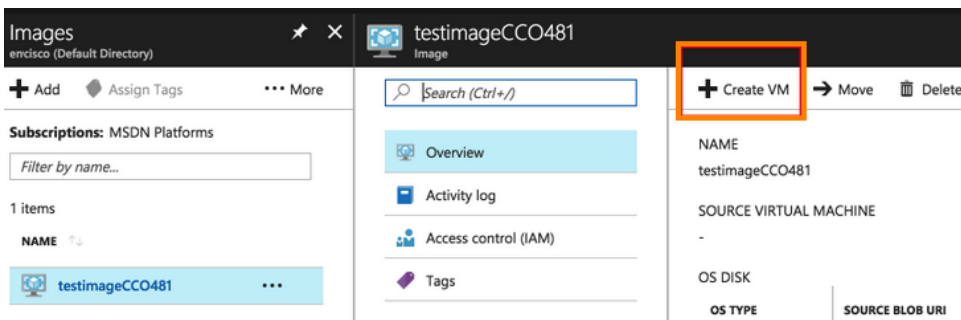
2. Click **Add** to create a Linux image from the uploaded VHD file as displayed in the following screenshot.



3. Select the disk name that you created in Step 2. The following screenshot displays a disk name called **CCO4100UploadedVHD.vhd**.



4. Click **Create VM** to spin up a VM using the created image from Azure console as displayed in the following screenshot.



You have now setup the installer for an Azure cloud.

# GCP Appliance Setup

## GCP Appliance Setup


- [Overview](#)
- [Cloud Storage Bucket](#)
- [Create the Image](#)
- [Create the Instance](#)

Setting up the GCP appliance, is a multi-step process:


- Address the prerequisite permissions
- Create a storage [bucket](#) using the tar.gz file provided by Cisco
- Create the image
- Create the instance

To upload Cisco's tar.gz file to the GCP bucket, follow this process.

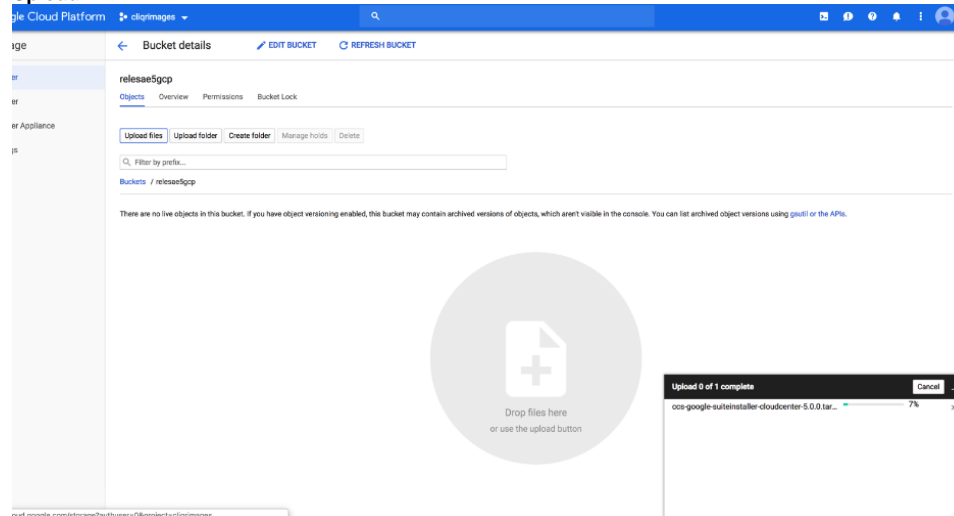
1. Open the Cloud Storage browser in the Google Cloud Platform Console as displayed in the following screenshot.



2. Click **Create bucket** and complete the required information for your environment. The following screenshot provides a sample setup.



- Upload the the tar.gz file provided by Cisco by dragging and dropping the file to the main pane as visible in the following screenshot or by clicking **Upload**.

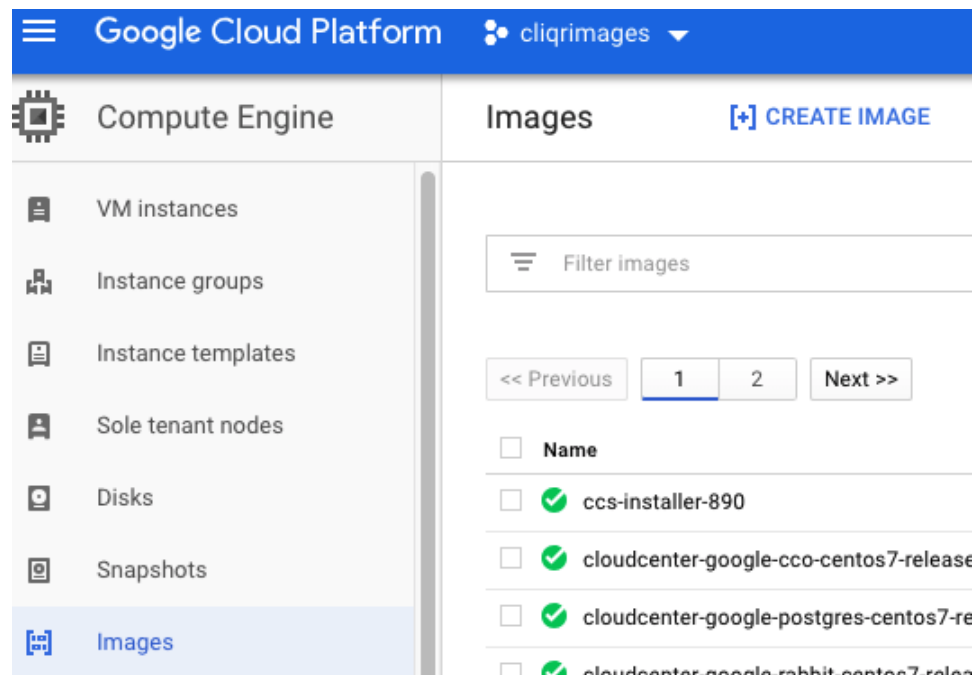


Uploading the file might take some time based on your network speed.

- After the upload is complete, use the same bucket to create the image as described in the next section.

To create an image, follow this process.

- Login to Google Cloud Platform.
- Create a Service Account with the following permissions:
  - Kubernetes Engine (Admin)
  - Compute Engine (Admin)
  - Service Account (User)
- Select **Compute Engine**.
- Click on **Images**.
- Click on **Create Image** as displayed in the following screenshot.



- Provide a **Name** for the new instance, select **Cloud Storage File** as the **Source**, browse and select the *image* file from the cloud storage bucket (uploaded in Step 2 above) for your environment and click **Create** to create an image as displayed in the following screenshot.

Google Cloud Platform cliqrimages

Compute Engine

← Create an image

Name <sup>?</sup>  
cloudcentersuite-v1

Family (Optional) <sup>?</sup>

Description (Optional)

Labels <sup>?</sup> (Optional)  
+ Add label

Encryption  
Data is encrypted automatically. Select an encryption key management solution.

Google-managed key  
No configuration required

Customer-managed key  
Manage via Google Cloud Key Management Service

Customer-supplied key  
Manage outside of Google Cloud

Source <sup>?</sup>  
Cloud Storage file

Cloud Storage file <sup>?</sup>  
Your image source must use the .tar.gz extension and the file inside the archive must be named disk.raw. [Learn more](#)

bucket/folder/file

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

7. Select the bucket where you uploaded the Cisco provided tar.gz file as displayed in the following screenshot.

Google Cloud Platform cliqrimages

Compute Engine

← Images

suite-950

Labels  
None

Creation time  
Dec 15, 2018, 4:27:39 AM

Encryption type  
Google managed

Equivalent [REST](#)

1. Navigate to the **GCP > Compute Engine > VM Instances** section and click **Create an Instance** as displayed in the following screenshot.

Google Cloud Platform cliqrimages

← Create an instance

To create a VM instance, select one of the options:

- New VM instance**  
Create a single VM instance from scratch
- New VM instance from template**  
Create a single VM instance from an existing template
- Marketplace**  
Deploy a ready-to-go solution onto a VM instance

**Name**  
instance-1

**Region** us-west1 (Oregon) **Zone** 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 20 GB standard persistent disk  
Image suite-950 [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](#)

**Create** **Cancel**

2. Select appropriate values for the new instance and click **Create**.



- Check the button to Allow HTTP or HTTPS access
- Change ports should list 443, 5671

3. Once the instance is created use the assigned public IP for this instance to access the suite installer UI.

You have now setup the installer for an GCP cloud.

# OpenStack Appliance Setup


## OpenStack Appliance Setup

To setup infrastructure for OpenStack clouds, follow this process.



The exact VM size really depends on the instance type configuration in your environment! See [Prepare Infrastructure](#) > *Resource Requirements for CloudCenter Suite Modules* for additional details.

1. Download the CloudCenter Suite QCOW2 file to your local machine.
2. Login into your OpenStack datacenter to perform this task.
  - a. Click **Images**.
  - b. Click the **Create Image** button.
  - c. Enter a valid name.
  - d. Click the **File Browse** button.
  - e. Select the QCOW2 file stored in your local machine as displayed in the following screenshot.



3. In the **Format** dropdown, select QCOW2.
4. To share this image with other users, select **Public** in the **Image Sharing Visibility** field.

5. Click **Next** and then click the **Create Image** button as displayed in the following screenshot.



The image import will take some time depending on the network speed. During this time, do not close the browser/application/tab.

6. Create the instance for each component using the imported images:
- Follow the standard OpenStack procedure to create the instance from an image.
  - Create the security group(s) with Port 80 and 443 (optionally 22 if you need SSH access) open for Ingress and Outbound communication.
  - You may need to assign floating IP to your VM after you create the VM is created.
7. Select a new or existing key pair to log into each instance – if multiple key pairs are available, you must *select one* to be used for the CloudCenter instance as displayed in the following screenshot.



If you do not select a key pair, you will not be able to log into the component VM!

You have now setup the installer for an OpenStack cloud.



# VMware vSphere Appliance Setup

## VMware vSphere Appliance Setup

To setup infrastructure using CloudCenter appliances for VMware vSphere clouds, follow this process.

1. **Configure Network Time Protocol (NTP) on the VMware ESXi hosts – this is important as the CloudCenter Suite installation can fail, if NTP is not configured or if it is wrongly configured.**

See [https://kb.vmware.com/s/article/57147?lang=en\\_US](https://kb.vmware.com/s/article/57147?lang=en_US) for additional details.



Note the value that you enter in this field for later use. You will need to enter the same values for the **NTP Servers** or **NTP Pools** fields in the Placement Properties page (see [VMware vSphere Installation](#) > Advanced Installation Process > Step 6).

Identical NTP values are required to ensure that the NTP communication between the installer and CloudCenter Suite master/worker VMs are in sync so the certificates generated by the installer for CloudCenter Suite are also in sync.

2. Download the OVA image file from [software.cisco.com](https://software.cisco.com) to your local machine.



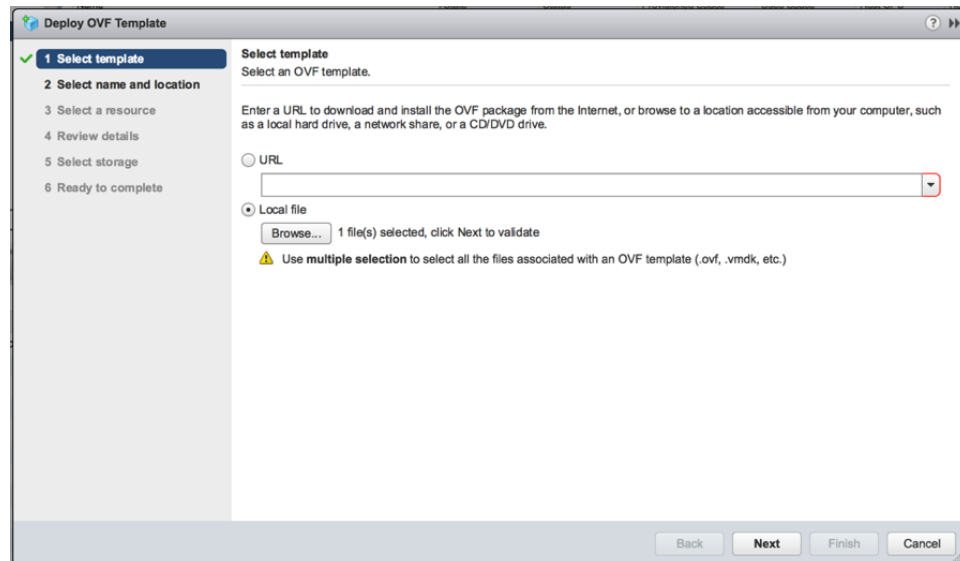
The installer appliance has/requires a minimum resource requirement of 4 vCPUs and 75 GB storage (root disk).

3. Log into the VMware Datacenter console and click on the **VMs and Templates** section.
4. Deploy an OVA template (right-click and select Deploy OVA Template option).

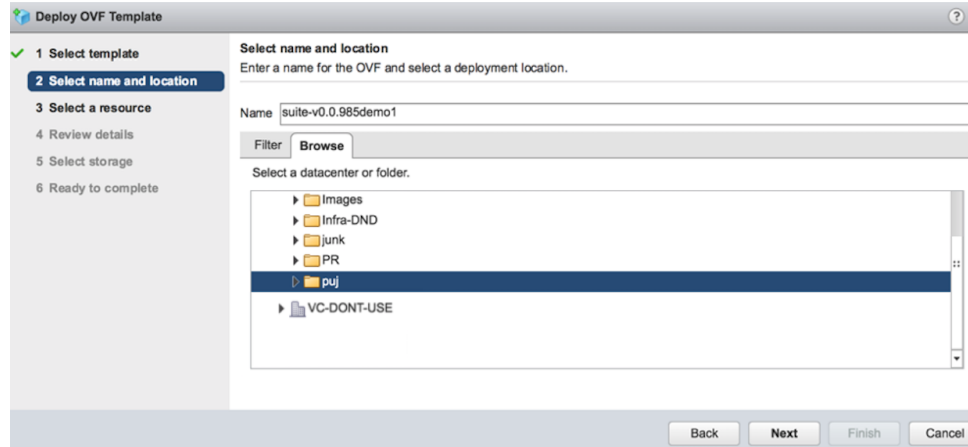
- a. If DHCP *is* installed, follow these steps.

Follow these steps **ONLY** if DHCP *is* installed.

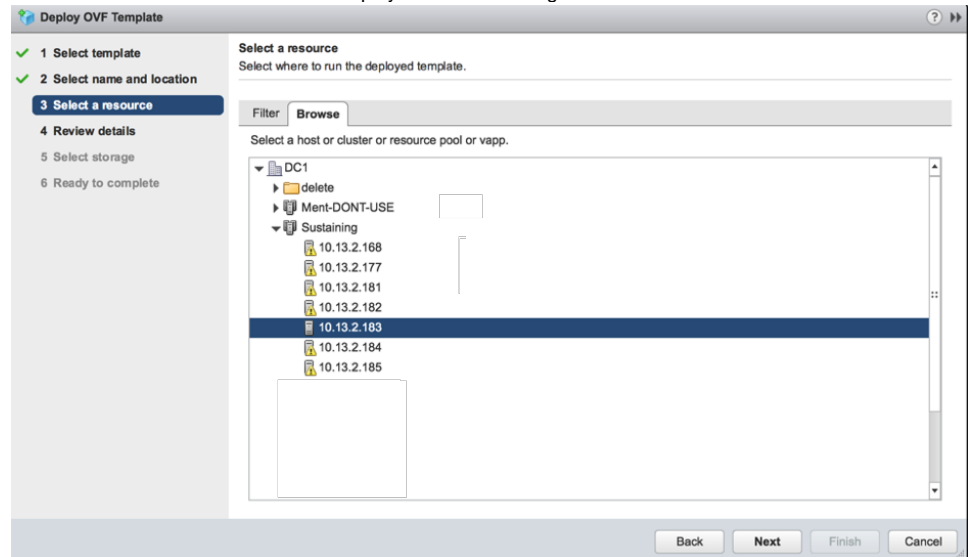
- i. Click the **Local file** option, click **Browse** to provide the location for the downloaded OVA file, ensure the file is selected, and then click **Next** as displayed in the following screenshot.



- ii. Provide a suitable name and select the target folder where you need to create the Template as displayed in the following screenshot.



- iii. Select a suitable host and cluster as displayed in the following screenshot.



- iv. Review the details as displayed in the following screenshot.



- v. Select the storage location as displayed in the following screenshots.



Use **Thin Provision** as the storage format so it has the flexibility to optimize the storage location. The following screenshots displays views from two different datacenters to provide a point of context.

Select storage

Select the datastore in which to store the configuration and disk files

Select virtual disk format:

- Same format as source
- Thick Provision Lazy Zeroed
- Thick Provision Eager Zeroed
- Thin Provision**

Configure per disk

VM Storage Policy:

Keep existing VM storage policies

Name	Capacity	Provisioned	Free	T
Storage Compatibility: Compatible				
datastore26-1	7.26 TB	370.6 GB	7.09 TB	

Compatibility

✓ Compatibility checks succeeded.

CANCEL

BACK

NEXT

Select storage

Select the datastore in which to store the configuration and disk files

Select virtual disk format:

- Thin Provision
- Same format as source
- Thick Provision Lazy Zeroed
- Thick Provision Eager Zeroed
- Thin Provision

VM storage policy:

The following datastores are available. Select the destination datastore for the virtual machine configuration files.

Name	Capacity	Used	Free	Type	Cluster
hx-scale	128 TB	14.24 TB	121.89 TB	NFS v3	
SpringpathDS-WZP223202DA	216 GB	9.89 GB	206.11 GB	VMFS 5	

Advanced >>

Compatibility

Back


Next

Finish

Cancel


- vi. Select a destination network as displayed in the following screenshot.

- vii. Enter the information identified below in the Customize vApp Properties page displayed in the following screenshot.

 Do not customize your setup credentials at this point or any other point during the installation. You can do so after you complete the installation process.

#### Customize vApp properties


Edit the vApp properties

 All properties have valid values [Show next...](#) [Collapse all...](#)

Property	Description
Uncategorized	6 settings
Encoded user-data	In order to fit into a XML attribute, this value is base64 encoded. It will be decoded, and then prc
SSH public keys	This field is optional, but indicates that the instance should populate the default user "authorized
Default user's password	If set, the default user password will be set to this value to allow password based login. The pas
A unique ID for this VM instance	Specifies the instance ID. This is required and is used to determine if the machine should take "l
Hostname	Specifies the hostname of the VM instance.
URL to seed instance data from	This field is optional, but indicates that the instance should "seed" user-data and meta-data from

[Back](#) [Next](#) [Finish](#) [Cancel](#)

1. The public SSH key.
2. The default user's password to SSH from the vSphere console.
3. The unique ID and hostname – ensure that these credentials are unique to avoid duplication issues.

 Use lowercase characters when providing the installer hostname in the Customize vApp Properties page.

- viii. Customize the template as required for your environment and review the completed information as displayed in the following screenshot.

**Ready to complete**  
Review configuration data.

Name	suite-v0.0.985demo
Source VM name	suite-v0.0.985
Download size	4.9 GB
Size on disk	32 GB
Folder	puj
Resource	10.13.72.163
Storage mapping	1
Network mapping	1
IP allocation settings	IPv4, Static - Manual
Properties	A unique ID for this VM instance = default-instance-id Default user's password = Encoded user-data = Hostname = default-hostname SSH public keys = URL to seed instance data from =

Buttons: Back, Next, Finish, Cancel

- ix. Click **Finish** to start deploying the VM from the template inside the target folder.

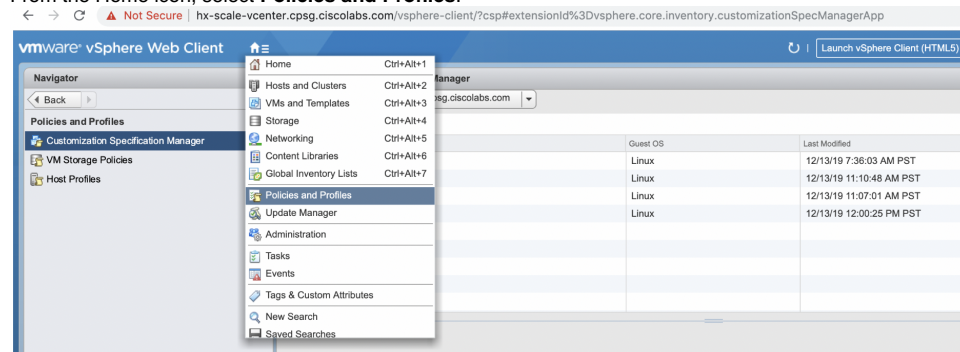
- b. If DHCP *is not* installed, follow these steps

Follow these steps **ONLY** if DHCP *is not* installed – use your static IP as the VMware customization specification is needed to attach the IP to the installer VM.

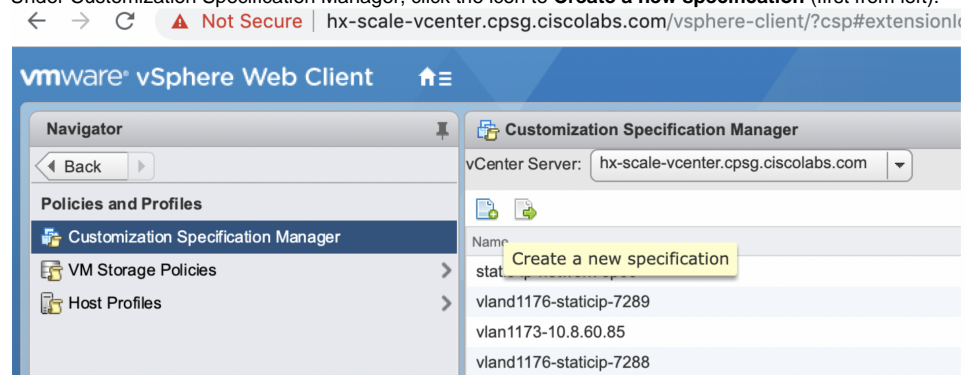
The details attached in the Customization *Specification* (term specific to vSphere), like the IP, DNS, Gateway, and so forth are assigned to the VM, when it is powered on.

IPs cannot be attached to the VM when it is Powered ON automatically and you must follow the instructions provided below to create an installation VM using the *Customization Specification* (specific to vSphere) which is used to create a template or custom profile with IP details, when attached to the VM.

- i. Login to vSphere.  
ii. From the Home icon, select **Policies and Profiles**.



- iii. Under Customization Specification Manager, click the icon to **Create a new specification** (first from left).



- iv. For the *Target VM OS*, select **Linux**.  
v. Set the *Computer Name* to any suitable name.

- vi. Enter **cpsg.ciscolabs.com** as the *Domain name*.

The screenshot shows the 'New VM Guest Customization Spec' wizard at the 'Set Computer Name' step. The left sidebar shows steps 1 through 6, with 'Set Computer Name' selected. The main area is titled 'Computer Name' and contains the following options:

- Enter a name: 
  - The name cannot exceed 63 characters.
  - Append a numeric value to ensure uniqueness. The name will be truncated if combined with the numeric value, it exceed 63 characters.
- Use the virtual machine name. If the name exceeds 63 characters, it will be truncated.
- Enter a name in the Clone/Deploy wizard.
  - Generate a name using the custom application configured with the vCenter Server.
    - Argument:

The 'Domain name' field at the bottom is filled with 'cpsg.ciscolabs.com'. Buttons for 'Back', 'Next', 'Finish', and 'Cancel' are at the bottom right.

- vii. To configure the network, select the button to **Manually select custom settings** for to ensure Static IP allocation so that you can manually enter the Static IP details.



Select the option to **use standard network....** if you are using a *DHCP* setup.

The screenshot shows the 'New VM Guest Customization Spec' wizard at the 'Configure Network' step. The left sidebar shows steps 1 through 6, with 'Configure Network' selected. The main area is titled 'Configure Network' and contains the following options:

- Use standard network settings for the guest operating system, including enabling DHCP on all network interfaces.
- Manually select custom settings

Below the options is a table with columns for 'Description', 'IPv4 Address', and 'IPv6 Address':

Description	IPv4 Address	IPv6 Address
NIC1	Use DHCP	Not used

Buttons for 'Back', 'Next', 'Finish', and 'Cancel' are at the bottom right.

- viii. Enter other details in subsequent screens to complete the wizard requirements.  
 ix. Wait for the installer VM to start – when it does, the Static IP assigned by the custom specification will be assigned to the VM.



Currently, an existing VMware issue does not save the check box setting. To workaround this issue, click the **Edit** settings on the VM, and check it again, and save your changes to assign the static IP.

- x. Click the **Edit** Button.

**NIC1 - Edit Network**

**IPv4**  
Specify IPv4 settings for the virtual network adapter.

**IPv6**

Use DHCP to obtain an IP address automatically.  
 Prompt the user for an address when the specification is used  
 Use an application configured on the vCenter Server to generate an IP address  
 Argument:

Use the following IP settings:  
 IP Address:   
 Subnet Mask:   
 Default Gateway:   
 Alternate Gateway:

OK Cancel

- xi. Click **OK** and then click the **Enter DNS and Domain Settings**.  
 xii. In the DNS search path enter **cpsg.ciscolabs.com** and click **OK**.

**New VM Guest Customization Spec**

**Enter DNS and Domain Settings**  
Enter the DNS and domain information for this new virtual machine.

Primary DNS:   
 Secondary DNS:   
 Tertiary DNS:

DNS Search Path

Add Delete Move Up Move Down

Back Next Finish Cancel

- xiii. Click **Next** and then **Finish**.

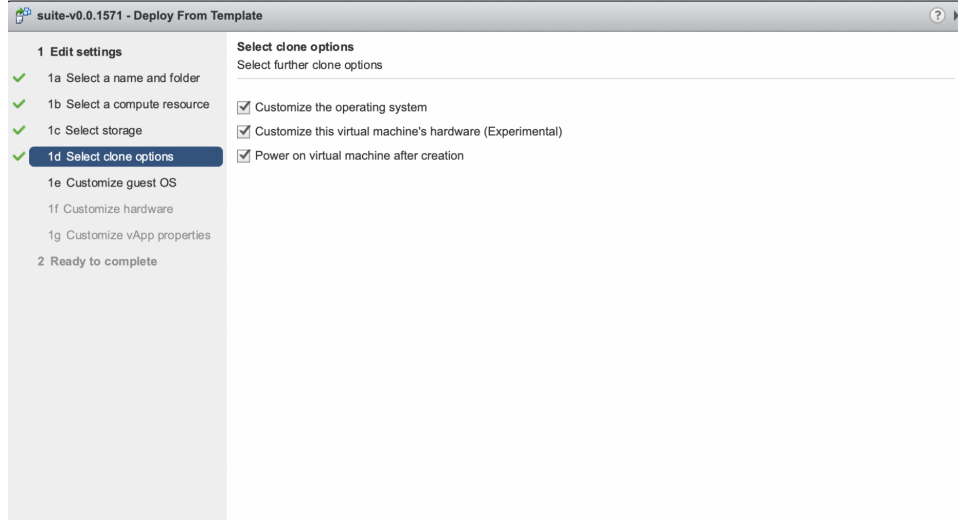
VMware vSphere Web Client

Customization Specification Manager

vCenter Server: hx-scale-vcenter.cpsg.ciscolabs.com

Name	Guest OS	Last Modified
static-ip-network-spec	Linux	12/13/19 7:36:03 AM PST
vlan1176-staticip-7289	Linux	12/13/19 11:10:48 AM PST
vlan1173-10.8.60.85	Linux	12/13/19 11:07:01 AM PST
vlan1176-staticip-7288	Linux	12/13/19 12:00:25 PM PST
Custom-Spec-test1571-1	Linux	12/18/19 10:35:48 AM PST

- xiv. **Create a New Installer VM using this customization spec.** Start creating the VM installer from the installer template, in the wizard section Select the **Clone** option, make sure to check the **Customize the Operating System** box so that you can select the custom specification in the next screen.



1 Edit settings

- ✓ 1a Select a name and folder
- ✓ 1b Select a compute resource
- ✓ 1c Select storage
- ✓ 1d Select clone options
- 1e Customize guest OS
- 1f Customize hardware
- 1g Customize vApp properties

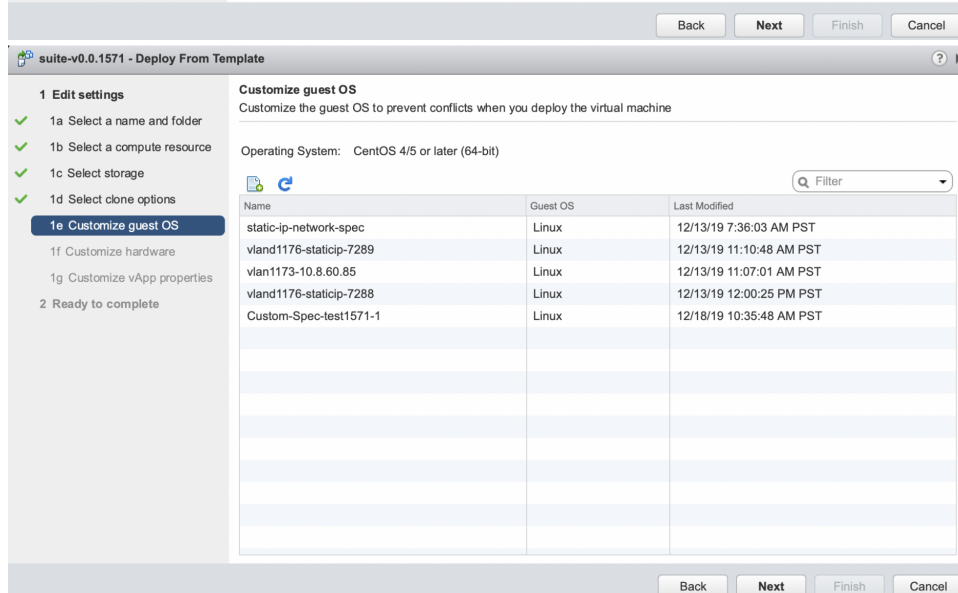
2 Ready to complete

Select clone options  
Select further clone options

- Customize the operating system
- Customize this virtual machine's hardware (Experimental)
- Power on virtual machine after creation

Back Next Finish Cancel

---



1 Edit settings

- ✓ 1a Select a name and folder
- ✓ 1b Select a compute resource
- ✓ 1c Select storage
- ✓ 1d Select clone options
- 1e Customize guest OS
- 1f Customize hardware
- 1g Customize vApp properties

2 Ready to complete

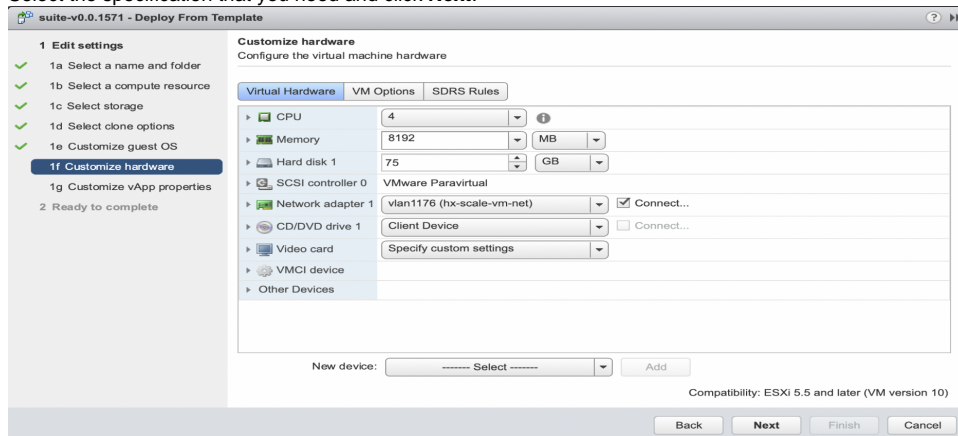
Customize guest OS  
Customize the guest OS to prevent conflicts when you deploy the virtual machine

Operating System: CentOS 4/5 or later (64-bit)

Name	Guest OS	Last Modified
static-ip-network-spec	Linux	12/13/19 7:36:03 AM PST
vlan1176-staticip-7289	Linux	12/13/19 11:10:48 AM PST
vlan1173-10.8.60.85	Linux	12/13/19 11:07:01 AM PST
vlan1176-staticip-7288	Linux	12/13/19 12:00:25 PM PST
Custom-Spec-test1571-1	Linux	12/18/19 10:35:48 AM PST

Back Next Finish Cancel

- xv. **Select the specification that you need and click Next.**



1 Edit settings

- ✓ 1a Select a name and folder
- ✓ 1b Select a compute resource
- ✓ 1c Select storage
- ✓ 1d Select clone options
- ✓ 1e Customize guest OS
- 1f Customize hardware
- 1g Customize vApp properties

2 Ready to complete

Customize hardware  
Configure the virtual machine hardware

Virtual Hardware VM Options SDRS Rules

- CPU: 4
- Memory: 8192 MB
- Hard disk 1: 75 GB
- SCSI controller 0: VMware Paravirtual
- Network adapter 1: vlan1176 (hx-scale-vm-net)  Connect...
- CD/DVD drive 1: Client Device  Connect...
- Video card: Specify custom settings
- VMCI device
- Other Devices


New device: ----- Select ----- Add

Compatibility: ESXi 5.5 and later (VM version 10)


Back Next Finish Cancel

- xvi. Enter other details in subsequent screens, to complete the wizard. Wait for the installer VM to start, the Static IP assigned by the custom specification will be assigned to the VM.
- xvii. Wait for the installer VM to start – when it does, the Static IP assigned by the custom specification will be assigned to the VM.


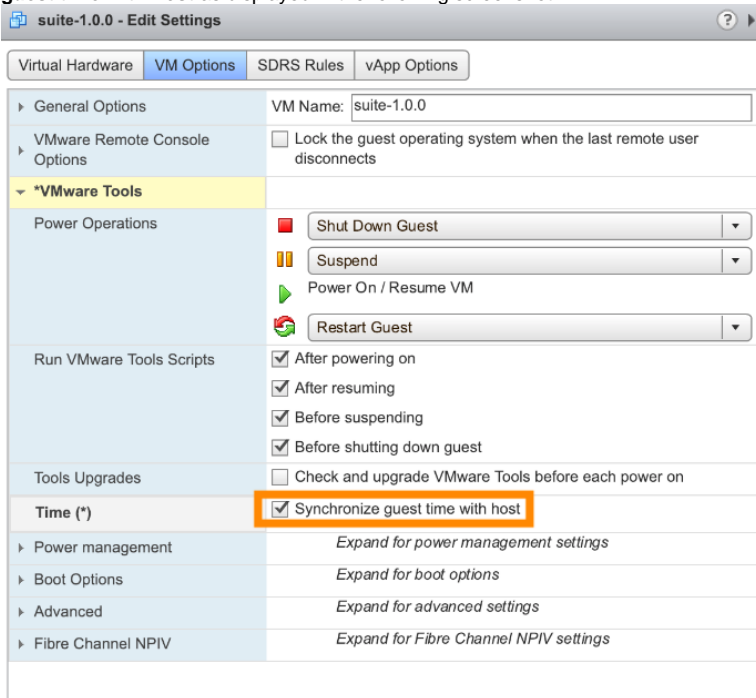


 Currently, an existing VMware issue does not save the check box setting. To workaround this issue, click the **Edit** settings on the VM, and check it again, and save your changes to assign the static IP.

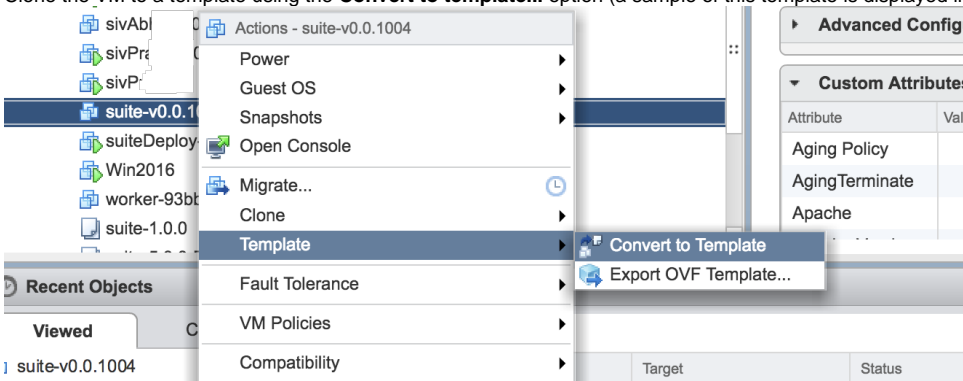
- Wait for some time so the VM is cloned and created, then refresh the VM page to view the powered off VM – The OVA is imported as a VM (powered off) on vSphere.

 When you import the OVA as a VM, ensure that it is powered **off** on vSphere.


- Right-click to edit the VM Settings for the powered off VM. Click the *VM Options* tab. Under *VMware Tools*, select the checkbox to **Synchronize guest time with host** as displayed in the following screenshot.




- Clone the VM to a template using the **Convert to template...** option (a sample of this template is displayed in the following screenshot).



- Once the VM is converted to template, it should appear as identified by the orange arrow in the following screenshot.



9. Right click this template name and select the **New VM from This Template** option as displayed in the following screenshot – this template will also be used as the value for the *vSphere Template Name* cloud setting, when you provide the details to install the Suite Admin.



10. Edit the **1e Customize vApp properties** to ensure that the VM has unique values for **A unique ID for this VM instance**, **Hostname**, **Default user's password**, and **SSH public keys** for this VM instance.



For the password and/or the public key to take effect when deploying the VMware OVA for the CloudCenter Suite installer, you **must** change the *default-instance-id* to something else than *default-instance-id* or *the hostname*!

#### suite-5.2.2.ova - Deploy From Template

- ✓ 1 Select a name and folder
- ✓ 2 Select a compute resource
- ✓ 3 Select storage
- ✓ 4 Select clone options
- ✓ 5 Customize hardware
- 6 Customize vApp properties
- 7 Ready to complete

#### Customize vApp properties

Edit the vApp properties

Uncategorized	6 settings
Encoded user-data	
SSH public keys	Provide-SSH-Pub-key
Default user's password	Provide-SecurePassword
A unique ID for this VM instance	Provide-Unique-ID-Here
Hostname	Provide-Hostname-Here
URL to seed instance data from	

11. After the VM is created from the template, power it on.  
 12. Use this IP address to access the CloudCenter Suite UI (displayed in the following screenshot), go to the newly created VM's IP using HTTPS protocol in a supported browser (see [Browser Compatibility](#)).

You have now setup the installer for a VMware cloud.

# Prepare Infrastructure


## Prepare Infrastructure

- [General Compatibility](#)
- [Resource Requirements for CloudCenter Suite Modules](#)
- [Number of VMs](#)
- [IP Pool Requirements](#)
- [NTP Requirements](#)
- [The Suite Installer Dashboard](#)
- [Without Internet Access](#)

See [Browser Compatibility](#) for additional details.

CloudCenter Suite supports Kubernetes 1.15.4 and earlier releases.

The CloudCenter Suite requires Tiller v2.12.3 to be installed. Refer to the [Helm documentation](#) for additional details.

 Installers are already incorporated in the CloudCenter Suite SaaS offer, see [SaaS Access](#) for additional details.

The following table lists the minimum resource requirements assuming that you install all available modules.

Module <sup>1,2</sup>	Public Cloud <sup>5</sup>			Private Cloud <sup>3</sup>		
	vCPU	Memory (GB)	Storage (GB)	vCPU	Memory (GB)	Storage (GB)
Suite Admin	16	37	300	16	37	300
Workload Manager <sup>4</sup> and Cost Optimizer	15	68	230 <sup>6</sup>	15	68	230 <sup>6</sup>
Action Orchestrator	5	6	60	5	6	60
Kubernetes Cluster (3 primary servers)	na	na	na	9	24	120
<b>Total</b>	<b>36</b>	<b>111</b>	<b>590</b>	<b>45</b>	<b>135</b>	<b>590</b>

<sup>1</sup> Update only one module at a time. If you simultaneously update more than one module, your update process may fail due to limited resource availability.

<sup>2</sup> Before updating any module, verify that you have un-allocated CPU/Memory in your cluster to ensure that your environment has free CPU/Memory – a module-update scenario requires additional resources for the old pod to continue running until the new pod initializes and takes over. This additional resource requirement is temporary and only required while a module update is in Progress. After the module is updated, the additional resources are no longer needed.

<sup>3</sup> On private clouds (vSphere and OpenStack), each of the 3 primary server instances require 3 vCPU and 8 GB memory and 40 GB storage (root disk), hence the difference in the additional requirement of 9 vCPU, 24 GB memory, and 120 GB storage (root disk). See the Number of VMs section below for additional details. Similarly, each worker instances require 3 vCPU and 8 GB memory and 40 GB storage (root disk) – however, the number of workers changes dynamically at install time. Installer VMs require a minimum of 4 vCPUs and 8 GB RAM.

<sup>4</sup> Workload Manager numbers include considerations for 4 Cloud Regions in the same instance. To support additional cloud regions, you must scale your cluster by adding Kubernetes worker nodes. You will need 1 CPU and 3 GB memory for each additional region.


<sup>5</sup> Public clouds do not support auto-scaling – the number of nodes might differ if scaled on an auto-scaling enabled node group.

<sup>6</sup> The storage is 230 GB just to enable [StatefulSet](#) migration. In reality, only 115 GB is being used for operation of services.

A CloudCenter Suite installation launches a highly available Kubernetes cluster which consists of primary server(s) and worker(s) instances.

 The number of worker nodes (for both private and public cloud) vary based on the instance type selected during the installation process.

For private clouds, a redundant cluster requires a minimum of 2 out of 3 primary server nodes to be running at any point, so the cluster can function as designed.

 If you plan to scale up at a later date, be aware that the worker instance type selected at installation time will also be used for the scaled nodes.

The CloudCenter Suite requires that the underlying disks for Kubernetes disk attachments be redundant and available. Most public clouds already provide built-in redundancy for their block disks (AWS EBS, GCP Persistent Disks, and so forth). Be sure to verify that the Datastores/Datastore Clusters are also on redundant, non-local storage (NFS, NetApp) before you begin the installation process.

You must select IP address to ensure that each IP endpoints is available, accessible, and not used by any other resource.

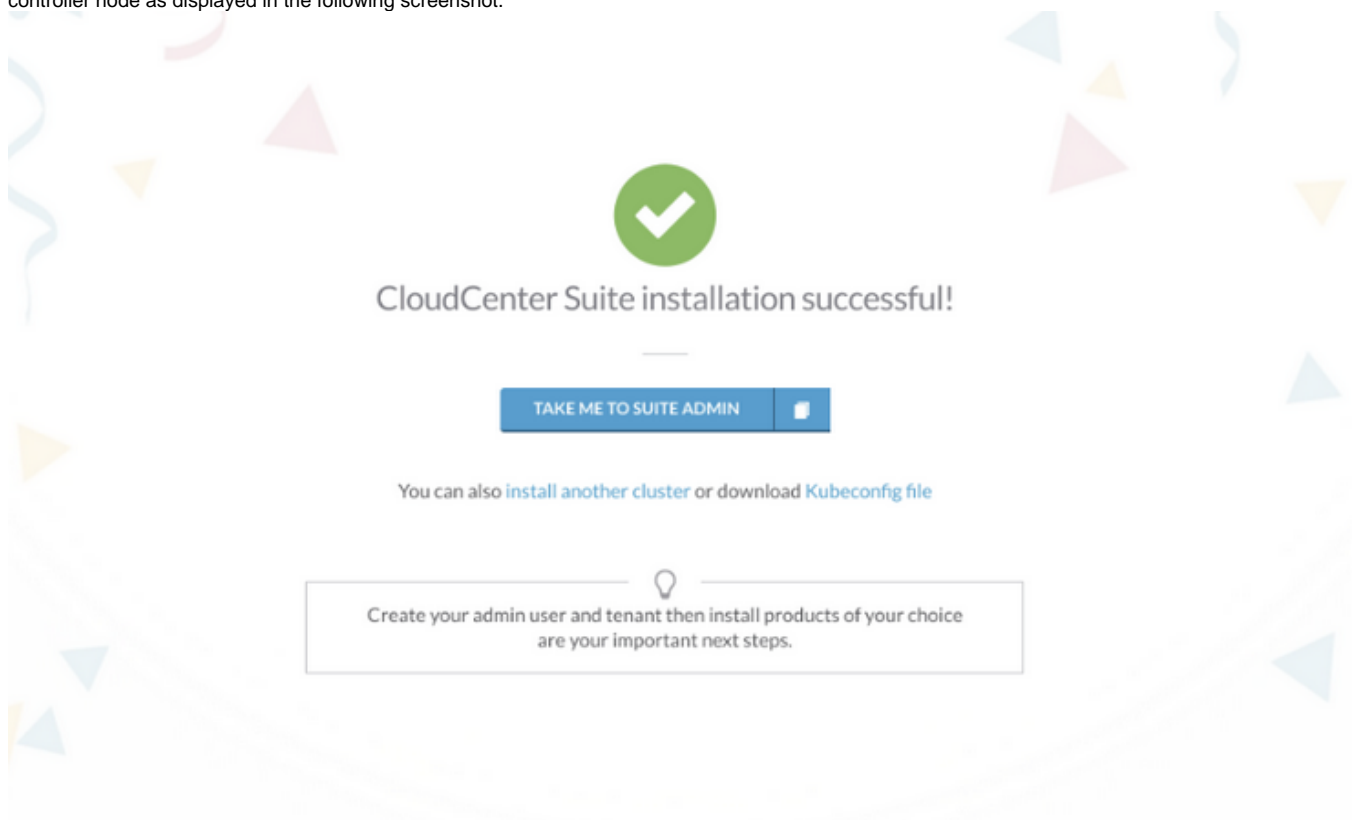
When configuring or modifying you pool of IP addresses, be aware of the following requirements:

- Verify if the IP pool can accommodate additional workloads.
- Select your instance type according to the following dependencies – based on your instance type selection, the installer displays the error or success information in the UI.
  - The CloudCenter Suite setup requires 3 primary servers.
  - The CloudCenter Suite dynamically calculates the number of application VMs (workers).
- Do not use 172.18.0.1/16 for the installer instance as this IP address is used by the Docker/Kubernetes setup.
- 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 details.

You must either set the Network Time Protocol (NTP) time at the datacenter level or at the time of installation.

If set at installation time, then verify that the network can access the NTP server.

The time for all worker and primary server nodes is synced with the *primary* controller node. The *primary* controller node is the instance used to launch the CloudCenter Suite – identified by the link that takes you to the Suite Admin UI (Take Me to Suite Admin). This link contains the IP address of the primary controller node as displayed in the following screenshot.



After launching the installer, navigate to the IP address of your VM in a supported browser. This presents the Suite Installer Dashboard. The Suite Installer Dashboard has the following options:

- [New Cluster Installation](#)
- [Existing Cluster Installation](#)
- [Upgrade Kubernetes Cluster](#)

The Cisco Repository is used to host Cisco-related files and packages for various purposes. You may need to install the CloudCenter Suite in an environment that does not have internet access. If so, you need to set up the offline repository. See [Offline Repository](#) to sync your offline repository with the Cisco repository.

As you will be shutting down the installer VM after the installation, you can reuse that VM to set up the offline repository.

# New Cluster Installation

## Install the CloudCenter Suite on a New Kubernetes Cluster

Once you access the Suite Installer Dashboard (see [Prepare Infrastructure](#)), you can install a new cluster and launch nodes for the new Kubernetes cluster

- [Amazon EKS Installation](#)
- [Azure AKS Installation](#)
- [Google GKE Installation](#)
- [OpenStack Installation](#)
- [VMware vSphere Installation](#)

# Amazon EKS Installation

## Amazon EKS Installation

- [Amazon Nuances](#)
- [Module Details](#)
- [Minimum Permissions Needed](#)
- [Installation Process](#)

Be aware of the following requirements when installing the CloudCenter Suite:

- **Maximum Supported Version:** EKS Version 1.13.7 and below.
- **Unavailable Resources:** The following resources will not be available until the upgrade completes:
  - EKS cluster
  - Suite admin cluster
- **Resources:** Amazon creates the following resources for the AWS account:
  - An EKS Cluster with user-provided specifications.
  - All resources remain in the same region as the cluster.
  - A new CloudFormation stack with the same number of instances, security groups, subnets, and roles that are used to connect to the cluster.
    - VPC Name: *cluster\_name*-VPC
    - Role Name for VPC: *cluster\_name*-Role
    - Role Name for Workers: *cluster\_name*-NodeInstanceRole
    - New CFN stack Name: *cluster\_name*-New-Workers-*random\_UUID32*
    - Auto Scaling Group for worker nodes as part of cloud formation workers stack
- **The Delete API:**



You cannot trigger a Delete call by deleting the Amazon cluster from either the AWS console or the AWS CLI. Instead, use the Delete API.

Additionally, refer to your module documentation for module-specific dependencies as specified in the following table.

Module	Documentation
Workload Manager	<a href="#">Cloud Overview</a>
Action Orchestrator	<a href="#">Add Cloud Account</a>
Cost Optimizer	<a href="#">Cloud Overview</a>

The following IAM policies are required for the CloudCenter Suite to access the EKS and create a new cluster on AWS.

- AmazonEC2FullAccess
- IAMFullAccess
- AutoScalingFullAccess
- AmazonEKSClusterPolicy
- AmazonEKSWorkerNodePolicy
- AmazonVPCFullAccess
- AmazonEKSServicePolicy
- AmazonEKS\_CNI\_Policy
- AmazonRoute53FullAccess
- Inline\_Policy\_EKS\_Cluster = an inline policy allowing the following actions on the EKS service to an IAM user:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "cloudformation:CreateStack",
        "cloudformation:DescribeStacks",
        "cloudformation:DescribeStackEvents",
        "cloudformation:DescribeStackResources",
        "cloudformation:DescribeStackResource",
        "cloudformation:GetTemplate",
        "cloudformation:ValidateTemplate",
        "cloudformation>DeleteStack",
        "eks:UpdateClusterVersion",
        "cloudformation:UpdateStack",
        "eks:ListUpdates",
        "eks:DescribeUpdate",
        "eks:DescribeCluster",
        "eks:ListClusters",
        "eks:CreateCluster",
        "eks>DeleteCluster"
      ],
      "Resource": "*"
    }
  ]
}

```

To install the CloudCenter Suite on a new Amazon cluster, perform the following procedure.

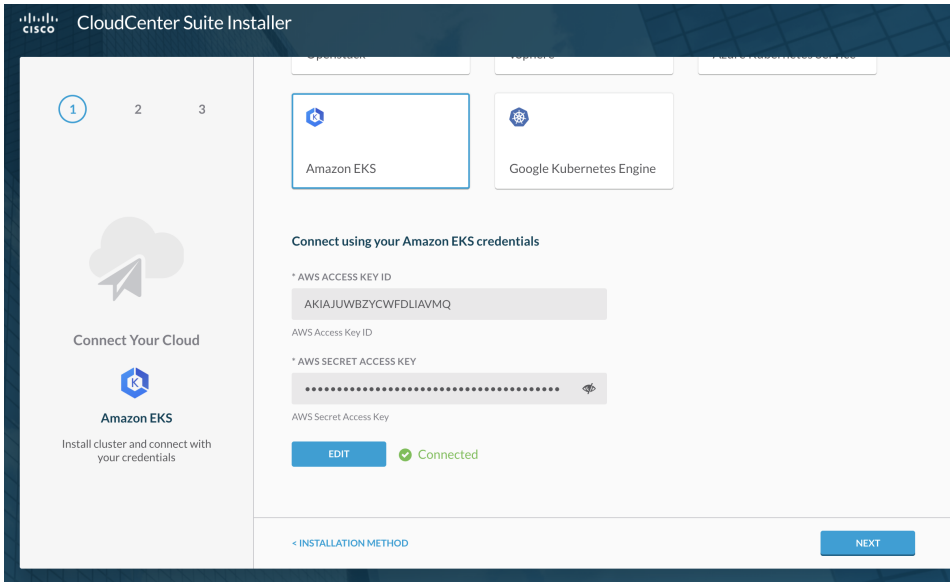
1. Verify that you have prepared your environment as listed in the *Amazon Nuances* section above.
2. Navigate to the Suite Installer Dashboard.
3. Click **New Cluster**.
4. Select **Amazon EKS**.
5. To connect using Amazon cloud credentials, enter the EKS details specified in the following table.

EKS Details	Description
<b>AWS Access Key ID</b>	AWS access key ID for the account
<b>EKS Secret Access Key</b>	AWS secret access key

6. Click **Connect** as displayed in the following screenshot.

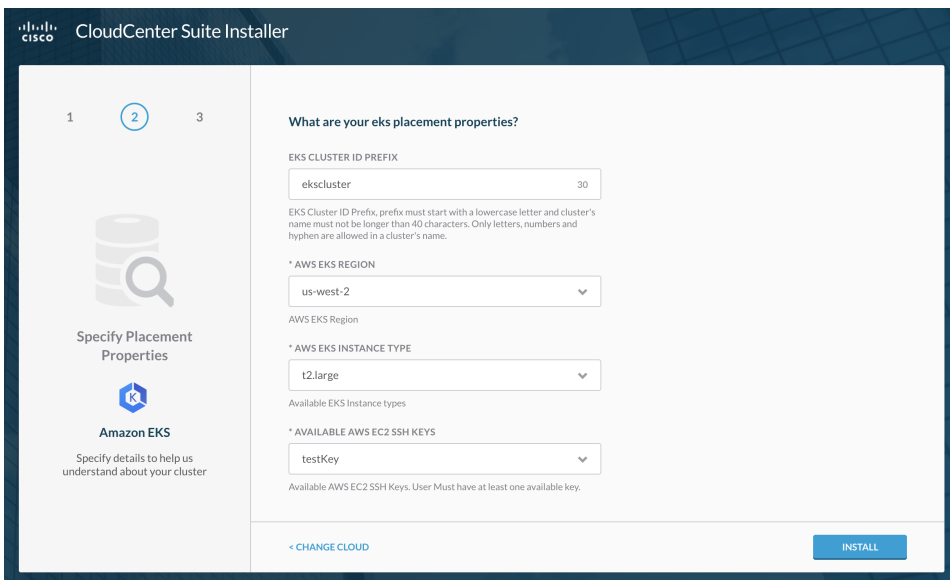
The screenshot shows the 'CloudCenter Suite Installer' interface. On the left, there is a progress indicator with steps 1, 2, and 3. Step 1 is active, showing a 'Connect Your Cloud' section with an Amazon EKS icon and the text 'Install cluster and connect with your credentials'. On the right, there are three installation method options: 'Openstack', 'vSphere', and 'Azure Kubernetes Service'. Below these, 'Amazon EKS' and 'Google Kubernetes Engine' are also visible. The 'Amazon EKS' option is selected. Underneath, there is a section titled 'Connect using your Amazon EKS credentials' with two input fields: 'AWS ACCESS KEY ID' (containing 'AKIAJ54EFW3BNGDLYWA') and 'AWS SECRET ACCESS KEY' (containing 'hAPs1vl8vAvXfA6EunXyZgukUezdIX+QeJ7EK13'). A 'CONNECT' button is located below the input fields. At the bottom, there are navigation buttons: '< INSTALLATION METHOD' and 'NEXT'.

7. Once the connection is validated, click **Next** as displayed in the following screenshot.



8. To specify the cloud properties, enter the EKS details listed in the following table and displayed in the following screenshot.

EKS Details	Description
<b>EKS Cluster ID Prefix</b>	EKS Cluster ID Prefix, the prefix must start with a lowercase letter and cluster's name must not be longer than 40 characters. Only letters, numbers and hyphen are allowed in a cluster's name.
<b>AWS EKS Region</b>	Select region to launch the cluster.
<b>EKS Instance Type</b>	Select the type of instance of worker nodes.
<b>Available EC2 SSH Keys</b>	Select the SSH key, account must have at least one key.



9. Click **Install**. The installation progress is visible on screen.





If the Suite Admin is installed in EKS, the you cannot use the config file immediately after downloading it from the Suite installer success page. To access the Kubernetes cluster, access your command window to install AWS-IAM-AUTHENTICATOR and execute the following commands:

```
brew install kubernetes-cli
curl -Lo aws-iam-authenticator https://github.com/kubernetes-sigs/aws-iam-authenticator/releases/download/v0.3.0/heptio-authenticator-aws_0.3.0_darwin_amd64
chmod +x aws-iam-authenticator
sudo mv aws-iam-authenticator /usr/local/bin
```

10. Once successful, you see the following message.

```
CloudCenter Suite installation successful!
```

11. You have the following options at this point:

- a. Click **Take Me To Suite Admin** to launch and set up the [Suite Admin](#).
- b. Click **Install Another Cluster** to start another installation and go back to the homepage (Installer Dashboard).
- c. Download **Kubeconfig file** to connect to the launched cluster using the [kubectl](#) tool.

12. Be sure to switch off the installer VM. You can reuse this VM for any other purpose, for example, as an Offline Repository.

# Azure AKS Installation

## Azure AKS Installation

- [Azure Nuances](#)
- [Module Details](#)
- [Installation Process](#)

Be aware of the following requirements to install CloudCenter Suite:

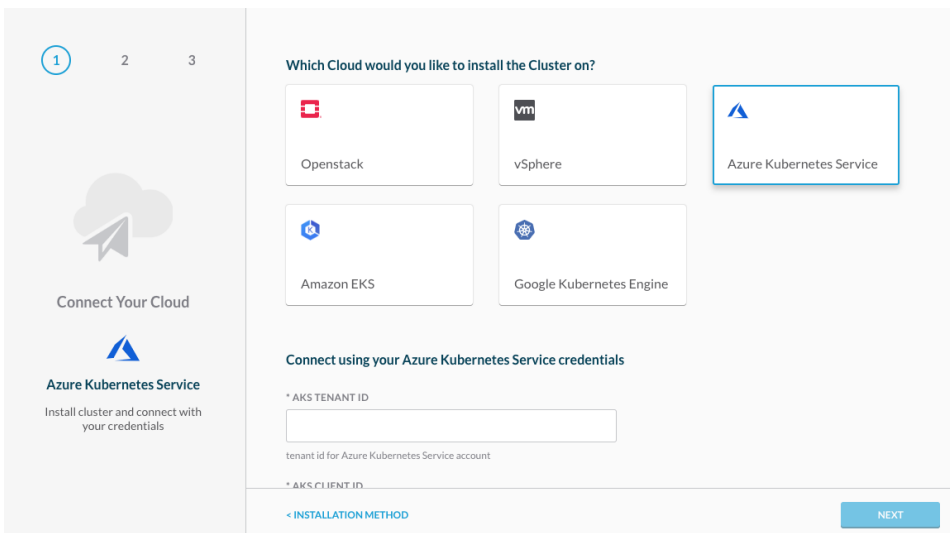
- **Maximum Supported Version:** AKS Version 1.12 and below.
- **Valid Azure Account:** A valid service account that allows you to use sufficient resource quota. See <https://docs.microsoft.com/en-us/azure/aks/container-service-quotas> for additional details.
- **Resource Group:** Create the resource group in a cloud region that supports Azure.

Additionally, refer to your module documentation for module-specific dependencies as displayed in the following table.

Module	Documentation
Workload Manager	<a href="#">Cloud Overview</a>
Action Orchestrator	<a href="#">Add Cloud Account</a>
Cost Optimizer	<a href="#">Cloud Overview</a>

To install the CloudCenter Suite on a new Azure AKS cluster, perform the following procedure.

1. Verify that you have prepared your environment as listed in the *Azure Nuances* section above.
2. Navigate to the Suite Installer Dashboard.
3. Click **New Cluster**.
4. Select **Azure Kubernetes Service** as displayed in the following screenshot.



5. To connect using Azure Kubernetes Service cloud credentials, enter the details identified in the following table and displayed in the following screenshot, and click **Connect**.

AKS Details	Description
AKS TENANT ID	The AKS account tenant ID.
AKS CLIENT ID	The AKS account client ID.
AKS CLIENT SECRET	The AKS account client secret.
AKS SUBSCRIPTION ID	The AKS subscription ID.



Refer to <https://docs.microsoft.com/en-us/azure/aks/kubernetes-service-principal> to learn about how to setup service principles with Azure Kubernetes Service (AKS), and use the credentials to populate the above fields.

6. Once the connection is validated, click **Next** as displayed in the following screenshot.

7. To specify the placement properties, enter the details identified in the following table and displayed in the following screenshot.

AKS Placement Property	Description
Resource Group	The AKS resource group to launch the cluster.
VM Size	The VM size of the cluster node.
AKS Cluster ID Prefix	<ul style="list-style-type: none"> <li>The prefix must begin with a lowercase letter.</li> <li>The entire name that you enter for this cluster must not be longer than 12 characters.</li> <li>Only letters, numbers and hyphens are allowed in this field.</li> </ul>

1 2 3

Specify Placement Properties

Azure Kubernetes Service

Specify details to help us understand about your cluster

**What are your aks placement properties?**

\* RESOURCE GROUP  
 Installer-West-3

List of supported resource groups in AKS

\* VM SIZE  
 Standard\_DS4\_v2

List of supported VM sizes in AKS

AKS CLUSTER ID PREFIX  
 mycluster01

AKS cluster ID prefix, prefix must start with a lowercase letter and cluster's name must not be longer than 12 characters. Only letters, numbers are allowed in a cluster's name.

< CHANGE CLOUD INSTALL

8. Click **Install**. The installation progress is visible on screen.
9. Once successful, you see the following message:

```
CloudCenter Suite installation successful!
```

10. You have the following options at this point:
  - a. Click **Take Me To Suite Admin** to launch and set up the [Suite Admin](#).
  - b. Click **Install Another Cluster** to start another installation and go back to the homepage (Installer Dashboard).
  - c. Download **Kubeconfig file** to connect to the launched cluster using the [kubectl](#) tool.
11. Be sure to switch off the installer VM. You can reuse this VM for any other purpose, for example, as an Offline Repository.

# Google GKE Installation

## Google GKE Installation

- [Google Nuances](#)
- [Module Details](#)
- [Installation Process](#)

Be aware of the following requirements when installing the CloudCenter Suite:

- **Maximum Supported Version:** GKE Version 1.12 and below.
- **Permissions:** Verify that the person upgrading the cluster has the following minimum permissions (roles) as displayed in the screenshot:  
A service account represents a Google Cloud service identity, such as code running on Compute Engine VMS, App Engine apps, or systems running outside Google.

Service account name  
**ad**

Display name for this service account

Service account ID  
**ad-581** @wakanda-214819.iam.gserviceaccount.com

**Project role**

Role  
**Service Account User**  
Create VMs and other GCP tasks with a service account. Users cannot impersonate the account directly as they can with Service Account Actor role.

Role  
**Kubernetes Engine Admin**  
Full management of Kubernetes Clusters and their Kubernetes API objects.

Role  
**Compute Admin**  
Full control of all Compute Engine resources.

[+ ADD ANOTHER ROLE](#)

**Furnish a new private key**  
Downloads a file that contains the private key. Store the file securely because this key can't be recovered if lost.

**Enable G Suite Domain-wide Delegation**  
Allows this service account to be authorized to access all users' data on a G Suite domain without manual authorization on their parts. [Learn more](#)


- Service Account User
- Kubernetes Engine Admin
- Compute Engine Admin

Additionally, refer to your module documentation for module-specific dependencies as identified in the following table:

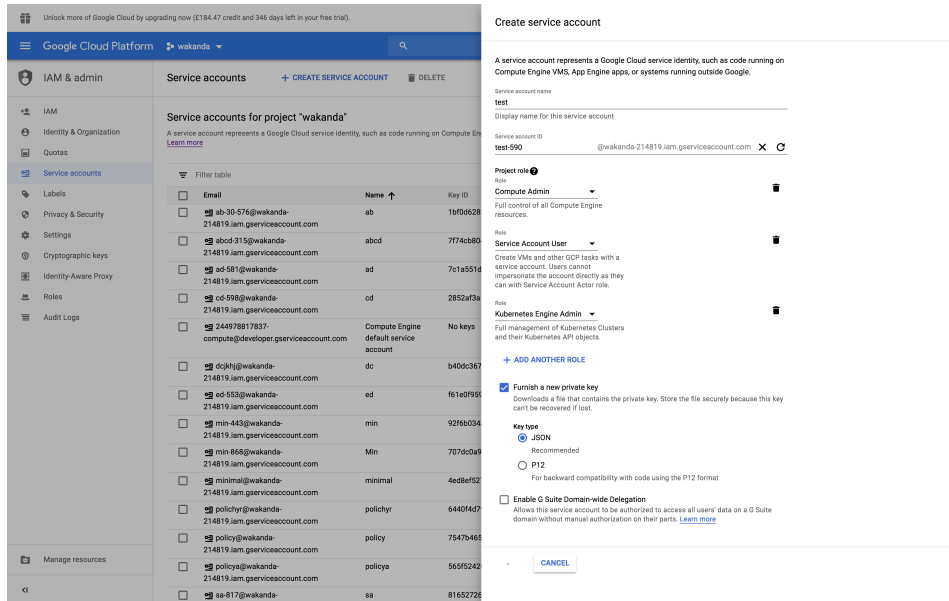
Module	Documentation
Workload Manager	<a href="#">Cloud Overview</a>
Action Orchestrator	<a href="#">Add Cloud Account</a>
Cost Optimizer	<a href="#">Cloud Overview</a>

To install the CloudCenter Suite on a new GKE Kubernetes cluster, perform the following procedure.

1. Verify that you have prepared your environment as listed in the *Google Nuances* section above.
2. Navigate to the Suite Installer Dashboard.
3. Click **New Cluster**.
4. Select the cloud of your choice (GKE in this case).
5. Generate a service account JSON file with the following minimum required permissions in the GKE console – be sure to check the "Furnish a new private key" checkbox for the JSON file to generate the key.

-  If you check the **Furnish a new private key** checkbox the resulting JSON file from the service account automatically contains a key when you download the file.

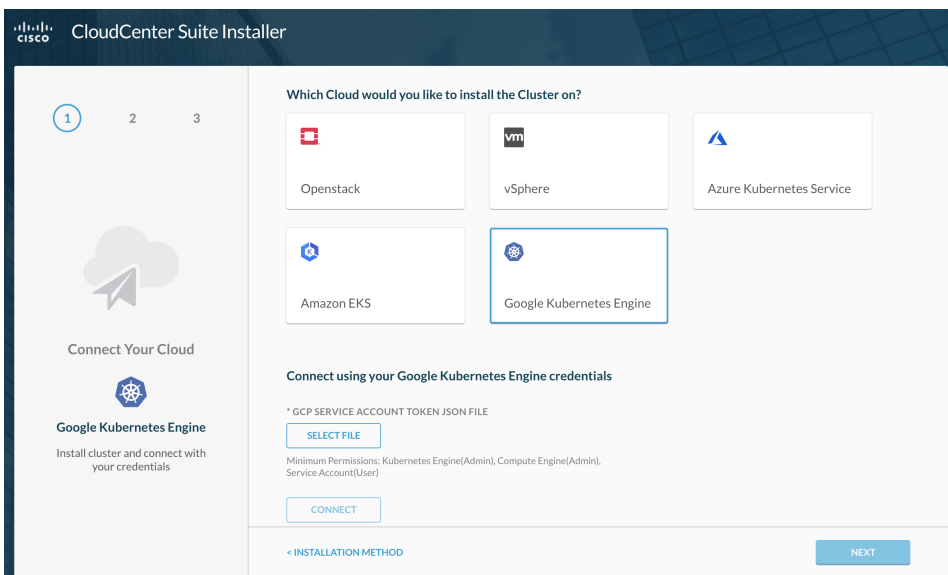
- Kubernetes Engine (Admin)
- Compute Engine (Admin)
- Service Account (User) as displayed in the following screenshot



The screenshot shows the Google Cloud Platform IAM & admin console. On the left, the 'IAM & admin' sidebar is visible. The main area displays 'Service accounts for project 'wakanda'' with a table of existing accounts. On the right, the 'Create service account' dialog is open, showing the 'test' service account with the role 'Service Account User' and the 'Furnish a new private key' checkbox checked.

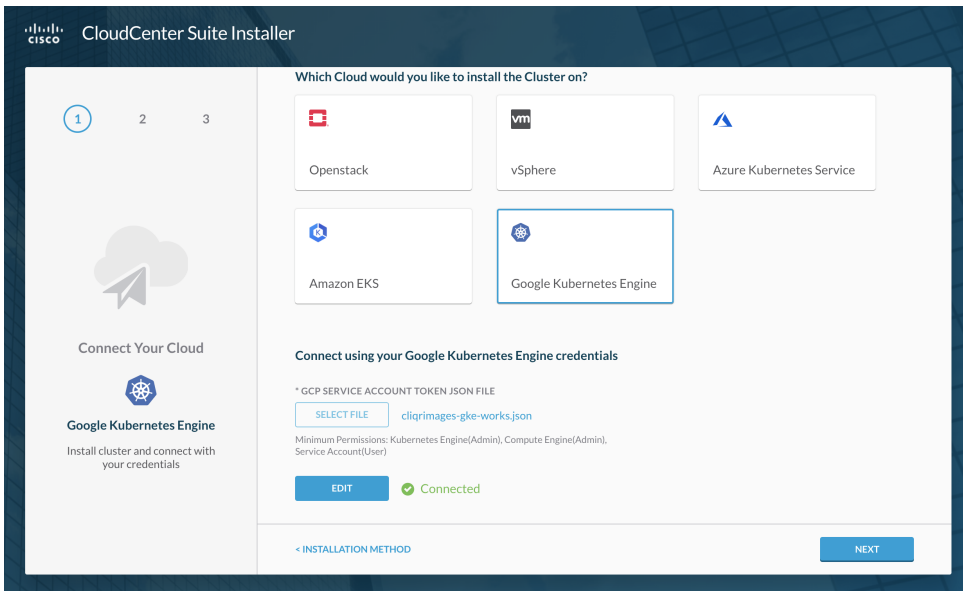
Email	Name	Key ID
ab-30-576@wakanda-214819.iam.gserviceaccount.com	ab	13f05628
abcd-315@wakanda-214819.iam.gserviceaccount.com	abcd	7774cb80
ad-581@wakanda-214819.iam.gserviceaccount.com	ad	7c1a5519
cd-59@wakanda-214819.iam.gserviceaccount.com	cd	2852af58
compute@developer.gserviceaccount.com	Compute Engine default service account	No keys
dc@wakanda-214819.iam.gserviceaccount.com	dc	b40dc367
ed-593@wakanda-214819.iam.gserviceaccount.com	ed	161e0f98
min-443@wakanda-214819.iam.gserviceaccount.com	min	92f6b034
min-866@wakanda-214819.iam.gserviceaccount.com	Min	70760049
minimal@wakanda-214819.iam.gserviceaccount.com	minimal	4ed8ef52
polichyr@wakanda-214819.iam.gserviceaccount.com	polichyr	64404d7
policy@wakanda-214819.iam.gserviceaccount.com	policy	7547b466
poliya@wakanda-214819.iam.gserviceaccount.com	poliya	565f9242
sa-817@wakanda-	sa	81652728

- To connect using Google cloud credentials, download the **Google service account token in JSON** format that you created in the previous step.
- Upload the JSON file mentioned in the previous step and click **Connect** to validate the credentials as displayed in the following screenshot.




The screenshot shows the Cisco CloudCenter Suite Installer interface. The main panel is titled 'Which Cloud would you like to install the Cluster on?' and displays several cloud provider options: Openstack, vSphere, Azure Kubernetes Service, Amazon EKS, and Google Kubernetes Engine. The 'Google Kubernetes Engine' option is selected and highlighted with a blue border. Below the options, there is a section for 'Connect using your Google Kubernetes Engine credentials' which includes a 'SELECT FILE' button for the 'GCP SERVICE ACCOUNT TOKEN JSON FILE' and a 'CONNECT' button. The minimum permissions required are listed as: Kubernetes Engine(Admin), Compute Engine(Admin), and Service Account(User). Navigation buttons for '< INSTALLATION METHOD' and 'NEXT' are visible at the bottom.

- Once the connection is validated, click **Next** as displayed in the following screenshot.




9. Enter the GKE details to specify the cloud properties as identified in the following table and as displayed in the following screenshot.

GKE Details	Description
<b>GKE Cluster ID Prefix</b>	<ul style="list-style-type: none"> <li>The prefix must begin with a lowercase letter.</li> <li>The entire name for this cluster must not be longer than 40 characters.</li> <li>Only letters, numbers and hyphens are allowed in this field.</li> </ul>
<b>GKE Zone</b>	The Google cloud zone to launch the cluster.
<b>GKE Instance Type</b>	Select the minimum resource requirements based on your environment setup.



10. Click **Install**. The installation progress is visible on screen as displayed in the following screenshot.



11. Once successful, you see the following message.

```
CloudCenter Suite installation successful!
```

12. You have the following options at this point:

- Click **Take Me To Suite Admin** to launch and set up the [Suite Admin](#).
- Click **Install Another Cluster** to start another installation and go back to the homepage (Installer Dashboard).
- Download **Kubeconfig file** to connect to the launched cluster using the [kubectl](#) tool.

13. Be sure to switch off the installer VM. You can reuse this VM for any other purpose, for example, as an Offline Repository.



# OpenStack Installation

## OpenStack Installation

- [OpenStack Nuances](#)
- [Module Details](#)
- [Installation Process](#)

Verify the following OpenStack nuances:

- OpenStack newton release with at least the following service versions:
  - Cinder v2
  - Keystone v3
  - OpenStack Nova v2
  - OpenStack Networking v2
  - OpenStack Glance v2
- Ensure to add Port 6443 to the default security group as the security group created for the cluster is not automatically assigned to the load balancer created for the cluster.
- The tenant and project requirements for OpenStack Cloud are identified in the following table.

Model	Quota	Description
For all cases	2 (primary server group, worker group)	Server Groups
	Number of workers + number of primary servers	Server Group Members
	3 (API load balancers)	Load Balancers
	6 (2 for each load balancer)	Health Monitors
	6 (2 for each load balancer)	Pools
	6 (2 for each load balancer)	Listeners
	3 (1 for the cluster VMs, 2 for the Kubernetes load balancer services)	Security Groups
	18	Security Group Rules
	See <a href="#">Prepare Infrastructure</a> for additional details	Volume GB
	Number of workers + number of primary servers + 3 for each load balancer	Ports
	Number of workers + number of primary servers	Instances
	16 GB (recommended for each worker and each primary server)	RAM
	32 (recommended for each workers and each primary server)	vCPUs
Tenant network	Floating IPs = 3	1 for each load balancer
	Networks = 1	For the tenant network
	Subnet = 1	For the tenant network
	Router = 1	For the tenant network to public network connection
Provider network	Number of workers + number of primary servers + 3 load balancers	Free IPs in the provider network

- **Network Time Protocol (NTP) must be configured – this is important as the CloudCenter Suite installation can fail, if NTP is not configured or if it is wrongly configured.**



If you setup CloudCenter Suite in offline mode, you must provide valid NTP server details before you save your configuration.

Additionally, refer to your module documentation for module-specific dependencies as identified in the following table:

Module	Documentation
Workload Manager	<a href="#">Cloud Overview</a>

Action Orchestrator	<a href="#">Add Cloud Account</a>
Cost Optimizer	<a href="#">Cloud Overview</a>

To install the CloudCenter Suite on a new OpenStack cluster, perform the following procedure.

1. Verify that you have prepared your environment as listed in the *OpenStack Nuances* section above.
2. Navigate to the Suite Installer Dashboard.
3. Click **New Cluster**.
4. Click the OpenStack card.
5. To connect using OpenStack cloud credentials, enter the OpenStack Placement Property details identified in the following table.

OpenStack Placement Properties	Description
<b>OpenStack Authentication URL</b>	The OpenStack authentication service URL.
<b>OpenStack Region</b>	The OpenStack cloud region.
<b>OpenStack Domain Name</b>	The OpenStack account domain name.
<b>OpenStack Project</b>	The OpenStack project name.
<b>OpenStack Username</b>	The OpenStack account username.
<b>OpenStack Password</b>	The OpenStack account password.
<b>OpenStack CA Certificate</b>	The CA certificate that is required to verify an OpenStack HTTPS URL. This field is mandatory using a HTTPS URL and is not required if using a HTTP URL.


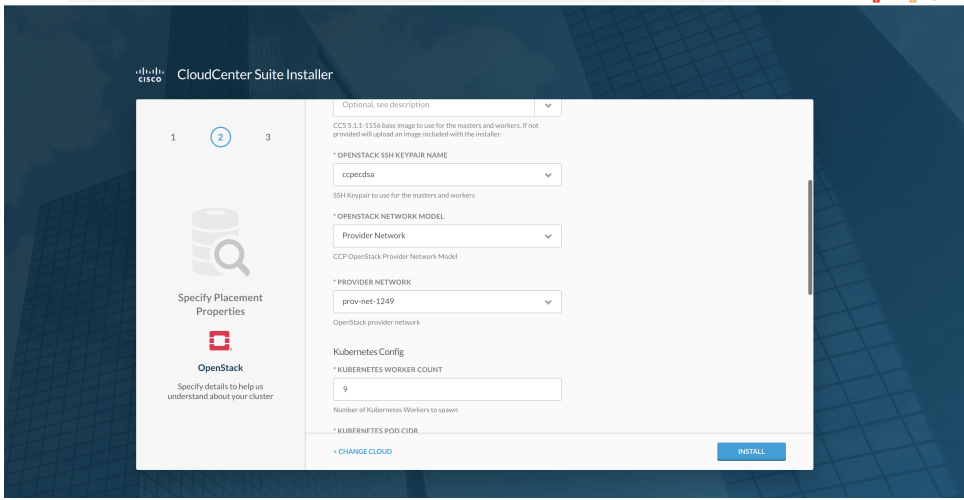
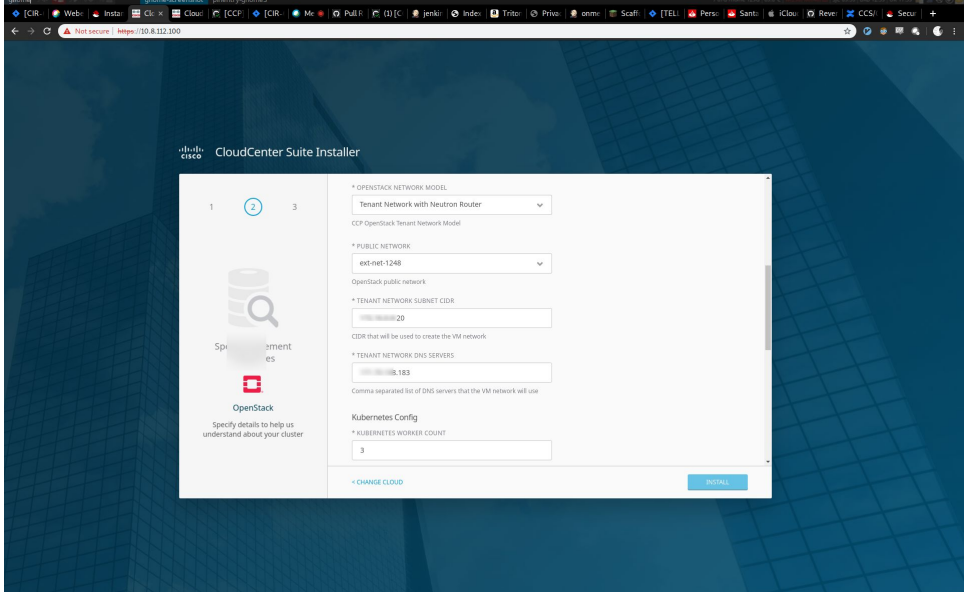
6. Click **Connect**.
7. Once the connection is validated, click **Next**.


To specify the placement properties, enter the following details.



If you setup CloudCenter Suite in offline mode, you must provide valid NTP server details before you save your configuration.

OpenStack Placement Properties	Description
<b>Control Plane Cluster Prefix</b>	Select the OpenStack project to which the Kubernetes cluster is deployed.
<b>OpenStack Details</b>	
<b>OpenStack Flavor UUID</b>	Select one of the existing flavors or VMs. Based on your selection, the recommended number of workers is calculated and displayed in the <b>Kubernetes Worker Count</b> field.
<b>OpenStack Image UUID</b>	<p>Different images will be used for the installer and the cluster launched by the installer. The installer includes a default Kubernetes cluster image (called, <i>CCS-version-Base-Image</i>) with a configurable option to override the use of this default image. The <i>CCS-version-Base-Image</i> image included in the installer is selected if you do not override the setting.</p> <p>To override the <i>CCS-version-Base-Image</i> image used by the Suite installer, be sure to add the applicable image in the <b>OpenStack</b> console and selected the applicable <b>QCOW2</b> image from the dropdown list in this field.</p> <p>If you use the <b>OVA</b> installer to launch the cluster in a vSphere environment, be sure to override this field and select the applicable <b>QCOW2</b> <i>CCS-version-Base-Image</i>.</p>
If you install the CloudCenter Suite using any image other than <i>CCS-version-Base-Image</i> , the installation will fail.	

<p><b>OpenStack SSH Keypair Name</b></p>	<p>Only SSH keys of type <code>ssh-ed25519</code> or <code>ecdsa-sha2-nistp256</code> are supported.</p> <div style="border: 1px solid #ccc; padding: 5px; background-color: #fff9c4; margin-top: 10px;"> <p> You must have at least one existing SSH-key in the selected OpenStack environment to begin the installation.</p> </div>
<p><b>OpenStack Network Model</b></p>	<p>The functional networking model for OpenStack. See <a href="https://docs.openstack.org/security-guide/networking/architecture.html">https://docs.openstack.org/security-guide/networking/architecture.html</a> for additional context.</p>
<p><b>Provider Network or Tenant Network</b></p>	<p><b>Provider Network</b> – Created by the OpenStack administrator on behalf of tenants and can be dedicated to a particular tenant, shared by a subset of tenants, or shared by all tenants. Refer to <a href="https://docs.openstack.org/liberty/networking-guide/intro-os-networking-overview.html">https://docs.openstack.org/liberty/networking-guide/intro-os-networking-overview.html</a> for additional details.</p>  <p><b>Tenant Network</b> – Created by tenants for use by their instances and cannot be shared (based upon default policy settings). Refer to <a href="https://docs.openstack.org/liberty/networking-guide/intro-os-networking-overview.html">https://docs.openstack.org/liberty/networking-guide/intro-os-networking-overview.html</a> for additional details.</p> 
<p><b>Kubernetes Configuration</b></p>	
<p><b>Kubernetes Worker Count</b></p>	<p>This field is auto-populated with the recommended number of worker VMs. While you can change the recommended number, be sure to verify that the worker count is adequate to accommodate the modules that you want to install. See <a href="#">Prepare Infrastructure</a> for additional details.</p>


<b>Kubernetes Pod CIDR</b>	Floating IP pool from which IP addresses are assigned to pods.  <div style="border: 1px solid green; padding: 5px; display: inline-block;">  Verify that this IP does not conflict with the node/VM IP address. </div>
<b>Proxy Configuration</b>	
<b>HTTP Proxy</b>	The hostname or IP address of the proxy host along with the port.
<b>HTTPS Proxy</b>	The hostname or IP address of the secure proxy host along with the port.
<b>NTP Configuration</b>	
<b>NTP Servers</b>	A comma-separated list of IP addresses or FQDNs of your NTP server(s) – to be used to sync VM clocks.
<b>NTP Pools</b>	A comma-separated list of IP addresses or FQDNs of your NTP cluster(s) – to be used to sync VM clocks.

8. Click **Install**. The installation progress is visible on screen.
9. Once successful, you see the following message.

```
CloudCenter Suite installation successful!
```

10. You have the following options at this point:

- a. Click **Take Me To Suite Admin** to launch and set up the [Suite Admin](#).
- b. Click **Install Another Cluster** to start another installation and go back to the homepage (Installer Dashboard).
- c. Download **Kubeconfig file** to connect to the launched cluster using the [kubectl](#) tool.
- d. After the installation is complete, use the following command to SSH into the workers/primary servers as **ubuntu** and use the private SSH key of the public key (provided when you configured the Placement Properties details above).

 Ensure that Port 22 is open on the primary server/worker node so you can provide communication security via Security Groups/Firewall rules for OpenStack environments.

```
#Sample command to SSH into a worker/primary server

ssh -i <private key> ubuntu@<primary server/worker IP>
```

11. Be sure to switch off the installer VM. You can reuse this VM for any other purpose, for example, as an Offline Repository or to upgrade the Kubernetes cluster or to upgrade the tenant image on the nodes.


# VMware vSphere Installation

## VMware vSphere Installation

- [Trial User Installation Procedure and Settings](#)
- [Advanced Prerequisites](#)
- [Advanced VMware Nuances](#)
- [Module Details](#)
- [Advanced Installation Process](#)

In some cases, you may merely want to try out the installation to check if it works. In these cases, try the installation with the following settings, regardless of your environment:

1. Upload the tenant image manually to the root folder and prefix the file with **CCS** (*all upper case*) before you begin the installation.
2. Do not convert the tenant image to be a template.
3. If you are new to Cloud Center Suite, installing CloudCenter Suite for the first time in a VMware environment or if you not sure of your vSphere capacity, then select the following settings in the **Placement properties** page as follows to ensure a successful installation:

Placement Properties Field	Settings and Description
<b>VM Template</b>	Select the image uploaded in as mentioned in <b>Step 1 above</b> .
<b>Resource Pool</b>	Create a <b>new resource pool</b> in your VMware environment and select this new resource pool.
<b>CIDR Network</b>	Placement properties has 2 types of networks: <b>vSphere Network</b> and <b>Kubernetes POD CIDR</b> .  <div style="border: 1px solid #ccc; padding: 10px; background-color: #fff9c4;">  The values for both these networks must be different.  If you select the same network for both settings, the installation will not succeed as the IP that is being assigned will be the same for both networks and thus cause a conflict. </div>
<b>Master VIP</b>	Make sure it is available and not allocated to any other environment <b>before entering the information in this field</b> .
<b>Static IP</b>	Make sure all values are correct and the range is wide enough and available. <b>The number of primary servers, workers, and load balancers must be included in this count.</b>
<b>Number of worker nodes</b>	<b>Reduce the Worker count to 2</b> (even if this field defaults to 5) for an environment that uses 8 CPU 32 GB memory. At a later point ( after your installation/registration is complete), you can increase this count by using the scale up procedure.
<b>Datastore</b>	The updated tenant image and the destination CCS_ image folder <b>MUST</b> have the same Datastore value – to verify this, note the datastore value when you upload the image and then use the same value to enter in this field.

The following **Advanced** sections are intended for users who would like to perform the installation using their environment-specific VMware settings.

If you are using a proxy requirement, be sure to verify that the proxy does not have a username or password restriction.



If you have credentials in place, you will see a field validation error below each Proxy field.



The installation process assumes internet connectivity to certain domains. When installing CloudCenter Suite into environments residing behind a proxy, please ensure the following domains are entirely accessible. Remember the proxy information - this will be used during the installation of CloudCenter Suite.



**Note:** The Installer VM supports HTTP and HTTPS proxies, with or without username and password. The proxy must support TLS 1.2.



**Warning:** Several of the following links might perform redirects. Please ensure your proxy and firewall are configured to allow redirects of the following URLs.

Proxy URL	Description
<a href="https://devhub.cisco.com">https://devhub.cisco.com</a> <a href="http://devhub.cisco.com">http://devhub.cisco.com</a> <a href="https://devhub-docker.cisco.com">https://devhub-docker.cisco.com</a> <a href="http://devhub-docker.cisco.com">http://devhub-docker.cisco.com</a>	Repository for Cisco CloudCenter Suite Docker Charts
<a href="https://gcr.io">https://gcr.io</a> <a href="http://gcr.io">http://gcr.io</a>	Repository for Cisco CloudCenter Suite Helm Charts
<a href="https://storage.googleapis.com">https://storage.googleapis.com</a> <a href="http://storage.googleapis.com">http://storage.googleapis.com</a>	Repository for Cisco CloudCenter Suite Tiller Image
Other	The Suite Installer may require additional connections to the installation environment (for example, vCenter, Hyperflex Data Platform, AWS Console, and so forth) Please ensure your cloud target is reachable via the proxy!

#### A Note on Offline Clusters

In CloudCenter Suite 5.1 and earlier, if your environment has strict URL rules that redirects (for example, using a shorter URL that redirects to <https://storage.googleapis.com>) the configured URL, you may not be able to complete the installation as these kind of redirects may not be allowed if you have installed the repository in an offline cluster. As the offline solution is not completely air gapped in CloudCenter Suite 5.0 and 5.1, you must add these URLs to your allowed lists behind the firewall so you can access these sites.


Verify the following VMware nuances:

- Ensure to use Version 6.0 and higher.
- Verify that you have sufficient shared storage between hosts.
- You must have privileges to launch a VM and access the selected DC/Datastore.
- The datastore clusters are not supported
- The vSphere datastore must reside outside the datastore cluster.
- If vSphere is slow:
  - Upload the VM template manually – in the same datastore where you are going to install CloudCenter Suite.
  - Initially select fewer number of workers than suggested – for example, if 5 workers are recommended, just enter 2 instead of 5. This helps prevent a timeout issue when the workers are being created.
  - After the installation completes, login to CloudCenter Suite as the root tenant (admin) user, click on the **Cloud Management** icon, and scale up the worker node.
  - Static IP Consideration – Verify that you have sufficient IPs available in the Static IP range provided during installation for scale up.
- If vSphere has **more than one datacenter**, be sure to:
  - Create and select one resource pool, do not leave this resource pool selection blank.
  - Upload the tenant image manually to vSphere, under root folder as provided in the following procedure.
    - Download the tenant image tar.gz file from [software.cisco.com](https://software.cisco.com).


- Extract the tenant image. The extracted folder contains the tenant image, rename it by including a CCS prefix. For example: ccp-tenant-image-1.13.5-ubuntu18-4.1.1.ova, rename it to CCS-tenant-image-1.13.5-ubuntu18-4.1.1.ova
- Next, upload this renamed image to your root folder, make sure to select the same data store where you will be installing CloudCenter Suite.
- The image will be displayed in the **VM Template** dropdown of the Placement Properties page.

ccp-tenant-image-1.13.5-ubuntu18-4.1.1.ova 3	--	Folder
└─ README	2 KB	TextEdit
└─ verify	6 KB	Unix executab
└─ ee.pem	2 KB	printabl...arc
└─ sub_ca.pem	2 KB	printabl...arc
└─ root_ca.pem	1 KB	printabl...arc
└─ ccp-tenant-image-1.13.5-ubuntu18-4.1.1.ova.signature	512 bytes	Document
└─ ccp_image_signing_release_v1_pubkey.der	550 bytes	certificate
└─ ccp-tenant-image-1.13.5-ubuntu18-4.1.1.ova	3.52 GB	Document


- Be sure to verify that the image is not converted to the template after uploading to vSphere.
- If vSphere has **only one datacenter**, then it is not mandatory to select a resource pool.
- Your datacenter must exist at the root level.

 Be aware that CloudCenter Suite does not support folders at the root level.

- **Network Time Protocol (NTP) must be configured – this is important as the CloudCenter Suite installation can fail, if NTP is not configured or if it is wrongly configured.**

 If you setup CloudCenter Suite in offline mode, you must provide valid NTP server details before you save your configuration.

- For CloudCenter Suite to use a particular user account in VMware, that account must have the permissions identified in the following table.

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  <div style="border: 1px solid green; border-radius: 10px; padding: 5px; display: inline-block;"> Create this folder under the root folder and be sure to select this path at installation time.</div>
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	<p>Add the following roles and permissions so the tenant image can be uploaded to vSphere under Datacenter during the installation for the given user:</p> <ul style="list-style-type: none"> <li>• Create a role by providing below privileges to this role.</li> <li>• Datastore.Allocate space</li> <li>• Datastore.Browse datastore</li> <li>• Datastore.Low level file operations</li> <li>• Datastore.Remove file</li> <li>• Folder. Create folder</li> <li>• Global.Manage Custom Attributes</li> <li>• Global.Set custom attribute</li> <li>• Network.Assign network</li> <li>• Resource.Apply recommendation</li> <li>• Resource.Apply vApp to resource pool</li> <li>• Resource.Apply virtual machine to resource pool</li> <li>• Storage views. View</li> <li>• Tasks.Create task</li> <li>• Tasks.Update task</li> <li>• Virtual machine (Check all the permissions under this Privilege).</li> <li>• vApp.Import</li> <li>• vApp.Power off</li> <li>• vApp.Power on</li> <li>• vApp.Suspend</li> <li>• vApp.vApp application configuration</li> <li>• vApp.vApp instance configuration</li> <li>• vApp.vApp managedBy configuration</li> <li>• vApp.vApp resource configuration</li> </ul>
Global Role	Set Custom Attributes	To add custom attributes on virtual machines
	Manage Custom Attributes	

Additionally, refer to your module documentation for module-specific dependencies identified in the following table.

Module	Documentation
Workload Manager	<a href="#">Cloud Overview</a>
Action Orchestrator	<a href="#">Add Cloud Account</a>
Cost Optimizer	<a href="#">Cloud Overview</a>

To install the CloudCenter Suite on a new vSphere cluster, perform the following procedure.

1. Verify that you have prepared your environment as listed in the *VMware Nuances* section above.
2. Navigate to the Suite Installer Dashboard.
3. Click **Get Started** in the New Kubernetes Cluster tile to create a new cluster and install the Suite Admin on it.
4. Click vSphere and enter your vSphere credentials identified in the following table and click **Connect**.





vCenter Details	Description
<b>vCenter Server</b>	The DNS address or IP address of the vCenter server.
<b>vCenter Port</b>	The egress endpoint for the vCenter server. For example, Port 443.
<b>vCenter Username</b>	The username to be used for the vCenter setup
<b>vCenter Password</b>	The password to be used for the vCenter setup




5. Once the connection is validated, click **Next**.

To specify the placement properties, enter the vCenter details identified in the following table.

vCenter Details	Description
<b>Datacenter</b>	The name of the vSphere datacenter from where this cluster will be launched
<b>Cluster</b>	The cluster to deploy the node in the above datacenter.
<b>Resource Pool</b>	The resource pool used to deploy the node.



<b>Datastore</b>	The datastore cluster to associate with the node.
<b>Network</b>	The network cluster to associate with the node.
<b>VM Template</b>	<p>Different images will be used for the installer and the cluster launched by the installer. The installer includes a default Kubernetes cluster image (called, <i>CCS-<code>version</code>-Base-Image</i>) with a configurable option to override the use of this default image. The <i>CCS-<code>version</code>-Base-Image</i> image included in the installer is selected if you do not override the setting.</p> <p>To override the <i>CCS-<code>version</code>-Base-Image</i> image used by the Suite installer, be sure to add the applicable image in the vSphere console and selected the applicable OVA from the dropdown list in this field.</p> <p>If you use the OVA installer to launch the cluster in an OpenStack environment, be sure to override this field and select the applicable QCOW2 <i>CCS-<code>version</code>-Base-Image</i>.</p> <div style="border: 1px solid #f96; padding: 5px; margin-top: 10px;">  If you install the CloudCenter Suite using any image other than <i>CCS-<code>version</code>-Base-Image</i>, the installation will fail. </div>
<b>Cluster Folder</b>	The folder which contains the Kubernetes cluster nodes.
<b>Kubernetes Cluster Configuration</b>	
<b>Worker Instance Type</b>	The memory and CPU usage that is required for the workers in your environment. See <a href="#">Prepare Infrastructure &gt; Resource Requirements for CloudCenter Suite Modules</a> for additional context. Based on your selection, the recommendation for the number of nodes in the cluster is also updated (right below this field).
<b>Kubernetes Worker Count</b>	<p>The recommendation from the <i>Worker Instance Type</i> field should be the number that you enter in this field. If you opt to reduce or increase the workers, see <a href="#">Prepare Infrastructure &gt; Resource Requirements for CloudCenter Suite Modules</a> for additional context.</p> <div style="border: 1px solid #f96; padding: 10px; margin-top: 10px;">  The <i>Number of Worker VMs</i> depends on the selected instance type. For example: <ul style="list-style-type: none"> <li>If the instance type is large (8 CPU, 32GB memory), then 5 workers are created and the total static IPs required for this environment are 7 IPs (4 worker VMs, and 3 primary servers).</li> <li>If the instance type is large (8 CPU, 24GB memory), then 5 workers are created and the total static IPs required for this environment are 8 IPs (5 worker VMs, and 3 primary server).</li> <li>If the instance type is large (8 CPU, 16GB memory), then 7 workers are created and the total static IPs required for this environment are 9 IPs (6 worker VMs, and 3 primary server).</li> <li>If the instance type is smaller (4 CPU, 16GB memory), then 9 workers are created, so the static IPs required for this environment are 11 IPs (8 worker VMs, and 3 primary server).</li> </ul> <p>Accordingly, select the IP range by taking into consideration the <i>Number of Worker VMs</i> that will be created based on instance type.</p> <p>This is just an example, be aware that different datacenters will have different instance types configurations and dependencies for each release.</p> </div> <div style="border: 1px solid #90EE90; padding: 10px; margin-top: 10px;">  To determine the number of workers select the instance type in the CloudCenter Suite installer, the <i>Number of Worker VMs</i> are calculated and displayed at the bottom of the instance field as displayed in the following screenshot. </div>
<b>Kubernetes Pod CIDR</b>	<p>The IP address of the pod's Classless Inter-Domain Routing (CIDR) block.</p> <div style="border: 1px solid #90EE90; padding: 10px; margin-top: 10px;">  Verify that this IP does not conflict with the node/VM IP address. </div>
<b>Cluster Prefix</b>	The name of each node in the cluster is prefixed with the information identified in this field to let users know that this node is part of the cluster. This prefix is defined by the user.

<b>IP Allocation Mode</b>	<p>This switch allows you to select the mode:</p> <ul style="list-style-type: none"> <li>• <b>DHCP:</b> This strategy allows the IP to be allocated by the DHCP server to the instance on server boot up. <ul style="list-style-type: none"> <li>• <b>Master VIP:</b> The IP address for the <b>Take Me to Suite Admin</b> link – Users can determine the IP address that should have the primary server role for the <b>Take Me to Suite Admin</b> link. <div data-bbox="462 294 1485 451" style="border: 1px solid #c8e6c9; padding: 10px; margin: 10px 0;"> <p> This should be a unique IP and should not be assigned to any other resource.</p> <p>This should be a unique IP that is not assigned to any other resource. Also, make sure the IP is not in the same range of IPs generated by any DHCP server in your vSphere environment – this will ensure that those IPs are not assigned by the DHCP server to any other node at installation time.</p> </div> </li> </ul> </li> <li>• <b>Static IP:</b> 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. <div data-bbox="418 588 1485 667" style="border: 1px solid #c8e6c9; padding: 10px; margin: 10px 0;"> <p> All IPs should be unique and should not be assigned to any other resource.</p> </div> </li> </ul> <ul style="list-style-type: none"> <li>• <b>Static IP Pool Start IP:</b> The first IP address of the static IP range. If you need to scale up nodes after setting up the Suite Admin, then you must ensure a wider range. The total number of IPs = the total number of nodes required in the cluster (with the scale requirements factored into this number) + 3 IPs for ingress controllers.</li> <li>• <b>Static IP Pool End IP:</b> The last IP address for the static IP range.</li> <li>• <b>Subnet Mask:</b> The netmask corresponding the the specified IP range.</li> <li>• <b>DNS Server List:</b> The comma separated list of DNS server IP addresses.</li> <li>• <b>Gateway List:</b> The comma separated list of Gateway server IP addresses.</li> </ul>
<b>SSH Configuration</b>	
<b>SSH Username</b>	<p>This is a user-assigned field to identify the user for SSH access into worker(s)/primary server(s).</p> <div data-bbox="373 1060 1485 1144" style="border: 1px solid #ffcdd2; padding: 10px; margin: 10px 0;"> <p> Do not use root as a username in this field, any other valid name is acceptable.</p> </div>

<b>SSH Public Key</b>	<p>This field only accepts one of the following keys:</p> <ul style="list-style-type: none"> <li>• ecdsa</li> <li>• ed25519</li> </ul> <p>For either key, you must use the following format:</p> <pre>#ssh-ed25519 ssh-ed25519 &lt;public key&gt; KEY-BODY &lt;username&gt;@&lt;hostname&gt;  #for example ssh-ed25519 AAA*\$...vI48 user@checkmachine  #The UI does not accept keys without &lt;username&gt;@&lt;hostname&gt; -- this is applicable for both ecdca and ssh-ed25519 keys ssh-ed25519 &lt;public key&gt; KEY-BODY &lt;username&gt;@&lt;hostname&gt;  #Example for ed25519 user@checkmachine ssh-ed25519 AAA...vI48 #Example for ecdsa user@checkmachine ssh-ecdsa AAA*\$...vI48  #ecdsa ssh-ecdsa &lt;public key&gt; KEY-BODY &lt;username&gt;@&lt;hostname&gt;  #for example ssh-ecdsa AAA*\$...vI48 diffuser@checkmachine</pre>
-----------------------	--

6. Specify the NTP Configuration details identified in the following table:



If you setup CloudCenter Suite in offline mode, you must provide valid NTP server details before you save your configuration.



When you enter values for the NTP Servers or NTP Pools fields, make sure to enter the NTP value that was assigned to the ESX Host where the CloudCenter Suite installer was created. This NTP value is available in the Placement Properties page at the time of installation, see [VMware vSphere Appliance Setup](#) > Step 1 for details.

Identical NTP values are required to ensure that the NTP communication between the installer and CloudCenter Suite primary server /worker VMs are in sync so the certificates generated by the installer for CloudCenter Suite are also in sync.

NTP Details	Description
<b>NTP Servers</b>	The list of IP addresses or FQDNs of your NTP server(s) – to be used to sync VM clocks.
<b>NTP Pools</b>	The list of IP addresses or FQDNs of your NTP pools.

7. If you are in an environment that uses Proxy connections to access the internet, you need to configure the settings identified in the following table.



If you use proxy values for both HTTP and HTTPS, *enter any one of the two values, not both*. If you enter both the HTTP and HTTPS values, then the UI dashboard may not display modules.

Proxy details	Description
<b>HTTP Proxy</b>	The IP addresses and port of the HTTP proxy server.
<b>HTTPS Proxy</b>	The IP addresses and port of the HTTPS proxy server.

8. Click **Install**. The installation progress is visible on screen.

9. Once successful, you see the following message.

```
CloudCenter Suite installation successful!
```

10. You have the following options at this point:

- a. Click **Take Me To Suite Admin** to launch and set up the [Suite Admin](#).
- b. Click **Install Another Cluster** to start another installation and go back to the homepage (Installer Dashboard).
- c. Download **KubeConfig file** to connect to the launched cluster using the [kubectl](#) tool.
- d. After the installation is complete, use the following command to SSH (using the SSH credentials configured during installation.) into the workers/primary servers as **cloud-user** and use the private SSH key or the public key (provided when you configured the Properties details above).

```
#Sample command to SSH into a worker/primary server

• ssh -I <private key> cloud-user@<Installer IP>

#or

• ssh -I <private key> ssh-user@<worker/primary server IP>
```

11. Be sure to switch off the installer VM. You can reuse this VM for any other purpose, for example, as an Offline Repository or to upgrade the Kubernetes cluster or to upgrade the tenant image on the nodes.

# Existing Cluster Installation

## Install the CloudCenter Suite on an Existing Kubernetes Cluster

- [Overview](#)
- [Restrictions](#)
- [Prerequisites](#)
- [Procedure](#)

Once you access the Suite Installer Dashboard (see [Prepare Infrastructure](#)), you can choose to install the Suite Admin on an existing cluster.

Before proceeding with section, adhere to the following restrictions:

- **AWS:** The CloudCenter Suite does not currently support a Suite Admin installation on an existing AWS cluster.
- **Permission:** Admin-level permissions for the cluster are mandatory for a user to install the Suite Admin in an existing cluster.

Verify that the cluster adheres to the following requirements:

- **Kubernetes Version:** The existing Kubernetes cluster must be of Version v1.14.x or and later.
- **Kubernetes Add Ons:** Install Cert-manager version v0.7.0 (required) using the following command (refer to <https://cert-manager.readthedocs.io/en/latest/> for details):

```
kubectl apply -f https://raw.githubusercontent.com/jetstack/cert-manager/release-0.5/contrib/manifests/cert-manager/with-rbac.yaml
```

- **Instance Type:** The instance type for GKE is should be n1-standard-8 or higher. Verify that it is large enough to accommodate the installation of Suite Admin and other CloudCenter Suite modules.
- **Basic Authentication:** When creating the GKE cluster, go to **Security** and check the box to **Enable Basic Authentication**.
- **Storage Class:** The default storageClass must be configured.
- **Kubeconfig:** The kubeconfig user must have cluster-admin permission in the kubeconfig namespace.
  - If the cluster does not support Load Balancer.
  - GCP: You must remove auth provider and use the admin user password.
- **RBAC** - Must be enabled.
- **Pod Priority:** Define the PriorityClass for suite-high/suite-medium/suite-low.
  - Refer to <https://kubernetes.io/docs/concepts/configuration/pod-priority-preemption/> for details.
  - The commands to define PriorityClass are listed in the following code block.

```
# create pod priority class: suite-high/suite-medium/suite-low
##### begin create pod priority

cat <<EOF | kubectl apply -f -

apiVersion: scheduling.k8s.io/v1beta1

kind: PriorityClass

metadata:
  name: suite-high

value: 1000000

globalDefault: false

description: "High priority"

---

apiVersion: scheduling.k8s.io/v1beta1

kind: PriorityClass

metadata:
  name: suite-medium

value: 10000

globalDefault: false

description: "Medium priority"

---

apiVersion: scheduling.k8s.io/v1beta1

kind: PriorityClass

metadata:
  name: suite-low

value: 100

globalDefault: false

description: "Low priority"

EOF

##### end create pod priority
```

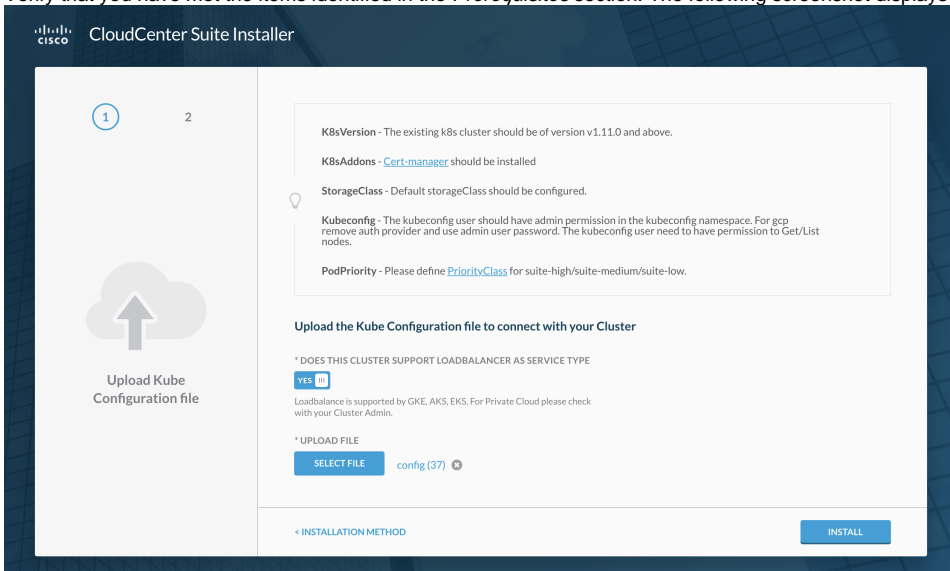
To install the CloudCenter Suite on an existing cluster, perform the following procedure.

1. Navigate to the Suite Installer Dashboard.

2. Click **Existing Cluster** to get started as displayed in the following screenshot.



3. Verify that you have met the items identified in the *Prerequisites* section. The following screenshot displays these items as well.



4. Identify if your cluster supports **load balancer as the service type** – accordingly, turn this toggle
- YES** – Toggle ON if supported (public clouds generally support load balancers)
  - NO** – Toggle OFF if not supported (private clouds generally do not support load balancers)
5. Upload the Kubeconfig file.

Click **Install**. The installation progress is visible on screen. Once successful, you see the following message .

```
CloudCenter Suite installation successful!
```

6. You have the following options at this point:
- Click **Take Me To Suite Admin** to launch and set up the **Suite Admin**.
  - Click **Install Another Cluster** to start another installation on the same cluster.

You have now installed the Suite Admin on an existing cluster.

# Upgrade Kubernetes Cluster

## Upgrade Kubernetes Cluster

Access the Suite Installer Dashboard (see [Prepare Infrastructure](#)) to install a new cluster and launch nodes for the new Kubernetes cluster

- [Upgrade Approach](#)
- [Amazon EKS Upgrade](#)
- [Azure AKS Upgrade](#)
- [Google GKE Upgrade](#)
- [OpenStack Upgrade](#)
- [VMware vSphere Upgrade](#)



# Upgrade Approach

## Upgrade Approach

- [Overview](#)
- [Restrictions](#)
- [Prerequisites](#)
- [Process](#)

This section provides details on restrictions, prerequisites, and the process to upgrade the Kubernetes cluster. During this upgrade, the software upgrades the cluster and migrates the pods to new worker instances.



If you restart any worker node, be sure to wait for approximately 10 minutes before logging into the CloudCenter Suite – this timeline is determined by the pods taking about 10 minutes to startup.

Before proceeding with an upgrade, adhere to the following restrictions:

- **Usage:** To upgrade the Kubernetes cluster to a new version, you can do so from CloudCenter Suite 5.1.0 and later releases.
  - You cannot use the CloudCenter Suite 5.1 upgrader to upgrade a CloudCenter Suite 5.0 cluster. You can only use the CloudCenter Suite 5.1 upgrader effective CloudCenter Suite 5.1.1 to upgrade to a later release.
  - As an upgrader is not available to upgrade from CloudCenter Suite 5.0 to CloudCenter Suite 5.1, you must use the [Backup and Restore](#) procedure to upgrade to a CloudCenter Suite 5.1 cluster.
  - Even if you update the Suite Admin to Suite Admin 5.1, the underlying cluster will not have the capability to be upgraded as it is still using CloudCenter Suite 5.0.
  - **Public Clouds:**
    - By upgrading the cluster, you upgrade to the applicable Kubernetes version.
  - **Private Clouds:**
    - By upgrading the cluster, you are performing a rolling upgrade on each base image in the cluster.
    - A rolling upgrade may or may not include a change in the Kubernetes version – it may merely apply an OS patch or address vulnerabilities depending on the image version that you use.
    - The installer includes a default Kubernetes cluster image (called, *CCS-version-Base-Image*). The VM Template contains a list of tenant images with a CCS-version-Base-Image name format. If you want to upgrade to a version other than the default version provided by the installer, then upload that CCS-version-Base-Image under the root folder, so that it will display in this dropdown list. You can use this option to upgrade the cluster across private clouds.
- **Suite Admin-level Permissions:** The Suite Admin-level permissions are mandatory for a user to upgrade the cluster.
- **New Clusters Only:** You can upgrade a cluster that is created (from the Suite Installer) using the **New cluster** option.



If you created your cluster by clicking the **Existing cluster** option (using the KubeConfig file), then you cannot upgrade this cluster using the process provided in this section.

Verify that the cluster adheres to the following requirements:

- **Backup Environment:** Back up your environment before initiating the upgrade. See [Backup](#) for additional details.
- **Schedule Downtime:** Schedule a suitable downtime during off-peak hours to minimize the impact to your users and or customers. Communicate the downtime as the CloudCenter Suite will not be accessible during the upgrade.
- **Verify Kubernetes Version:** Verify that the existing Kubernetes cluster is Version v1.11.0 and above.

This is the generic process to upgrade a Kubernetes cluster for a cloud that is supported by the CloudCenter Suite.

1. Navigate to the Suite Installer Dashboard (see [Prepare Infrastructure](#)).



2. Click **Upgrade** in the *Upgrade Kubernetes Cluster* section to specify the credentials for your cluster as displayed in the following screenshot.

3. Enter the Suite Admin URL (or DNS), username, password, and Tenant ID for the admin account.
4. Identify if this is **An Amazon EKS Cluster** by toggling the switch (default is No). If it is, provide the Access Key and Secret Key details.

See the individual cloud upgrade pages for additional notes and nuances.

5. Click **Connect** to validate your credentials.
6. At this point, you have multiple scenarios:

- You will be able to click **Next** and select the desired Kubernetes version from the dropdown list for this upgrade. Proceed to Step 8.
- If an upgrade is not available for your cluster as displayed in the following screenshot, some possible reasons are:
  - An upgrade is not currently available as the cluster is already at the latest available version of Kubernetes.

- You may have provided the wrong cluster credentials (in this case, you will not see the *Connected* status update when you try to connect). If so, enter the right credentials and try again.

- Once Connected, you see the cloud type and other information on the left side of the screen as visible in the following screenshot (sample of a GKE environment):
- If an upgrade is available, select the **Desired K8s version** for the upgrade.
- Click **Upgrade** to upgrade the Kubernetes cluster as well as the master and worker nodes once the upgrade is complete. A progress bar with relevant status messages is displayed.



An upgrade operation can take more than one hour depending on the number of nodes to be upgraded and cloud response time.

- At this point, you can:
  - Download the latest logs to track the upgrade process.
  - Wait for cluster to finish upgrading.
- The installation progress and success is visible on the screen.



See the individual cloud upgrade pages for which of these options are available and for additional notes and nuances.

- You have the following options at this point – depending on your cloud environment:
  - Click **Take Me To Suite Admin** to launch and set up the [Suite Admin](#).
  - Click **Install Another Cluster** to start another installation on the same cluster.
  - Download the **Kubeconfig file**.
  - Download the **SSH private key**.
  - Re-purpose the installer server.
- Login to CloudCenter Suite using valid credentials and verify that your information is preserved and that the cluster was upgraded.

# Amazon EKS Upgrade

## Amazon EKS Upgrade

- [Overview](#)
- [Amazon Nuances](#)
- [Module Details](#)
- [Minimum Permissions Needed](#)
- [Installation Process](#)

See [Upgrade Approach](#) for details on permissions and prerequisites.

Be aware of the following requirements when installing the CloudCenter Suite:

- **Maximum Supported Version:** EKS Version 1.13.7 and below.
- **Unavailable Resources:** The following resources will not be available until the upgrade completes:
  - EKS cluster
  - Suite admin cluster
- **Resources:** Amazon creates the following resources for the AWS account:
  - An EKS Cluster with user-provided specifications.
  - All resources remain in the same region as the cluster.
  - A new CloudFormation stack with the same number of instances, security groups, subnets, and roles that are used to connect to the cluster.
    - VPC Name: *cluster\_name-VPC*
    - Role Name for VPC: *cluster\_name-Role*
    - Role Name for Workers: *cluster\_name-NodeInstanceRole*
    - New CFN stack Name: *cluster\_name-New-Workers-random\_UUID32*
    - Auto Scaling Group for worker nodes as part of cloud formation workers stack
- **The Delete API:**



You cannot trigger a Delete call by deleting the Amazon cluster from either the AWS console or the AWS CLI. Instead, use the Delete API.

Additionally, refer to your module documentation for module-specific dependencies as specified in the following table.

Module	Documentation
Workload Manager	<a href="#">Cloud Overview</a>
Action Orchestrator	<a href="#">Add Cloud Account</a>
Cost Optimizer	<a href="#">Cloud Overview</a>

The following IAM policies are required for the CloudCenter Suite to access the EKS and create a new cluster on AWS.

- AmazonEC2FullAccess
- IAMFullAccess
- AutoScalingFullAccess
- AmazonEKSClusterPolicy
- AmazonEKSWorkerNodePolicy
- AmazonVPCFullAccess
- AmazonEKSServicePolicy
- AmazonEKS\_CNI\_Policy
- AmazonRoute53FullAccess
- Inline\_Policy\_EKS\_Cluster = an inline policy allowing the following actions on the EKS service to an IAM user:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "cloudformation:CreateStack",
        "cloudformation:DescribeStacks",
        "cloudformation:DescribeStackEvents",
        "cloudformation:DescribeStackResources",
        "cloudformation:DescribeStackResource",
        "cloudformation:GetTemplate",
        "cloudformation:ValidateTemplate",
        "cloudformation>DeleteStack",
        "eks:UpdateClusterVersion",
        "cloudformation:UpdateStack",
        "eks:ListUpdates",
        "eks:DescribeUpdate",
        "eks:DescribeCluster",
        "eks:ListClusters",
        "eks:CreateCluster",
        "eks>DeleteCluster"
      ],
      "Resource": "*"
    }
  ]
}

```

To upgrade the cluster for an Amazon EKS Kubernetes environment, perform the following procedure.

1. Navigate to the Suite Installer Dashboard (see [Prepare Infrastructure](#)).
2. Click **Upgrade** in the *Upgrade Kubernetes Cluster* section to specify the credentials for your cluster as displayed in the following screenshot.
3. Enter the Suite Admin DNS (or URL), username, password, and Tenant ID for the admin account.
4. Identify if this is an **Amazon EKS Cluster** by toggling the switch (the default is **No**).
5. Provide the Access Key and Secret Key details for the Amazon EKS Cluster as visible in the following screenshot.

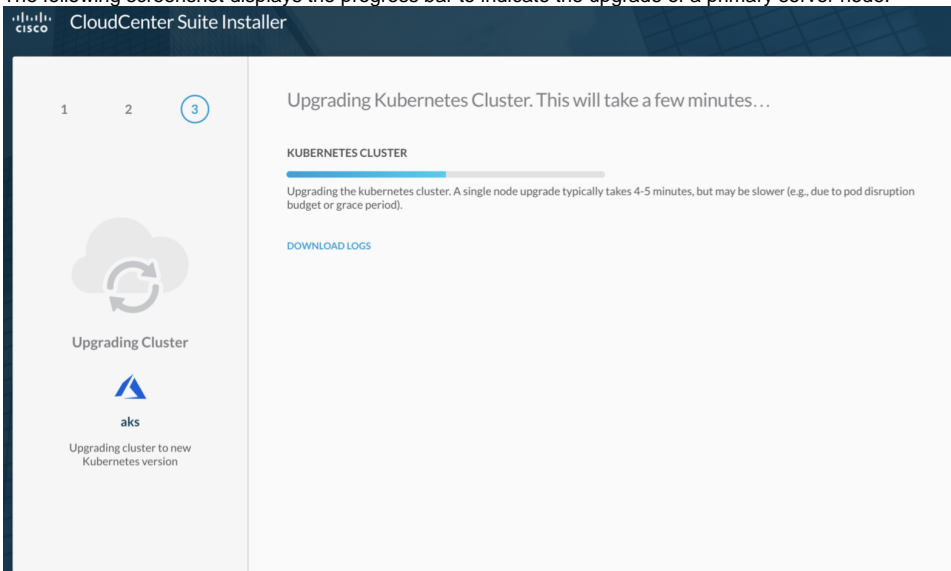


The CloudCenter Suite validates the EKS credentials to ensure that the EKS cluster is available to this user.

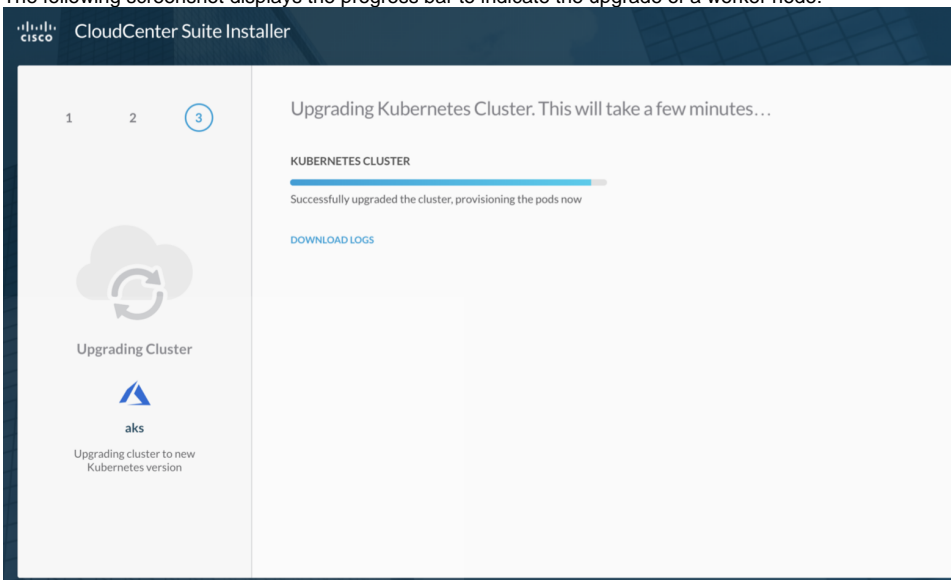
6. Click **Connect** to validate your credentials. Once Connected, you see the cloud type and other information on the left side of the screen.
7. Click **Next** and select the desired Kubernetes version from the dropdown list for this upgrade.
8. If an upgrade is available, select the **Desired K8s version** for the upgrade.
9. Click **Upgrade** to upgrade the Kubernetes cluster as well as the primary server and worker nodes once the upgrade is complete. A progress bar with relevant status messages is displayed.

- ✓ An upgrade operation can take more than one hour depending on the number of nodes to be upgraded and cloud response time.

- a. The following screenshot displays the progress bar to indicate the upgrade of a primary server node:



- b. The following screenshot displays the progress bar to indicate the upgrade of a worker node:



10. At this point, you can:
- Download the latest logs to track the upgrade process.
  - Wait for cluster to finish upgrading.
11. The installation progress is visible on screen. Once successful, you see the success message displayed.

```
CloudCenter Suite installation successful!
```

12. You have the following options at this point:
- Click **Take Me To Suite Admin** to launch and set up the [Suite Admin](#).
  - Click **Install or Upgrade Another Cluster** to start another installation on the same cluster.
  - Download the **Kubeconfig file**.
  - Re-purpose the installer server.
13. Login to CloudCenter Suite using valid credentials and verify that your information is preserved and that the cluster was upgraded.

You have now upgraded the cluster on the EKS cloud. Verify your Suite Admin and tenant data.

# Azure AKS Upgrade

## Azure AKS Upgrade

- [Overview](#)
- [Azure Nuances](#)
- [Module Details](#)
- [Installation Process](#)

See [Upgrade Approach](#) for details on permissions and prerequisites.

Be aware of the following requirements to install CloudCenter Suite:

- **Maximum Supported Version:** AKS Version 1.12 and below.
- **Valid Azure Account:** A valid service account that allows you to use sufficient resource quota. See <https://docs.microsoft.com/en-us/azure/aks/container-service-quotas> for additional details.
- **Resource Group:** Create the resource group in a cloud region that supports Azure.

Additionally, refer to your module documentation for module-specific dependencies as displayed in the following table.

Module	Documentation
Workload Manager	<a href="#">Cloud Overview</a>
Action Orchestrator	<a href="#">Add Cloud Account</a>
Cost Optimizer	<a href="#">Cloud Overview</a>

To upgrade the cluster for an Azure AKS Kubernetes environment, perform the following procedure.

1. Navigate to the Suite Installer Dashboard (see [Prepare Infrastructure](#)).
2. Click **Upgrade** in the *Upgrade Kubernetes Cluster* section to specify the credentials for your cluster as displayed in the following screenshot.
3. Enter the Suite Admin URL (or DNS), username, password, and Tenant ID for the admin account.
4. Identify if this is **An Amazon EKS Cluster** by toggling the switch (default is No). If it is, provide the Access Key and Secret Key details.
5. Click **Connect** to validate your credentials. Once Connected, you see the cloud type and other information on the left side off the screen as visible in the following screenshot.
6. Click **Next** and select the desired Kubernetes version from the dropdown list for this upgrade.
7. If an upgrade is available, select the **Desired K8s version** for the upgrade.
8. Click **Upgrade** to upgrade the Kubernetes cluster as well as the master and worker nodes once the upgrade is complete. A progress bar with relevant status messages is displayed.



An upgrade operation can take a long time depending on the number of nodes to be upgraded and cloud response time.

9. At this point, you can:
  - a. Download the latest logs to track the upgrade process.
  - b. Wait for cluster to finish upgrading.
10. The installation progress is visible on screen. Once successful, you see the success message displayed.

```
CloudCenter Suite installation successful!
```

11. You have the following options at this point:
  - a. Click **Take Me To Suite Admin** to launch and set up the [Suite Admin](#).
  - b. Click **Install or Upgrade Another Cluster** to start another installation on the same cluster.
  - c. Download the **Kubeconfig file**.
  - d. Download the **SSH private key**.
  - e. Re-purpose the installer server.
12. Log in to CloudCenter Suite using valid credentials and verify that your information is preserved and that the cluster was upgraded.

You have now upgraded the cluster on the AKS cloud. Verify your Suite Admin and tenant data.



# Google GKE Upgrade

## Google GKE Upgrade

- [Overview](#)
- [Google Nuances](#)
- [Module Details](#)
- [Installation Process](#)

See [Upgrade Approach](#) for details on permissions and prerequisites.

Be aware of the following requirements when installing the CloudCenter Suite:

- **Maximum Supported Version:** GKE Version 1.12 and below.
- **Permissions:** Verify that the person upgrading the cluster has the following minimum permissions (roles) as displayed in the screenshot:

A service account represents a Google Cloud service identity, such as code running on Compute Engine VMS, App Engine apps, or systems running outside Google.

The screenshot shows the configuration page for a service account named 'ad'. It lists three roles assigned to the account:

- Service Account User:** Create VMs and other GCP tasks with a service account. Users cannot impersonate the account directly as they can with Service Account Actor role.
- Kubernetes Engine Admin:** Full management of Kubernetes Clusters and their Kubernetes API objects.
- Compute Admin:** Full control of all Compute Engine resources.

Below the roles, there are two checkboxes:

- Furnish a new private key**: Downloads a file that contains the private key. Store the file securely because this key can't be recovered if lost.
- Enable G Suite Domain-wide Delegation**: Allows this service account to be authorized to access all users' data on a G Suite domain without manual authorization on their parts. [Learn more](#)

- Service Account User
- Kubernetes Engine Admin
- Compute Engine Admin

Additionally, refer to your module documentation for module-specific dependencies as identified in the following table:

Module	Documentation
Workload Manager	<a href="#">Cloud Overview</a>
Action Orchestrator	<a href="#">Add Cloud Account</a>
Cost Optimizer	<a href="#">Cloud Overview</a>

To upgrade the cluster for a GKE Kubernetes environment, perform the following procedure.

1. Navigate to the Suite Installer Dashboard (see [Prepare Infrastructure](#)).
2. Click **Upgrade** in the *Upgrade Kubernetes Cluster* section to specify the credentials for your cluster as displayed in the following screenshot.
3. Enter the Suite Admin URL (or DNS), username, password, and Tenant ID for the admin account.
4. Identify if this is an **Amazon EKS Cluster** by toggling the switch (default is No). If it is, provide the Access Key and Secret Key details.
5. Click **Connect** to validate your credentials. Once Connected, you see the cloud type and other information on the left side of the screen as visible in the following screenshot.

6. Click **Next** and select the desired Kubernetes version from the dropdown list for this upgrade.

CloudCenter Suite Installer

1 2 3

Specify Cluster Credentials

gke

Connect to your cluster to check if an upgrade is possible

\* SUITE ADMIN ENDPOINT FOR THE CLUSTER TO BE UPGRADED  
https://34.68.248.60  
Suite Admin URL, with port(if required) - Eg. https://example.com:1234

\* EMAIL ADDRESS(USERNAME) OF THE SUITE ADMIN  
admin@cisco.com  
Suite Admin Email Address(Username)

\* PASSWORD FOR THE SUITE ADMIN  
\*\*\*\*\*  
Suite Admin Password

\* TENANT ID FOR THE SUITE ADMIN  
cisco  
Suite Admin Tenant ID

\* IS THIS AN AMAZON EKS CLUSTER?  
 NO  
Toggle to provide credentials of Amazon Elastic Kubernetes Service cluster

EDIT ✔ Connected

< BACK TO MENU NEXT >

7. If an upgrade is available, select the **Desired K8s version** for the upgrade.
8. Click **Upgrade** to upgrade the Kubernetes cluster as well as the primary server and worker nodes once the upgrade is complete. A progress bar with relevant status messages is displayed.



An upgrade operation can take more than one hour depending on the number of nodes to be upgraded and cloud response time.

- a. The following screenshot displays the progress bar to indicate the upgrade of a primary server node:

CloudCenter Suite Installer

1 2 3

Upgrading Cluster

gke

Upgrading cluster to new Kubernetes version


Upgrading Kubernetes Cluster. This will take a few minutes...

KUBERNETES CLUSTER

Upgrading kubernetes cluster's masters

[DOWNLOAD LOGS](#)

- b. The following screenshot displays the progress bar to indicate the upgrade of a worker node:

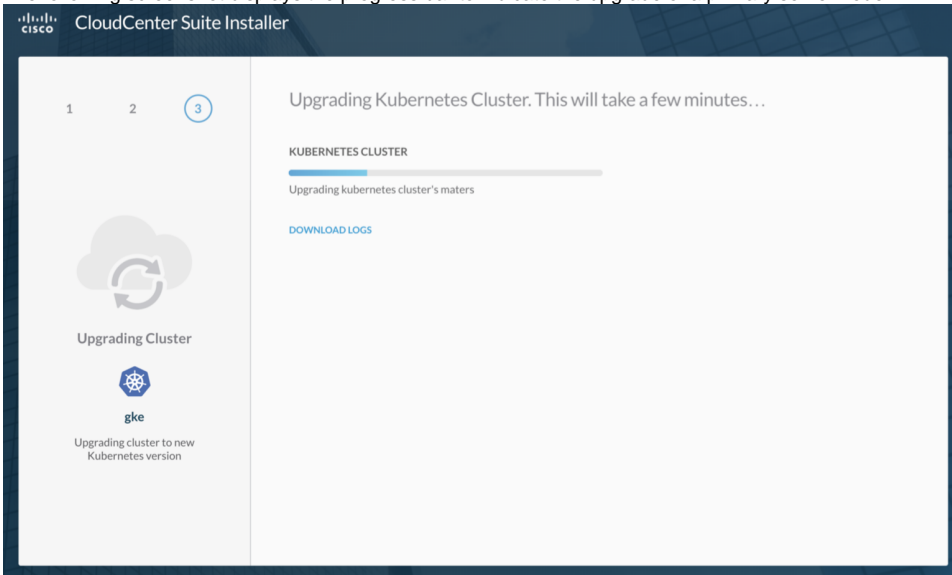


9. At this point, you can:


- Download the latest logs to track the upgrade process.
- Wait for cluster to finish upgrading.

10. The installation progress is visible on screen.

- The following screenshot displays the progress bar to indicate the upgrade of a primary server node:



- The following screenshot displays the progress bar to indicate the upgrade of a worker node:



Once successful, you see the success message displayed.

```
CloudCenter Suite installation successful!
```

11. You have the following options at this point:

- a. Click **Take Me To Suite Admin** to launch and set up the [Suite Admin](#).
- b. Click **Install or Upgrade Another Cluster** to start another installation on the same cluster.
- c. Download the **Kubeconfig file**.
- d. Re-purpose the installer server.

12. Login to CloudCenter Suite using valid credentials and verify that your information is preserved and that the cluster was upgraded.

You have now upgraded the cluster on the GKE cloud. Verify your Suite Admin and tenant data.

# OpenStack Upgrade

## OpenStack Upgrade

- [Overview](#)
- [OpenStack Nuances](#)
- [Module Details](#)
- [Installation Process](#)

See [Upgrade Approach](#) for details on permissions and prerequisites.

Verify the following OpenStack nuances:

- OpenStack newton release with at least the following service versions:
  - Cinder v2
  - Keystone v3
  - OpenStack Nova v2
  - OpenStack Networking v2
  - OpenStack Glance v2
- Ensure to add Port 6443 to the default security group as the security group created for the cluster is not automatically assigned to the load balancer created for the cluster.
- The tenant and project requirements for OpenStack Cloud are identified in the following table.

Model	Quota	Description
For all cases	2 (primary server group, worker group)	Server Groups
	Number of workers + number of primary servers	Server Group Members
	3 (API load balancers)	Load Balancers
	6 (2 for each load balancer)	Health Monitors
	6 (2 for each load balancer)	Pools
	6 (2 for each load balancer)	Listeners
	3 (1 for the cluster VMs, 2 for the Kubernetes load balancer services)	Security Groups
	18	Security Group Rules
	See <a href="#">Prepare Infrastructure</a> for additional details	Volume GB
	Number of workers + number of primary servers + 3 for each load balancer	Ports
	Number of workers + number of primary servers	Instances
	16 GB (recommended for each worker and each primary server)	RAM
32 (recommended for each workers and each primary server)	vCPUs	
Tenant network	Floating IPs = 3	1 for each load balancer
	Networks = 1	For the tenant network
	Subnet = 1	For the tenant network
	Router = 1	For the tenant network to public network connection
Provider network	Number of workers + number of primary servers + 3 load balancers	Free IPs in the provider network

- **Network Time Protocol (NTP) must be configured – this is important as the CloudCenter Suite installation can fail, if NTP is not configured or if it is wrongly configured.**



If you setup CloudCenter Suite in offline mode, you must provide valid NTP server details before you save your configuration.

Additionally, refer to your module documentation for module-specific dependencies as identified in the following table:

Module	Documentation
--------	---------------

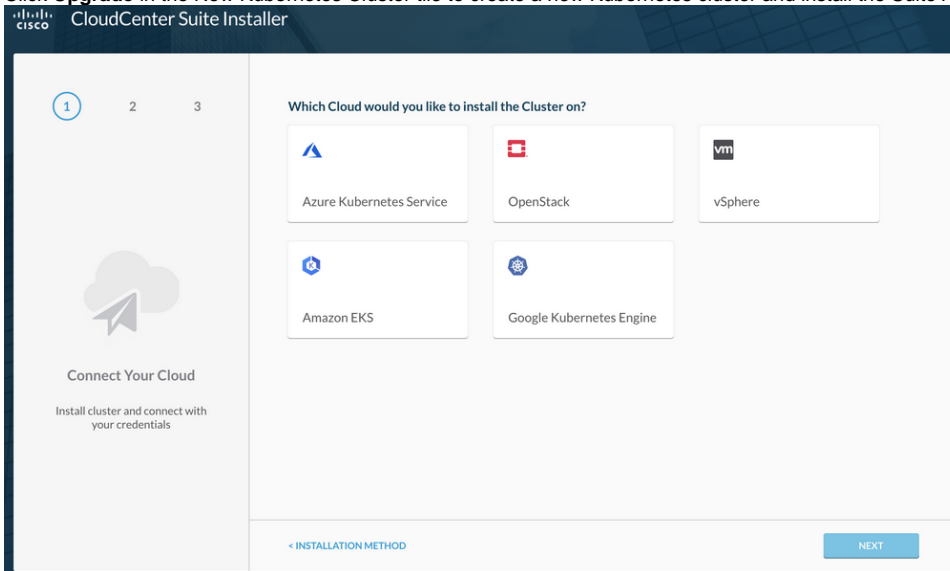
Workload Manager	<a href="#">Cloud Overview</a>
Action Orchestrator	<a href="#">Add Cloud Account</a>
Cost Optimizer	<a href="#">Cloud Overview</a>

To upgrade the cluster for an OpenStack Kubernetes environment, perform the following procedure.

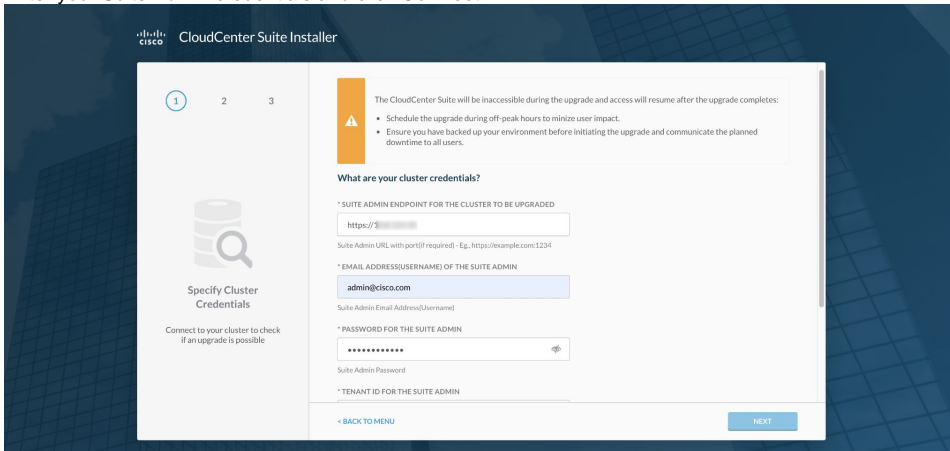
1. Verify that you have prepared your environment as listed in the *OpenStack Nuances* section above.
2. Navigate to the Suite Installer Dashboard.



3. Click **Upgrade** in the New Kubernetes Cluster tile to create a new Kubernetes cluster and install the Suite Admin on it.

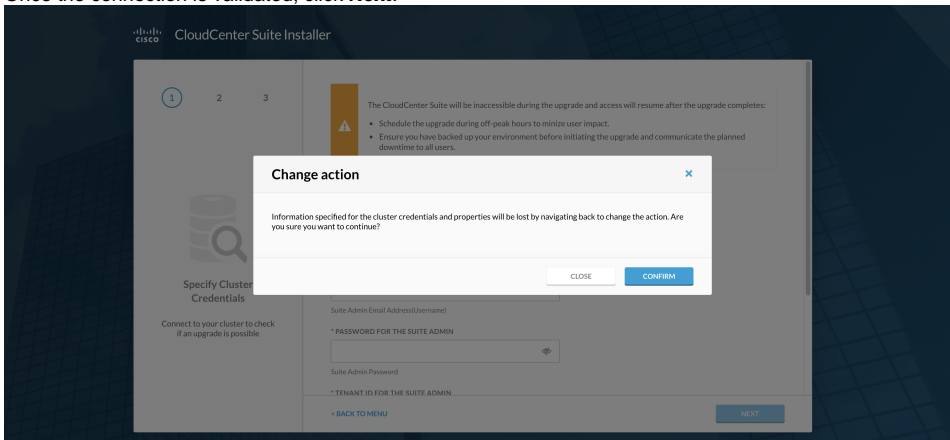


4. Click the **OpenStack** tile. You see the Specify Cluster Credentials page.

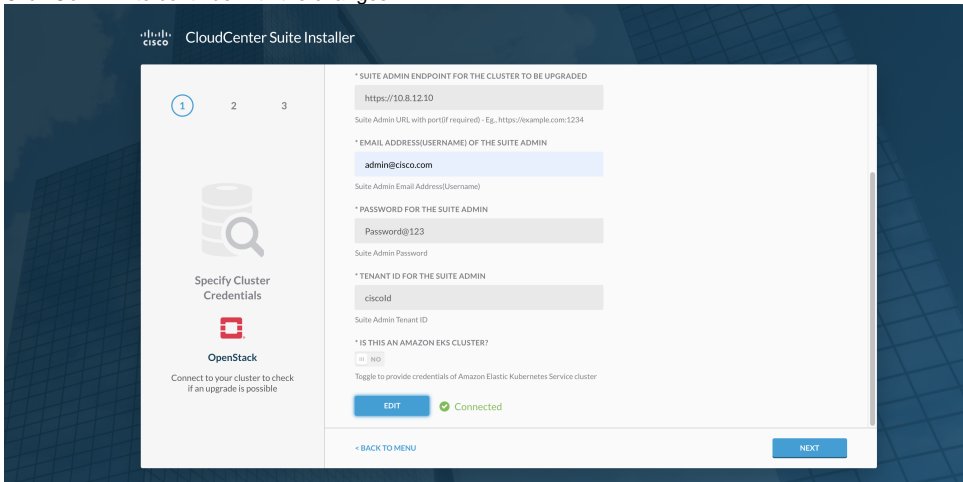
5. Enter your Suite Admin credentials and click **Connect**.

OpenStack Details	Description
<b>Suite Admin Endpoint for the Cluster to be Upgraded</b>	The DNS address or IP address of the vCenter server where you launch the Suite Admin.
<b>Email Address (Username) of the Suite Admin</b>	The email address of Suite Admin (the <a href="#">Initial Administrator</a> ) who setup the Suite Admin.
<b>Password for the Suite Admin</b>	The password for the Suite Admin (the <a href="#">Initial Administrator</a> ) who setup the Suite Admin.
<b>Tenant ID for the Suite Admin</b>	The Tenant ID for the Suite Admin (the <a href="#">Initial Administrator</a> ) who setup the Suite Admin.
<b>Is This an Amazon EKS Cluster</b>	Toggle the switch (default = No). If it is, provide the Access Key and Secret Key details.

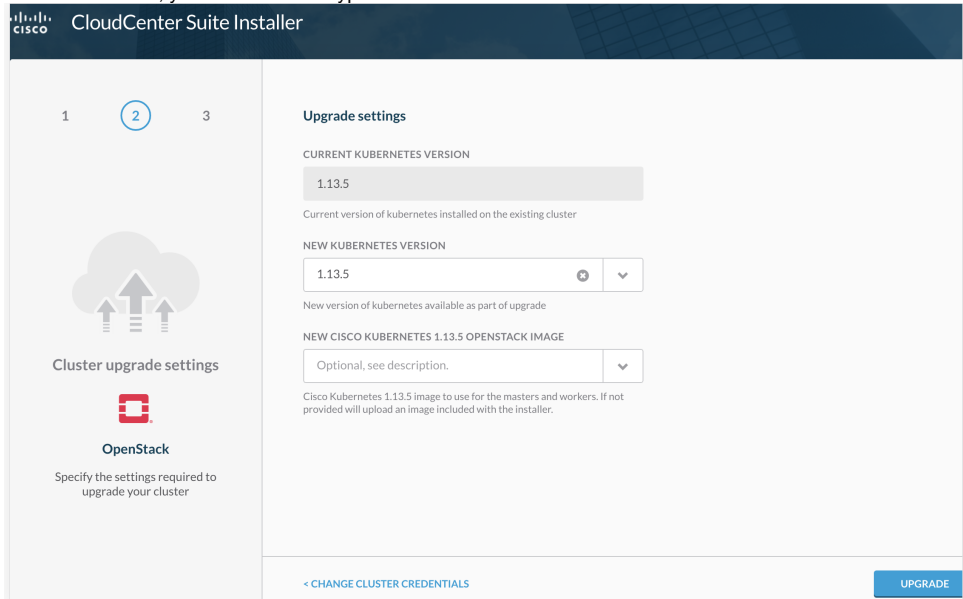
The CloudCenter Suite validates the OpenStack credentials to ensure that the cluster is available to this user.

6. Once the connection is validated, click **Next**.

7. Click **Confirm** to continue with the changes.

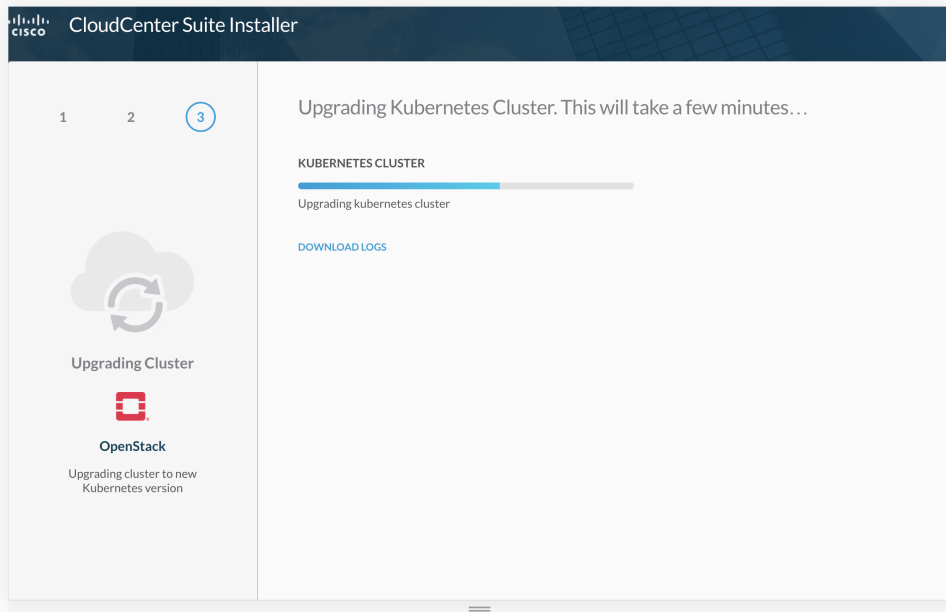
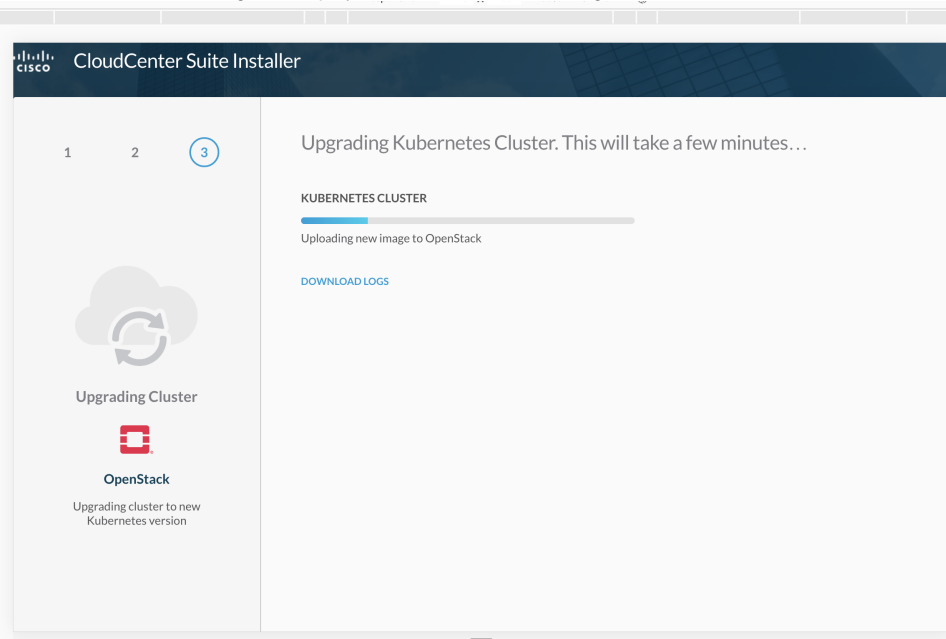


8. When Connected, you see the cloud type and other information on the left side off the screen – enter the information in the Upgrade settings fields



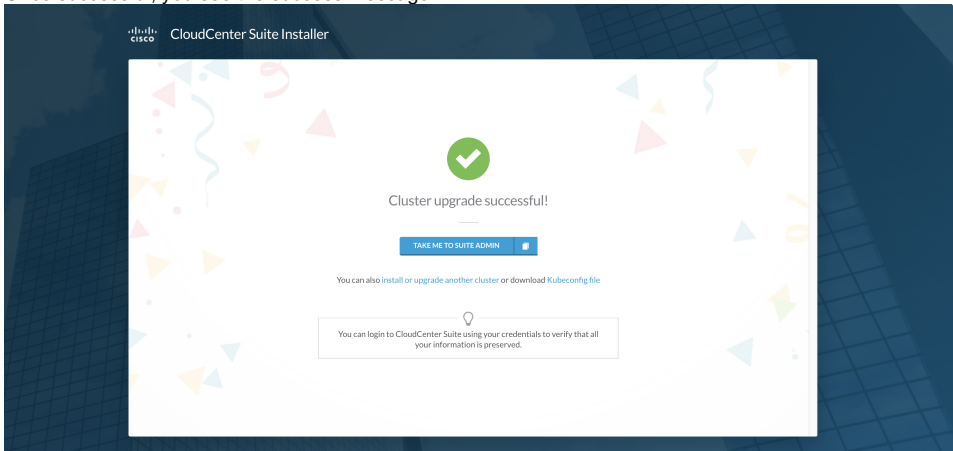


9. Click **Upgrade** to upgrade the Kubernetes cluster as well as the primary server and worker nodes once the upgrade is complete. A progress bar with relevant status messages is displayed as visible in the following screenshots.



10. At this point, you can:
- Download the latest logs to track the upgrade process.
  - Wait for cluster to finish upgrading.

11. Once successful, you see the success message.



You have the following options at this point:

- a. Click **Take Me To Suite Admin** to launch and set up the [Suite Admin](#).
  - b. Click **Install Another Cluster** to start another installation and go back to the homepage (Installer Dashboard).
  - c. Download **Kubeconfig file** to connect to the launched cluster using the [kubectl](#) tool.
12. After the installation is complete, use the following command to SSH into the workers/primary servers as **cloud-user** and use the private SSH key or the public key (provided when you configured the Placement Properties details above).

```
#Sample command to SSH into a worker/primary server  
ssh -I <private key> cloud-user@<primary server/worker IP>
```

13. Login to CloudCenter Suite using valid credentials and verify that your information is preserved and that the cluster was upgraded.

You have now upgraded the cluster on the OpenStack cloud. Verify your Suite Admin and tenant data.

# VMware vSphere Upgrade


## VMware vSphere Upgrade

- [Overview](#)
- [Trial User Installation Procedure and Settings](#)
- [Advanced Prerequisites](#)
- [Advanced VMware Nuances](#)
- [Module Details](#)
- [Installation Process](#)

See [Upgrade Approach](#) for details on permissions and prerequisites.

In some cases, you may merely want to try out the installation to check if it works. In these cases, try the installation with the following settings, regardless of your environment:

1. Upload the tenant image manually to the root folder and prefix the file with **CCS** (*all upper case*) before you begin the installation.
2. Do not convert the tenant image to be a template.
3. If you are new to Cloud Center Suite, installing CloudCenter Suite for the first time in a VMware environment or if you not sure of your vSphere capacity, then select the following settings in the **Placement properties** page as follows to ensure a successful installation:

Placement Properties Field	Settings and Description
<b>VM Template</b>	Select the image uploaded in as mentioned in <b>Step 1 above</b> .
<b>Resource Pool</b>	Create a <b>new resource pool</b> in your VMware environment and select this new resource pool.
<b>CIDR Network</b>	Placement properties has 2 types of networks: <b>vSphere Network</b> and <b>Kubernetes POD CIDR</b> .  <div style="border: 1px solid #ffc107; padding: 10px; margin-top: 10px;">  The values for both these networks must be different.  If you select the same network for both settings, the installation will not succeed as the IP that is being assigned will be the same for both networks and thus cause a conflict. </div>
<b>Master VIP</b>	Make sure it is available and not allocated to any other environment <b>before entering the information in this field</b> .
<b>Static IP</b>	Make sure all values are correct and the range is wide enough and available. <b>The number of primary servers, workers, and load balancers must be included in this count</b> .
<b>Number of worker nodes</b>	<b>Reduce the Worker count to 2</b> (even if this field defaults to 5) for an environment that uses 8 CPU 32 GB memory. At a later point ( after your installation/registration is complete), you can increase this count by using the scale up procedure.
<b>Datastore</b>	The updated tenant image and the destination CCS_ image folder <b>MUST</b> have the same Datastore value – to verify this, note the datastore value when you upload the image and then use the same value to enter in this field.

The following **Advanced** sections are intended for users who would like to perform the installation using their environment-specific VMware settings.

If you are using a proxy requirement, be sure to verify that the proxy does not have a username or password restriction.



If you have credentials in place, you will see a field validation error below each Proxy field.



The installation process assumes internet connectivity to certain domains. When installing CloudCenter Suite into environments residing behind a proxy, please ensure the following domains are entirely accessible. Remember the proxy information - this will be used during the installation of CloudCenter Suite.



**Note:** The Installer VM supports HTTP and HTTPS proxies, with or without username and password. The proxy must support TLS 1.2.



**Warning:** Several of the following links might perform redirects. Please ensure your proxy and firewall are configured to allow redirects of the following URLs.

Proxy URL	Description
<a href="https://devhub.cisco.com">https://devhub.cisco.com</a> <a href="http://devhub.cisco.com">http://devhub.cisco.com</a> <a href="https://devhub-docker.cisco.com">https://devhub-docker.cisco.com</a> <a href="http://devhub-docker.cisco.com">http://devhub-docker.cisco.com</a>	Repository for Cisco CloudCenter Suite Docker Charts
<a href="https://gcr.io">https://gcr.io</a> <a href="http://gcr.io">http://gcr.io</a>	Repository for Cisco CloudCenter Suite Helm Charts
<a href="https://storage.googleapis.com">https://storage.googleapis.com</a> <a href="http://storage.googleapis.com">http://storage.googleapis.com</a>	Repository for Cisco CloudCenter Suite Tiller Image
Other	The Suite Installer may require additional connections to the installation environment (for example, vCenter, Hyperflex Data Platform, AWS Console, and so forth) Please ensure your cloud target is reachable via the proxy!

#### A Note on Offline Clusters

In CloudCenter Suite 5.1 and earlier, if your environment has strict URL rules that redirects (for example, using a shorter URL that redirects to <https://storage.googleapis.com>) the configured URL, you may not be able to complete the installation as these kind of redirects may not be allowed if you have installed the repository in an offline cluster. As the offline solution is not completely air gapped in CloudCenter Suite 5.0 and 5.1, you must add these URLs to your allowed lists behind the firewall so you can access these sites.


Verify the following VMware nuances:

- Ensure to use Version 6.0 and higher.
- Verify that you have sufficient shared storage between hosts.
- You must have privileges to launch a VM and access the selected DC/Datastore.
- The datastore clusters are not supported
- The vSphere datastore must reside outside the datastore cluster.
- If vSphere is slow:
  - Upload the VM template manually – in the same datastore where you are going to install CloudCenter Suite.
  - Initially select fewer number of workers than suggested – for example, if 5 workers are recommended, just enter 2 instead of 5. This helps prevent a timeout issue when the workers are being created.
  - After the installation completes, login to CloudCenter Suite as the root tenant (admin) user, click on the **Cloud Management** icon, and scale up the worker node.
  - Static IP Consideration – Verify that you have sufficient IPs available in the Static IP range provided during installation for scale up.
- If vSphere has **more than one datacenter**, be sure to:
  - Create and select one resource pool, do not leave this resource pool selection blank.
  - Upload the tenant image manually to vSphere, under root folder as provided in the following procedure.
    - Download the tenant image tar.gz file from [software.cisco.com](https://software.cisco.com).


- Extract the tenant image. The extracted folder contains the tenant image, rename it by including a CCS prefix. For example: ccp-tenant-image-1.13.5-ubuntu18-4.1.1.ova, rename it to CCS-tenant-image-1.13.5-ubuntu18-4.1.1.ova
- Next, upload this renamed image to your root folder, make sure to select the same data store where you will be installing CloudCenter Suite.
- The image will be displayed in the **VM Template** dropdown of the Placement Properties page.

ccp-tenant-image-1.13.5-ubuntu18-4.1.1.ova 3	--	Folder
└─ README	2 KB	TextEdit
└─ verify	6 KB	Unix executab
└─ ee.pem	2 KB	printabl...arc
└─ sub_ca.pem	2 KB	printabl...arc
└─ root_ca.pem	1 KB	printabl...arc
└─ ccp-tenant-image-1.13.5-ubuntu18-4.1.1.ova.signature	512 bytes	Document
└─ ccp_image_signing_release_v1_pubkey.der	550 bytes	certificate
└─ <b>ccp-tenant-image-1.13.5-ubuntu18-4.1.1.ova</b>	<b>3.52 GB</b>	<b>Document</b>


- Be sure to verify that the image is not converted to the template after uploading to vSphere.
- If vSphere has **only one datacenter**, then it is not mandatory to select a resource pool.
- Your datacenter must exist at the root level.

 Be aware that CloudCenter Suite does not support folders at the root level.

- **Network Time Protocol (NTP) must be configured – this is important as the CloudCenter Suite installation can fail, if NTP is not configured or if it is wrongly configured.**

 If you setup CloudCenter Suite in offline mode, you must provide valid NTP server details before you save your configuration.

- For CloudCenter Suite to use a particular user account in VMware, that account must have the permissions identified in the following table.

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
<div style="border: 1px solid #90EE90; border-radius: 10px; padding: 10px; display: inline-block;">  Create this folder under the root folder and be sure to select this path at installation time.         </div>		
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	<p>Add the following roles and permissions so the tenant image can be uploaded to vSphere under Datacenter during the installation for the given user:</p> <ul style="list-style-type: none"> <li>• Create a role by providing below privileges to this role.</li> <li>• Datastore.Allocate space</li> <li>• Datastore.Browse datastore</li> <li>• Datastore.Low level file operations</li> <li>• Datastore.Remove file</li> <li>• Folder. Create folder</li> <li>• Global.Manage Custom Attributes</li> <li>• Global.Set custom attribute</li> <li>• Network.Assign network</li> <li>• Resource.Apply recommendation</li> <li>• Resource.Apply vApp to resource pool</li> <li>• Resource.Apply virtual machine to resource pool</li> <li>• Storage views. View</li> <li>• Tasks.Create task</li> <li>• Tasks.Update task</li> <li>• Virtual machine (Check all the permissions under this Privilege).</li> <li>• vApp.Import</li> <li>• vApp.Power off</li> <li>• vApp.Power on</li> <li>• vApp.Suspend</li> <li>• vApp.vApp application configuration</li> <li>• vApp.vApp instance configuration</li> <li>• vApp.vApp managedBy configuration</li> <li>• vApp.vApp resource configuration</li> </ul>
Global Role	Set Custom Attributes	To add custom attributes on virtual machines
	Manage Custom Attributes	

Additionally, refer to your module documentation for module-specific dependencies identified in the following table.

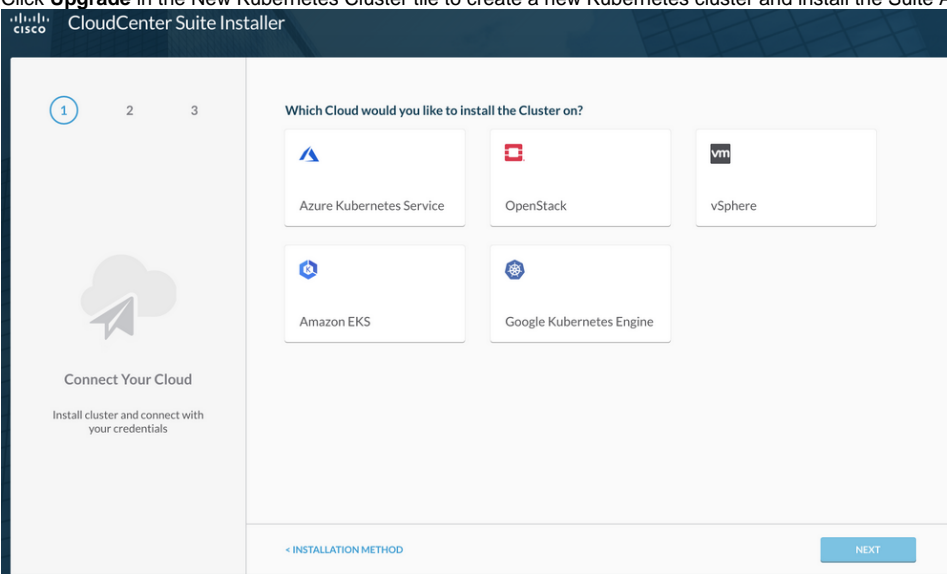
Module	Documentation
Workload Manager	<a href="#">Cloud Overview</a>
Action Orchestrator	<a href="#">Add Cloud Account</a>
Cost Optimizer	<a href="#">Cloud Overview</a>

To install the CloudCenter Suite on a new vSphere cluster, perform the following procedure.

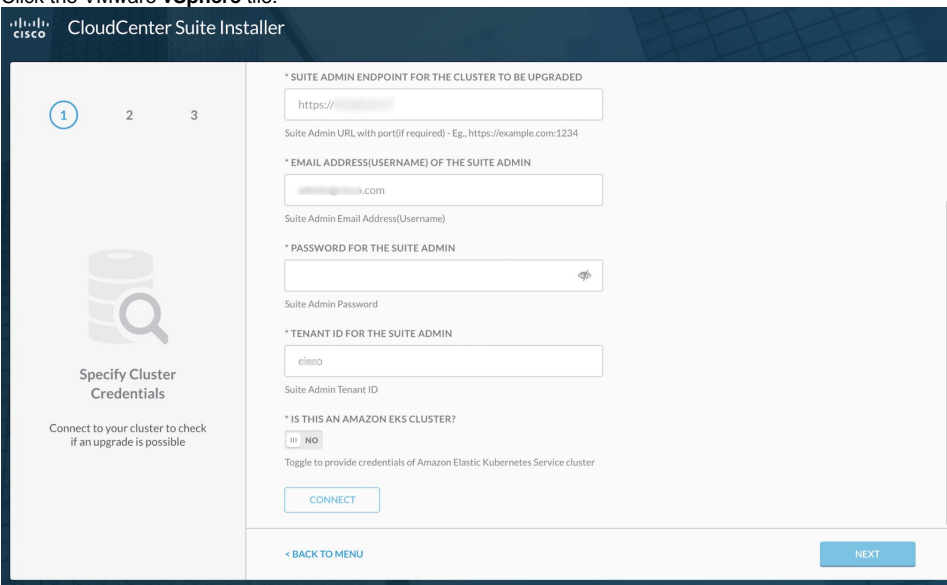
1. Verify that you have prepared your environment as listed in the *VMware Nuances* section above.
2. Navigate to the Suite Installer Dashboard.



3. Click **Upgrade** in the New Kubernetes Cluster tile to create a new Kubernetes cluster and install the Suite Admin on it.



4. Click the VMware **vSphere** tile.

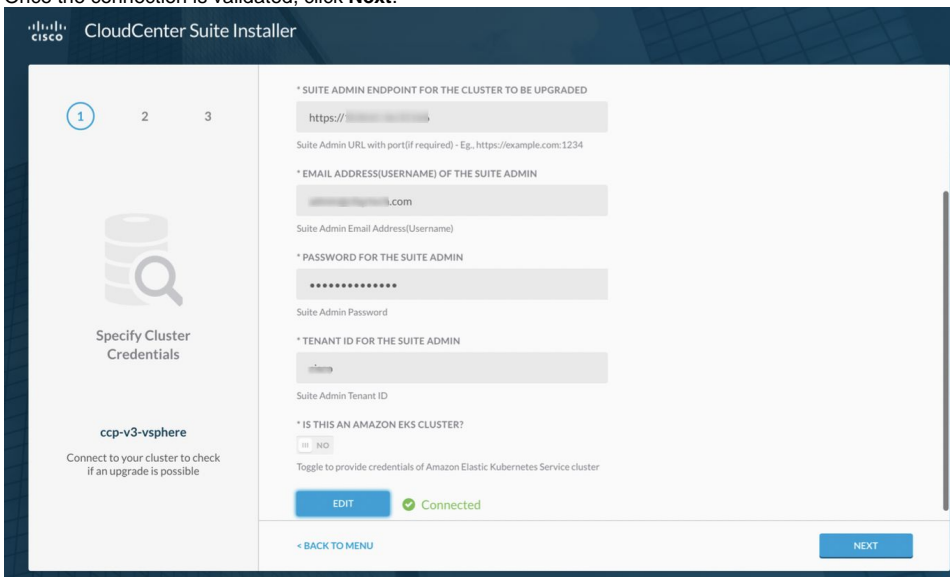


5. Enter your Suite Admin credentials and click **Connect**.

vSphere Details	Description
<b>Suite Admin Endpoint for the Cluster to be Upgraded</b>	The DNS address or IP address of the vCenter server where you launch the Suite Admin.
<b>Email Address (Username) of the Suite Admin</b>	The email address of Suite Admin (the <a href="#">Initial Administrator</a> ) who setup the Suite Admin.
<b>Password for the Suite Admin</b>	The password for the Suite Admin (the <a href="#">Initial Administrator</a> ) who setup the Suite Admin.
<b>Tenant ID for the Suite Admin</b>	The Tenant ID for the Suite Admin (the <a href="#">Initial Administrator</a> ) who setup the Suite Admin.
<b>Is This an Amazon EKS Cluster</b>	Toggle the switch (default is No). If it is, provide the Access Key and Secret Key details.

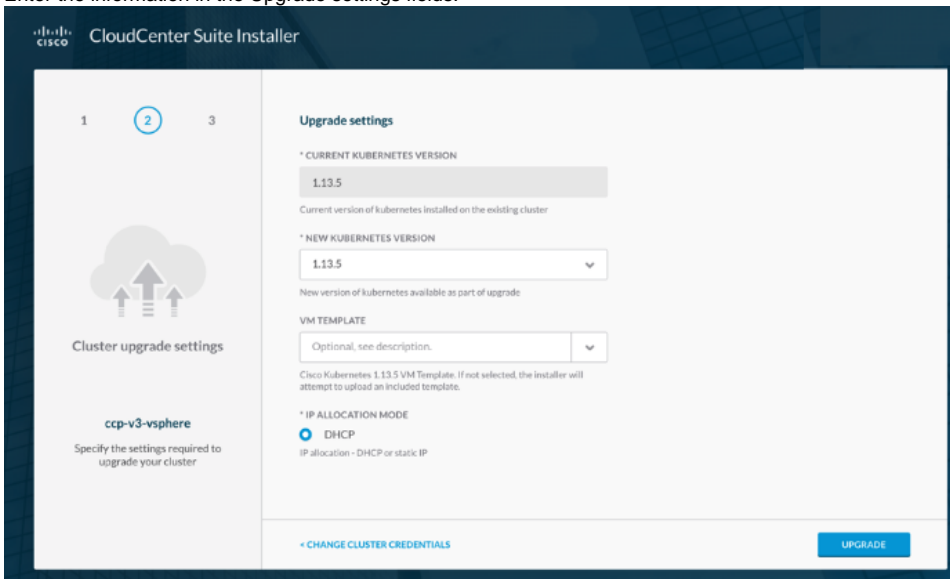
The CloudCenter Suite validates the vSphere credentials to ensure that the cluster is available to this user.

6. Once the connection is validated, click **Next**.



Once Connected, you see the cloud type and other information on the left side off the screen

7. Enter the information in the Upgrade settings fields.




Upgrade Settings Field	Description
<b>Current Kubernetes Version</b>	The current version for your Kubernetes setup is pre-populated in this field.
<b>New Kubernetes Version</b>	If an upgrade is available, it is listed in this dropdown list. Select the <b>Desired K8s version</b> for the upgrade.



**VM  
Template**

Different images will be used for the installer and the cluster launched by the installer as visible in the following screenshot.



The installer includes a default Kubernetes cluster image (called, *CCS-*version*-Base-Image*). The VM Template contains a list of tenant images with a *CCS-*version*-Base-Image* name format. If you want to upgrade to a version other than the default version provided by the installer, then upload that *CCS-*version*-Base-Image* under the root folder, so that it will display in this dropdown list.

The *CCS-*version*-Base-Image* image included in the installer is selected if you do not override the setting.

To override the *CCS-*version*-Base-Image* image used by the Suite installer, be sure to add the applicable image in the vSphere console and selected the applicable **OVA** from the dropdown list in this field.

If you use the **OVA** installer to launch the cluster in an OpenStack environment, be sure to override this field and select the applicable **QCOW2** *CCS-*version*-Base-Image*.



If you install the CloudCenter Suite using any image other than *CCS-*version*-Base-Image*, the installation will fail.

**IP  
Allocation  
Mode**

This switch allows you to select the mode. Currently, only DHCP is supported.

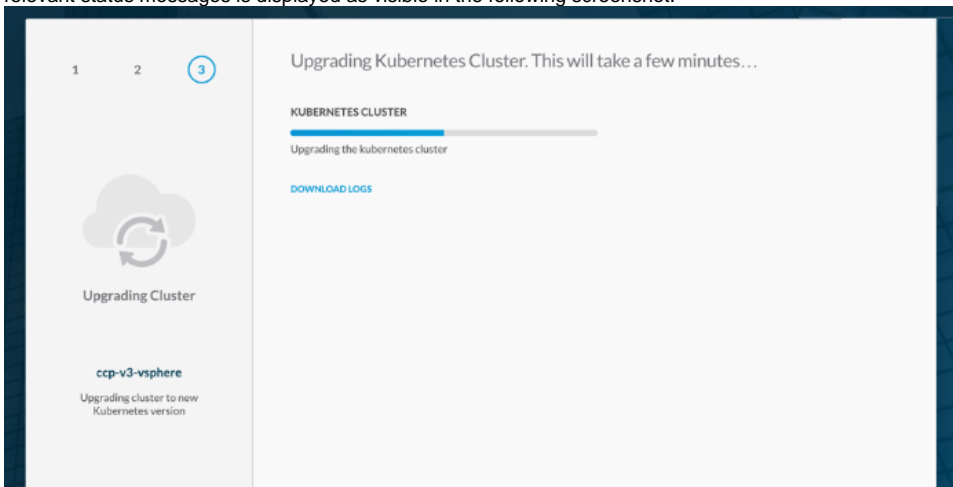
- **DHCP:** This strategy allows the IP to be allocated by the DHCP server to the instance on server boot up.
  - **Master VIP:** The IP address for the **Take Me to Suite Admin** link – Users can determine the IP address that should have the primary role for the **Take Me to Suite Admin** link.



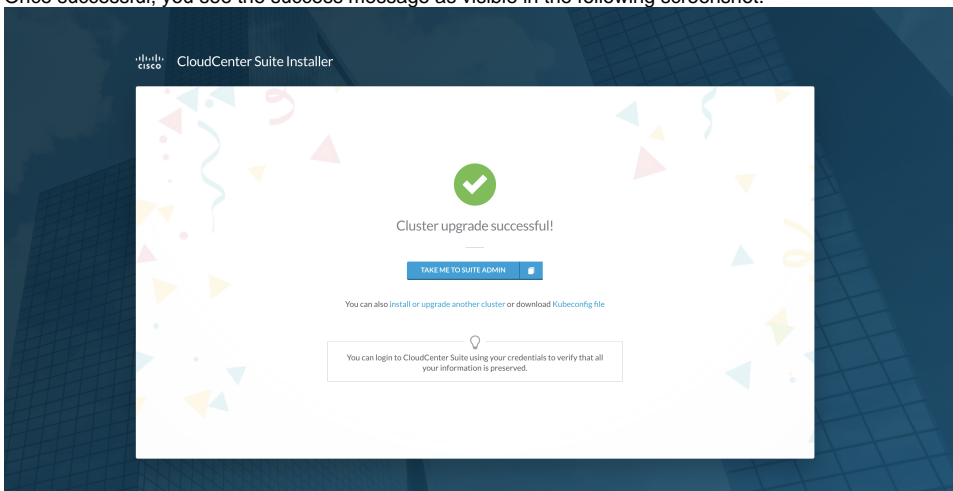
This should be a unique IP and should not be assigned to any other resource.

- **Static IP:** 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.
  - **Static IP Pool Start IP:** The first IP address of the static IP range. If you need to scale up nodes after setting up the Suite Admin, then you must ensure a wider range.
  - **Static IP Pool End IP:** The last IP address for the static IP range.
  - **Subnet Mask:** The netmask corresponding the the specified IP range.
  - **DNS Server List:** The comma separated list of DNS server IP addresses.
  - **Gateway List:** The comma separated list of Gateway server IP addresses.

8. Click **Upgrade** to upgrade the Kubernetes cluster as well as the primary and worker nodes once the upgrade is complete. A progress bar with relevant status messages is displayed as visible in the following screenshot.



9. At this point, you can:
- Download the latest logs to track the upgrade process.
  - Wait for cluster to finish upgrading.
10. Once successful, you see the success message as visible in the following screenshot.



11. You have the following options at this point:
- Click **Take Me To Suite Admin** to launch and set up the [Suite Admin](#).
  - Click **Install or Upgrade Another Cluster** to start another installation and go back to the homepage (Installer Dashboard).
  - Download **KubeConfig file** to connect to the launched cluster using the [kubectl](#) tool.
  - After the installation is complete, use the following command to SSH into the workers/primary servers as **cloud-user** and use the private SSH key or the public key (provided when you configured the Placement Properties details above).

```
#Sample command to SSH into a worker/primary server
• ssh -I <private key> cloud-user@<primary server/worker IP>
```

12. Be sure to switch off the installer VM. You can reuse this VM for any other purpose, for example, as an Offline Repository.

You have now upgraded the cluster on the VMware cloud. Verify your Suite Admin and tenant data.

# Offline Repository

## Offline Repository

- [Introduction](#)
- [Prerequisites to Configure the Offline Repository](#)
- [Setup the Offline Repository](#)

A repository connection enables access from one of the CloudCenter Suite VMs to the default **Cisco Products Repository**. This default repository is only accessible if you have **direct internet access**.

The CloudCenter Suite will try to connect to the Cisco Products Repository to install and upgrade the product modules.



- The default Cisco Products Repository is only accessible if your underlying Kubernetes cluster has direct internet access.
- If you are behind a proxy environment, you must provide the proxy settings in the installers and you will not need an offline repository (for private clouds).

If you do not have internet access, you must connect the CloudCenter Suite to the offline repository (see [Offline Repository Configuration](#) for additional details).

After you create a VM from the OVA, you have the option to use the VM as an offline repository server.

The offline repository connects to the default Cisco Products Repository and allows you to install or upgrade within the Suite Administration.



The offline repository is the same for all supported clouds – and is only supported for OpenStack and VMware clouds.

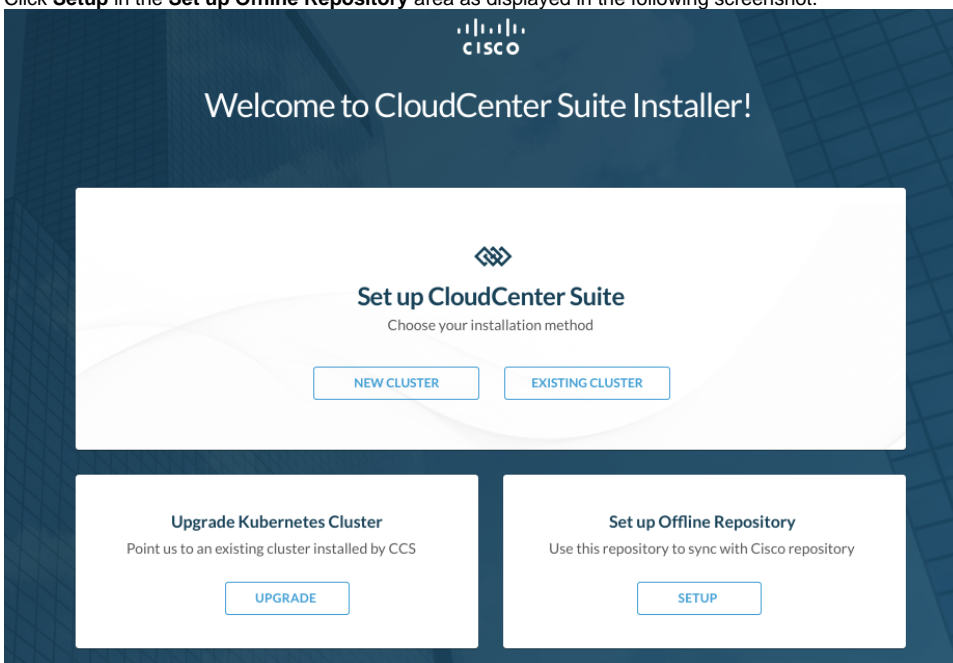
- The offline repository VM must have access to the Cisco repository at [devhub-docker.cisco.com](https://devhub-docker.cisco.com) and [devhub.cisco.com](https://devhub.cisco.com).
- You must manually set up a valid DNS name for the repository VM.
- You must get a valid certificate and a private key pair for the DNS name (self-signed certificates are *not* acceptable, you must get a [Certificate Authority](#) to sign the certificate)
- The repository VM must be accessible from the Kubernetes cluster through the domain name.
- *Optional.* If your offline repository server requires a proxy to connect to the Internet, you must have the proxy configuration ready.
- A VM that was used for the installation can also be used as an offline repository after the installation completes.




Once converted to an offline repository, this VM can no longer be used as the installer VM.

To setup an offline repository, follow this procedure.


1. Click **Setup** in the **Set up Offline Repository** area as displayed in the following screenshot.




2. Click **Select File** to upload the certificate and the private key as displayed in the following screenshot.

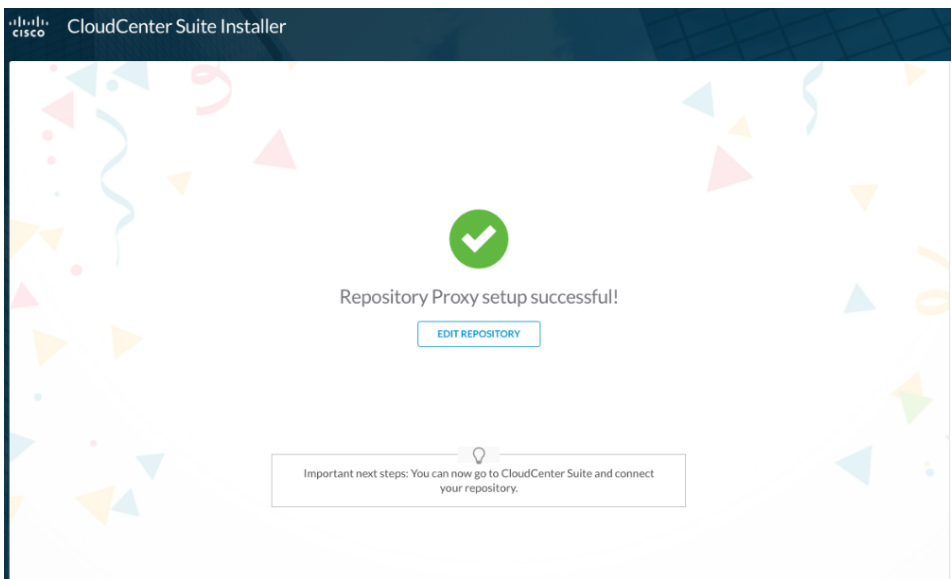
 Verify that the certificate and private key files have been assigned 755 permissions (full permissions for the owner, and read/execute permission for others).

3. Enter the DNS name in the **Offline Repository Domain Name** field.
4. (Optional) Enter the proxy IP or DNS name in the HTTPS Proxy Host field..

 This step is required only if you connect to the Internet via a proxy.

5. Click **Done** to complete the installation.
6. Navigate back to CloudCenter Suite to connect to the repository and perform further actions in CloudCenter Suite as displayed in the following screenshot.


 Once you setup offline repository, note its DNS name so you can re-launch the CloudCenter Suite Installer Repository Proxy success page so you can edit the repository details at a later date.



## Edit Offline Repository

Once you have set up the repository, you can click the **Edit Repository** link to change the certificates, DNS name, and proxy settings.

If you are editing the repository at a later time, use the DNS address (that you noted down after you *Setup the Offline Directory* as described in the previous section) to re-launch the CloudCenter Suite Installer Repository Proxy success page and click the **Edit Repository** link as displayed in the following screenshots.



**Upload Files**

Upload certificate and private key to configure repository

**Configure your offline repository.**

\* UPLOAD CERTIFICATE


\* UPLOAD PRIVATE KEY

\* OFFLINE REPOSITORY DOMAIN NAME  
  
valid domain name matching the provided certificate

HTTPS PROXY HOST  
  
proxy host IP address or DNS name, e.g. myproxy.com

PORT  
  
proxy port, allowed value range is from 1 to 65535

PROXY REQUIRES PASSWORD  
 YES

 If you change the proxy password in the proxy instance, wait for at least 30 seconds for the new password to take effect, before updating the new password in the **Edit Repository** page.

### Upload Files

Upload certificate and private key to configure repository

valid domain name matching the provided certificate

HTTPS PROXY HOST

proxy host IP address or DNS name, e.g. myproxy.com

PORT

proxy port, allowed value range is from 1 to 65535

PROXY REQUIRES PASSWORD

YES

\* USER NAME

username to be used to authenticate to proxy

\* PASSWORD

password to be used to authenticate to proxy

# Backup and Restore

## Backup and Restore

- [With Internet Access](#)
  - [Backup](#)
  - [Restore](#)
    - [Restore without Proxy](#)
    - [Restore with Proxy](#)
- [Without Internet Access](#)

# With Internet Access

## With Internet Access

- [Backup](#)
- [Restore](#)
  - [Restore without Proxy](#)
  - [Restore with Proxy](#)





# Backup

## Backup Approach

- [Overview](#)
- [Limitations](#)
- [What Data Is Backed Up?](#)
- [Requirements](#)
- [Process](#)
- [Actions after Configuring the Backup](#)


You may sometimes need to backup your CloudCenter Suite setup so you have the option to recover the data when required. When you have a cluster running, it can go into a bad state for a number of reasons (resource shortage, application unavailability, infrastructure changes, undependable state and so forth). In these cases, backing up the data allows you a to recover data when required.

 The backup/restore feature is only available on *new* CloudCenter Suite clusters installed using CloudCenter Suite installers and *not on* existing Kubernetes clusters.


 For isolated, air gap, environments, that do not have internet access, or to back up to a local system, a manual backup procedure is available – see [Without Internet Access](#) for additional details.

Before proceeding with a backup, adhere to the following limitations:

- **Supported Clouds:** You can backup data to one of the following locations:
  - Google Cloud Storage (use the procedure below)
  - AWS S3 (use the procedure below)
- **Switching between Clouds and Cloud Accounts:**
  - While editing the storage location in the CloudCenter Suite, if you switch to a new cloud type or cloud account within the same cloud type, be aware that backups in the previously configured storage location will no longer be accessible from the CloudCenter Suite.
  - The backup files from the previously configured storage location will continue to be available via your cloud console.
- **Restoring to a Different Cluster:**
  - This feature is only supported for clusters launched by the CloudCenter Suite installer.
  - You cannot backup from and restore to the same cluster – you **can only** backup to one cluster and restore to a different cluster.
  - The backed up cluster and the target restore cluster should both be on the same cloud.
- **User Credentials:**
  - The credentials are specific to your service account in the cloud and only the user with those credentials can configure and initiate the backup.
  - If you change the credentials you will see a warning message to indicate that you cannot access previous backups.

 The CloudCenter Suite does NOT provide a granular option to backup Kubernetes resources or application-specific databases. Additionally, you CANNOT take volume snapshots.

The CloudCenter Suite uses the *latest* cloud/cloud account and bucket configurations to retrieve the list of existing backups, displayed in the table in the **Admin > Backup** page (under the Data Recovery section in the Suite Admin UI).

 If you update the existing configuration for any reason, users cannot manage the backups from the earlier cloud/cloud account and bucket configuration.

The backup action backs up the ENTIRE *cisco* namespace.

- **Backed Up:** Any data under the Cisco (*cisco*) name space. This includes but is not restricted to the Kubernetes resources with associated application data, pod data, secrets, PersistentVolumeClaim (PVC) data, PersistentVolume (PV) data, and other relevant data associated with these sub-systems
- **Not Backed Up:** Any data that is not under the Cisco (*cisco*) name space.

Before proceeding with a backup, adhere to the following limitations:

- **General:** When configuring a backup for the first time, verify that the storage bucket is empty before scheduling any backups.
- **GCP:**
  - Configure a Storage Bucket with the required permissions: The following screenshot displays a sample storage bucket in a GCP environment:

<input type="checkbox"/>	Name	Size	Type	Storage class	Last modified	Public access
<input type="checkbox"/>	ab1-backup-20190723/	—	Folder	—	—	Per object
<input type="checkbox"/>	ab2-backup-20190723/	—	Folder	—	—	Per object
<input type="checkbox"/>	backup-20190724/	—	Folder	—	—	Per object
<input type="checkbox"/>	backup-2019072402/	—	Folder	—	—	Per object

- The cloud account used to configure the backup must have an empty **storage.bucket.list**.
  - The bucket must have its ACL set to **storage.objects(create,delete,get,list)**.
- **AWS:**
    - The storage bucket in your AWS S3 environment must be empty with the applicable ACL permission.
    - The IAM user permissions define the user privilege on the S3 bucket as listed in the following screenshot:



In the following code block, the bucket name is defined as *velero-cisco*— this is just an example! Be sure to change this value to reflect the name of your own bucket!

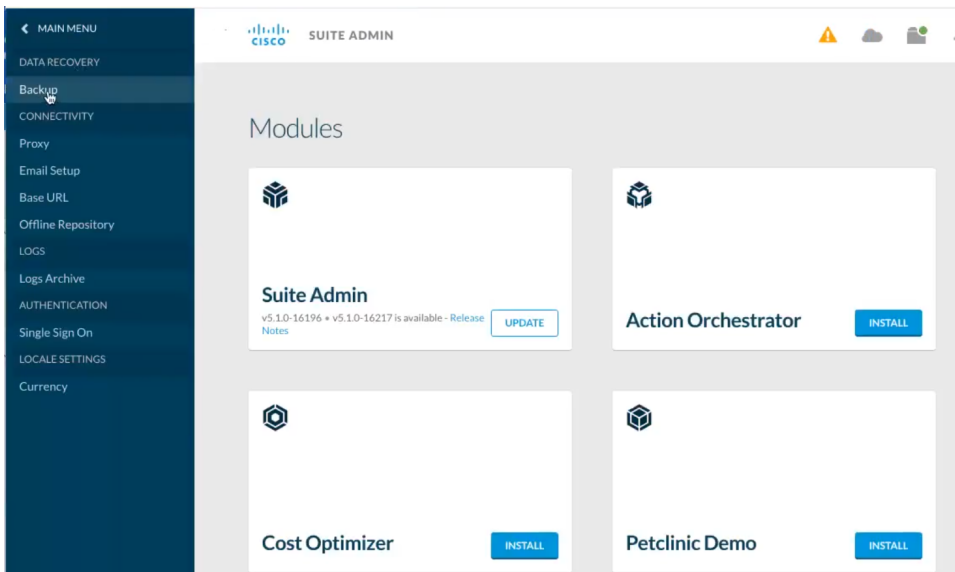
```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:DescribeRegions",
        "ec2:DescribeVolumes",
        "ec2:DescribeSnapshots",
        "ec2:CreateTags",
        "ec2:CreateVolume",
        "ec2:CreateSnapshot",
        "ec2>DeleteSnapshot"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "s3:GetObject",
        "s3:DeleteObject",
        "s3:PutObject",
        "s3:AbortMultipartUpload",
        "s3:ListMultipartUploadParts"
      ],
      "Resource": [
        "arn:aws:s3:::velero-cisco/*"
      ]
    },
    {
      "Effect": "Allow",
      "Action": [
        "s3:ListBucket"
      ],
      "Resource": [
        "arn:aws:s3:::velero-cisco"
      ]
    },
    {
      "Effect": "Allow",
      "Action": "s3:ListAllMyBuckets",
      "Resource": [
        "arn:aws:s3:::*"
      ]
    }
  ]
}


```

To backup the CloudCenter Suite data, follow this procedure.


1. Navigate to the [Suite Admin Dashboard](#).
2. Click **Admin > Backup** (under the Data Recovery section) to access the Backup page as displayed in the following screenshot.



3. Click the  **cog**  icon in the Backup page (as displayed in the following screenshot) to configure a new backup storage location.



4. Select the required cloud in the Configure a Backup Storage Location page as displayed in the following screenshot.

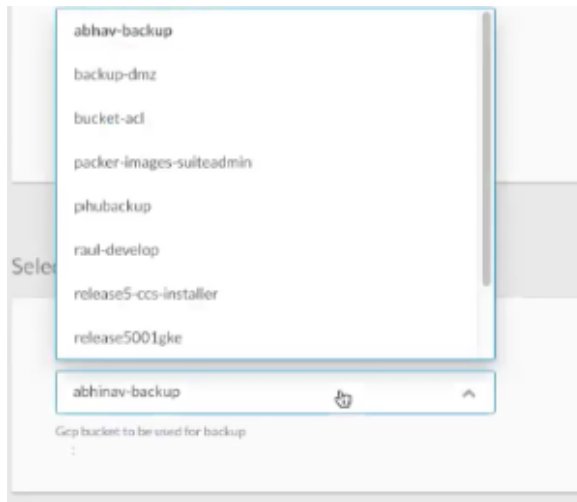


5. Depending on the selected cloud, the Add Credential section differs:


- GCP:
  - a. Select the file containing the credentials is displayed in the following screenshot.



b. Select the Storage bucket as displayed in the following screenshot.




- c. Click **Done** to save the backup configuration as displayed in the following screenshot.




- AWS S3:

- a. Select the file containing the credentials as displayed in the following screenshot.




- b. Select the **Storage bucket** as displayed in the following screenshot.




- c. Click **Done** to save the backup configuration as displayed in the following screenshot.

Select Cloud



Google Cloud Storage



AWS S3

Add Credential

\* AWS ACCESS KEY ID

AWS Access Key ID

\* AWS SECRET ACCESS KEY

AWS Secret Access Key

\* AWS REGION

AWS Region (optional)

AWS S3 ENDPOINT

AWS S3 Endpoint (optional)

✔ Connected

6. Once configured, click **Backup** in the Backup page to initiate the data backup. Until you initiate the first backup, this page will be empty. Once you have initiated one or more backups, they are automatically listed in this page as visible in the following screenshot.

CISCO SUITE ADMIN ⚠️ ☁️ 📄 🔔 ⚙️ | AC Welcome, Admin ▾

### Backup

NAME	CREATED DATE	CREATED BY	LOCATION	ACTIONS
ab1-backup-20190723	2 days ago	Admin Cllgitech	gcp > abnav-backup	<input type="button" value="⌵"/>
ab2-backup-20190723	2 days ago	Admin Cllgitech	gcp > abnav-backup	<input type="button" value="⌵"/>

7. In the Backup Name popup, assign a unique name (by default, the current date is listed) for this backup task and click **OK** as displayed in the following screenshot.

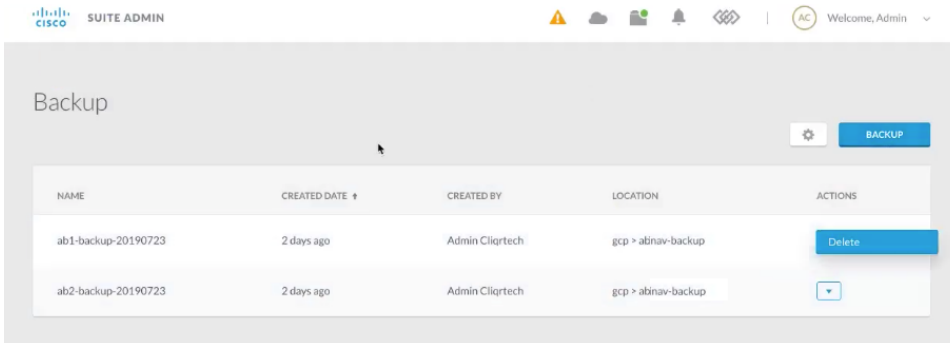
### Backup Name

\* BACKUP NAME

You have now backed up the CloudCenter Suite data to a cloud of choice.

Once you have configure one or more backup settings in the Backup page, you may see the following actions in the Actions column.

- **Delete:** You can delete the configured backup as visible in the following screenshot:



- **Cancel:** You will only see the **Cancel** option when you are in the process of backing up a storage location. After you create the location, the only option you will see is **Delete**.

**Back to:** [With Internet Access](#)

# Restore

## Restore

- [Restore without Proxy](#)
- [Restore with Proxy](#)

Back to: [With Internet Access](#)




# Restore without Proxy

## Restore without Proxy


- [Overview](#)
- [Limitations](#)
- [Requirements](#)
  - [1. Launch the Target Cluster](#)
  - [2. Download the KubeConfig Files](#)
  - [3. Download Velero](#)
  - [4. Download JQ](#)
  - [5. Pre-Restore Procedure](#)
  - [6. Restore Procedure](#)
  - [7. Post-Restore Procedure](#)
  - [8. Workload Manager-Specific Post-Restore Procedure](#)
    - [8a. Understand the Workload Manager Restore Context](#)
    - [8b. Retrieve the Port Numbers from the NEW Restored Cluster](#)
    - [8c. Retrieve the IP Address of the NEW Restored Cluster](#)
    - [8d. Change the IP Address and Port Numbers for the NEW Restored Cluster](#)
    - [8e. Perform the Pre-Migrate Activities](#)
    - [8f. Migrate Deployments from the OLD Cluster to the NEW Cluster](#)

To restore data, the CloudCenter Suite requires that you launch a new cluster.

 The backup/restore feature is only available on CloudCenter Suite clusters installed using CloudCenter Suite installers and not on existing Kubernetes clusters.

If you configured the old cluster using a DNS, be sure to update the new IP address (from the restored cluster) that is mapped to the DNS entry. Once you update the DNS entry of your new cluster, these services will continue to work as designed.

Additionally, be aware that you may need to update the DNS for the [Base URL Configuration](#) and [SSO Setup](#) (both ADFS and SP).

 Reconfiguration of Base URL and SSO are only applicable for backup & restore functions IF the source cluster is created using the CloudCenter Suite **5.0.x installer** and the destination cluster is freshly created using the CloudCenter Suite **5.1.1 installer**.


Before proceeding with a restore, adhere to the following limitations:

- The Velero tool must be installed. Velero Version 0.11.0 – refer to <https://velero.io/docs/v0.11.0/> for details.
- Launch a new cluster to restore the data.
- You will need to execute multiple scripts as part of these procedures. Make sure to use the 755 permission to execute each script mentioned in this section.


### 1. Launch the Target Cluster

To launch CloudCenter Suite on a new target cluster and access the Suite Admin UI for this cluster.

1. Navigate to the [Suite Admin Dashboard](#) for the new cluster.
2. Configure the identical backup configuration that you configured in your old cluster. See [Backup](#) > *Process* additional details. When you provide the credentials, the new cluster automatically connects to the cloud storage location.

 This step is REQUIRED to initiate the connection and fetch the backup(s).

3. Wait for a few minutes (at least 5 Mins, maybe more) for the Velero service in the new cluster to be synced up with the cloud storage location. At this point return to your local command window (shell console or terminal window) to perform the remaining steps in this process.

 If both your clusters are accessible from your local machine, the scripts used in the following steps can be executed as designed.

If either one of your clusters uses proxy access or if you cannot recover/download the KubeConfig file from your old cluster, follow the instructions provided in the [Restore with Proxy](#) section.

### 2. Download the KubeConfig Files

You must download the KubeConfig file from the Suite Admin Kubernetes cluster management page for your source and target clusters to your local machine via a local command window (shell console or terminal window):

- From the source cluster, download the KubeConfig file and name it **KUBECONFIG\_OLD**.
- From the target cluster, download the KubeConfig file and name it **KUBECONFIG\_NEW**.

See [Kubernetes Cluster Management](#) for additional details on accessing the KubeConfig file as displayed in the following screenshot.

NAME	IP ADDRESS	STATUS	CPU	MEMORY(GB)	RUNTIME
ab1473-8ca90e44-4e3f-4f2d-adf6-52a18809bf9a-mg-1-master-0	10.10.98.247	Up	2	16.82	1h
ab1473-8ca90e44-4e3f-4f2d-adf6-52a18809bf9a-mg-1-master-1	10.10.97.97	Up	2	16.82	1h
ab1473-8ca90e44-4e3f-4f2d-adf6-52a18809bf9a-mg-1-master-2	10.10.99.7	Up	2	16.82	1h

### 3. Download Velero

The restore process requires Velero and must be performed on a local command window (shell console or terminal window).

To download Velero, use one of the following options:

- OSX option:

```
$ cd <VELERO_DIRECTORY>
$ curl -L -O https://github.com/heptio/velero/releases/download/v0.11.0/velero-v0.11.0-darwin-amd64.tar.gz
$ tar -xvf velero-v0.11.0-darwin-amd64.tar.gz
```

- CentOS Option:

```
$ mkdir -p /velero-test && cd /velero-test
$ curl -LO https://github.com/heptio/velero/releases/download/v0.11.0/velero-v0.11.0-linux-amd64.tar.gz
$ tar -xvf velero-v0.11.0-linux-amd64.tar.gz && rm -rf velero-v0.11.0-linux-amd64.tar.gz
$ cp /velero-test/velero /usr/local/bin/
```

After you download Velero, export the KubeConfig file of the target (restore) cluster using the downloaded file:

```
export KUBECONFIG=<KUBECONFIG_PATH>
```

### 4. Download JQ

The restore process requires that you install JQ on your machine. Refer to <https://stedolan.github.io/jq/download> for additional details.

```
# To install jq on MacOS
$ brew install jq

# To install jq on Debian and Ubuntu
$ sudo apt-get install jq

# To install jq on CentOS
$ sudo yum install epel-release -y
$ sudo yum install jq -y
$ sudo jq --version
```

## 5. Pre-Restore Procedure

The pre-restore script creates the storageclass, if it does not exist on destination cluster, and saves the nginx-ingress-controller YAML file as well as the config maps for the following Suite Admin services:

- The suite-k8 service
- The suite-prod service

To execute the pre-restore script, run the **pre-restore.sh** script with the provided parameters:

```
# Command to execute the bashscript
$ ./pre-restore.sh <ccs_installer_version> </pathTo/oldCluster/kube_config> </pathTo/targetCluster/kube_config>

#<ccs_installer_version> is the CloudCenter Suite version without any characters inbetween. For example, "510"
or"502"or"5101"
#</pathTo/oldCluster/kube_config> is the path to the OLD KubeConfig file downloaded in Step 2.
#</pathTo/targetCluster/kube_config> is the path to the NEW KubeConfig file downloaded in Step 2.
```



Make sure that the backup folder does not exist at `~/backup` on the device in which you are execute these scripts. If a `~/backup` exists, delete it using the following command:

```
rm -rf ~/backup
```

The following code block includes the **pre-restore.sh** script:

```
#!/bin/bash
INSTALLER_VERSION_OLD=$1
KUBECONFIG_OLD=$2
KUBECONFIG_NEW=$3

declare INSTALLER_STORAGECLASS
INSTALLER_STORAGECLASS["500"]="thin"
INSTALLER_STORAGECLASS["501"]="thin"
INSTALLER_STORAGECLASS["502"]="thin"
INSTALLER_STORAGECLASS["51"]="standard"
INSTALLER_STORAGECLASS["510"]="standard"

if [[ ( ($KUBECONFIG_OLD == "" && $INSTALLER_VERSION_OLD == "" ) || $KUBECONFIG_NEW == "" ) ]]; then
    echo "Missing Paths for kubeconfigs"
    echo "Quitting"
    exit 0
else
    export KUBECONFIG_SAVED=$KUBECONFIG
    export KUBECONFIG=$HOME/.kube/config

    mkdir $HOME/backup
    cp $HOME/.kube/config $HOME/backup/saved_config

    if [[ $KUBECONFIG_OLD != "" ]]; then

        # Fetching the storage class name for the old(backup) cluster and storing it in variable
```

```

STORAGECLASS_NAME_OLD
cp $KUBECONFIG_OLD $HOME/.kube/config
STORAGECLASS_NAME_OLD=$(kubectl get storageclass -o json | jq '.items[0].metadata.name' | sed -e 's/^"
//' -e 's/"$//') # Extracting the storage class name from the json file of old cluster
echo "Creating storage class "${STORAGECLASS_NAME_OLD} "in the target cluster."

else
echo "Creating storage class "${INSTALLER_STORAGECLASS[$INSTALLER_VERSION_OLD]} "in the target cluster."
STORAGECLASS_NAME_OLD=${INSTALLER_STORAGECLASS[$INSTALLER_VERSION_OLD]}
fi

# Creating a storage class with the name STORAGECLASS_NAME_OLD in the target(restore) cluster
cp $KUBECONFIG_NEW $HOME/.kube/config
kubectl get storageclass -o json | jq --arg inpl $STORAGECLASS_NAME_OLD '.items[0].metadata.name=$inpl' >
$HOME/backup/storageclass.json
cat $HOME/backup/storageclass.json | kubectl create -f -

#Scripts to backup ingress service spec, k8s and prod-mgmt configmaps on the target cluster
mkdir -p $HOME/backup/configmap
mkdir -p $HOME/backup/service
mkdir -p $HOME/backup/sshkeys

kubectl get svc -n cisco common-framework-nginx-ingress-controller -o json > $HOME/backup/service/ingress.
json

for cm in $(kubectl get configmaps -n cisco -o custom-columns=:metadata.name --no-headers=true | grep "k8s-
mgmt")
do
kubectl get configmap $cm -n cisco -o yaml > $HOME/backup/configmap/$cm
done

for cm in $(kubectl get configmaps -n cisco -o custom-columns=:metadata.name --no-headers=true | grep "prod-
mgmt")
do
kubectl get configmap $cm -n cisco -o yaml > $HOME/backup/configmap/$cm
done

kubectl get configmap suite.key -n cisco -o yaml > $HOME/backup/sshkeys/suite.key
kubectl get configmap suite.pub -n cisco -o yaml > $HOME/backup/sshkeys/suite.pub

cp $HOME/backup/saved_config $HOME/.kube/config

export KUBECONFIG=$KUBECONFIG_SAVED

fi

echo 'Successfull!'

```

## 6. Restore Procedure

To restore the backed up data to the target cluster, run the following Velero commands from your local machine.

1. List available backups.

```
$ ./<VELERO_DIRECTORY>/velero backup get
```



Verify if the backups are listed BEFORE proceeding to the next step.

2. Make sure the backed up *cisco* namespace does not exist in the target cluster. Be sure to delete the *cisco* name space, if it exists, before you restore.

```
$ kubectl delete ns cisco
```

- Restore from one of the listed backups.

```
$ ./velero restore create --from-backup <BACKUPNAME>
```

You have now restored the CloudCenter Suite data to the new cluster.

## 7. Post-Restore Procedure

At this stage, you must restore the config maps for the following Suite Admin services:

- The suite-k8 service
- The suite-prod service

If the new cluster is accessible (from the local device) using the KubeConfig file, execute the following post-restore.sh script.

### With Internet Access - The post-restore.sh script

```
#!/bin/bash

KUBECONFIG_NEW=$1

if [[ ( $KUBECONFIG_NEW == "" ) ]]; then
    echo "Missing Paths for kubeconfig"
    echo "Quitting"
    exit 0
else
    export KUBECONFIG_SAVED=$KUBECONFIG
    export KUBECONFIG=$HOME/.kube/config

    cp $HOME/.kube/config $HOME/backup/saved_config
    cp $KUBECONFIG_NEW $HOME/.kube/config

    kubectl delete svc -n cisco common-framework-nginx-ingress-controller
    cat $HOME/backup/service/ingress.json | kubectl create -f -

    for cm in $(ls $HOME/backup/configmap)
    do
        kubectl delete configmap $cm -n cisco
    done

    for cm in $(ls $HOME/backup/configmap)
    do
        cat $HOME/backup/configmap/$cm | kubectl create -f -
    done


    kubectl delete configmap suite.key -n cisco
    kubectl delete configmap suite.pub -n cisco
    cat $HOME/backup/sshkeys/suite.key | kubectl create -f -
    cat $HOME/backup/sshkeys/suite.pub | kubectl create -f -

    cp $HOME/backup/saved_config $HOME/.kube/config
    export KUBECONFIG=$KUBECONFIG_SAVED


    rm -r $HOME/backup/
fi

echo 'Successfull!'
```

## 8. Workload Manager-Specific Post-Restore Procedure

 This migration procedure only applies to **Running** deployments.

Be sure to verify that you are only migrating deployment in the **Running** state.


 The first few steps differ based on your use of private clouds or public clouds. Be sure to use the procedure applicable to your cloud environment.

## 8a. Understand the Workload Manager Restore Context

If you have installed the Workload Manager module, you must perform this procedure to update the DNS/IP address for the private cloud resources listed below and displayed in the following image:


- The Worker AMQP IP
- The Guacamole Public IP and Port
- The Guacamole IP Address and Port for Application VMs

Cloud endpoint accessible from CloudCenter Suite	Yes
CloudCenter Suite AMQP reachable from worker VM's	Yes
CloudCenter Suite AMQP accessible from cloud	Yes
Remote AMQP IP	
Worker AMQP IP	10.8.1.140:26642
Guacamole Public IP and Port	10.8.1.140:708
Guacamole IP Address and Port for Application VMs	10.8.1.140:32941
Blade Name	cloudcenter-blade-vmware-1-2033

 As public clouds use load balancers and static IP ports, these resource details may differ accordingly. Be sure to use the resources applicable to your cloud environment.

## 8b. Retrieve the Port Numbers from the NEW Restored Cluster

The Kubernetes cluster contains the information that is required to update the Workload Manager UI. This section provides the commands required to retrieve this information.

 As public clouds use load balancers and static IP ports, these resource details may differ accordingly. Be sure to use the resources applicable to your cloud environment.

To retrieve the port numbers from the new cluster for private clouds, follow this procedure.

1. The port numbers for each component will differ.
  - a. Run the following command on the new cluster (login to the KubeConfig of the new cluster) to locate the new port numbers for the **Worker AMQP IP**.

```
kubectl get service -n cisco | grep rabbitmq-ext | awk '{print $5}'

# In the resulting response, locate the port corresponding to Port 443 and use that port number!

443:26642/TCP,15672:8902/TCP
```

- b. Run the following command on the new cluster to retrieve the port number for the **Guacamole Public IP and Port**.

```
kubectl get service -n cisco | grep cloudcenter-guacamole | awk '{print $5}'

# In the resulting response, locate the port corresponding to Port 443 and use that port number
for the Guacamole port!

8080:2376/TCP,7788:25226/TCP,7789:32941/TCP,443:708/TCP
```

- c. Run the following command on the new cluster to retrieve the port number for the **Guacamole IP Address and Port for Application VMs**.


```
kubectl get service -n cisco | grep cloudcenter-guacamole | awk '{print $5}'

# In the resulting response, locate the port corresponding to Port 7789 and use that port number
for the Guacamole port!

8080:2376/TCP,7788:25226/TCP,7789:32941/TCP,443:708/TCP
```


### 8c. Retrieve the IP Address of the NEW Restored Cluster

Use the IP address of one of the masters of the NEW restored Kubernetes cluster for all the resources where the IP address needs to be replaced.

 As public clouds use load balancers and static IP ports, these resource details may differ accordingly. Be sure to use the resources applicable to your cloud environment.

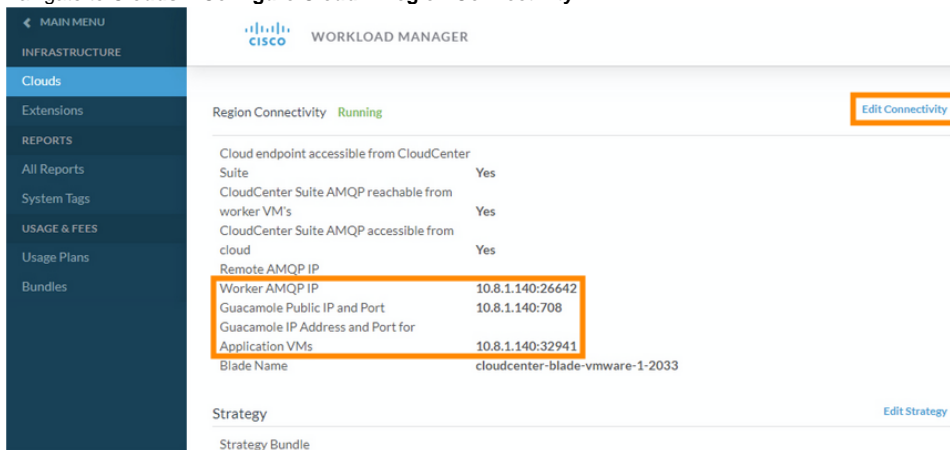
### 8d. Change the IP Address and Port Numbers for the NEW Restored Cluster

The IP addresses and port numbers are not updated automatically in the Workload Manager UI and you must explicitly update them using this procedure.

 As public clouds use load balancers and static IP ports, these resource details may differ accordingly. Be sure to use the resources applicable to your cloud environment.

To configure the IP address and port number in the new cluster, follow this procedure.


1. Access the Workload Manager module.
2. Navigate to **Clouds > Configure Cloud > Region Connectivity**.



The screenshot shows the 'Region Connectivity' settings in the Workload Manager UI. The 'Worker AMQP IP' field is highlighted with an orange box, showing the value 10.8.1.140:26642. Other fields include 'Guacamole Public IP and Port' (10.8.1.140:708) and 'Application VMs' (10.8.1.140:32941). An 'Edit Connectivity' button is also visible.

Cloud endpoint accessible from CloudCenter Suite	Yes
CloudCenter Suite AMQP reachable from worker VM's	Yes
CloudCenter Suite AMQP accessible from cloud	Yes
Remote AMQP IP	
Worker AMQP IP	10.8.1.140:26642
Guacamole Public IP and Port	10.8.1.140:708
Guacamole IP Address and Port for Application VMs	10.8.1.140:32941
Blade Name	cloudcenter-blade-vmware-1-2033

3. Click **Edit Connectivity** in the Region Connectivity settings.
4. In the Configure Region popup, change the 3 fields mentioned above to ensure that the IP and port details are updated to the NEW restored VM.

 DO NOT MAKE ANY OTHER CONFIGURATION CHANGES!

Configure Region
✕

IS CLOUD END POINT DIRECTLY ACCESSIBLE?

YES

SHOULD WORKER VMS DIRECTLY CONNECT WITH CLOUDCENTER SUITE?

YES

WORKER AMQP IP ADDRESS

GUACAMOLE PUBLIC IP AND PORT

GUACAMOLE IP ADDRESS AND PORT FOR APPLICATION VMS

- Click **OK** to save your changes.



Saving your changes may not automatically update the information in the Region Connectivity settings. Be sure to refresh the page to see the saved information.

You have now updated the DNS/IP/Port for the restored WM for this particular cloud. If you have configured other clouds in this environment, be sure to repeat this procedure for each cloud. Once you complete this procedure for all configured clouds, you can resume new deployment activities using the Workload Manager.

## 8e. Perform the Pre-Migrate Activities

Before you migrate the deployment details you need to ensure that you can connect to both clusters and have the required files to perform the migration.

To perform the pre-migrate activities, follow this procedure.

- Verify that the OLD cluster VMs can reach the NEW cluster. The remaining steps in this procedure are dependent on this connectivity in your environment.
- Save the contents of the following **actions.json** file using the same name and file extension to your local directory with a file type JSON format.

### The actions.json file

```
{
  "repositories": [],
  "actions": {
    "resource": null,
    "size": 2,
    "pageNumber": 0,
    "totalElements": 2,
    "totalPages": 1,
    "actionJaxbs": [
      {
        "id": "57",
        "resource": null,
        "name": "AgentReConfig_Linux",
        "description": "",
        "actionType": "EXECUTE_COMMAND",
        "category": "ON_DEMAND",
        "lastUpdatedTime": "2019-09-19 22:14:54.245",
        "timeOut": 1200,
        "enabled": true,
        "encrypted": false,
        "explicitShare": false,
        "showExplicitShareFeature": false,
        "deleted": false,
        "systemDefined": false,
        "bulkOperationSupported": true,
        "isAvailableToUser": true,
        "currentlyExecuting": false,
        "owner": 1,
        "actionParameters": [
          {
            "paramName": "downloadFromBundle",
            "paramValue": "true",
            "customParam": false,
            "required": true,
            "useDefault": false,
            "preference": "VISIBLE_UNLOCKED"
          },
          {
            "paramName": "bundlePath",
            "paramValue": "http://10.0.0.3/5.1-release/ccs-bundle-artifacts-5.1.0-20190819/agent.zip",
            "customParam": false,
            "required": true,
            "useDefault": false,
            "preference": "VISIBLE_UNLOCKED"
          },
          {
            "paramName": "script",
            "paramValue": "agent/agentReconfig.sh",
            "customParam": false,
            "required": true,
            "useDefault": false,
            "preference": "VISIBLE_UNLOCKED"
          },
          {
            "paramName": "executeOnContainer",
            "paramValue": "false",
            "customParam": false,
            "required": true,
            "useDefault": false,
            "preference": "VISIBLE_UNLOCKED"
          },
          {
            "paramName": "rebootInstance",
            "paramValue": "false",
            "customParam": false,
            "required": true,
            "useDefault": false,
            "preference": "VISIBLE_UNLOCKED"
          }
        ],
        "actionResourceMappings": [
          {
            "type": "VIRTUAL_MACHINE",
            "actionResourceFilters": [
              {
                "cloudRegionResource": null,
                "serviceResource": null,
                "applicationProfileResource": null,
                "deploymentResource": null,
                "vmResource": {
                  "type": "DEPLOYMENT_VM",
                  "appProfiles": [
                    "all"
                  ],
                  "cloudRegions": [
                    "all"
                  ],
                  "services": [
                    "all"
                  ],
                  "osTypes": [],
                  "cloudFamilyNames": [],
                  "nodeStates": [],
                  "cloudResourceMappings": []
                },
                "isEditable": true
              }
            ],
            "actionResource": null,
            "serviceResource": null,
            "applicationProfileResource": null,
            "deploymentResource": null,
            "vmResource": {
              "type": "IMPORTED_VM",
              "appProfiles": [],
              "cloudRegions": [
                "all"
              ],
              "cloudAccounts": [
                "all"
              ],
              "services": [],
              "osTypes": [
                "all"
              ],
              "cloudFamilyNames": [],
              "nodeStates": []
            }
          }
        ]
      }
    ]
  }
}
```




```

cloudResourceMappings": [], "isEditable": true}}], "actionResourceMappingAncillaries": [], "
actionCustomParamSpecs": [{"paramName": "brokerHost", "displayName": "BrokerHost", "helpText": "Ip Address or
HostName of Rabbit MQ cluster", "type": "string", "valueList": null, "defaultValue": "", "confirmValue": "", "
pathSuffixValue": "", "userVisible": true, "userEditable": true, "systemParam": false, "exampleValue": null, "
dataUnit": null, "optional": false, "deploymentParam": false, "multiselectSupported": false, "useDefault": true, "
valueConstraint": {"minValue": 0, "maxValue": 255, "maxLength": 255, "regex": null, "allowSpaces": true, "
sizeValue": 0, "step": 0, "calloutWorkflowName": null}, "scope": null, "webserviceListParams": {"url": "", "
protocol": "", "username": "", "password": "", "requestType": null, "contentType": null, "commandParams": null, "
requestBody": null, "resultString": null}, "secret": null, "tabularTypeData": null, "collectionList": [], "
preference": "VISIBLE_UNLOCKED"}], {"paramName": "brokerPort", "displayName": "BrokerPort", "helpText": "
RabbitMQ Port number", "type": "string", "valueList": null, "defaultValue": "", "confirmValue": "", "
pathSuffixValue": "", "userVisible": true, "userEditable": true, "systemParam": false, "exampleValue": null, "
dataUnit": null, "optional": false, "deploymentParam": false, "multiselectSupported": false, "useDefault": true, "
valueConstraint": {"minValue": 0, "maxValue": 255, "maxLength": 255, "regex": null, "allowSpaces": true, "
sizeValue": 0, "step": 0, "calloutWorkflowName": null}, "scope": null, "webserviceListParams": {"url": "", "
protocol": "", "username": "", "password": "", "requestType": null, "contentType": null, "commandParams": null, "
requestBody": null, "resultString": null}, "secret": null, "tabularTypeData": null, "collectionList": [], "
preference": "VISIBLE_UNLOCKED"}], {"id": "58", "resource": null, "name": "AgentReConfig_Win", "
description": "", "actionType": "EXECUTE_COMMAND", "category": "ON_DEMAND", "lastUpdatedTime": "2019-09-19 22:
15:02.311", "timeout": 1200, "enabled": true, "encrypted": false, "explicitShare": false, "
showExplicitShareFeature": false, "deleted": false, "systemDefined": false, "bulkOperationSupported": true, "
isAvailableToUser": true, "currentlyExecuting": false, "owner": 1, "actionParameters": [{"paramName": "
downloadFromBundle", "paramValue": "true", "customParam": false, "required": true, "useDefault": false, "
preference": "VISIBLE_UNLOCKED"}, {"paramName": "bundlePath", "paramValue": "http://10.0.0.3/5.1-release/ccs-
bundle-artifacts-5.1.0-20190819/agent.zip", "customParam": false, "required": true, "useDefault": false, "
preference": "VISIBLE_UNLOCKED"}, {"paramName": "script", "paramValue": "agent\\agentReconfig.ps1", "
customParam": false, "required": true, "useDefault": false, "preference": "VISIBLE_UNLOCKED"}, {"paramName": "
executeOnContainer", "paramValue": "false", "customParam": false, "required": true, "useDefault": false, "
preference": "VISIBLE_UNLOCKED"}, {"paramName": "rebootInstance", "paramValue": "false", "customParam": false, "
required": true, "useDefault": false, "preference": "VISIBLE_UNLOCKED"}, {"paramName": "refreshInstanceInfo", "
paramValue": "false", "customParam": false, "required": true, "useDefault": false, "preference": "
VISIBLE_UNLOCKED"}], "actionResourceMappings": [{"type": "VIRTUAL_MACHINE", "actionResourceFilters":
[{"cloudRegionResource": null, "serviceResource": null, "applicationProfileResource": null, "
deploymentResource": null, "vmResource": {"type": "DEPLOYMENT_VM", "appProfiles": ["all"], "cloudRegions":
["all"], "cloudAccounts": ["all"], "services": ["all"], "osTypes": [], "cloudFamilyNames": [], "nodeStates": [], "
cloudResourceMappings": []}, {"type": "IMPORTED_VM", "appProfiles": [], "cloudRegions": ["all"], "cloudAccounts": ["all"], "services": [], "osTypes": ["all"], "
cloudFamilyNames": [], "nodeStates": [], "cloudResourceMappings": []}], "isEditable": true}}], "
actionResourceMappingAncillaries": [], "actionCustomParamSpecs": [{"paramName": "brokerHost", "displayName": "
BrokerHost", "helpText": "Ip Address or Hostname of Rabbit MQ cluster", "type": "string", "valueList": null, "
defaultValue": "", "confirmValue": "", "pathSuffixValue": "", "userVisible": true, "userEditable": true, "
systemParam": false, "exampleValue": null, "dataUnit": null, "optional": false, "deploymentParam": false, "
multiselectSupported": false, "useDefault": true, "valueConstraint": {"minValue": 0, "maxValue": 255, "maxLength":
255, "regex": null, "allowSpaces": true, "sizeValue": 0, "step": 0, "calloutWorkflowName": null}, "scope": null, "
webserviceListParams": {"url": "", "protocol": "", "username": "", "password": "", "requestType": null, "
contentType": null, "commandParams": null, "requestBody": null, "resultString": null}, "secret": null, "
tabularTypeData": null, "collectionList": [], "preference": "VISIBLE_UNLOCKED"}, {"paramName": "brokerPort", "
displayName": "BrokerPort", "helpText": "RabbitMQ Port number", "type": "string", "valueList": null, "
defaultValue": "", "confirmValue": "", "pathSuffixValue": "", "userVisible": true, "userEditable": true, "
systemParam": false, "exampleValue": null, "dataUnit": null, "optional": false, "deploymentParam": false, "
multiselectSupported": false, "useDefault": true, "valueConstraint": {"minValue": 0, "maxValue": 255, "maxLength":
255, "regex": null, "allowSpaces": true, "sizeValue": 0, "step": 0, "calloutWorkflowName": null}, "scope": null, "
webserviceListParams": {"url": "", "protocol": "", "username": "", "password": "", "requestType": null, "
contentType": null, "commandParams": null, "requestBody": null, "resultString": null}, "secret": null, "
tabularTypeData": null, "collectionList": [], "preference": "VISIBLE_UNLOCKED"}], "
repositoriesMappingRequired": false, "actionTypesCounts": [{"key": "EXECUTE_COMMAND", "value": "2"}]}

```

3. Access Workload Manager in your OLD cluster and navigate to the Actions Library page.
4. Import the actions.json file that you saved in Step 2 above. You should see two files (**AgentReconfig\_Linux** and **AgentReconfig\_Win**) as displayed in the following screenshot.



5. The files are disabled by default (OFF) – enable both files by toggling each switch to **ON**.
6. Save the following script to a file in your local directory and name it **agentReconfig.sh**. This is the file to use for Linux environments.

**The agentReconfig.sh file**

```

#!/bin/bash

#Write to system log as well as to terminal
logWrite()
{
    msg=$1
    echo "$(date) ${msg}"
    logger -t "OSMOSIX" "${msg}"
    return 0
}

logWrite "Starting agent migrate..."

env_file="/usr/local/osmosix/etc/userenv"
if [ -f $env_file ];
then
    logWrite "Source the userenv file..."
    . $env_file
fi

if [ -z $brokerHost ];
then
    logWrite "Broker Host / Rabbit Server Ip not passed as action parameter"
    exit 3;
fi

if [ -z $brokerPort ];
then
    logWrite "Broker Port / Rabbit Server Port not passed as action parameter"
    exit 4
fi

replaceUserdataValue() {
    key=$1
    value=$2

    if [ -z $key ] || [ -z $value ];
    then
        logWrite "Command line arguments missing to update user-data file, key: $key, value:$value"
        return
    fi

    user_data_file="/usr/local/agentlite/etc/user-data"
    if [ -f $user_data_file ];
    then
        json_content=`cat $user_data_file`
        old_value=`echo $json_content | awk -F $key '{print $2}' | awk -F \" '{print $3}'`
        sed -i 's@"$old_value"@'"$value"@g' $user_data_file
    fi
}

export AGENT_HOME="/usr/local/agentlite"

logWrite "Updating the user data file"
replaceUserdataValue "brokerClusterAddresses" "$brokerHost:$brokerPort"

logWrite "Updating config.json file"
sed -i '/AmqpAddress/c\    "AmqpAddress": "'${brokerHost}:${brokerPort}'"', '$AGENT_HOME/config/config.json'

cd $AGENT_HOME
echo "sleep 10" > execute.sh
echo "/usr/local/agentlite/bin/agent-stop.sh" >> execute.sh
echo "/usr/local/agentlite/bin/agent-start.sh" >> execute.sh
chmod a+x execute.sh

```

```
nohup bash execute.sh > /dev/null 2>&1 &

exit 0
```

7. Save the following script to a file in your local directory and name it **agentReconfig.ps1**. This is the file to use for Windows environments.

#### The agentReconfig.ps1 file

```
param (
    [string]$brokerHost = "$env:brokerHost",
    [string]$brokerPort = "$env:brokerPort"
)

$SERVICE_NAME = "AgentService"
$SYSTEM_DRIVE = (Get-WmiObject Win32_OperatingSystem).SystemDrive
. "$SYSTEM_DRIVE\temp\userenv.ps1"

if ($brokerHost -eq 0 -or $brokerHost -eq $null -or $brokerHost -eq "") {
    echo "Variable brokerHost not available in the env file"
    exit 1
}

if ($brokerPort -eq 0 -or $brokerPort -eq $null -or $brokerPort -eq "") {
    echo "Variable brokerPort not available in the env file"
    exit 2
}

$AGENTGO_PARENT_DIR = "$SYSTEM_DRIVE\opt"

echo "Check if AgentGo Parent directory exists. If not create it: '$AGENTGO_PARENT_DIR'"
if (-not (Test-Path $AGENTGO_PARENT_DIR)) {
    echo "Create $AGENTGO_PARENT_DIR..."
    mkdir $AGENTGO_PARENT_DIR
}
else {
    echo "$AGENTGO_PARENT_DIR already exists."
}

$AGENT_CONFIG="{0}\agentlite\config\config.json" -f $AGENTGO_PARENT_DIR
if (Test-Path $AGENT_CONFIG) {
    echo "Changing the config.json file with the new broker host $env:brokerHost and port $env:
brokerPort"
    $confJson = get-content $AGENT_CONFIG | out-string | convertfrom-json
    $confJson.AmqpAddress = "$($env:brokerHost): $($env:brokerPort)"
    $confJson | ConvertTo-Json | set-content $AGENT_CONFIG
}

$USER_DATA_FILE = "{0}\agentlite\etc\user-data" -f $AGENTGO_PARENT_DIR
if (Test-Path $USER_DATA_FILE) {
    echo "Changing user-data file with new broker host $env:brokerHost and port $env:brokerPort"
    $userDataJson = get-content $USER_DATA_FILE | out-string | convertfrom-json
    $userDataJson.brokerClusterAddresses = "$($env:brokerHost): $($env:brokerPort)"
    $userDataJson | ConvertTo-Json | set-content $USER_DATA_FILE
}

$AGENT_SERVICE_NAME = "AgentService"
echo "Stop-Service $AGENT_SERVICE_NAME" > $AGENTGO_PARENT_DIR\exec.ps1
echo "sleep 10" >> $AGENTGO_PARENT_DIR\exec.ps1
echo "Start-Service $AGENT_SERVICE_NAME" >> $AGENTGO_PARENT_DIR\exec.ps1

echo "Restarting agent"
Start-Process -filepath "powershell" -argumentlist "-executionpolicy bypass -noninteractive -file
`"$AGENTGO_PARENT_DIR\exec.ps1`""

echo "Agent set to restart after config changes"
```

8. Add these two files to a folder called **agent** (just an example) and compress the folder to create **agent.zip** with the same structure displayed here.

**agent**

**agentReconfig.ps1**

**agentReconfig.sh**

9. Move the **agent.zip** folder to an HTTP repository in your local environment that is accessible from the OLD and NEW clusters.



This procedure uses the following URL as an example:

**http://10.0.0.3/repo/agent.zip**

You have now ensured cluster connectivity and saved the required files for the migration procedure.

## 8f. Migrate Deployments from the OLD Cluster to the NEW Cluster

To migrate the deployment details from the old cluster to the new cluster, follow this procedure.

1. Navigate to the Workload Manager **Actions Library** page and edit the **AgentReconfig\_Linux** action. This procedure continues to use the Linux file going forward.
2. Scroll to the **Actions Definition** section and update the URL as displayed in the following screenshot.

Action Definition

\* EXECUTE FROM BUNDLE  
 YES

\* LOCATION      \* URL  
URL      http://10.0.0.3/repo/agent.zip

\* SCRIPT FROM BUNDLE  
agent/agentReconfig.sh



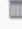


The URL and Script from Bundle fields in the above screenshot are in accordance with the steps above.

3. Scroll to the **Custom Fields** section and change the default value of the **Broker Host** to use the NEW cluster IP.

### 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

  **BrokerHost** 

\* DISPLAY NAME

\* PARAMETER NAME



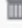
HELP TEXT

\* TYPE MAX LENGTH

DEFAULT VALUE

REQUIRED FIELD ?  
 YES

4. Scroll down to the **Broker Port** and change the default to use the NEW Worker AMQP IP port (for example, 26642 in Step 8 above).

  **BrokerPort** 

\* DISPLAY NAME

\* PARAMETER NAME

HELP TEXT

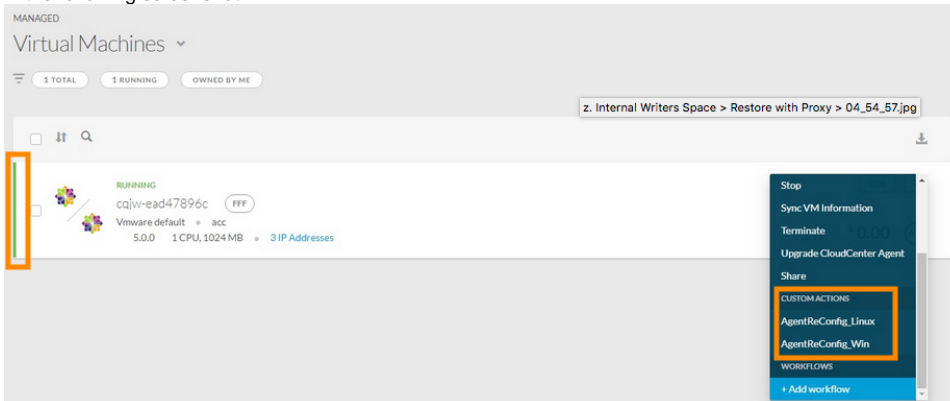
\* TYPE MAX LENGTH

DEFAULT VALUE

REQUIRED FIELD ?  
 YES

5. Click **Done** to save your default configuration changes in the OLD cluster.
6. Navigate to the **Virtual Machines** page and locate the VM to migrate to the new cluster.

7. Click the **Actions** dropdown and verify if your newly modified actions are visible under the Custom Actions section in the dropdown list as visible in the following screenshot.



8. Click one of the actions and verify that the configured defaults are displayed in the Broker host and Broker port fields as indicated earlier.
9. Click **Submit** to migrate this VM to the new cluster.
10. Verify that the migration is complete by going to the Deployment page in your NEW cluster and the VM is listed as RUNNING (green line).
11. Repeat Steps 6 through 10 for each VM that needs to be migrated to the NEW cluster.

You have now migrated the deployment details from the old cluster to the new cluster

**Back to:** [With Internet Access](#)

# Restore with Proxy

## Restore with Proxy

- [Overview](#)
- [Limitations](#)
- [Requirements](#)
  - [1. Launch the Target Cluster](#)
  - [2. Download the KubeConfig Files](#)
  - [3. Download Velero](#)
  - [4. Download JQ](#)
  - [5. Pre-Restore Procedure](#)
  - [6. Restore Procedure](#)
  - [7. Post-Restore Procedure](#)
  - [8. Workload Manager-Specific Post-Restore Procedure](#)
    - [8a. Understand the Workload Manager Restore Context](#)
    - [8b. Retrieve the Port Numbers from the NEW Restored Cluster](#)
    - [8c. Retrieve the IP Address of the NEW Restored Cluster](#)
    - [8d. Change the IP Address and Port Numbers for the NEW Restored Cluster](#)
    - [8e. Perform the Pre-Migrate Activities](#)
    - [8f. Migrate Deployments from the OLD Cluster to the NEW Cluster](#)

To restore data, the CloudCenter Suite requires that you launch a new cluster.



The backup/restore feature is only available on CloudCenter Suite clusters installed using CloudCenter Suite installers and not on existing Kubernetes clusters.

If you configured the old cluster using a DNS, be sure to update the new IP address (from the restored cluster) that is mapped to the DNS entry. Once you update the DNS entry of your new cluster, these services will continue to work as designed.

Additionally, be aware that you may need to update the DNS for the [Base URL Configuration](#) and [SSO Setup](#) (both ADFS and SP).



Reconfiguration of Base URL and SSO are only applicable for backup & restore functions IF the source cluster is created using the CloudCenter Suite **5.0.x installer** and the destination cluster is freshly created using the CloudCenter Suite **5.1.1 installer**.

Before proceeding with a restore, adhere to the following limitations:

- The Velero tool must be installed. Velero Version 0.11.0 – refer to <https://velero.io/docs/v0.11.0/> for details.
- Launch a new cluster to restore the data.
- You will need to execute multiple scripts as part of these procedures. Make sure to use the 755 permission to execute each script mentioned in this section.

### 1. Launch the Target Cluster

To launch CloudCenter Suite on a new target cluster and access the Suite Admin UI for this cluster.

1. Navigate to the [Suite Admin Dashboard](#) for the new cluster.
2. Configure the identical backup configuration that you configured in your old cluster. See [Backup](#) > *Process* additional details. When you provide the credentials, the new cluster automatically connects to the cloud storage location.



This step is REQUIRED to initiate the connection and fetch the backup(s).

3. Wait for a few minutes (at least 5 Mins, maybe more) for the Velero service in the new cluster to be synced up with the cloud storage location. At this point return to your local command window (shell console or terminal window) to perform the remaining steps in this process.



If both your clusters are accessible from your local machine, the scripts used in the following steps can be executed as designed.

If either one of your clusters uses proxy access or if you cannot recover/download the KubeConfig file from your old cluster, follow the instructions provided in the [Restore with Proxy](#) section.

### 2. Download the KubeConfig Files

You must download the KubeConfig file from the Suite Admin Kubernetes cluster management page for your source and target clusters to your local machine via a local command window (shell console or terminal window):

- From the source cluster, download the KubeConfig file and name it **KUBECONFIG\_OLD**.
- From the target cluster, download the KubeConfig file and name it **KUBECONFIG\_NEW**.

See [Kubernetes Cluster Management](#) for additional details on accessing the KubeConfig file as displayed in the following screenshot.

The screenshot shows the Cisco Suite Admin interface for a Kubernetes Cluster. The cluster is named 'Kubernetes Cluster' and has a version of v1.13.5 with 7 nodes installed on 05 Aug 2019. The cluster status is '7/7 Nodes Running'. A blue button labeled 'Download KubeConfig File' is highlighted with an orange box. Below the cluster information, there is a table listing the nodes:

NAME	IP ADDRESS	STATUS	CPU	MEMORY(GB)	RUNTIME
ab1473-8ca90e44-4e3f-4f2d-adf6-52a18809bf9a-mg-1-master-0	10.10.98.247	Up	2	16.82	1h
ab1473-8ca90e44-4e3f-4f2d-adf6-52a18809bf9a-mg-1-master-1	10.10.97.97	Up	2	16.82	1h
ab1473-8ca90e44-4e3f-4f2d-adf6-52a18809bf9a-mg-1-master-2	10.10.99.7	Up	2	16.82	1h

### 3. Download Velero

The restore process requires Velero and must be performed on a local command window (shell console or terminal window).

To download Velero, use one of the following options:

- OSX option:

```
$ cd <VELERO_DIRECTORY>
$ curl -L -O https://github.com/heptio/velero/releases/download/v0.11.0/velero-v0.11.0-darwin-amd64.tar.gz
$ tar -xvf velero-v0.11.0-darwin-amd64.tar.gz
```

- CentOS Option:

```
$ mkdir -p /velero-test && cd /velero-test
$ curl -LO https://github.com/heptio/velero/releases/download/v0.11.0/velero-v0.11.0-linux-amd64.tar.gz
$ tar -xvf velero-v0.11.0-linux-amd64.tar.gz && rm -rf velero-v0.11.0-linux-amd64.tar.gz
$ cp /velero-test/velero /usr/local/bin/
```

After you download Velero, export the KubeConfig file of the target (restore) cluster using the downloaded file:

```
export KUBECONFIG=<KUBECONFIG_PATH>
```

### 4. Download JQ

The restore process requires that you install JQ on your machine. Refer to <https://stedolan.github.io/jq/download> for additional details.

```
# To install jq on MacOS
$ brew install jq

# To install jq on Debian and Ubuntu
$ sudo apt-get install jq

# To install jq on CentOS
$ sudo yum install epel-release -y
$ sudo yum install jq -y
$ sudo jq --version
```



## 5. Pre-Restore Procedure

If either one of your clusters uses proxy access or if you cannot recover/download the KubeConfig file from your old cluster, follow the instructions provided in this section.

1. SSH into one of the VMs in your old cluster and retrieve the storageclass names.



This step is required because of changes in the storageclass name between CloudCenter Suite 5.0.0 and 5.1.0.

```
$ kubectl get storageclass -o json | grep '"name"' | cut -d ':' -f 2 | sed 's/"/\\"/g' | sed 's/[,]/ /g'
```

For example:

### Example

```
$ kubectl get storageclass -o json | grep '"name"' | cut -d ':' -f 2 | sed 's/"/\\"/g' | sed 's/[,]/ /g' "thin"
```

2. SSH into one of the VMs in your new cluster and retrieve the storageclass names:

```
$ kubectl get storageclass -o json | grep '"name"' | cut -d ':' -f 2 | sed 's/"/\\"/g' | sed 's/[,]/ /g'
```

For example:

### Example

```
$ kubectl get storageclass -o json | grep '"name"' | cut -d ':' -f 2 | sed 's/"/\\"/g' | sed 's/[,]/ /g'
"standard"
```

3. Copy the contents of storageclass from the new cluster using the command below: (use the storageclass\_name retrieve using the above step). You need to run the following command, copy the output, and save the output to a file called backupStorageclass.yaml.

```
$ kubectl get storageclass <storageclass_name> -o yaml
```

For example:

```

cloud-user@ab21461-fcc43751-1381-4e98-8d45-934bb965edfe-mg-1-master-0:~$ kubectl get storageclass
standard -o yaml
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  annotations:
    kubectl.kubernetes.io/last-applied-configuration: |
      {"apiVersion":"storage.k8s.io/v1beta1","kind":"StorageClass","metadata":{"annotations":
{"storageclass.beta.kubernetes.io/is-default-class":"true"},"name":"standard"},"parameters":
{"diskformat":"thin"},"provisioner":"kubernetes.io/vsphere-volume"}
      storageclass.beta.kubernetes.io/is-default-class: "true"
  creationTimestamp: "2019-07-31T23:26:57Z"
  name: standard
  resourceVersion: "605"
  selfLink: /apis/storage.k8s.io/v1/storageclasses/standard
  uid: b045d700-b3ea-11e9-9b1d-0050569f28fd
parameters:
  diskformat: thin
  provisioner: kubernetes.io/vsphere-volume
  reclaimPolicy: Delete
  volumeBindingMode: Immediate

```

4. Create a new file backupStorageclass.yaml and paste the contents copied from the previous step.
5. Replace the field **name** in the backupStorageclass.yaml file with the OLD storage\_classname from the old cluster from Step 1.

For example:

```

apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  annotations:
    kubectl.kubernetes.io/last-applied-configuration: |
      {"apiVersion":"storage.k8s.io/v1beta1","kind":"StorageClass","metadata":{"annotations":{"storageclass.beta.kubernetes.io/is-default-class":"t
storageclass.beta.kubernetes.io/is-default-class: "true"
  creationTimestamp: "2019-07-31T23:26:57Z"
→ name: thin
  resourceVersion: "605"
  selfLink: /apis/storage.k8s.io/v1/storageclasses/standard
  uid: b045d700-b3ea-11e9-9b1d-0050569f28fd
parameters:
  diskformat: thin
  provisioner: kubernetes.io/vsphere-volume
  reclaimPolicy: Delete
  volumeBindingMode: Immediate

```

6. Create a new storageclass in the new cluster using the command below

```
$ cat /path/backupStorageclass.yaml | kubectl create -f -
```

7. Create a backup of the Kubernetes config maps of the following services by executing the script provided in this step.
  - The suite-k8 service
  - The suite-prod service
8. Run the command to execute the backup\_configmap.sh script

```

#Execute the script as sudo user
$ sudo /path/to/script/backup_configmap.sh.sh

```

The backup\_configmap.sh script

**backup\_configmap.sh**

```
#!/bin/bash

#Scripts to backup k8s and prod-mgmt configmaps on the target cluster
mkdir -p $HOME/backup/configmap
mkdir -p $HOME/backup/service
mkdir -p $HOME/backup/sshkeys

kubectl get svc -n cisco common-framework-nginx-ingress-controller -o json > $HOME/backup/service/ingress.json

for cm in $(kubectl get configmaps -n cisco -o custom-columns=:metadata.name --no-headers=true | grep "k8s-mgmt")
do
    kubectl get configmap $cm -n cisco -o yaml > $HOME/backup/configmap/$cm
done

for cm in $(kubectl get configmaps -n cisco -o custom-columns=:metadata.name --no-headers=true | grep "prod-mgmt")
do
    kubectl get configmap $cm -n cisco -o yaml > $HOME/backup/configmap/$cm
done

kubectl get configmap suite.key -n cisco -o yaml > $HOME/backup/sshkeys/suite.key
kubectl get configmap suite.pub -n cisco -o yaml > $HOME/backup/sshkeys/suite.pub

echo 'Successfull!'
```

## 6. Restore Procedure

1. List available backups.



Verify if the backups are listed BEFORE proceeding to the next step.

```
$ ./<VELERO_DIRECTORY>/velero backup get
```

2. Make sure the backed up namespace does not exist in the target cluster (for example, if the *cisco* namespace was backed up it shouldn't be here on the cluster).

```
$ kubectl delete ns cisco
```

3. Restore from one of the listed backups.

```
$ ./velero restore create --from-backup <BACKUPNAME>
```

You have now restored the CloudCenter Suite data to the new cluster.

## 7. Post-Restore Procedure

At this stage, you must restore the config maps for the following Suite Admin services:

- The suite-k8 service
- The suite-prod service

If the new cluster is NOT accessible (from the local device) using kubeconfig, execute the following script from the remote device after the restore process is complete.

```
#Execute the script as sudo user
$ sudo /path/to/script/post-restore.sh
```

**Without Internet Access - The post-restore.sh script**

```
#!/bin/bash
kubectrl delete svc -n cisco common-framework-nginx-ingress-controller
cat $HOME/backup/service/ingress.json | kubectrl create -f -

for cm in $(ls $HOME/backup/configmap)
do
    kubectrl delete configmap $cm -n cisco
done

for cm in $(ls $HOME/backup/configmap)
do
    cat $HOME/backup/configmap/$cm | kubectrl create -f -
done

kubectrl delete configmap suite.key -n cisco
kubectrl delete configmap suite.pub -n cisco
cat $HOME/backup/sshkeys/suite.key | kubectrl create -f -
cat $HOME/backup/sshkeys/suite.pub | kubectrl create -f -

rm -r $HOME/backup/configmap

echo 'Successfull!'
```

You have now restored the Suite Admin data to the new cluster.

**8. Workload Manager-Specific Post-Restore Procedure**

This migration procedure only applies to **Running** deployments.

Be sure to verify that you are only migrating deployment in the **Running** state.



The first few steps differ based on your use of private clouds or public clouds. Be sure to use the procedure applicable to your cloud environment.

**8a. Understand the Workload Manager Restore Context**

If you have installed the Workload Manager module, you must perform this procedure to update the DNS/IP address for the private cloud resources listed below and displayed in the following image:

- The Worker AMQP IP
- The Guacamole Public IP and Port
- The Guacamole IP Address and Port for Application VMs

Cloud endpoint accessible from CloudCenter Suite	Yes
CloudCenter Suite AMQP reachable from worker VM's	Yes
CloudCenter Suite AMQP accessible from cloud	Yes
Remote AMQP IP	
Worker AMQP IP	10.8.1.140:26642
Guacamole Public IP and Port	10.8.1.140:708
Guacamole IP Address and Port for Application VMs	10.8.1.140:32941
Blade Name	cloudcenter-blade-vmware-1-2033



As public clouds use load balancers and static IP ports, these resource details may differ accordingly. Be sure to use the resources applicable to your cloud environment.

**8b. Retrieve the Port Numbers from the NEW Restored Cluster**

The Kubernetes cluster contains the information that is required to update the Workload Manager UI. This section provides the commands required to retrieve this information.



As public clouds use load balancers and static IP ports, these resource details may differ accordingly. Be sure to use the resources applicable to your cloud environment.

To retrieve the port numbers from the new cluster for private clouds, follow this procedure.

1. The port numbers for each component will differ.
  - a. Run the following command on the new cluster (login to the KubeConfig of the new cluster) to locate the new port numbers for the **Worker AMQP IP**.

```
kubectl get service -n cisco | grep rabbitmq-ext | awk '{print $5}'

# In the resulting response, locate the port corresponding to Port 443 and use that port number!

443:26642/TCP,15672:8902/TCP
```

- b. Run the following command on the new cluster to retrieve the port number for the **Guacamole Public IP and Port**.

```
kubectl get service -n cisco | grep cloudcenter-guacamole | awk '{print $5}'

# In the resulting response, locate the port corresponding to Port 443 and use that port number
for the Guacamole port!

8080:2376/TCP,7788:25226/TCP,7789:32941/TCP,443:708/TCP
```

- c. Run the following command on the new cluster to retrieve the port number for the **Guacamole IP Address and Port for Application VMs**.

```
kubectl get service -n cisco | grep cloudcenter-guacamole | awk '{print $5}'

# In the resulting response, locate the port corresponding to Port 7789 and use that port number
for the Guacamole port!

8080:2376/TCP,7788:25226/TCP,7789:32941/TCP,443:708/TCP
```

### 8c. Retrieve the IP Address of the NEW Restored Cluster

Use the IP address of one of the masters of the NEW restored Kubernetes cluster for all the resources where the IP address needs to be replaced.



As public clouds use load balancers and static IP ports, these resource details may differ accordingly. Be sure to use the resources applicable to your cloud environment.

### 8d. Change the IP Address and Port Numbers for the NEW Restored Cluster

The IP addresses and port numbers are not updated automatically in the Workload Manager UI and you must explicitly update them using this procedure.



As public clouds use load balancers and static IP ports, these resource details may differ accordingly. Be sure to use the resources applicable to your cloud environment.

To configure the IP address and port number in the new cluster, follow this procedure.

1. Access the Workload Manager module.

2. Navigate to **Clouds > Configure Cloud > Region Connectivity**.

Region Connectivity Running [Edit Connectivity](#)

Cloud endpoint accessible from CloudCenter Suite	Yes
CloudCenter Suite AMQP reachable from worker VM's	Yes
CloudCenter Suite AMQP accessible from cloud	Yes
Remote AMQP IP	
Worker AMQP IP	10.8.1.140:26642
Guacamole Public IP and Port	10.8.1.140:708
Guacamole IP Address and Port for Application VMs	10.8.1.140:32941
Blade Name	cloudcenter-blade-vmware-1-2033

Strategy [Edit Strategy](#)

Strategy Bundle

3. Click **Edit Connectivity** in the Region Connectivity settings.

## 4. In the Configure Region popup, change the 3 fields mentioned above to ensure that the IP and port details are updated to the NEW restored VM.

DO NOT MAKE ANY OTHER CONFIGURATION CHANGES!

Configure Region ✕

IS CLOUD END POINT DIRECTLY ACCESSIBLE?  
 YES

SHOULD WORKER VMS DIRECTLY CONNECT WITH CLOUDCENTER SUITE?  
 YES

WORKER AMQP IP ADDRESS  
10.8.1.140:26642

GUACAMOLE PUBLIC IP AND PORT  
10.8.1.140:708

GUACAMOLE IP ADDRESS AND PORT FOR APPLICATION VMS  
10.8.1.140:32941

5. Click **OK** to save your changes.

Saving your changes may not automatically update the information in the Region Connectivity settings. Be sure to refresh the page to see the saved information.

You have now updated the DNS/IP/Port for the restored WM for this particular cloud. If you have configured other clouds in this environment, be sure to repeat this procedure for each cloud. Once you complete this procedure for all configured clouds, you can resume new deployment activities using the Workload Manager.

## 8e. Perform the Pre-Migrate Activities

Before you migrate the deployment details you need to ensure that you can connect to both clusters and have the required files to perform the migration.

To perform the pre-migrate activities, follow this procedure.

1. Verify that the OLD cluster VMs can reach the NEW cluster. The remaining steps in this procedure are dependent on this connectivity in your environment.
2. Save the contents of the following **actions.json** file using the same name and file extension to your local directory with a file type JSON format.

## The actions.json file

```
{
  "repositories": [],
  "actions": {
    "resource": null,
    "size": 2,
    "pageNumber": 0,
    "totalElements": 2,
    "totalPages": 1,
    "actionJaxbs": [
      {
        "id": "57",
        "resource": null,
        "name": "AgentReConfig_Linux",
        "description": "",
        "actionType": "EXECUTE_COMMAND",
        "category": "ON_DEMAND",
        "lastUpdatedTime": "2019-09-19 22:14:54.245",
        "timeOut": 1200,
        "enabled": true,
        "encrypted": false,
        "explicitShare": false,
        "showExplicitShareFeature": false,
        "deleted": false,
        "systemDefined": false,
        "bulkOperationSupported": true,
        "isAvailableToUser": true,
        "currentlyExecuting": false,
        "owner": 1,
        "actionParameters": [
          {
            "paramName": "downloadFromBundle",
            "paramValue": "true",
            "customParam": false,
            "required": true,
            "useDefault": false,
            "preference": "VISIBLE_UNLOCKED"
          },
          {
            "paramName": "bundlePath",
            "paramValue": "http://10.0.0.3/5.1-release/ccs-bundle-artifacts-5.1.0-20190819/agent.zip",
            "customParam": false,
            "required": true,
            "useDefault": false,
            "preference": "VISIBLE_UNLOCKED"
          },
          {
            "paramName": "script",
            "paramValue": "agent/agentReconfig.sh",
            "customParam": false,
            "required": true,
            "useDefault": false,
            "preference": "VISIBLE_UNLOCKED"
          },
          {
            "paramName": "executeOnContainer",
            "paramValue": "false",
            "customParam": false,
            "required": true,
            "useDefault": false,
            "preference": "VISIBLE_UNLOCKED"
          },
          {
            "paramName": "rebootInstance",
            "paramValue": "false",
            "customParam": false,
            "required": true,
            "useDefault": false,
            "preference": "VISIBLE_UNLOCKED"
          },
          {
            "paramName": "refreshInstanceInfo",
            "paramValue": "false",
            "customParam": false,
            "required": true,
            "useDefault": false,
            "preference": "VISIBLE_UNLOCKED"
          }
        ],
        "actionResourceMappings": [
          {
            "type": "VIRTUAL_MACHINE",
            "actionResourceFilters": [
              {
                "cloudRegionResource": null,
                "serviceResource": null,
                "applicationProfileResource": null,
                "deploymentResource": null,
                "vmResource": {
                  "type": "DEPLOYMENT_VM",
                  "appProfiles": [
                    "all"
                  ],
                  "cloudRegions": [
                    "all"
                  ],
                  "cloudAccounts": [
                    "all"
                  ],
                  "services": [
                    "all"
                  ],
                  "osTypes": [],
                  "cloudFamilyNames": [],
                  "nodeStates": [],
                  "cloudResourceMappings": [],
                  "isEditable": true
                },
                "cloudRegionResource": null,
                "serviceResource": null,
                "applicationProfileResource": null,
                "deploymentResource": null,
                "vmResource": {
                  "type": "IMPORTED_VM",
                  "appProfiles": [],
                  "cloudRegions": [
                    "all"
                  ],
                  "cloudAccounts": [
                    "all"
                  ],
                  "services": [
                    "all"
                  ],
                  "osTypes": [
                    "all"
                  ],
                  "cloudFamilyNames": [],
                  "nodeStates": [
                    "all"
                  ],
                  "cloudResourceMappings": [
                    "all"
                  ],
                  "isEditable": true
                }
              ]
            ],
            "actionResourceMappingAncillaries": [
              {
                "paramName": "brokerHost",
                "displayName": "BrokerHost",
                "helpText": "Ip Address or Hostname of Rabbit MQ cluster",
                "type": "string",
                "valueList": null,
                "defaultValue": "",
                "confirmValue": "",
                "pathSuffixValue": "",
                "userVisible": true,
                "userEditable": true,
                "systemParam": false,
                "exampleValue": null,
                "dataUnit": null,
                "optional": false,
                "deploymentParam": false,
                "multiselectSupported": false,
                "useDefault": true,
                "valueConstraint": {
                  "minValue": 0,
                  "maxValue": 255,
                  "maxLength": 255,
                  "regex": null,
                  "allowSpaces": true,
                  "sizeValue": 0,
                  "step": 0,
                  "calloutWorkflowName": null
                },
                "scope": null,
                "webserviceListParams": {
                  "url": "",
                  "protocol": "",
                  "username": "",
                  "password": "",
                  "requestType": null,
                  "contentType": null,
                  "commandParams": null,
                  "requestBody": null,
                  "resultString": null,
                  "secret": null,
                  "tabularTypeData": null,
                  "collectionList": []
                },
                "preference": "VISIBLE_UNLOCKED"
              },
              {
                "paramName": "brokerPort",
                "displayName": "BrokerPort",
                "helpText": "RabbitMQ Port number",
                "type": "string",
                "valueList": null,
                "defaultValue": "",
                "confirmValue": "",
                "pathSuffixValue": "",
                "userVisible": true,
                "userEditable": true,
                "systemParam": false,
                "exampleValue": null,
                "dataUnit": null,
                "optional": false,
                "deploymentParam": false,
                "multiselectSupported": false,
                "useDefault": true,
                "valueConstraint": {
                  "minValue": 0,
                  "maxValue": 255,
                  "maxLength": 255,
                  "regex": null,
                  "allowSpaces": true,
                  "sizeValue": 0,
                  "step": 0,
                  "calloutWorkflowName": null
                },
                "scope": null,
                "webserviceListParams": {
                  "url": "",
                  "protocol": "",
                  "username": "",
                  "password": "",
                  "requestType": null,
                  "contentType": null,
                  "commandParams": null,
                  "requestBody": null,
                  "resultString": null,
                  "secret": null,
                  "tabularTypeData": null,
                  "collectionList": []
                },
                "preference": "VISIBLE_UNLOCKED"
              }
            ],
            "id": "58",
            "resource": null,
            "name": "AgentReConfig_Win",
            "description": "",
            "actionType": "EXECUTE_COMMAND",
            "category": "ON_DEMAND",
            "lastUpdatedTime": "2019-09-19 22:15:02.311",
            "timeOut": 1200,
            "enabled": true,
            "encrypted": false,
            "explicitShare": false,
            "showExplicitShareFeature": false,
            "deleted": false,
            "systemDefined": false,
            "bulkOperationSupported": true,
            "isAvailableToUser": true,
            "currentlyExecuting": false,
            "owner": 1,
            "actionParameters": [
              {
                "paramName": "downloadFromBundle",
                "paramValue": "true",
                "customParam": false,
                "required": true,
                "useDefault": false,
                "preference": "VISIBLE_UNLOCKED"
              },
              {
                "paramName": "bundlePath",
                "paramValue": "http://10.0.0.3/5.1-release/ccs-bundle-artifacts-5.1.0-20190819/agent.zip",
                "customParam": false,
                "required": true,
                "useDefault": false,
                "preference": "VISIBLE_UNLOCKED"
              },
              {
                "paramName": "script",
                "paramValue": "agent\\agentReconfig.ps1",
                "customParam": false,
                "required": true,
                "useDefault": false,
                "preference": "VISIBLE_UNLOCKED"
              },
              {
                "paramName": "executeOnContainer",
                "paramValue": "false",
                "customParam": false,
                "required": true,
                "useDefault": false,
                "preference": "VISIBLE_UNLOCKED"
              },
              {
                "paramName": "rebootInstance",
                "paramValue": "false",
                "customParam": false,
                "required": true,
                "useDefault": false,
                "preference": "VISIBLE_UNLOCKED"
              },
              {
                "paramName": "refreshInstanceInfo",
                "paramValue": "false",
                "customParam": false,
                "required": true,
                "useDefault": false,
                "preference": "VISIBLE_UNLOCKED"
              }
            ],
            "actionResourceMappings": [
              {
                "type": "VIRTUAL_MACHINE",
                "actionResourceFilters": [
                  {
                    "cloudRegionResource": null,
                    "serviceResource": null,
                    "applicationProfileResource": null,
                    "deploymentResource": null,
                    "vmResource": {
                      "type": "DEPLOYMENT_VM",
                      "appProfiles": [
                        "all"
                      ],
                      "cloudRegions": [
                        "all"
                      ],
                      "cloudAccounts": [
                        "all"
                      ],
                      "services": [
                        "all"
                      ],
                      "osTypes": [],
                      "cloudFamilyNames": [],
                      "nodeStates": [],
                      "cloudResourceMappings": [],
                      "isEditable": true
                    },
                    "cloudRegionResource": null,
                    "serviceResource": null,
                    "applicationProfileResource": null,
                    "deploymentResource": null,
                    "vmResource": {
                      "type": "IMPORTED_VM",
                      "appProfiles": [
                        "all"
                      ],
                      "cloudRegions": [
                        "all"
                      ],
                      "cloudAccounts": [
                        "all"
                      ],
                      "services": [
                        "all"
                      ],
                      "osTypes": [
                        "all"
                      ],
                      "cloudFamilyNames": [],
                      "nodeStates": [
                        "all"
                      ],
                      "cloudResourceMappings": [
                        "all"
                      ],
                      "isEditable": true
                    }
                  ]
                },
                "actionResourceMappingAncillaries": [
                  {
                    "paramName": "brokerHost",
                    "displayName": "BrokerHost",
                    "helpText": "Ip Address or Hostname of Rabbit MQ cluster",
                    "type": "string",
                    "valueList": null,
                    "defaultValue": "",
                    "confirmValue": "",
                    "pathSuffixValue": "",
                    "userVisible": true,
                    "userEditable": true,
                    "systemParam": false,
                    "exampleValue": null,
                    "dataUnit": null,
                    "optional": false,
                    "deploymentParam": false,
                    "multiselectSupported": false,
                    "useDefault": true,
                    "valueConstraint": {
                      "minValue": 0,
                      "maxValue": 255,
                      "maxLength": 255,
                      "regex": null,
                      "allowSpaces": true,
                      "sizeValue": 0,
                      "step": 0,
                      "calloutWorkflowName": null
                    },
                    "scope": null,
                    "webserviceListParams": {
                      "url": "",
                      "protocol": "",
                      "username": "",
                      "password": "",
                      "requestType": null,

```

```
contentType":null,"commandParams":null,"requestBody":null,"resultString":null},"secret":null,"
tabularTypeData":null,"collectionList":[],"preference":"VISIBLE_UNLOCKED"},{"paramName":"brokerPort",
displayName":"BrokerPort","helpText":"RabbitMQ Port number","type":"string","valueList":null,"
defaultValue":"","confirmValue":"","pathSuffixValue":"","userVisible":true,"userEditable":true,"
systemParam":false,"exampleValue":null,"dataUnit":null,"optional":false,"deploymentParam":false,"
multiselectSupported":false,"useDefault":true,"valueConstraint":{"minValue":0,"maxValue":255,"maxLength":
255,"regex":null,"allowSpaces":true,"sizeValue":0,"step":0,"calloutWorkflowName":null},"scope":null,"
webserviceListParams":{"url":"","protocol":"","username":"","password":"","requestType":null,"
contentType":null,"commandParams":null,"requestBody":null,"resultString":null},"secret":null,"
tabularTypeData":null,"collectionList":[],"preference":"VISIBLE_UNLOCKED"}]]},"
repositoriesMappingRequired":false,"actionTypesCounts":[{"key":"EXECUTE_COMMAND","value":"2"}]}
```

3. Access Workload Manager in your OLD cluster and navigate to the Actions Library page.
4. Import the actions.json file that you saved in Step 2 above. You should see two files (**AgentReconfig\_Linux** and **AgentReconfig\_Win**) as displayed in the following screenshot.

AgentReConfig_Linux ON DEMAND	Command or Script	September 19, 2019 4:48 PM	Virtual Machines	<input checked="" type="checkbox"/>	<input type="checkbox"/>
AgentReConfig_Win ON DEMAND	Command or Script	September 19, 2019 4:48 PM	Virtual Machines	<input checked="" type="checkbox"/>	<input type="checkbox"/>

5. The files are disabled by default (OFF) – enable both files by toggling each switch to **ON**.
6. Save the following script to a file in your local directory and name it **agentReconfig.sh**. This is the file to use for Linux environments.

#### The agentReconfig.sh file

```
#!/bin/bash

#Write to system log as well as to terminal
logWrite()
{
    msg=$1
    echo "$(date) ${msg}"
    logger -t "OSMOSIX" "${msg}"
    return 0
}

logWrite "Starting agent migrate..."

env_file="/usr/local/osmosix/etc/userenv"
if [ -f $env_file ];
then
    logWrite "Source the userenv file..."
    . $env_file
fi

if [ -z $brokerHost ];
then
    logWrite "Broker Host / Rabbit Server Ip not passed as action parameter"
    exit 3;
fi

if [ -z $brokerPort ];
then
    logWrite "Broker Port / Rabbit Server Port not passed as action parameter"
    exit 4
fi

replaceUserDataValue() {
    key=$1
    value=$2

    if [ -z $key ] || [ -z $value ];
    then
        logWrite "Command line arguments missing to update user-data file, key: $key, value:$value"
        return
    fi

    user_data_file="/usr/local/agentlite/etc/user-data"
```



```

    if [ -f $user_data_file ];
    then
        json_content=`cat $user_data_file`
        old_value=`echo $json_content | awk -F $key '{print $2}' | awk -F \" '{print $3}'`
        sed -i 's@"$old_value"@'"$value"@g' $user_data_file
    fi
}

export AGENT_HOME="/usr/local/agentlite"

logWrite "Updating the user data file"
replaceUserDataValue "brokerClusterAddresses" "$brokerHost:$brokerPort"

logWrite "Updating config.json file"
sed -i '/AmqpAddress/c\      "AmqpAddress": "'${brokerHost}:${brokerPort}','','' "$AGENT_HOME/config/config.json"

cd $AGENT_HOME
echo "sleep 10" > execute.sh
echo "/usr/local/agentlite/bin/agent-stop.sh" >> execute.sh
echo "/usr/local/agentlite/bin/agent-start.sh" >> execute.sh
chmod a+x execute.sh
nohup bash execute.sh > /dev/null 2>&1 &

exit 0

```

7. Save the following script to a file in your local directory and name it **agentReconfig.ps1**. This is the file to use for Windows environments.

**The agentReconfig.ps1 file**

```

param (
    [string]$brokerHost = "$env:brokerHost",
    [string]$brokerPort = "$env:brokerPort"
)

$SERVICE_NAME = "AgentService"
$SYSTEM_DRIVE = (Get-WmiObject Win32_OperatingSystem).SystemDrive
. "$SYSTEM_DRIVE\temp\userenv.ps1"

if ($brokerHost -eq 0 -or $brokerHost -eq $null -or $brokerHost -eq "") {
    echo "Variable brokerHost not available in the env file"
    exit 1
}

if ($brokerPort -eq 0 -or $brokerPort -eq $null -or $brokerPort -eq "") {
    echo "Variable brokerPort not available in the env file"
    exit 2
}

$AGENTGO_PARENT_DIR = "$SYSTEM_DRIVE\opt"

echo "Check if AgentGo Parent directory exists. If not create it: '$AGENTGO_PARENT_DIR'"
if (-not (Test-Path $AGENTGO_PARENT_DIR)) {
    echo "Create $AGENTGO_PARENT_DIR..."
    mkdir $AGENTGO_PARENT_DIR
}
else {
    echo "$AGENTGO_PARENT_DIR already exists."
}

$AGENT_CONFIG="{0}\agentlite\config\config.json" -f $AGENTGO_PARENT_DIR
if (Test-Path $AGENT_CONFIG) {
    echo "Changing the config.json file with the new broker host $env:brokerHost and port $env:
brokerPort"
    $confJson = get-content $AGENT_CONFIG | out-string | convertfrom-json
    $confJson.AmqpAddress = "$($env:brokerHost): $($env:brokerPort)"
    $confJson | ConvertTo-Json | set-content $AGENT_CONFIG
}

$USER_DATA_FILE = "{0}\agentlite\etc\user-data" -f $AGENTGO_PARENT_DIR
if (Test-Path $USER_DATA_FILE) {
    echo "Changing user-data file with new broker host $env:brokerHost and port $env:brokerPort"
    $userDataJson = get-content $USER_DATA_FILE | out-string | convertfrom-json
    $userDataJson.brokerClusterAddresses = "$($env:brokerHost): $($env:brokerPort)"
    $userDataJson | ConvertTo-Json | set-content $USER_DATA_FILE
}

$AGENT_SERVICE_NAME = "AgentService"
echo "Stop-Service $AGENT_SERVICE_NAME" > $AGENTGO_PARENT_DIR\exec.ps1
echo "sleep 10" >> $AGENTGO_PARENT_DIR\exec.ps1
echo "Start-Service $AGENT_SERVICE_NAME" >> $AGENTGO_PARENT_DIR\exec.ps1

echo "Restarting agent"
Start-Process -filepath "powershell" -argumentlist "-executionpolicy bypass -noninteractive -file
`"$AGENTGO_PARENT_DIR\exec.ps1`""

echo "Agent set to restart after config changes"

```


8. Add these two files to a folder called **agent** (just an example) and compress the folder to create **agent.zip** with the same structure displayed here.

**agent**

**agentReconfig.ps1**

**agentReconfig.sh**

9. Move the **agent.zip** folder to an HTTP repository in your local environment that is accessible from the OLD and NEW clusters.

 This procedure uses the following URL as an example:

**http://10.0.0.3/repo/agent.zip**

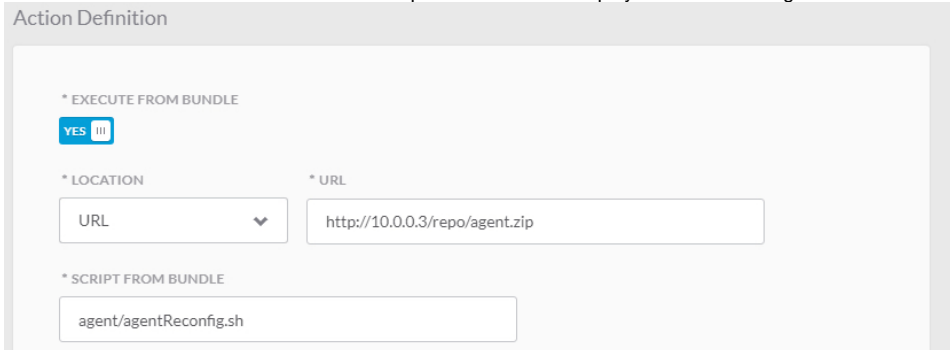
You have now ensured cluster connectivity and saved the required files for the migration procedure.

**8f. Migrate Deployments from the OLD Cluster to the NEW Cluster**

To migrate the deployment details from the old cluster to the new cluster, follow this procedure.

1. Navigate to the Workload Manager **Actions Library** page and edit the **AgentReconfig\_Linux** action. This procedure continues to use the Linux file going forward.
2. Scroll to the **Actions Definition** section and update the URL as displayed in the following screenshot.

Action Definition



\* EXECUTE FROM BUNDLE  
 YES

\* LOCATION  
 URL

\* URL

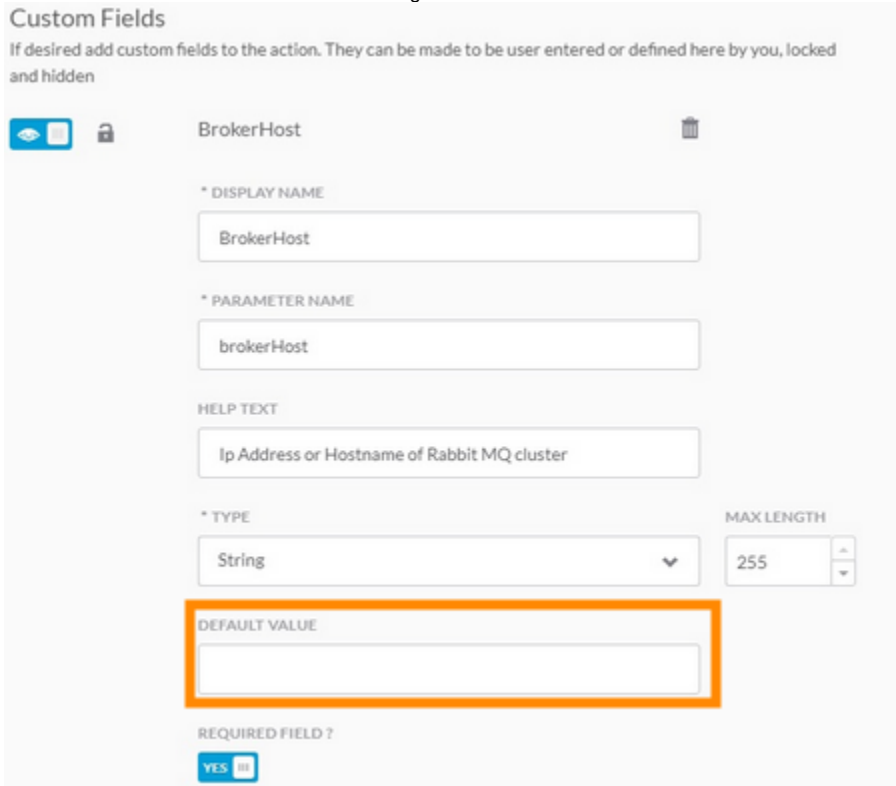
\* SCRIPT FROM BUNDLE

 The URL and Script from Bundle fields in the above screenshot are in accordance with the steps above.

3. Scroll to the **Custom Fields** section and change the default value of the **Broker Host** to use the NEW cluster IP.

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



BrokerHost

\* DISPLAY NAME

\* PARAMETER NAME

HELP TEXT

\* TYPE

MAX LENGTH

DEFAULT VALUE

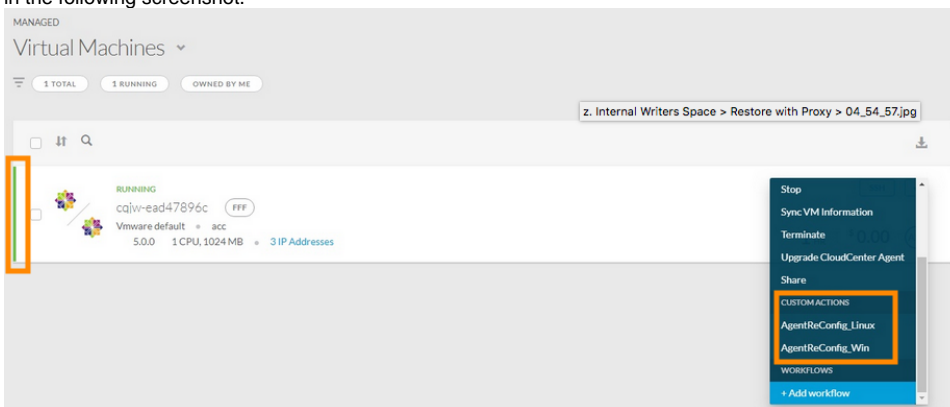
REQUIRED FIELD?  
 YES

4. Scroll down to the **Broker Port** and change the default to use the NEW Worker AMQP IP port (for example, 26642 in Step 8 above).

The screenshot shows the configuration page for a parameter named "BrokerPort". The form contains the following fields and options:

- \* DISPLAY NAME:** BrokerPort
- \* PARAMETER NAME:** brokerPort
- HELP TEXT:** RabbitMQ Port number
- \* TYPE:** String
- MAX LENGTH:** 255
- DEFAULT VALUE:** (This field is highlighted with an orange border and is currently empty.)
- REQUIRED FIELD?:** YES (checked)

- Click **Done** to save your default configuration changes in the OLD cluster.
- Navigate to the **Virtual Machines** page and locate the VM to migrate to the new cluster.
- Click the **Actions** dropdown and verify if your newly modified actions are visible under the Custom Actions section in the dropdown list as visible in the following screenshot.



- Click one of the actions and verify that the configured defaults are displayed in the Broker host and Broker port fields as indicated earlier.
- Click **Submit** to migrate this VM to the new cluster.
- Verify that the migration is complete by going to the Deployment page in your NEW cluster and the VM is listed as RUNNING (green line).
- Repeat Steps 6 through 10 for each VM that needs to be migrated to the NEW cluster.

You have now migrated the deployment details from the old cluster to the new cluster

**Back to:** [With Internet Access](#)

# Without Internet Access

## Isolated (Air Gap) Environment Setup

- [Overview](#)
- [Minio Server Setup](#)
- [Backup and Restore Process](#)
- [Sample Commands Using Fictional Names](#)

You may sometimes need to work in an environment that is completely behind the firewall. This section addresses the backup and restore procedures for those environments.

See [Backup](#) for restrictions and limitations.

You need to set up a Minio server to configure a S3-compatible backup storage location. Refer to <https://min.io/download#/macos> to setup the Minio server.

Once the Minio server is setup, use YOUR Minio server credentials to login to your Minio server.

- Minio server URL
- Minio server username
- Minio server password



The script provided as part of this process uses publicly available **Velero** and **Minio** tools to complete the manual backup and restore process in isolated environments.

To backup and restore the CloudCenter Suite data in an air gap environment, follow this procedure.

1. Create a bucket on the Minio server and provide a meaningful name. This example, uses **velero**. See [Backup](#) for details.
2. Before installing Velero, annotate all the pods in your cluster by using Velero-specific annotations that are provided in the script below.

```
kubectl -n YOUR_POD_NAMESPACE annotate pod/YOUR_POD_NAME backup.velero.io/backup-
volumes=YOUR_VOLUME_NAME_1,YOUR_VOLUME_NAME_2,...
```

To make things simpler here is a utility that does it for you. Be sure to save the following script contents to a file called **pod\_vol\_restic\_scan.py** to your local system.

### The pod\_vol\_restic\_scan.py script

```
# This utility is used to annotate pods for Velero backups

import random
import logging
import string
import os
import time
import datetime
from argparse import ArgumentParser
import sys
import zipfile
import shutil
import subprocess
import re
from pprint import pprint as pp
import yaml

__copyright__ = "Copyright 2019, abmitra"
__license__ = "Cisco Systems"

def script_run_time(seconds):
    min, sec = divmod(seconds, 60)
    hrs, min = divmod(min, 60)
    timedatastring = "%d:%02d:%02d" % (hrs, min, sec)
    return timedatastring
```

```

def random_char(y):
    return ''.join(random.choice(string.ascii_letters) for x in range(y))

def border_print(symbol, msg):
    line = " " + msg + " "
    totalLength = len(line) + 50
    logger.info("")
    logger.info(symbol * totalLength)
    logger.info(line.center(totalLength, symbol))
    logger.info(symbol * totalLength)
    logger.info("")

def setup_custom_logger(name, tcStartTime, fileBaseName, inputName=""):
    if inputName == "" or inputName == None:
        st = datetime.datetime.fromtimestamp(tcStartTime).strftime('%Y-%m-%d-%H-%M-%S')
        filename = fileBaseName + "-" + st + '.log'
        dirName = "po-scan" + st
        dirPath = os.path.abspath(os.path.join(os.path.dirname(__file__), '..', dirName))
        logfilename = os.path.join(dirPath, filename)
        if not os.path.isdir(dirPath):
            os.makedirs(dirPath)
    else:
        logfilename = inputName

    # print(logfilename)
    formatter = logging.Formatter(fmt='%(asctime)s %(levelname)-8s %(message)s',
                                  datefmt='%Y-%m-%d %H:%M:%S')
    handler = logging.FileHandler(logfilename, mode='w')
    handler.setFormatter(formatter)
    screen_handler = logging.StreamHandler(stream=sys.stdout)
    screen_handler.setFormatter(formatter)
    logger = logging.getLogger(name)
    logger.setLevel(logging.DEBUG)
    logger.addHandler(handler)
    logger.addHandler(screen_handler)
    return logger, logfilename

def shell_cmd(cmd):
    logger.info("Shell cmd execution >>> {}".format(cmd))
    p = subprocess.Popen(cmd, shell=True, stdout=subprocess.PIPE, universal_newlines=True)
    output = p.communicate()[0]
    p_status = p.wait()
    return output.split("\n")

def zipdir(path, ziph):
    # ziph is zipfile handle
    for root, dirs, files in os.walk(path):
        for file in files:
            # print(file)
            ziph.write(os.path.join(root, file))

def create_zip():
    st = datetime.datetime.fromtimestamp(tcStartTime).strftime('%Y-%m-%d-%H-%M-%S')
    dirName = "ccs-log" + st
    zipFileName = dirName + ".zip"
    zipFilePath = os.path.abspath(os.path.join(os.path.dirname(__file__)))
    logger.info("Generating zip file '{}' at {}".format(zipFileName, zipFilePath))
    zipf = zipfile.ZipFile(zipFileName, 'w', zipfile.ZIP_DEFLATED)
    zipdir(dirName, zipf)
    zipf.close()
    shutil.rmtree(dirName)

```

```

if __name__ == "__main__":

    fileName = os.path.basename(__file__).split(".")[0]
    tcStartTime = time.time()
    timeStamp = datetime.datetime.fromtimestamp(tcStartTime).strftime('%Y%m%d%H%M%S')

    parser = ArgumentParser()
    parser.add_argument("-n", "--namespace", dest="namespace", help="Kubernetes Namespace", required=True)
    args = parser.parse_args()
    namespace = args.namespace.strip()
    logger, logFileName = setup_custom_logger("Cloudcenter K8 Debug", tcStartTime, fileName)

    cmd = "kubectl get pod -n " + namespace + " | grep -v NAME | awk '{print $1}'"
    pod_name_list = shell_cmd(cmd)
    pod_pvc_dict = {}
    pod_vol_dict = {}

    for pod in pod_name_list:
        if pod != "":
            cmd = "kubectl get pod {} -n {} -o yaml > temp.yaml".format(pod, namespace)
            data = shell_cmd(cmd)
            temp_file = open("temp.yaml", "r")
            with open('temp.yaml', 'r') as temp_file:
                try:
                    file_contents = (yaml.load(temp_file))
                    #print("Pod Name = {}".format(pod.strip()))
                    for vol in file_contents['spec']['volumes']:
                        #pp(vol)
                        try:
                            pvc = vol['persistentVolumeClaim']
                            pod_vol_dict[pod.strip()] = vol['name'].strip()
                            #print("Vol Name = {}".format(vol['name']))
                        except:
                            pass
                    except yaml.YAMLError as exc:
                        logger.error("Error in reading YAML file.")
                        logger.error(exc)
                os.remove('temp.yaml')

            #pp(pod_vol_dict)
            border_print("+","Applying POD annotations")
            for pod in pod_vol_dict.keys():
                cmd = "kubectl -n {} annotate --overwrite pod {} backup.velero.io/backup-volumes={}".format(
namespace,pod,pod_vol_dict[pod])
                data = shell_cmd(cmd)

```

- From where you have saved the pod\_vol\_restic\_scan.py script, run the following command.

```

#Needs Python3
python pod_vol_restic_scan.py -n cisco

```

- Install Velero Version 0.11.0 – refer to <https://velero.io/docs/v0.11.0/> for details.
- Create a credential file to store your credentials. This example, uses the following URL and credentials – *this is only an example!*

#### Contents of the credentials-minio file

```

[default]
aws_access_key_id = <your Minio username>
aws_secret_access_key = <your Minio password>

```

- On the existing Kubernetes cluster, you must deploy Velero and configure it with the AWS compatible bucket location, in this example, *minio*.

### ✓ Velero and Minio Usage

This process uses Velero to backup the Kubernetes data to a Minio server.

Once you finish this task you can configure the AWS S3 storage provider using the Minio server credentials as specified below. Configuring Minio is similar to configuring an AWS S3 environment, the difference is that you must provide the region and endpoint details when adding the Minio server as AWS S3 storage. You can verify the data from Minio server GUI or command line. The following steps are an example to verify the data from the Minio command line.

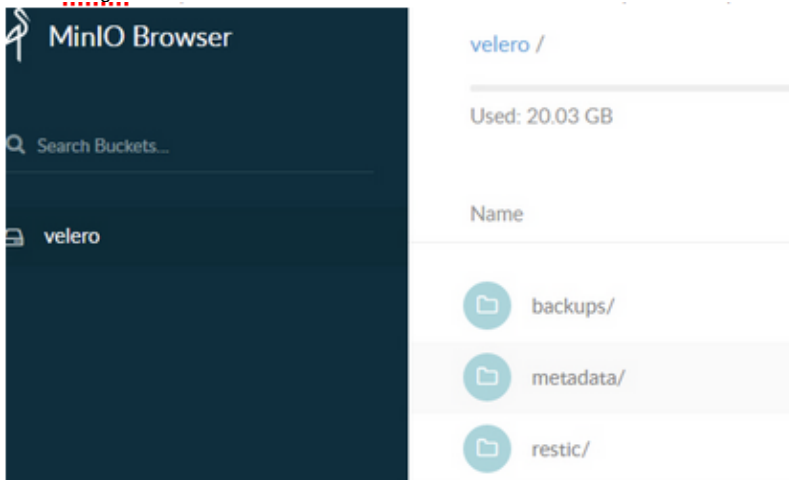
Refer to <https://docs.min.io/docs/aws-cli-with-minio.html> for additional details.

```
velero install \
  --provider aws \
  --bucket <Minio bucket name from Step 1 above> \
  --secret-file <Fully qualified path of the Minio credentials file> \
  --use-volume-snapshots=false \
  --backup-location-config region=minio,s3ForcePathStyle="true",s3Url=<your Minio server URL> \
  --use-restic \
  --wait
```

7. Start a backup using the following command.

```
velero backup create <Minio backup name> --include-namespaces=cisco --wait
```

8. Wait for the backup to complete and watch the logs. Once the backup is complete, the Minio output should look like the information displayed in the following screenshot.



9. To restore the backup to a different cluster or a fresh cluster (assuming that the *cisco* namespace is not present).

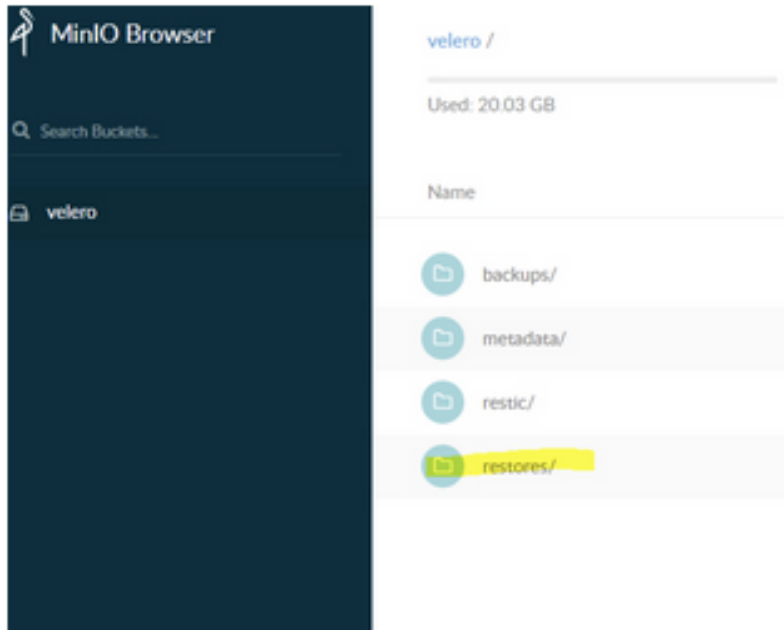
```
velero install \
  --provider aws \
  --bucket <Minio bucket name from Step 1 above> \
  --secret-file <Fully qualified path of the Minio credentials file> \
  --use-volume-snapshots=false \
  --backup-location-config region=minio,s3ForcePathStyle="true",s3Url=<your Minio server URL> \
  --use-restic \
  --wait
```

10. Start the restore process:

```
velero restore create --from-backup <Minio backup name>
```



11. The Minio output should look like the information displayed in the following screenshot – you will see an additional restore folder as displayed in the following screenshot



You have now backed up and restored the CloudCenter Suite to an isolated environment using the Minio server.



The following commands are only examples and need to be run using the names that you have assigned to resources in your environment.

#### Deploying Velero on the Existing Cluster

```
velero install \
  --provider aws \
  --bucket velero \
  --secret-file ./credentials-velero \
  --use-volume-snapshots=false \
  --backup-location-config region=minio,s3ForcePathStyle="true",s3Url=http://12.16.1.1:9000 \
  --use-restic \
  --wait
```

#### Backup Command

```
velero backup create minio-backup --include-namespaces=cisco --wait
```

#### Deploying Velero on the New Cluster

```
velero install \
  --provider aws \
  --bucket velero \
  --secret-file ./credentials-velero \
  --use-volume-snapshots=false \
  --backup-location-config region=minio,s3ForcePathStyle="true",s3Url=http://12.16.1.1:9000 \
  --use-restic \
  --wait
```

**Restore Command**

```
velero restore create --from-backup minio-backup
```

# Troubleshooting

## Troubleshooting

- [Expired Certificates](#)
- [Finding Kubernetes Resources](#)
- [A Pod has unbound PersistentVolumeClaims](#)
- [Error during the Suite Installation Process](#)
- [Error in Creating Cluster](#)
- [The Progress bar for a Kubernetes Cluster is stuck at Launching cluster nodes on the cloud or Configuring the primary cluster](#)
- [The Kubernetes Cluster is installed successfully, but the progress bar for Suite Administration is stuck at Waiting for product to be ready](#)
- [Installation Failed: Failed to copy <script-name.sh> to remote host or any error related to SSH connection failure](#)
- [DHCP IP Allocation Mode](#)
- [After using Suite Admin for a while, users cannot login to Suite Admin if any of the cluster nodes are in a Not Ready state](#)
- [When one of the workers is down a worker node scale up operation is stuck](#)
- [Download Logs](#)
- [Velero Issues](#)

Sometimes, the certificates may have expired and cannot renew automatically – you will typically see this problem when you can no longer login to your CloudCenter Suite cluster and you start receiving a networking error or similar error. If you review the *AUTH pod logs* you will may issues with accessing the certificate or its location. If you review the certification, you will see that it is failing due to an auto-renew setting with an error similar to the following code block:

```
kubectl -n cisco describe cert suite-auth-tls
```

These kind of error are caused by changes in either the certificate manager or the certificate failing due to an auto-renewal setting and the cluster is down.

To address this issue fix your cluster by following this process – use the following scripts with *caution*.



These scripts were only tested in a GCP environment where this error was first seen.

1. Export the current certificates and secrets to YAML files.

```
#!/bin/bash
namespace="cisco"
mkdir -p $namespace
for n in $(kubectl -n $namespace get secrets -o custom-columns=:metadata.name | grep -v 'service-account')
do
    echo "Saving $namespace/secret_$.n..."
    kubectl -n $namespace get secret $n -o yaml > $namespace/secret_$.n.yaml
done
for n in $(kubectl -n $namespace get cert -o custom-columns=:metadata.name)
do
    echo "Saving $namespace/cert_$.n..."
    kubectl -n $namespace get cert $n -o yaml > $namespace/cert_$.n.yaml
done
```

2. Delete the old certificates and secrets.

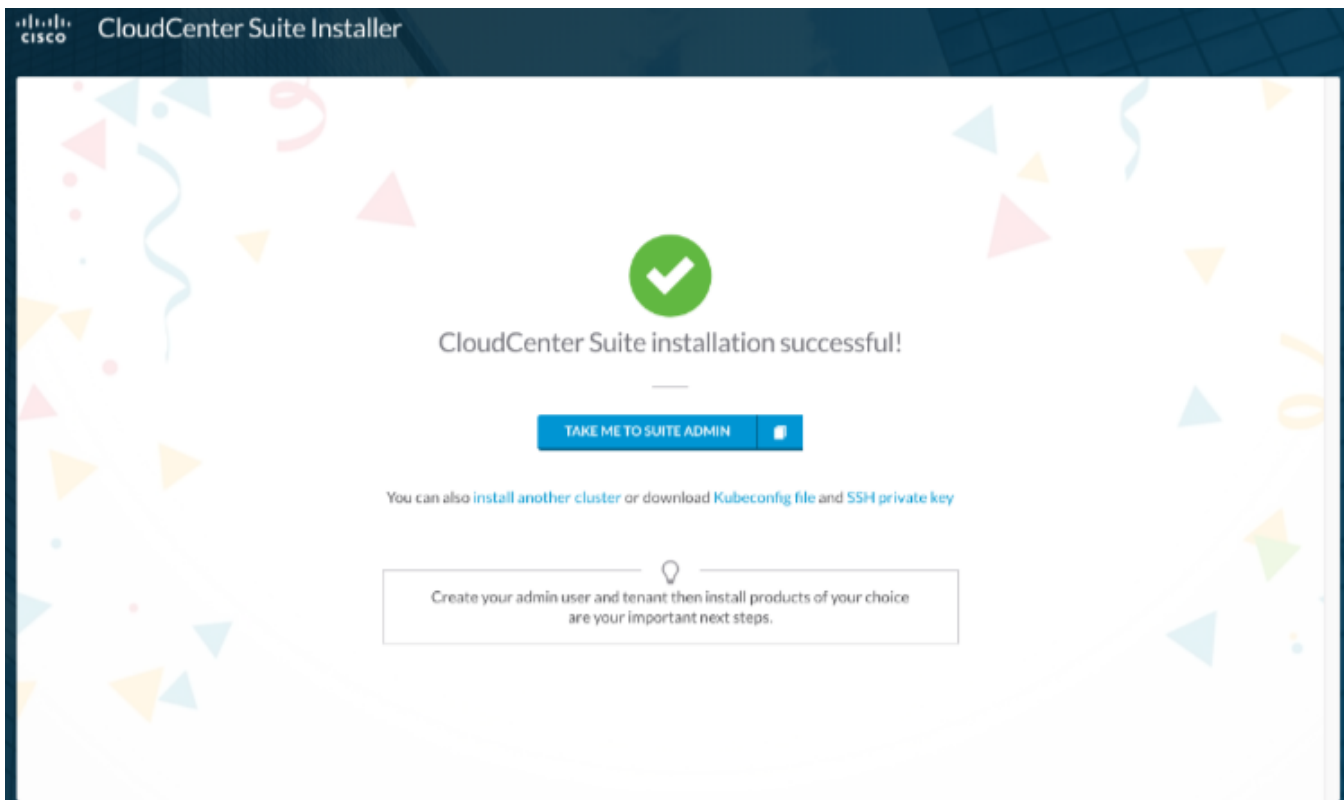
```
#!/bin/bash
commit=1
for n in $(kubectl -n cisco get secrets -o custom-columns=:metadata.name | grep -v 'service-account')
do
    echo "Deleting $n..."
    if [[ $commit==1 ]]; then
        kubectl -n cisco delete secret "$n"
    fi
done
for n in $(kubectl -n cisco get cert -o custom-columns=:metadata.name)
do
    echo "Deleting $n..."
    if [[ $commit==1 ]]; then
        kubectl -n cisco delete cert "$n"
    fi
done
```

3. Restore the certificates from their respective YAML files to the cluster.

```
#!/bin/bash
namespace=cisco
echo "Restoring Opaque secrets..."
kubectl apply -f $namespace/secret_ca-key-pair.yaml
kubectl apply -f $namespace/secret_suite-fluentd-s3-config.yaml
kubectl apply -f $namespace/secret_suite-fluentd-s3-config-original.yaml
kubectl apply -f $namespace/secret_suite-gateway-external-tls-secrets.yaml
kubectl apply -f $namespace/secret_suite-random-password.yaml
kubectl apply -f $namespace/secret_suite-image-pull-secret.yaml
kubectl apply -f $namespace/secret_action-orchestrator-jwt-secret.yaml
echo "Restoring Certs via YAML"
for n in $namespace/*.yaml; do
    [ -f "$n" ] || break
    if [[ $n =~ "cert" ]]; then
        echo "Restoring Cert via yaml file $n..."
        kubectl apply -f "$n"
    fi
done
echo "Restarting Cert Manager Pod..."
kubectl delete --all pods --namespace=cert-manager
echo "Restarting all CCS Pods..."
kubectl delete --all pods --namespace=$namespace
```

4. Restore the opaque and other non-TLS based secrets.
5. Restart the cert-manager.
6. Restart all the CloudCenter Suite cluster pods.

For private clouds, the download link for the **Kubeconfig file** is available on the last page of the installer UI as displayed in the following screenshot.



While you may see this file for successful installations in the above screen, you will not be able to access this file if your installation was not successful. This file is required to issue any command listed in the <https://kubernetes.io/docs/reference/kubectl/cheatsheet/> section of the Kubernetes documentation.

By default, the `kubectl` command looks for the Kubeconfig file in the `$HOME/.kube` folder.

- **Successful installation:** Copy the downloaded Kubeconfig file to your `$HOME/.kube` folder and then issue any of the `kubectl` commands listed in the Kubernetes cheatsheet link above.
- **Stalled Installation:**
  - Private clouds and most public clouds: SSH into one of the primary server nodes and copy the Kubeconfig file from `/etc/kubernetes/admin.conf` to the `/root/.kube` folder.
  - GCP: Login to GCP, access the Kubernetes Engine, locate your cluster, click **Connect** to Connect to the cluster, and click the **copy** icon as displayed in the following screenshot. You should have already installed `gcloud` in order to view this icon.

## Connect to the cluster

You can connect to your cluster via command-line or using a dashboard.

**Command-line access**

Configure `kubectl` command line access by running the following command:

```
$ gcloud container clusters get-credentials pujanrc221-7220e62e-ca6f-4f08-963c-9e49b --zone us-east1-b --project [ ]
```

[Run in Cloud Shell](#)

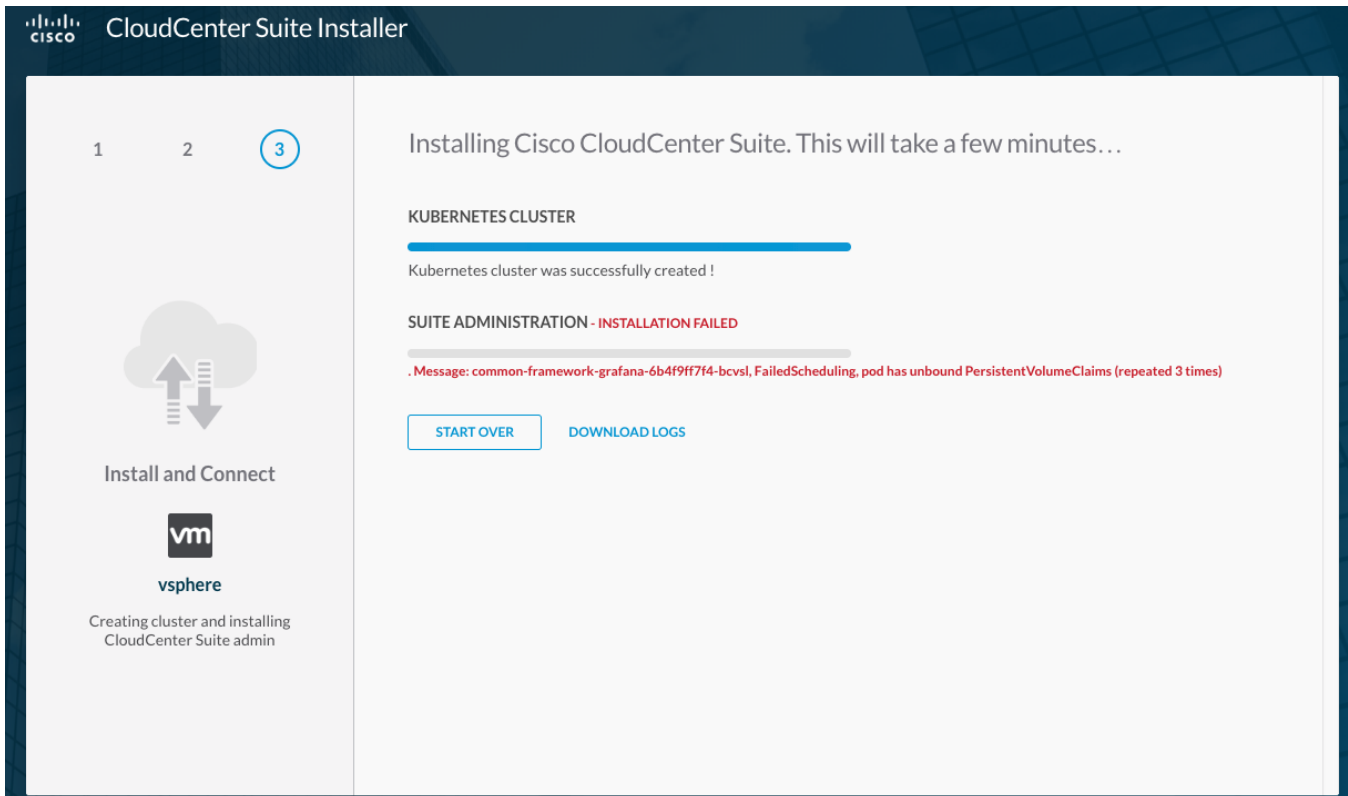
**Cloud Console dashboard**

You can view the workloads running in your cluster in the Cloud Console [Workloads dashboard](#).

[Open Workloads dashboard](#)

OK

The problem displayed in the following screenshot is usually caused when the cloud user does not have permissions to the configured storage. For example, a vSphere user may not have permissions to the selected datastore.



At any time, if your installation stalls due to a lack of resources, perform this procedure to analyze the error logs.

To fetch the logs for this pod run :

1. Locate the actual name of the container by running the following command:

```
kubectl get pods -all-namespaces | grep common-framework-suite-prod-mgmt-xxxx
```

2. Click the **Download Logs** link to download the installation logs for the failed service in case of an installation failure.
3. **View the** Logs for the container: common-framework-suite-prod-mgmt ...
4. Run the following command to view the error:

```
kubectl logs -f common-framework-suite-prod-mgmt-xxxx -n cisco
```

In case of failure (due to a quota availability issue) during the installation process, an error message similar to the one displayed in the following screenshot appears.

CloudCenter Suite Installer

1 2 3

Specify Placement Properties

Google Kubernetes Engine

Specify details to help us understand about your cluster

What are your gke placement properties?

GKE CLUSTER ID PREFIX

gkecluster 30

GKE Cluster ID Prefix, prefix must start with a lowercase letter and cluster's name must not be longer than 40 characters. Only letters, numbers and hyphen are allowed in a cluster's name.

\* GKE ZONE

asia-east2-a

GKE Zone

\* GKE INSTANCE TYPE

n1-standard-2

GKE Instance Type

Error in Creating Cluster! rpc error: code = PermissionDenied desc = Insufficient project quota to satisfy request: resource "ROUTES": request requires '3.0' and is short '2.0', project has a quota of '300.0' with '1.0' available.

< CHANGE CLOUD

ERROR IN CREATING CLUSTER! RPC ERROR: CODE...

INSTALL

The issue displayed in the following screenshot could be an issue with the cloud environment. Refer to your cloud documentation for possible issues.

CloudCenter Suite Installer

1 2 3

Install and Connect

Amazon EKS

Creating cluster and installing CloudCenter Suite admin

Installing Cisco CloudCenter Suite. This will take a few minutes...

KUBERNETES CLUSTER - INSTALLATION FAILED

Launching cluster on the cloud

SUITE ADMINISTRATION

START OVER

DOWNLOAD LOGS

Other examples:

- If the target cloud is vSphere, check if the cloud account being used has permissions to launch a VM and if the VM is configured with a valid IPv4 address.
- If the cluster nodes are configured to use static IP, verify if the IP pool used is valid and if all the launched nodes have a unique IP from the pool.

This issue indicates that the CloudCenter Suite installation has some issue. Use the downloaded SSH key to SSH into one of the primary server nodes. To check the status of the pods, run **kubectl get pods --all-namespaces** for each pod. If the status does not display **Running**, run the following commands to debug further:

```
kubectl describe pod <pod-name> -n cisco
```

or

```
kubectl logs -f <pod-name> -n cisco
```

Use the downloaded SSH key to SSH into each cluster node and check if the system clock is synchronized on all nodes. Even if the NTP servers were initially synchronized verify if they are still active by using the following command.

```
ntpdate <ntp_server>
```

If any of the nodes are **Not Ready** state, then run the following command on the node:

```
kubectl describe node <node-name>
```

This issue can occur when the installer node cannot SSH/SCP into launched cluster nodes. Verify if all the launched nodes have a valid IPv4 address and if the installer network can communicate with the Kubernetes cluster network (if they are on different networks). Also verify that the cluster nodes are able to connect to vSphere.

If none of the above methods work, retry the installation or contact your CloudCenter Suite admin.

During installation if you select DHCP IP allocation mode, you may see an error when you start the installation (assuming other values are appropriate). In this case, check your installer VM's /etc/resolve.conf file, and comment or remove the entry containing the following keyword.

```
searchdomain
```

This entry adds a search domain entry in the /etc/resolve.conf file (required). This addition finds IPs for Nginx services from external locations, and maps them to the Nginx service within the CloudCenter Suite. However, CCP services do not need these maps as the internal IP map to the nginx service is the only required mapping. As such, you must remove the errant entries from the /etc/resolve.conf file.

To correct the error, you (sudo permissions required) must restart all the installer PODs which contain using the following command.



Execute **sudo -i** or prefix the command with **sudo**

```
kubectl delete pod $(kubectl get pods -n ccp | grep suite | awk '{print $1}') -n ccp
```

Now, wait for a minute to ensure that the PODs have started running before restarting the installation process.

This issue may be the result of any of the following situations:

- Are all the cluster nodes up and running with a valid IP address?
- If the nodes are running, then use the downloaded SSH key to SSH into one of the primary server nodes.
- Run the following command on the primary server to verify if all the nodes are in the **Ready** state.

```
kubectl get nodes
```

When one of the workers is down, and you try to scale up the worker node, the node does not scaled up. The scale up operation remains stuck in scaling state.

Restart the operator POD of your environment by using the following command. The following example displays vSphere, and the corresponding vSphere operator. Similarly, if you are working in an OpenStack environment, use the OpenStack operator as applicable.



```
kubectl delete pod kaas-ccp-vsphere-operator-<dynamic alphanumeric characters> -n ccp  
  
#or  
  
kubectl delete pod kaas-ccp-openstack-operator-<dynamic alphanumeric characters> -n ccp
```

By restarting this service on any worker node, you will start the shutdown VM and scale up the new node that was stuck during the scale operation.

Click the **Download Logs Download** link to download the installation logs for the failed service in case of an installation failure. See [Monitor Modules > Download Logs](#) for additional information.

Refer to <https://heptio.github.io/velero/v0.11.0/> for Velero troubleshooting information.

# Suite Admin Workflow

## Suite Admin Workflow

The following table identifies the tasks to be performed on the Suite Admin once you install the CloudCenter Suite.

#	Required?	Goal	Task	Description
1	Yes	Onboarding	Create the suite administrator and root tenant.	See <a href="#">Initial Administrator Setup</a>
			Navigate to the Suite Admin Dashboard.	See <a href="#">Suite Admin Dashboard</a>
2	No	Language selection	Select your language choices to view the CloudCenter Suite UI.	See <a href="#">UI Language Availability</a>
3	Yes	Module installation	Install module(s) of choice based on the list available in the Dashboard.  This is optional, however, you cannot configure resources other than users/tenants/groups/roles/admin menu settings if you don't install modules!	See <a href="#">Install Module</a>
4	Yes	User management	Create users	See <a href="#">Create and Manage Users</a>
5	Yes	Group Management	Assign users to default groups.  When the suite administrator installs any module, additional, default out-of-box groups become available. These groups vary based on the module.	See <a href="#">Create and Assign Groups</a>
	Optional		Create a custom group  If the out-of-box groups don't meet your requirements, you can create custom groups.	See <a href="#">Custom Groups by Admin</a>
	Yes		Assign roles to a group  For each custom group, you must assign at least one role.	See <a href="#">Understand Roles</a>
6	Yes	Admin Management	Set up the base URL	See <a href="#">Base URL Configuration</a>
	Yes		Set up email communication	See <a href="#">Email Settings</a>
	Optional		Configure a dedicated alias hostname and use an external IdP to authenticate its users.	See <a href="#">SSO Setup</a>
	Optional		Set up the proxy server	See <a href="#">Proxy Settings</a>
7	Yes	Product Registration	Configure a license	See <a href="#">Configure Smart Licenses</a>
8	Optional	Cluster Management	Modify the size of the cluster	See <a href="#">Manage Clusters</a>
9	Optional	Troubleshooting	<ul style="list-style-type: none"> <li>View log archives</li> <li>Download logs for troubleshooting purposes</li> </ul>	<ul style="list-style-type: none"> <li>See <a href="#">Log Archive</a></li> <li>See <a href="#">Monitor Modules</a></li> </ul>
10	Optional	Tenant/Sub-tenant Management	Manage your own tenant or create additional sub-tenants	See <a href="#">Manage Tenants</a>
			Add users as additional tenant administrators to a group	See <a href="#">Create and Assign Groups</a>
11	Optional	Admin Management	Backup CloudCenter Suite	See <a href="#">Backup</a>
			Restore CloudCenter Suite	See <a href="#">Restore</a>
			Setup Isolated (Air Gap) environment	See <a href="#">Without Internet Access</a>



# Initial Administrator Setup

## Initial Administrator Setup

- [Overview](#)
- [The Suite Administrator](#)
- [Configure an Admin User and Tenant](#)

Once the Suite Admin is installed you must perform the following tasks:

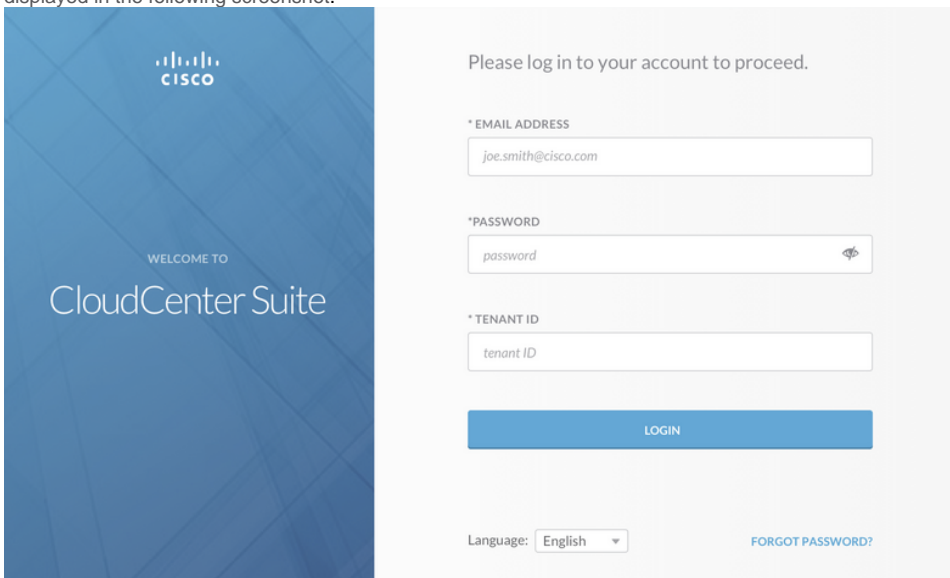
- Note or bookmark the IP address for the Suite Admin console.
- Set up the credentials for the Suite administrator.
- Configure a Root tenant.

As the administrator for the Suite Admin, you can perform the following tasks from the Suite Admin dashboard:

- [Install Module\(s\)](#)
- [Create and Manage Users](#), including tenants and tenant administrators
- [Create and Assign Groups](#), including user-group(s) association
- [Configure Smart Licenses](#)
- [Manage Clusters](#), if the cluster was created by the suite administrator

To configure the admin user and tenant, follow this procedure:

1. Navigate to the Suite Admin console and complete the **Admin User and Tenant Credentials** form to enter details for the root user and tenant as displayed in the following screenshot.



The screenshot shows a login form for the Cisco CloudCenter Suite. On the left, there is a blue sidebar with the Cisco logo and the text 'WELCOME TO CloudCenter Suite'. The main content area is white and contains the following elements:

- A heading: 'Please log in to your account to proceed.'
- Three input fields:
  - \* EMAIL ADDRESS: Contains the text 'joe.smith@cisco.com'.
  - \* PASSWORD: Contains the text 'password' and has an eye icon for toggling visibility.
  - \* TENANT ID: Contains the text 'tenant ID'.
- A blue button labeled 'LOGIN'.
- At the bottom left, a 'Language: English' dropdown menu.
- At the bottom right, a blue link labeled 'FORGOT PASSWORD?'.

2. Besides the First and Last Name, Email Address, Password, Company Name, and Company Logo (defaults to the Cisco logo), you must enter a Tenant ID of your choice so you can log into the Suite Admin using this Tenant ID and password.
3. Click **Done** to save your settings and launch the Suite Admin Dashboard as displayed in the following screenshot.

The screenshot displays the Cisco Suite Admin dashboard. At the top, the Cisco logo and 'SUITE ADMIN' are visible on the left, and a navigation bar on the right contains a warning icon, a cloud icon, a mail icon, a bell icon, a refresh icon, and a user profile icon labeled 'PT Welcome, Puj'. A dark teal sidebar on the left contains icons for user management, organization, network, search, and settings. The main content area is titled 'Modules' and contains four white cards, each representing a module. Each card features a cube icon, the module name, and the version 'v5.1.0' with an installation date of 'Aug 07, 2019'. The 'Suite Admin' card includes a small blue dropdown arrow in the top right corner.

Module Name	Version	Installation Date
Suite Admin	v5.1.0	Aug 07, 2019
Workload Manager	v5.1.0	Aug 07, 2019
Cost Optimizer	v5.1.0	Aug 07, 2019
Action Orchestrator	v5.1.0	Aug 07, 2019

# Kubernetes Cluster Management

## Kubernetes Cluster Management

- [Cluster Status](#)
- [Manage Clusters](#)

# Cluster Status

## Cluster Status

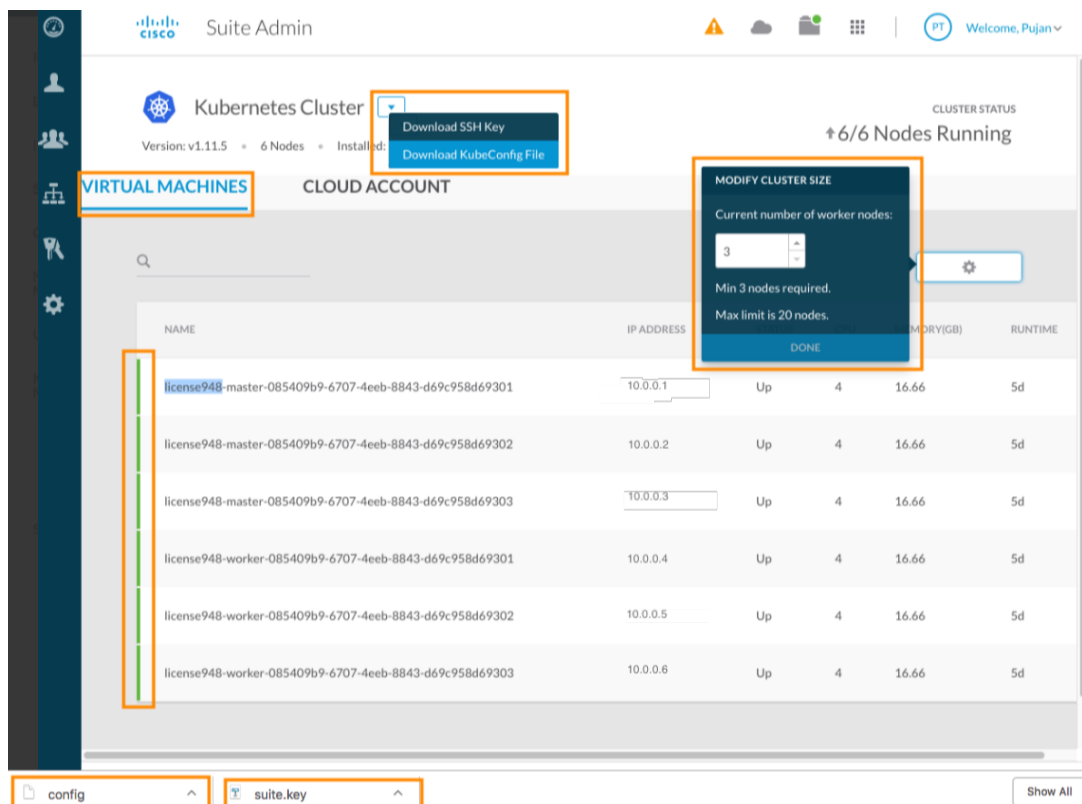
- [Overview](#)
- [Requirements](#)
- [The Cloud Icon Details](#)
- [Kubernetes Cluster Actions](#)
- [Modify Cluster Size](#)
- [Virtual Machines](#)

You can view the status of a Kubernetes cluster by clicking the *cloud* icon located in the header of the [Suite Admin Dashboard](#). The Cluster status popup displays. Click **View Details** to view detailed information about each node in the cluster.

 Kubernetes Cluster Management is already incorporated in the CloudCenter Suite SaaS offer, see [SaaS Access](#) for additional details.

For private clouds, the HA cluster requires a minimum of 2 out of 3 master nodes to be running at any point, for the cluster to function as designed.

Click the *cloud* icon to view and verify the number of nodes in the Kubernetes cluster. The **View Details** page displays detailed information about each node in the cluster. This information is retrieved from the Kubernetes cluster after you install the CloudCenter Suite. The following screenshot displays details within this page.



The screenshot shows the Suite Admin interface for a Kubernetes Cluster. The cluster is named 'Kubernetes Cluster' and has version v1.11.5, 6 nodes, and is installed on a cloud account. The cluster status is '6/6 Nodes Running'. A 'VIRTUAL MACHINES' tab is selected, showing a table of nodes. A 'MODIFY CLUSTER SIZE' dialog box is open, showing the current number of worker nodes (3) and a minimum requirement of 2 nodes. The cluster status is '6/6 Nodes Running'.

NAME	IP ADDRESS	STATUS	CPU	MEMORY (GB)	RUNTIME
license948-master-085409b9-6707-4eeb-8843-d69c958d69301	10.0.0.1	Up	4	16.66	5d
license948-master-085409b9-6707-4eeb-8843-d69c958d69302	10.0.0.2	Up	4	16.66	5d
license948-master-085409b9-6707-4eeb-8843-d69c958d69303	10.0.0.3	Up	4	16.66	5d
license948-worker-085409b9-6707-4eeb-8843-d69c958d69301	10.0.0.4	Up	4	16.66	5d
license948-worker-085409b9-6707-4eeb-8843-d69c958d69302	10.0.0.5	Up	4	16.66	5d
license948-worker-085409b9-6707-4eeb-8843-d69c958d69303	10.0.0.6	Up	4	16.66	5d

The cluster-level actions allow you to download the following files.

- The SSH key file is used to connect to the cluster.
- The KubeConfig file is used to view cluster information.

Based on your environment requirements, you can modify the Kubernetes cluster size from the Suite Admin. See [Manage Clusters](#) for additional details.

This tab displays the VMs that make up the Kubernetes cluster accessed from this instance of CloudCenter Suite.

The colored status indicators identify the state of each VM in your Kubernetes cluster as described in the following table.

Cluster Status Color	Indication
Green	The node is functioning.
Red	The node is not functioning.

The color merely indicates the health of your Kubernetes cluster so you can make the required changes to your Kubernetes setup as required by your environment.




# Manage Clusters

## Manage Clusters

- [Overview](#)
- [Scale Up](#)
- [Scale Down](#)
- [Reconfigure Cloud Credentials](#)


If a cluster was created by the suite administrator as described in [Initial Administrator Setup](#), then this suite administrator can manage those clusters. Managing a cluster includes the following tasks.

- Scale this cluster.
- Monitor the cluster by viewing alerts.

 Suite administrators can only manage clusters that they installed.

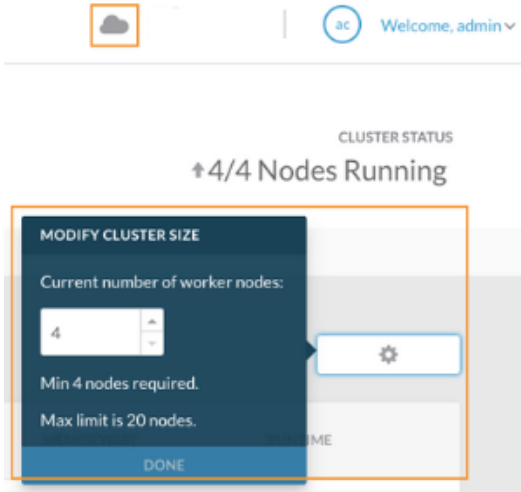
The suite administrator's ability to view a cluster is indicated by the green circle on the **cloud icon**. Clicking this icon provides additional information as displayed in the following screenshot.

 Kubernetes Cluster Management is already incorporated in the CloudCenter Suite SaaS offer, see [SaaS Access](#) for additional details.

 If you setup the CloudCenter Suite using static IPs, verify that the static IP range has free IPs available to support scale up operations. If IPs are not available in the static IP range (defined during installation) then the scale up process will not take place.

To increase the number of nodes in your cluster, perform this procedure.


1. Navigate to the [Suite Admin Dashboard](#) > **Tenants** page.
2. Click the **cloud icon** to access the [Cluster Status](#) > **View Details** page.
3. In the Kubernetes Cluster page, click the wheel icon to display the **Modify Cluster Size** popup as displayed in the following screenshot.



4. Increase the number as required in the **Current number of worker nodes:** field. You will see the status bar list a *Scaling operation successful* alert. It takes a few minutes to increase the node count.
  - Initially, the node will be in the red state while it is still initializing. Once it has initialized, it will turn green.
  - The Runtime displays the length of time that this node has been running:
    - h = Upto 24 hours
    - d = Any number of days
  - The Status can only be up (red) or down (green).
  - The memory and CPU details are displayed as available in the Kubernetes cluster.
  - When complete, you see a subsequent alert notifying you of the Cluster node being added.

You have now increased the number of nodes in your cluster.

While you can scale up the number of nodes in the Kubernetes cluster from the Suite Admin, you *cannot* scale down using this process.

 **OpenStack**

If you installed CloudCenter Suite 5.1.1 as a fresh installation, this feature is not available in OpenStack environments.

If you upgraded CloudCenter Suite from 5.0.x to 5.1.0 or 5.1.1, the Cloud Account section is preserved and you can update the password.

 **vSphere**

If you have updated your password in the vSphere console, be sure to update it in the Cloud Accounts tab (in the Kubernetes Cluster page), before the vSphere lockout period takes effect.

If you do not update the password, be aware that the vSphere policy will prevent you from proceeding with your CloudCenter Suite configuration and CloudCenter Suite will continue with its polling attempts with vSphere.

The Cloud Accounts tab, provides a way to change your cloud credentials for the [cloud where the CloudCenter Suite is installed](#).

You can change your cloud account password based on your cloud credentials for each supported cloud as listed in [New Cluster Installation](#).

# Configure Smart Licenses

## Configure Smart Licenses

- [Overview](#)
- [Cisco Smart Software Manager](#)
  - [Virtual Accounts](#)
  - [Smart Call Home](#)
- [Configuring Cisco Smart Software Licensing](#)
  - [Request a Smart Account](#)
  - [Adding Users to a Smart Account](#)
- [License Usage and Compliance](#)
- [Workflow of Cisco Smart Software Licensing](#)
  - [Generating a Registration Token](#)
  - [Configuring Transport Settings](#)
  - [Registering a CloudCenter Suite License](#)
  - [Renewing Authorization](#)
  - [Re-Registering a CloudCenter Suite License](#)
  - [De-Registering a CloudCenter Suite License](#)
- [Enable for Production](#)
- [Troubleshooting Licensing Issues](#)
  - [Invalid Token](#)
  - [Download Logs](#)

CloudCenter Suite integrates with the [Cisco Smart Software Licensing](#) solution. The CloudCenter Suite is available for a 90-day evaluation period after which, you must register with Cisco Smart Software Manager.

The number of licenses required depends on your deployment scenario. For example, the Workload Manager and Cost Optimizer define entitlements based on features used in those modules. These entitlements may apply to the use of a specific public/private cloud, the number of management units used when deploying applications (VMs and containers), the options purchased (essentials, advanced, premium), and so forth.



Smart licenses are already incorporated in the CloudCenter Suite SaaS offer, see [SaaS Access](#) for additional details.

Cisco Smart Software Manager (Cisco SSM) enables the management of software licenses and Smart Account from a single portal. This interface allows you to activate your product, manage entitlements, renew and upgrade software. You must have a functioning Smart Account to complete the registration process and will need to exchange three key elements with the Cisco Smart Software Manager over HTTPS:

- **Trusted Unique Identifier** – This is the Product ID (SUDI/SUVI/ID).
- **Organizational Identifier** – In a numerical format to associate product with a Smart / Virtual Account.
- **Licenses consumed** – Allows the Cisco Smart Software Manager to understand the license type and level of consumption.

## Virtual Accounts

A Smart Account provides a single location for all Smart enabled products and entitlements. It assists to speed procurement, deployment and maintenance of Cisco Software. When creating a Smart Account the submitter must have the authority to represent the requesting organization. After submitting the request goes through a brief approval.

A Virtual Account exists as a sub-account within the Smart Account. Virtual Accounts are a customer defined structure based on organizational layout, business function, geography or any defined hierarchy. They are created and maintained by the Smart Account administrator(s).

## Smart Call Home

Smart Call Home is feature to communicate with the Cisco Smart Software Manager. By default, Smart Call Home is enabled when you configure Smart Software Licensing. Smart Call Home creates a Cisco TAC-1 profile and sends associated Smart Call Home messages after the enablement. For platforms with Smart Software Licensing enabled by default, call-home is also enabled by default with associated messages.

You need to configure Cisco Smart Software Licensing to easily procure, deploy, and manage licenses for your CloudCenter Suite.

Smart Licensing is a cloud-based approach to licensing. The solution simplifies the purchase, deployment and management of Cisco software assets. Entitlements are purchased through your Cisco account via Cisco Commerce Workspace (CCW) and immediately deposited into a *Virtual Account* for usage. This process eliminates the need to install license files on every device using the product. Products that are smart enabled communicate directly to Cisco to report consumption. A single location is available to customers to manage Cisco software licenses – the Cisco SSM. License ownership and consumption are readily available to help make better purchase decision based on consumption or business need.

Cisco SSM enables you to manage your Cisco Smart Software Licenses from one centralized website. With Cisco SSM, you can organize and view your licenses into *Virtual Account* groups. You can also use Cisco SSM to transfer licenses between virtual accounts as needed. You can access Cisco SSM from the Cisco Software Central homepage at [software.cisco.com](https://software.cisco.com), under Smart Licensing.

If you do not want to manage licenses using Cisco SSM, either for policy reasons or network availability reasons, you can choose to install Cisco SSM Satellite at your premises. CloudCenter Suite registers and reports license consumption to the Cisco SSM Satellite as it does to Cisco SSM. Cisco SSM Satellite coordinates with the Cisco Smart Software Manager to manage software licenses on premises. Devices register locally to report license ownership and consumption.



Ensure that you use Cisco SSM Satellite version 5.0 or later. For more information on installing and configuring Cisco SSM Satellite, refer to <http://www.cisco.com/go/smartsatellite>.

## Request a Smart Account

The creation of a new Smart Account is a one-time event and subsequent management of users is a capability provided through the tool. To request a Smart Account, visit [software.cisco.com](http://software.cisco.com) and follow this process.

1. After logging in, select **Request a Smart Account** in the Administration section as displayed in the following screenshot.



2. Select the type of Smart Account to create using one of two options as displayed in the following screenshot.

### Create Account

Would you like to create the Smart Account now?

- Yes, I have authority to represent my company and want to create the Smart Account.
- No, the person specified below will create the account:

\* Email Address:

Message to Creator:

- Individual Smart Account requiring agreement to represent your company. By creating this Smart Account you agree to authorize, create, and manage product and service entitlements, users, and roles on behalf of your organization.
- Create the account on someone else's behalf

3. Provide the required domain identifier and the preferred account name as displayed in the following screenshot.

### Account Information

The Account Domain Identifier will be used to **uniquely identify the account**. It is based on the email address of the person creating the account by default and must belong to the company that will own this account. [Learn More](#)

The screenshot shows the 'Account Information' form. It has two main fields: 'Account Domain Identifier' with the value 'domainidentifier.com' and an 'Edit' link, and 'Account Name' with a text input field. The entire form is highlighted with a red box.

4. The account request requires approval for the Account Domain Identifier as displayed in the following screenshot. An email will be sent to the requester to complete the setup process.

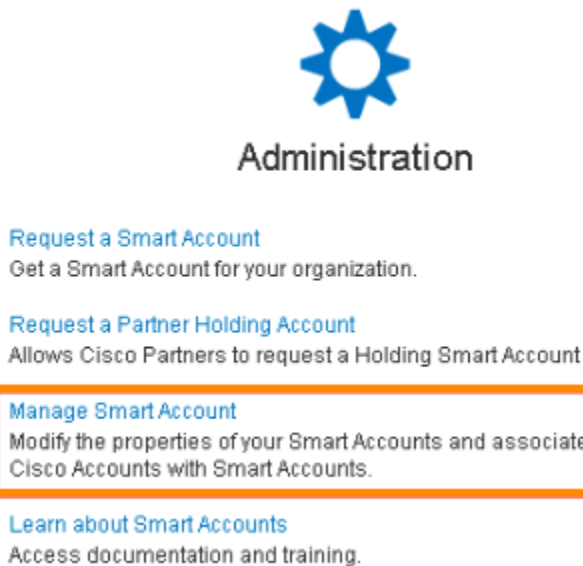
**Smart Account Request Pending**

The account setup process is pending approval of an Account Domain Identifier. You will receive an email confirmation and a Cisco representative will contact you at the number provided below.


## Adding Users to a Smart Account

Smart Account user management is available in the Administration section of [software.cisco.com](https://software.cisco.com). To add a new user to a Smart Account, follow this process.

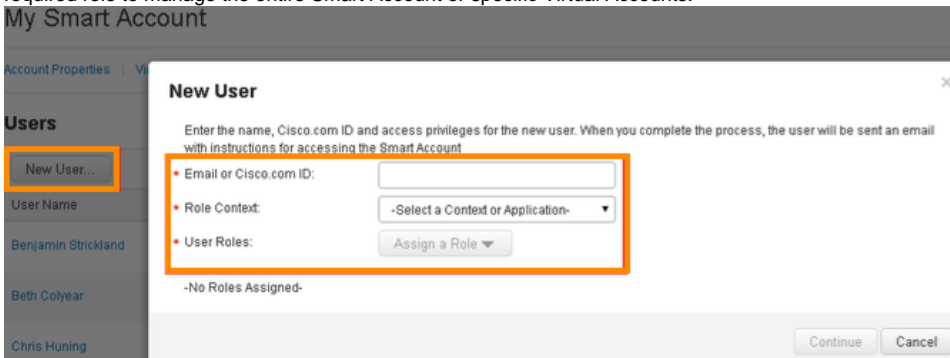
1. After logging in, select **Manage Smart Account** in the Administration section as displayed in the following screenshot.



2. Select the **Users** tab as displayed in the following screenshot.



3. Select **New User** and provide the required email address, cisco.com ID, and role as displayed in the following screenshot. You can select the required role to manage the entire Smart Account or specific Virtual Accounts.



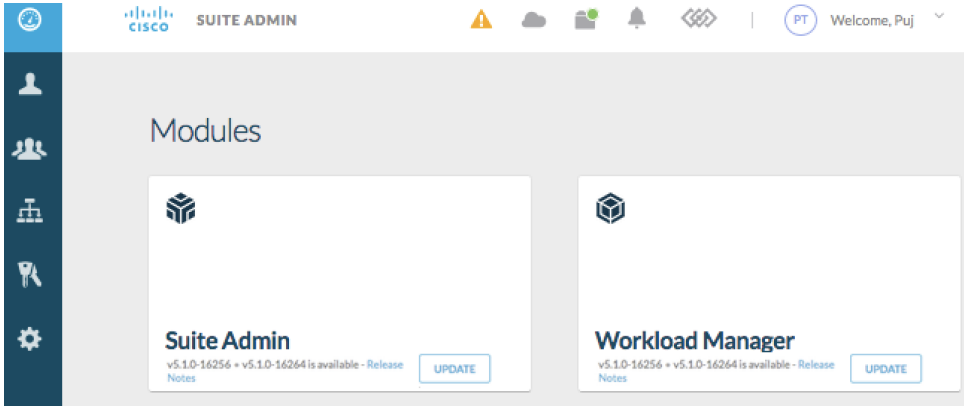
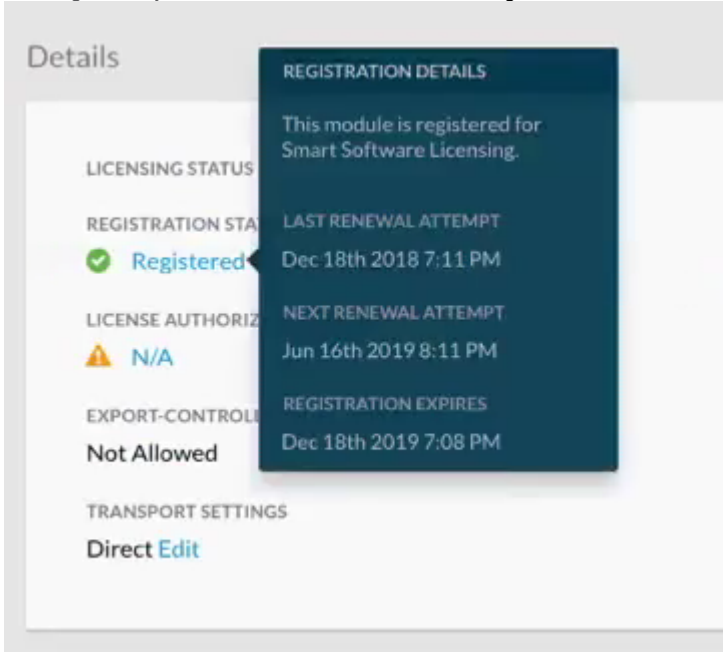
4. Click **Continue** to complete the process.

Once you register CloudCenter Suite with Cisco SSM, you will receive the CloudCenter Suite License.

If you use specific resources, the CloudCenter Suite reports each usage to the Cisco SSM to tally the number of times that this resource was used and report it in the **Count** column. By verifying this usage count, Cisco SSM calculates the license usage and compliance.

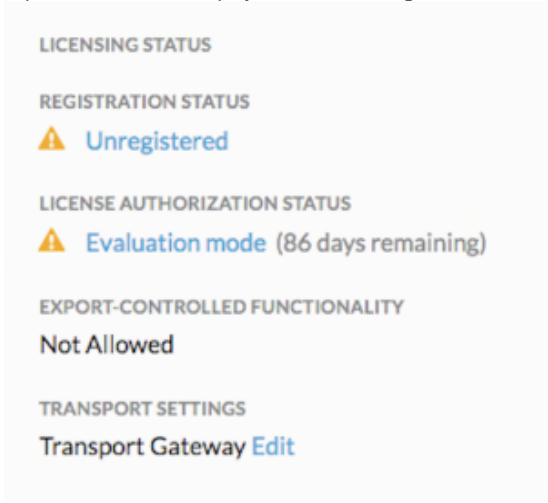
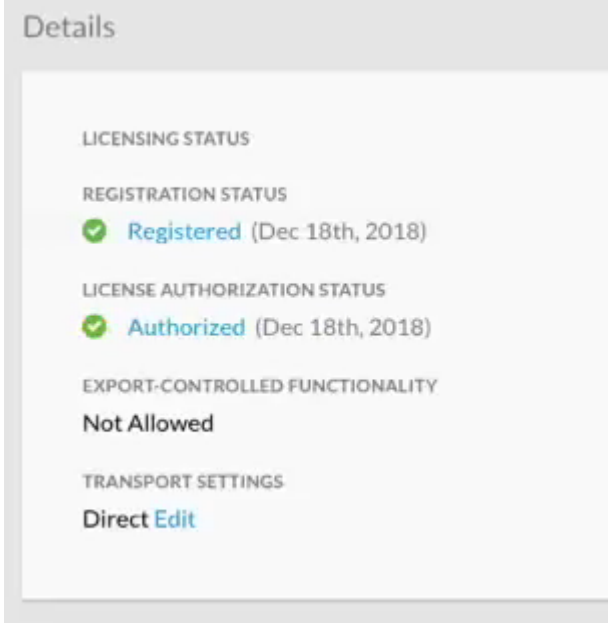
Cisco SSM or Cisco SSM Satellite totals the license requirements for all your CloudCenter Suite instances and compares the total license usage to the number of licenses purchased, on a daily basis.

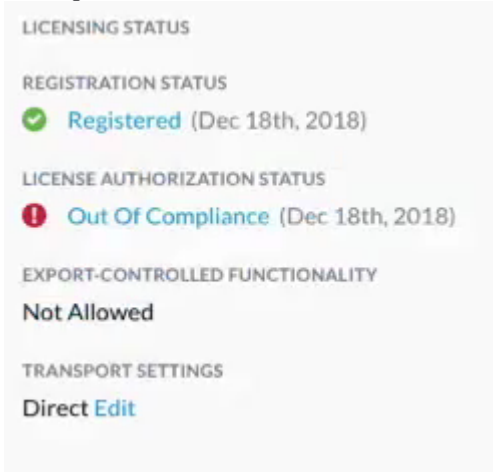
After the data synchronization, your CloudCenter Suite instance displays one of the **Registration Status** indicators listed in the following table.

Registration Status	Description
<p><b>Unregistered</b></p>	<p>The Smart Software Licensing is running in Evaluation mode and you have not yet registered the CloudCenter Suite. This status is identified in the following screenshot when you click on the Licensing icon – the orange, exclamation icon in the following screenshot.</p> 
<p><b>Registered</b></p>	<p>The product registration was completed and an ID certificate was received and will be used for future communication with the Cisco licensing authority. This status is identified in the following screenshot.</p> 

After the data synchronization, your CloudCenter Suite instance displays one of the **Licensing Authorization Status** indicators as explained in the following table.

License Authorization Status	Description

<p><b>Evaluation Mode</b> (countdown from 90 days)</p>	<p>You must register your CloudCenter Suite instance with Cisco SSM or Cisco SSM Satellite before the 90-day evaluation period expires. This state is displayed in the following screenshot.</p>  <p>LICENSING STATUS</p> <p>REGISTRATION STATUS ⚠ Unregistered</p> <p>LICENSE AUTHORIZATION STATUS ⚠ Evaluation mode (86 days remaining)</p> <p>EXPORT-CONTROLLED FUNCTIONALITY Not Allowed</p> <p>TRANSPORT SETTINGS Transport Gateway <a href="#">Edit</a></p>
<p><b>Authorized</b></p>	<p>The number of licenses purchased is sufficient – Registration is complete and valid and the license consumption has started. This state indicates compliance and is displayed in the following screenshot.</p>  <p>Details</p> <p>LICENSING STATUS</p> <p>REGISTRATION STATUS ✔ Registered (Dec 18th, 2018)</p> <p>LICENSE AUTHORIZATION STATUS ✔ Authorized (Dec 18th, 2018)</p> <p>EXPORT-CONTROLLED FUNCTIONALITY Not Allowed</p> <p>TRANSPORT SETTINGS Direct <a href="#">Edit</a></p>
<p><b>Authorization Expired</b></p>	<p>The product has not communicated with Cisco SSM or Cisco SSM Satellite for a period of 90 days.</p> <p>The product has been unable to communicate with the Cisco SSM for an extended period of time. This state is due to non-communication with Cisco SSM or Cisco SSM Satellite for more than 90 days. The product will attempt to contact the Cisco SSM every hour in order to renew the authorization until the registration period expires.</p>

<b>Out of Compliance</b>	<p>The number of licenses is insufficient – Consumption exceeds available licenses in the Virtual Account. This state is displayed in the following screenshot.</p>  <p>The screenshot shows a licensing status page with the following sections:</p> <ul style="list-style-type: none"> <li><b>LICENSING STATUS</b></li> <li><b>REGISTRATION STATUS</b>: Registered (Dec 18th, 2018) with a green checkmark icon.</li> <li><b>LICENSE AUTHORIZATION STATUS</b>: Out Of Compliance (Dec 18th, 2018) with a red exclamation mark icon.</li> <li><b>EXPORT-CONTROLLED FUNCTIONALITY</b>: Not Allowed</li> <li><b>TRANSPORT SETTINGS</b>: Direct Edit</li> </ul>
--------------------------	--

The following table describes the workflow of Cisco Smart Software Licensing.

Task	See the Related Section
Generate a product instance registration token in your virtual account	<a href="#">Generating a Registration Token</a>
Configure the transport settings using which CloudCenter Suite connects to Cisco SSM or Cisco SSM Satellite	<a href="#">Configuring Transport Settings</a>
Register the CloudCenter Suite instance with Cisco SSM or Cisco SSM Satellite	<a href="#">Registering a CloudCenter Suite License</a>
Manage licenses	<ul style="list-style-type: none"> <li>• <a href="#">Renewing Authorization</a></li> <li>• <a href="#">Re-Registering a CloudCenter Suite License</a></li> <li>• <a href="#">De-Registering a CloudCenter Suite License</a></li> </ul>


## Generating a Registration Token

You need to generate a registration token from Cisco SSM or Cisco SSM Satellite to register the CloudCenter Suite instance.

- ✔ Ensure that you have set up a Smart Account and a Virtual account on Cisco SSM or Cisco SSM Satellite.

To generate a registration token, follow this procedure.

1. Log in to your Smart Account using **Cisco SSM** or Cisco SSM Satellite.
2. Navigate to the Virtual account using which you want to register the CloudCenter Suite instance.
3. If you want to enable higher levels of encryption for the products registered using the registration token, check the **Allow export-controlled** functionality on the products registered with this token check box.

 This option is available only if your smart account is enabled for Export Control.

4. Click **New Token** to generate a registration token.
5. Copy and save the token so you can use it when you register your CloudCenter Suite instance.
6. For more information on registering your CloudCenter Suite instance, see [Registering a CloudCenter Suite License](#).

## Configuring Transport Settings

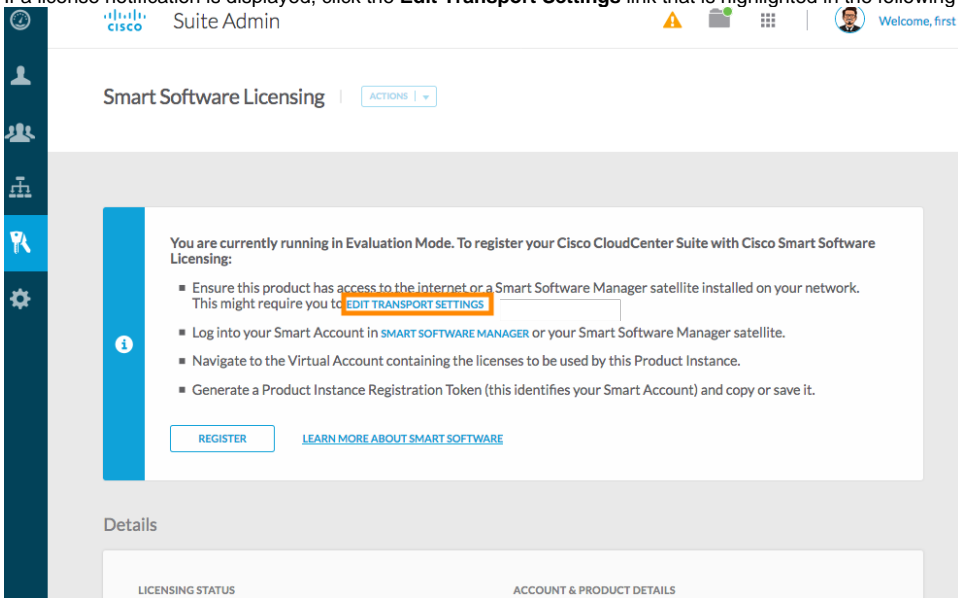
By default, CloudCenter Suite directly communicates with the Cisco SSM. You can modify the mode of communication by configuring the transport settings.

- ✔ Ensure that you have obtained the registration token for the CloudCenter Suite instance.


To configure the transport settings, follow this procedure.



1. Navigate to the [Suite Admin Dashboard](#).
2. Click **Licensing** in the left tree pane. If you are running CloudCenter Suite in the Evaluation mode, a license notification is displayed on the Smart Software Licensing pane.
3. If a license notification is displayed, click the **Edit Transport Settings** link that is highlighted in the following screenshot.



- Alternatively, click the **Licensing Status** tab, and then click the **View/Edit** link that appears under Transport Settings.
4. In the Transport Settings dialog box displayed in the following screenshot, perform one of these steps:



- To configure CloudCenter Suite to send the license usage information to Cisco SSM using the Internet (default):
    - a. Click the **Direct** switch to communicate directly using the Cisco URL.
    - b. Configure a DNS on CloudCenter Suite to resolve [tools.cisco.com](https://tools.cisco.com).
  - To configure CloudCenter Suite to send the license usage information to Cisco SSM using the Cisco SSM Satellite:
    - a. Click the **Transport Gateway** button.
    - b. Enter the URL of the Cisco SSM Satellite.
  - To configure CloudCenter Suite to send the license usage information to Cisco SSM using a proxy server. For example, an off-the-shelf proxy, such as Cisco Transport Gateway or Apache:
    - a. Toggle the **HTTP/HTTPS Proxy** switch.
    - b. Enter the IP address and port number of the proxy server.
5. Click **Save**.

## Registering a CloudCenter Suite License

You need to register your CloudCenter Suite instance with Cisco SSM or Cisco SSM Satellite before the 90-day evaluation period expires.

- ✔ Ensure that you have configured the transport settings.

To register the CloudCenter Suite license, follow this procedure.

1. Navigate to the [Suite Admin Dashboard](#).
2. Click **Licensing** in the left tree pane.
3. In the license notification, click **Register**. The Smart Software Licensing Product Registration dialog box appears.
4. In the Product Instance Registration Token field, paste the registration token that you generated using the Cisco SSM or Cisco SSM Satellite. For more information on generating a registration token, see [Generating a Registration Token](#).
5. Click **Register** to complete the registration process. The CloudCenter Suite sends a request to Cisco SSM or Cisco SSM Satellite to check the registration status and Cisco SSM or Cisco SSM Satellite reports back the status to CloudCenter Suite, on a daily basis. If registering the token fails, you can re-register the CloudCenter Suite instance using a new token. For more information on re-registering CloudCenter Suite, see [Re-Registering a CloudCenter Suite License](#).

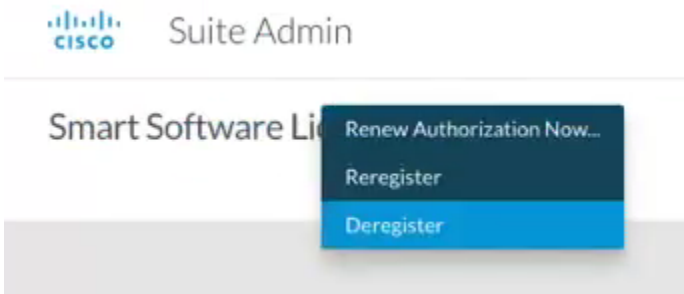
## Renewing Authorization

By default, the authorization is automatically renewed every 30 days. However, CloudCenter Suite allows a user to manually initiate the authorization renew in case the automatic renewal process fails. The authorization expires if CloudCenter Suite is not connected to Cisco SSM or Cisco SSM Satellite for 90 days and the licenses consumed by CloudCenter Suite are reclaimed and put back to the license pool.

- ✔ Ensure that the CloudCenter Suite instance is registered with Cisco SSM or Cisco SSM Satellite.

To renew authorization, follow this procedure.

1. Navigate to the [Suite Admin Dashboard](#).
2. Click **Licensing** in the left tree pane.
3. From the Actions drop-down list, choose **Renew Authorization Now** as displayed in the Actions dropdown in the following screenshot.



4. Click **OK** in the Renew Authorization dialog box to confirm authorization renewal. The CloudCenter Suite synchronizes with Cisco SSM or Cisco SSM Satellite to check the license authorization status and Cisco SSM or Cisco SSM Satellite reports back the status to CloudCenter Suite, on a daily basis.

## Re-Registering a CloudCenter Suite License

You can re-register CloudCenter Suite with Cisco SSM or Cisco SSM Satellite by de-registering it and registering it again, or by using a register force option.


- ✔ Ensure that you have obtained a new registration token from Cisco SSM or Cisco SSM Satellite

To re-register CloudCenter Suite license, follow this procedure.

1. Navigate to the [Suite Admin Dashboard](#).
2. Click **Licensing** in the left tree pane.
3. From the Actions drop-down list, choose **Reregister**.
4. In the **Product Instance Registration Token** field of the Smart Software Licensing Product Re-registration dialog box, enter the registration token that you generated using Cisco SSM or Cisco SSM Satellite. For more information on generating a registration token, see [Generating a Registration Token](#).
5. Click **Register** to complete the registration process. The CloudCenter Suite sends a request to Cisco SSM or Cisco SSM Satellite to check the registration status and Cisco SSM or Cisco SSM Satellite reports back the status to CloudCenter Suite, on a daily basis.

## De-Registering a CloudCenter Suite License

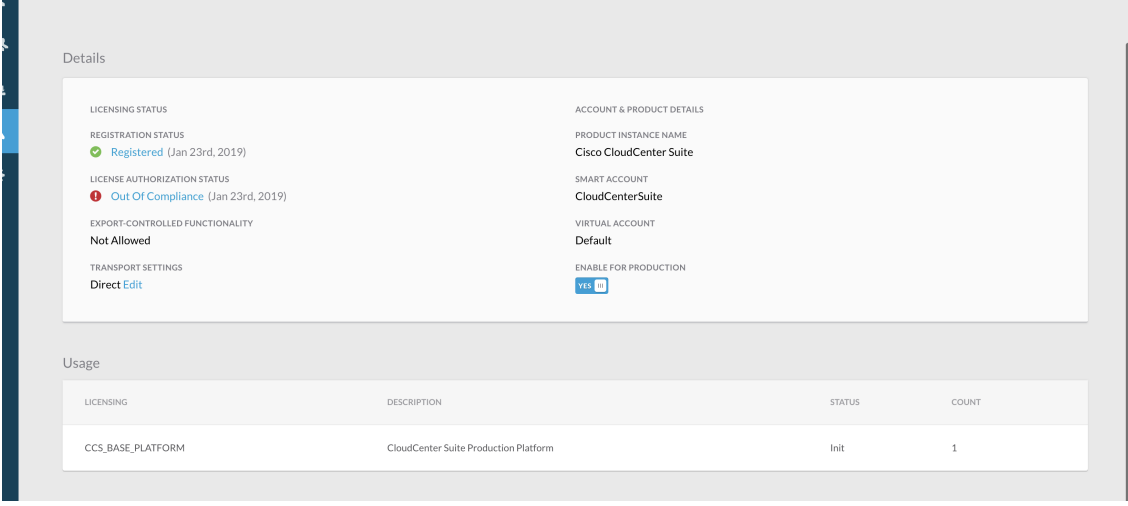
You can de-register the CloudCenter Suite instance from Cisco SSM or Cisco SSM Satellite to release all the licenses from the current Virtual account and the licenses are available for use by other products in the virtual account. De-registering disconnects CloudCenter Suite from Cisco SSM or Cisco SSM Satellite.

 Ensure that the CloudCenter Suite instance is registered with Cisco SSM or Cisco SSM Satellite.

To de-register CloudCenter Suite license, follow this procedure.

1. Navigate to the [Suite Admin Dashboard](#).
2. Click **Licensing** in the left tree pane.
3. From the Actions drop-down list, choose **Deregister**.
4. Click **Deregister** in the confirmation dialog box. The CloudCenter Suite sends a request to Cisco SSM or Cisco SSM Satellite to check the de-registration status and Cisco SSM or Cisco SSM Satellite reports back the status to CloudCenter Suite, on a daily basis.

Toggle the **Enable for Production** switch to use the license in production mode displayed in the following screenshot. When you purchase one license for the CloudCenter Suite, you automatically receive a free non-production license as well. Both modes are independent of each other and you can switch from one mode to the other any number of times.



The screenshot displays the 'Details' section of the CloudCenter Suite Admin interface. It is divided into two main columns: 'LICENSING STATUS' and 'ACCOUNT & PRODUCT DETAILS'.

**LICENSING STATUS:**

- REGISTRATION STATUS:** Registered (Jan 23rd, 2019) with a green checkmark icon.
- LICENSE AUTHORIZATION STATUS:** Out Of Compliance (Jan 23rd, 2019) with a red warning icon.
- EXPORT-CONTROLLED FUNCTIONALITY:** Not Allowed.
- TRANSPORT SETTINGS:** Direct Edit.

**ACCOUNT & PRODUCT DETAILS:**

- PRODUCT INSTANCE NAME:** Cisco CloudCenter Suite
- SMART ACCOUNT:** CloudCenterSuite
- VIRTUAL ACCOUNT:** Default
- ENABLE FOR PRODUCTION:** YES (with a toggle switch icon).

**Usage:**

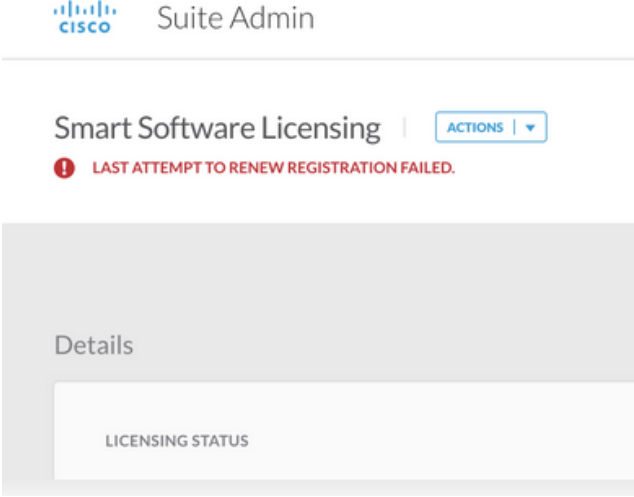
LICENSING	DESCRIPTION	STATUS	COUNT
CCS_BASE_PLATFORM	CloudCenter Suite Production Platform	Init	1

When the CloudCenter Suite is in non-production mode, the entitlement tags do not validate the license for usage, in which case, you can use it for development, testing, or staging purposes.

This section identifies issues that you may encounter when dealing with licenses.

## Invalid Token

When you see the message displayed in the following screenshot for your instance, verify if your token is still valid and if it needs to be renewed.



The screenshot shows the 'Suite Admin' interface. At the top, there is a 'Smart Software Licensing' header with an 'ACTIONS' dropdown menu. Below the header, a red error message is displayed: 'LAST ATTEMPT TO RENEW REGISTRATION FAILED.' Below the error message, there is a 'Details' section with a 'LICENSING STATUS' sub-section.

## Download Logs

If you have any issues with Smart Licenses, download the logs files by using the UI (see [Monitor Modules > Download Logs](#)) or the suite-logs/v2/api-docs (see [Logs Service API Calls](#)) and contact the [Smart License team](#).

# Module Lifecycle Management

## Module Lifecycle Management

- [Install Module](#)
- [Update Module](#)
- [Monitor Modules](#)

# Install Module

## Install Module

- [Overview](#)
- [Requirements](#)
- [Process](#)
- [Free License](#)
- [Module Actions](#)
- [Uninstall a Module](#)
- [Module States](#)

The [Suite Admin Dashboard](#) lists the available modules in the Display pane. If you are installing each module for the first time, you will see the **Install** button enabled. Once installed, each module may be in various lifecycle phases as described in this section.

 Module lifecycle management is already incorporated in the CloudCenter Suite SaaS offer, see [SaaS Access](#) for additional details.

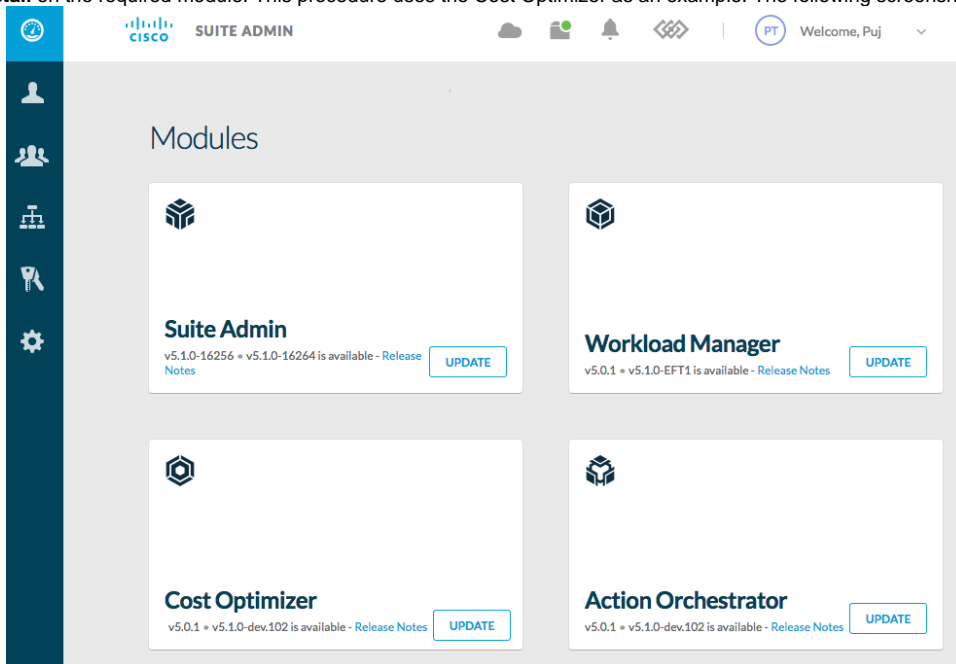
Be sure to adhere to the following requirements:

- If your current cluster does not have sufficient resources to meet the minimum requirements mentioned in the [Prepare Infrastructure](#) section, then the installation process will be blocked and you will need to resolve these issues by scaling up to these requirements (see [Manage Clusters > Scale Up](#) for details).
- Only a suite administrator can install a module. By installing the module, this suite administrator automatically inherits the module admin role as well.
- Be sure to synchronize the server time for all instances running the CloudCenter Suite as this can potentially cause module install or upgrade to fail.

You can install multiple modules simultaneously.


To install a module, follow this procedure.

1. Navigate to the [Suite Admin Dashboard](#).
2. Click **Install** on the required module. This procedure uses the Cost Optimizer as an example. The following screenshot displays the available



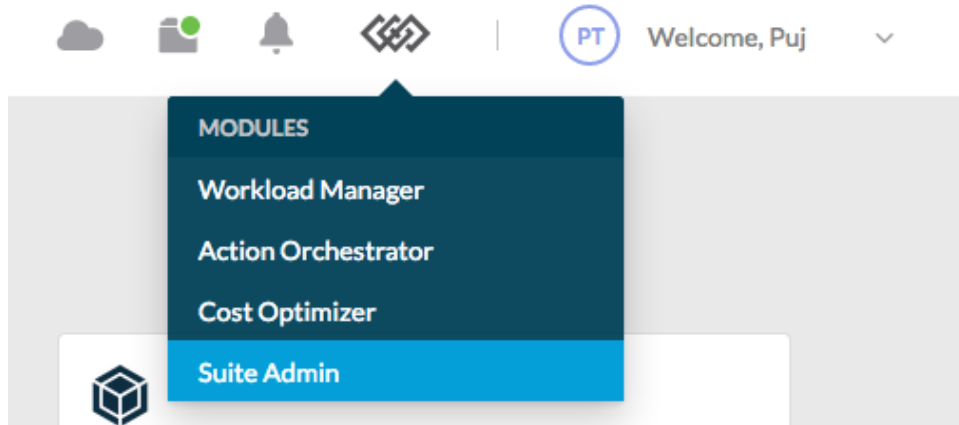
modules.

3. In the **You're updating *module name*** popup, select the required version from the dropdown list.

 Once installed, you cannot revert to a previous version.

4. The module starts its installation process and displays a progress bar indicator.
5. Once Installed, you can perform the following actions:
  - Click a module to [Monitor Modules](#).
  - Open the module or uninstall the module (see the section below).

- Navigate back and forth to other modules and the Suite Admin using the navigation icon in the header as displayed in the following screenshot.



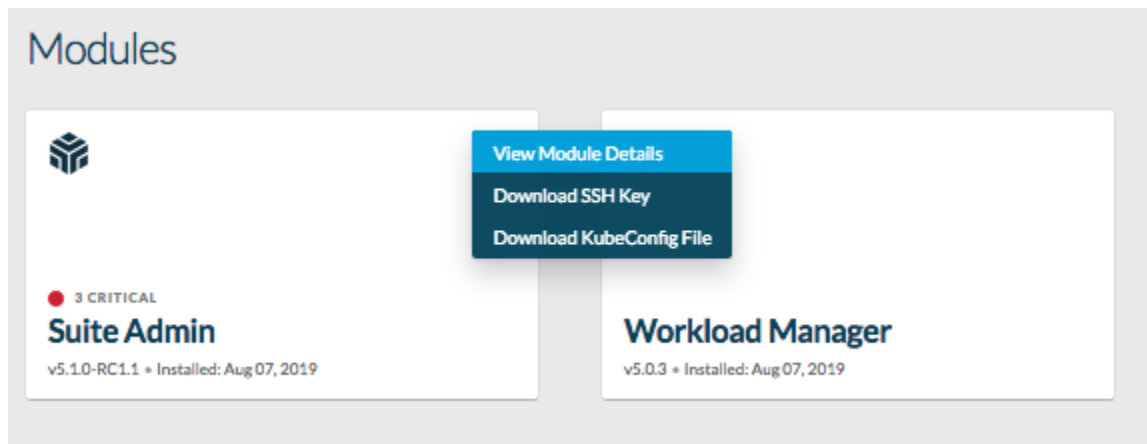
You have now installed one of the modules in the CloudCenter Suite.

When you install any module, you see the countdown for the 90-day free license time remaining for the license in the top left portion of the module. See [Configure Smart Licenses](#) for details.

Once installed, the suite administrator can perform the following actions on a module:

- [Update Module](#)
- [Monitor Modules](#)
- [Configure Smart Licenses](#)
- [Manage Module-Specific Content](#)

The Suite Admin module allows the additional actions displayed in the following screenshot:



- Download SSH Key (used to connect to the cluster).
- Download KubeConfig file (used to view cluster information).
- See [Cluster Status](#) for additional context.

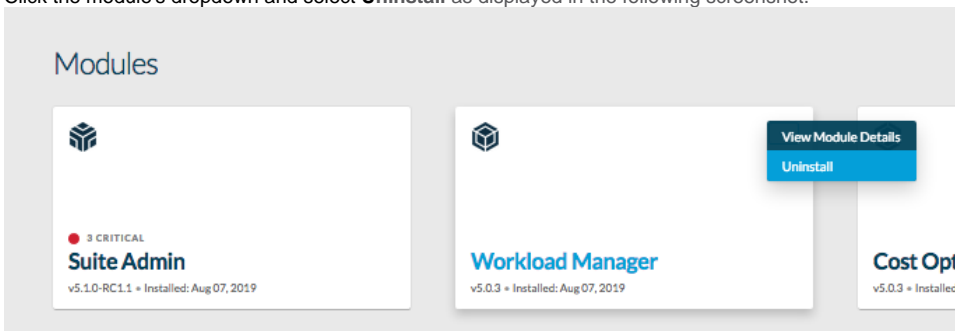


After you uninstall any module, verify that all dependent resources have been deleted.

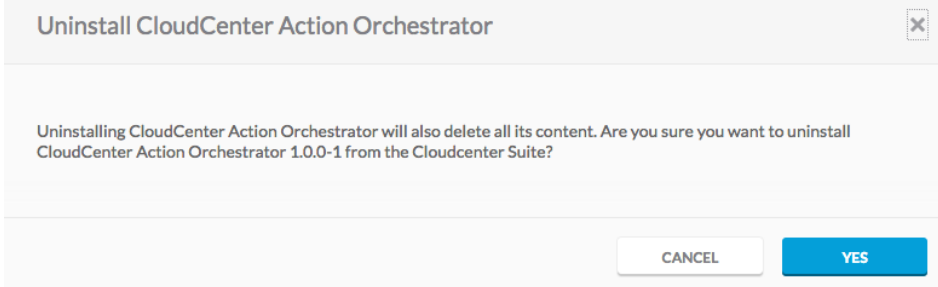
Before re-installing a module that was previously installed, verify that the volumes, secrets, and other dependent details have been cleaned up.

To uninstall a module, follow this procedure.

1. Click the module's dropdown and select **Uninstall** as displayed in the following screenshot.



2. Confirm your intention to uninstall as all your content will be deleted as displayed in the following screenshot.



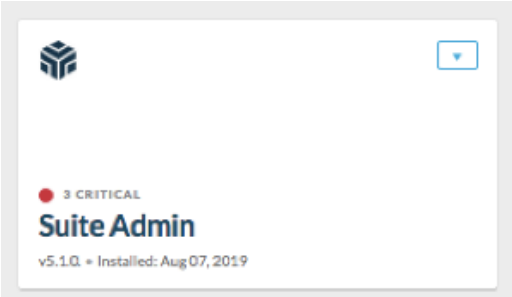



3. The module starts its uninstallation process. Uninstallation takes a few minutes as the CloudCenter Suite cleans up all aspects of the installation.

The following table provides details on the various module states.

State and Screenshot	Description
<p><b>New Installation</b></p>	<p>A new module is available for installation in the <a href="#">Suite Admin Dashboard</a>.</p>
<p><b>Installing (or updating)</b></p>	<p>The module is being installed/updated and the installation process displays a progress bar indicator.</p>



<p><b>Licensed</b></p>  <p><b>Workload Manager</b> v 5.0 - Installed: 05 July 2017</p>	<p>This screenshot identifies a module that is installed, registered, and licensed. See <a href="#">Configure Smart Licenses</a> for details.</p>
<p><b>Update Available</b></p>  <p><b>Suite Admin</b> v5.1.0-16256 - v5.1.0-16264 is available - <a href="#">Release Notes</a> <span>UPDATE</span></p>	<p>Once a new software version becomes available, the module displays the new version availability and provides a link to the documentation website. See <a href="#">Update Module</a> for details.</p> <p>The release notes link for the available release is directly linked to the release notes for each module.</p> <p>The dropdown list also provides additional options for each module</p>
<p><b>Alerts</b></p> 	<p>When alerts are generated, they are displayed in the Suite Admin Dashboard (dropdown list for this module) &gt; <b>View Module Details</b> &gt; <b>Alerts</b> tab.</p> <p>The number of alerts are also identified in the corresponding module tile that are displayed in the Suite Admin Dashboard (the screenshot identifies that 3 Warning alerts are available for this module)</p> <p>See <a href="#">Monitor Modules</a> for details.</p>
<p><b>Validation Error</b></p>  <p><b>Workload Manager</b> Failed to install - Please Try again</p>	<p>The module installation resulted in an error. See <a href="#">Troubleshoot Suite Admin</a> for additional details.</p>

# Update Module

## Update Module

- [Overview](#)
- [Considerations](#)
- [Limitations](#)
- [Process](#)
- [Configuring Memory Limits for Modules](#)
- [Module Actions](#)

The suite administrator can only upgrade the module to later versions of the software and will not be able to revert to an earlier version of the software.

 Module lifecycle management is already incorporated in the CloudCenter Suite SaaS offer, see [SaaS Access](#) for additional details.

Before updating a module, see the following module considerations:

- [Workload Manager Installation Overview](#) > *Module Update Considerations*
- [Cost Optimizer Overview](#) > *Module Update Considerations*

Only a suite administrator can update a module.

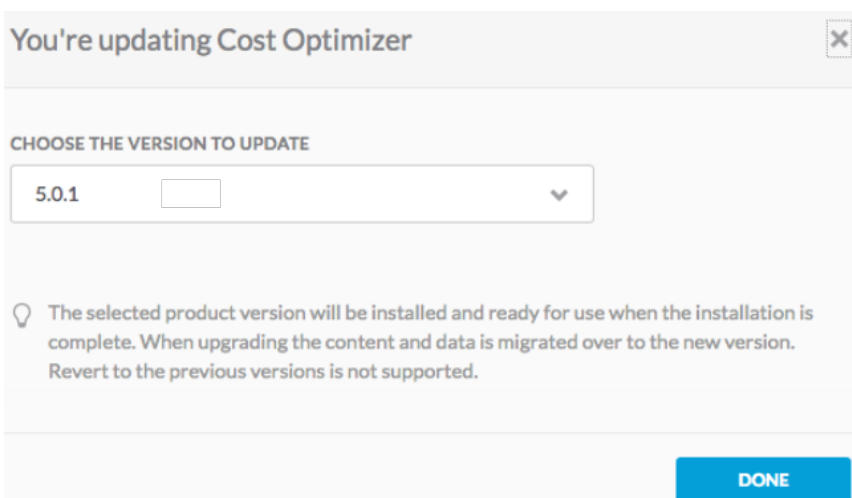
Once a new software version becomes available, the module displays the new version availability and provides a link to the documentation website.



- Before updating any module, verify that you have un-allocated CPU/Memory in your cluster to ensure that your environment has free CPU/Memory – a module-update scenario requires additional resources for the old pod to continue running until the new pod initializes and takes over. This additional resource requirement is temporary and only required while a module update is in Progress. After the module is updated, the additional resources are no longer needed.
- You must update the Suite Admin module before you update any other CloudCenter Suite module.
- **Update only one module at a time. If you simultaneously update more than one module, your update process may fail due to limited resource availability. See [Prepare Infrastructure](#) for additional context.**
- You may see one or more error messages during the update process. Be aware that these messages will not affect the update itself.

To update a module, follow this process.

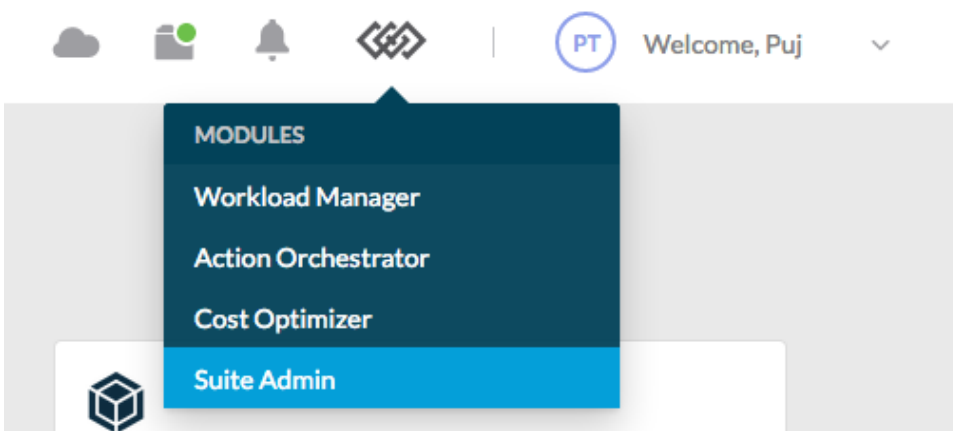
1. Navigate to the [Suite Admin Dashboard](#).
2. Select the required version and click **Done** to upgrade this module. The following screenshot displays Cost Optimizer as an example. All available releases are displayed in the dropdown list in descending order with the latest version at the start of the list.



3. The module starts its upgrade process and displays a progress bar indicator.
4. Once Installed, you can click the module to access the details of that module

or

Navigate to other modules using the module navigation icon in the header as displayed in the following screenshot.



You have now updated the modules in the CloudCenter Suite.

In some Cloud Center Suite 5.x environments it may be necessary to increase CPU and memory limits for the *common-framework-suite-prod-mgmt* pod prior to upgrade of any CloudCenter Suite module. The instructions below explain how to configure the new limits.

1. From the Suite Admin Dashboard download the KubeConfig file for your CloudCenter Suite deployment. Save to your local machine.

2. Create a file named `ccs-upgrade.yaml` with the following content. Save to your local machine.

```

apiVersion: extensions/v1beta1
kind: Deployment
metadata:
  name: common-framework-suite-prod-mgmt
spec:
  template:
    metadata:
      labels:
        app: suite-prod-mgmt
        release: common-framework
    spec:
      containers:
        -
          name: suite-prod-mgmt
          resources:
            limits:
              cpu: 200m
              memory: 256Mi
            requests:
              cpu: 200m
              memory: 256Mi

```

3. Verify if kubectl is installed and has connectivity to the CloudCenter Suite deployment by executing the following command to list all pods.

```
kubectl get pods -n cisco --kubeconfig=<PATH_TO_KUBECONFIG>
```

4. Now apply the new CPU and memory limits defined in the yaml file created in Step 2.

```
kubectl apply -f ccs-upgrade.yaml -n cisco --kubeconfig=<PATH_TO_KUBECONFIG>
```

The output from the command will be:

```
deployment.extensions/common-framework-suite-prod-mgmt configured
```

5. Optionally, execute the command to verify that the CPU and memory limits have been configured.

```

kubectl
  get deployment common-framework-suite-prod-mgmt -n cisco
  --kubeconfig=<PATH_TO_KUBECONFIG> -o yaml --export >
  common-framework-suite-prod-mgmt-deployment.yaml

```

6. Open the file *common-framework-suite-prod-mgmt-deployment.yaml* and verify the values have been changed to:

```

resources:
  limits:
    cpu: 200m
    memory: 256Mi
  requests:
    cpu: 200m
    memory: 256Mi

```


Once a module is upgraded, the suite administrator can perform the following actions on a module:

- [Monitor Modules](#)
- [Configure Smart Licenses](#)
- [Manage Module-Specific Content](#)

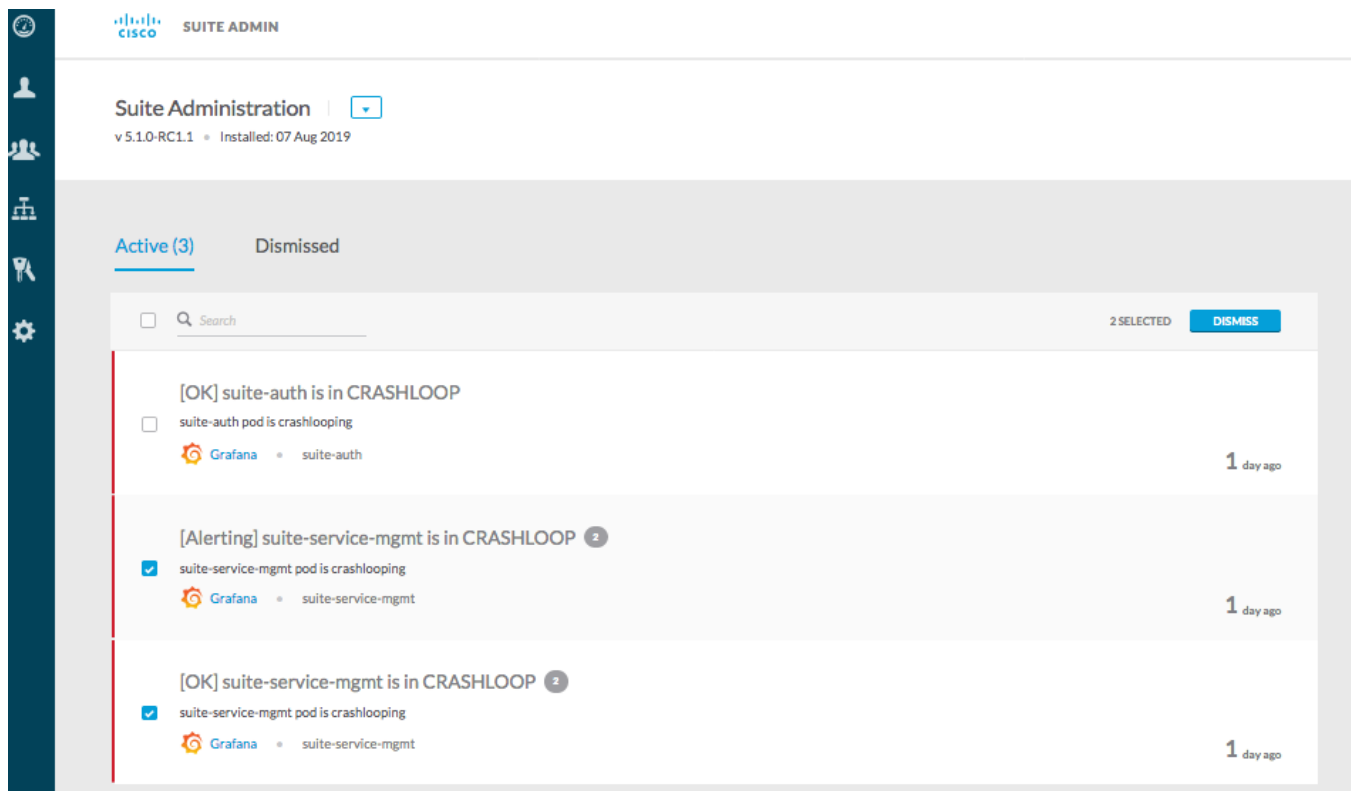
# Monitor Modules

## Monitor Modules

- [Overview](#)
- [Accessing a Module](#)
- [View Logs in Kibana](#)
- [Download Logs](#)
- [The Grafana Dashboard Alert](#)
- [Default Alert Categories](#)
- [Type of Alerts](#)
- [Alert Severities](#)
- [Viewing Alerts in Grafana](#)

 For SaaS customers, Module lifecycle management is managed by the CloudCenter Suite operations teams and not exposed publicly; see [SaaS Access](#) for additional details.

Once Installed, you can click a module to access the **Module Details** page displayed in the following screenshot. If you click the Workload Manager, the following screenshot displays the corresponding page to monitor this module.

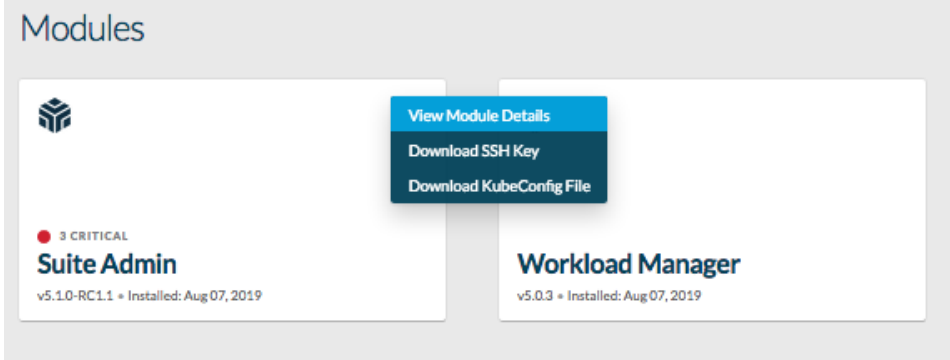


The screenshot shows the Cisco Suite Administration interface. The top navigation bar includes the Cisco logo and 'SUITE ADMIN'. Below this, the 'Suite Administration' section is visible, showing version 'v 5.1.0-RC1.1' and installation date '07 Aug 2019'. The main content area is divided into 'Active (3)' and 'Dismissed' tabs. Under the 'Active (3)' tab, there is a search bar and a '2 SELECTED' indicator with a 'DISMISS' button. The active alerts are listed as follows:

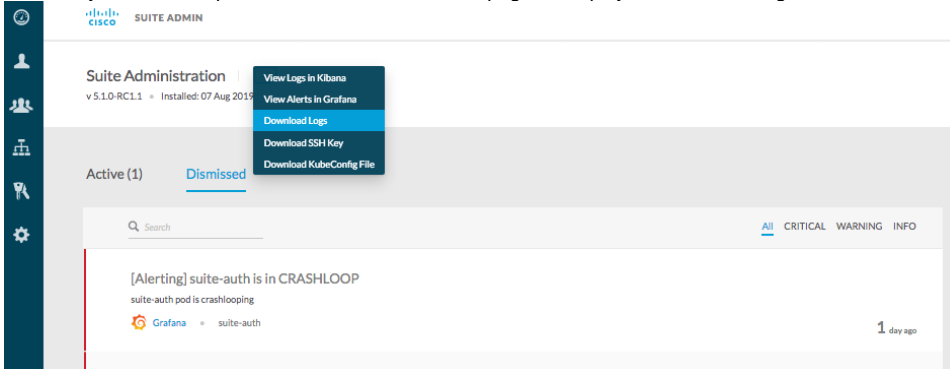
Alert Status	Alert Message	Source	Time
[OK]	suite-auth is in CRASHLOOP	Grafana - suite-auth	1 day ago
[Alerting]	suite-service-mgmt is in CRASHLOOP	Grafana - suite-service-mgmt	1 day ago
[OK]	suite-service-mgmt is in CRASHLOOP	Grafana - suite-service-mgmt	1 day ago

The module name displays at the top of the page and you can perform the following actions on this page:

- Perform one of the actions listed in the **Dropdown** next to the *Module* name as displayed in the following screenshot:



Alternately, click the dropdown from the Module Alerts page as displayed in the following screenshot:



- View the **Alerts** Tab – See the *Understand Dashboard Alerts* section below.
- Access the **License Usage** Tab

There are numerous ways for you to access a module in the CloudCenter Suite. However, your [User Levels](#) determine if you can access the module!


**Kibana** is a web interface that can be used to search and view the logs for **any of the CloudCenter Suite modules**.

CloudCenter Suite log file use the standard log format:

- Where relevant, modules display the user and tenant information.
- You can search by *userid* or *tenantid* when users view logs in Kibana.
- The log files support JSON format.

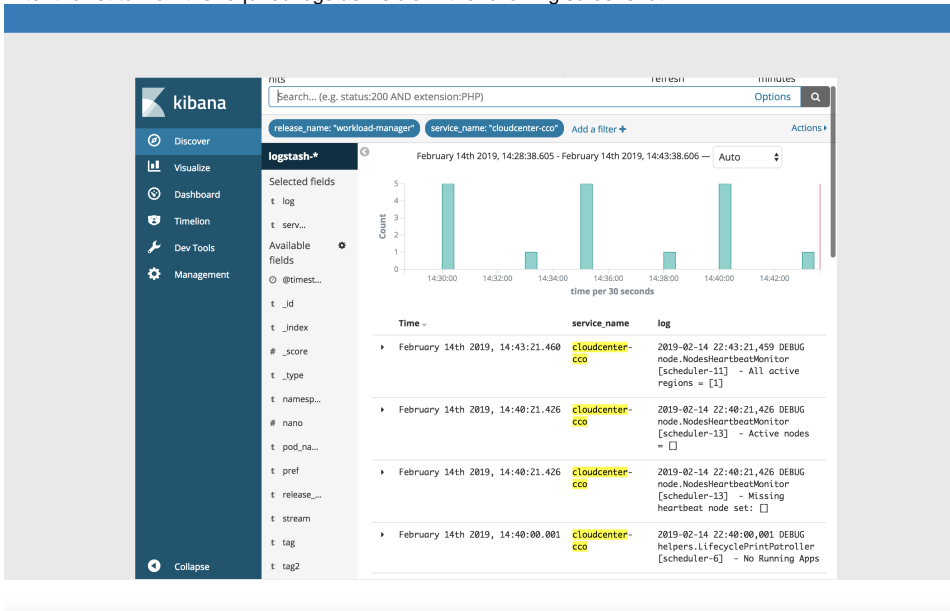
To view the Kibana logs, follow this procedure.

- Click the module dropdown and select **View Logs in Kibana** from the dropdown to display the Kibana dashboard visible in the following screenshot.



- Click **Discover > Open** to list and filter the available logs for this module.

3. Filter the list to view the required logs as visible in the following screenshot.



An alternative to viewing logs in Kibana is to download the log files by clicking a module and selecting **Download Logs** from the dropdown as displayed in the following screenshot.

The screenshot shows the 'Download Logs' dialog box with the following elements:

- Title Bar:** Download Logs (with a close button)
- Section:** Choose the time period to download the logs
- TIME PERIOD:** 1W, 30D (selected), 60D, 90D, YTD, CUSTOM
- Buttons:** VIEW LOGS IN KIBANA, DOWNLOAD

**Grafana** is an open source visualization tool that allows you to create and edit dashboards.

Modules can create their own services to write custom alerts or create alerts in Grafana for services that they wish to monitor.

When alerts are generated, they are displayed in the Suite Admin's *module* details page > **Alerts** tab. When you acknowledge active alerts, they are move to the Dismissed tab and stored there for 60 days before they are deleted.

The **Alerts** tab lists two categories of alerts which are driven from Grafana.

- Active Alerts: Each active alert lists the following details:
  - A color-coded alert category
  - The alert title – click the alert link to open the chart in Grafana using authorized credentials
  - An alert count – only displayed when there is more than one alert
  - A brief description of the alert
  - The alert source
  - The impacted component
  - A snapshot of the chart in Grafana – not available for application alerts
  - The timestamp when this alert was issued – hovering over this timestamp displays the exact time
  - The option to multi-select multiple alertes – the **Dismiss** button becomes visible when you multi-select alerts
- Dismissed Alerts

Alert types are described in the following table.

Alert Type	Description
Infrastructure	These alerts pertain to network, disk, CPU, and memory usage derived from module configured Grafana dashboards.
Application	These alerts are derived from application endpoints that provide the current health of the system.

You can filter alerts based on the type. Alert types are described in the following table.

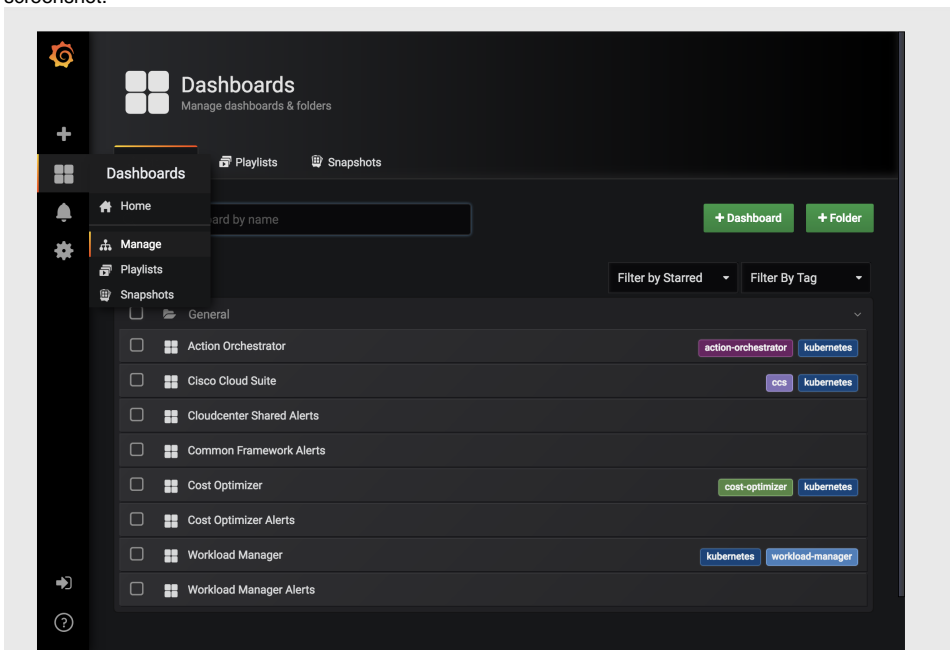
Alert Type	Color	Description
Critical	Red	Red bar on the side. VM launch failure rate is increasing on the configured cloud.
Warning	Orange	The connection to the AMQP server is not stable and has been dropped t times in the last 45 minutes.
Info	Blue	Updates based on endpoint reports.

When you access the Grafana dashboard, you will see the following sections:

- **System metrics:** CPU usage, memory usage, and crash loops. You can also configure additional alerts in this section, refer to <http://docs.grafana.org/alerting/rules/>.
- **Visualization metrics:** Cluster health, deployments, nodes, pods (number of pods and pods status), containers, and jobs. You cannot configure additional alerts in this section.

To view the Grafana alerts, follow this procedure.

1. Click the module dropdown and select **View Alerts in Grafana** from the dropdown to display the Grafana dashboard visible in the following screenshot.



2. Click **Dashboard > Manage** to list and filter the available alerts for this module.



3. Filter the list to view the required alerts as visible in the following screenshot.

