



## Cloud onRamp for IaaS

---

Cloud onRamp for IaaS extends the fabric of the Cisco SD-WAN overlay network into public clouds, allowing branches with vEdge routers to connect directly to public-cloud application providers. By eliminating the need for a physical data center, Cloud onRamp for IaaS improves the performance of IaaS applications.

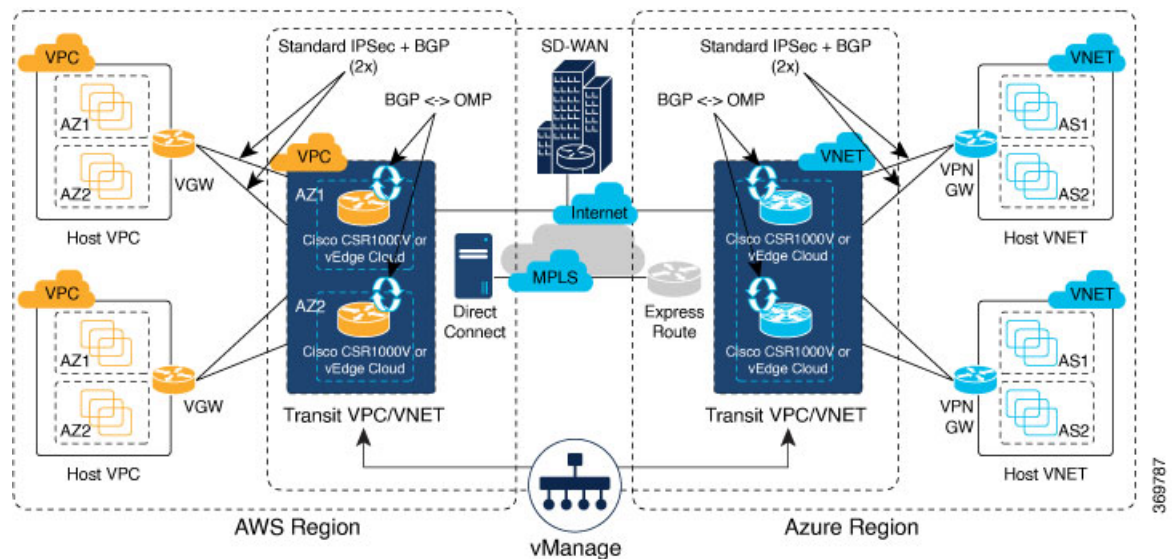
The connection between the overlay network and a public-cloud application is provided by two or four pairs of redundant vEdge Cloud routers for AWS, which act together as a transit between the overlay network and the application. By using redundant routers to form the transit offers path resiliency to the public cloud. In addition, having redundant routers improves the availability of public-cloud applications. Together, the two routers can remediate in the event of link degradation. You create these routers as part of the Cloud onRamp workflow.

Cloud onRamp for IaaS discovers any already existing private cloud instances in geographical cloud regions and allows you to select which of them to make available for the overlay network. In such a scenario, Cloud onRamp for IaaS allows simple integration between legacy public-cloud connections and the Cisco SD-WAN overlay network.

You configure and manage Cloud onRamp for IaaS through the vManage NMS server. A configuration wizard in the vManage NMS automates the bring-up of the transit to a your public cloud account and automates the connections between public-cloud applications and the users of those applications at branches in the overlay network.

A configuration wizard in vManage automates the bring-up of the transit to a your public cloud account and automates the connections between public-cloud applications and the users of those applications at branches in the overlay network.

The Cloud onRamp for IaaS works in conjunction with AWS virtual private clouds (VPCs) and Azure virtual networks (VNETs). The following image provides a high level overview of multi-cloud onRamp for IaaS.



### Supported Routers

Cloud onRamp for IaaS is supported on Cisco Cloud vEdge and Cisco Cloud Services Routers (CSRs). In this topic, supported routers are referred to collectively as *cloud routers*.

- [Prerequisites to Configure Cloud onRamp for IaaS, on page 2](#)
- [Configure Cloud onRamp for IaaS for Amazon Web Services, on page 6](#)
- [Configure Cloud onRamp for IaaS for Azure, on page 13](#)
- [Troubleshoot Cloud onRamp for IaaS, on page 18](#)

## Prerequisites to Configure Cloud onRamp for IaaS

Before you configure Cloud onRamp for IaaS, ensure that you provision the vManage NMS, AWS, and Azure.

### 1. Provision vManage

Before you can configure Cloud onRamp for IaaS, you must properly provision vManage NMS.

- a. Ensure that your vManage server has access to the internet and that it has a DNS server configured so that it can reach AWS. To configure a DNS server, in the vManage VPN feature configuration template, enter the IP address of a DNS server, and then reattach the configuration template to the vManage server.
- b. Ensure that two cloud routers that are to be used to bring up the Cloud OnRamp for IaaS have been added to the vManage NMS and have been attached to the appropriate configuration template. (These two routers are deployed in AWS in their own VPC, and together they form the transit VPC, which is the bridge between the overlay network and AWS cloud applications.) Ensure that the configuration for these routers includes the following:
  - Hostname
  - IP address of vBond orchestrator
  - Site ID

- Organization name
  - Tunnel interface configuration on the eth1 interface
- c. Ensure that the vManage NMS is synchronized with the current time. To check the current time, click the Help (?) icon in the top bar of any vManage screen. The Timestamp field shows the current time. If the time is not correct, configure the vManage server's time to point to an NTP time server, such as the Google NTP server. To do this, in the vManage NTP feature configuration template, enter the hostname of an NTP server, and then reattach the configuration template to the vManage server. The Google NTP servers are time.google.com, time2.google.com, time3.google.com, and time4.google.com

## 2. Provision Amazon Web Services (AWS)

Before you can configure Cloud OnRamp for IaaS, ensure that you provision AWS properly.

- a. Ensure that you have subscribed to the Viptela marketplace Amazon machine images (AMIs) and the Cisco CSR AMIs in your AWS account. See [Subscribe to Cisco SD-WAN AMIs](#).
- b. Ensure that at least one user who has administrative privileges has the AWS API keys for your AWS account. For Cloud OnRamp for IaaS, these keys are used to authenticate the vManage server with AWS and to bring up the VPC and Elastic Compute Cloud (EC2) instances.
- c. Check the AWS limits associated with your account (in the Trusted Advisor section of AWS) to ensure that the following resources can be created in your account:
  - 1 VPC, which is required for creating the transit VPC
  - 6 Elastic IP addresses associated with each pair of transit vEdge router




---

**Note** vEdge Cloud routers support C3 and C4 compute-intensive families.

---

- 1 AWS virtual transit (VGW) for each host VPC
- 4 VPN connections for mapping each host VPC




---

**Note** Cisco XE SD-WAN devices use VRFs in place of VPNs. When you complete the VPN configuration, the system automatically maps the VPN configurations to VRF configurations.

---

## 3. Subscribe to Cisco SD-WAN AMIs

To use the Cloud OnRamp for IaaS and other Cisco SD-WAN services, you must subscribe to the Amazon Machine Image (AMI) for your router in AWS. When you subscribe, you can complete the following tasks:

- Launch a cloud router AMI instance
- Generate a key pair to use for the instance
- Use the key pair to subscribe to the cloud router instance.

You subscribe to the vEdge Cloud router AMI only once, when you first create a Cisco SD-WAN Viptela AMI instance.

To create a new AMI subscription, generate and upload a key pair:

- a. In AWS, search to locate a cloud router AMI for your devices.
- b. Select and launch an EC2 instance with the AMI instance. For more information, see [Create Cisco IOS XE SD-WAN Cloud VM Instance on AWS](#).
- c. Generate a key pair. For full instructions, see [Set Up the Cisco IOS XE SD-WAN Cloud VM Instance](#).
- d. Click **Download Key Pair**. The key pair then downloads to your local computer as a .pem file.
- e. Click **Launch Instance**. A failure message displays, because you now need to upload the key pair to complete the subscription process.
- f. To upload the key pair, in [AWS Marketplace](#), search for your router AMI.
- g. Click **Continue**.
- h. Click **Key Pair** to bring up a vEdge router instance. In the option to enter the key pair, upload the .pem file from your local computer. This is the file that you had generated in Step c when creating a new AMI subscription.

#### 4. Provision Azure

Before you can configure Cloud OnRamp for IaaS, you must properly provision Azure.

- a. Ensure that you have accepted the terms and conditions for the Cisco Cloud vEdge Router in the Azure Marketplace.

##### Accept the Azure Terms of Service

To use a Cisco cloud router as part of the Cloud onRamp workflow, you must accept marketplace terms for using a virtual machine (VM). You can do accept the Azure Terms of Service in one of the following ways:

- Spin up the cloud router on the portal manually, and accept the terms as part of the final page of the bringup wizard.
  - In the Azure APIs or Powershell/Cloud Shell, use the Set-AzureRmMarketplaceTerms command.
- b. Ensure that you create an App Registration in Azure and retrieve the credentials for your Azure account. For Cloud OnRamp for IaaS, these credentials are used to authenticate the vManage server with Azure and bring up the VNet and the Virtual Machine instances.

##### Create and Retrieve Azure Credentials

To create and retrieve Azure credentials, you must create an App Registration in Azure with Contributor privileges:

1. Launch the Microsoft Azure portal.
2. Create an application ID:
  - a. In the left pane of the Azure portal, click **Azure Active Directory**.
  - b. In the sub-menu, click **App registrations**.
  - c. Click **New application registration**. The system displays the Create screen.



- c. Check the Azure limits associated with your account (by going to your subscription in the portal and checking Usage + Quotas) to ensure that the following resources can be created in your account:
  1. 1 VNet, which is required for creating the transit VNet
  2. 1 Availability set, required for Virtual Machine distribution in the transit VNet
  3. 6 Static Public IP addresses associated with the transit cloud routers
  4. 1 Azure Virtual Network Gateway and 2 Static Public IP Addresses for each host VNet
  5. 4 VPN connections for mapping each host VNet




---

**Note** Cisco XE SD-WAN devices use VRFs in place of VPNs. When you complete the VPN configurations, the system automatically maps the VPN configurations to VRF configurations.

---

- d. F-Series Azure VMs (F4 and F8) are supported on the cloud routers.

## Configure Cloud onRamp for IaaS for Amazon Web Services

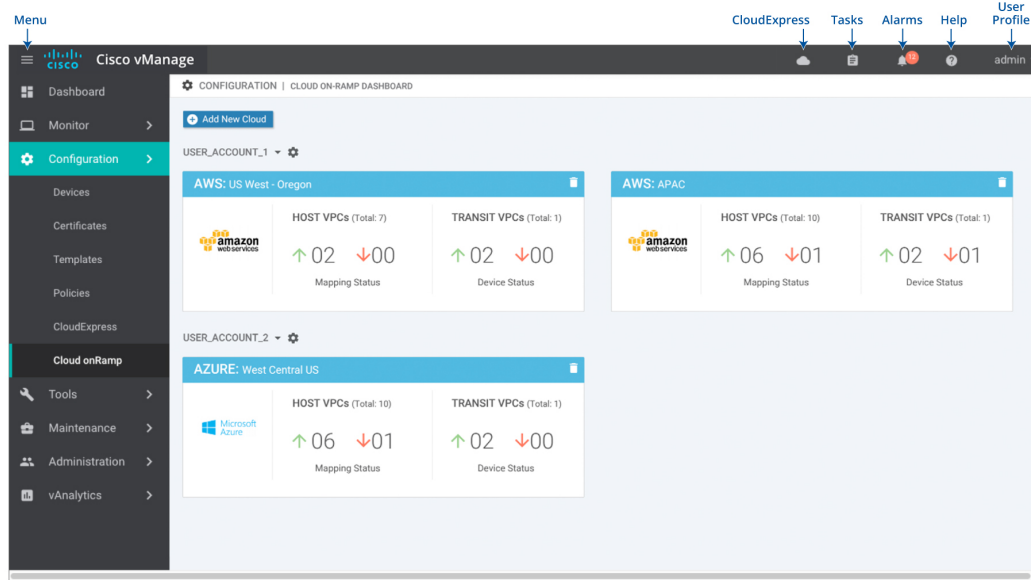
### Before you begin

A series of considerations are essential to configure Cloud onRamp for IaaS for AWS.

- Transit VPCs provide the connection between the Cisco overlay network and the cloud-based applications running on host VPCs. Each transit VPC consists of up to four pairs of cloud routers that reside in their own VPC. Multiple routers are used to provide redundancy for the connection between the overlay network and cloud-based applications. On each of these two cloud routers, the transport VPN (VPN 0) connects to a branch router, and the service-side VPNs (any VPN except for VPN 0 and VPN 512) connect to applications and application providers in the public cloud.
- Cloud onRamp supports auto-scale for AWS. To use auto-scale, ensure that you associate two to four pairs of cloud routers to a transit VPC. Each of the devices that are associated with the transit VPC for auto-scale should have a device template attached to it.
- Host VPCs are virtual private clouds in which your cloud-based applications reside. When a transit VPC connects to an application or application provider, it is simply connecting to a host VPC.
- All host VPCs can belong to the same account, or each host VPC can belong to a different account. A host that belongs one account can be mapped to a transit VPC that belongs to a completely different account. You configure cloud instances by using a configuration wizard.

---

**Step 1** In Cisco vManage, select **Configuration > Cloud onRamp for IaaS**.



368703

**Step 2** Click **Add New Cloud Instance**.

**Step 3** In the Add Cloud Instance – log in to a Cloud Server popup:

- a) In the **Cloud** drop-down, select the **Amazon Web Services** radio button.
- b) Click **IAM Role** or **Key** to log in to the cloud server. It is recommended that you use IAM Role.
- c) If you select **IAM Role**:
  1. In the **Role ARN** field, enter the role ARN of the IAM role.
  2. In the **External ID** field, enter external ID created for the role ARN. It is recommended that the external ID include 10 to 20 characters in random order. To authenticate to the vManage NMS using an IAM role, vManage NMS must be hosted by Cisco on AWS and have the following attributes:
    - Trusts the AWS account, 200235630647, that hosts the vManage NMS.
    - Have all permissions for EC2 and VPC resources.
    - A default timeout of at least one hour.

If Cisco vManage is not hosted by Cisco on AWS, assign an IAM role with permissions to AssumeRole to the vManage server running the Cloud OnRamp process. Refer to the AWS documentation for details.

- d) If you select **Key**:
  1. In the **API Key** field, enter your Amazon API key.
  2. In the **Secret Key** field, enter the password associated with the API key.

**Step 4** Click **Login** to log in to the cloud server.



The cloud instance configuration wizard opens. This wizard consists of three screens that you use to select a region and discover host VPCs, add transit VPC, and map host VPCs to transit VPCs. A graphic on the right side of each wizard screen illustrates the steps in the cloud instance configuration process. The steps that are not yet completed are shown in

light gray. The current step is highlighted within a blue box. Completed steps are indicated with a green checkmark and are shown in light orange.

**Step 5** Select a region:

- a) In the **Choose Region** drop-down, choose a geographical region.
- b) Click **Save and Finish** to create a transit VPC or optionally click **Proceed to Discovery and Mapping** to continue with the wizard.

**Step 6** Add a transit VPC:

- a) In the **Transit VPC Name** field, type a name for the transit VPC.  
The name can be up to 128 characters and can contain only uppercase and lowercase letters, the digits 0 through 9, hyphens (-), and underscores (\_). It cannot contain spaces or any other characters.
- b) Under **Device Information**, enter information about the transit VPC:
  1. In the **WAN Edge Version** drop-down, select the software version of the cloud router to run on the transit VPC.
  2. In the **Size of Transit WAN Edge** drop-down, choose how much memory and how many CPUs to use for each of the cloud routers that run on the transit VPC.
  3. In the **Max. Host VPCs per Device Pair** field, select the maximum number of host VPCs that can be mapped to each device pair for the transit VPC. Valid values are 1 through 32.
  4. In the **Device Pair 1#** fields, select the serial numbers of each device in the pair. To remove a device serial number, click the **X** that appears in a field.  
The devices that appear in this field have been associated with a configuration template and support the WAN Edge Version that you selected.
  5. To add additional device pairs, click .  
To remove a device pair, click .  
A transit VPC can be associated with one to four device pairs. To enable the Cloud onRamp auto-scale feature for AWS, you must associate at least two device pairs with the transit VPC.
  6. Click **Save and Finish** to complete the transit VPC configuration or optionally click **Proceed to Discovery and Mapping** to continue with the wizard.
  7. Click **Advanced** if you wish to enter more specific configuration options:
    - a. In the **Transit VPC CIDR** field, enter a custom CIDR that has a network mask in the range of 16 to 25. If you choose to leave this field empty, the Transit VPC is created with a default CIDR of 10.0.0.0/16.
    - b. In the **SSH PEM Key** drop-down, select a PEM key pair to log in to an instance. Note that the key pairs are region-specific. Refer to the AWS documentation for instructions on creating key pairs.
  8. Click **Save and Finish** to complete the transit VPC configuration, or optionally click **Proceed to Discovery and Mapping** to continue with the wizard.
  9. Select hosts to discover:
    - a. In the **Select an account to discover** field, select a host to map to this transit VPC.
    - b. Click **Discover Host VPCs**.
    - c. In the table that displays, choose one or more hosts to map to this transit VPC.



You can use the search field and options to display only host VPCs that mention specific search criteria.


You can click the **Refresh** icon to update the table with current information.


You can click the **Show Table Columns** icon to specify which columns display in the table.

d. Click **Next**.

1. In the **WAN Edge Version** drop-down, select the software version of the cloud router to run on the transit VPC.
2. In the **Size of Transit WAN Edge** drop-down, choose how much memory and how many CPUs to use for each of the cloud routers that run on the transit VPC.
3. In the **Max. Host VPCs per Device Pair** field, select the maximum number of host VPCs that can be mapped to each device pair for the transit VPC. Valid values are 1 through 32.
4. In the **Device Pair 1#** fields, select the serial numbers of each device in the pair. To remove a device serial number, click the **X** that appears in a field.

The devices that appear in this field have been associated with a configuration template and support the WAN Edge Version that you selected.

5. To add additional device pairs, click .

To remove a device pair, click .

A transit VPC can be associated with one to four device pairs. To enable the Cloud onRamp auto-scale feature for AWS, you must associate at least two device pairs with the transit VPC.

6. Click **Save and Finish** to complete the transit VPC configuration or optionally click **Proceed to Discovery and Mapping** to continue with the wizard.
7. Click **Advanced** if you wish to enter more specific configuration options:
  - a. In the **Transit VPC CIDR** field, enter a custom CIDR that has a network mask in the range of 16 to 25. If you choose to leave this field empty, the Transit VPC is created with a default CIDR of 10.0.0.0/16.
  - b. In the **SSH PEM Key** drop-down, select a PEM key pair to log in to an instance. Note that the key pairs are region-specific. Refer to the AWS documentation for instructions on creating key pairs.
8. Click **Save and Finish** to complete the transit VPC configuration, or optionally click **Proceed to Discovery and Mapping** to continue with the wizard.
9. Select hosts to discover:
  - a. In the **Select an account to discover** field, select a host to map to this transit VPC.
  - b. Click **Discover Host VPCs**.
  - c. In the table that displays, choose one or more hosts to map to this transit VPC.

You can use the search field and options to display only host VPCs that mention specific search criteria.

You can click the **Refresh** icon to update the table with current information.

You can click the **Show Table Columns** icon to specify which columns display in the table.
  - d. Click **Next**.

- Step 7** Map the host VPCs to transit VPCs:
- In the table of host VPCs, select the desired host VPCs.
  - Click **Map VPCs**. The Map Host VPCs popup opens.
  - In the **Transit VPC** drop-down, select the transit VPC to map to the host VPCs.
  - In the **VPN** drop-down, select the VPN in the overlay network in which to place the mapping.
  - Enable the **Route Propagation** option if you want Cisco vManage to automatically propagate routes to the host VPC routes table.
  - Click **Map VPCs**.
  - Click **Save and Complete**.

In the VPN feature configuration template for VPN 0, when configuring the two cloud routers that form the transit VPC, ensure that the color you assign to the tunnel interface is a public color, not a private color. Public colors are **3g**, **biz-internet**, **blue**, **bronze**, **custom1**, **custom2**, **custom3**, **default**, **gold**, **green**, **lte**, **metro-ethernet**, **mpls**, **public-internet**, **red**, and **silver**.

---

## Display Host VPCs

- Step 1** In the Cloud OnRamp Dashboard, click the pane for the desired VPC. The Host VPCs/Transit VPCs screen opens, and Host VPCs is selected by default. In the bar below this, Mapped Host VPCs is selected by default, and the table on the screen lists the mapping between host and transit VPCs, the state of the transit VPC, and the VPN ID.
- Step 2** To list unmapped host VPCs, click **Unmapped Host VPCs**. Then click **Discover Host VPCs**.
- Step 3** To display the transit VPCs, click **Transit VPCs**.

---

## Map Host VPCs to a Transit VPC

- Step 1** In the Cloud OnRamp Dashboard, click the pane for the desired VPC. The Host VPCs/Transit VPCs screen opens.
- Step 2** Click **Un-Mapped Host VPCs**.
- Step 3** Click **Discover Host VPCs**.
- Step 4** From the list of discovered host VPCs, select the desired host VPCs.
- Step 5** Click **Map VPCs**. The Map Host VPCs popup opens.
- Step 6** In the \ drop-down, choose the desired transit VPC.
- Step 7** In the **VPN** drop-down, choose the VPN in the overlay network in which to place the mapping.
- Step 8** Click **Map VPCs**.

---

## Unmap Host VPCs

- Step 1** In the Cloud OnRamp Dashboard, click the pane for the desired VPC. The Host VPCs/Transit VPCs screen opens.

- Step 2** Click **Mapped Host VPCs**.
- Step 3** From the list of VPCs, select the desired host VPCs.
- Step 4** Click **Unmap VPCs**.
- Step 5** Click **OK** to confirm the unmapping.

---

Unmapping host VPCs deletes all VPN connections to the VPN gateway in the host VPC, and then deletes the VPN gateway. When you make additional VPN connections to a mapped host VPC, they will be terminated as part of the unmapping process.

## Display Transit VPCs

- 
- Step 1** In the Cloud OnRamp Dashboard, click the pane for the desired VPC. The Host VPCs/Transit VPCs screen opens, and Host VPCs is selected by default.
  - Step 2** Click **Transit VPCs**.

---

The table at the bottom of the screen lists the transit VPCs

## Add Transit VPC

- 
- Step 1** In the Cloud onRamp Dashboard, click the pane for the desired VPC. The Host VPCs/Transit VPCs screen opens, and Host VPCs is selected by default.
  - Step 2** Click **Transit VPCs**.
  - Step 3** Click **Add Transit VPC**.
- To add a transit VPC, perform operations from step 6 of [Configure Cloud onRamp for IaaS for Amazon Web Services, on page 6](#).

## Delete Device Pair

### Before you begin

The device pair must be offline.

- 
- Step 1** Go to the Cloud onRamp Dashboard.
  - Step 2** Click a device pair ID.
  - Step 3** Verify that the status of the device pair is offline.
  - Step 4** To descale the device pairs, click the trash can icon in the Action column or click the **Trigger Autoscale** option.
-

## Delete Transit VPC



---

**Note** To delete the last pair of online device pairs, you must delete a transit VPC.

---

### Before you begin

Delete the device pairs that are associated with the transit VPC.

---

**Step 1** In the Cloud onRamp Dashboard, click the pane for the desired VPC. The Host VPCs/Transit VPCs screen opens, and Host VPCs is selected by default.

**Step 2** Click **Host VPCs**.

**Step 3** Select all host VPCs, and click **Unmap VPCs**.

Ensure that all host mappings with transit VPCs are unmapped.

**Step 4** Click **OK** to confirm the unmapping.

**Step 5** Click **Transit VPCs**.

**Step 6** Click the trash icon to the left of the row for the transit VPC.

**Note** The trash icon is not available for the last device pair of transit VPC. Hence, to delete the last device pair, click **Delete Transit** drop-down list at the right corner. The trash icon is only available from the second device pair onwards.

**Step 7** Click **OK** to confirm.

---

## Add Device Pairs

---

**Step 1** Click **Add Device Pair**.

**Note** Ensure that the devices you are adding are already associated with a device template.

**Step 2** In the box, select a device pair.

**Step 3** Click the **Add** icon to add more device pairs.

You can add up to a total of four device pairs to the transit VPC.

**Step 4** Click **Save**.

---

## History of Device Pairs for Transit VPCs

---

**Step 1** To display the Transit VPC Connection History page with all the corresponding events, click **History for a device pair**.

**Step 2** View a histogram of events that have occurred in the previous one hour is displayed and a table of all events for the selected transit VPC. The table lists all the events generated in the transit VPC. The events can be one of the following:

- Device Pair Added
  - Device Pair Spun Up
  - Device Pair Spun Down
  - Device Pair Removed
  - Host Vpc Mapped
  - Host Vpc Unmapped
  - Host Vpc Moved
  - Transit Vpc Created
  - Transit Vpc Removed
- 

## Edit Transit VPC

---

**Step 1** Click **Edit Transit Details**. Provide a value for the maximum number of host VPCs per device pair to which the transit VPC can be mapped.

**Step 2** Click **OK**.

This operation can trigger auto-scale.

---

## Configure Cloud onRamp for IaaS for Azure

To configure Cloud onRamp for IaaS for Azure, you create Azure transit VNets, each of which consist of a pair of routers. You then map the host vNets to transit VNets that already exist in the Azure cloud. All VNets reside in the same resource group.

- Transit VNets provide the connection between the overlay network and the cloud-based applications running on host VNet. Each transit VNet consists of two routers that reside in their own VNet. Two routers are used to provide redundancy for the connection between the overlay network and cloud-based applications. On each of these two cloud routers, the transport VPN (VPN 0) connects to a branch router, and the service-side VPNs (any VPN except for VPN 0 and VPN 512) connect to applications and application providers in the public cloud.
- Host VNets are virtual private clouds in which your cloud-based applications reside. When a transit VNet connects to an application or application provider, it is simply connecting to a host VNet.

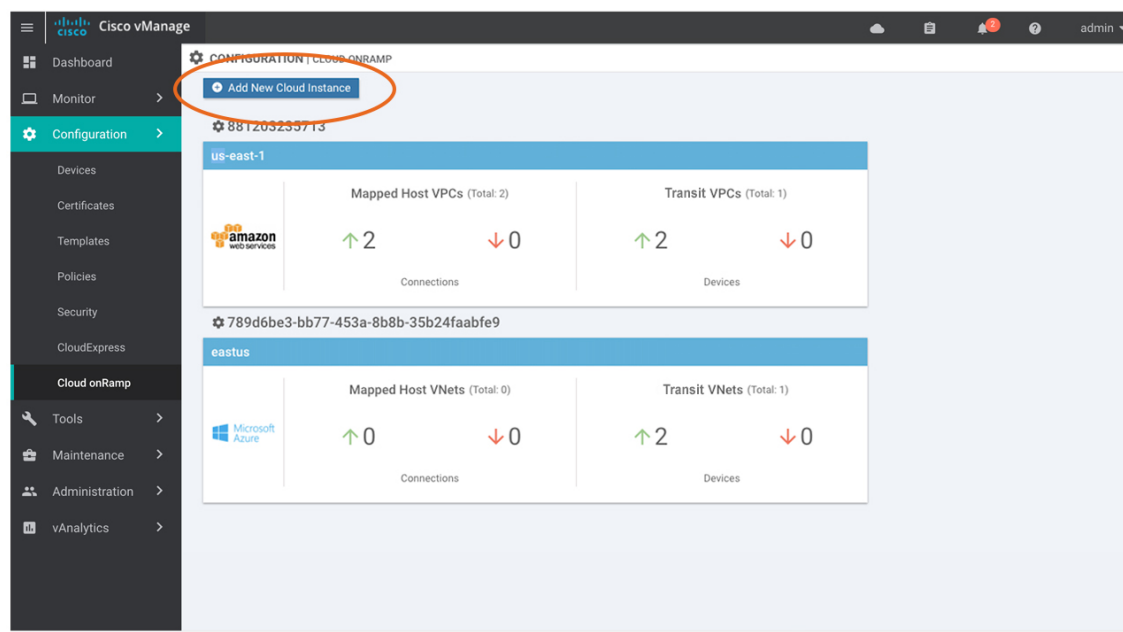
In the Cloud onRamp configuration process, you map one or more host VPCs or host VNets to a single transit VPC or transit VNet. In doing this, you are configuring the cloud-based applications that branch users are able to access.

The mapping process establishes IPsec and BGP connections between the transit VPC or transit VNet and each host VPC or host VNet. The IPsec tunnel that connects the transit and host VPC or VNet runs IKE to provide security for the connection. For AWS, the IPsec tunnel runs IKE Version 1. For Azure, the IPsec tunnel runs IKE version 2. The BGP connection that is established over the secure IPsec tunnel allows the transit and host VPC or VNet to exchange routes so that the transit VPC or VNet can direct traffic from the branch to the proper host VPC or VNet, and hence to the proper cloud-based application.

During the mapping process, the IPsec tunnels and BGP peering sessions are configured and established automatically. After you establish the mappings, you can view the IPsec and BGP configurations, in the VPN Interface IPsec and BGP feature configuration templates, respectively, and you can modify them as necessary. You can configure Cloud OnRamp for IaaS for Azure by using the configuration wizard:

### Create a Cloud Instance

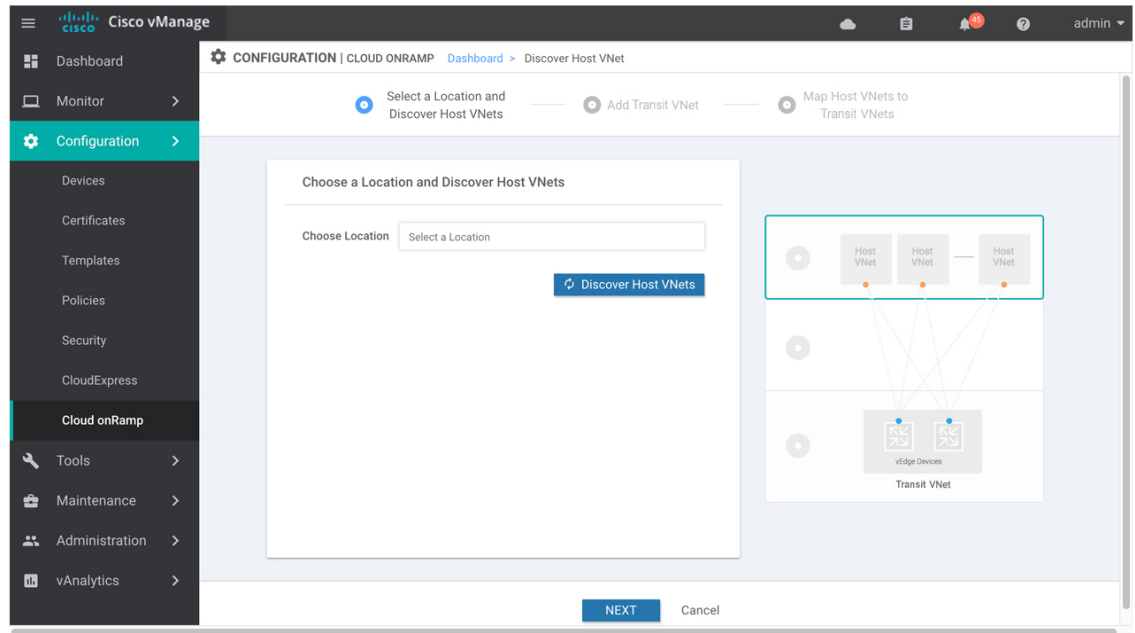
1. In vManage NMS, select the **Configuration > Cloud onRamp for IaaS** screen.
2. Click **Add New Cloud Instance**:



3. In the Add Cloud Instance–Log In to a Cloud Server popup:
  - a. In the **Cloud** drop-down, select **Azure** as the cloud type.
  - b. To give vManage programmatic access to your Azure Subscription, log in to the cloud server:
    1. In the **Subscription ID** field, enter the ID of the Azure subscription you want to use as part of the Cloud OnRamp workflow.
    2. In the **Client ID** field, enter the ID of an existing application or create a new application in Azure. To create a new application, go to your **Azure Active Directory > App Registrations > New Application Registration**.
    3. In the **Tenant ID** field, enter the ID of your Azure account. To find the tenant ID, go to your Azure Active Directory and click **Properties**.
    4. In the **Secret Key** field, enter the password associated with the client ID.

4. Click **Log In**. The cloud instance configuration wizard opens.


This wizard consists of three screens that you use to select a location and discover host VNets, add transit VNet, and map host VNets to transit VNets. A graphic on the right side of each wizard screen illustrates the steps in the cloud instance configuration process. Steps not yet completed are shown in light gray. The current step is highlighted within a blue box. Completed steps are indicated with a green checkmark and are shown in light orange.




5. Select a location and discover host VNets:

- a. In the **Choose Location** drop-down, select a geographical location.
- b. Click **Save and Finish** to create a transit VNet or optionally click **Proceed to Discovery and Mapping** to continue with the wizard.

6. Add a transit VNet:

- a. In the **Transit VNet Name** field, type a name for the transit VNet.  
The name can be up to 32 characters and can contain only uppercase and lowercase letters, the digits 0 through 9, hyphens (-), and underscores (\_). It cannot contain spaces or any other characters.
- b. Under **Device Information**, enter information about the transit VNet:
  1. In the **WAN Edge Version** drop-down, select the software version to run on the VNet transit.  
The drop-down lists the published versions of the Viptela software in the Azure marketplace.
  2. In the **Size of Transit VNet** drop-down, select how much memory and how many CPUs to create on the VNet transit.
  3. In the **Device 1** drop-down, select the serial number to use.
  4. In the **Device 2** drop-down, select the serial number to use.
  5. To add additional device pairs, click .

To remove a device pair, click .

6. Click **Save and Finish** to complete the transit VNet configuration or optionally click **Proceed to Discovery and Mapping** to continue with the wizard.
  7. Click **Advanced** if you wish to enter more specific configuration options.
  8. In the **Transit VNet CIDR** field, enter a custom CIDR that has a network mask in the range of 16 to 25. If you choose to leave this field empty, the Transit VPC is created with a default CIDR of 10.0.0.0/16.
- c. Click **Save and Finish** to complete the transit VPC configuration, or optionally click **Proceed to Discovery and Mapping** to continue with the wizard.
7. Map the host VNets to transit VNets:
    - a. In the table of host VNets, select the desired host VNet.
    - b. Click **Map VNets**. The Map Host VNets popup opens.
    - c. In the **Transit VNet** drop-down, choose the transit VNet to map to the host VNets.
    - d. In the **VPN** drop-down, choose the VPN in the overlay network in which to place the mapping.
    - e. In the IPsec Tunnel CIDR section, enter two pairs of interface IP addresses for each Cisco vEdge device and a pair of loopback IP addresses to configure IPsec tunnels to reach the Azure virtual network transit. The IP addresses must be network addresses in the /30 subnet, be unique across the overlay network, and not be a part of the host VNet CIDR. If they are part of the host VNet CIDR, Azure will return an error while attempting to create VPN connections to the transit VNet.

As Azure supports single Virtual Private Gateway (VGW) configuration over IPsec tunnels with redundancy provided over a single tunnel, Cloud OnRamp for IaaS supports two VGWs for redundancy. During a planned maintenance or an unplanned event of a VGW, the IPsec tunnel from the VGW to the cloud routers get disconnected. This loss of connectivity causes the cloud routers lose BGP peering with vManage over IPsec tunnel. To enable BGP peering with cloud routers rather than IP address of the IPsec tunnel, provide the loopback addresses for each cloud router.




---

**Note**

The loopback option for BGP peering supports single and multiple Virtual Gateway or Customer Gateway configuration or both on Azure cloud. This option applies only to the new host VNets mapped to transit VNets and not on the existing VNets.

---

- f. In the Azure Information section:
  1. In the **BGP ASN** field, enter the ASN that will be configured on the Azure Virtual Network Gateway that is spun up within the host VNet. Use an ASN that is not part of an existing configuration on Azure. For acceptable ASN values, refer to Azure documentation.
  2. In the **Host VNet Gateway Subnet** field, enter a host VNet subnet in which the Virtual Network Gateway can reside. It is recommended you use a /28 subnet or higher. You must not provide a subnet that is already created in the VNet.
- g. Click **Map VNets**.



- h. Click **Save** and **Complete**.

In the VPN feature configuration template for VPN 0, when configuring the two cloud routers that form the transit VNet, ensure that the color you assign to the tunnel interface is a public color, not a private color. Public colors are **3g**, **biz-internet**, **blue**, **bronze**, **custom1**, **custom2**, **custom3**, **default**, **gold**, **green**, **lte**, **metro-ethernet**, **mpls**, **public-internet**, **red**, and **silver**.

### Display Host VNets

1. In the Cloud onRamp Dashboard, click the pane for the desired VNet. The Host VNets/Transit VNets screen opens, and Host VNets is selected by default. In the bar below this, Mapped Host VNets is selected by default, and the table on the screen lists the mapping between host and transit VNets, the state of the transit VNet, and the VPN ID.
2. To list unmapped host VNets, click **Unmapped Host VNets**.
3. To display the transit VNets, click **Transit VNets**.

### Map Host VNets to an Existing Transit VNet

1. In the Cloud onRamp Dashboard, click the pane for the desired location of the required account. The Host VNets/Transit VNets screen opens.
2. Click **Unmapped Host VNets**.
3. Click **Discover Host VNets**.
4. From the list of discovered host VNets, select the desired host VNet.
5. Click **Map VNets**. The Map Host VNets popup opens.
6. In the **Transit VNet** drop-down, select the desired transit VNet.
7. In the **VPN** drop-down, choose the VPN in the overlay network in which to place the mapping.
8. Click **Map VNets**.

### Unmap Host VNets

1. In the Cloud onRamp Dashboard, click the pane for the desired VNet. The Host VNets/Transit VNets screen opens.
2. Click **Mapped Host VNets**.
3. From the list of VNets, select the desired host VNets. It is recommended that you unmap one vNet at a time. If you want to unmap multiple vNets, do not select more than three in a single unmapping operation.
4. Click **Unmap VNets**.
5. Click **OK** to confirm the unmapping.

### Display Transit VNets

1. In the Cloud onRamp Dashboard, click the pane for the desired VNets. The Host VNets/Transit VNets screen opens, and Host VNets is selected by default.

2. Click **Transit VNets**.

The table at the bottom of the screen lists the transit VNets.

#### Add a Transit VNet

1. In the Cloud onRamp Dashboard, click the pane for the desired VNet. The Host VNets/Transit VNets screen opens, and Host VNets is selected by default.
2. Click **Transit VNets**.
3. Click **Add Transit VNet**.

#### Delete a Transit VNet

1. In the Cloud onRamp Dashboard, click the pane for the desired VNet. The Host VNets/Transit VNets screen opens, and Host VNets is selected by default.
2. Click **Mapped Host VNets**.
3. Select the desired host VNet, and click **Unmap VNets**.
4. Click **OK** to confirm the unmapping.
5. Click **Transit VNets**.
6. Click the trash icon to the left of the row for the transit VNet.
7. Click **OK** to confirm.

## Troubleshoot Cloud onRamp for IaaS

This section describes how to troubleshoot common problems with Cloud onRamp for IaaS.

### Two vEdge Routers are Not Available

In vManage NMS, when you select the **Configuration > Cloud onRamp for IaaS** screen, and click **Add New Cloud instance**, you see an error message indicating that two vEdge routers are not available.

#### Resolve the Problem

The vManage NMS does not have two vEdge Cloud routers that are running licensed Cisco SD-WAN software. Contact your operations team so that they can create the necessary vEdge Cloud routers.

If the vEdge routers are present and the error message persists, the two vEdge Cloud routers are not attached to configuration templates. Attach these templates in the vManage **Configuration > Templates** Device screen. Select the vEdge Cloud router, and then select **Attach Devices** from the More Actions icon to the right of the row.

### Required Permissions for API

#### Problem Statement

When you enter your API keys, you get an error message indicating that this user does not have the required permissions.

### Resolve the Problem

Ensure that the vManage server can reach the internet and has a DNS server configured so that it can reach AWS or Azure. To configure a DNS server, in the vManage VPN feature configuration template, enter the IP address of a DNS server, and then reattach the configuration template to the vManage server.

For AWS, check the API keys belonging to your AWS account. If you think you have the wrong keys, generate another pair of keys.

For AWS, if you are entering the correct keys and the error message persists, the keys do not have the required permissions. Check the user permissions associated with the key. Give the user the necessary permissions to create and edit VPCs and EC2 instances.

If the error message persists, check the time of the vManage server to ensure that it is set to the current time. If it is not, configure the vManage server's time to point to the Google NTP server. In the vManage NTP feature configuration template, enter a hostname of time.google.com, time2.google.com, time3.google.com, or time4.google.com. Then reattach the configuration template to the vManage server.

### No vEdge Software Versions Appear in the Drop-Down

#### Problem Statement

When you are trying to configure transit VPC parameters for the transit VPC, no vEdge Cloud software versions are listed in the drop-down.

#### Resolve the Problem

Ensure that your customer account has subscribed to the Cisco SD-WAN vEdge Cloud routers.

Ensure that the vEdge cloud router is running software Release 17.2.0 or later.

### No VPNs Appear in Drop-Down

#### Problem Statement

When you select the host VPCs or VNets to map, no VPNs are listed in the drop-down.

#### Resolve the Problem

This problem occurs when the device configuration template attached to the cloud router includes no service-side VPNs. Service-side VPNs (VPNs other than VPN 0 and VPN 512) are required to configure the IPsec connection between the two cloud routers selected for the transit and host VPCs or VNets.

This problem can also occur if the two cloud routers selected for the transit VPC or VNet have no overlapping service-side VPNs. Because the two vEdge routers form an active-active pair, the same service-side VPNs must be configured on both of them.

To configure service-side VPNs, in the vManage VPN feature configuration template, configure at least one service-side VPN. Ensure that at least one of the service-side VPNs is the same on both routers. Then reattach the configuration template to the routers.

### Cloud onRamp Task Fails

#### Problem Statement

After you have completed mapping the host VPCs to the transit VPCs, or host VNets to transit VNets, the Cloud OnRamp tasks fail.

#### Resolve the Problem

Review the displayed task information that is displayed on the screen to determine why the task failed. If the errors are related to AWS or Azure resources, ensure that all required resources are in place.

### **Cloud onRamp Task Succeeds, But Routers Are Down**

#### **Problem Statement**

The Cloud OnRamp task was successful, but the cloud routers are still in the Down state.

#### **Resolve the Problem**

Check the configuration templates:

- Check that all portions of the cloud router configuration, including policies, are valid and correct. If the configuration are invalid, they are not applied to the router, so the router never comes up.
- Check that the configuration for the vBond orchestrator is correct. If the DNS name or IP address configured of the vBond orchestrator is wrong, the vEdge router is unable to reach it and hence is unable to join the overlay network.

After you have determined what the configuration issues are:

1. Delete the Cloud OnRamp components:
  - a. Unmap the host VPNs and the transit VPCs or VNets.
  - b. Delete the transit vEdge routers.
2. Edit the configuration templates and reattach them to the cloud routers.
3. Repeat the Cloud OnRamp configuration process.

### **Desired Routes Not Exchanged**

#### **Problem Statement**

The Cloud OnRamp configuration workflow is successful, the Cloud vEdge routers are up and running, but the desired routes are not getting exchanged.

#### **Resolve the Problem**

In vManage NMS, check the BGP configuration on the transit cloud routers. During the mapping process when you configure Cloud OnRamp service, BGP is configured to advertise the network 0.0.0.0/0. Make sure that the service-side VPN contains an IP route that points to 0.0.0.0/0. If necessary, add a static route in the VPN feature configuration template, and then reattach the configuration to the two cloud routers that you selected for the transit VPC or VNet.

On AWS, go to the host VPC and check its route table. In the route table, click the option **Enable route propagation** to ensure that the VPC receives the routes.

### **End-to-End Ping Is Unsuccessful**

#### **Problem Statement**

Routing is working properly, but an end-to-end ping is not working.

#### **Resolve the Problem**

On AWS, check the security group rules of the host VPC. The security group rules must allow the source IP address range subnets of the on-premises or branch-side devices, to allow traffic from the branch to reach AWS.

