



Configuring Cisco Cloud Network Controller Using the Setup Wizard

- [Configuring and Deploying Inter-Site Connectivity](#), on page 1
- [Gathering On-Premises Configuration Information](#), on page 2
- [Understanding Limitations for Number of Sites, Regions and CCRs](#), on page 2
- [Cloud Resources Naming](#), on page 3
- [Locating the Cisco Cloud Network Controller IP Address](#), on page 7
- [Configuring Cisco Cloud Network Controller Using the Setup Wizard](#), on page 9
- [Verifying the Cisco Cloud Network Controller Setup Wizard Configurations](#), on page 19

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 Cisco Cloud Routers deployed by Cisco Cloud Network Controller in Microsoft Azure. See [Components of Extending Cisco ACI Fabric to the Public Cloud](#) 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 documentation: 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 clouds 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 eight CCRs per managed region is supported, for a total of 32 CCRs per cloud site. For more information on increasing the number of CCRs, see the *Cisco Cloud Network Controller for Azure User Guide*.



Note The number of CCRs per managed region differs between AWS and Azure, with four CCRs per region supported for AWS and eight CCRs per region supported for Azure.

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

Cloud Resources Naming

You can create a global naming policy on the Cisco Cloud Network Controller, which allows you to define a custom cloud resources naming convention for all objects deployed from the Cisco Cloud Network Controller into the Azure cloud. You can define custom naming rules for all cloud resources during the first time setup wizard of the Cisco Cloud Network Controller, with the exception of the **Resource group** name used for the Cisco Cloud Network Controller ARM template deployment. The resource group name for the template is defined when you first deploy it and cannot be changed after. In addition to the global policy, you can also explicitly define the names of the cloud resources created from each Cisco Cloud Network Controller object using the REST API.

For Layer 4 to Layer 7 service deployments, you can provide custom names to cloud resources, such as Network Load Balancers, Application Load Balancers and Device Application Security Groups.



Note Keep in mind that even with custom naming policy, once a cloud resource is created, you will not be able to modify the name. If you want to change the name of an existing cloud resource, you would need to delete all configured cloud resources and recreate them. Cloud resources to be deleted include overlay-2 CIDR and subnets, Cisco Cloud Router deployed by Cisco Cloud Network Controller and therefore IPSec tunnels from the CCRs to every remote site.

Variables Available for Naming Rules

When creating your cloud resources naming policy, you can use the following variables to dynamically define the name of the cloud resource based on the Cisco Cloud Network Controller objects:

- `${tenant}` – the resource will include the name of the Tenant
- `${ctx}` – the resource will include the name of the VRF
- `${ctxprofile}` – the resources will include the cloud context profile, which is a VRF deployed in a given cloud region
- `${subnet}` – the resource will include the string `subnet` followed by the subnet IP address
- `${app}` – the resource will include the name of the application profile.
- `${epg}` – the resource will include the name of the EPG.
- `${contract}` – the resource will include the name of the contract
- `${region}` – the resource will include the name of the cloud region
- `${priority}` – the resource will include the name of the network security group (NSG) rule priority. This number is allocated automatically to ensure that each NSG rule name is unique
- `${serviceType}` – the resource will include an abbreviation of the service Type (only valid for private endpoint resources)
- `${resourceName}` – the resource will include the name of the target resource (only valid for private endpoint resources)
- `${device}` – the resource will include the name of the Layer 4 to Layer 7 device.
- `${interface}` – the resource will include the name of the Layer 4 to Layer 7 device interface.
- `${deviceInterfaceDn}` – the resource will include the DN of the Layer to Layer 7 device interface.

For private endpoints, the combination of the

`${app}-${svcepg}-${subnet}-${serviceType}-${resourceName}` makes the private endpoint name unique. Removing any of these variables might form a name of a private endpoint that already exists. This would result in a fault raised by the Cisco Cloud Network Controller. Also, the max length requirements vary from Azure service to service.

When you define a global naming policy using one or more of the above variables, Cisco Cloud Network Controller validates the string to ensure that all mandatory variables are present and no invalid string is specified.

There is a maximum name length limit in Azure. If the length of the name exceeds the length supported by the cloud provider, it rejects the config and Cisco Cloud Network Controller raises a fault that the resource creation failed. You can then check the fault for details and correct the naming rules. The maximum length limits at the time of Cisco Cloud Network Controller, Release 5.0(2) are listed below, for the latest up-to-date information and any changes to the length limit, consult the Azure documentation.

The following table provides a summary of which cloud resources support each of the naming variables above. Cells denoted with an asterisk (*) indicate variables that are mandatory for that type of cloud resource. Cells denoted with a plus sign (+) indicate that at least one of these variables is mandatory for that type of cloud resource; for example, for VNET resources you can provide `${ctx}`, or `${ctxprofile}`, or both.

Table 1: Supported Variables for Cloud Resources

Azure Resource	<code>\${tenant}</code>	<code>\${ctx}</code>	<code>\${ctxprofile}</code>	<code>\${subnet}</code>	<code>\${app}</code>	<code>\${epg}</code>	<code>\${contract}</code>	<code>\${region}</code>	<code>\${priority}</code>
Resource Group Max Length: 90	Yes*	Yes*						Yes*	
Virtual Network (VNET) Max Length: 64	Yes	Yes+	Yes+					Yes	
Subnet Max Length: 80	Yes	Yes	Yes	Yes*				Yes	
Application Security Group (ASG) Max Length: 80	Yes				Yes*	Yes*		Yes	
Network Security Group (NSG) Max Length: 80	Yes				Yes*	Yes*		Yes	

Azure Resource	\${tenant}	\${ctx}	\${ctxprofile}	\${subnet}	\${app}	\${epg}	\${contract}	\${region}	\${priority}
Network Security Group Rule Max Length: 80	Yes						Yes		Yes* (auto)

Table 2: Supported Variables for Cloud Resources (Layer 4 to Layer 7 device services)

Azure Resource	\${tenant}	\${region}	\${ctxprofile}	\${device}	\${interface}	\${deviceInterfaceID}
Internal Network Load Balancer Max Length: 80	Yes	Yes	Yes	Yes*		
Internet-facing Network Load Balancer Max Length: 80	Yes	Yes	Yes	Yes*		
Internal Application Load Balancer Max Length: 80	Yes	Yes	Yes	Yes*		
Internet-facing Application Load Balancer Max Length: 80	Yes	Yes	Yes	Yes*		
Device ASG Max Length: 80	Yes	Yes		Yes*	Yes*	Yes*

Naming Rules Guidelines and Limitations

When configuring custom rules for naming cloud resources, the following restrictions apply:

- You define global naming policy during the Cisco Cloud Network Controller 's first time setup using two sets of naming rules:
 - Hub Resource Naming Rules** define names for the Hub Resource Group, Hub VNET, Overlay-1 CIDR, Overlay-2 CIDR subnet in the Infra Tenant, as well as the subnet prefixes for subnets that are created automatically by the system in the Infra tenant.

- **Cloud Resource Naming Rules** define the names of the Network Security Group (NSG), Application Security Group (ASG), Network Load Balancer, Application Load Balancer, Device Application Security Group, and subnets you create in the Infra Tenant, as well as the names of all resources (Resource Groups, Virtual Networks, Subnets, NSG, ASG, Network Load Balancer, Application Load Balancer) in user Tenants.

After you define the naming rules, you will be required to review and confirm them. Keep in mind that you must confirm the naming rules before any cloud resources are deployed.

- Once a cloud resource is created, its name cannot be changed and the naming policy cannot be updated in the GUI. If you upgrade your Cisco Cloud Network Controller to Release 5.0(2) with some resources already deployed in Azure, you will also not be able to change the global custom naming rules.

If you want to change the names of the existing cloud resources or the policy, you would need to delete the deployed resources before being able to update the global naming policy in the GUI.

In these cases you can use the REST API to explicitly assign custom names to any new resources you create.

- When updating cloud resources naming via REST API, we recommend you do not import configuration at the same time.

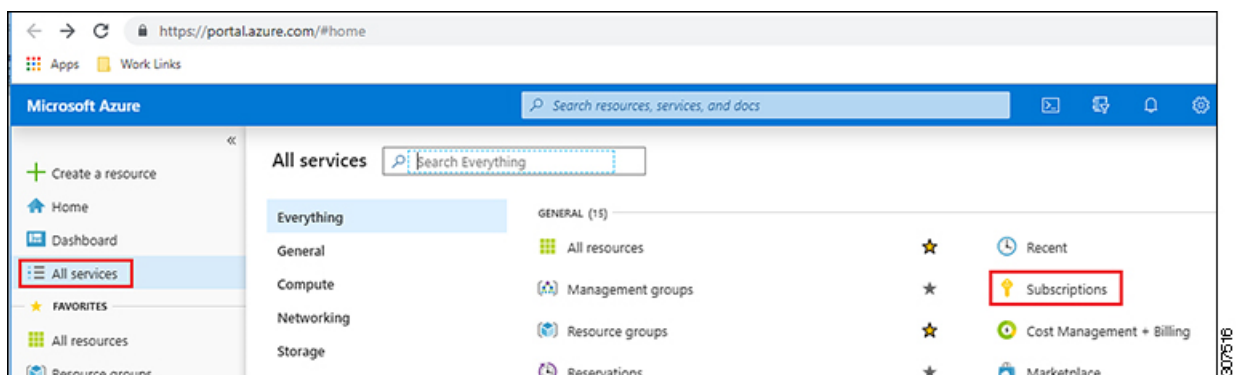
We recommend you define any naming rules first. Then any tenant configuration.

We recommend that you do not change the naming policy after the tenant configuration is deployed.

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 Azure site.

- Step 1** From the main Azure management portal page, click the **All services** link in the left nav bar, then click the **Subscriptions** link.



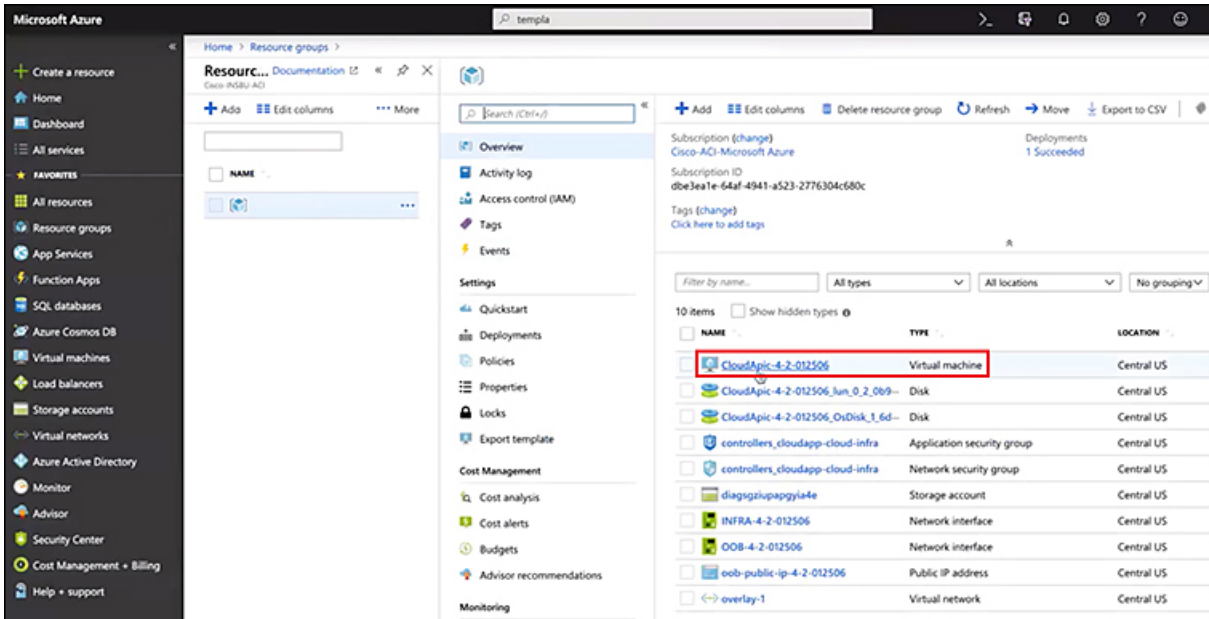
- Step 2** In the **Subscriptions** page in the Azure management portal, click the subscription account that you just created. The overview information for that subscription is displayed.

- Step 3** From the overview page for that subscription, locate the **Resource groups** link in the left nav bar and click that link. The resource groups for that subscription is displayed.

Locating the Cisco Cloud Network Controller IP Address


Step 4 Choose the resource group that you chose or created in [Deploying the Cisco Cloud Network Controller in Azure](#). The overview information for that resource group is displayed.

Step 5 In the overview page for the resource group, locate your Cisco Cloud Network Controller VM instance (shown as **Virtual machine** under the TYPE column), and click the link for that VM instance.



The overview information for the Cisco Cloud Network Controller VM instance is displayed.

Step 6 Locate the entry in the **Public IP address** field in this page and copy that IP address entry.



This is the Cisco Cloud Network Controller IP address that you will use to log into the Cisco Cloud Network Controller

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 Azure constructs and the necessary Catalyst 8000Vs.

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](#) before proceeding with the tasks in this section.
- You have successfully completed the procedures that are provided in [Deploying the Cisco Cloud Network Controller in Azure](#).

Step 1 Locate the IP address for your Cisco Cloud Network Controller.

See [Locating the Cisco Cloud Network Controller IP Address, on page 7](#) 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 to log into the Cisco Cloud Network Controller.
- **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 and NTP Servers**
- **Region Management**
- **Advanced Settings**

- **Smart Licensing**

Step 6 In the **DNS and NTP 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 10](#) if you want to configure an NTP server and you do not want to configure a DNS server.
- If you want to use a specific DNS server, under the **DNS Servers** area, click **+Add DNS Provider**.
 - Enter the IP address for the DNS servers and, if necessary, check the box next to Preferred DNS Provider.
 - Click the check mark next to the DNS server, and repeat for any additional DNS servers that you want to add.
 - Under the **NTP Servers** area, click **+Add Providers**.
 - Enter the IP address for the NTP servers and, if necessary, check the box next to Preferred NTP Provider.
 - 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 Verify that the **Virtual Network Peering** in the **Connectivity for Internal Network** area is automatically enabled.

VNet peering at the global level is set in the **Connectivity for Internal Network** area, which enables VNet peering at the Cisco Cloud Network Controller level, deploying NLBs in all the regions with a Catalyst 8000V. For release 5.1(2) and later, VNet peering at the global level is enabled by default and cannot be disabled. See [Configuring VNet Peering for Cloud APIC for Azure](#) for more information.

Step 11 In the **Regions to Manage** area, verify that the Cisco Cloud Network Controller home region is selected.

The region that you selected when you were configuring your cloud site 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.

Note Because Azure VNet peering is enabled automatically in [#unique_23 unique_23_Connect_42_step_ktp_hxc_knb](#), you must also check the box in the **Cloud Routers** column for the Cisco Cloud Network Controller home region, if it is not checked already.

Step 12 Select additional regions if you want the Cisco Cloud Network Controller to manage additional regions, and to possibly deploy Catalyst 8000Vs to have inter-VNet communication and Hybrid-Cloud, Hybrid Multi-Cloud, or Multi-Cloud connectivity on those other regions.

The Catalyst 8000V can manage up to 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 manages two regions, those two regions are considered a single site by Cisco ACI.

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-VNet 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 2](#) for more information.

Step 13 When you have selected all the appropriate regions, click **Next** at the bottom of the page.

The **General Connectivity** page appears.

Step 14 Enter the following information on the **General Connectivity** page.

a) Under the **General** area, in the **Subnet Pools for Cloud Routers** field, click **Add Subnet Pool for Cloud Routers** if you want to add additional subnets for Catalyst 8000Vs.

The first subnet pool is automatically populated (shown as `System Internal`). 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. Subnet pools added in this field must be a valid IPv4 subnet with mask /24.

Add additional subnets for Catalyst 8000Vs in this step in these situations:

- If you have a Catalyst 8000V deployed in the Cisco Cloud Network Controller home region, add one additional subnet pool in addition to the `System Internal` subnet pool that is automatically generated
- If you selected additional regions to be managed by Cisco Cloud Network Controller in the previous page:
 - Add *one* additional subnet pool for every managed region with 2-4 Catalyst 8000Vs per managed region (if you enter **2**, **3**, or **4** in the **Number of Routers Per Region** field in [14.f, on page 13](#))
 - Add *two* additional subnet pools for every managed region with five or more Catalyst 8000Vs per managed region (if you enter between **5** and **8** in the **Number of Routers Per Region** field in [14.f, on page 13](#))

For example:

- Assume you have only the Cisco Cloud Network Controller home region selected in the previous page, and you have a Catalyst 8000V deployed in the Cisco Cloud Network Controller home region. You will need two subnet pools (the automatically-populated **System Internal** subnet pool and one additional subnet pool that is created by you).
- Next, assume you selected two additional regions for Cisco Cloud Network Controller to manage in the previous page, and you have Catalyst 8000Vs deployed in both additional regions. In addition, assume you select between 2-4 Catalyst 8000Vs to be deployed in each managed region in the **Number of Routers Per Region** field ([14.f, on page 13](#)). In this case, you would need to add two additional subnet pools (one subnet pool for every region with Catalyst 8000Vs that you selected in the previous page), for a total of four subnet pools (one automatically populated as **System Internal** and three additional ones that are created by you).
- Finally, assume that you decide to increase the number of Catalyst 8000Vs in each managed region to eight at a later date, where you return to this page and you change the value to **8** in the **Number of Routers Per Region** field ([14.f, on page 13](#)). Because you have three regions selected in the previous screen (the Cisco Cloud Network Controller home region and two additional regions that you selected to have the Cisco Cloud Network Controller to manage), and you increased the number of Catalyst 8000Vs per managed region above four, you would need to add three additional subnet pools again, one for every managed region that has more than four Catalyst 8000Vs, for a total of seven subnet pools:
 - One automatically populated as **System Internal**

- Two for the Catalyst 8000Vs in the home region (one subnet pool created by you previously, and the other one created by you here when you increased the number of Catalyst 8000Vs to 8 per managed region)
- Four for the Catalyst 8000Vs in the two additional regions that you selected to have managed by the Cisco Cloud Network Controller (two subnet pools created by you previously, and the other two created by you here when you increased the number of Catalyst 8000Vs to 8 per managed region)

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

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

- c) 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.

- d) Under the **Catalyst 8000Vs** area, in the **BGP Autonomous System Number for C8kVs** field, enter the BGP autonomous system number (ASN) that is unique to this site.

The BGP autonomous system number can be in the range of 1 - 65534.

Note the following Microsoft Azure ASN restrictions:

- Do not use 64518 as the autonomous system number in this field.
- Do not use 32-bit ASNs. Azure VPN Gateways support 16-Bit ASNs at this time.
- The following ASNs are reserved by Azure for both internal and external peerings:
 - Public ASNs: 8074, 8075, 12076
 - Private ASNs: 65515, 65517, 65518, 65519, 65520

You cannot specify these ASNs for your on-premises VPN devices when connecting to Azure VPN gateways.

- The following ASNs are [reserved by IANA](#) and cannot be configured on your Azure VPN Gateway: 23456, 64496-64511, 65535-65551 and 429496729

- e) In the **Assign Public IP to C8kV 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 C8kV Interface** option determines whether public IP addresses will also be assigned to the Catalyst 8000V interfaces or not.

The Catalyst 8000V interface IP addresses are used for the following purposes:

- Allows you to configure the Catalyst 8000V through the Management Interface in the Cisco Cloud Network Controller GUI
- Allows you to cross-program the interfaces across sites for multi-cloud and hybrid cloud connectivity through the Cisco Nexus Dashboard Orchestrator

- For the Catalyst 8000Vs for both control plane and data plane traffic

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 a public IP is not assigned to a Catalyst 8000V, only the private IP is displayed.

- f) In the **Number of Routers Per Region** field, choose the number of C8kVs that will be used in each region.

See [Understanding Limitations for Number of Sites, Regions and CCRs, on page 2](#) for more information on any limitations on the number of Catalyst 8000Vs per region.

Note If you change the value in this field to increase or decrease the number of Catalyst 8000Vs 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 Catalyst 8000Vs, wait long enough for those Catalyst 8000Vs to get deleted before changing the value in this field again.
- If you are increasing the number of Catalyst 8000Vs, wait long enough for those Catalyst 8000Vs to get deployed before changing the value in this field again.

- g) In the **Username**, enter the username for the Cisco Cloud Router.

Note Do not use `admin` as a username for the Cisco Cloud Router when connecting to an Azure cloud site.

- h) In the **Password** field, enter the password for the Cisco Cloud Router.

Enter the password again in the **Confirm Password** field.

- i) 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 Azure 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 Cisco Cloud Router.

The Cisco Catalyst 8000V supports tier-based (T0/T1/T2/T3) throughput options. The following table lists what Azure VM sizes are needed for different router throughput settings for the Cisco Catalyst 8000V:

Catalyst 8000V Throughput	Azure VM Size
T0 (up to 15M throughput)	DS3_v2
T1 (up to 100M throughput)	DS3_v2
T2 (up to 1G throughput)	DS3_v2
T3 (up to 10G throughput)	F16s_v2

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

Changing the value in this field changes the size of the Catalyst 8000V instance that is deployed. Choosing a higher value for the throughput results in a larger VM being deployed.

Note the following:

- The licensing of the Catalyst 8000V 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 Azure Public Cloud](#) for more information.
- Cloud routers should be undeployed from all regions before changing the router throughput or login credentials.

If you wish to change this value at some point in the future, you must delete the Catalyst 8000V, 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.

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 to all cloud router interfaces, including data Gigabit Ethernet interfaces, IPSec tunnel interfaces of cloud routers, and VPN tunnel interfaces toward cloud, on-premises, 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 Cisco Cloud Router.

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. See [Cisco Cloud Network Controller Licensing](#) for more information.

Note If you assigned private IP addresses to the Catalyst 8000Vs in [14.e, on page 12](#), the only supported option is **Direct connect to Cisco Smart Software Manager (CSSM)** when registering smart licensing for Catalyst 8000Vs with private IP addresses. You must provide reachability to the CSSM through express route in this case.

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

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

The Cisco Cloud Network Controller supports a range of VM types. The table below shows the various instances of the VM types available along with their capacity.

VmName on Azure	Memory	vCPUs	NetworkBw
DS3V2	14GiB	4	Up to 3 Gigabit
DS4V2	28GiB	8	Up to 6 Gigabit
F16SV2	32GiB	16	Up to 12.5 Gigabit
F32SV2	64GiB	32	Up to 16 Gigabit

Note If you wish to change this value at some point in the future, you must delete the Catalyst 8000V, then repeat the processes in this chapter again and select the new value that you would like in the same VM field.

Changing the value in this field changes the other factors of the Catalyst 8000V 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 15 Click the appropriate button, depending on whether you are configuring inter-site connectivity or not.

- If you are not configuring inter-site connectivity (if you did not select **Inter-Site Connectivity** when you were selecting regions to manage in the **Region Management** page), click **Save and Continue**. The **Let's Configure the Basics** page appears again. Skip to [Configuring Cisco Cloud Network Controller Using the Setup Wizard, on page 9](#).
- If you are configuring inter-site connectivity (if you selected **Inter-Site Connectivity** when you were selecting regions to manage in the **Region Management** page), click **Next** at the bottom of the page. The **Inter-Site Connectivity** page appears.

Step 16 Enter the following information in the **Inter-Site Connectivity** page:

- **IPSec Tunnels to Inter-Site Routers:** This field is necessary only for on-premises connectivity to cloud sites. There is no need to enter information in this field if you don't have an on-premises site.

In this area, click the + button next to the **Add Public IP of IPsec Tunnel Peer** field.

- Enter the peer IP address for the IPsec tunnel termination to the on-premises device.
- Click the check mark to add this peer IP address.
- **OSPF Area for Inter-Site Connectivity:** Enter the underlay OSPF area ID that will be used with on-premises ISN peering (for example, 0.0.0.1)

- Under the **External Subnets for Inter-Site Connectivity** heading, click the + button next to the **+Add External Subnet** field.
 - Enter the subnet tunnel endpoint pool (the cloud TEP) that will be used in Azure. It must be a valid IPv4 subnet with a mask between /16 and /22 (for example, 30.29.0.0/16). This subnet will be used to address the IPsec tunnel interfaces and loopbacks of the Cloud Routers used for on-premises connectivity, and cannot overlap with other on-premises TEP pools.
 - Click the check mark after you have entered in the appropriate subnet pools.

Step 17 When you have configured all the connectivity options, click **Next** at the bottom of the page.

The **Cloud Resource Naming Rules** page appears.

Step 18 Choose **Cloud Resource Naming** mode.

You can create a global naming policy on the Cisco Cloud Network Controller, which allows you to define a custom cloud resources naming convention for all objects deployed from the Cisco Cloud Network Controller into the Azure cloud. Additional details about naming rules, available object name variables, guidelines, and limitations are available in the earlier [Cloud Resources Naming, on page 3](#) section of this chapter.

You can choose one of the following:

- **Default**, in which case the cloud resources created by the Cisco Cloud Network Controller in Azure will be assigned names that are derived from the names of the ACI objects. For example, resource groups' names will be based on the Tenant, VRF, and region: `CAPIC_<tenant>_<vrf>_<region>`.
- **Custom**, in which case you can define your own rules for how each of the cloud resources is named.

When you select the custom naming, an **Edit** icon appears next to each cloud resource. You can click the edit icon to define the naming convention for one or more of the displayed resources.

The variables that are available for this type of resource are listed under the naming rule text box. The variables are divided into the **Required Keyword** and **Optional Keywords**, you must include all the required keywords for the rule you are updating. For example, when defining the naming rule for Azure's Resource Groups, you must include the Tenant name, VRF name, and the Region keywords.

Step 19 Confirm you have reviewed and accept the global resource naming policy.

Once a cloud resource is created, its name cannot be changed. As such, you must review and accept the global naming policy you have defined in the previous step before any cloud resources can be deployed. When ready, enable the **Deploy cloud resources based on these naming rules** checkbox.

Note that you can leave the checkbox unchecked and choose to proceed, in which case any changes you have made will be saved but no configuration will be deployed. You would need to come back to this screen to accept the naming policy to deploy.

Step 20 When you have entered all the necessary information on this page, click **Save and Continue** at the bottom of the page.

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

Step 21 In the **Advanced Settings** row, click **Edit Configuration**.

The **Advanced Settings** page appears.

Step 22 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 23 Click **Save and Continue**.

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

Step 24 In the **Smart Licensing** row, click **Register**.

The **Smart Licensing** page appears.

Step 25 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 26 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 27 Verify the information on the **Summary** page, then click **Finish**.

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.

Step 28 Verify that the Catalyst 8000Vs were successfully deployed.

- a) From the main Azure management portal page, click the **All services** link in the left nav bar, then click the **Subscriptions** link.



- b) In the **Subscriptions** page in the Azure management portal, click the subscription account that you created. The overview information for that subscription is displayed.
- c) From the overview page for that subscription, locate the **Resource groups** link in the left nav bar and click that link. The resource groups for that subscription is displayed.
- d) Choose the resource group that you chose or created in the **Custom deployment** page in [Deploying the Cisco Cloud Network Controller in Azure](#). The overview information for that resource group is displayed.
- e) In the overview page for the resource group, locate your Catalyst 8000V VM instance (shown as **Virtual machine** under the TYPE column), and click the link for that VM instance.

The Catalyst 8000V VM instance will have a name with a `ct_routerp_region_x_0` format, where:

- *region* is the managed region (for example, `westus`, `westus2`, `centralus`, or `eastus`)
- *x* is the Catalyst 8000V count, starting from zero

For example: `ct_routerp_centralus_0_0` or `ct_routerp_centralus_1_0`

The overview information for the Catalyst 8000V VM instance is displayed.

- f) Locate the **Status** field at the top left area in the page.
- If you see the text **Creating** in the **Status** field, then the Catalyst 8000Vs are not fully deployed yet.

- If you see the text **Running** in the **Status** field, then the Catalyst 8000Vs are fully deployed.

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 (if you selected the **Inter-Site Connectivity** option in the **Region Management** page), go to [Managing Cisco Cloud Network Controller Through Multi-Site](#).
- 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 (if you selected only the **Cloud Routers** option in the **Region Management** page), you will not need to use the Multi-Site for additional configurations. However, you will have additional configurations that you must perform in the Cisco Cloud Network Controller GUI in this case.

You also need to create a tenant using the Cisco Cloud Network Controller GUI using the instructions in [Creating a Tenant Using the Cisco Cloud Network Controller GUI](#).

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](#) and [Configuring Cisco Cloud Network Controller Components](#) 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 **Inter-Site 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](#).