



Cisco Cloud Network Controller for AWS Installation Guide, Release 26.0(x)

First Published: 2021-09-21

Last Modified: 2023-02-28

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



Trademarks

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS REFERENCED IN THIS DOCUMENTATION ARE SUBJECT TO CHANGE WITHOUT NOTICE. EXCEPT AS MAY OTHERWISE BE AGREED BY CISCO IN WRITING, ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS DOCUMENTATION ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED.

The Cisco End User License Agreement and any supplemental license terms govern your use of any Cisco software, including this product documentation, and are located at:

<http://www.cisco.com/go/softwareterms>. Cisco product warranty information is available at <http://www.cisco.com/go/warranty>. US Federal Communications Commission Notices are found here <http://www.cisco.com/c/en/us/products/us-fcc-notice.html>.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any products and features described herein as in development or available at a future date remain in varying stages of development and will be offered on a when-and-if-available basis. Any such product or feature roadmaps are subject to change at the sole discretion of Cisco and Cisco will have no liability for delay in the delivery or failure to deliver any products or feature roadmap items that may be set forth in this document.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

The documentation set for this product strives to use bias-free language. For the purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on RFP documentation, or language that is used by a referenced third-party product.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: [www.cisco.com go trademarks](http://www.cisco.com/go/trademarks). Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)



CONTENTS

PREFACE	Trademarks iii
----------------	-----------------------

CHAPTER 1	New and Changed Information 1
	New and Changed Information 1

CHAPTER 2	Overview 3
	Extending the Cisco ACI Fabric to the Public Cloud 3
	Components of Extending Cisco ACI Fabric to the Public Cloud 4
	Supported Cloud Computing Platforms and Connectivity Options 7
	Support for AWS Organizations and Organization User Tenant 8
	Policy Terminology 10
	Cisco Cloud Network Controller Licensing 10
	Cisco Cloud Network Controller Related Documentation 12

CHAPTER 3	Preparing for Installing Cisco Cloud Network Controller 13
	Requirements for Extending the Cisco ACI Fabric to the Public Cloud 13
	Requirements for the On-Premises Data Center 13
	Requirements for the AWS Public Cloud 14
	Cisco Cloud Network Controller Communication Ports 17
	Cisco Cloud Network Controller Installation Workflow 17

CHAPTER 4	Configuring the Cloud Formation Template Information for the Cisco Cloud Network Controller 19
	Deploying the Cisco Cloud Network Controller in AWS 19
	Resolving Subnet Conflict Issue With Infra Subnet 22
	Setting Up the AWS Account for the User Tenant 24
	Setting Up the AWS Account for a Trusted User Tenant Using the CFT 25

Setting Up the AWS Account for an Untrusted User Tenant Using the AWS Access Key ID and Secret Access Key 27

Setting Up the AWS Account for an Organization User Tenant 28

CHAPTER 5	Configuring Cisco Cloud Network Controller Using the Setup Wizard	29
	Configuring and Deploying Inter-Site Connectivity	29
	Gathering On-Premises Configuration Information	29
	Understanding Limitations for Number of Sites, Regions and CCRs	30
	Locating the Cisco Cloud Network Controller IP Address	31
	Configuring Cisco Cloud Network Controller Using the Setup Wizard	32
	Verifying the Cisco Cloud Network Controller Setup Wizard Configurations	40

CHAPTER 6	Managing Cisco Cloud Network Controller Through Multi-Site	41
	About Cisco Cloud Network Controller and Multi-Site	41
	Adding the Cisco Cloud Network Controller Site to Multi-Site	42
	Configuring the Intersite Infrastructure	42
	Enabling Connectivity Between the Cisco Cloud Network Controller and the ISN Devices	43
	Configuring a Shared Tenant	47
	Creating a Schema	49
	Configuring an Application Profile and the EPGs	50
	Creating and Associating a Bridge Domain with a VRF	50
	Creating a Filter for a Contract	51
	Creating a Contract	51
	Adding Sites to the Schema	52
	Configuring Instances in AWS	52
	Adding an Endpoint Selector	54
	Verifying the Multi-Site Configurations	58

CHAPTER 7	Understanding the Cisco Cloud Network Controller GUI	61
	Navigating the Cisco Cloud Network Controller GUI	61
	Configuring Cisco Cloud Network Controller Components	61

CHAPTER 8	Performing a System Upgrade, Downgrade or Recovery	63
	Important Notes	63

Upgrading the Software	68
Policy-Based Upgrade	68
Backing Up Your Existing Configuration	68
Downloading an Image	69
Upgrading the Software Using the Policy-Based Upgrade Process	70
Migration-Based Upgrade	71
Upgrading Your Cisco Cloud Network Controller Software Using Migration Procedures	72
Downgrading the Software	76
Downgrading the Software: Release 25.0(1) to 5.2(1)	76
Downgrading the Software: Release 25.0(2) to 25.0(1) or 5.2(1)	81
Downgrading the Software: Release 25.0(3) to 25.0(2), 25.0(1), or 5.2(1)	86
Performing a System Recovery	91
Triggering an Upgrade of the CCRs	91
Triggering an Upgrade of the CCRs	91
Triggering an Upgrade of the CCRs Using the Cisco Cloud Network Controller GUI	92
Triggering an Upgrade of the CCRs Using the REST API	93
<hr/>	
APPENDIX A	AWS Resources and Naming Conventions 95
	AWS Resources and Naming Conventions 95
<hr/>	
APPENDIX B	AWS IAM Roles and Permissions 97
	AWS IAM Roles and Permissions 97
<hr/>	
APPENDIX C	Tenant-Region Management 105
	Tenant-Region Management 105
<hr/>	
APPENDIX D	Locating CCR and Tenant Information 107
	Locating CCR and Tenant Information 107



CHAPTER 1

New and Changed Information

- [New and Changed Information](#), on page 1

New and Changed Information

The following table provides an overview of the significant changes to the organization and features in this guide up to this current release. The table does not provide an exhaustive list of all changes made to the guide or of the new features up to this release.

Table 1: New Features and Changed Behavior in Cisco Cloud Network Controller for Release 26.0(2)

Feature or Change	Description	Where Documented
Multi Account support per tenant	Beginning with 26.0(2), Cisco Cloud Network Controller supports association of multiple cloud accounts under a given tenant.	
New consumption with Pay-as-you-go (PAYG)	Beginning with 26.0(2), Cisco Cloud Network Controller is offered on the AWS Marketplace as a Pay-as-you-go (PAYG) image.	
Multiple VPCs/VNETs in a single VRF	This release supports deploying multiple VPCs/VNETs belonging to one VRF in the same region	

Table 2: New Features and Changed Behavior in Cisco Cloud Network Controller for Release 26.0(1)

Feature or Change	Description	Where Documented
UI new look and feel	The screenshots and diagrams in this document have been updated to reflect the recent UI changes. While the look and feel of the new UI is different, the arrangement of the screens and configuration options remains the same.	



CHAPTER 2

Overview

- [Extending the Cisco ACI Fabric to the Public Cloud, on page 3](#)
- [Components of Extending Cisco ACI Fabric to the Public Cloud, on page 4](#)
- [Supported Cloud Computing Platforms and Connectivity Options, on page 7](#)
- [Support for AWS Organizations and Organization User Tenant, on page 8](#)
- [Policy Terminology, on page 10](#)
- [Cisco Cloud Network Controller Licensing, on page 10](#)
- [Cisco Cloud Network Controller Related Documentation, on page 12](#)

Extending the Cisco ACI Fabric to the Public Cloud

Cisco Application Centric Infrastructure (ACI) customers who own a private cloud sometimes may run part of their workload on a public cloud. However, migrating workload to the public cloud requires working with a different interface and learning different ways to set up connectivity and define security policies. Meeting these challenges can result in increased operational cost and loss of consistency.

Cisco ACI can use Cisco Cloud Network Controller to extend a multi-site fabric to Amazon Web Services (AWS), Microsoft Azure, and Google Cloud public clouds.

What Cisco Cloud Network Controller Is

Cisco Cloud Network Controller is a software deployment of Cisco APIC that can be deployed on a cloud-based virtual machine (VM). Cisco Cloud Network Controller provides the following features:

- Provides an interface that is similar to the existing Cisco APIC to interact with the Amazon AWS, Microsoft Azure, or Google Cloud public clouds.
- Automates the deployment and configuration of cloud deployment.
- Configures the cloud router control plane.
- Configures the data path between the on-premises Cisco ACI fabric and the cloud site.
- Translates Cisco ACI policies to cloud native policies.
- Discovers endpoints.

How Users Can Benefit from Cisco ACI Extension to the Public Cloud

Cisco Cloud Network Controller is a key part of Cisco ACI extension to the public cloud. Cisco Cloud Network Controller provides consistent policy, security, and analytics for workloads deployed either on or across on-premises data centers and the public cloud.

Cisco ACI extension to the public cloud also provides an automated connection between on-premises data centers and the public cloud with easy provisioning and monitoring. It also provides a single point for managing, monitoring, and troubleshooting policies across on-premises data centers and the public cloud or between cloud sites.

AWS GovCloud Support

Cisco Cloud Network Controller supports AWS GovCloud in the us-gov-west and us-gov-east regions. Cisco CCRs can also be deployed in the us-gov-east region.

Note that these areas have a unique configuration when you deploy a Cisco Cloud Network Controller on AWS GovCloud:

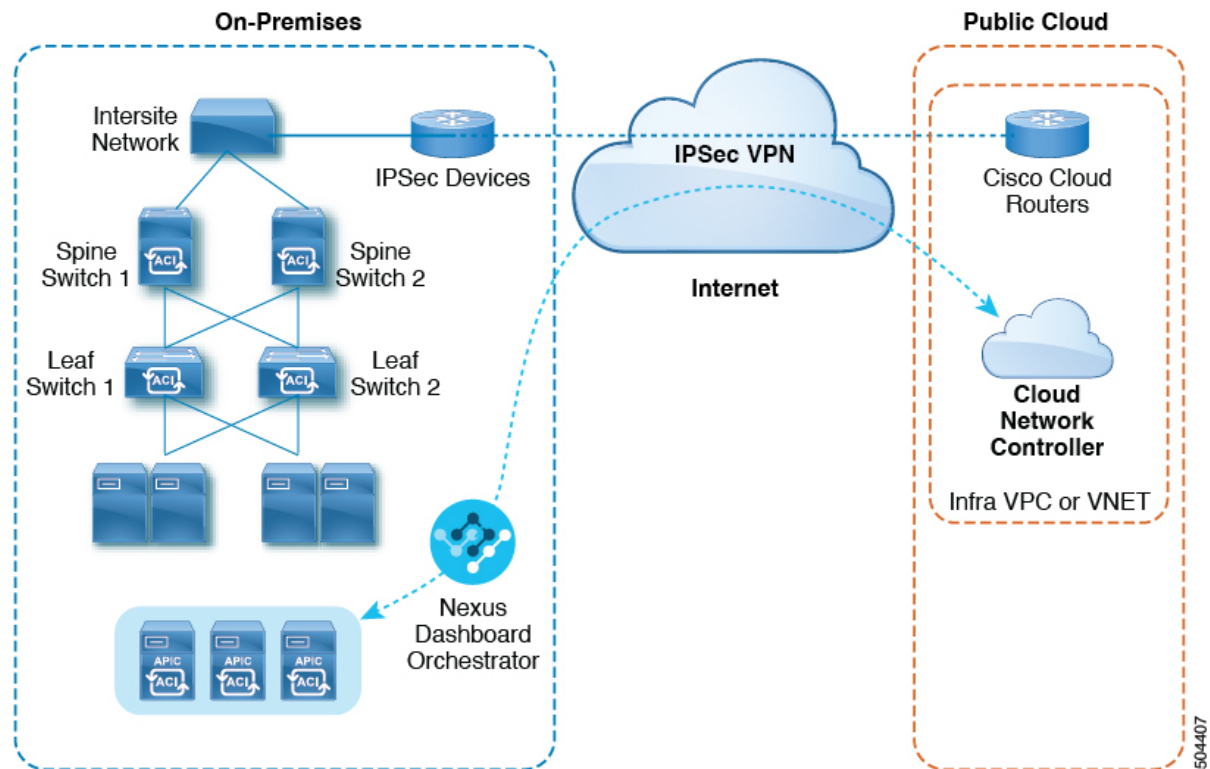
- You will subscribe to the CCR on the commercial account.
- You will subscribe to the Cisco Cloud Network Controller on the commercial account.
- You will launch the Cloud Formation template from the commercial account, which redirects the request to AWS GovCloud for the login.

Components of Extending Cisco ACI Fabric to the Public Cloud

Several components—each with its specific role—are required to extend the Multi-Site fabric to the public cloud.

The following illustration shows the architecture of Cisco Cloud Network Controller.

Figure 1: Cisco Cloud Network Controller Architecture



On-Premises Data Center Components

Cisco ACI Fabric and Cisco APIC

The Cisco ACI allows application requirements to define the network. This architecture simplifies, optimizes, and accelerates the entire application deployment life cycle. Cisco Application Policy Infrastructure Controller (APIC) is a key component of Cisco ACI. It enables applications to directly connect with a secure, shared, high-performance resource pool that includes network, compute, and storage capabilities.

Multi-Site and Multi-Site Orchestrator/Cisco Nexus Dashboard Orchestrator

Multi-Site is an architecture that allows the application to define the networking requirements in a programmatic way. This architecture simplifies, optimizes, and accelerates application deployment. You must have Multi-Site installed to use Cisco Cloud Network Controller to extend the fabric into the public cloud.

For more information, see the [Nexus Dashboard documentation](#) on Cisco.com and the section [Managing Cisco Cloud Network Controller Through Multi-Site, on page 41](#) in this guide.

Cisco Nexus Dashboard Orchestrator (NDO) manages multiple instances of Cisco Application Policy Infrastructure Controller (APICs) in multiple fabrics (sites).

When extending the Cisco ACI fabric to the public cloud, Cisco Nexus Dashboard Orchestrator creates connectivity between the on-premises data center and the public cloud. Use Multi-Site to create tenants across the on-premises data center and the public cloud.



Note You must configure the on-premises Cisco ACI fabric: Create a Fabric Ext Connection Policy and define the overlay TEP and other information required for Multi-Site. You also must add the on-premises Cisco ACI fabric to the Multi-Site architecture. See the [Nexus Dashboard Configuration Guide](#) on Cisco.com.

For more information, see the [Nexus Dashboard documentation](#) on Cisco.com and the section [Managing Cisco Cloud Network Controller Through Multi-Site, on page 41](#) in this guide.

IP Security (IPsec) Router

A router capable of Internet Protocol Security (IPsec) is required to establish IPsec connections between the on-premises site and the public cloud site.

AWS Public Cloud Components

Cisco Cloud Network Controller

Cisco Cloud Network Controller performs the following actions:

- Defines a site on the public cloud, provisions the cloud infra virtual private clouds (VPCs) or virtual networks (VNETs) and manages the Cisco Cloud Router (CCR) across all regions.
- Renders the Cisco ACI policy model in the public cloud, and manages cloud health.

For more information, see *Cisco Cloud Network Controller Release Notes*. Also see the sections [Deploying the Cisco Cloud Network Controller in AWS, on page 19](#) and [Configuring Cisco Cloud Network Controller Using the Setup Wizard, on page 32](#) in this guide.

Cisco Cloud Router

The Cisco Cloud Router (CCR) is a virtual router that delivers comprehensive WAN gateway and network services into virtual and cloud environments. The CCR enables enterprises to extend their WANs into provider-hosted clouds. Two CCRs are required for Cisco Cloud Network Controller solution.

For release 25.0(3) and later, Cisco Cloud Network Controller uses the **Cisco Catalyst 8000V** as the cloud services router. For more information on this CCR, see the [CCR 8000v documentation](#).

AWS public cloud

AWS is a cloud-based platform that provides on-demand services such as compute, storage, network, and databases. Subscribers to AWS have access through the Internet to virtual computers where they can run their workloads.

For more information, see the documentation on the AWS website.

Connections Between the On-Premises Data Center and the Public Cloud

IPsec VPN

You need Internet connectivity with a VPN from the IPsec router, including a publicly routable IP address and with sufficient bandwidth for AWS or Microsoft Azure connectivity.


Management Connection

You need a management connection between the Nexus Dashboard Orchestrator in the on-premises data center and Cisco Cloud Network Controller in the public cloud.

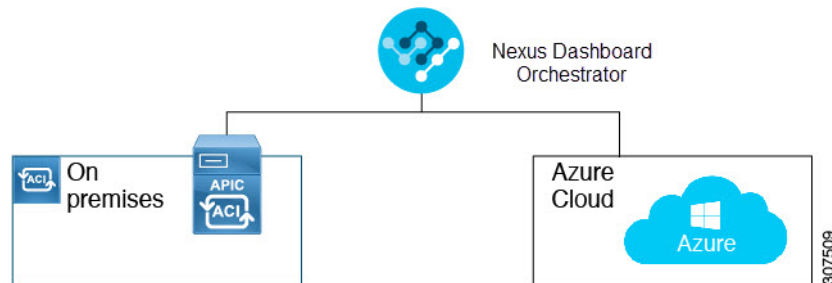
Supported Cloud Computing Platforms and Connectivity Options

Cisco Cloud Network Controller is supported on the following cloud computing platforms:

- As part of the initial release of the Cisco Cloud Network Controller in release 4.1(1), support is provided for on-premises-to-cloud connectivity, or Hybrid-Cloud, where you could use the Cisco Cisco Nexus Dashboard Orchestrator to extend an on-premises Cisco ACI site to Amazon AWS public clouds.




- Beginning in release 4.2(1), support is available for using the Cisco Cisco Nexus Dashboard Orchestrator to extend an on-premises Cisco ACI site to Microsoft Azure public clouds.



- Support is available for using the Cisco Cisco Nexus Dashboard Orchestrator to extend an on-premises Cisco ACI site to Google Cloud public clouds.


You can also use the Cisco Nexus Dashboard Orchestrator to establish connectivity between the following components:

- On-premises-to-cloud connectivity:
 - Connectivity for these public cloud sites:
 - On-premises Cisco ACI and Amazon AWS public cloud sites
 - On-premises Cisco ACI and Microsoft Azure public cloud sites
 - On-premises Cisco ACI and Google Cloud public cloud sites
 - On-premises-to-single cloud site connectivity (Hybrid-Cloud)
 - On-premises-to-multiple cloud sites connectivity (Hybrid Multi-Cloud)



- Cloud site-to-cloud site connectivity (Multi-Cloud):

- Between Amazon AWS public cloud sites (Amazon AWS public cloud site-to-Amazon AWS public cloud site)
- Between Microsoft Azure public cloud sites (Microsoft Azure public cloud site-to-Microsoft Azure public cloud site)
- Between Google Cloud public cloud sites (Google Cloud public cloud site-to-Google Cloud public cloud site)
- Between Amazon AWS, Microsoft Azure, and Google Cloud public cloud sites



In addition, support is also available for the single-cloud configuration (Cloud First).

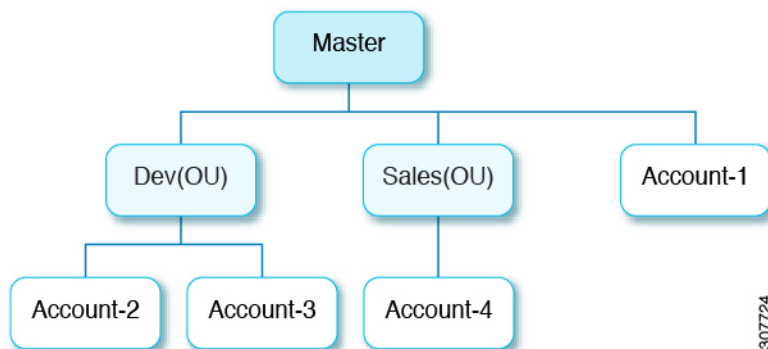
Support for AWS Organizations and Organization User Tenant

With multiple accounts in an organization, it is not easy to control access policies and permissions for various accounts individually, whereas it is easier to do so at the organizational level or at a sub-organizational level within the organization.

Using AWS Organizations, an enterprise might have multiple AWS accounts managed in an organization, as explained here:

<https://aws.amazon.com/organizations/>

This control of the access policies for accounts (or sub-accounts) in the organization is done by the master account of the organization, which is at the root of accounts hierarchy in the organization. The figure below shows an example setup of accounts in an organization.



There are two ways that AWS accounts become part of an AWS Organization:

- **Created:** Within the existing organization in the master account, you can create an AWS account that is automatically part of your AWS organization using the AWS GUI or the AWS API.
- **Invited:** For accounts that are created outside the organization but need to be joined to the organization, an invitation needs to be sent by the master account to the account owner. After accepting the invitation, the invited account becomes a sub-account within the organization.

If you are using AWS Organizations to consolidate and manage your AWS accounts, you will use AWS Organizations to set up your organization and add the created or invited accounts, as you would normally. See [Creating an Organization](#) for more information.

Once you have added the created or invited accounts to your organization through AWS, you will then make the necessary Cisco Cloud Network Controller configurations so that the Cisco Cloud Network Controller recognizes the AWS Organization configurations that you've made through AWS. The Cisco Cloud Network Controller uses the `OrganizationAccountAccessRole` IAM role to manage policies for AWS Organization tenants.

- If you **created** an AWS account within the existing organization in the master account, the `OrganizationAccountAccessRole` IAM role is automatically assigned to that created AWS account. You do not have to manually configure the `OrganizationAccountAccessRole` IAM role in AWS in this case.
- If the master account **invited** an existing AWS account to join the organization, then you must manually configure the `OrganizationAccountAccessRole` IAM role in AWS. Configure the `OrganizationAccountAccessRole` IAM role in AWS for the organization tenant and verify that it has Cisco Cloud Network Controller -related permissions available.

The `OrganizationAccountAccessRole` IAM role, together with the SCP (Service Control Policy) used for the organization or the account, must have the minimum permissions that are required by the Cisco Cloud Network Controller to manage policies for the tenants. The access policy requirement is the same as the requirement for the trusted or untrusted tenants.

For more information, see the "Configure a Tenant AWS Provider" section in the [Cisco Cloud Network Controller for AWS User Guide](#)

You can then assign the Organization tag to tenants through the Cisco Cloud Network Controller GUI using procedures described in [Configuring a Shared Tenant, on page 47](#).

Policy Terminology

A key feature of Cisco Cloud Network Controller is translation of Cisco Application Centric Infrastructure (ACI) policy to the native constructs of the public cloud.

The following table lists Cisco ACI policy terms and the equivalent terms in Amazon Web Services (AWS).

Cisco ACI	AWS
Tenant	User account
AAA user, security domain	Identity and Access Management (IAM)
Virtual Routing and Forwarding (VRF)	VPC
BD subnet	Virtual Private Cloud (VPC) subnet (CIDR)
ACI infra (or ACI infra tenant)	VPC (named Infra VPC by Cisco Cloud Network Controller)
Contract, filter	Security Group Rule
Taboo	Network access list
EPG	Security group
EP-to-EPG mapping	Tag, label
Endpoint	Network adapter on EC2 instances

Cisco Cloud Network Controller Licensing

This section lists the licensing requirements to use Cisco Cloud Network Controller.

Cisco Catalyst 8000V

The Cisco Catalyst 8000V on Cisco Cloud Network Controller supports the following licensing models:

1. **Bring Your Own License (BYOL)** Licensing Model
2. **Pay As You Go (PAYG)** Licensing Model

BYOL Licensing Model

The BYOL licensing model on Cisco Catalyst 8000V which requires you to purchase your Catalyst 8000V Cisco DNA license from Cisco and deploy it in the cloud.

- For instructions on subscribing to one of the tier-based Cisco Catalyst 8000V licenses, see [Cisco Catalyst 8000V Edge Software](#).
- For more information on different throughputs based on the tiers, see the "Throughput" section in "About the Cisco Catalyst 8000V" in the [Cisco Cloud Network Controller for AWS User Guide](#).

PAYG Licensing Model

Beginning with the 25.0(4) release, Cisco Cloud Network Controller supports Pay-As-You-Go (PAYG) Licensing Model on Cisco Catalyst 8000V which allows users to deploy a Catalyst 8000V instance in the cloud based on the VM size and purchase the usage on an hourly basis.

As you completely depend on the VM size to get the throughput, the PAYG licensing model can be enabled only by first un-deploying the current Cisco Catalyst 8000V and then re-deploying it using the First Time Set Up with the new VM size. For more information, see [Configuring Cisco Cloud Network Controller Using the Setup Wizard, on page 32](#).



Note The procedure for switching between licenses can also be used if you would like to switch between the two licensing types available.



Note There are two PAYG options for consuming licenses in the AWS marketplace: **Catalyst 8000V Cisco DNA Essentials** and **Catalyst 8000V Cisco DNA Advantage**. Cisco Cloud Network Controller will make use of **Catalyst 8000V Cisco DNA Advantage**. For features supported by the “Cisco DNA Advantage” subscription, see [Cisco DNA Software SD-WAN and Routing Matrices](#)

Cisco Cloud Network Controller and On-Premises ACI Licensing Summary

- Licensing requirements for all leaf switches on the on-premises Cisco ACI sites:
 - If the Cisco ACI on-premises site is a single site, then use the Essentials license tier (or higher) for the on-premises leaf switches
 - If the Cisco ACI on-premises site is a multi-site, then use the Advantage license tier (or higher) for the on-premises leaf switches
- Licensing requirements for all VM instances managed by Cisco Cloud Network Controller instances:
 - If the Cisco ACI on the cloud has only one Cisco Cloud Network Controller, then use the Essentials Cloud license tier (or higher) for Cisco Cloud Network Controller
 - If the Cisco ACI on the cloud has more than one Cisco Cloud Network Controller, then use the Advantage Cloud license tier (or higher) for Cisco Cloud Network Controller

Amazon Web Services (AWS)

Prior to 26.0(2), Cisco Cloud Network Controller was offered in the AWS marketplace only as a BYOL image. Beginning with release 26.0(2), you can choose between **Bring Your Own License (BYOL)** and **Pay As You Go (PAYG)** licensing mode to launch Cisco Cloud Network Controller in AWS. CCNC is now offered as a PAYG image so that customers can directly consume the product from the AWS marketplace without having to get a license from Cisco. On AWS Marketplace Cisco Cloud Network Controller will be listed as PAYG and BYOL. Based on the product you select in the marketplace, the Cisco Cloud Network Controller license is automatically set to BYOL/PAYG.

- For **BYOL Licensing Model** - User will enter the license token. Cisco provides a license using Cisco Smart License manager and users are billed based on the number of endpoints managed by CCNC.

- For **PAYG Licensing Model** - There will be no prompt to apply the license token. All the data will be automatically reported to the AWS for billing based on the dimensions.

The dimensions for the PAYG Licensing Model are used to send statistics of the CCNC usage to the AWS marketplace to bill accordingly. The dimensions are shared every hour and the customer is billed at the end of every month. The dimensions are :

- **Essential EPs** - This indicates that Cisco Cloud Network Controller is on Essential tier and provides the count of instances when it is managed by CCNC without NDO.
- **Advantage EPs** - This indicates that the Cisco Cloud Network Controller is on Advantage tier and provides the count of instances when it is managed by CCNC. CCNC is put on the Advantage tier when it is managed by NDO.



Note Only new deployments support PAYG licensing model image. For deployments prior to 26.0(2) to be used with the PAYG image, you will have to configure a new deployment in PAYG and then import the older data from the BYOL environment. Upgrade from BYOL to PAYG is not supported.

You must subscribe through the AWS Marketplace:

- For **BYOL Cisco Catalyst 8000V Licensing Model**, subscribe to [Cisco Catalyst 8000V Edge Software-BYOL](#).
- For **PAYG Cisco Catalyst 8000V Licensing Model**, subscribe to [Cisco Catalyst 8000V Edge Software – PAYG](#)

Cisco Cloud Network Controller Related Documentation

You can find information about Cisco Cloud Network Controller, Nexus Dashboard, and Amazon Web Services (AWS) from different resources.

Cisco Documentation

You can find documentation for Cisco products on Cisco.com:

- [Cisco Cloud Network Controller documentation](#)
Includes videos, release notes, fundamentals, installation, configuration, and user guides.
- [Nexus Dashboard documentation](#)
Includes videos, release notes, installation, configuration, and user guides.
- [Cisco Cloud Router documentation](#)
Includes release notes, command reference, data sheets, installation, upgrade, and configuration guides.

AWS Documentation

You can find documentation, including user guides, FAQs, case studies, and white papers, on the AWS website.



CHAPTER 3

Preparing for Installing Cisco Cloud Network Controller

- [Requirements for Extending the Cisco ACI Fabric to the Public Cloud, on page 13](#)
- [Cisco Cloud Network Controller Communication Ports, on page 17](#)
- [Cisco Cloud Network Controller Installation Workflow, on page 17](#)

Requirements for Extending the Cisco ACI Fabric to the Public Cloud

Before you can extend the Cisco Application Centric Infrastructure (ACI) to the public cloud, you must meet requirements for the Cisco ACI on-premises datacenter and the Amazon Web Services (AWS) deployment.

Requirements for the On-Premises Data Center

This section lists the on-premises data center requirements for extending the Cisco Application Centric Infrastructure (ACI) fabric to the public cloud.

- Ensure that the Cisco ACI fabric is installed with the following components:
 - At least two Cisco Nexus EX or FX spine switches, or Nexus 9332C and 9364C spine switches, running Cisco Nexus 9000 Series ACI Mode switch software release 14.1 or later.
 - At least two Cisco Nexus pre-EX, EX, or FX leaf switches running the Cisco Nexus 9000 Series ACI Mode switch software release 14.1 or later.



Note Even though Cisco Nexus pre-EX leaf switches are supported, we recommend using later-generation leaf switches, such as EX or FX leaf switches, due to the End-of-Life announcement for these older pre-EX leaf switches as described in [End-of-Sale and End-of-Life Announcement for the Cisco Nexus 9372PX and 9372TX Switches](#).

- At least one on-premises Cisco Application Policy Infrastructure Controller (APIC) running release 4.1 or later and Cisco Nexus Dashboard Orchestrator (NDO) Release 2.2(x) or later.

- Cisco Nexus Dashboard Orchestrator 2.2(x) deployed with basic configuration.
- A router capable of terminating Internet Protocol Security (IPsec).
- You need to make sure that you have enough bandwidth for tenant traffic between on-premises and cloud sites.
- Verify that all leaf switches on the on-premises sites have the appropriate Cisco ACI license:
 - If the Cisco ACI on-premises site is a single site, then use the Essentials license tier (or higher) for the on-premises leaf switches
 - If the Cisco ACI on-premises site is a multi-site, then use the Advantage license tier (or higher) for the on-premises leaf switches



Note These licensing requirements for the on-premises data center are independent of the number of Cisco Cloud Network Controllers deployed on public clouds. For Cisco Cloud Network Controller licensing requirements, see [Cisco Cloud Network Controller and On-Premises ACI Licensing Summary, on page 11](#).

- Workloads that are connected to the Cisco ACI fabric.
- An intersite network (ISN) that is configured between the Cisco ACI fabric (spine) and the IP Security (IPsec) termination device.

For information about creating an ISN, see the "Multipod" chapter of the [Cisco APIC Layer 3 Networking Configuration Guide](#).

- Certain firewall ports must be permitted if you are deploying firewalls between your on-premises and AWS deployments. These include HTTPS access for the Cisco Cloud Network Controller, IPsec ports for each AWS CCR, and SSH connectivity for AWS CCR remote management.

These firewall ports are described in more detail in [Cisco Cloud Network Controller Communication Ports, on page 17](#) in this guide.

Requirements for the AWS Public Cloud

This section lists the Amazon Web Services (AWS) requirements for extending the Cisco Application Centric Infrastructure (ACI) fabric to the public cloud.

AWS Accounts

You need one AWS account for the infra tenant, and you need at least one AWS account for each user tenant.



Note You can run only one Cloud Network Controller in the infra account. Running multiple Cloud Network Controllers in the same infra account is not supported.

For example, if you want to create two user tenants, you need at least three AWS accounts. You must have one account for each user tenant and one account for the infra tenant. The user tenant can be trusted or untrusted. For details, see the section [Setting Up the AWS Account for the User Tenant, on page 24](#) in this guide.



Note Beginning with 26.0(2), support is now available for having multiple cloud accounts under a single tenant. For more information, see [Tenants](#).

AWS Resources

You need the following resources as part of the AWS deployment:

- Access to the Cisco APIC 5.0 Amazon Machine Image (AMI).



Note To have access to the AMI, you must subscribe to the Cisco Cloud Network Controller in the Amazon Marketplace.

- Two instances of Elastic Cloud Computer (EC2), which function as virtual machines (VM) for applications running in the cloud.
- Virtual Private Clouds (VPCs), subnets, a virtual private gateway (VGW), an Internet gateway (IGW), security groups, and resources that are based on tasks you plan to perform.

CCR

There are two types of licensing models available:

- BYOL (Bring Your Own License)
- PAYG (Pay as You Go)

BYOL

Subscribe to the CCR Bring Your Own License (BYOL) through the AWS Marketplace. See [Cisco Cloud Network Controller Licensing, on page 10](#) for more information.

Deploy the CCRs in the appropriate size, depending on the bandwidth requirement defined during the Cisco Cloud Network Controller setup.

The value for the throughput of the routers determines the size of the CCR instance that you deploy; a higher value for the throughput results in the deployment of a larger VM. CCR licensing is based on the throughput configuration that you set as part of the Cisco Cloud Network Controller setup process. You need the equivalent or higher license in your Smart account and the AX feature set for compliance.

Make sure that your AWS account has an allowed limit to deploy the instances. You can check your account instance limits in the AWS Management Console: **Services** > **EC2** > **Limits**.

The Cisco Catalyst 8000V supports tier-based (T0/T1/T2/T3) throughput options. The following table lists what AWS EC2 instance is used for different router throughput settings for the Cisco Catalyst 8000V:

CCR Throughput	AWS EC2 Instance
T0 (up to 15M throughput)	c5.xlarge
T1 (up to 100M throughput)	c5.xlarge
T2 (up to 1G throughput)	c5.xlarge

CCR Throughput	AWS EC2 Instance
T3 (up to 10G throughput)	c5.9xlarge

Tier2 (T2) is the default throughput supported by Cisco Cloud Network Controller.

PAYG

Cisco Cloud Network Controller will support a range of AWS EC2 instances for cloud networking needs powered by Cisco's Catalyst 8000V virtual router. The table below shows the cloud instance type supported by Cisco Cloud Network Controller on AWS.

AWS EC2 Instance	CCR Throughput	vCPUs	Memory
c5.xlarge	up to 5 Gigabit throughput	4	8 GiB
c5.2xlarge	up to 10 Gigabit throughput	8	16 GiB
c5.4xlarge	up to 10 Gigabit throughput	16	32 GiB
c5.9xlarge	up to 10 Gigabit throughput	36	72 GiB
c5n.xlarge	up to 25 Gigabit throughput	4	10.5 GiB
c5n.2xlarge	up to 25Gigabit throughput	8	21 GiB
c5n.4xlarge	up to 25 Gigabit throughput	16	42 GiB
c5n.9xlarge	up to 50 Gigabit throughput	36	96 GiB

Changing the value in the **VM Type** field in the First Time Setup changes the other factors of the CCR as listed in the table above. Choosing a higher value for the VM size results in higher throughput.

Elastic IP Addresses

Make sure that you have at least nine elastic IP addresses in the region where the infra VPC is deployed.

You need one elastic IP address for Cisco Cloud Network Controller and four for each CCR. Make sure that your account in the region of deployment is allowed nine or more elastic IP addresses. If it is not, raise an AWS case to increase the number of elastic IP addresses. We recommend ten or more.



Note The addresses must not be disassociated elastic IP address. You need enough resources for nine new elastic IP addresses. If you have unused elastic IP addresses, you can release them.

Cisco Cloud Network Controller

Cisco Cloud Network Controller is deployed using the m5.2xlarge AWS instance.

Make sure that your account has limits that are allowed to deploy this instance. You can check the limits in the AWS Management Console: **Services > EC2 > Limits**.

You can also see how many elastic IP addresses that are used in the AWS Management Console: **Services > EC2 > NETWORK & SECURITY > Elastic IPs**.

Cisco Cloud Network Controller Communication Ports

When configuring your Cisco Cloud Network Controller environment, keep in mind that the following ports are required for network communications:

- For communication between the Cisco Nexus Dashboard Orchestrator and the Cisco Cloud Network Controller: HTTPS (TCP Port 443 inbound/outbound)

For the Cisco Cloud Network Controller, use the same Cisco Cloud Network Controller management IP address that you will use to log into the Cisco Cloud Network Controller at the beginning of [Configuring Cisco Cloud Network Controller Using the Setup Wizard, on page 32](#).

- For communication between the on-premises IPsec device and the CCRs deployed by Cisco Cloud Network Controller in AWS: Standard IPsec ports (UDP port 500 and permit IP protocol numbers 50 and 51 inbound/outbound)

For the two Amazon Web Services CCRs, the public IPsec peering IP uses the elastic IP address of the third network interface, as described in [Locating CCR and Tenant Information, on page 107](#) or as provided if you download the ISN device configuration files using the instructions in [Configuring the Intersite Infrastructure, on page 42](#).

- If you want to connect and manage the CCRs deployed by Cisco Cloud Network Controller in AWS, allow port TCP 22 inbound/outbound to the public IP address of each CCR.
- For license registration (towards `tools.cisco.com`): Port 443 (outbound) is required
- For DNS: UDP Port 53 outbound
- For NTP: UDP Port 123 outbound
- If remote authentication is used (LDAP, Radius, TACACS+, SAML), open the proper ports
- If a certificate authority is used, open the proper ports

Cisco Cloud Network Controller Installation Workflow

This section provides a high-level description of the tasks that are required to install and deploy Cisco Cloud Network Controller. You perform installation tasks through AWS Management Console, the AWS Cloud Formation template, the Cisco Cloud Network Controller Setup Wizard, and Nexus Dashboard Orchestrator.

1. Fulfill all prerequisites, which include tasks in the on-premises data center and the public cloud.

See the section "[Requirements for Extending the Cisco ACI Fabric to the Public Cloud, on page 13](#)."

2. Deploy Cisco Cloud Network Controller through the AWS Cloud Formation template.

This task includes creating a stack, uploading a template (or providing an AWS template URL), configuring template parameters, and submitting the template. You then capture the Cisco Cloud Network Controller IP address.

You also must create an Amazon EC2 SSH keypair and subscribe to Cisco Cloud Network Controller in the AWS Marketplace.

See the section "[Deploying the Cisco Cloud Network Controller in AWS, on page 19.](#)"

3. Configure Cisco Cloud Network Controller using the Setup Wizard.

This task includes logging into Cisco Cloud Network Controller and configuring the fabric managed by the Cisco Cloud Network Controller for connecting to the public cloud. You also add the AWS region selection. You provide the Border Gateway Protocol (BGP) autonomous system number (ASN) and OSPF area ID for intersite network (ISN) peering and add an external subnet. You then add the IPsec peer address.

See the section "[Configuring Cisco Cloud Network Controller Using the Setup Wizard, on page 32.](#)"

4. Configure Cisco Cloud Network Controller using Nexus Dashboard Orchestrator.

This task includes logging into the Nexus Dashboard Orchestrator GUI, adding the on-premises and cloud site, configuring the fabric connectivity infra, and configuring the properties for the on-premises site. You then configure the Cisco ACI spines, BGP peering, and enable the connectivity between the on-premises site and the AWS Cisco Cloud Network Controller sites.

See the section "[Managing Cisco Cloud Network Controller Through Multi-Site, on page 41.](#)"

5. Use Cisco Cloud Network Controller to extend Cisco ACI policy into the AWS public cloud.

See the sections "[Navigating the Cisco Cloud Network Controller GUI, on page 61](#)" and "[Configuring Cisco Cloud Network Controller Components, on page 61.](#)"



CHAPTER 4

Configuring the Cloud Formation Template Information for the Cisco Cloud Network Controller

- [Deploying the Cisco Cloud Network Controller in AWS, on page 19](#)
- [Setting Up the AWS Account for the User Tenant, on page 24](#)

Deploying the Cisco Cloud Network Controller in AWS

Before you begin

- Verify that you have met the requirements that are outlined in [Requirements for Extending the Cisco ACI Fabric to the Public Cloud, on page 13](#) before proceeding with the tasks in this section. For example, verify that you have the correct number of elastic IP addresses and that you have checked the limits allowed to deploy the instances.
- Verify that you have the full Administrator Access on AWS, because specific AWS IAM roles and permissions are required for the installation and operation of the Cisco Cloud Network Controller.

When installing Cisco Cloud Network Controller using the CloudFormation template (CFT), we recommend installation by a user who has the full Administrator Access on AWS (for example, by a user who has the permission policy ARN `arn:aws:iam::aws:policy/AdministratorAccess` attached to it, either directly, by using a role policy, or with a user group). However, if there is no one with AWS Administrator Access available, the person installing Cisco Cloud Network Controller must have a minimum set of permissions. Beginning with 26.0(2), IAM access has been optimized. See [AWS IAM Roles and Permissions, on page 97](#) for more information on these AWS IAM roles and permissions.

- If you are using AWS Organizations to control access policies and permissions for various accounts and you want to use Cisco Cloud Network Controller to manage these accounts, verify that the AWS account where you are deploying the Cisco Cloud Network Controller in these procedures (the Cisco Cloud Network Controller infra tenant) is the master account for that AWS organization. When the Cisco Cloud Network Controller is deployed in the master account for an AWS organization, you can add any AWS accounts that are part of the organization as tenants through the Cisco Cloud Network Controller GUI. See [Support for AWS Organizations and Organization User Tenant, on page 8](#) and [Configuring a Shared Tenant, on page 47](#) for more information.

- If you are deploying Cisco Cloud Network Controller on AWS GovCloud, review the information provided in the section "AWS GovCloud Support" in [Extending the Cisco ACI Fabric to the Public Cloud, on page 3](#) for information specific to those deployments.

Step 1 Log into your Amazon Web Services account for the Cisco Cloud Network Controller infra tenant and go to the AWS Management Console, if you are not there already:

<https://signin.aws.amazon.com/>

<https://console.aws.amazon.com/>

Step 2 In the upper right corner of the AWS Management Console screen, locate the area that shows a region, and choose the region in AWS that you want to have managed by Cisco Cloud Network Controller (where the Cisco Cloud Network Controller AMI image will be brought up).

Step 3 Create an Amazon EC2 SSH key pair:

- Click the **Services** link at the top left area of the screen, then click the **EC2** link.

The **EC2 Dashboard** screen appears.

- In the **EC2 Dashboard** screen, click the **Key Pairs** link.

The **Create Key Pair** screen appears.

- Click **Create Key Pair**.

- Enter a unique name for this key pair (for example, `CloudNetControllerKeyPair`), then click **Create**.

A screen is displayed that shows the public key that is stored in AWS. In addition, a Privacy Enhanced Mail (PEM) file is downloaded locally to your system with the private key.

- Move the private key PEM file to a safe location on your system and note the location.

You will navigate back to the private key PEM file in this location in a step later in these procedures.

Step 4 Go to the Cisco Cloud Network Controller page on the AWS Marketplace:

<http://cs.co/capic-aws>

Step 5 Click **Subscribe**.

Step 6 Review and accept the End User License Agreement (EULA) by clicking the **Accept Terms** button.

Step 7 After a minute, you should see the message `Subscription should be processed`. Click the **Continue to Configuration** button.

The **Configure this software** page appears.

Step 8 Select the following parameters:

- **Fulfillment Option:** Cisco Cloud Network Controller Cloud Formation Template (selected by default)
- **Software Version:** Select the appropriate version of the Cisco Cloud Network Controller software
- **Region:** Region where Cisco Cloud Network Controller will be deployed

Step 9 Click the **Continue to Launch** button.

The **Launch this software** page appears, which shows a summary of your configuration and lets you launch the cloud formation template.

Step 10 Click **Launch** to go directly to the CloudFormation service in the correct region, with the correct Amazon S3 template URL already populated.

Step 11 Click **Next** at the bottom of the screen.

The **Specify Details** page appears within the **Create stack** page.

Step 12 Enter the following information on the **Specify Details** page.

- **Stack name:** Enter the name for this Cisco Cloud Network Controller configuration.
- **Fabric name:** Leave the default value as-is or enter a fabric name. This entry will be the name for this Cisco Cloud Network Controller.
- **Infra VPC Pool:** The VPC (Virtual Private Cloud) CIDR. This field is automatically populated from the CFT with a default value of 10.10.0.0/24. Change the value in this field if the default value overlaps with your infra pool from your on-premises fabric. This entry must be a /24 subnet.

Note We recommend that you do not use any subnet from 172.17.0.0/16 (for example, 172.17.10.0/24) as the infra VPC CIDR, as this might cause a conflict with the Docker bridge IP subnet, as described in [Resolving Subnet Conflict Issue With Infra Subnet, on page 22](#).

- **Availability Zone:** Select an availability zone for the Cisco Cloud Network Controller subnets from the scroll-down menu.

The availability zone options that are presented will be based on the region that you selected in [Step 2, on page 20](#). Select the lowest availability zone from the list. For example, if you see `us-west-1a` and `us-west-1b` as the availability zone options, select `us-west-1a`.

- **Password/Confirm Password:** Enter and confirm an admin password. This entry is the password that you will use to log into the Cisco Cloud Network Controller after you have enabled SSH access.
- **SSH Key Pair:** Choose the name of the SSH key pair that you created in [Step 3, on page 20](#).
You will use this SSH key pair to log into the Cisco Cloud Network Controller.
- **Access Control:** Enter the IP addresses and subnets of the external networks that you will allow to connect to Cisco Cloud Network Controller (for example, 192.0.2.0/24). Only the IP addresses from this subnet are allowed to connect to Cisco Cloud Network Controller. Entering a value of `0.0.0.0/0` means that anyone is allowed to connect to Cisco Cloud Network Controller.
- **Other parameters: Assign Public IP address:** Select whether to assign a public IP address to the Out-of-Band (OOB) management interface for the Cisco Cloud Network Controller or not.

By default, a private IP address is assigned to the management interface of the Cisco Cloud Network Controller and assigning a public IP address is optional. For more information, see the "Private IP Address Support for Cisco Cloud Network Controller and CCR" topic in the *Cisco Cloud Network Controller for AWS User Guide*, Release 25.0(5) or later.

- **true:** Assigns a public IP address to the Out-of-Band (OOB) management interface for the Cisco Cloud Network Controller.
- **false:** Disables the public IP address and assigns a private IP address to the Out-of-Band (OOB) management interface for the Cisco Cloud Network Controller.

Step 13 Click **Next** at the bottom of the screen.

The **Options** page appears within the **Create stack** page.

Step 14 Accept all the default values in the **Options** screen.

There is a **Permissions: IAM Role** area on this page. An IAM role is an IAM entity that defines a set of permissions for making Amazon Web Services service requests. You can use roles to delegate access to users, applications, or services that don't normally have access to your Amazon Web Services resources.

There is no need for IAM role information with regards to the Cisco Cloud Network Controller, but if you want to assign an IAM role for another reason, choose the appropriate role in the **IAM Role** field.

Step 15 Click **Next** at the bottom of the **Options** screen.

The **Review** page appears within the **Create stack** page.

Step 16 Verify that all the information on the **Review** page is correct.

If you see any errors on the **Review** page, click the **Previous** button to go back to the page with the incorrect information.

Step 17 When you have verified that all the information on the **Review** page is correct, check the box next to the **I acknowledge that AWS CloudFormation might create IAM resources with custom names** area.

Step 18 Click the **Create** button at the bottom of the page.

The **CloudFormation** page reappears, and the Cisco Cloud Network Controller template that you created is displayed with the text **CREATE_IN_PROGRESS** displayed in the Status column.

The system now uses the information that you provided in the template to create the Cisco Cloud Network Controller instance. This process takes 5-10 minutes to complete. You can monitor the progress of the creation process by checking the box next to the name of your Cisco Cloud Network Controller template, then clicking on the Events tab. The text **CREATE_IN_PROGRESS** is displayed in the Status column under the Events tab.

Step 19 When the **CREATE_COMPLETE** message is shown, verify that the instance is ready before proceeding.

a) Click the **Services** link at the top of the screen, then click the **EC2** link.

The **EC2 Dashboard** screen appears.

b) In the EC2 Dashboard screen, you should see text displaying the number of running instances in the **Resources** area (for example, **1 Running Instances**). Click this running instances link.

The **Instances** screen appears.

c) Wait until you see that instance is ready before proceeding.

You will see the new instance going through the **Initializing** stage under Status Checks. Wait until you see the **2/2 Checks Passed** message under Status Checks before proceeding.

What to do next

Go to [Setting Up the AWS Account for the User Tenant, on page 24](#) to set up the AWS account for the user tenant.

Resolving Subnet Conflict Issue With Infra Subnet

In some situations, you might encounter an issue with a subnet conflict with your Cisco Cloud Network Controller. This issue might occur when the following conditions are met:

- Your Cisco Cloud Network Controller is running on release 25.0(2)

- The infra VPC subnet for your Cisco Cloud Network Controller is configured within the 172.17.0.0/16 CIDR (for example, if you entered 172.17.10.0/24 in the **Infra VPC Pool** field as part of the procedures in [Deploying the Cisco Cloud Network Controller in AWS, on page 19](#))
- There is something else configured that overlaps with the 172.17.0.0/16 CIDR that you are using for the infra VPC subnet for your Cisco Cloud Network Controller (for example, if the Docker bridge IP subnet is configured with 172.17.0.0/16, which is the default subnet in Cisco Cloud Network Controller).

In this situation, your Cisco Cloud Network Controller might not be able to reach the CCR private IP address due to this subnet conflict and the Cisco Cloud Network Controller will raise an SSH connectivity fault for the affected CCR.

You could determine if there might be a possible conflict by logging in as root into the Cisco Cloud Network Controller and entering the `route -n` command:

```
[root@ACI-Cloud-Fabric-1 ~]# route -n
```

Assume that you see output similar to the following:

```
Kernel IP routing table
Destination      Gateway         Genmask         Flags Metric Ref    Use Iface
0.0.0.0          172.17.0.17    0.0.0.0         UG    16     0      0 oobmgmt
169.254.169.0    0.0.0.0        255.255.255.0   U     0     0      0 bond0
169.254.254.0    0.0.0.0        255.255.255.0   U     0     0      0 lxcbr0
172.17.0.0      0.0.0.0        255.255.0.0     U     0     0      0 docker0
172.17.0.12     0.0.0.0        255.255.255.252 U     0     0      0 bond0
172.17.0.16     0.0.0.0        255.255.255.240 U     0     0      0 oobmgmt
```

In this example output, the highlighted text shows that a Docker bridge is configured with 172.17.0.0/16.

Because this overlaps with the 172.17.0.0/16 CIDR that you used for the infra VPC subnet for your Cisco Cloud Network Controller, you might see an issue where you lose connectivity to the CCR, where you are not able to SSH into the CCR, and you see a Host Unreachable message when you try to ping the CCR (such as in the following example, where 172.17.0.84 is the private IP address of the CCR):

```
[root@ACI-Cloud-Fabric-1 ~]# ping 172.17.0.84
PING 172.17.0.84 (172.17.0.84) 56(84) bytes of data.
From 172.17.0.1 icmp_seq=1 Destination Host Unreachable
From 172.17.0.1 icmp_seq=2 Destination Host Unreachable
From 172.17.0.1 icmp_seq=3 Destination Host Unreachable
From 172.17.0.1 icmp_seq=5 Destination Host Unreachable
From 172.17.0.1 icmp_seq=6 Destination Host Unreachable
^C
--- 172.17.0.84 ping statistics ---
 9 packets transmitted, 0 received, +5 errors, 100% packet loss, time 8225ms
pipe 4
[root@ACI-Cloud-Fabric-1 ~]#
```

To resolve the conflict in this situation, enter a REST API post similar to the following to change the IP address for the other area that is causing the conflict:

```
https://{apic}/api/plgnhandler/mo/.xml
<apPluginPolContr>
  <apContainerPol containerBip="<new-IP-address>" />
</apPluginPolContr>
```

For example, to move the Docker bridge IP address out from under the 172.17.0.0/16 CIDR, which was shown in the example scenario above, you might enter a REST API post such as the following:

```
https://{apic}/api/plgnhandler/mo/.xml
<apPluginPolContr>
  <apContainerPol containerBip="172.19.0.1/16" />
</apPluginPolContr>
```

where 172.19.0.1/16 is the new subnet for the Docker bridge. This moves the Docker bridge IP address under the 172.19.0.0/16 CIDR, where there is no longer a conflict with the infra VPC subnet for your Cisco Cloud Network Controller that is configured within the 172.17.0.0/16 CIDR.

You can use the same commands as before to verify that there is no longer a conflict:

```
[root@ACI-Cloud-Fabric-1 ~]# route -n
Kernel IP routing table
Destination      Gateway         Genmask        Flags Metric Ref    Use Iface
0.0.0.0          172.17.0.17    0.0.0.0        UG    16     0      0 oobmgmt
169.254.169.0    0.0.0.0        255.255.255.0  U     0     0      0 bond0
169.254.254.0    0.0.0.0        255.255.255.0  U     0     0      0 lxcbr0
172.17.0.12      0.0.0.0        255.255.255.252 U     0     0      0 bond0
172.17.0.16      0.0.0.0        255.255.255.240 U     0     0      0 oobmgmt
172.19.0.0      0.0.0.0        255.255.0.0    U     0     0      0 docker0
```

In this example output, the highlighted text shows that a Docker bridge is configured with the IP address 172.19.0.0. Because there is no overlap with the 172.17.0.0/16 CIDR that you are using for the infra VPC subnet for your Cisco Cloud Network Controller, there is no issue with connectivity with the CCR:

```
[root@ACI-Cloud-Fabric-1 ~]# ping 172.17.0.84
PING 172.17.0.84 (172.17.0.84) 56(84) bytes of data.
64 bytes from 172.17.0.84: icmp_seq=1 ttl=255 time=1.15 ms
64 bytes from 172.17.0.84: icmp_seq=2 ttl=255 time=1.01 ms
64 bytes from 172.17.0.84: icmp_seq=3 ttl=255 time=1.03 ms
64 bytes from 172.17.0.84: icmp_seq=4 ttl=255 time=1.03 ms
64 bytes from 172.17.0.84: icmp_seq=5 ttl=255 time=1.09 ms
64 bytes from 172.17.0.84: icmp_seq=6 ttl=255 time=1.06 ms
64 bytes from 172.17.0.84: icmp_seq=7 ttl=255 time=1.03 ms
64 bytes from 172.17.0.84: icmp_seq=8 ttl=255 time=1.05 ms
^C
--- 172.17.0.84 ping statistics ---
8 packets transmitted, 8 received, 0% packet loss, time 7005ms
rtt min/avg/max/mdev = 1.014/1.061/1.153/0.046 ms
[root@ACI-Cloud-Fabric-1 ~]#
```

Setting Up the AWS Account for the User Tenant

You can set up the AWS account for the user tenant using one of the following methods:

- Where the user tenant in Cisco Cloud Network Controller is trusted, using the CFT. See [Setting Up the AWS Account for a Trusted User Tenant Using the CFT, on page 25](#).
- Where the user tenant in Cisco Cloud Network Controller is untrusted, using the AWS access key ID and secret access key. See [Setting Up the AWS Account for an Untrusted User Tenant Using the AWS Access Key ID and Secret Access Key, on page 27](#).
- Where you can manage policies for AWS Organization accounts through the Cisco Cloud Network Controller. See [Setting Up the AWS Account for an Organization User Tenant, on page 28](#).

Setting Up the AWS Account for a Trusted User Tenant Using the CFT

Using the tenant role Cloud Formation template (CFT) in the tenant account establishes a trust relationship between the tenant and the account where the Cisco Cloud Network Controller is deployed.

Use the following procedures to set up the AWS account for the user tenant using the tenant role CFT.

Before you begin

Following are the rules and restrictions for configuring the Cisco Cloud Network Controller user tenant:

- You cannot use the same AWS account for the infra tenant and the user tenant.
- You need one AWS account for each user tenant.

Step 1 Log into your Amazon Web Services account for the user tenant:

<https://signin.aws.amazon.com/>

Note Do not use the infra tenant account for the user tenant.

Step 2 Click the **Services** link at the top of the screen, then click the **CloudFormation** link.

The **CloudFormation** screen appears.

Step 3 Click the **Create Stack** button.

Note Do not choose any options from the drop-down list next to the **Create Stack** button. Click directly on the **Create Stack** button instead.

The **Select Template** page appears within the **Create stack** page.

Step 4 Determine how you will select the template to use for the IAM role for the user tenant configuration.

- If you want to download the tenant role CFT from your AWS account, or if you downloaded it from your cisco.com account (formerly CCO), follow these procedures:
 - a. If you want to download the tenant role CFT from your AWS account, locate the tenant role CFT. The tenant role CFT is located in the S3 bucket in the AWS account for the Cisco Cloud Network Controller infra tenant. The name of the S3 bucket is `capic-common-[CloudNetworkControllerAccountId]-data` and the tenant role CFT object is `tenant-cft.json` in that bucket. The `CloudNetworkControllerAccountId` is the AWS account number for the Cisco Cloud Network Controller infra tenant, which is the account in which Cisco Cloud Network Controller is deployed.
 - b. Download the tenant role CFT to a location on your computer.

For security reasons, public access to this S3 bucket in AWS is not allowed, so you must download this file and use it in the tenant account.
 - c. In AWS, in the **Choose a template** area, click the circle next to **Upload a template to Amazon S3**, then click the **Choose File** button.
 - d. Navigate to the location on your computer where you saved the JSON-formatted tenant role CFT that you received from Cisco (for example, `tenant-cft.json`) and select that template file.

- If you were given a tenant role CFT URL from Cisco, in the **Choose a template** area, click the circle next to **Specify an Amazon S3 template URL**, then enter the tenant role CFT URL that you received from Cisco into the field below the text.

Step 5 Click **Next** at the bottom of the screen.

The **Specify Details** page appears within the **Create stack** page.

Step 6 Enter the following information on the **Specify Details** page.

- **Stack name:** Enter the name for this IAM role for the user tenant configuration (for example, `IAM-Role`).
- **infraAccountId:** If you see this field, enter the AWS account for the infra tenant as described in [Deploying the Cisco Cloud Network Controller in AWS, on page 19](#).

Note that this field is displayed if you downloaded and used the tenant role CFT from your cisco.com account. It is not displayed if you downloaded and used the tenant role CFT from your AWS account because the `infraAccountId` information is pre-populated in the CFT when it is downloaded from the S3 bucket in the infra AWS account.

Step 7 Click **Next** at the bottom of the screen.

The **Options** page appears within the **Create stack** page.

Step 8 Accept all the default values in the **Options** screen, if applicable, then click **Next** at the bottom of the screen.

The **Review** page appears within the **Create stack** page.

Step 9 In the **Review** page, check the box next to the **I acknowledge that AWS CloudFormation might create IAM resources with custom names** area, then click the **Create** button at the bottom of the page.

The **CloudFormation** page reappears, and the Cisco Cloud Network Controller template that you created is displayed with the text **CREATE_IN_PROGRESS** displayed in the Status column.

The system now uses the information that you provided in the template to create the IAM role for the user tenant. This process takes 5-10 minutes to complete. You can monitor the progress of the creation process by checking the box next to the name of the template, then clicking on the Events tab. The text **CREATE_IN_PROGRESS** is displayed in the Status column under the Events tab.

CREATE_COMPLETE is shown when the process is completed.

Step 10 When the **CREATE_COMPLETE** is shown, navigate to the appropriate area to verify that the IAM role for the user tenant was created successfully.

- Click the **Services** link at the top of the screen, then click the **IAM** link.
- Click **Roles**.

An entry with the name **ApicTenantRole** should appear under the Role name.

What to do next

Go to [Configuring Cisco Cloud Network Controller Using the Setup Wizard, on page 29](#) to continue setting up the Cisco Cisco Cloud Network Controller.

Setting Up the AWS Account for an Untrusted User Tenant Using the AWS Access Key ID and Secret Access Key

Use the following procedures if you want to set up the AWS account for an untrusted user using the AWS access key ID and secret access key, where you will manually set up the AWS account for an untrusted user tenant and assign the appropriate permissions through AWS IAM.

Before you begin

Following are the rules and restrictions for configuring the Cisco Cloud Network Controller user tenant:

- You cannot use the same AWS account for the infra tenant and the user tenant.
- You need one AWS account for each user tenant.

-
- Step 1** Log into your Amazon Web Services account for the user tenant:
<https://signin.aws.amazon.com/>
- Note** Do not use the infra tenant account for the user tenant.
- Step 2** Go to the AWS Management Console:
<https://console.aws.amazon.com/>
- Step 3** Click the **Services** link at the top of the screen, then click the **IAM** link.
- Step 4** In the left pane, click **Users**, then click the **Add user** button.
The **Add User** page appears.
- Step 5** In the **User name** field, enter a unique name for this AWS user account, such as `user1`.
- Step 6** In the **Access type** field, check **Programmatic access**.
- Step 7** Click the **Next: Permissions** button at the bottom of the page.
- Step 8** In the **Set permissions** area, select **Attach existing policies directly**.
The screen expands to display **Filter policies** information.
- Step 9** Check the box next to **Administrator Access**, then click the **Next: Tags** button at the bottom of the page.
- Step 10** Leave the information in the **Add tags** page as-is and click the **Next: Review** button at the bottom of the page.
- Step 11** Click the **Create User** button at the bottom of the page.
Ignore the warning that states **This user has no permissions** if that warning appears.
An access key is created for you at this point.
- Step 12** Make a note of the Access Key ID and Secret Access Key information for this AWS account.
- Copy the Access Key ID and the Secret Access Key information for the user tenant to the appropriate rows in [Locating CCR and Tenant Information, on page 107](#).
 - Download the .csv file or copy the information from the **Access key ID** and **Secret access key** fields to a file.
- Step 13** Click the **Close** button at the bottom of the page.

Step 14 Repeat the steps in this topic for additional user accounts, if necessary.

What to do next

Go to [Configuring Cisco Cloud Network Controller Using the Setup Wizard, on page 29](#) to continue setting up the Cisco Cisco Cloud Network Controller.

Setting Up the AWS Account for an Organization User Tenant

As described in [Support for AWS Organizations and Organization User Tenant, on page 8](#), you can manage policies for AWS Organization accounts through the Cisco Cloud Network Controller.

To set up the AWS account for an organization tenant, you must have the following configurations in order to use this feature:

- The Cisco Cloud Network Controller must be deployed in the master account. Earlier in this document, when you deployed the Cisco Cloud Network Controller in AWS using the instructions provided in [Deploying the Cisco Cloud Network Controller in AWS, on page 19](#), verify that you deployed the Cisco Cloud Network Controller (the Cisco Cloud Network Controller infra tenant) in the master account for this AWS organization.
- Later in this document, you will assign the Organization tag to tenants through the Cisco Cloud Network Controller GUI, using procedures described in [Configuring a Shared Tenant, on page 47](#).



CHAPTER 5

Configuring Cisco Cloud Network Controller Using the Setup Wizard

- [Configuring and Deploying Inter-Site Connectivity](#) , on page 29
- [Gathering On-Premises Configuration Information](#), on page 29
- [Understanding Limitations for Number of Sites, Regions and CCRs](#), on page 30
- [Locating the Cisco Cloud Network Controller IP Address](#), on page 31
- [Configuring Cisco Cloud Network Controller Using the Setup Wizard](#), on page 32
- [Verifying the Cisco Cloud Network Controller Setup Wizard Configurations](#), on page 40

Configuring and Deploying Inter-Site Connectivity

Before you can begin to configure and deploy your Cisco Cloud Network Controller, you must first configure and deploy your Multi-Site and your on-premises Cisco ACI, if you are connecting an on-premises site to cloud sites. The actual configuration for each varies, depending on your requirements and setup. If you are connecting an on-premises site to cloud sites, you will also need to configure and deploy an on-premises IPsec termination device to connect to the Cloud Services Router deployed by Cisco Cloud Network Controller in AWS. See [Components of Extending Cisco ACI Fabric to the Public Cloud](#), on page 4 for more information.

Following are documents that will aid you in the process of configuring and deploying these components:

- Cisco ACI documentation: Available at [Cisco Application Policy Infrastructure Controller \(APIC\) documentation](#), such as [Operating Cisco Application Centric Infrastructure](#) and [Cisco APIC Basic Configuration Guide](#).
- Nexus Dashboard documentation: Available at [Nexus Dashboard documentation](#), such as [Nexus Dashboard Orchestrator Installation and Upgrade Guide](#).
- Cisco Catalyst 8000v Edge Software: Available at [Cisco Catalyst 8000v Edge software documentation](#).

Gathering On-Premises Configuration Information



Note You do not have to gather any information in this section if you are only configuring cloud site-to-cloud site connectivity for your Cisco Cloud Network Controller.

Use the following list to gather and record the necessary on-premises configuration information that you will need throughout these procedures to set up your Cisco Cloud Network Controller:

Necessary On-Premises Information	Your Entry
On-premises IPsec device public IP address	
IPsec termination device to CCR OSPF area	
On-premises APIC IP address	
Cisco Cloud Network Controller IP address	

Understanding Limitations for Number of Sites, Regions and CCRs

Throughout this document, you will be asked to decide on various configurations for sites, regions and CCRs. Following is a list of limitations for each that you should keep in mind as you're making configuration decisions for each.

Sites

The total number of sites that you can have with Cisco Cloud Network Controller depends on the type of configuration that you are setting up:

- **On-premises ACI site-to-cloud site configuration (AWS or Azure):** Multi-Site multi-cloud deployments support any combination of one or two cloud sites (AWS or Azure) and one or two on-premises sites for a maximum total of four sites. The connectivity options are:
 - Hybrid-Cloud: On-premises-to-single cloud site connectivity
 - Hybrid Multi-Cloud: On-premises-to-multiple cloud sites connectivity
- **Multi-Cloud: Cloud site-to-cloud site connectivity (AWS or Azure):** Multi-Site multi-cloud deployments support a combination of:
 - Two cloud sites in EVPN deployment mode (AWS and Azure only)
 - Three cloud sites in BGP IPv4 deployment mode (AWS, Azure, and Google Cloud)

Google Cloud to Google Cloud is not yet supported, either with BGP IPv4 or BGP EVPN.

- **Cloud First: Single-Cloud Configuration:** Multi-Site multi-cloud deployments support a single cloud site (AWS, Azure, or Google Cloud)

Regions

The supported region limits are:

- Sixteen regions can be managed in AWS and Azure clouds. Of the 16, only 4 regions can be external connectivity. All 16 regions can be used for workload deployment.

- All regions can be managed in the Google Cloud. Sixteen regions can be used for workload deployments, but only 4 regions can be used for external connectivity.

CCRs

You can have a certain number of CCRs within some regions, with the following limitations:

- You must have at least one region with CCRs deployed to have inter-VNET (Azure), inter-VPC (AWS), or inter-VRF communications.
- You do not have to have CCRs in every region.
- For regions with CCRs deployed to enable connectivity:
 - CCRs can be deployed on all four managed regions.
 - A maximum of four CCRs per managed region is supported, for a total of 16 CCRs per cloud site.



Note The number of CCRs per managed region differs between AWS and Azure, with four CCRs per region supported for AWS (for a total of 16 CCRs per cloud site) and eight CCRs per region supported for Azure (for a total of 32 CCRs per cloud site).

- CCR deployment in Google Cloud by Cisco Cloud Network Controller is not yet supported.

Locating the Cisco Cloud Network Controller IP Address

These procedures describe how to locate the IP address for the Cisco Cloud Network Controller through the AWS site.

Step 1 Go to the AWS account for the Cisco Cloud Network Controller infra tenant.

Step 2 Click the **Services** link at the top of the screen, then click the **EC2** link.

The **EC2 Dashboard** screen appears.

Step 3 In the EC2 Dashboard screen, you should see text displaying the number of running instances in the **Resources** area (for example, **1 Running Instances**). Click this running instances link.

The **Instances** screen appears.

Step 4 Choose the Cisco Cloud Network Controller instance and copy the IP address that is shown in the **IPv4 Public IP** column. This is the Cisco Cloud Network Controller IP address that you will use to log into the Cisco Cloud Network Controller.

Note You can also get the Cisco Cloud Network Controller IP address by going back to the **CloudFormation** page, clicking on the box next to the Cisco Cloud Network Controller and then clicking on the **Outputs** tab. The Cisco Cloud Network Controller IP address is shown in the **Value** column.

Configuring Cisco Cloud Network Controller Using the Setup Wizard

Follow the procedures in this topic to set up the cloud infrastructure configuration for your Cisco Cloud Network Controller. Cisco Cloud Network Controller will automatically deploy the required AWS constructs and the necessary CCRs.

Before you begin

Following are the prerequisites for this task:

- You have met the requirements that are outlined in [Requirements for Extending the Cisco ACI Fabric to the Public Cloud, on page 13](#) before proceeding with the tasks in this section.
- You have successfully completed the procedures that are provided in [Configuring the Cloud Formation Template Information for the Cisco Cloud Network Controller](#), on page 19.

Step 1 In the AWS site, get the Cisco Cloud Network Controller IP address.

See [Locating the Cisco Cloud Network Controller IP Address, on page 31](#) for those instructions.

Step 2 Open a browser window and, using the secure version of HTTP (`https://`), paste the IP address into the URL field, then press Return to access this Cisco Cloud Network Controller.

For example, `https://192.168.0.0`.

If you see a message asking you to **Ignore Risk and Accept Certificate**, accept the certificate to continue.

Step 3 Enter the following information in the login page for the Cisco Cloud Network Controller:

- **Username:** Enter **admin** for this field.
- **Password:** Enter the password that you provided on the Specify Details page from [Step 12, on page 21](#) in the [Deploying the Cisco Cloud Network Controller in AWS, on page 19](#) procedures.
- **Domain:** If you see the **Domain** field, leave the default Domain entry as-is.

Step 4 Click **Login** at the bottom of the page.

Note If you see an error message when you try to log in, such as `REST Endpoint user authentication datastore is not initialized- Check Fabric Membership Status of this fabric node`, wait for several minutes, then try again after a few minutes. You might also have to refresh the page in order to log in.

The Welcome to Cisco Cloud Network Controller setup wizard page appears.

Step 5 Click **Begin Set Up**.

The **Let's Configure the Basics** page appears, with these areas to be configured:

- **DNS Servers**
- **Region Management**

- **Advanced Settings**
- **Licensing**

- Step 6** In the **DNS Servers** row, click **Edit Configuration**.
The **DNS and NTP** page appears.
- Step 7** In the **DNS and NTP** page, add the DNS, if necessary, and NTP servers.
- A DNS server is already configured by default. Add a DNS server if you want to use a specific DNS server.
 - An NTP server is not configured by default, however, so we recommend that you configure an NTP server. Skip to [7.d, on page 33](#) if you want to configure an NTP server and you do not want to configure a DNS server.
- a) If you want to use a specific DNS server, under the **DNS Servers** area, click **+Add DNS Provider**.
 - b) Enter the IP address for the DNS servers and, if necessary, check the box next to Preferred DNS Provider.
 - c) Click the check mark next to the DNS server, and repeat for any additional DNS servers that you want to add.
 - d) Under the **NTP Servers** area, click **+Add Providers**.
 - e) Enter the IP address for the NTP servers and, if necessary, check the box next to Preferred NTP Provider.
 - f) Click the check mark next to the NTP server, and repeat for any additional NTP servers that you want to add.
- Step 8** When you have finished adding the DNS and NTP servers, click **Save and Continue**.
The **Let's Configure the Basics** page appears again.
- Step 9** In the **Region Management** row, click **Begin**.
The **Region Management** page appears.
- Step 10** Determine if you want to use AWS Transit Gateway.
Use Transit Gateway to avoid using VPN tunnels for connectivity within a region and across the regions where TGW peering is supported. For more information, see the [Increasing Bandwidth Between VPCs by Using AWS Transit Gateway or AWS Transit Gateway Connect](#) document.
In the **Use Transit Gateway** area, click the checkbox next to **Enable** if you want to use AWS Transit Gateway.
- Step 11** In the **Regions to Manage** area, verify that the Cisco Cloud Network Controller home region is selected.
The region that you selected in [Step 2, on page 20](#) in [Deploying the Cisco Cloud Network Controller in AWS, on page 19](#) is the home region and should be selected already in this page. This is the region where the Cisco Cloud Network Controller is deployed (the region that will be managed by Cisco Cloud Network Controller), and will be indicated with the text `Cloud Network Controller deployed` in the Region column.
- Step 12** Select additional regions if you want the Cisco Cloud Network Controller to manage additional regions, and to possibly deploy CCRs to have inter-VPC communication and Hybrid-Cloud, Hybrid Multi-Cloud, or Multi-Cloud connectivity on those other regions.
The CCR can manage four regions, including the home region where Cisco Cloud Network Controller is deployed.
A Cisco Cloud Network Controller can manage multiple cloud regions as a single site. In a typical Cisco ACI configuration, a site represents anything that can be managed by an APIC cluster. If a Cisco Cloud Network Controller cluster manages two regions, those two regions are considered a single site by Cisco ACI.
- Step 13** To deploy cloud routers locally to a region, click to place a check mark in the **Catalyst 8000Vs** check box for that region.

You must have at least one region with Catalyst 8000Vs deployed to have inter-VPC communications. However, if you choose multiple regions in this page, you do not have to have Catalyst 8000Vs in every region that you choose. See [Understanding Limitations for Number of Sites, Regions and CCRs, on page 30](#) for more information.

Step 14 When you have selected all the appropriate regions, click **Next** at the bottom of the page.

The **General Connectivity** page appears.

Step 15 Enter the following information on the **General Connectivity** page.

- If you enabled the AWS Transit Gateway Connect feature in [Step 10, on page 33](#), then the Hub Network fields will be available in this window. Go to [15.a, on page 34](#).
- If you did not enable the AWS Transit Gateway Connect feature in [Step 10, on page 33](#), skip to [15.f, on page 34](#).

a) In the **Hub Network** area, click **Add Hub Network**.

The **Add Hub Network** window appears.

b) In the **Name** field, enter a name for the hub network.

c) In the **BGP Autonomous System Number** field, enter a zero for AWS to choose a number, or enter a value between 64512 and 65534, inclusive, for each hub network, and then click the check mark next to the field.

To configure your own BGP autonomous number, enter a value between 64512 and 65534 for each hub network.

We recommend that you use different numbers for different instances of AWS Transit Gateway.

d) In the **TGW Connect** field, click the checkbox next to **Enable** if you want to use AWS Transit Gateway Connect.

See [Increasing Bandwidth Between VPCs by Using AWS Transit Gateway or AWS Transit Gateway Connect](#) for more information.

e) In the **CIDRs** area, click **Add CIDR**.

This will be the AWS Transit Gateway Connect CIDR block, which will be used as the connect peer IP address (the GRE outer peer IP address) on the Transit Gateway side.

1. In the **Region** field, select the appropriate region.
2. In the **CIDR Block Range** field, enter the CIDR block that will be used as the connect peer IP address on the Transit Gateway side.
3. Click the checkmark to accept these values for this CIDR block.
4. For every managed region that will be using the AWS Transit Gateway Connect feature, repeat these steps to add CIDR blocks to be used for each of those managed regions.

f) To add a subnet pool for the CCRs, click **Add Subnet Pool for Cloud Routers** and enter the subnet in the text box.

The first subnet pool for the first two regions is automatically populated. If you selected more than two regions, you will need to add a subnet for the cloud router to the list for the additional two regions. Addresses from this subnet pool will be used for inter-region connectivity for any additional regions that are added that need to be managed by the Cisco Cloud Network Controller after the first two regions. This must be a valid IPv4 subnet with mask /24.

Note The /24 subnet provided during the Cisco Cloud Network Controller deployment would be sufficient for up to two cloud sites. If you need to manage more than two cloud sites, you need to add more subnets.

- g) In the **IPSec Tunnel Subnet Pool** area, click **Add IPSec Tunnel Subnet Pools**.

The **Add IPSec Tunnel Subnet Pools** window appears.

- h) Enter the subnet pool to be used for IPsec tunnels, if necessary.

This subnet pool is used to create an IPsec tunnel between your cloud router and the router on the branch office or external network. This subnet will be used to address the IPsec tunnel interfaces and loopbacks of the cloud routers used for external connectivity.

You can add more subnets to be used for IPsec tunnels in this area, or delete entries in this area if subnets are not used by any tunnels.

Click the check mark after you have entered in the appropriate subnet pools.

- i) In the **CCRs** area, enter a value in the **BGP Autonomous System Number for CCRs** field.

The BGP ASN can be in the range of 1–65534.

Note Do not use **64512** as the autonomous system number in this field.

- j) In the **Assign Public IP to CCR Interface** field, determine if you want to assign public IP addresses to the Catalyst 8000V interfaces.

Private IP addresses are assigned to the Catalyst 8000V interfaces by default. The **Assign Public IP to CCR Interface** option determines whether public IP addresses will also be assigned to the Catalyst 8000V interfaces or not.

By default, the **Enabled** check box is checked. This means that public IP addresses can be assigned to the Catalyst 8000Vs.

- If you want *public* IP addresses assigned to the Catalyst 8000Vs in addition to the private IP addresses, leave the check in the box next to **Enabled**.
- If you want only *private* IP addresses assigned to the Catalyst 8000Vs, remove the check in the box next to **Enabled** to disable this option.

Note that changing the Catalyst 8000V connectivity from private to public, or vice versa, may cause disruption in your network.

Note Both the public and private IP addresses assigned to a Catalyst 8000V are displayed with the other details of the router in the **Cloud Resources** area. If public IP addresses are not assigned to a Catalyst 8000V, only the private IP addresses are displayed.

- k) In the **Number of Routers Per Region** field, choose the number of CCRs that will be used in each region.

See [Understanding Limitations for Number of Sites, Regions and CCRs, on page 30](#) for more information on any limitations on the number of CCRs per region.

Note If you change the value in this field to increase or decrease the number of CCRs that will be used in each region, wait long enough for the operation to complete before changing the value in this field again to allow time for the registration in the smart license server to synchronize properly.

- If you are decreasing the number of CCRs, wait long enough for those CCRs to get deleted before changing the value in this field again.
- If you are increasing the number of CCRs, wait long enough for those CCRs to get deployed before changing the value in this field again.

- l) In the **Username**, enter the username for the CCR.
- m) In the **Password** field, enter the password for the CCR.
- n) In the **Pricing Type** field, select one of the two types of licensing models:

Note There are two PAYG options for consuming licenses in the AWS marketplace: **Catalyst 8000V Cisco DNA Essentials** and **Catalyst 8000V Cisco DNA Advantage**. Cisco Cloud Network Controller will make use of **Catalyst 8000V Cisco DNA Advantage**.

1. **BYOL**
2. **PAYG**

For the **BYOL Pricing Type**, the steps are as follows:

1. In the **Throughput of the routers** field, choose the throughput of the CCR.

The Cisco Catalyst 8000V supports tier-based (T0/T1/T2/T3) throughput options. The following table lists what AWS EC2 instance is used for different router throughput settings for the Cisco Catalyst 8000V:

CCR Throughput	AWS EC2 Instance
T0 (up to 15M throughput)	c5.xlarge
T1 (up to 100M throughput)	c5.xlarge
T2 (up to 1G throughput)	c5.xlarge
T3 (up to 10G throughput)	c5.9xlarge

Tier2 (T2) is the default throughput supported by Cisco Cloud Network Controller.

Changing the value in this field changes the size of the CCR instance that is deployed. Choosing a higher value for the throughput results in a larger VM being deployed.

Note If you wish to change this value at some point in the future, you must delete the CCR, then repeat the processes in this chapter again and select the new value that you would like in the same **Throughput of the routers** field.

In addition, the licensing of the CCR is based on this setting. You will need the equivalent or higher license in your Smart account for it to be compliant. See [Requirements for the AWS Public Cloud, on page 14](#) for more information.

Note Cloud routers should be undeployed from all regions before changing the router throughput or login credentials.

2. Enter the necessary information in the **TCP MSS** field, if applicable.

The **TCP MSS** option is available to configure the TCP maximum segment size (MSS). This value will be applied all cloud router interfaces, including VPN tunnels towards the cloud and external tunnels towards the on-premises site or other cloud sites. For VPN tunnels towards the cloud, if the cloud provider's MSS value is less than the value that you enter in this field, then the lower value is used; otherwise, the value that you enter in this field is used.

The MSS value affects only TCP traffic, and has no impact on other types of traffic, such as ping traffic.

3. In the **License Token** field, enter the license token for the CCR.

This is the Product Instance Registration token from your Cisco Smart Software Licensing account. To get this license token, go to <http://software.cisco.com>, then navigate to **Smart Software Licensing > Inventory > Virtual Account** to find the Product Instance Registration token.

Note If the public IP addresses are disabled to the CCRs in 15.j, on page 35, the only supported option is **AWS Direct Connect or Azure Express Route to Cisco Smart Software Manager (CSSM)** when registering smart licensing for CCRs with private IP addresses (available by navigating to **Administrative > Smart Licensing**). You must provide reachability to the CSSM through AWS Direct Connect or Azure Express Route in this case. When the public IP addresses are disabled, public internet cannot be used because private IP addresses are being used. The connectivity should therefore use Private Connection, which is AWS Direct Connect or Azure Express Route.

For the **PAYG Pricing Type**, the steps are as follows:

1. In the **VM Type** field, select one of the AWS EC2 Instances as per your requirement.

Cisco Cloud Network Controller will support a range of AWS EC2 instances for cloud networking needs powered by Cisco's Catalyst 8000V virtual router. The table below shows the cloud instance type supported by Cisco Cloud Network Controller on AWS.

AWS EC2 Instance	CCR Throughput	vCPUs	Memory
c5.xlarge	up to 5 Gigabit throughput	4	8 GiB
c5.2xlarge	up to 10 Gigabit throughput	8	16 GiB
c5.4xlarge	up to 10 Gigabit throughput	16	32 GiB
c5.9xlarge	up to 10 Gigabit throughput	36	72 GiB
c5n.xlarge	up to 25 Gigabit throughput	4	10.5 GiB
c5n.2xlarge	up to 25Gigabit throughput	8	21 GiB
c5n.4xlarge	up to 25 Gigabit throughput	16	42 GiB
c5n.9xlarge	up to 50 Gigabit throughput	36	96 GiB

Changing the value in this field changes the other factors of the CCR as listed in the table above. Choosing a higher value for the VM size results in higher throughput.

2. Enter the necessary information in the **TCP MSS** field, if applicable.

The **TCP MSS** option is available to configure the TCP maximum segment size (MSS). This value will be applied all cloud router interfaces, including VPN tunnels towards the cloud and external tunnels towards the on-premises site or other cloud sites. For VPN tunnels towards the cloud, if the cloud provider's MSS value is less than the value that you enter in this field, then the lower value is used; otherwise, the value that you enter in this field is used.

The MSS value affects only TCP traffic, and has no impact on other types of traffic, such as ping traffic.

Note User need not provide the License token on selecting PAYG.

Note All the features supported in BYOL will be supported by PAYG.

Step 16 Click **Save and Continue**.

The **Let's Configure the Basics** page appears again.

Step 17 In the **Advanced Settings** row, click **Edit Configuration**.

The **Advanced Settings** page appears.

Step 18 Make the necessary configurations in the **Advanced Settings** page.

- **Contract Based Routing:** The **Contract Based Routing** setting reflects the current internal VRF route leak policy, which is a global policy under the infra tenant where a Boolean flag is used to indicate whether contracts can drive routes in the absence of route maps:
 - **Off** (no check is in the **yes** box): The default setting. Indicates that routes are not leaked based on contracts, and are leaked based on route maps instead.
 - **On** (a check is in the **yes** box): Indicates that routes are leaked based on contracts in the absence of route maps. When enabled, contracts drive routing when route maps are not configured. When route maps exist, route maps always drives routing.

- **Cloud Network Controller Access Privilege:** Set to **Routing & Security** by default.

If you want to change the access policy, click the scroll-down menu in the **Cisco Cloud Network Controller Access Privilege** field and choose one of the access policies to apply at the VPC (cloud context profile) level.

- **Routing & Security:** The default access policy. If you do not assign an access policy to the Cisco Cloud Network Controller, then the Cisco Cloud Network Controller has the Routing & Security access policy applied to it by default.

Assigning a Routing & Security access policy to a Cisco Cloud Network Controller means that it has full permissions, where it is able to control routing and security.

- **Routing Only:** Assigning a routing-only access policy to a Cisco Cloud Network Controller means that it can control only the routing policy and the network connectivity.

Step 19 Click **Save and Continue**.

You are returned to the **Let's Configure the Basics** page.

Step 20 Depending on the registration mode selected in AWS, the **Licensing** row, shows the **Registration Mode** as **PAYG** or **Smart Licensing**.

Step 21 If the registration mode is shown as **Smart Licensing**, click **Register** and continue to step 22. If the registration mode is shown as **PAYG**, go directly to step 24.

Step 22 Enter the necessary information in the **Smart Licensing** page.

Cisco Smart Licensing is a unified license management system that manages software licenses across Cisco products. To register your Cisco Cloud Network Controller with Cisco Smart Software Licensing, do the following

- Ensure that this product has access to the internet or a Smart Software Manager satellite installed on your network.

- Log in to Smart Account:
 - Smart Software Manager: <https://software.cisco.com/>
 - Smart Software Manager Satellite: <https://www.cisco.com/c/en/us/buy/smart-accounts/software-manager-satellite.html>
- Navigate to the Virtual Account containing the licenses to be used by this Product Instance.
- Generate a Product Instance Registration Token (this identifies your Smart Account) and copy or save it.

To learn more about Smart Software Licensing, visit <https://www.cisco.com/go/smartlicensing>.

Step 23 Click **Register** at the bottom of the page if you entered the necessary licensing information on this page, or click **Continue in Evaluation Mode** if you want to continue in evaluation mode instead.

The **Summary** page appears.

Step 24 Verify the information on the **Summary** page, then click **Close**.

At this point, you are finished with the internal network connectivity configuration for your Cisco Cloud Network Controller.

If this is the first time that you are deploying your Cisco Cloud Network Controller, this process might take quite a bit of time, possibly 30 minutes or so before the process is successfully completed.

What to do next

Determine if you are managing additional sites along with the Cisco Cloud Network Controller site or not:

- If you are managing additional sites (an on-premises site or cloud sites) along with the Cisco Cloud Network Controller site, go to [Managing Cisco Cloud Network Controller Through Multi-Site, on page 41](#).
- If you are setting up a Cloud First configuration, where you are not managing any other sites along with the Cisco Cloud Network Controller site, you will not need to use the Cisco Cisco Nexus Dashboard Orchestrator for additional configurations. However, you will have additional configurations that you must perform in the Cisco Cloud Network Controller GUI in this case. Use the Global Create option in the Cisco Cloud Network Controller GUI to configure the following components:
 - Tenant
 - Application Profile
 - EPG

See [Navigating the Cisco Cloud Network Controller GUI, on page 61](#) and [Configuring Cisco Cloud Network Controller Components, on page 61](#) for more information.

Verifying the Cisco Cloud Network Controller Setup Wizard Configurations

Use the procedures in this topic to verify that the configuration information that you entered in the Cisco Cloud Network Controller Setup Wizard are applied correctly.

In Cisco Cloud Network Controller, verify the following settings:

- Under **Cloud Resources**, click on **Regions** and verify that the regions that you selected are shown as **managed** in the Admin State column.
- Under **Infrastructure**, click on **Inter-Region Connectivity** and verify the information in this screen is correct.
- Under **Infrastructure**, click on **On Premises Connectivity** and verify the information in this screen is correct.
- Click on Dashboard and use the information in the On Premises Connectivity Status and the Inter-Region Connectivity Status boxes to verify that the setup wizard and tunnel configurations were done properly.

What to do next

Complete the multi-site configuration using the procedures provided in [Managing Cisco Cloud Network Controller Through Multi-Site](#), on page 41.



CHAPTER 6

Managing Cisco Cloud Network Controller Through Multi-Site

- [About Cisco Cloud Network Controller and Multi-Site, on page 41](#)
- [Adding the Cisco Cloud Network Controller Site to Multi-Site, on page 42](#)
- [Configuring the Intersite Infrastructure, on page 42](#)
- [Enabling Connectivity Between the Cisco Cloud Network Controller and the ISN Devices, on page 43](#)
- [Configuring a Shared Tenant, on page 47](#)
- [Creating a Schema, on page 49](#)
- [Configuring an Application Profile and the EPGs, on page 50](#)
- [Creating and Associating a Bridge Domain with a VRF, on page 50](#)
- [Creating a Filter for a Contract, on page 51](#)
- [Creating a Contract, on page 51](#)
- [Adding Sites to the Schema, on page 52](#)
- [Configuring Instances in AWS, on page 52](#)
- [Adding an Endpoint Selector, on page 54](#)
- [Verifying the Multi-Site Configurations, on page 58](#)

About Cisco Cloud Network Controller and Multi-Site

If you selected the **Inter-Site Connectivity** option in the **Region Management** page when configuring Cisco Cloud Network Controller using the setup wizard, you will use Multi-Site to manage another site, such as an on-premises site or cloud sites, along with the Cisco Cloud Network Controller site. You do not need the Multi-Site if you selected only the **Cloud Routers** option in the **Region Management** page in the Setup Wizard for Cisco Cloud Network Controller.

Several new pages have been introduced in the Cisco Nexus Dashboard Orchestrator that are used specifically for the management of the Cisco Cloud Network Controller. The topics in this chapter provide information on these new Cisco Cloud Network Controller management pages. Once you have entered the necessary information in these Cisco Cloud Network Controller management pages, the Cisco Cloud Network Controller essentially becomes another site that you manage through the Multi-Site.

If you are managing an on-premises site along with the Cisco Cloud Network Controller site, we recommend that you set up your on-premises site before beginning these procedures, if it is not set up already. See the for those procedures, located here: [Nexus Dashboard Orchestrator Installation and Upgrade Guide](#)

Adding the Cisco Cloud Network Controller Site to Multi-Site

- Step 1** Log in to the Cisco Nexus Dashboard Orchestrator, if you aren't already logged in.
- Step 2** In the Main menu, click **Sites**.
- Step 3** In the **Sites List** page, click **ADD SITES**.
- Step 4** In the **Connection Settings** page, perform the following actions:
- In the **NAME** field, enter the site name.
For example, `cloudsite1`.
 - (Optional) In the **LABELS** field, choose or create a label.
 - In the **APIC CONTROLLER URL** field, enter the URL of the Cisco Cloud Network Controller. This is the public IP address allocated by Amazon Web Services, which is the same public IP address that you used to log into the Cisco Cloud Network Controller at the beginning of the procedures for configuring Cisco Cloud Network Controller using the setup wizard.
For example, `https://192.0.2.1`.
 - In the **USERNAME** field, enter a username.
For example, `admin`. Note that you can also register with any account that has the same privilege as `admin`.
 - In the **PASSWORD** field, enter the password.
 - In the **APIC SITE ID** field, enter a unique site ID, if this field is not already populated automatically.
The site ID must be a unique identifier of the Cisco Cloud Network Controller site. The range must be from 1 to 127.
 - Click **SAVE**.
- Step 5** Verify that Cisco Cloud Network Controller site was added correctly.
- If you are managing multiple sites, all sites should be displayed in the Sites screen in the Cisco Nexus Dashboard Orchestrator. The Cisco Nexus Dashboard Orchestrator automatically detects if the site is an on-premises or a Cisco Cloud Network Controller site.
-

What to do next

Go to [Configuring the Intersite Infrastructure, on page 42](#).

Configuring the Intersite Infrastructure

- Step 1** In the **Sites** screen, click **CONFIGURE INFRA**.
The **Fabric Connectivity Infra** page appears.
- Step 2** In the left pane, under **SITES**, click on the cloud site.

Almost all of the information in the cloud site area is automatically populated and cannot be changed, with the exception of the BGP Password field, described in the next step.

Step 3 Determine if you want to configure a password between your on-premises site and your cloud site:

- If you do *not* want to configure a password between your on-premises site and your cloud site, skip to [Step 4, on page 43](#).
- If you want to configure a password between your on-premises site and your cloud site:
 - a) In the right pane, click on the **BGP Password** field and enter a password.
 - b) Click the Refresh icon at the upper right corner of the CloudSite window.

All of the cloud properties are automatically fetched from the Cisco Cloud Network Controller. A `Site refreshed successfully` message appears, verifying that all the cloud properties were successfully fetched from the Cisco Cloud Network Controller.

Step 4 Click the **Multi-Site** button to toggle this on to enable Multi-Site connectivity in the cloud site.

Step 5 Choose the type of deployment that you would like to use to configure the intersite infrastructure.

When you click the **Deploy** button at the top right of the screen, it shows the following scroll-down menu options:

- **Deploy Only:** Select this option if you are configuring Multi-Cloud (cloud site-to-cloud site) connectivity. This option pushes the configuration to the cloud sites and the Cisco Cloud Network Controller site and enables the end-to-end interconnect connectivity between the cloud sites.
- **Deploy & Download IPN Device config files:** Pushes the configuration to both the on-premises APIC site and the Cisco Cloud Network Controller site and enables the end-to-end interconnect connectivity between the on-premises and the cloud site. In addition, this option downloads a zip file that contains configuration information that you will use to enable connectivity between the CCR deployed in AWS and the on-premises IPsec termination device. A followup screen appears that allows you to select all or some of the configuration files to download.
- **Download IPN Device config files only:** Downloads a zip file that contains configuration information that you will use to enable connectivity between the CCR deployed in AWS and the on-premises IPsec termination device. A followup screen appears that allows you to select all or some of the configuration files to download.

Enabling Connectivity Between the Cisco Cloud Network Controller and the ISN Devices



Note Follow the procedures in this section only if you are enabling connectivity between the on-premises site and the cloud site. If you do not have an on-premises site, skip these procedures and go to [Configuring a Shared Tenant, on page 47](#).

Follow these procedures to manually enable connectivity between CCR deployed in Amazon Web Services and the on-premises IPsec termination device.

By default, the Cisco Cloud Network Controller will deploy a pair of redundant CCRs. The procedures in this section creates two tunnels, one IPsec tunnel from the on-premises IPsec device to each of these CCRs.

The following information provides commands for CCR as your on-premises IPsec termination device. Use similar commands if you are using a different device or platform.

Step 1 Gather the necessary information that you will need to enable connectivity between the CCRs deployed in AWS and the on-premises IPsec termination device.

- If you selected either the **Deploy & Download IPN Device config files** or the **Download IPN Device config files only** option in Cisco Nexus Dashboard Orchestrator as part of the procedures provided in [Configuring the Intersite Infrastructure, on page 42](#), locate the zip file that contains the configuration files for the ISN devices.
- If you are manually locating the information that you need to enable connectivity between the CCRs deployed in AWS and the on-premises IPsec termination device, gather the CCR and Tenant information, as described in the Appendix of the *Cisco Cloud Network Controller Installation Guide*.

Step 2 Log into the on-premises IPsec device.

Step 3 Configure the tunnel for the *first* CCR.

If you downloaded the configuration files for the ISN devices through Cisco Nexus Dashboard Orchestrator, locate the configuration information for the first CCR and enter that configuration information.

Following is an example of what the configuration information for the first CCR might look like:

```
crypto isakmp policy 1
  encryption aes
  authentication pre-share
  group 2
  lifetime 86400
  hash sha
exit

crypto keyring infra:overlay-1-<first-CCR-tunnel-ID>
  pre-shared-key address <first-CCR-elastic-IP-address> key <first-CCR-preshared-key>
exit

crypto isakmp profile infra:overlay-1-<first-CCR-tunnel-ID>
  local-address <interface>
  match identity address <first-CCR-elastic-IP-address>
  keyring infra:overlay-1-<first-CCR-tunnel-ID>
exit

crypto ipsec transform-set infra:overlay-1-<first-CCR-tunnel-ID> esp-aes esp-sha-hmac
  mode tunnel
exit

crypto ipsec profile infra:overlay-1-<first-CCR-tunnel-ID>
  set pfs group2
  set security-association lifetime seconds 86400
exit

interface tunnel <first-CCR-tunnel-ID>
  ip address <peer-tunnel-for-onprem-IPsec-to-first-CCR> 255.255.255.252
  ip virtual-reassembly
  tunnel source <interface>
  tunnel destination <first-CCR-elastic-IP-address>
  tunnel mode ipsec ipv4
```

```

tunnel protection ipsec profile infra:overlay-1-<first-CCR-tunnel-ID>
ip mtu 1476
ip tcp adjust-mss 1460
ip ospf <process-id> area <area-id>
no shut
exit

```

Where:

- <first-CCR-tunnel-ID> is a unique tunnel ID that you assign to this tunnel.
- <first-CCR-elastic-IP-address> is the elastic IP address of the third network interface of the first CCR.
- <first-CCR-preshared-key> is the preshared key of the first CCR.
- <interface> is the interface that is used for connecting to the CCR deployed in Amazon Web Services.
- <peer-tunnel-for-onprem-IPsec-to-first-CCR> is the peer tunnel IP address for the on-premises IPsec device to the first cloud CCR.
- <process-id> is the OSPF process ID.
- <area-id> is the OSPF area ID.

For example:

```

crypto isakmp policy 1
  encryption aes
  authentication pre-share
  group 2
  lifetime 86400
  hash sha
exit

crypto keyring infra:overlay-1-1000
  pre-shared-key address 192.0.2.20 key 123456789009876543211234567890
exit

crypto isakmp profile infra:overlay-1-1000
  local-address GigabitEthernet1
  match identity address 192.0.2.20
  keyring infra:overlay-1-1000
exit

crypto ipsec transform-set infra:overlay-1-1000 esp-aes esp-sha-hmac
  mode tunnel
exit

crypto ipsec profile infra:overlay-1-1000
  set pfs group2
  set security-association lifetime seconds 86400
exit

interface tunnel 1000
  ip address 30.29.1.2 255.255.255.252
  ip virtual-reassembly
  tunnel source GigabitEthernet1
  tunnel destination 192.0.2.20
  tunnel mode ipsec ipv4
  tunnel protection ipsec profile infra:overlay-1-1000
  ip mtu 1476

```

```

ip tcp adjust-mss 1460
ip ospf 1 area 1
no shut
exit

```

Step 4 Configure the tunnel for the *second* CCR.

If you downloaded the configuration files for the ISN devices through Cisco Nexus Dashboard Orchestrator, locate the configuration information for the second CCR and enter that configuration information.

Following is an example of what the configuration information for the second CCR might look like:

```

crypto isakmp policy 1
  encryption aes
  authentication pre-share
  group 2
  lifetime 86400
  hash sha
exit

crypto keyring infra:overlay-1-<second-CCR-tunnel-ID>
  pre-shared-key address <second-CCR-elastic-IP-address> key <second-CCR-preshared-key>
exit

crypto isakmp profile infra:overlay-1-<second-CCR-tunnel-ID>
  local-address <interface>
  match identity address <second-CCR-elastic-IP-address>
  keyring infra:overlay-1-<second-CCR-tunnel-ID>
exit

crypto ipsec transform-set infra:overlay-1-<second-CCR-tunnel-ID> esp-aes esp-sha-hmac
  mode tunnel
exit

crypto ipsec profile infra:overlay-1-<second-CCR-tunnel-ID>
  set pfs group2
  set security-association lifetime seconds 86400
exit

interface tunnel <second-CCR-tunnel-ID>
  ip address <peer-tunnel-for-onprem-IPsec-to-second-CCR> 255.255.255.252
  ip virtual-reassembly
  tunnel source <interface>
  tunnel destination <second-CCR-elastic-IP-address>
  tunnel mode ipsec ipv4
  tunnel protection ipsec profile infra:overlay-1-<second-CCR-tunnel-ID>
  ip mtu 1476
  ip tcp adjust-mss 1460
  ip ospf <process-id> area <area-id>
  no shut
exit

```

For example:

```

crypto isakmp policy 1
  encryption aes
  authentication pre-share
  group 2
  lifetime 86400

```

```

    hash sha
  exit

crypto keyring infra:overlay-1-1001
  pre-shared-key address 192.0.2.21 key 123456789009876543211234567891
exit

crypto isakmp profile infra:overlay-1-1001
  local-address GigabitEthernet1
  match identity address 192.0.2.21
  keyring infra:overlay-1-1001
exit

crypto ipsec transform-set infra:overlay-1-1001 esp-aes esp-sha-hmac
  mode tunnel
exit

crypto ipsec profile infra:overlay-1-1001
  set pfs group2
  set security-association lifetime seconds 86400
exit

interface tunnel 1001
  ip address 30.29.1.6 255.255.255.252
  ip virtual-reassembly
  tunnel source GigabitEthernet1
  tunnel destination 192.0.2.21
  tunnel mode ipsec ipv4
  tunnel protection ipsec profile infra:overlay-1-1001
  ip mtu 1476
  ip tcp adjust-mss 1460
  ip ospf 1 area 1
  no shut
exit

```

Step 5 Repeat these steps for any additional CCRs that you need to configure.

Step 6 Verify that the tunnels are up on your on-premises IPsec device.

For example:

```

ISN_CCR# show ip interface brief | include Tunnel
Interface          IP-Address      OK? Method Status Protocol
Tunnel1000         30.29.1.2      YES manual up      up
Tunnel1001         30.29.1.4      YES manual up      up

```

If you do not see that both tunnels are shown as up, verify the information that you entered in the steps in this section to determine where you might have an issue. Do not proceed to the next section until you see that both tunnels are shown as up.

Configuring a Shared Tenant

Use the procedures in this section to configure a tenant that is shared between the on-premises site and the Cisco Cloud Network Controller site.

Step 1 In the Cisco Nexus Dashboard Orchestrator:

- a) In the Main menu, click **Tenants**.
- b) In the Tenants List area, click **ADD TENANT**.
- c) In the Tenant Details pane, perform the following actions:
 - In the **DISPLAY NAME** field, enter the tenant name.
 - **Optional:** In the **DESCRIPTION** field, enter the a brief description of the tenant.
 - In the **Associated Sites** section, choose the on-premises and the cloud sites.
 - In the **Associated Users** section, choose the users if they are not already selected.
 - Click **SAVE**.

Step 2 Log into the Cisco Cloud Network Controller site and configure the Amazon Web Services account details for this tenant:

- a) On the main Cisco Cloud Network Controller page, under **Application Management**, click **Tenants**.
- b) On the Tenants page, click on the tenant that you just created through the Cisco Nexus Dashboard Orchestrator in the previous step.
- c) Click the expand button at the top right of the screen.

This is the button with the square and up-right-pointing arrow next to the close (X) button.

- d) On the Tenant page, click the Edit button at the top right of the screen. This is the button with the pencil icon next to the Actions field.
- e) On the Edit Tenant page, scroll to the Settings area and enter the necessary information, depending on the access type for the user tenant:

- If the user tenant in Cisco Cloud Network Controller is trusted (if you set up the AWS account for Trusted Tenant using CFT), enter the following information in this page:
 - **AWS Account ID:** Enter the AWS account number for the user tenant (the AWS account that you logged into when setting up the AWS account for Trusted Tenant using the CFT).
 - **Access Type:** Select **Trusted** in this field.

Note The **Cloud Access Key ID** and the **Cloud Secret Access Key** fields are not displayed when you select **Trusted** as the **Access Type**. These fields are not needed for a trusted tenant.

- If the user tenant in Cisco Cloud Network Controller is untrusted (if you set up the AWS account for an Untrusted User Tenant using the AWS access key ID and secret access key), enter the following information in this page:
 - **AWS Account ID:** Enter the AWS account number for the user tenant in this field.
 - **Access Type:** Select **Untrusted** in this field.
 - **Cloud Access Key ID:** Enter the AWS access key ID information for the user tenant in this field.
 - **Cloud Secret Access Key:** Enter the AWS secret access key information for the user tenant in this field.
- If the user tenant in Cisco Cloud Network Controller is a member of an AWS Organization (if you used AWS Organizations to set up your organization and added accounts to this organization either by creating accounts within the organization or by inviting accounts into the organization), and you have deployed Cisco Cloud

Network Controller in the master account of the organization, enter the following information to assign the Organization tag to this tenant:

- **AWS Account ID:** Enter the AWS account number for the user tenant in this field.
- **Access Type:** Select **Organization** in this field.

Note The following applies if you are assigning the Organization tag to this tenant:

- If the **Organization** option is grayed out in this field, that means that you did not deploy the Cisco Cloud Network Controller (the infra tenant) in the master account for an AWS organization. You cannot assign the Organization tag to a tenant if the Cisco Cloud Network Controller (the infra tenant) was not deployed in the master account for an AWS organization. See [Deploying the Cisco Cloud Network Controller in AWS, on page 19](#) for more information.
- If the master account **invited** an existing AWS account to join the organization, verify that you have the `OrganizationAccountAccessRole` IAM role configured in AWS for the organization tenant and that it has Cisco Cloud Network Controller -related permissions available. See [Support for AWS Organizations and Organization User Tenant, on page 8](#) for more information.

Note The **Cloud Access Key ID** and the **Cloud Secret Access Key** fields are not displayed when you select **Organization** as the **Access Type**. These fields are not needed for an organization tenant.

f) Click **Save** at the bottom of the screen.

What to do next

Go to [Creating a Schema, on page 49](#).

Creating a Schema

There are several general Multi-Site procedures that are not specific to the Cisco Cloud Network Controller, but that must be performed as part of the overall Cisco Cloud Network Controller setup if you are managing an on-premises site and a Cisco Cloud Network Controller site through Multi-Site. The following topics provide these general Multi-Site procedures that are part of the overall Cisco Cloud Network Controller setup.

Follow the instructions in this section if you want to create a new schema for the Cisco Cloud Network Controller site.

If you already have a schema that you want to use for the Cisco Cloud Network Controller site, you can skip these steps and go straight to [Adding Sites to the Schema, on page 52](#).

Step 1 In the Main menu, click **Schemas**.

Step 2 On the Schema page, click the **Add Schema** button.

Step 3 On the Untitled Schema page, replace the text `Untitled Schema` at the top of the page with a name for the schema that you intend to create (for example, `Cloudbursting-Schema`).

- Step 4** In the left pane, click **Template 1**.
- Step 5** In the middle pane, click the area **To build your schema please click here to select a tenant**.
- Step 6** In the right pane, access the **Select A Tenant** dialog box and select the tenant that you created in [Configuring a Shared Tenant, on page 47](#) from the drop-down menu.

Configuring an Application Profile and the EPGs

This procedure describes how to configure an application profile and add two EPGs, one for cloud site and one for the on-premises site, where the provider contract is associated with one EPG and the consumer contract is associated with the other EPG.

- Step 1** In the middle pane, locate the Application Profile area, then click + **Application Profile**.
- Step 2** In the right pane, enter the Application Profile name in the **DISPLAY NAME** field.
- Step 3** In the middle pane, click + **Add EPG** to create an EPG for the cloud site.
- Step 4** In the right pane, enter an EPG name in the **DISPLAY NAME** field (for example, `epg1`).
- Step 5** In the middle pane, click + **Add EPG** again, if you want to create an EPG for the on-premises site.
- Step 6** In the right pane, enter an EPG name in the **DISPLAY NAME** field (for example, `epg2`).
- Step 7** Create a VRF:
- In the middle pane, scroll down until you see the VRF area, then click the + in the dotted box.
 - In the right pane, enter the VRF name in the **DISPLAY NAME** field (for example, `vrf1`).
- Step 8** Click **SAVE**.

Creating and Associating a Bridge Domain with a VRF

Follow the procedures in this section to create a bridge domain for the on-premises site and associate it with the VRF. Note that these procedures are not necessary for a cloud-only schema.

- Step 1** In the middle pane, scroll back up to **EPG** and click on the EPG that you created earlier for the on-premises site.
- Step 2** In the right pane, in the **ON-PREM PROPERTIES** area, under **BRIDGE DOMAIN**, create a new bridge domain by typing a name in the field (for example, `bd1`), then click the **Create** area.
- Step 3** In the middle pane, click the bridge domain that you just created.
- Step 4** In the **Virtual Routing & Forwarding** field, select the VRF that you created in [Configuring an Application Profile and the EPGs, on page 50](#).
- Step 5** Scroll down to the **SUBNETS** area and click on the + next to **SUBNET** under the **GATEWAY** heading.
- Step 6** On the **Add Subnet** dialog, enter the **Gateway IP** address and a description for the subnet you plan to add. The Gateway IP address is the on-premises subnet.
- Step 7** In the **Scope** field, select **Advertised Externally**.

Step 8 Click **SAVE**.

Creating a Filter for a Contract

Step 1 In the middle pane, scroll down until you see the Filter area, then click + in the dotted box.

Step 2 In the right pane, enter a name for the filter in the **DISPLAY NAME** field.

Step 3 Click + **Entry** to provide information for your schema filter on the **Add Entry** display:

- Enter a name for the schema filter entry in the **Name** field on the **Add Entry** dialog.
- Optional. Enter a description for the filter in the **Description** field.
- Enter the details as appropriate to filter EPG communication.

For example, to add an entry allowing HTTPS traffic through a filter, choose:

TYPE: IP, IP PROTOCOL: TCP, and DESTINATION PORT RANGE FROM and DESTINATION PORT RANGE TO: https.

- Click **SAVE**.
-

Creating a Contract

Step 1 In the middle pane, scroll down until you see the Contract area, then click + in the dotted box.

Step 2 In the right pane, enter a name for the contract in the **DISPLAY NAME** field.

Step 3 In the **SCOPE** area, leave the selection at VRF.

Step 4 In the **FILTER CHAIN** area, click + **FILTER**.

The Add Filter Chain screen appears.

Step 5 In the **NAME** field, select the filter that you created in [Creating a Filter for a Contract, on page 51](#).

Step 6 In the middle pane, scroll back up to **EPG** and click on the EPG that you created for the cloud site.

Step 7 In the right pane, click + **CONTRACT**.

The Add Contract screen appears.

Step 8 In the **CONTRACT** field, select the contract that you created earlier in this procedure.

Step 9 In the **TYPE** field, select either **CONSUMER** or **PROVIDER**.

Step 10 Scroll to the **CLOUD PROPERTIES** area, then, in the **VIRTUAL ROUTING & FORWARDING** area, choose the VRF that you created in [Configuring an Application Profile and the EPGs, on page 50](#).

Step 11 Click **SAVE**.

Step 12 In the middle pane, scroll back up to **EPG** and click on the EPG that you created for the on-premises site.

Step 13 In the right pane, click + **CONTRACT**.

The Add Contract screen appears.

- Step 14** In the **CONTRACT** field, select the same contract that you created earlier in this procedure.
- Step 15** In the **TYPE** field, select either **CONSUMER** or **PROVIDER**, whatever you did not select for the previous EPG. For example, if you selected **PROVIDER** for the first EPG, select **CONSUMER** for the second EPG.
- Step 16** Scroll to the **CLOUD PROPERTIES** area, then, in the **VIRTUAL ROUTING & FORWARDING** area, choose the same VRF that you created in [Configuring an Application Profile and the EPGs, on page 50](#).

Adding Sites to the Schema

- Step 1** In the left pane, click the + next to **Sites**.
- Step 2** On the **Add Sites** page, add the on-premises and cloud sites to the schema by checking the box next to each, then click **Save**.
- Step 3** Click on the template underneath the cloud site in the left pane to configure the site local properties for the template.
- Step 4** In the middle pane, click on the VRF.
- Step 5** In the right pane, in the **SITE LOCAL PROPERITES** area, enter the following information:
- In the **REGIONS** field, select the Amazon Web Services region that this VRF will be deployed on.
 - In the **CIDRS** field, click **+CIDR**.

The **ADD CLOUD CIDR** dialog appears. Enter the following information:

- **CIDR** — Enter the VPC CIDR information. For example, 11.11.0.0/16.

The CIDR includes the scope of all subnets that are going to be available to an Amazon Web Services VPC.

Note The VPC CIDR information that you enter in this field cannot overlap with the infra VPC CIDR. Verify that the CIDR information that you enter in this field does not overlap with the infra VPC CIDR information that you entered in the **Infra VPC Pool** field in [Step 12, on page 21](#) in [Deploying the Cisco Cloud Network Controller in AWS, on page 19](#).

- **CIDR TYPE** — Select Primary or Secondary. If this is your first CIDR, select Primary for the CIDR type.
- **ADD SUBNETS** — Enter the subnet information and select the zone, then click the check mark. For example, 11.11.1.0/24

Allocate a subnet within the range of the CIDR block for each availability zone.

- Click **SAVE** in the window.

Configuring Instances in AWS

When you configure endpoint selectors for Cisco Cloud Network Controller, either through the Cisco Cloud Network Controller GUI or through the Cisco Nexus Dashboard Orchestrator GUI, you will also need to configure the instances that you will need in AWS that will correspond with the endpoint selectors that you configure for Cisco Cloud Network Controller.

This topic provides the instructions for configuring the instances in AWS. You can use these procedures to configure the instances in AWS either before you configure the endpoint selectors for Cisco Cloud Network Controller or afterward. For example, you might go to your account in AWS and create a custom tag or label in AWS first, then create an endpoint selector using a custom tag or label in Cisco Nexus Dashboard Orchestrator afterward. Or you might create an endpoint selector using a custom tag or label in Cisco Nexus Dashboard Orchestrator first, then go to your account in AWS and create a custom tag or label in AWS afterward.

Step 1 Determine if you configured the cloud context profile through the Cisco Nexus Dashboard Orchestrator GUI or through the Cisco Cloud Network Controller GUI.

You must configure a cloud context profile as part of the AWS instance configuration process, where the cloud context profile, in conjunction with a VRF and a region, represents the AWS VPC in that region. When you configure a cloud context profile using the Cisco Cloud Network Controller GUI, the configurations, such as the VRF and region settings, are pushed out to AWS afterward. A similar action takes place when you configure a Cisco Cloud Network Controller through the Cisco Nexus Dashboard Orchestrator GUI, where these cloud context profile settings are pushed out to AWS as part of the Cisco Cloud Network Controller configuration process through the Cisco Nexus Dashboard Orchestrator GUI.

- If you are configuring the Cisco Cloud Network Controller through the Cisco Nexus Dashboard Orchestrator GUI, then you do not have to manually configure a cloud context profile. Certain cloud context profile configuration settings, such as the VRF and region settings, are pushed out to AWS as part of the Cisco Cloud Network Controller configuration process through the Cisco Nexus Dashboard Orchestrator GUI that you performed in previous sections.
- If you are configuring the cloud context profile through the Cisco Cloud Network Controller GUI, follow the procedures in the *Cisco Cloud Network Controller User Guide, Release 4.1(x)* to configure the cloud context profile, either through the GUI or through the REST API.

Step 2 Review your cloud context profile configuration settings and determine which settings you will use with your AWS instance.

- a) Log in to your Cisco Cloud Network Controller, if you are not already logged in.
- b) From the **Navigation** menu, choose the **Application Management** tab.

When the **Application Management** tab expands, a list of subtab options appear.

- c) Choose the **Cloud Context Profiles** subtab option.

A list of the cloud context profiles that you have created for your Cisco Cloud Network Controller are displayed.

- d) Select the cloud context profile that you will use as part of this AWS instance configuration process.

Various configuration parameters are displayed for this cloud context profile, such as the region, VRF, IP address and subnets. Use the information displayed in this window when you configure the AWS instance.

Step 3 Log in to the Amazon Web Services account for the Cisco Cloud Network Controller user tenant, if you are not logged in already.

Step 4 Go to **Services > EC2 > Instances > Launch Instance**.

Step 5 In the **Choose an Amazon Machine Image (AMI)** page, select an Amazon Machine Image (AMI).

Step 6 In the **Choose an Instance Type** page, select an instance type, then click **Configure Instance Details**.

Step 7 In the **Configure Instance Details** page, enter the necessary information in the appropriate fields.

- In the **Network** field, select your Cisco Cloud Network Controller VRF.

This would be the VRF that is associated with the cloud context profile that you are using as part of this AWS instance configuration process.

- In the **Subnet** field, select the subnet.
- In the **Auto-assign Public IP** field, if you want to have a public IP, select **Enable** from the scroll-down menu.

Step 8 When you have finished entering the necessary information into the **Configure Instance Details** page, click **Add Storage**.

Step 9 In the **Add Storage** page, accept the default values or configure the storage in this page, if necessary, and click **Add Tags**.

Step 10 In the **Add Tags** page, click **Add Tag** and enter the necessary information in the appropriate fields in this page.

Note If you will be using IP Address, Region or Zone for the type of endpoint selector later in these procedures, you do not have to enter any information in this page. In those situations, when you start the instance in AWS, the IP address, region or zone will be discovered by the Cisco Cloud Network Controller and the endpoint will be assigned to the EPG.

- **Key:** Enter the key that you will use when you create a custom tag for the type of endpoint selector that you are adding later in these procedures.
- **Value:** Enter the value that you will be using for this key.
- **Instances:** Check the box for this field.
- **Volumes:** Check the box for this field.

For example, if you are planning on creating a custom tag for a specific building for your endpoint selector later in these procedures (such as building6), you might enter the following values in these fields on this page:

- **Key:** Location
- **Value:** building6

Step 11 Click **Review and Launch**.

The **Select an existing key pair or create a new key pair** page appears. Use the information in this page if you want to ssh to the instance later on.

Adding an Endpoint Selector

On the Cisco Cloud Network Controller, a cloud EPG is a collection of endpoints that share the same security policy. Cloud EPGs can have endpoints in one or more subnets and are tied to a VRF.

The Cisco Cloud Network Controller has a feature called endpoint selector, which is used to assign an endpoint to a Cloud EPG. The endpoint selector is essentially a set of rules run against the cloud instances assigned to the AWS VPC managed by Cisco ACI. Any endpoint selector rules that match endpoint instances will assign that endpoint to the Cloud EPG. The endpoint selector is similar to the attribute-based microsegmentation available in Cisco ACI.

You can configure the endpoint selector either through the Cisco Cloud Network Controller GUI or through the Cisco Nexus Dashboard Orchestrator GUI. There are slight differences in the options available between

the two GUIs, but the general concept and overall procedures to add endpoint selectors is essentially the same between the two.

The procedures in this section describe how to set up the endpoint selectors using the Cisco Nexus Dashboard Orchestrator GUI. For information on setting up the endpoint selectors using the Cisco Cloud Network Controller GUI, see the *Cisco Cloud Network Controller User Guide*.

Step 1 Gather the necessary information from the Amazon Web Services site that you could use for your Cisco Cloud Network Controller endpoint selector.

See [Configuring Instances in AWS, on page 52](#) for those instructions.

Note These steps assume that you are configuring the instance in AWS first, then adding an endpoint selector for Cisco Cloud Network Controller afterward; however, as described in [Configuring Instances in AWS, on page 52](#), you can also add an endpoint selector in Cisco Cloud Network Controller first, then perform this AWS instance configuration step afterward, at the end of these endpoint selector procedures.

Step 2 Log into the Cisco Nexus Dashboard Orchestrator, if you aren't already logged in.

Step 3 In the left pane, click **Schemas**, then select the schema that you created earlier.

Step 4 Determine how you want to create the endpoint selector.

- If you want to create an endpoint selector that could be applied to any additional cloud site in the future, follow these procedures:
 - a. In the left pane, leave the template selected.
Do not select a specific site for these procedures.
 - b. In the middle pane, select the EPG that you created for the cloud site.
 - c. In the right pane, in the **CLOUD PROPERITES** area, click + next to **SELECTORS** to configure the endpoint selector.
 - d. In the **Add New End Point Selector** dialog, enter a name in the **END POINT SELECTOR NAME** field, based on the classification that you use for this endpoint selector.
 - e. Click + **Expression**, then select the type of endpoint selector.
For an endpoint selector created this way, the only option available under the Key field is EPG.
 - f. Go to [Step 5, on page 56](#).
- If you want to create an endpoint selector specifically for this cloud site, follow these procedures:
 - a. In the left pane, select the cloud site.
 - b. In the middle pane, select the EPG that you created for the cloud site.
 - c. In the right pane, in the **SITE LOCAL PROPERITES** area, under the **SELECTORS** area, click + next to **SELECTOR** to configure the endpoint selector.
 - d. In the **Add New End Point Selector** dialog, enter a name in the **END POINT SELECTOR NAME** field, based on the classification that you use for this endpoint selector.

For example, for an endpoint selector with the IP Subnet classification, you might use a name such as `IP-Subnet-EPSelector`.

- e. Click + **Expression**, then select the key that you want to use for the endpoint selector.
- **IP Address**: Used to select by the IP address or subnet.
 - **Region**: Used to select by the AWS region of the endpoint.
 - **Zone**: Used to select by the AWS availability zone of the endpoint.
 - If you want to create a custom tag for the endpoint selector, start typing in the **Type to search or create field** to enter the custom tag or label, then click **Create** on the new field to create a new custom tab or label.
- Using the example earlier in these procedures when you were adding a tag in AWS, you might create the custom tag `Location` in this field, to match the `Location` tag that you added in AWS earlier.

Step 5 In the **Operator** field, choose the operator that you want to use for the endpoint selector.

The options are:

- **Equals**: Used when you have a single value in the Value field.
- **Not Equals**: Used when you have a single value in the Value field.
- **In**: Used when you have multiple comma-separated values in the Value field.
- **Not In**: Used when you have multiple comma-separated values in the Value field.
- **Has Key**: Used if the expression contains only a key.
- **Does Not Have Key**: Used if the expression contains only a key.

Step 6 In the **Value** field, choose which value that you want to use for the endpoint selector, based on the choices that you made for the two previous fields. You can have multiple comma-separated entries in the **Value** field, where a logical OR exists between the entries in this field.

Note The Value field is not displayed if **Has Key** or **Does Not Have Key** is selected for the Operator field.

For example, if you want to have a specific Amazon Web Services availability zone for the endpoint selector, such as `us-west-1a`, you might make the following selections in this screen:

- **Key**: Zone
- **Operator**: Equals
- **Value**: `us-west-1a`

As another example, assume that you used the following values in these fields:

- **Key**: IP
- **Operator**: Has Key
- **Value**: Not available because Has Key was used in the Operator field.

The EPG rules will be applied to all endpoints with an IP address in this situation.

As a final example, assume that you used the following values in these fields:

- **Key**: custom tag: `Location`

- **Operator:** Has Key
- **Value:** Not available because Has Key was used in the Operator field.

In this situation, the EPG rules will be applied to all endpoints with the AWS tag key Location, regardless of the location value.

Step 7 Click the checkmark when you have finished creating this endpoint selector expression.

Step 8 Determine if you want to create additional endpoint selector expressions.

If you create more than one expression under a single endpoint selector, a logical AND exists between those expressions. For example, assume you created two sets of expressions under a single endpoint selector:

- Endpoint selector 1, expression 1:
 - **Key:** Zone
 - **Operator:** Equals
 - **Value:** us-west-1a
- Endpoint selector 1, expression 2:
 - **Key:** IP
 - **Operator:** Equals
 - **Value:** 192.0.2.1/24

In this case, if *both* of these expressions are true (if the availability zone is us-west-1a AND if the IP address belongs to subnet 192.0.2.1/24), then that endpoint will be assigned to the Cloud EPG.

Click the checkmark after every additional expression that you want to create under this endpoint selector.

Step 9 When you have finished creating the expressions for this endpoint selector, click **SAVE** in the lower right corner of the **Add New End Point Selector**.

If you create more than one endpoint selector under an EPG, a logical OR exists between those endpoint selectors. For example, assume you had created endpoint selector 1 as described in the previous step, and then you created a second endpoint selector as described below:

- Endpoint selector 2, expression 1:
 - **Key:** Region
 - **Operator:** In
 - **Value:** us-east-1, us-east-2

In this case:

- If the availability zone is us-west-1a AND the IP address belongs to the 192.0.2.1/24 subnet (endpoint selector 1 expressions)
- OR
- If the region is either us-east-1 or us-east-2 (endpoint selector 2 expression)

Then that end point is assigned to the Cloud EPG.

Step 10 When you have finished creating the endpoint selectors, click **SAVE** in the upper right corner.

Step 11 Click on the **DEPLOY TO SITES** button at the top right corner of the screen to deploy the schema to the sites.

You should see a message saying `Successfully Deployed` at this point.

What to do next

Verify that the Multi-Site areas were configured correctly using the instructions in [Verifying the Multi-Site Configurations, on page 58](#).

Verifying the Multi-Site Configurations

Use the procedures in this topic to verify that the configurations that you entered in the Cisco Nexus Dashboard Orchestrator are applied correctly.

Step 1 Log into the Cisco Cloud Network Controller and verify the following:

- a) Click on Dashboard and use the information in the On Premises Connectivity Status and the Inter-Region Connectivity Status boxes to verify the following:
 - That the tunnels are up from the CCR on AWS to the ISN (IPsec termination point) on-premises and to the VGWs in the user VPCs.
 - That the OSPF neighbors are coming up between the CCR and the ISN on-premises devices.
 - That the BGP EVPN routes for the VRF show the cloud and on-premises routes, and that the cloud routes are populated through the BGP EVPN in the ACI spine switch.
- b) Click on Application Management → Tenants and verify that the tenants were configured correctly.
- c) Click on Application Management → Application Profiles and verify that the application profiles were configured correctly.
- d) Click on Application Management → EPGs and verify that the EPGs were configured correctly.
- e) Click on Application Management → Contracts and verify that the contracts were configured correctly.
- f) Click on Application Management → VRFs and verify that the VRFs were configured correctly.
- g) Click on Application Management → Cloud Context Profiles and verify that the cloud context profiles were configured correctly.
- h) Click on Cloud Resources → Regions and verify that the regions were configured correctly.
- i) Click on Cloud Resources → VPCs and verify that the VPCs were configured correctly.
- j) Click on Cloud Resources → Cloud Endpoints and verify that the cloud endpoints were configured correctly.
- k) Click on Cloud Resources → Routers and verify that the CCRs were configured correctly.

Step 2 Log into on-premises APIC site and verify the schema in APIC.

You should see the shared tenant that you configured in the Cisco Nexus Dashboard Orchestrator is displayed in the tenants area in APIC and the VRF and EPG deployed from the Cisco Nexus Dashboard Orchestrator schema is configured in the on-premises APIC.

Step 3 From a command line, verify that the VRFs were created properly on the CCR on AWS:

```
show vrf
```

If the tenant `t1` and the VRF `v1` is deployed from the Cisco Nexus Dashboard Orchestrator, the CCR output will be similar to the following:

Name	Default RD	Protocols	Interfaces
t1:v1	64514:3080192	ipv4	BD1 Tu4 Tu5

Step 4 From a command line, verify that the tunnels are up between the Cisco Cloud Routers on AWS and the ISN on-premises devices.

You can run the following command on either the CCR on AWS or on the ISN on-premises devices.

```
show ip interface brief | inc Tunnel
```

Output similar to the following should appear:

Interface	IP-Address	OK?	Method	Status	Protocol
Tunnel1	1.2.3.22	YES	manual	up	up
Tunnel2	1.2.3.30	YES	manual	up	up
Tunnel3	1.2.3.6	YES	manual	up	up
Tunnel4	1.2.3.14	YES	manual	up	up

Step 5 From a command line, verify that the OSPF neighbors are up between the CCR on AWS and the ISN on-premises devices:

```
show ip ospf neighbor
```

Output similar to the following should appear:

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.200.10.201	0	FULL/-	00:00:36	1.2.3.13	Tunnel4
20.30.40.50	0	FULL/-	00:00:36	1.2.3.29	Tunnel2
10.202.101.202	0	FULL/-	00:00:38	1.2.3.5	Tunnel3

Step 6 From a command line, verify that the on-premises BGP EVPN neighbors are present in the CCR:

```
show bgp l2vpn evpn summary
```

Output similar to the following should appear:

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.1.1.2	4	100	139	137	99	0	0	01:30:36	6

Step 7 From a command line, verify that the BGP routes for the VRF show both the cloud and on-premises routes.

Note In the current Cisco Cloud Network Controller workflow, a VRF will not be configured on the CCR until the corresponding VPC is created in AWS.

```
show ip route vrf t1:v1
```

Output similar to the following should appear:

```
B 129.1.1.5/32[20/0] via 10.11.0.34, 01:12:41, BD|1
```

```
B    130.1.0.0/16[20/100] via 131.254.4.5, 01:09:55
```



CHAPTER 7

Understanding the Cisco Cloud Network Controller GUI

- [Navigating the Cisco Cloud Network Controller GUI, on page 61](#)
- [Configuring Cisco Cloud Network Controller Components, on page 61](#)

Navigating the Cisco Cloud Network Controller GUI

After you install Cisco Cloud Network Controller, you can use it for extending Cisco Application Centric Infrastructure (ACI) policy to the Amazon Web Services (AWS) or Microsoft Azure public cloud. You do so through the Cisco Cloud Network Controller GUI.

In the Cisco Cloud Network Controller GUI, you can create a tenant, configure application profiles, endpoint groups (EPGs), contracts, filters, and VRFs. You can also view Cisco Cloud Network Controller topology, configurations, and resources.

You perform configuration steps with the **Intent** feature. For instructions on using the **Intent** feature, see the section [Configuring Cisco Cloud Network Controller Components, on page 61](#). Also see the section "Understanding the Cisco Cloud Network Controller GUI Icons" in the *Cisco Cloud Network Controller User Guide*.

The steps for performing basic tasks in Cisco Cloud Network Controller differ from the steps in regular Cisco APIC. However, the functions of the tenant, application profile, and other elements of Cisco APIC are the same. For more information, see the [Cisco Application Centric Infrastructure Fundamentals Guide](#) on Cisco.com.

You view configurations and other information with the left navigation pane. You can choose **Dashboard** (the default view), **Topology**, **Application Management**, **Cloud Resources**, **Operations**, **Infrastructure**, and **Administrative**.

For information about the icons, see the section "Understanding the Cisco Cloud Network Controller GUI Icons" in the *Cisco Cisco Cloud Network Controller User Guide* on Cisco.com.

Configuring Cisco Cloud Network Controller Components

This section provides an overview of performing key tasks in Cisco Cloud Network Controller, including creating a tenant, application profile, and endpoint group (EPG).

Before you begin

You must have installed Cisco Cloud Network Controller. See the previous installation sections in this guide.

-
- Step 1** Log into Cisco Cloud Network Controller.
- Step 2** At the upper right of the **Dashboard** pane, click the icon with an arrow pointing to a bull's-eye.
This icon might be referred to as the **Intent** icon or feature.
- Step 3** In the **What do you want to do?** window, type a term in the search window to bring up a list of options.
For example, if you want to configure a tenant, type the word **tenant** in the search window. The search returns a list of tasks that are related to creating and configuring tenants.
- Step 4** Click a task and perform the configuration steps in the windows that open.
-

What to do next

You can view the configuration in the left navigation pane. Expand the pane by clicking the hamburger icon at the upper left of the **Dashboard** pane. Expand the appropriate heading to view the configurations.

For example, if you've configured a tenant, expand **Application Management** and click **Tenants**. Information about tenants appears in the central work pane.



CHAPTER 8

Performing a System Upgrade, Downgrade or Recovery

- [Important Notes, on page 63](#)
- [Upgrading the Software, on page 68](#)
- [Downgrading the Software, on page 76](#)
- [Performing a System Recovery, on page 91](#)
- [Triggering an Upgrade of the CCRs, on page 91](#)

Important Notes

- [Important Notes For Release 25.0\(3\), on page 63](#)
- [General Important Notes, on page 66](#)

Important Notes For Release 25.0(3)

Following are important notes for release 25.0(3) regarding the installation, upgrade or downgrade procedures for the Cisco Cloud Network Controller:

- Because of the move from the Cisco Cloud Services Router 1000v to the Cisco Catalyst 8000V, you must add the necessary policy before upgrading from a release prior to 25.0(3) to release 25.0(3) or later:
 1. Go to the infra tenant in the AWS portal.
 2. Click **IAM > Policies**.
 3. In the **Policies** window, click the **ApicAdminFullAccess** policy.
The **Summary** page for this policy is displayed.
 4. Click **Edit Policy**.
 5. Click the **JSON** tab.
 6. Copy the entry below and paste it into the policy:

```
{  
  "Effect": "Allow",  
  "Action": "ssm:*",
```

```

    "Resource": "*"
  }

```

7. Click **Review Policy**, then click **Save Changes**.

- The Cisco Catalyst 8000V supports subscription-based licensing. Before upgrading from a release prior to 25.0(3) to release 25.0(3), you must first subscribe to one of the tier-based Cisco Catalyst 8000V licenses.
 - For instructions on subscribing to one of the tier-based Cisco Catalyst 8000V licenses, see [Cisco Catalyst 8000V Edge Software](#).
 - For more information on different throughputs based on the tiers, see [Requirements for the AWS Public Cloud, on page 14](#).

Cisco Cloud Network Controller makes use of the “Cisco DNA Advantage” subscription. For features supported by the “Cisco DNA Advantage” subscription, see [Cisco DNA SoftwareSD-WAN and Routing Matrices](#).

- When you upgrade your Cisco Cloud Network Controller to release 25.0(3), you should then upgrade the CCRs as soon after the Cisco Cloud Network Controller upgrade as possible. For those instructions, see:
 - [Upgrading the Software, on page 68](#)
 - [Triggering an Upgrade of the CCRs, on page 91](#)

Following are examples of how you would go through these upgrade processes:

- **Single-Site Upgrade:** You normally would not need to have CCRs for a single-site AWS deployment. However, if you do have CCRs deployed in this situation, once the Cisco Cloud Network Controller has completed the upgrade to release 25.0(3) and reached the ready state, you must then start the upgrade of the older CCRs (the Cisco Cloud Services Router 1000v) to the newer CCRs (the Cisco Catalyst 8000V) before making any configuration changes.
- **Multi-Cloud/Hybrid-Cloud Upgrade:** As an example of this upgrade process, assume that you have the following setup:
 - Site 1: AWS
 - Site 2: Azure
 - Site 3: On-premises site

You would then upgrade these sites the following way:

1. Upgrade Nexus Dashboard Orchestrator to the 3.7(1) release.
2. Upgrade site 1 (AWS site) to the Cisco Cloud Network Controller release 25.0(3) using the procedures in [Upgrading the Software, on page 68](#).
Wait until this upgrade has reached the steady state before proceeding to the next step.
3. Upgrade the CCRs on site 1 (AWS site) from the older CCRs (the Cisco Cloud Services Router 1000v) to the newer CCRs (the Cisco Catalyst 8000V) using the procedures in [Triggering an Upgrade of the CCRs, on page 91](#).

Wait until the CCRs are fully upgraded to the newer Cisco Catalyst 8000Vs before proceeding to the next step.

4. Once the CCRs on site 1 (AWS site) are fully upgraded, repeat these steps for site 2 (Azure site), where you will first upgrade the Cisco Cloud Network Controller software to release 25.0(3). After that upgrade has reached the steady state, then you will upgrade the CCRs on site 2 to the newer Cisco Catalyst 8000Vs.
- Prior to Cisco Cloud Network Controller release 25.0(3), the older Cisco Cloud Services Router 1000v routers were configured with number-based throughput, as described in [Requirements for the AWS Public Cloud, on page 14](#). Since the Cisco Catalyst 8000V routers will only support tier-based throughput options, during upgrades to release 25.0(3), the Cisco Cloud Network Controller will map the throughput values from the number-based throughput used by the older Cisco Cloud Services Router 1000v routers to the tier-based throughput used by the newer Cisco Catalyst 8000V routers.

The following table shows the mapping of throughput from the older Cisco Cloud Services Router 1000v routers to the newer Cisco Catalyst 8000V routers during an upgrade:

Throughput on Cisco Cloud Services Router 1000v	Throughput on Cisco Catalyst 8000V
10M	T0 (up to 15M throughput)
50M	T1 (up to 100M throughput)
100M	T1 (up to 100M throughput)
250M	T2 (up to 1G throughput)
500M	T2 (up to 1G throughput)
1G	T2 (up to 1G throughput)
2.5G	T3 (up to 10G throughput)
5G	T3 (up to 10G throughput)
7.5G	T3 (up to 10G throughput)
10G	T3 (up to 10G throughput)

For the Cisco Catalyst 8000V routers, end user should choose IOS XE Autonomous boot up mode. The exact license PID for the routers supported in Cisco Cloud APIC is **IOSXE-AUTO-MODE-PF**. The requirements for the license are:

- License Tier : DNA Advantage
- Duration : 3Y, 5Y or 7Y
- Throughput : 25M, 200M, 1G, 10G
- Deployment mode : Autonomous

Below are examples for PID supported in Cisco Cloud APIC for the CCR routers. End users need to use a similar PID by choosing the tier and the number of years for the license.

Throughput on Cisco Cloud Services Router 8000v	PID License
25M (T0)	DNA-P-T0-A-3Y
200M (T1)	DNA-P-T1-A-3Y
1G (T2)	DNA-P-T2-A-3Y
10G (T3)	DNA-P-T3-A-3Y

In the above table, P stands for On-prem, T is the tier, A stands for Advantage, and Y stands for year.

When migrating from the older Cisco Cloud Services Router 1000v routers to the newer Cisco Catalyst 8000V routers during an upgrade, the Cisco Cloud Network Controller will migrate the comparable bandwidth as described above. When these Cisco Catalyst 8000V routers come up, they will try to register for that bandwidth to the smart licensing account. If the smart licensing server does not have these licenses, then the Cisco Catalyst 8000V will fall back to the default bandwidth and will fail to service the existing workload traffic. So you must procure and provision the required Cisco Catalyst 8000V licenses in your smart account before migrating from the older Cisco Cloud Services Router 1000v routers to the newer Cisco Catalyst 8000V routers during an upgrade.

- Similarly, when downgrading from release 25.0(3) to an earlier release, the Cisco Cloud Network Controller will map the throughput values from the tier-based throughput used by the newer Cisco Catalyst 8000V routers to the number-based throughput used by the older Cisco Cloud Services Router 1000v routers.

The following table shows the mapping of throughput from the newer Cisco Catalyst 8000V routers to the number-based throughput used by the older Cisco Cloud Services Router 1000v routers during a downgrade:

Throughput on Cisco Catalyst 8000V	Throughput on Cisco Cloud Services Router 1000v
T0 (up to 15M throughput)	10M
T1 (up to 100M throughput)	100M
T2 (up to 1G throughput)	1G
T3 (up to 10G throughput)	10G



Note Do not make any configuration changes when the Cisco Cloud Network Controller and the CCRs are in incompatible mode. When upgrading to release 25.0(3), verify that both the Cisco Cloud Network Controller and the CCRs are upgraded to that latest release before making any configuration changes.

General Important Notes

Following are general important notes:

- Cisco Cloud Network Controller supports policy-based upgrades for the following upgrade paths:
 - Release 5.2(1) to 25.0(5)

- Release 25.0(1) to 25.0(5)
 - Release 25.0(2) to 25.0(5)
 - Release 25.0(3) to 25.0(5)
 - Release 25.0(4) to 25.0(5)
- For Cisco Cloud Network Controller services compatibility information, see the [Nexus Dashboard and Services Compatibility Matrix](#).
 - When you downgrade from release 5.0(x) to a previous release, as the CCRs downgrade to a lower release, you could see some of the tunnels in a “down” state in the CCRs. This could occur due to stale VPN resources in the AWS accounts that did not get cleaned up.

To correct this issue, manually clean up the stale VPN connections.

- As noted in [Requirements for the AWS Public Cloud, on page 14](#), the supported instance type for the Cisco Cloud Network Controller deployment has changed for release 5.0(x) or later:
 - For releases prior to release 5.0(x), Cisco Cloud Network Controller is deployed using the m4.2xlarge instance.
 - For release 5.0(x) and later, Cisco Cloud Network Controller is deployed using the m5.2xlarge instance.

When upgrading from a 4.2(x) release to release 5.0(x) or later, policy-based upgrades are not supported because you cannot change the instance type through a policy-based upgrade; instead, for these upgrades, you must upgrade using a migration procedure, as provided in [Migration-Based Upgrade, on page 71](#).

- There is an issue with the upgrade process where an upgrade from release 5.2(1g) to any later release will fail.

To work around this issue, enable the **Ignore Compatibility Check** option:

1. Follow the normal upgrade instructions provided in [Upgrading the Software Using the Policy-Based Upgrade Process, on page 70](#) until you get to the **Ignore Compatibility Check** step in the **Schedule Upgrade** window.
2. Enter a check mark in the box next to the **Ignore Compatibility Check** field to enable the **Ignore Compatibility Check** option.

Enabling the **Ignore Compatibility Check** option allows this specific upgrade to proceed normally.

3. Complete the upgrade to the post-5.2(1g) release.
4. Once you have completed the upgrade to the post-5.2(1g) release, return to the **Schedule Upgrade** window and remove the check mark in the box next to the **Ignore Compatibility Check** field.

This disables the **Ignore Compatibility Check** option, which is the default setting for this field.

- Due to the issue described in the previous bullet, if you are upgrading from a release prior to release 5.2(1) to a 5.2(1) release, we recommend that you upgrade directly to release 5.2(1h) and not release 5.2(1g).

Upgrading the Software

The following sections provide information on upgrading the Cisco Cloud Network Controller software using either a policy-based upgrade or a migration-based upgrade.



Note If the policy-based upgrade does not work for some reason, you can upgrade using the migration-based process as described in [Migration-Based Upgrade, on page 71](#).

Upgrading the CCRs

Regardless of the method that you use to upgrade your Cisco Cloud Network Controller software, the CCRs must also be upgraded whenever the Cisco Cloud Network Controller software is upgraded.

- Prior to release 5.2(1), the CCRs are upgraded automatically whenever you trigger an upgrade for the Cisco Cloud Network Controller.
- Beginning with release 5.2(1), you can trigger upgrades to the CCRs and monitor those CCR upgrades, independent from the Cisco Cloud Network Controller upgrades. This is useful to reduce traffic loss by allowing you to split up the upgrades for the management plane (Cisco Cloud Network Controller) and the data plane (CCRs).

See [Triggering an Upgrade of the CCRs, on page 91](#) for more information.

Policy-Based Upgrade

Use the procedures in the following sections to perform a policy-based upgrade of your Cisco Cloud Network Controller software.

Backing Up Your Existing Configuration

We recommend that you back up your existing configuration before performing any policy-based upgrades.

If you decide to downgrade back to that previous release at some point afterward using the procedures provided in [Downgrading the Software, on page 76](#), you will need the backed-up configuration files in order to perform that downgrade successfully.

-
- Step 1** Enable Global AES encryption before performing the backup.
- In your Cisco Cloud Network Controller GUI, navigate to **Infrastructure > System Configuration**.
You should see the **General** tab selected by default; if not, click the **General** tab.
 - Click the pencil icon at the upper right part of the **Global AES Encryption** area.
The **Global AES Encryption Settings** window appears.
 - Click the box next to the **Encryption: Enabled** area, enter a passphrase in the **Passphrase/Confirm Passphrase** fields, then click **Save** at the bottom of the window.

Make a note of the passphrase that you entered in this step, as you will need it if you need as part of the backup restoration process.

Step 2 Make a note of the infra VPC pool that you configured during the stack deployment.

For the infra VPC pool, you might have multiple infra subnet pools, so be sure to locate the information for the infra subnet that was used when you launched the original Cisco Cloud Network Controller.

a) Log into your AWS account for the infra tenant:

<https://signin.aws.amazon.com/>

b) Click the **Services** link at the top of the screen, then click the **CloudFormation** link.

The **CloudFormation** screen appears.

c) On the AWS **CloudFormation** dashboard, click your existing Cisco Cloud Network Controller stack.

The **Stack details** window appears for your Cisco Cloud Network Controller stack.

d) Click the **Parameters** tab in the **Stack details** window.

e) Locate the **pInfraVPCPool** line in the **Parameters** table.

Make a note of the entry in the **pInfraVPCPool** line. This is infra VPC pool that you configured during the stack deployment.

Step 3 Back up your existing configuration.

a) Navigate to **Operations > Backup & Restore**.

b) Click the **Backup Policies** tab.

c) Click **Actions > Create Backup Configuration**.

d) Back up your existing configuration.

For more information on the options available in the **Create Backup Configuration**, see the "Creating a Backup Configuration Using the Cisco Cloud Network Controller GUI" procedure in the *Cisco Cloud Network Controller for AWS User Guide*.

Downloading an Image

Step 1 Log in to your Cisco Cloud Network Controller, if you aren't logged in already.

Step 2 From the **Navigation** menu, choose **Operations > Firmware Management**.

The **Firmware Management** window appears.

Step 3 Click the **Images** tab in the **Firmware Management** window.

Step 4 Click **Actions**, then choose **Add Firmware Image** from the scroll-down menu.

The **Add Firmware Image** pop-up appears.

Step 5 Determine if you want to add the firmware image from a local or a remote location.

- If you want to add the firmware image from a *local* location, click the **Local** radio button in the **Image Location** field. Click the **Choose File** button, then navigate to the folder on your local system with the firmware image that you want to import and select the file. Go to [Step 6, on page 70](#).

- If you want to import the firmware image from a *remote* location, click the **Remote** radio button in the **Image Location** field, then perform the following actions:
 - a) In the **Protocol** field, click either the **HTTP** or the **SCP** radio button.
 - b) In the **URL** field, enter the URL from where the image will be downloaded.
 - If you selected the **HTTP** radio button in the previous step, enter the http source that you want to use to download the software image. An example URL is `10.67.82.87:/home/<username>/cloud-network-controller-dk9.25.0.5f.iso`. Go to [Step 6, on page 70](#).
 - If you selected the **SCP** radio button in the previous step, enter the Secure Copy Protocol (SCP) source that you want to use to download the software image, using the format `<SCP server>:/<path>`. An example URL is `10.67.82.87:/home/<username>/cloud-network-controller-dk9.25.0.5f.iso`.
 - c) In the **Username** field, enter your username for secure copy.
 - d) In the **Authentication Type** field, select the type of authentication for the download. The type can be:
 - **Password**
 - **SSH Key**

The default is **Password**.
 - e) If you selected **Password**, in the **Password** field, enter your password for secure copy. Go to [Step 6, on page 70](#).
 - f) If you selected **SSH Key**, enter the following information:
 - **SSH Key Content** — The SSH Key Content is used to create the SSH Key File which is required when creating a Remote location for the download.

Note The public key is generated at the time of the transfer. After the transfer the key files that were generated in the background are deleted. The temporary key files are stored in dataexport directory of the Cisco Cloud Network Controller.
 - **SSH Key Passphrase** — The SSH Key Passphrase is used to create the SSH Key File which is required when creating a Remote location for the download.

Note The Passphrase field can remain empty.

Step 6 Click **Select**.
Wait for the Cisco Cloud Network Controller firmware images to download.

Upgrading the Software Using the Policy-Based Upgrade Process

Use the procedures in the following sections to perform a policy-based upgrade of your Cisco Cloud Network Controller software.

Before you begin

Verify that you have downloaded an image using the procedures provided in [Downloading an Image, on page 69](#).

-
- Step 1** Back up your existing configuration before performing the policy-based upgrade.
- We recommend that you back up the configuration for your existing release using the information provided in [Backing Up Your Existing Configuration, on page 68](#) before performing a policy-based upgrade.
- After you have completed the policy-based upgrade, if you decide to downgrade back to the previous release at some point afterward using the procedures provided in [Downgrading the Software, on page 76](#), you will need the backed-up configuration files from the previous release in order to perform that downgrade successfully.
- Step 2** In the Cisco Cloud Network Controller GUI, from the **Navigation** menu, choose the **Operations > Firmware Management**.
- The **Firmware Management** window appears.
- Step 3** Click **Schedule Upgrade**.
- The **Schedule Upgrade** pop-up appears.
- If you see a message that says that faults are present in your fabric, we recommend that you resolve these faults before performing an upgrade. See "Viewing Health Details Using the Cisco Cloud Network Controller GUI" in the *Cisco Cisco Cloud Network Controller for AWS User Guide* for more information.
- Step 4** In the **Target Firmware** field, choose a firmware image from the scroll-down menu.
- Step 5** In the **Upgrade Start Time** field, determine if you want to begin the upgrade now or later.
- Click **Now** if you want to schedule the upgrade for now. Go to [Step 6, on page 71](#).
 - Click **Later** if you want to schedule the upgrade for a later date or time, then select the date and time from the pop-up calendar for the scheduled upgrade.
- Step 6** In the **Ignore Compatibility Check** field, leave the setting in the default off (unchecked) setting, unless you are specifically told to disable the compatibility check feature.
- In Cisco Cloud Network Controller, there is a compatibility check feature that verifies if an upgrade path from the currently-running version of the system to a specific newer version is supported or not. The **Ignore Compatibility Check** setting is set to off by default, so the system automatically checks the compatibility for possible upgrades by default.
- Note** If you choose to disable the compatibility check feature by entering a check mark in the box next to the **Ignore Compatibility Check** field, you run the risk of making an unsupported upgrade to your system, which could result in your system going to an unavailable state.
- Step 7** Click **Schedule Upgrade**.
- You can monitor the progress of the upgrade in the main **Firmware Management** window, under the **Upgrade Status** area.
-

Migration-Based Upgrade

The following section provides migration-based upgrade procedures, which will allow you to upgrade without losing traffic flow.

Upgrading Your Cisco Cloud Network Controller Software Using Migration Procedures

This section provides the migration-based upgrade procedures for your Cisco Cloud Network Controller. There should be no effect on traffic with this migration.

Step 1

Enable the encryption passphrase control, if it is not enabled already.

- a) In your Cisco Cloud Network Controller GUI, navigate to **Infrastructure > System Configuration**.

You should be underneath the **General** tab by default; if not, click the **General** tab.

- b) Determine if the encrypted passphrase control is enabled already.

- In the **Global AES Encryption** area, if you see **Yes** underneath the **Encryption** and **Key Configured** fields, then you have the encrypted passphrase control enabled already. Go to [Step 2, on page 72](#).
- If you do not see **Yes** underneath the **Encryption** and **Key Configured** fields:
 1. Click the pencil icon at the upper right part of the **Global AES Encryption** area.
The **Global AES Encryption Settings** window appears.
 2. Click the box next to the **Encryption: Enabled** area, enter a passphrase in the **Passphrase/Confirm Passphrase** fields, then click **Save** at the bottom of the window.

Step 2

Back up your existing Cisco Cloud Network Controller configuration.

There are a number of different ways that you can back up your Cisco Cloud Network Controller configuration. See the [Cisco Cloud Network Controller for AWS Users Guide](#) for more information. Note that if you want to use a remote backup, you will also need to add a remote location first.

Step 3

Terminate the Cisco Cloud Network Controller EC2 instance from the AWS infra account.

- a) Log into your Amazon Web Services account for the Cisco Cloud Network Controller infra tenant and go to the AWS Management Console, if you are not there already:

<https://signin.aws.amazon.com/>

<https://console.aws.amazon.com/>

- b) Go into **Instances** in the EC2 Dashboard in the AWS Management Console.

- c) Locate the Cisco Cloud Network Controller instance.

- For releases prior to release 5.0(x), Cisco Cloud Network Controller is deployed using the m4.2xlarge instance.
- For release 5.0(x) and later, Cisco Cloud Network Controller is deployed using the m5.2xlarge instance.

- d) Check the box next to the Cisco Cloud Network Controller instance to select it, then click **Actions > Instance State > Terminate**.

In the **Terminate Instances** popup window, select **Yes, Terminate** to terminate this instance.

The **Instances** window reappears and the status changes to **shutting-down** in the **Instance State** row for the Cisco Cloud Network Controller instance. Even though you are terminating the Cisco Cloud Network Controller instance here, there should be no traffic drop for your Cisco Cloud Network Controller.

Step 4

Go to the Cisco Cloud Network Controller page on the [AWS Marketplace](#)

Step 5

Click **Continue to Subscribe**.

- Step 6** In the **Subscribe to this software** page, click the **Continue to Configuration** button.
The **Configure this software** page appears.
- Step 7** Select the following parameters:
- **Delivery Method:** Cisco Cloud Network Controller Cloud Formation Template (selected by default)
 - **Software Version:** Select the appropriate version of the Cisco Cloud Network Controller software.
 - **Region:** Region where Cisco Cloud Network Controller will be deployed
- Step 8** Click the **Continue to Launch** button.
The **Launch this software** page appears, which shows a summary of your configuration and lets you launch the cloud formation template.
- Step 9** In the **Choose Action** field, choose **Launch CloudFormation**, then click **Launch** to go directly to the CloudFormation service in the correct region, with the correct Amazon S3 template URL already populated. The **Specify template** page appears within the **Create stack** page.
- Step 10** In the **Specify template** page, make the following selections:
- **Prerequisite- Prepare template** field: Leave the default **Template is ready** option selected.
 - **Specify template** area:
 - In the **Template source** field, leave the default **Amazon S3 URL** option selected.
 - In the **Amazon S3 URL** field, leave the automatically-generated entry as-is.
 - Click **View in designer**.
- Step 11** In the **template1** area in the lower half of the screen:
- Leave the **Choose template language** selection as **JSON**.
 - Place your cursor at the very beginning of the text string on line 1, press the Shift key and scroll down to the bottom of the window to select the entire text string in the window, then copy all of the text in this window (press Ctrl+C, or right-click and select **Copy**).
- Step 12** On your local computer, navigate to an appropriate folder and create a text file, giving it a unique name, and paste the text string that you just copied into the text file.
This will be the Cisco Cloud Network Controller CFT, which has the m5.2xlarge instance type.
- Step 13** Save and close the text file.
- Step 14** Upload the Cisco Cloud Network Controller CFT to AWS.
- a) Log in to the AWS CloudFormation console:
<https://console.aws.amazon.com/cloudformation>
 - b) On the AWS CloudFormation dashboard, click your existing Cisco Cloud Network Controller stack, then click **Update**.
 - c) In the **Update Stack** wizard, in the **Prepare template** screen, select **Replace current template**.
The **Specify template** area appears.

- d) In the **Update Stack** wizard, on the **Specify template** area, select **Upload a template file**.
The **Upload a template file** option appears.
- e) Click **Choose file** underneath the **Upload a template file** option and navigate to the area where you created the Cisco Cloud Network Controller CFT.
- f) Select the Cisco Cloud Network Controller CFT and then click **Next**.
- g) In the **Specify stack details** screen, verify that the instance type shown in the **Other parameters** area at the bottom of the screen is correctly set to **m5.2xlarge**, then click **Next**.
Do not change the instance type to **m4.2xlarge** in this step.
- h) In the **Configure stack options** screen, click **Next**.
- i) In the **Review** screen, click **Update stack**.

The following actions take place at this point:

- The AWS infra detects three IAM resources that will be updated (shown as **False** in the Replacement column).
- The AWS infra detects one EC2 instance that will be replaced (shown as **True** in the Replacement column).

Action	Logical ID	Physical ID	Resource type	Replacement
Modify	rApicAdminFullAccess Policy	arn:aws:iam::70289519:7007:policy/ApicAdminFullAccess	AWS::IAM::ManagedPolicy	False
Modify	rApicAdminReadOnly Role	ApicAdminReadOnly	AWS::IAM::Role	False
Modify	rApicAdminRole	ApicAdmin	AWS::IAM::Role	False
Modify	rCAPICInstance	i-0a767732513c1010c	AWS::EC2::Instance	True

This will bring up the new Cisco Cloud Network Controller instance with the release image, with the same public IP address as you had previously. You can check the progress of the new Cisco Cloud Network Controller instance coming up by navigating back to **Instances** in the EC2 Dashboard in the AWS Management Console.

Step 15 When you see the **Instance State** change to **running**, you can then log into your Cisco Cloud Network Controller as you did previously.

The Cisco Cloud Network Controller will come up with no configurations at this point.

Note If you see an error message when you try to log in, such as **REST Endpoint user authentication datastore is not initialized- Check Fabric Membership Status of this fabric node**, wait for several minutes, then try again after a few minutes. You might also have to refresh the page in order to log in.

Step 16 Enable the same encryption passphrase.

- a) In your Cisco Cloud Network Controller GUI, navigate to **Infrastructure > System Configuration**.
You should be underneath the **General** tab by default; if not, click the **General** tab.
- b) In the **Global AES Encryption** area, click the pencil icon at the upper right part of the **Global AES Encryption** area.

The **Global AES Encryption Settings** window appears.

- c) Click the box next to the **Encryption: Enabled** area, enter the same passphrase in the **Passphrase/Confirm Passphrase** fields that you used in [Step 1, on page 72](#), then click **Save** at the bottom of the window.

Step 17

Import the configuration that you backed up in [Step 2, on page 72](#).

If you configured a remote location when you backed up your configuration, you might have to create the remote location again to access the backup.

- a) In your Cisco Cloud Network Controller GUI, navigate to **Operations > Backup & Restore**.
- b) In the **Backup & Restore** window, click the **Backups** tab.
- c) Click the **Actions** scroll-down menu, then choose **Restore Configuration**.

The **Restore Configuration** window appears.

- d) Enter the necessary information to restore the configuration that you backed up in [Step 2, on page 72](#).

Use the following settings:

- In the **Restore Type** field, choose **Merge**.
- In the **Restore Mode** field, choose **Best Effort**.

Click **Restore Configuration** when you have entered the necessary information in this window. Click the **Job Status** tab in the **Backup & Restore** window to get the status of the backup restore.

Step 18

Run the CapicTenantRole update to change the set for all trusted tenants.

- a) Locate the tenant role CFT.

The tenant role CFT is located in the S3 bucket in the AWS account for the Cisco Cloud Network Controller infra tenant. The name of the S3 bucket is `capic-common-[cloud-network-controller-AccountId]-data` and the tenant role CFT object is `tenant-cft.json` in that bucket. The `cloud-network-controller-AccountId` is the AWS account number for the Cisco Cloud Network Controller infra tenant, which is the account in which Cisco Cloud Network Controller is deployed.

- b) Click the tenant role CFT link.

The **Overview** page for this tenant role CFT appears.

- c) Click the box next to the **tenant-cft.json** entry on the **Overview** page.

A slide-in pane appears for this JSON-formatted tenant role CFT.

- d) Click **Download** to download the tenant role CFT to a location on your computer.

For security reasons, public access to this S3 bucket in AWS is not allowed, so you must download this file and use it in the tenant account.

- e) In AWS, go to the user account of the trusted tenants, then click **CloudFormation**.
- f) On the AWS CloudFormation dashboard, locate the trusted tenant stack and click on the stack name for that trusted tenant.

The stack properties page appears for this particular stack.

- g) Click the **Change sets** tab.
- h) In the **Change sets** area, click **Create change set**.
- i) In the Create change set window for this stack, click **Replace current template**.
- j) In the **Specify template** area, click the circle next to **Upload a template file**, then click the **Choose File** button.

- k) Navigate to the location on your computer where you downloaded the tenant role CFT and select that template file.
- l) Click **Next** in the Create change set window for this stack.

The **Create Change Set** pop-up appears.

- m) Click **Create Change Set** in the **Create Change Set** pop-up window.

The Status will show as **CREATE_PENDING** for a period of time, then will change to **CREATE_COMPLETE**.

- n) Repeat these steps for each trusted tenant.

On each trusted tenant, use this **tenant-cft.json** file to create a change set and run that change set.

Step 19

In your Cisco Cloud Network Controller GUI, verify that all the configurations that you previously had for your Cisco Cloud Network Controller prior to the migration are present now.

Note that for releases prior to 5.2(1), the CCRs will also get upgraded automatically, from the 16.x version to the 17.x version. You can verify this by navigating to **Instances** in the EC2 Dashboard in the AWS Management Console and locating the CCR instances to verify that they are also upgraded.

For release 5.2(1) and later, CCRs are no longer upgraded automatically when the Cisco Cloud Network Controller is upgraded, so you must trigger the CCR upgrades separately after the Cisco Cloud Network Controller has finished upgrading. See [Triggering an Upgrade of the CCRs, on page 91](#) for more information.

Downgrading the Software

The following sections provide the necessary information that you will need to successfully downgrade your Cisco Cloud Network Controller software.

Downgrading the Software: Release 25.0(1) to 5.2(1)

These procedures describe how to downgrade the software from release 25.0(1) to release 5.2(1).

These procedures assume the following scenario:

1. At some point previously, you were running release 5.2(1) and you decided to upgrade to release 25.0(1). Before you performed that upgrade, however, you backed up your release 5.2(1) configuration and saved that backed-up configuration file.
2. You then performed a policy-based upgrade to release 25.0(1) and, at some pointer later on, decided to revert back to release 5.2(1).

These procedures describe how to revert back to release 5.2(1), but you will need that backed-up release 5.2(1) configuration file in order for these downgrade procedures to work.

Step 1

Verify that you have the backed-up release 5.2(1) configuration file, as described in [Backing Up Your Existing Configuration, on page 68](#).

Do not use these procedures to downgrade from release 25.0(1) if you do not have that backed-up release 5.2(1) configuration file available. You will need that backup configuration file for these downgrade procedures.

Step 2 Verify that there are non-home region CCRs configured.

Step 3 Remove the CCRs from the home region.

There will be an intersite traffic loss for around 3-5 minutes while the home region CCRs are getting deleted and the traffic flow switches to the non-home region CCRs.

- a) In your Cisco Cloud Network Controller GUI, click the Intent icon (the icon with an arrow pointing into several circles) and choose **Cisco Cloud Network Controller Setup**.
- b) In the Region Management area, click **Edit Configuration**.

The **Regions to Manage** window appears.

- c) Unselect (remove checks from boxes) in the **Cloud Routers** column for the home region (the region that has the text **Cisco Cloud Network Controller Deployed**).
- d) Click **Next**, then enter the necessary information in the following page and click **Save and Continue**.

The process of removing the CCRs might take 5-10 minutes. You can monitor the process of the CCR removal by looking at the virtual machines in AWS portal.

Note Do not proceed to the next step until the CCRs in the home region have been completely removed.

Step 4 From the infra account in the AWS portal, manually delete all infra VPC peering connections between the home region VPC and any remote region VPCs.

- a) In the navigation pane, choose **Peering connections**.
- b) Select the VPC peering connection, then choose **Actions > Delete VPC peering connection**.
- c) Inside the **Delete VPC peering connection** dialog box, review the connection details, check the **Delete related route table entries** checkbox to remove the necessary routes, then choose **Yes, Delete** to delete the selected VPC peering connection.

Do not alter any VPC peering connections from a remote region VPC to other remote region VPCs.

Step 5 Wait 10-15 minutes for the remaining configurations to be deleted automatically.

The following configurations should be deleted automatically after 10-15 minutes:

- The connect peers for the transit gateway connect attachment in the home region
- The transit gateway connect attachment
- The transit gateway attachment to the infra VPC

If they do not delete automatically, delete them manually as follows:

- a) For the transit gateway connect attachment in the home region, delete the connect peers.
 1. In the navigation pane, choose **Transit Gateway Attachments**.
 2. Select the Connect attachment.
 3. In the **Connect peers** tab, select the Transit Gateway Connect peer and choose **Actions > Delete Connect peer**.
 4. In the confirmation dialog box, choose **Yes, Delete**.
 5. Repeat these steps to delete additional connect peers for the transit gateway connect attachment in the home region.

- b) Delete the transit gateway connect attachment.
 1. In the navigation pane, choose **Transit Gateway Attachments**.
 2. Select the Connect attachment.
 3. Choose **Actions > Delete**.
 4. When prompted for confirmation, choose **Delete**.
- c) Delete the transit gateway attachment to the infra VPC.
 1. In the navigation pane, choose **Transit Gateway Attachments**.
 2. Select the infra VPC attachment only.

There may be other user VPC attachments, so verify that you are selecting the infra VPC attachment for this procedure.
 3. Choose **Actions > Delete**.
 4. When prompted for confirmation, choose **Delete**.

Step 6 Delete the stack.

- a) In the AWS console, navigate to **Services > CloudFormation > Stacks**.
- b) Select the stack that you want to delete.
- c) Click **Delete Stack**.

This will delete the Cisco Cloud Network Controller VM and will attempt to delete the other resources.

Step 7 Wait 15-20 minutes for the stack to be deleted.

If the stack deletion is stuck in *Delete in Progress*, then delete the infra VPC manually in the home region:

- a) In the AWS console, navigate to **Services > Virtual Private Cloud > Your VPCs**.
- b) Select the infra VPC.
- c) Choose **Actions > Delete VPC**.
The **Delete VPC** window appears.
- d) Type `delete` in the **To confirm deletion, type delete in the field** area, then click **Delete**.

Step 8 Recreate a fresh stack with the cloud formation template for the release image that you're downloading to.

Note Alternatively, you can deploy a cloud formation template from the AWS Marketplace in place of steps a-c below.

- a) In the AWS console, navigate to **Services > CloudFormation > Stacks**.
- b) Click **Create Stack > With new resources (standard)**.
The **Create stack** window appears.
- c) In the **Specify template** area, click the circle next to **Upload a template file**, then click the **Choose File** button.
- d) Navigate to the location on your computer with the appropriate JSON-formatted tenant role CFT and select that template file, then click **Next**.

The **Specify Details** page appears within the **Create stack** page.

- e) Enter the necessary information on the **Specify Details** page.
 - **Stack name:** Enter the name for this Cisco Cloud Network Controller configuration.

- **Fabric name:** Leave the default value as-is or enter a fabric name. This entry will be the name for this Cisco Cloud Network Controller.

- **Infra VPC Pool:** Use the same infra VPC pool information that you originally had when you first deployed your Cisco Cloud Network Controller.

You should have noted this infra VPC pool information as part of the procedures in [Backing Up Your Existing Configuration, on page 68](#).

- **Availability Zone:** Select an availability zone for the Cisco Cloud Network Controller subnets from the scroll-down menu.

- **Instance Type:** Select the EC2 instance type.

- **Password/Confirm Password:** Enter and confirm an admin password. This entry is the password that you will use to log into the Cisco Cloud Network Controller after you have enabled SSH access.

- **SSH Key Pair:** Choose the name of the SSH key pair.

You will use this SSH key pair to log into the Cisco Cloud Network Controller.

- **Access Control:** Enter the IP addresses and subnets of the external networks that you will allow to connect to Cisco Cloud Network Controller (for example, 192.0.2.0/24). Only the IP addresses from this subnet are allowed to connect to Cisco Cloud Network Controller. Entering a value of 0.0.0.0/0 means that anyone is allowed to connect to Cisco Cloud Network Controller.

- **Assign Public IP address:** Select whether to assign a public IP address to the Out-of-Band (OOB) management interface for the Cisco Cloud Network Controller or not.

Prior to release 5.2(1), the management interface of the Cisco Cloud Network Controller was assigned a public IP address and a private IP address. Beginning with release 5.2(1), a private IP address is assigned to the management interface of the Cisco Cloud Network Controller and assigning a public IP address is optional. For more information, see the "Private IP Address Support for Cisco Cloud Network Controller and CCR" topic in the *Cisco Cloud Network Controller for AWS User Guide*, Release 5.2(1) or later.

- **true:** Assigns a public IP address to the Out-of-Band (OOB) management interface for the Cisco Cloud Network Controller.

- **false:** Disables the public IP address and assigns a private IP address to the Out-of-Band (OOB) management interface for the Cisco Cloud Network Controller.

f) Click **Next** at the bottom of the screen.

The **Options** page appears within the **Create stack** page.

g) Accept all the default values in the **Options** screen, then click **Next** at the bottom of the **Options** screen.

The **Review** page appears within the **Create stack** page.

h) Verify that all the information on the **Review** page is correct.

If you see any errors on the **Review** page, click the **Previous** button to go back to the page with the incorrect information.

i) When you have verified that all the information on the **Review** page is correct, check the box next to the **I acknowledge that AWS CloudFormation might create IAM resources with custom names** area.

j) Click the **Create stack** button at the bottom of the page.

The **CloudFormation** page reappears, and the Cisco Cloud Network Controller template that you created is displayed with the text **CREATE_IN_PROGRESS** displayed in the Status column.

The system now uses the information that you provided in the template to create the Cisco Cloud Network Controller instance. This process takes 5-10 minutes to complete. You can monitor the progress of the creation process by checking the box next to the name of your Cisco Cloud Network Controller template, then clicking on the Events tab. The text **CREATE_IN_PROGRESS** is displayed in the Status column under the Events tab.

- k) When the **CREATE_COMPLETE** message is shown, verify that the instance is ready before proceeding.
1. Click the **Services** link at the top of the screen, then click the **EC2** link.
The **EC2 Dashboard** screen appears.
 2. In the EC2 Dashboard screen, you should see text displaying the number of running instances in the **Resources** area (for example, **1 Running Instances**). Click this running instances link.
The **Instances** screen appears.
 3. Wait until you see that instance is ready before proceeding.
You will see the new instance going through the **Initializing** stage under **Status check**. Wait until you see the **2/2 Checks Passed** message under **Status check** before proceeding.

Step 9 Enable Global AES encryption using the same passphrase that you noted when you backed up your configuration in [Backing Up Your Existing Configuration, on page 68](#).

- a) In your Cisco Cloud Network Controller GUI, navigate to **Infrastructure > System Configuration**.
It should be underneath the **General** tab by default; if not, click the **General** tab.
- b) Click the pencil icon at the upper right part of the **Global AES Encryption** area.
The **Global AES Encryption Settings** window appears.
- c) Click the box next to the **Encryption: Enabled** area, then enter the passphrase that you noted in [Backing Up Your Existing Configuration, on page 68](#) in the **Passphrase/Confirm Passphrase** fields.
- d) Click **Save** at the bottom of the window.

Step 10 Import the release 5.2(1) configuration that you backed up before you upgraded to release 25.0(1) and verify that the previous configurations converge.

Use the following settings when importing the release 5.2(1) configuration that you backed up:

- In the **Restore Type** field, choose **Merge**.
- In the **Restore Mode** field, choose **Best Effort**.

The home region CCR creation will automatically begin after this step.

Step 11 If the site is managed by ACI Multi-Site Orchestrator/Nexus Dashboard Orchestrator, update the new Cisco Cloud Network Controller VM IP address.

- a) Log into ACI Multi-Site Orchestrator/Nexus Dashboard.
- b) Edit and reregister the site.
 1. In Nexus Dashboard, navigate to **Sites** and click on the correct site.
 2. Click the Details icon to bring up the Overview window.
 3. Click on the pencil icon to edit the information for this site.

4. Click the box next to **Re-register Site** and enter the necessary information to update with the new Cisco Cloud Network Controller VM IP address.
 5. Click **Save**.
- c) Go into ACI Multi-Site Orchestrator/Nexus Dashboard Orchestrator and verify that the site is still managed.
1. In Nexus Dashboard Orchestrator, navigate to **Sites**.
 2. Locate your site and verify that **Managed** is displayed in the **State** column.
- d) Perform a cloud site refresh.
1. In Nexus Dashboard Orchestrator, navigate to **Infrastructure > Infra Configuration**, then click **Configure Infra**.
 2. Select the site in the left nav bar, then click **Refresh**.
Click **Yes** in the confirmation window to continue with the cloud site refresh.
- e) Click **DEPLOY > Deploy Only** to deploy the infra configuration.
-

Downgrading the Software: Release 25.0(2) to 25.0(1) or 5.2(1)

These procedures describe how to downgrade the software from release 25.0(2) to 25.0(1) or 5.2(1).

These procedures assume the following scenario:

1. At some point previously, you were running release 25.0(1) or 5.2(1) and you decided to upgrade to release 25.0(2). Before you performed that upgrade, however, you backed up your release 25.0(1) or 5.2(1) configuration and saved that backed-up configuration file, as described in [Backing Up Your Existing Configuration, on page 68](#).
2. You then performed a policy-based upgrade to release 25.0(2) and, at some pointer later on, decided to revert back to release 25.0(1) or 5.2(1).

These procedures describe how to revert back to that previous release, but you will need that backed-up configuration file for that previous release in order for these downgrade procedures to work.

Step 1 Verify that you have the backed-up configuration file from the previous release, as described in [Backing Up Your Existing Configuration, on page 68](#).

Do not use these procedures to downgrade from release 25.0(2) if you do not have that backed-up configuration file from the previous release available. You will need that backup configuration file for these downgrade procedures.

Step 2 Create a duplicate of the SSH key with the same contents (the same public or private key).

- a) Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
- b) In the navigation pane, choose **Key Pairs**.
- c) Choose **Import key pair**.



Name	Type	Public Key
capic_downgrade	rsa	e5:06:7b:0d:fd:f9:ff:4a:53:ef:70:5a:42:...
<input checked="" type="checkbox"/> capic_upgrade	rsa	d4:db:17:e2:ff:dc:f9:ce:a0:da:12:39:13:...
capic_cisco	rsa	f3:b0:47:b6:6e:42:55:45:ef:5b:39:9f:f4:...

- d) For **Name**, enter a descriptive name for the public key. The name can include up to 255 ASCII characters. It can't include leading or trailing spaces.

Note When you connect to your instance from the EC2 console, the console suggests this name for the name of your private key file.

- e) Either choose **Browse** to navigate to and select your public key, or paste the contents of your public key into the **Public key contents** field.
- f) Choose **Import key pair**.
- g) Verify that the public key that you imported appears in the list of key pairs.


Note If the **Import key pair** process doesn't work for any reason, you can create a new key pair using the **Create key pair** option, and use that in **Step 7**, if necessary.

Step 3

Navigate to the EC2 instance area and terminate the Cisco Cloud Network Controller VM instance.

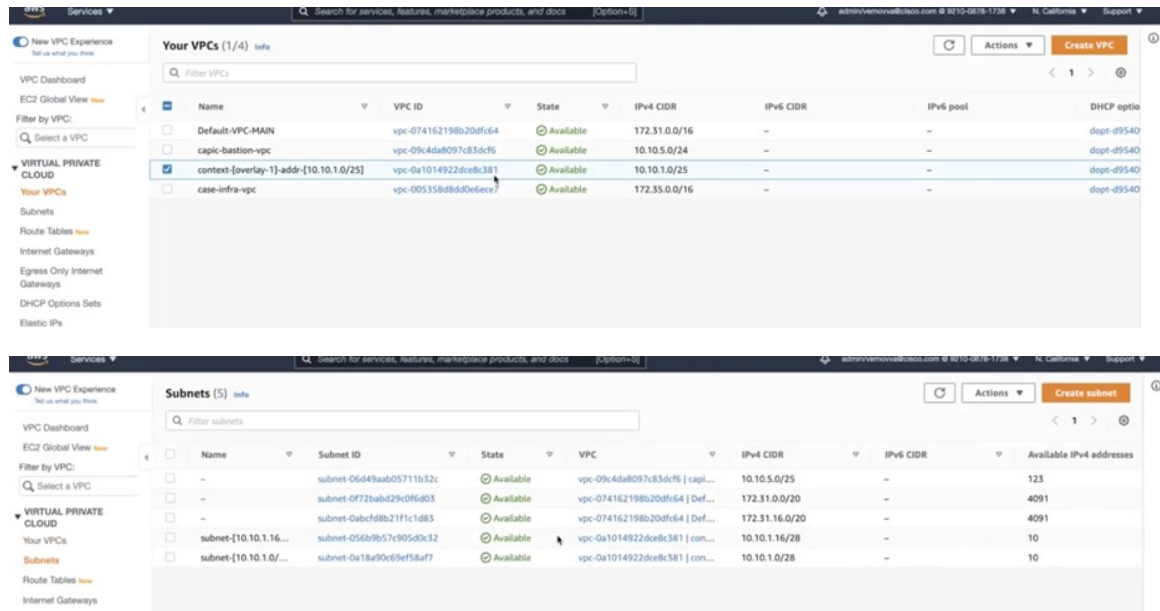
- In the navigation pane, choose **Instances**.
- Check the box next to the Cisco Cloud Network Controller VM instance.
- Right-click on the line for the Cisco Cloud Network Controller VM instance and choose **Instance State > Terminate**.

It will take a few minutes for the Cisco Cloud Network Controller VM instance to terminate.



After you the Cisco Cloud Network Controller VM instance is terminated, the two interfaces associated with the VM will hang at this point. Once the new VM comes up as part of the upgrade, it will get reattached to the same interfaces.


Once the termination process is completed for the Cisco Cloud Network Controller VM, you will note that the VPCs and other network resources, such as the CIDRs and subnets, are still intact.



Step 4 When the termination process is completed for the Cisco Cloud Network Controller VM, go back to the stack and verify that it is still in the running state.

Navigate to the **CloudFormation** area and verify that the Cisco Cloud Network Controller stack is still in the running state.

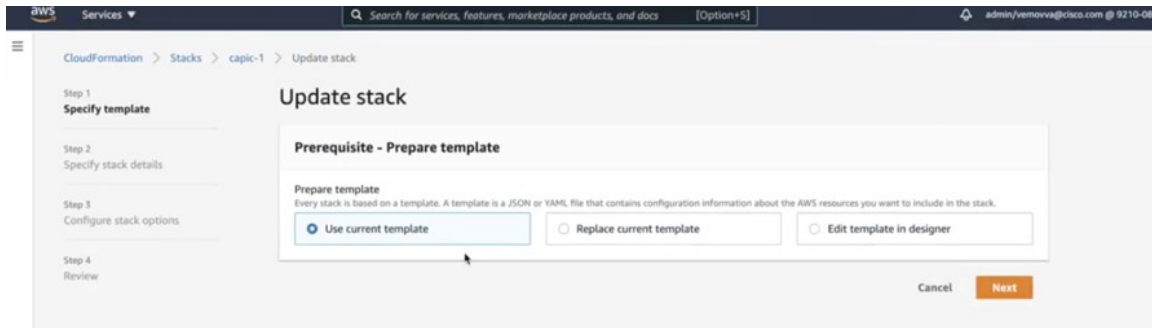
Step 5 Click the circle next to the Cisco Cloud Network Controller stack and click **Update**.



The **Update stack** window appears.

Step 6 Click **Use current template**, then click **Next**.

Because you are not changing anything in the template, you will choose the **Use current template** option in this window.

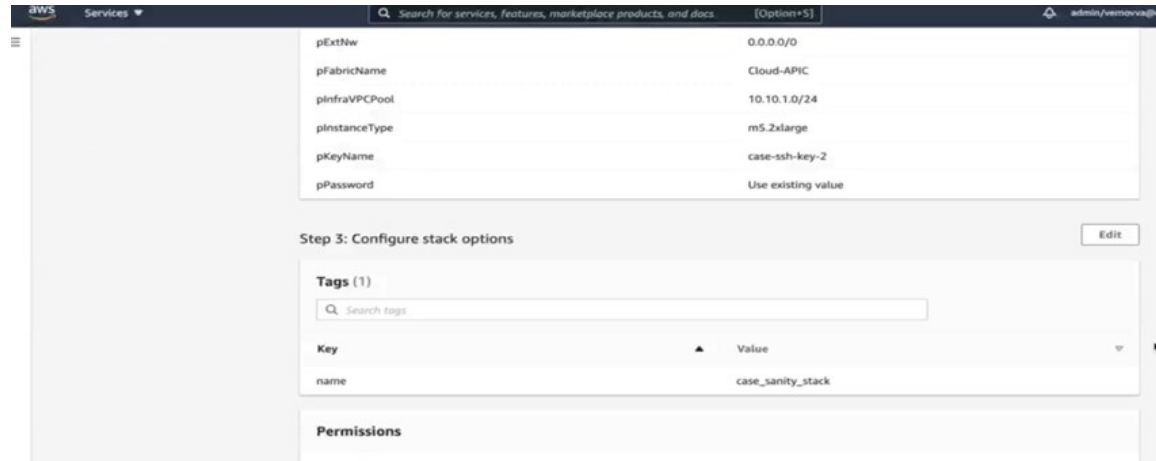


The **Specify stack details** window appears.

Step 7 In the **Specify stack details** window, leave all of the fields as-is, except for the **SSH Key Pair** field.

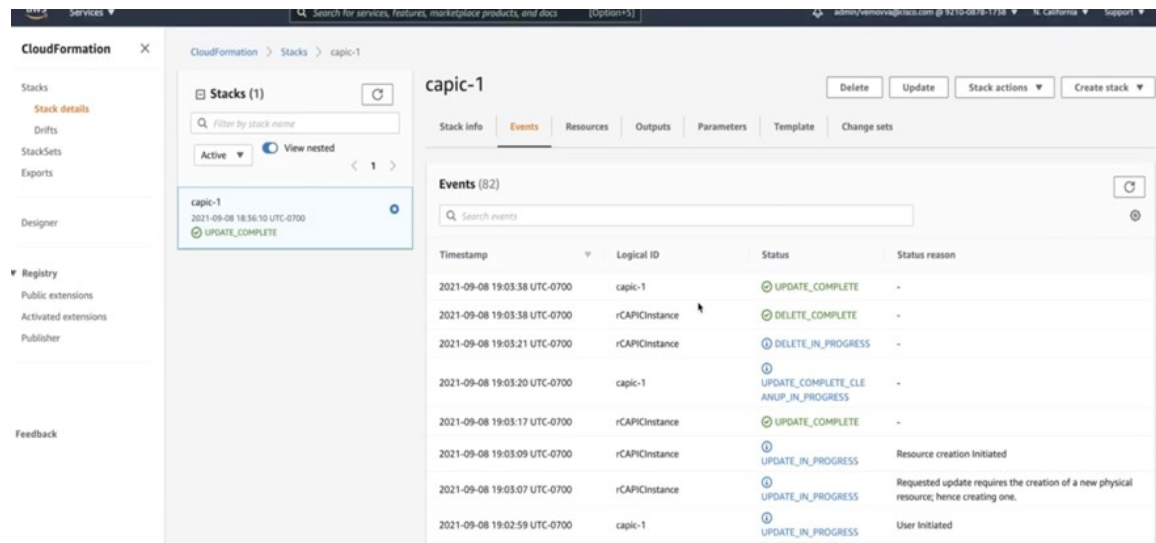
In the **SSH Key Pair** field, select the new SSH key file name that you configured in **Step 2**.

Step 8 Click **Next** at the bottom of the **Specify stack details** window, then navigate through the remaining windows in the **Update stack** windows, verifying that you see the new SSH key file name in the fields in those windows.



Step 9 Click on **Update stack** at the end of the process.

The update for the stack begins.



Step 10 Monitor the progress of the update for the stack.

The update for the stack will go through the following stages:

- AWS will first create a new Cisco Cloud Network Controller VM.
- As part of the stack update, it will try to delete the old Cisco Cloud Network Controller VM, which was already deleted manually.

- The Cisco Cloud Network Controller will be posted in the stack.

- Step 11** Wait until you see the **UPDATE_COMPLETE** message in the **Stacks** window, then navigate back to the **Instances** window.
- The Cisco Cloud Network Controller instance will have the new instance ID and will be using the new SSH key.
 - The old interfaces will be reattached to the new instance, and the CIDRs and subnets will all remain the same.
 - The Cisco Cloud Network Controller management IP address will also be the same.
- Step 12** After roughly 5-10 minutes, verify that the version is correct in the Cisco Cloud Network Controller. Log into your Cisco Cloud Network Controller using the management IP address. You should see the version of release that was running previously, before you upgraded to release 25.0(2).
- Step 13** Enable Global AES encryption using the same passphrase that you noted when you backed up your configuration in [Backing Up Your Existing Configuration, on page 68](#).
- a) In your Cisco Cloud Network Controller GUI, navigate to **Infrastructure > System Configuration**. It should be underneath the **General** tab by default; if not, click the **General** tab.
 - b) Click the pencil icon at the upper right part of the **Global AES Encryption** area. The **Global AES Encryption Settings** window appears.
 - c) Click the box next to the **Encryption: Enabled** area, then enter the passphrase that you noted in [Backing Up Your Existing Configuration, on page 68](#) in the **Passphrase/Confirm Passphrase** fields.
 - d) Click **Save** at the bottom of the window.
- Step 14** Import the configuration for the previous release that you backed up before you upgraded to release 25.0(2) and verify that the previous configurations converge. Use the following settings when importing the configuration for the previous release that you backed up:
- In the **Restore Type** field, choose **Merge**.
 - In the **Restore Mode** field, choose **Best Effort**.
- The home region CCR creation will automatically begin after this step.
- Step 15** If the site is managed by ACI Multi-Site Orchestrator/Nexus Dashboard Orchestrator, update the new Cisco Cloud Network Controller VM IP address.
- a) Log into ACI Multi-Site Orchestrator/Nexus Dashboard.
 - b) Edit and reregister the site.
 1. In Nexus Dashboard, navigate to **Sites** and click on the correct site.
 2. Click the Details icon to bring up the Overview window.
 3. Click on the pencil icon to edit the information for this site.
 4. Click the box next to **Re-register Site** and enter the necessary information to update with the new Cisco Cloud Network Controller VM IP address.
 5. Click **Save**.
 - c) Go into ACI Multi-Site Orchestrator/Nexus Dashboard Orchestrator and verify that the site is still managed.

1. In Nexus Dashboard Orchestrator, navigate to **Sites**.
 2. Locate your site and verify that **Managed** is displayed in the **State** column.
- d) Perform a cloud site refresh.
1. In Nexus Dashboard Orchestrator, navigate to **Infrastructure > Infra Configuration**, then click **Configure Infra**.
 2. Select the site in the left nav bar, then click **Refresh**.
Click **Yes** in the confirmation window to continue with the cloud site refresh.
- e) Click **DEPLOY > Deploy Only** to deploy the infra configuration.

Downgrading the Software: Release 25.0(3) to 25.0(2), 25.0(1), or 5.2(1)

These procedures describe how to downgrade the software from release 25.0(3) to 25.0(2), 25.0(1), or 5.2(1).

These procedures assume the following scenario:

1. At some point previously, you were running release 25.0(2), 25.0(1), or 5.2(1) and you decided to upgrade to release 25.0(3). Before you performed that upgrade, however, you backed up your release 25.0(2), 25.0(1), or 5.2(1) configuration and saved that backed-up configuration file, as described in [Backing Up Your Existing Configuration, on page 68](#).
2. You then performed a policy-based upgrade to release 25.0(3) and, at some pointer later on, decided to revert back to release 25.0(2), 25.0(1), or 5.2(1).

These procedures describe how to revert back to that previous release, but you will need that backed-up configuration file for that previous release in order for these downgrade procedures to work.

Step 1 Verify that you have the backed-up configuration file from the previous release, as described in [Backing Up Your Existing Configuration, on page 68](#).

Do not use these procedures to downgrade from release 25.0(3) if you do not have that backed-up configuration file from the previous release available. You will need that backup configuration file for these downgrade procedures.

Step 2 Create a duplicate of the SSH key with the same contents (the same public or private key).

- a) Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
- b) In the navigation pane, choose **Key Pairs**.
- c) Choose **Import key pair**.
- d) For **Name**, enter a descriptive name for the public key. The name can include up to 255 ASCII characters. It can't include leading or trailing spaces.

Note When you connect to your instance from the EC2 console, the console suggests this name for the name of your private key file.

- e) Either choose **Browse** to navigate to and select your public key, or paste the contents of your public key into the **Public key contents** field.
- f) Choose **Import key pair**.

g) Verify that the public key that you imported appears in the list of key pairs.

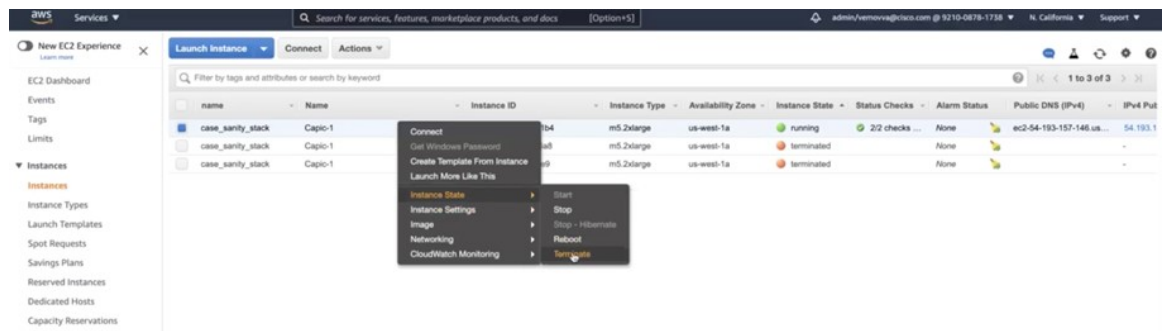
Note If the **Import key pair** process doesn't work for any reason, you can create a new key pair using the **Create key pair** option, and use that in [#unique_67 unique_67_Connect_42_step_it2_mtz_yrb](#), if necessary.

Step 3

Navigate to the EC2 instance area and terminate the Cisco Cloud Network Controller VM instance.

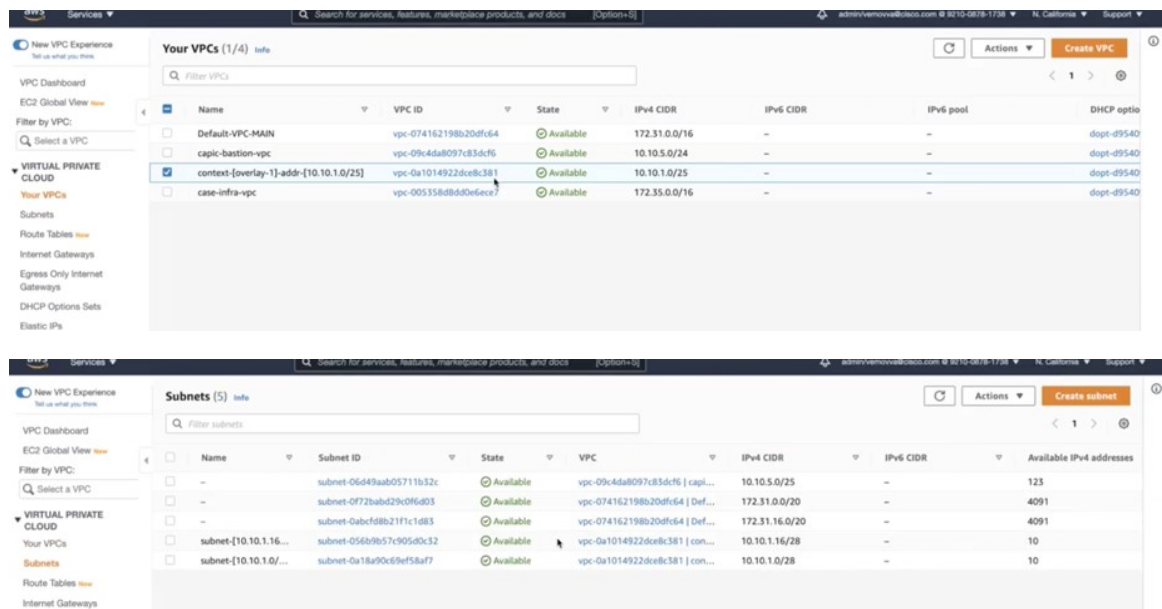
- In the navigation pane, choose **Instances**.
- Check the box next to the Cisco Cloud Network Controller VM instance.
- Right-click on the line for the Cisco Cloud Network Controller VM instance and choose **Instance State > Terminate**.

It will take a few minutes for the Cisco Cloud Network Controller VM instance to terminate.



After you the Cisco Cloud Network Controller VM instance is terminated, the two interfaces associated with the VM will hang at this point. Once the new VM comes up as part of the upgrade, it will get reattached to the same interfaces.

Once the termination process is completed for the Cisco Cloud Network Controller VM, you will note that the VPCs and other network resources, such as the CIDRs and subnets, are still intact.




Step 4

When the termination process is completed for the Cisco Cloud Network Controller VM, go back to the stack and verify that it is still in the running state.

Navigate to the **CloudFormation** area and verify that the Cisco Cloud Network Controller stack is still in the running state.

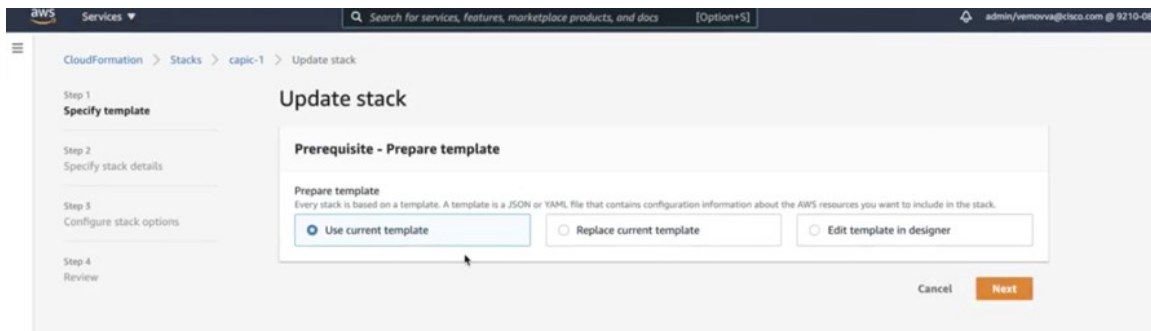
Step 5 Click the circle next to the Cisco Cloud Network Controller stack and click **Update**.



The **Update stack** window appears.

Step 6 Click **Use current template**, then click **Next**.

Because you are not changing anything in the template, you will choose the **Use current template** option in this window.

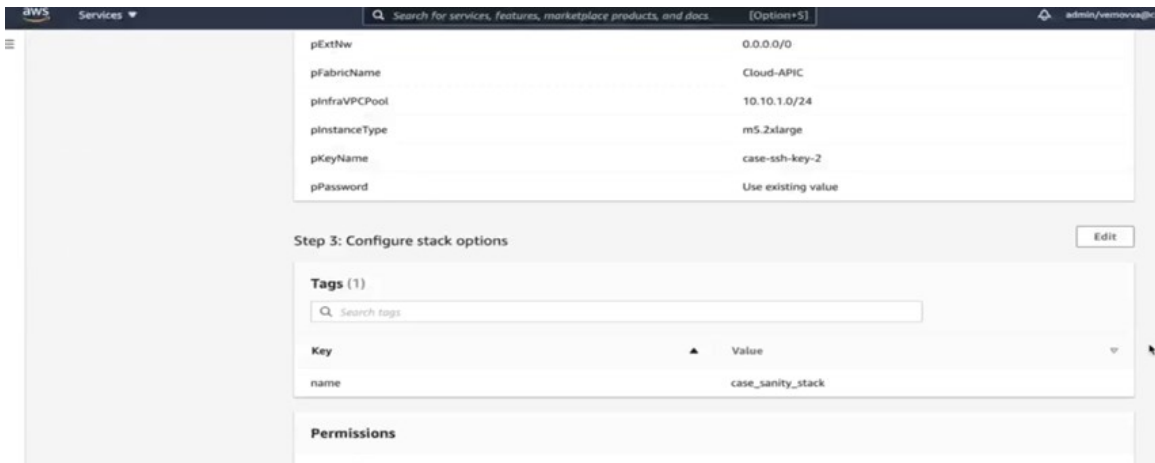


The **Specify stack details** window appears.

Step 7 In the **Specify stack details** window, leave all of the fields as-is, except for the **SSH Key Pair** field.


In the **SSH Key Pair** field, select the new SSH key file name that you configured in [#unique_67 unique_67_Connect_42_step_ayv_tsz_yrb](#).

Step 8 Click **Next** at the bottom of the **Specify stack details** window, then navigate through the remaining windows in the **Update stack** windows, verifying that you see the new SSH key file name in the fields in those windows.



Step 9 Click on **Update stack** at the end of the process.

The update for the stack begins.



Step 10 Monitor the progress of the update for the stack.

The update for the stack will go through the following stages:

- AWS will first create a new Cisco Cloud Network Controller VM.
- As part of the stack update, it will try to delete the old Cisco Cloud Network Controller VM, which was already deleted manually.
- The Cisco Cloud Network Controller will be posted in the stack.

Step 11 Wait until you see the **UPDATE_COMPLETE** message in the **Stacks** window, then navigate back to the **Instances** window.

- The Cisco Cloud Network Controller instance will have the new instance ID and will be using the new SSH key.
- The old interfaces will be reattached to the new instance, and the CIDRs and subnets will all remain the same.

- The Cisco Cloud Network Controller management IP address will also be the same.

Step 12 After roughly 5-10 minutes, verify that the version is correct in the Cisco Cloud Network Controller.

Log into your Cisco Cloud Network Controller using the management IP address. You should see the version of release that was running previously, before you upgraded to release 25.0(3).

Step 13 Trigger a CCR downgrade to the older Cisco Cloud Services Router 1000v.

As part of the upgrade to 25.0(3), you also moved from the older Cisco Cloud Services Router 1000v to the newer Cisco Catalyst 8000V. Downgrading from 25.0(3) to an earlier release therefore requires downgrading the CCR back to the older Cisco Cloud Services Router 1000v.

When the downgrade is completed, the system will recognize that the CCRs are now incompatible with the Cisco Cloud Network Controller. You will see a message saying that the CCRs and the Cisco Cloud Network Controller are incompatible and that any new policies configured for the Cisco Cloud Network Controller will not be applied to the CCRs until you've downgraded the CCRs.

You can begin the process of triggering the CCR downgrade using either of the following methods. Note that while the menu option is shown as **Upgrade CCRs** in both methods, you are actually downgrading the CCRs in this situation by selecting this option.

- In the banner at the top of the screen when your first log into the Cisco Cloud Network Controller, click on the **Upgrade CCRs** link, or
- Through the **CCRs** area in the **Firmware Management** page by navigating to:
Operations > Firmware Management
Click the **CCRs** tab, then choose **Upgrade CCRs**.

Step 14 Enable Global AES encryption using the same passphrase that you noted when you backed up your configuration in [Backing Up Your Existing Configuration, on page 68](#).

a) In your Cisco Cloud Network Controller GUI, navigate to **Infrastructure > System Configuration**.

It should be underneath the **General** tab by default; if not, click the **General** tab.

b) Click the pencil icon at the upper right part of the **Global AES Encryption** area.

The **Global AES Encryption Settings** window appears.

c) Click the box next to the **Encryption: Enabled** area, then enter the passphrase that you noted in [Backing Up Your Existing Configuration, on page 68](#) in the **Passphrase/Confirm Passphrase** fields.

d) Click **Save** at the bottom of the window.

Step 15 Import the configuration for the previous release that you backed up before you upgraded to release 25.0(3) and verify that the previous configurations converge.

Use the following settings when importing the configuration for the previous release that you backed up:

- In the **Restore Type** field, choose **Merge**.
- In the **Restore Mode** field, choose **Best Effort**.

The home region CCR creation will automatically begin after this step.

Step 16 If the site is managed by ACI Multi-Site Orchestrator/Nexus Dashboard Orchestrator, update the new Cisco Cloud Network Controller VM IP address.

- a) Log into ACI Multi-Site Orchestrator/Nexus Dashboard.
 - b) Edit and reregister the site.
 1. In Nexus Dashboard, navigate to **Sites** and click on the correct site.
 2. Click the Details icon to bring up the Overview window.
 3. Click on the pencil icon to edit the information for this site.
 4. Click the box next to **Re-register Site** and enter the necessary information to update with the new Cisco Cloud Network Controller VM IP address.
 5. Click **Save**.
 - c) Go into ACI Multi-Site Orchestrator/Nexus Dashboard Orchestrator and verify that the site is still managed.
 1. In Nexus Dashboard Orchestrator, navigate to **Sites**.
 2. Locate your site and verify that **Managed** is displayed in the **State** column.
 - d) Perform a cloud site refresh.
 1. In Nexus Dashboard Orchestrator, navigate to **Infrastructure > Infra Configuration**, then click **Configure Infra**.
 2. Select the site in the left nav bar, then click **Refresh**.
Click **Yes** in the confirmation window to continue with the cloud site refresh.
 - e) Click **DEPLOY > Deploy Only** to deploy the infra configuration.
-

Performing a System Recovery

The procedures for performing a system recovery is identical to the procedures for performing a migration-based upgrade. Refer to the section [Migration-Based Upgrade, on page 71](#) for those procedures.

Triggering an Upgrade of the CCRs

The following topics provide information and procedures for triggering an upgrade of the CCRs.

Triggering an Upgrade of the CCRs

You can trigger upgrades to the CCRs and monitor those CCR upgrades, independent from the Cisco Cloud Network Controller upgrades. This is useful to reduce traffic loss by allowing you to split up the upgrades for the management plane (Cisco Cloud Network Controller) and the data plane (CCRs).

This feature is enabled by default, where the default assumption is that you will be triggering the upgrades to the CCRs after you trigger an upgrade to the Cisco Cloud Network Controller. You cannot disable this feature once it's enabled.

When this feature is enabled, the proper upgrade sequence for the Cisco Cloud Network Controller and the CCRs is as follows.



Note Following are upper-level steps to describe the overall process for triggering upgrades to the CCRs. For specific step-by-step instructions, see [Triggering an Upgrade of the CCRs Using the Cisco Cloud Network Controller GUI, on page 92](#).

1. Upgrade Cisco Cloud Network Controller using the instructions provided in this chapter.
2. Wait for the Cisco Cloud Network Controller upgrade process to complete. When that upgrade is completed, the system will recognize that the CCRs are now incompatible with the Cisco Cloud Network Controller. You will then see a message saying that the CCRs and the Cisco Cloud Network Controller are incompatible and that any new policies configured for the Cisco Cloud Network Controller will not be applied to the CCRs until you've upgraded the CCRs.
3. View and accept the terms and conditions for the CCRs on the AWS portal.
4. Trigger the CCR upgrade so that it is now at a compatible version as the Cisco Cloud Network Controller.

You can begin the process of triggering the CCR upgrade using either of these two methods:

- In the banner at the top of the screen, click on the **Upgrade CCRs** link, or
- Through the **CCRs** area in the **Firmware Management** page. Navigate to:

Operations > Firmware Management

Click the **CCRs** tab, then choose **Upgrade CCRs**.

You can also trigger the CCR upgrade through the REST API. See [Triggering an Upgrade of the CCRs Using the REST API, on page 93](#) for those instructions.

Guidelines and Limitations

- After you have upgraded the Cisco Cloud Network Controller, if you do not see the message saying that the CCRs and the Cisco Cloud Network Controller are incompatible, you might have to refresh the browser for that message to appear.
- Trigger an upgrade to the CCRs *after* you have upgraded the Cisco Cloud Network Controller. Do not trigger an upgrade to the CCRs before you have upgraded the Cisco Cloud Network Controller.
- Once you have triggered an upgrade to the CCRs, it cannot be stopped.
- If you see any errors after you trigger an upgrade to the CCRs, check and resolve those errors. The CCR upgrade will continue automatically once those CCR upgrade errors have been resolved.

Triggering an Upgrade of the CCRs Using the Cisco Cloud Network Controller GUI

This section describes how to trigger an upgrade to the CCRs using the Cisco Cloud Network Controller GUI. For more information, see [Triggering an Upgrade of the CCRs, on page 91](#).

Step 1 Begin the process of triggering the CCR upgrade to a compatible CCR version.

You can begin the process of triggering the CCR upgrade using either of these two methods:

- In the banner at the top of the screen, click on the **Upgrade CCRs** link, or
- Through the **CCRs** area in the **Firmware Management** page. Navigate to:

Operations > Firmware Management

Click the **CCRs** tab, then choose **Upgrade CCRs**.

A warning appears after clicking **Upgrade CCRs**, stating that upgrading the CCRs will cause the CCRs to reboot, which may cause temporary disruption in traffic.

Step 2 If this is a good time to upgrade the CCRs and have a temporary disruption in traffic, click **Confirm Upgrade** in the warning message.

The CCR software upgrade begins. A banner appears at the top of the screen, saying that the CCR upgrade is in process. Click **View CCR upgrade status** in the message to view the status of the CCR upgrade.

Step 3 Fix any faults that might occur during the upgrade of the CCRs.

If a fault occurs during the upgrade, you can get more information on the fault by navigating to:

Operations > Event Analytics > Faults

Triggering an Upgrade of the CCRs Using the REST API

This section describes how to trigger an upgrade to the CCRs using the REST API. For more information, see [Triggering an Upgrade of the CCRs, on page 91](#).

Set the value for the `routerUpgrade` field to "true" in the cloud template to trigger an upgrade to the CCRs through the REST API (`routerUpgrade="true"`).

```
<polUni>
<fvTenant name="infra">
  <cloudtemplateInfraNetwork name="default" vrfName="overlay-1">
    <cloudtemplateProfile name="defaultxyz" routerUsername="SomeFirstName" routerPassword="SomePass"
routerUpgrade="true">
      </cloudtemplateProfile>
      <cloudtemplateExtSubnetPool subnetpool="10.20.0.0/16"/>
      <cloudtemplateIntNetwork name="default">
        <cloudRegionName provider="aws" region="us-west-1"/>
        <cloudRegionName provider="aws" region="us-west-2"/>
      </cloudtemplateIntNetwork>
      <cloudtemplateExtNetwork name="default">
        <cloudRegionName provider="aws" region="us-west-2"/>
        <cloudtemplateVpnNetwork name="default">
          <cloudtemplateIpSecTunnel peeraddr="23.2.1.1/32" />
          <cloudtemplateIpSecTunnel peeraddr="23.0.1.1/32" />
          <cloudtemplateIpSecTunnel peeraddr="23.1.1.1/32" />
          <cloudtemplateOspf area="0.0.0.1"/>
        </cloudtemplateVpnNetwork>
      <cloudtemplateBgpEvpn peeraddr="34.1.1.1/32" asn="63000" siteId="123" password="abcd1234">

```

```
    />  
      </cloudtemplateExtNetwork>  
    </cloudtemplateInfraNetwork>  
</fvTenant>  
</polUni>
```



APPENDIX **A**

AWS Resources and Naming Conventions

- [AWS Resources and Naming Conventions, on page 95](#)

AWS Resources and Naming Conventions

Following is a list of AWS resources created by the Cisco Cloud Network Controller when it is installed, and the naming conventions used in the Cisco Cloud Network Controller. Use the information in this list to better understand these AWS resources and to avoid using similar names.

Item	Number of Items Used	Naming Convention for Item
S3 buckets	<ul style="list-style-type: none"> • One global (used to store the CFT templates) • One per region (used to store the CloudTrail logs) 	Cisco Cloud Network Controller S3 buckets begin with the prefix <code>capic</code> . Avoid using buckets that begin with this prefix.
Tags	Minimum of two, maximum of eight	Following are the tag keys used: <ul style="list-style-type: none"> • <code>AciDnTag</code> • <code>AciOwnerTag</code> • Name (tag value contains object relative name, or RN) • <code>AciStaleTag</code> (present only if a resource is considered stale by Cisco Cloud Network Controller) • <code>AciResolvedObjDnTag</code> (only for VPC – it carries the Distinguished Name, or DN, for the resolved object) • <code>AciPeerDnTag</code> (only for VPC peering – it carries the DN for the peer VPC)

Item	Number of Items Used	Naming Convention for Item
		Avoid creating tags starting with <code>Aci</code> or <code>Capic</code> .
CloudTrails	One per region	Trail names begin with the prefix <code>capic</code> . Avoid creating trails that begin with this prefix.
CloudWatch events	Three per region	Rules begin with the prefix <code>capic</code> . Avoid creating rules that begin with this prefix.
Simple Queue Service (SQS) queues	One per region	Queue names begin with the prefix <code>capic</code> . Avoid creating queues that begin with this prefix.



APPENDIX **B**

AWS IAM Roles and Permissions

- [AWS IAM Roles and Permissions](#), on page 97

AWS IAM Roles and Permissions



Note Additional information on AWS IAM roles and permissions is available in the [Cisco Cloud Network Controller AWS User Guide](#), including how to configure an AWS provider as one of the following types of tenants:

- Trusted tenant
- Untrusted tenant
- Organization tenant, supported in Release 4.2(3) and later

Specific AWS IAM roles and permissions are required for the installation and operation of the Cisco Cloud Network Controller.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": "iam:*",
    "Resource": "*"
  },
  {
    "Effect": "Allow",
    "Action": "ec2:*",
    "Resource": "*"
  },
  {
    "Effect": "Allow",
    "Action": "cloudformation:*",
    "Resource": "*"
  },
  {
    "Effect": "Allow",
    "Action": "s3:*",
    "Resource": "*"
  },
  {
    "Effect": "Allow",
```

```

    "Action": "sns:*",
    "Resource": "*"
  }
]
}

```

When installing Cisco Cloud Network Controller using the CloudFormation template (CFT), we recommend installation by a user who has the full Administrator Access on AWS (for example, by a user who has the permission policy ARN **arn:aws:iam::aws:policy/AdministratorAccess** attached to it, either directly, by using a role policy, or with a user group). However, if there is no user with AWS Administrator Access available, the user installing Cisco Cloud Network Controller must have this minimum set of permissions:

The above permission set is necessary for a user who installs Cisco Cloud Network Controller using the CFT. Following are more detailed descriptions of each of the required permissions presented above, as shown in the **Action** lines:

- **iam Permissions:** The Cisco Cloud Network Controller instance is an AWS EC2 instance that runs with an AWS role called **ApicAdmin**. This role needs to be created by the CloudFormation stack. Running the Cisco Cloud Network Controller instance with the **ApicAdmin** role allows the Cisco Cloud Network Controller instance to get temporary credentials using the AWS metadata service. This frees the Cisco Cloud Network Controller instance from having to use fixed access key IDs and secret access keys for making AWS API calls.
- **ec2 Permissions:** Needed so that the stack can create the needed VPC, subnets, security groups, and so on. The stack creates the infra VPC, where the Cisco Cloud Network Controller instance is deployed.
- **cloudformation Permissions:** Needed to run the CFT itself.
- **s3 Permissions:** Needed so that the CFT is saved in an S3 bucket based on the needs of the AWS CloudFormation stack.
- **sns Permissions:** Needed to get notifications for running the CloudFormation stack.

For operations, Cisco Cloud Network Controller runs with **ApicAdmin** role. This role has two policies attached, and they get created as part of launching the CloudFormation template:

- **ApicAdminFullAccess Policy:** Permissions listed in this policy allows Cisco Cloud Network Controller to create and manage EC2 and VPC resources, S3 buckets, Resource Groups, account notifications and logs. Note that Cisco Cloud Network Controller only tries to manage the resources it creates. It does not deal with resources created by any other applications.

This policy should have the following permissions:

```

{
  "Version": "2012-10-17",
  "Statement": [{"Resources": {
    "rApicAdminFullAccessPolicy": {
      "Type": "AWS::IAM::ManagedPolicy",
      "Properties": {
        "Description": "Full
Access for ApicAdmin Role",
        "ManagedPolicyName":
"ApicAdminFullAccess",
        "Path": "/",
        "PolicyDocument": {
          "Version":
"2012-10-17",
          "Statement": [{

```

```

"Effect": "Allow",

"Action": "organizations:*",

"Resource": "*"
  }, {
    "Effect": "Allow",
    "Action": [
      "ec2:DeleteCustomerGateway",
      "ec2:DeleteInternetGateway",
      "ec2:DeleteKeyPair",
      "ec2:DeleteSecurityGroup",
      "ec2:DeleteSubnet",
      "ec2:DeleteTransitGateway*",
      "ec2:DeleteVpc*",
      "ec2:DeleteVpn*"
    ],
    "Resource": "*",
    "Condition": {
      "StringEquals": {
        "ec2:ResourceTag/CiscoAciCapic": ""
      }
    }
  }, {
    "Effect": "Allow",
    "Action": [
      "ec2:AcceptTransitGatewayPeeringAttachment",
      "ec2:AcceptVpcPeeringConnection",
      "ec2:AllocateAddress",
      "ec2:AssignPrivateIpAddresses",
      "ec2:AssociateAddress",
      "ec2:AssociateRouteTable",
      "ec2:AssociateTransitGatewayRouteTable",
      "ec2:AssociateVpcCidrBlock",
      "ec2:AttachInternetGateway",
      "ec2:AttachNetworkInterface",
      "ec2:AttachVpnGateway",
      "ec2:AuthorizeSecurityGroup*",
      "ec2:CreateCustomerGateway",
      "ec2:CreateFlowLogs",
      "ec2:CreateInternetGateway",
      "ec2:CreateKeyPair",
      "ec2:CreateNetwork*",
      "ec2:CreateRoute*",
      "ec2:CreateSecurityGroup",
      "ec2:CreateSubnet",
      "ec2:CreateTags",
      "ec2:CreateTransitGateway*",
      "ec2:CreateVpc*",
      "ec2:CreateVpn*",
      "ec2>DeleteFlowLogs",
      "ec2>DeleteNetwork*",
      "ec2>DeleteRoute*",
      "ec2>DeleteTags",
      "ec2:DetachInternetGateway",
      "ec2:DetachNetworkInterface",
      "ec2:DetachVpnGateway",
      "ec2:DisableTransitGatewayRouteTablePropagation",
      "ec2:DisassociateAddress",
      "ec2:DisassociateRouteTable",
      "ec2:DisassociateTransitGatewayRouteTable",
      "ec2:DisassociateVpcCidrBlock",
      "ec2:EnableTransitGatewayRouteTablePropagation",

```

```

        "ec2:EnableVgwRoutePropagation",
        "ec2:GetManagedPrefixListEntries",
        "ec2:GetTransitGatewayRouteTableAssociations",
        "ec2:ModifyInstanceAttribute",
        "ec2:ModifyNetworkInterfaceAttribute",
        "ec2:ModifySubnetAttribute",
        "ec2:ModifyTransitGateway",
        "ec2:ModifyTransitGatewayVpcAttachment",
        "ec2:ModifyVpcAttribute",
        "ec2:ModifyVpcEndpoint",
        "ec2:ReleaseAddress",
        "ec2:ReplaceNetworkAclAssociation",
        "ec2:RevokeSecurityGroup*",
        "ec2:RunInstances",
        "ec2:SearchTransitGatewayRoutes",
        "ec2:StartInstances",
        "ec2:TerminateInstances",
        "ec2:UnassignPrivateIpAddresses"
    ],
    "Resource": "*"
}, {
    "Effect": "Allow",
    "Action": "s3:*",
    "Resource": "*"
}, {
    "Effect": "Allow",
    "Action": [
        "sqs:CreateQueue",
        "sqs>DeleteMessage",
        "sqs>DeleteQueue",
        "sqs:GetQueueAttributes",
        "sqs:GetQueueUrl",
        "sqs:ListQueueTags",
        "sqs:ListQueues",
        "sqs:ReceiveMessage",
        "sqs:SetQueueAttributes",
        "sqs:TagQueue"
    ],
    "Resource": "*"
}, {
    "Effect": "Allow",
    "Action": [
        "elasticloadbalancing:AddListenerCertificates",
        "elasticloadbalancing:AddTags",
        "elasticloadbalancing:CreateListener",
        "elasticloadbalancing:CreateLoadBalancer",
        "elasticloadbalancing:CreateRule",
        "elasticloadbalancing:CreateTargetGroup",
        "elasticloadbalancing>DeleteListener",
        "elasticloadbalancing>DeleteLoadBalancer",
        "elasticloadbalancing>DeleteRule",
        "elasticloadbalancing>DeleteTargetGroup",
        "elasticloadbalancing:DeregisterTargets",
        "elasticloadbalancing:ModifyListener",
        "elasticloadbalancing:ModifyRule",
        "elasticloadbalancing:ModifyTargetGroup",
        "elasticloadbalancing:RegisterTargets",
        "elasticloadbalancing:RemoveListenerCertificates",
        "elasticloadbalancing:RemoveTags",
        "elasticloadbalancing:SetIpAddressType",
        "elasticloadbalancing:SetRulePriorities",
        "elasticloadbalancing:SetSecurityGroups",
        "elasticloadbalancing:SetSubnets"
    ],
}

```

```

    "Resource": "*"
  }, {
    "Effect": "Allow",
    "Action": [
      "acm:DeleteCertificate",
      "acm:ImportCertificate"
    ],
    "Resource": "*"
  }, {
    "Effect": "Allow",
    "Action": ["config:*"],
    "Resource": "*"
  }, {
    "Effect": "Allow",
    "Action": [
      "cloudtrail:AddTags",
      "cloudtrail:CreateTrail",
      "cloudtrail:GetTrailStatus",
      "cloudtrail:StartLogging",
      "cloudtrail>DeleteTrail"
    ],
    "Resource": "*"
  }, {
    "Effect": "Allow",
    "Action": [
      "cloudwatch:DeleteAlarms",
      "cloudwatch:GetMetricStatistics",
      "cloudwatch:PutMetricAlarm"
    ],
    "Resource": "*"
  }, {
    "Effect": "Allow",
    "Action": [
      "logs:CreateLogGroup",
      "logs:CreateLogStream",
      "logs>DeleteLogGroup",
      "logs>DeleteLogStream",
      "logs:FilterLogEvents",
      "logs:ListTagsLogGroup",
      "logs:PutRetentionPolicy",
      "logs:PutLogEvents",
      "logs:TagLogGroup"
    ],
    "Resource": "*"
  }, {
    "Effect": "Allow",
    "Action": [
      "resource-groups:CreateGroup",
      "resource-groups>DeleteGroup",
      "resource-groups:GetGroup",
      "resource-groups:GetGroupQuery",
      "resource-groups:UpdateGroupQuery"
    ],
    "Resource": "*"
  }, {
    "Sid": "CloudWatchEventsFullAccess",
    "Effect": "Allow",
    "Action": [
      "events>DeleteRule",
      "events:DisableRule",
      "events:EnableRule",
      "events:ListRuleNamesByTarget",
      "events:ListRules",
      "events:ListTargetsByRule",

```


- **ApicTenantsAccess Policy:** Permissions listed in this policy allows Cisco Cloud Network Controller to assume the role of tenant accounts and call AWS APIs on those tenant AWS accounts. This allows Cisco Cloud Network Controller to access tenant accounts without having to use the hard credentials of those tenant accounts.

This policy should have the following permissions:

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Action": "sts:AssumeRole",
    "Resource": "*",
    "Effect": "Allow"
  }]
}
```

Note that Cisco Cloud Network Controller itself does not need IAM permissions for its operation because it does not create any IAM policies or roles after its installation.

Cisco Cloud Network Controller will attempt to manage the AWS resources that are created by it, but it will not attempt to manage resources created by other applications, other than listing existing resources as inventory. At the same time, AWS IAM users in those accounts (both the infra account and other tenant accounts) should not interfere with the resources created by Cisco Cloud Network Controller. Therefore, all resources created by Cisco Cloud Network Controller on AWS have at least one of these two tags applied on them:

- **AciDnTag**
- **AciOwnerTag**

Therefore, when you create AWS IAM users who have permission to create, delete or update EC2, VPC and other resources, you must prevent these users from accessing or modifying the resources created and managed by Cisco Cloud Network Controller. Such restrictions should apply on both the infra and other user tenant accounts. AWS account administrators should use the above two tags to prevent users from accessing or modifying the resources created and managed by Cisco Cloud Network Controller.

For example, you might have an access policy similar to the following for an IAM user to prevent unintended access to resources managed by Cisco Cloud Network Controller:

```
{
  "Effect": "Deny",
  "Action": [
    "ec2:*"
  ],
  "Resource": "*",
  "Condition": {
    "StringLike": {
      "ec2:ResourceTag/AciDnTag": "*"
    }
  }
}
```




APPENDIX C

Tenant-Region Management

- [Tenant-Region Management](#), on page 105

Tenant-Region Management

Deploying Tenant Policies in Different Regions

Cisco Cloud Network Controller enforces ownership checks to prevent deployment of policies in the same tenant-region combination, done either intentionally or by mistake. For example, assume that one Cisco Cloud Network Controller (CNC1) is deployed in AWS account IA1 in the region R1, and you want to deploy a tenant in account TA1 in region R2. This tenant deployment (the account-region combination of TA1-R2) is now owned by IA1-R1 (CNC1). If another Cisco Cloud Network Controller (CNC2) attempts to manage the same tenant-region combination of TA1-R2 at some point in the future (for example, if CNC2 is deployed in AWS account IA2 in the region R3), this will not be allowed because the current owner for the deployment TA1-R2 is IA1-R1 (CNC1).

These restrictions are achieved using AWS Resource Groups. The following example provides several valid and invalid deployment combinations.

Cisco Cloud Network Controller	Tenant	Validity	Reason
IA1-R1 (CNC1)	TA1-R1	Valid	Tenant TA1-R1 is owned by IA1-R1 (CNC1)
IA1-R1 (CNC1)	TA1-R2	Valid	Tenant TA1-R2 is owned by IA1-R1 (CNC1)
IA1-R2 (CNC2)	TA1-R1	Invalid	Tenant TA1-R1 is already owned by IA1-R1 (CNC1)
IA1-R2 (CNC2)	TA1-R3	Valid	Tenant TA1-R3 is owned by IA1-R2 (CNC2)
IA2-R1 (CNC3)	TA1-R1	Invalid	Tenant TA1-R1 is already owned by IA1-R1 (CNC1)
IA2-R1 (CNC3)	TA1-R4	Valid	Tenant TA1-R4 is owned by IA2-R1 (CNC3)

Cisco Cloud Network Controller	Tenant	Validity	Reason
IA2-R1 (CNC3)	TA2-R4	Valid	Tenant TA2-R4 is owned by IA2-R1 (CNC3)

Deployment enforcement is done for the infra tenant as well as for user tenants. If CNC1 is deployed in the account IA1 in the region R1 and is also trying to manage the regions R2 and R3, another Cisco Cloud Network Controller (for example, CNC2) trying to manage the same account IA1 for regions R1, R2 and R3 would not be allowed.

The validation for the tenant-region ownership is done using AWS Resource Groups. For every tenant-region combination, a Resource Group is created using the syntax `CloudAPIC_TenantName_Region` (for example, the name `CNC_TA1_R2` would be created if a tenant is deployed in account TA1 in region R2). It would also have an ownership tag of `IA1_R1_TA1_R2`, if the Cisco Cloud Network Controller is deployed in account IA1 in region R1.

Following are examples of situations where there might be an AciOwnerTag mismatch, where existing tenant-region deployments would fail:

- If a Cisco Cloud Network Controller was initially installed in one account, was then torn down and the Cisco Cloud Network Controller was installed in a different account. In this case, all existing tenant-region deployments would fail if you try to manage the same tenant-region combinations again.
- If a Cisco Cloud Network Controller was initially installed in one region, was then torn down and the Cisco Cloud Network Controller is installed in a different region. In this case, all existing tenant-region deployments would fail.
- If another Cisco Cloud Network Controller is managing the same tenant-region.

In ownership mismatch cases, Cisco Cloud Network Controller does not perform a retry of the tenant-region setup again. To resolve ownership mismatch cases, if you are positive that no other Cisco Cloud Network Controller is managing the same tenant-region combination, log in to the tenant's AWS account and manually remove the affected Resource Group (for example, `CNC_123456789012_us-east-2`). Then either reload the Cisco Cloud Network Controller instance or delete the tenant from the Cisco Cloud Network Controller and add it again.



APPENDIX **D**

Locating CCR and Tenant Information

- [Locating CCR and Tenant Information, on page 107](#)

Locating CCR and Tenant Information

There are several pieces of CCR and tenant information that you need to enable connectivity between the Cisco Cloud Network Controller and the ISN devices. You should be able to get this information through Cisco Nexus Dashboard Orchestrator (**Sites > Configure Infra > Download IPN Device Config files only**). However, if you find that you need to manually gather the CCR and tenant information, the following sections provide instructions for locating this information.

- [Information for the CCR, on page 107](#)
- [Information for the Infra Tenant, on page 108](#)
- [Information for the User Tenant, on page 109](#)

Information for the CCR

Necessary AWS Information	Your Entry	How To Locate This Information in the AWS Site
Elastic IP address of the third network interface of a CCR		<ol style="list-style-type: none">1. Go into Instances in the EC2 Dashboard in the AWS Management Console.2. Choose a CCR instances (click the box next to a CCR instance).3. Scroll down until you see <code>Network interfaces</code> on the right side, then click the <code>eth2</code> link and locate the IP address shown in the <code>Public IP address</code> field.
Public IP address for a CCR		<ol style="list-style-type: none">1. Go into Instances in the EC2 Dashboard in the AWS Management Console.2. Locate a CCR instance.3. Copy the IP address shown in the <code>IPv4 Public IP</code> column for that CCR instance.

Necessary AWS Information	Your Entry	How To Locate This Information in the AWS Site
Preshared key for a CCR		<ol style="list-style-type: none"> Log into a CCR: <pre>ssh ip-address</pre> where <i>ip-address</i> is the public IP address for the CCR. Get the crypto keyring information: <pre>show running-config include pre-shared-key</pre> Output similar to the following appears, where the preshared key is highlighted: <pre>pre-shared-key address 192.0.2.15 key 12345678909876543211234567890</pre>
Peer tunnel IP address for the on-premises IPsec device to a CCR		<ol style="list-style-type: none"> Log into a CCR: <pre>ssh ip-address</pre> where <i>ip-address</i> is the public IP address for the CCR. Enter the following command: <pre>show ip interface brief include Tunnel2</pre> Output similar to the following appears: <pre>Tunnel2 30.29.1.1 YES NVRAM up down</pre> Take the IP address for this tunnel and increment the address by one to get the peer tunnel IP address for the on-premises IPsec device to the CCR. For example, if the IP address shown in the output is 30.29.1.1, then the peer tunnel IP address for the on-premises IPsec device to the CCR would be 30.29.1.2.

Information for the Infra Tenant

Necessary AWS Information	Your Entry	How To Locate This Information in the AWS Site
Cloud Account ID for infra tenant		Use the AWS account for the infra tenant as described in Deploying the Cisco Cloud Network Controller in AWS, on page 19 .
Cloud Access Key ID and Cloud Secret Access Key for infra tenant		<ol style="list-style-type: none"> Log into the Amazon Web Services account for the infra tenant. Go to IAM. In the left pane, select Users. Click the link for your admin account. On the Summary page, click the Security credentials tab. Click Create access key if you do not already have an Amazon Web Services access key ID. Locate the information from the Access key ID and Secret access key fields.

Information for the User Tenant

Necessary AWS Information	Your Entry	How To Locate This Information in the AWS Site
Cloud Account ID for Cisco Cloud Network Controller user tenant		Use the AWS account for the user tenant as described in Setting Up the AWS Account for the User Tenant, on page 24 .
Cloud Access Key ID and Cloud Secret Access Key for Cisco Cloud Network Controller user tenant		<ol style="list-style-type: none"> 1. Log into the Amazon Web Services account for the user account. 2. Go to IAM. 3. In the left pane, select Users. 4. Click the link for your Cisco Cloud Network Controller user tenant account. 5. On the Summary page, click the Security credentials tab. 6. Click Create access key if you do not already have an Amazon Web Services access key ID. 7. Locate the information from the Access key ID and Secret access key fields.

