



Managing Cisco Cloud Network Controller Through Multi-Site

- [About Cisco Cloud Network Controller and Multi-Site, on page 1](#)
- [Adding the Cisco Cloud Network Controller Site to Multi-Site, on page 2](#)
- [Configuring the Intersite Infrastructure, on page 2](#)
- [Enabling Connectivity Between the Cisco Cloud Network Controller and the ISN Devices, on page 3](#)
- [Configuring a Shared Tenant, on page 7](#)
- [Creating a Schema, on page 9](#)
- [Configuring an Application Profile and the EPGs, on page 10](#)
- [Creating and Associating a Bridge Domain with a VRF, on page 10](#)
- [Creating a Filter for a Contract, on page 11](#)
- [Creating a Contract, on page 11](#)
- [Adding Sites to the Schema, on page 12](#)
- [Configuring Instances in AWS, on page 12](#)
- [Adding an Endpoint Selector, on page 14](#)
- [Verifying the Multi-Site Configurations, on page 18](#)

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 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 2](#).

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 3](#).
- 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 7](#).

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 2](#), 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](#) 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](#) 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 9](#).

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 12](#).

-
- 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 7](#) 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 10](#).
- 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:

- a) Enter a name for the schema filter entry in the **Name** field on the **Add Entry** dialog.
- b) Optional. Enter a description for the filter in the **Description** field.
- c) 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.

- d) 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 11](#).

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 10](#).

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 10](#).

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 [12 in Deploying the Cisco Cloud Network Controller in AWS](#).

- **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 12](#) 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 12](#), 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 PROPERTIES** 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 16](#).
- 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 PROPERTIES** 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 18](#).

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:
- 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.
 - Click on Application Management → Tenants and verify that the tenants were configured correctly.
 - Click on Application Management → Application Profiles and verify that the application profiles were configured correctly.
 - Click on Application Management → EPGs and verify that the EPGs were configured correctly.
 - Click on Application Management → Contracts and verify that the contracts were configured correctly.
 - Click on Application Management → VRFs and verify that the VRFs were configured correctly.
 - Click on Application Management → Cloud Context Profiles and verify that the cloud context profiles were configured correctly.
 - Click on Cloud Resources → Regions and verify that the regions were configured correctly.
 - Click on Cloud Resources → VPCs and verify that the VPCs were configured correctly.
 - Click on Cloud Resources → Cloud Endpoints and verify that the cloud endpoints were configured correctly.
 - 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<br>Tu4<br>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
```

