Cloud OnRamp Overview

Cloud OnRamp for IaaS

Cloud OnRamp for IaaS extends the fabric of the Cisco Catalyst SD-WAN overlay network into public clouds, allowing branches with Cisco CSR1000V Cloud Services 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 Cisco CSR 1000V 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 Catalyst 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.

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.

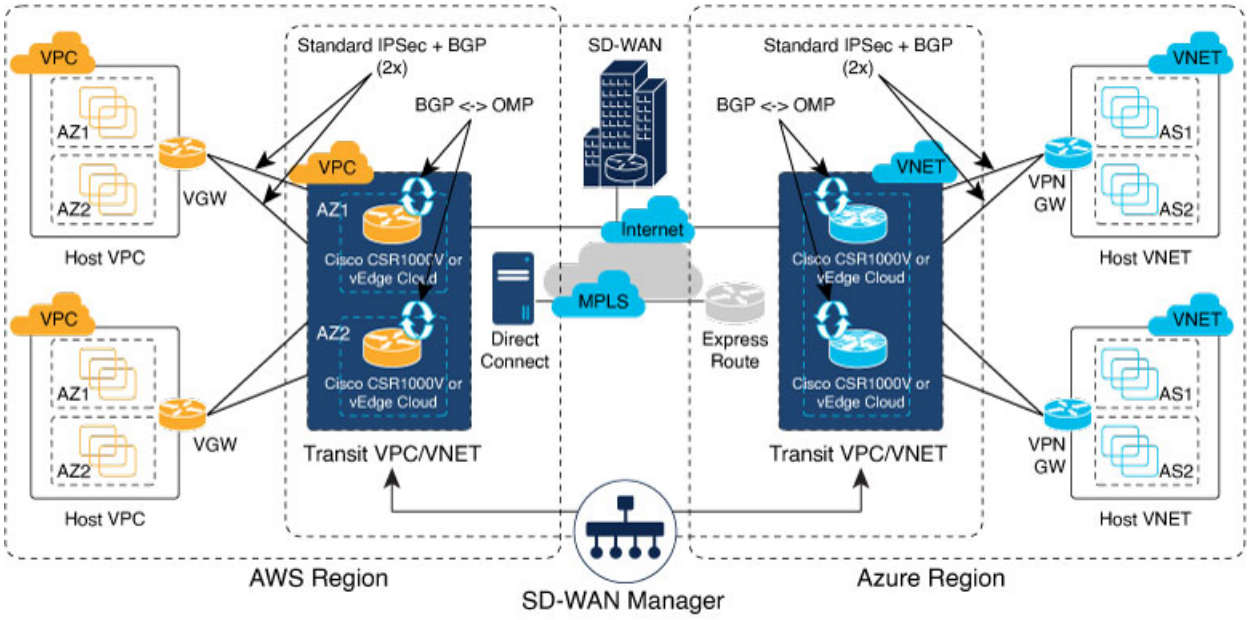

Supported Routers

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*.

Provision vManage for Cloud OnRamp for IaaS

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

vManage NMS Prerequisites

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

- 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.
- 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

- 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

AWS Prerequisites

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

- 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*.
- 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.
- 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 Cisco CSR 1000V routers
  - 1 AWS virtual transit (VGW) for each host VPC
  - 4 VPN connections for mapping each host VPC

Note Cisco IOS XE Catalyst 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.

Note Cisco CSR 1000V support C3 and C4 compute-intensive families.

Subscribe to Cisco Catalyst SD-WAN AMIs

To use the Cloud OnRamp for IaaS and other Cisco Catalyst 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 Cisco CSR 1000V AMI only once, when you first create a Viptela AMI instance.

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

1. In AWS, search to locate a cloud router AMI for your devices.
2. Select and launch an EC2 instance with the AMI instance. For more information, see *Create Cisco IOS XE Catalyst SD-WAN Cloud VM Instance on AWS*.
3. Generate a key pair. For full instructions, see *Set Up the Cisco Catalyst SD-WAN Cloud VM Instance*.
4. Click **Download Key Pair**. The key pair then downloads to your local computer as a .pem file.
5. Click **Launch Instance**. A failure message displays, because you now need to upload the key pair to complete the subscription process.

To upload the key pair:

1. In AWS Marketplace, search for your router AMI.
2. Click **Continue**.
3. Click **Key Pair** to bring up a Cisco CSR 1000V 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 3 when creating a new AMI subscription.

Azure Prerequisites

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

- Ensure that you have accepted the terms and conditions for the Cisco CSR 1000V Router in the Azure Marketplace. See *Accept the Azure Terms of Service*
- 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. See *Create and Retrieve Azure Credentials*.
- 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 VNet, which is required for creating the transit VNet
  - 1 Availability set, required for Virtual Machine distribution in the transit VNet
  - 6 Static Public IP addresses associated with the transit cloud routers
  - 1 Azure Virtual Network Gateway and 2 Static Public IP Addresses for each host VNet
  - 4 VPN connections for mapping each host VNet

Note Cisco IOS XE Catalyst 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.

- F-series VMs (F4 and F8) are supported on the cloud routers.

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 this 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.

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.
   d. In the **Name** field, enter a descriptive name such as CloudOnRampApp.
   e. In the **Application Type** field, select **Web app / API**
   f. In the **Sign-on URL** field, enter any valid sign-on URL; this URL is not used in Cloud OnRamp.
   g. Click **Create**. The system displays a summary screen with the Application ID.
3. Create a secret key for the Cloud OnRamp application:
   a. In the summary screen, click **Settings** in the upper-left corner.
   b. In the right pane, click **Keys**. The system displays the **Keys** > **Password** screen.
   c. On the Passwords screen:
      1. In the **Description** column, enter a description for your secret key.
      2. In the **Expires** column, from the **Duration** drop-down, select the duration for your secret key.
      3. Click **Save** in the upper-left corner of the screen. The system displays the secret key in the Value column but then hides it permanently, so be sure to copy and save the password in a separate location.
4. In the left pane of the Azure portal, click **Subscriptions** to view the subscription ID. If you have multiple subscriptions, copy and save the subscription ID which you are planning to use for configuring the Cloud OnRamp application.
5. View the Tenant ID:
   a. In the left pane of the Azure portal, click **Azure Active Directory**.
   b. Click **Properties**. The system displays the directory ID which is equivalent to the tenant ID.

6. Assign Contributor privileges to the application:
   a. In the left pane of the Azure portal, click **Subscriptions**.
   b. Click the subscription that you will be using for the Cloud OnRamp application.
   c. In the subscription pane, navigate to Access Control (IAM).
   d. Click **Add**. The system displays the Add Permissions screen.
   e. From the **Role** drop-down menu, select **Contributor**.
   f. From the **Assign Access To** drop-down, select the default value **Azure AD user**, **group**, or **application**.
   g. From the **Select** drop-down, select the application you just created for Cloud OnRamp.
   h. Click **Save**.

You can now log into the Cloud OnRamp application with the Azure credentials you just created and saved.

Configure Cloud OnRamp for IaaS for AWS

Configure Cloud OnRamp for IaaS for AWS

To configure Cloud OnRamp for IaaS for AWS, you create AWS transit VPCs, each of which consists of up to four pairs of Cisco IOS XE Catalyst SD-WAN devices. You then map the transit virtual private clouds (VPC)s to host VPCs that already exist in the AWS cloud.

- 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.

1. In vManage NMS, select the **Configuration** > **Cloud onRamp for IaaS** screen.

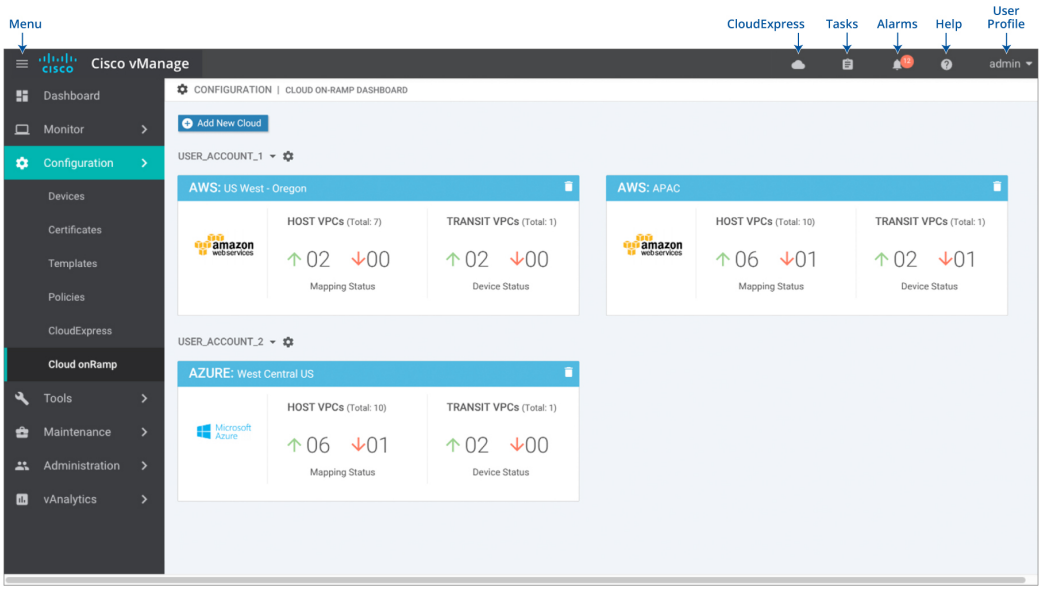

2. Click **Add New Cloud Instance**.
3. In the Add Cloud Instance – log in to a Cloud Server popup:
   a. In the **Cloud** drop-down, select the **Amazon Web Services** radio button.
   b. Click **IAM Role** or **Key** to log in to the cloud server. It is recommended that you use IAM Role.
   c. If you 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 vManage NMS 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.
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.

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.

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**.

7. Map the host VPCs to transit VPCs:
   a. In the table of host VPCs, select the desired host VPCs.
   b. Click **Map VPCs**. The Map Host VPCs popup opens.
   c. In the **Transit VPC** drop-down, select the transit VPC to map to the host VPCs.
   d. In the **VPN** drop-down, select the VPN in the overlay network in which to place the mapping.
   e. Enable the **Route Propagation** option if you want vManage to automatically propagate routes to the host VPC routes table.
   f. Click **Map VPCs**.
   g. 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

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.
2. To list unmapped host VPCs, click **Unmapped Host VPCs**. Then click **Discover Host VPCs**.
3. To display the transit VPCs, click **Transit VPCs**.

Map Host VPCs to a Transit VPC

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

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

Unmap Host VPCs

1. In the Cloud OnRamp Dashboard, click the pane for the desired VPC. The Host VPCs/Transit VPCs screen opens.
2. Click **Mapped Host VPCs**.
3. From the list of VPCs, select the desired host VPCs.
4. Click **Unmap VPCs**.
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

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.
2. Click **Transit VPCs**.

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

Add Transit VPC

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.
2. Click **Transit VPCs**.
3. Click **Add Transit VPC**.

   To add a transit VPC, perform operations from step 6 of Configure Cloud OnRamp for IaaS for AWS, on page 6.

Delete Device Pair

The device pair must be offline.

1. In the Cloud OnRamp Dashboard,

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

Delete Transit VPC

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

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

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.
2. Click **Host VPCs**.
3. Select all host VPCs, and click **Unmap VPCs**.

   Ensure that all host mappings with transit VPCs are unmapped.
4. Click **OK** to confirm the unmapping.
5. Click **Transit VPCs**.
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.

7. Click **OK** to confirm.

Add Device Pairs

1. Click **Add Device Pair**.

   Ensure that the devices you are adding are already associated with a device template.
2. In the box, select a device pair.
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.
4. Click **Save**.

History of Device Pairs for Transit VPCs

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

In this view, by default, 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

You can change the maximum number of host VPCs that can be mapped to a device pair.

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.
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.
- The Cloud onRamp for IaaS workflow uses a public IP address of the second WAN interface to set up the Customer Gateway for mapping (ipsec tunnels) the host VNets to a transit VNet. To add the public IP address of the WAN interface, configure the VPN Interface Ethernet template with GigabitEthernet2 for the devices used in Cloud onRamp for IaaS. In Cisco CSR1000V, the tunnel interface is on the GigabitEthernet2 interface.
- 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**:

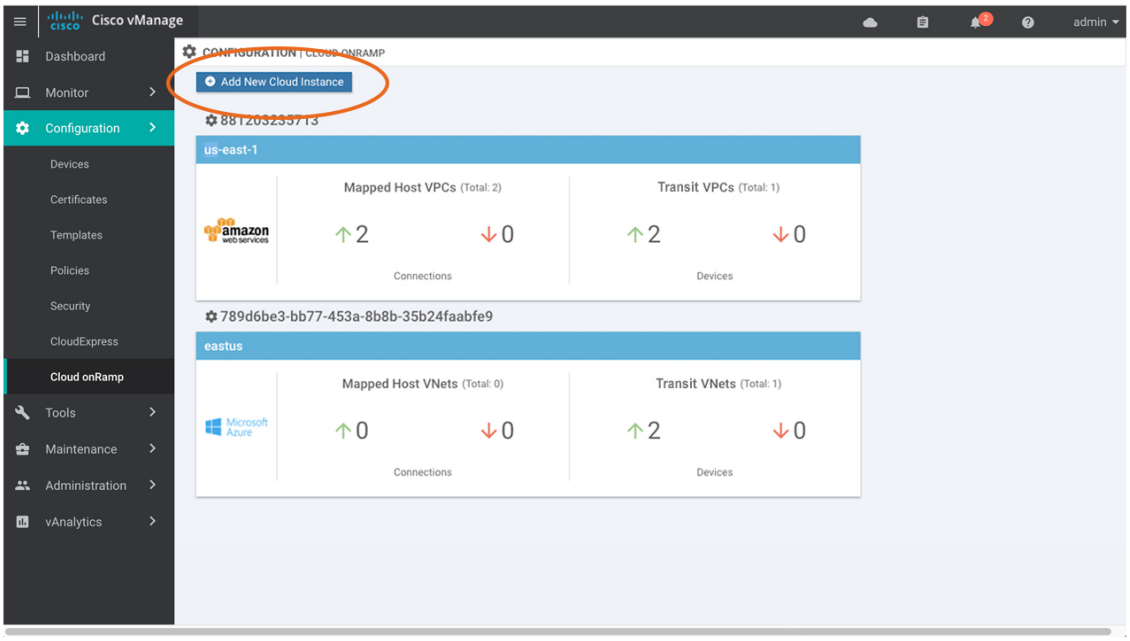

368405

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.

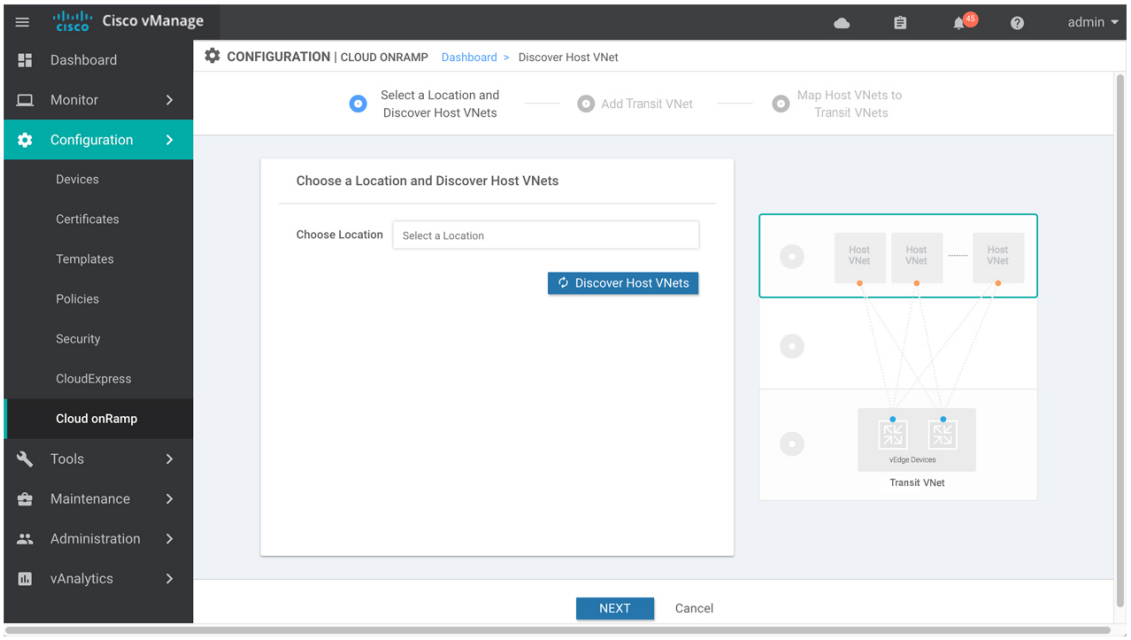

5. Select a location and discover host VNets:
   a. In the **Choose Location** drop-down, select a geographical location.
   b. Click **Save and Finish** to create a transit VNet or optionally click **Proceed to Discovery and Mapping** to continue with the wizard.

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 Router 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.
   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 Cisco Catalyst SD-WAN Cloud OnRamp for IaaS

This section describes how to troubleshoot common problems with Cisco Catalyst SD-WAN Cloud OnRamp for IaaS.

Two Cisco CSR1000V or Cisco Catalyst 8000V Devices are Not Available

From the Cisco SD-WAN Manager menu, choose **Configuration** > **Cloud onRamp for IaaS**. After you click **Add New Cloud Instance**, you see an error message indicating that two Cisco CSR1000V or Cisco Catalyst 8000V devices aren't available.

Resolve the Problem

The Cisco SD-WAN Manager server doesn't have two Cisco CSR1000V or Cisco Catalyst 8000V devices that are running licensed Cisco Catalyst SD-WAN software. Contact your operations team so that they can create the necessary Cisco CSR1000V or Cisco Catalyst 8000V devices.

If the Cisco CSR1000V or Cisco Catalyst 8000V devices are present and the error message persists, then the two devices aren't attached to configuration templates. Attach these templates in the Cisco SD-WAN Manager **Configuration** > **Templates** Device window. For the desired device templates, click **…** and choose **Attach Devices**.

Required API Permissions are Unavailable

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

Resolve the Problem

Ensure that the Cisco SD-WAN Manager server can reach the internet and has a DNS server configured so that it can reach AWS or Microsoft Azure. To configure a DNS server, in the Cisco SD-WAN Manager VPN feature configuration template, enter the IP address of a DNS server, and then reattach the configuration template to the Cisco SD-WAN Manager 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're entering the correct keys and the error message persists, the keys don't have the required permissions. Check the user permissions associated with the key. Give necessary permissions to the user to create and edit VPCs and EC2 instances.

If the error message persists, check the time of the Cisco SD-WAN Manager server to ensure that it's set to the current time. If it's not, configure the Cisco SD-WAN Manager server time to point to the Google NTP server. To configure the server time, in the Cisco SD-WAN Manager NTP feature configuration template, enter the hostname of an NTP server. Next, reattach the configuration template to the NTP feature using Cisco SD-WAN Manager. The Google NTP servers are time.google.com, time2.google.com,time3.google.com, and time4.google.com, and so on.

WAN Edge Router Software Versions don't Appear in the Drop-Down When Configuring for AWS

Problem Statement

When you're trying to configure transit VPC parameters for the transit VPC, Cisco CSR1000V and Cisco Catalyst 8000V devices software versions aren't listed in the drop-down list.

Resolve the Problem

Ensure that you subscribe to the Cisco CSR1000V or Cisco Catalyst 8000V devices Amazon machine image (AMI) in your account within the AWS Marketplace.

Ensure that the Cisco CSR1000V is using software Release 16.12.1b or later and Cisco Catalyst 8000V is using software Release 17.4.1a or later.

VPNs aren't Listed During Configuration

Problem Statement

After you select the host VPCs or VNets to map, VPNs aren't listed in the drop-down list.

Resolve the Problem

The problem occurs when the device configuration template attached to the Cisco Catalyst SD-WAN cloud devices doesn't include service-side VPNs. You require the service-side VPNs (VPNs other than VPN 0 and VPN 512) to configure the IPsec connection between the two Cisco Catalyst SD-WAN cloud devices that you select for the transit and host VPCs or VNets.

This problem can also occur if the two Cisco Catalyst SD-WAN cloud devices that you select for the transit VPC or VNet have no overlapping service-side VPNs. Because the two Cisco Cloud Services router 1000V or Cisco Catalyst 8000V devices form an active–active pair, configure the same service-side VPNs on both of them.

To configure service-side VPNs, in the Cisco SD-WAN Manager 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.

Cisco Catalyst SD-WAN Cloud OnRamp for IaaS Task Fails

Problem Statement

After you have completed mapping the host VPCs to the transit VPCs, or host VNets to transit VNets, the configuration of Cisco Catalyst SD-WAN Cloud OnRamp for IaaS fails.

Resolve the Problem

Review the displayed task information that appears 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.

Cisco Catalyst SD-WAN Cloud OnRamp for IaaS Task Succeeds, but Cisco Catalyst SD-WAN Cloud Devices Are Down

Problem Statement

The Cisco Catalyst SD-WAN Cloud OnRamp for IaaS task was successful, but the Cisco Catalyst SD-WAN cloud devices are still in the down state.

Resolve the Problem

Check the configuration templates:

- Check that all portions of the Cisco Catalyst SD-WAN cloud devices configuration, including policies, are valid and correct. If the configurations are invalid, they aren't applied to the router, and the router never comes up.
- Check that the configuration for the Cisco Catalyst SD-WAN Validator is correct. If the DNS name or IP address configured in the Cisco Catalyst SD-WAN Validator is wrong, the Cisco CSR1000V or Cisco Catalyst 8000V device are unable to reach the Cisco Catalyst SD-WAN Validator, and hence they are unable to join the overlay network.

After you have determined what the configuration issues are:

1. Delete the Cisco Catalyst SD-WAN Cloud OnRamp for IaaS components:
   a. Unmap the host VPCs or VNets and the transit VPCs or VNets.
   b. Delete the transit VPC for Cisco CSR1000V or Cisco Catalyst 8000V devices.
2. Edit the configuration templates and reattach them to the Cisco Catalyst SD-WAN cloud devices.
3. Repeat the Cisco Catalyst SD-WAN Cloud OnRamp for IaaS configuration process.

Desired Routes are Not Exchanged

Problem Statement

The Cisco Catalyst SD-WAN Cloud OnRamp for IaaS configuration workflow is successful, the Cisco CSR1000V or Cisco Catalyst 8000V devices are available and running, but the desired routes aren't getting exchanged.

Resolve the Problem

In Cisco SD-WAN Manager, check the BGP configuration on the transit cloud routers. During the mapping process, when you configureCisco Catalyst SD-WAN Cloud OnRamp for IaaS service, BGP is configured to advertise the network address, 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 the route table. In the route table, click **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 isn't working.

Resolve the Problem

On AWS, check the security group rules of the host VPC. On Azure, check the network security group rules of the host VNet. 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.

Cloud OnRamp for SaaS

Enterprise software providers deliver many applications as Software as a Service (SaaS) cloud applications, such as Dropbox, Microsoft Office365, and Salesforce. Latency and packet loss impact the performance of these applications, but in legacy networks, network administrators have little visibility into network characteristics between end users and SaaS applications. When a path is impaired in a legacy network, the manual process of shifting application traffic to an alternate path is complex, time consuming, and error prone.

Cloud OnRamp for SaaS (formerly called CloudExpress service) addresses these issues by optimizing performance for SaaS applications in the Cisco SD-WAN overlay network. From a central dashboard, Cloud OnRamp for SaaS provides clear visibility into the performance of individual cloud applications and automatically chooses the best path for each one. It responds to changes in network performance in real-time, intelligently re-routing cloud application traffic onto the best available path. The following image provdes a high level overview of OnRamp for SaaS.

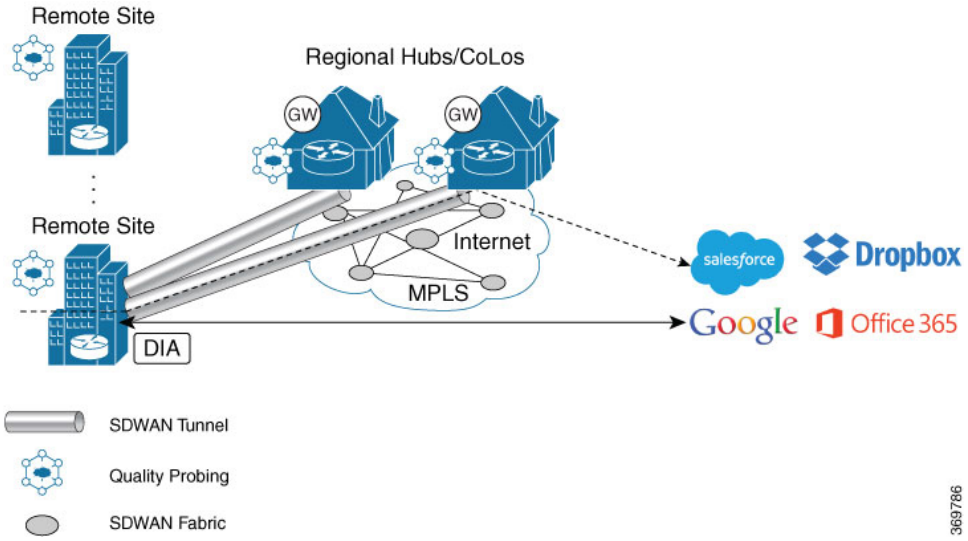

Cloud OnRamp for SaaS calculates a value called the Viptela Quality of Experience (vQoE). The vQoE value weighs loss and latency using a formula customized for each application. For example, email applications tolerate latency better than video applications, and video applications tolerate loss better than email applications. The vQoE value ranges from zero to ten, with zero being the worst quality and ten being the best. Cloud OnRamp for SaaS computes vQoE values for applications and paths, then assigns applications to the paths that best match their vQoE value. Cloud OnRamp for SaaS periodically recalculates vQoE values for paths to ensure ongoing optimal application performance.

Cloud OnRamp for SaaS supports the following enterprise applications:

- Amazon Web Service (AWS)
- Box

- Concur
- Dropbox
- Google Apps
- GoToMeeting
- Intuit
- Microsoft Office 365
- Oracle
- Salesforce
- SugarCRM
- Zendesk
- Zoho CRM

Cisco IOS XE Catalyst SD-WAN devices use VRFs in place of VPNs. So, when you complete the configuration, the system automatically maps the VPN configurations to VRF configurations.

Enable Cloud OnRamp for SaaS

You can enable Cloud OnRamp for SaaS in your Cisco SD-WAN overlay network on sites with Direct Internet Access (DIA) and on DIA sites that access the internet through a secure web gateway such as Zscaler or iboss. You can also enable Cloud OnRamp for SaaS on client sites that access the internet through another site in the overlay network, called a gateway site. Gateway sites can include regional data centers or carrier-neutral facilities. When you enable Cloud OnRamp for SaaS on a client site that accesses the internet through a gateway, you also enable Cloud OnRamp for SaaS on the gateway site.

All Cisco SD-WAN devices configured for Cloud OnRamp for SaaS must meet the following requirements:

- The devices must run Cisco SD-WAN Software Release 16.3 or higher.
- The devices must run in vManage mode.
- You must configure a DNS server address in VPN 0.
- You must configure local exit interfaces in VPN 0:
  - If the local interface list contains only physical interfaces, you must enable NAT on those interfaces. You can use normal default IP routes for next hops.
  - If the local interface list contains only GRE interfaces, you do not need to enable NAT on those interfaces. You can add default routes to the IP address of the GRE tunnel to the destination.

Enable Cloud OnRamp for SaaS

1. In vManage NMS, click **Administration** > **Settings**.
2. Click the **Edit** button to the right of the **Cloud onRamp for SaaS** bar.
3. In the **Cloud onRamp for SaaS** field, click **Enabled**.

4. Click **Save**.

Configure Cloud OnRamp for SaaS

Add Applications

1. In vManage NMS, select the **Configuration** > **Cloud onRamp for SaaS** screen. The Cloud OnRamp for SaaS Dashboard screen appears.

   To edit the VPN configured for an application, click the Edit icon for that application, then enter the new VPN. You can enter any VPN other than 0, which is the transport VPN, or 512, which is the management VPN.

2. To add applications, from the **Manage Cloud OnRamp for SaaS** drop-down, located to the right of the title bar, select **Applications** to add applications to the cloud onRamp configuration.
3. Click the **Add Applications and VPN** button. The Add Applications & VPN pop-up window appears.
4. In the **Applications** field, select an application.
5. In the **VPN** field, enter the service VPN in which that application runs. You can enter any VPN other than 0 and 512.
6. Click **Add**.
7. Repeat Steps 3 through 6 for each application you want to add.
8. Click **Save Changes**.

Configure Client Sites

To configure Cloud OnRamp for SaaS on client sites that access the internet through gateways, you must configure Cloud OnRamp for SaaS both on the client sites and on the gateway sites.

Note You cannot configure Cloud OnRamp for SaaS with Point-to-Point Protocol (PPP) interface on the gateway sites.

Client sites in Cloud onRamp service choose the best gateway site for each application to use for accessing the internet.

1. In vManage NMS, select the **Configuration** > **Cloud onRamp for SaaS** screen. The Cloud OnRamp for SaaS Dashboard screen opens.
2. From the **Manage Cloud OnRamp for SaaS** drop-down, located to the right of the title bar, select **Client Sites**. The screen changes and displays the following elements:
   - Attach Sites—Add client sites to Cloud onRamp for SaaS service.
   - Detach Sites—Remove client sites from Cloud onRamp for SaaS service.
   - Client sites table—Display client sites configured for Cloud onRamp for SaaS service.

3. In the Manage Sites screen, click the **Attach Sites** button. The Attach Sites screen displays all sites in the overlay network with available sites highlighted. For a site to be available, all devices at that site must be running in vManage mode.
4. In the Available Sites pane, select a client site to attach and click the right arrow. To remove a site, select it in the Selected Sites pane and click the left arrow.
5. Click **Attach**. vManage NMS pushes the feature template configuration to the devices. The Task View window displays a Validation Success message.
6. Select **Configuration** > **Cloud onRamp for SaaS** to return to the Cloud OnRamp for SaaS Dashboard screen.
7. From the **Manage Cloud OnRamp for SaaS** drop-down, located to the right of the title bar, choose **Gateways**. The screen changes and displays the following elements:
   - Attach Gateways—Attach gateway sites.
   - Detach Sites—Remove gateway sites from Cloud onRamp service.
   - Edit Sites—Edit interfaces on gateway sites.
   - Gateways table—Display gateway sites configured for Cloud onRamp service.
8. In the Manage Gateways screen, click the **Attach Gateways** button. The Attach Gateways popup window displays all sites in your overlay network with available sites highlighted. For a site to be available, all devices at that site must be running in vManage mode.
9. In the Available Gateways pane, select a gateway site to attach and click the right arrow. To remove a site, select it in the Selected Sites pane and click the left arrow.
10. If you do not specify interfaces for Cloud OnRamp for SaaS to use, the system selects a NAT-enabled physical interface from VPN 0. To specify GRE interfaces for Cloud OnRamp for SaaS to use:
    a. Click the link **Add interfaces** to selected sites (optional), located in the bottom right corner of the window.
    b. In the **Select Interfaces** drop-down, select GRE interfaces to add.
    c. Click **Save Changes**.
11. Click **Attach**. vManage NMS pushes the feature template configuration to the devices. The Task View window displays a Validation Success message.
12. To return to the Cloud OnRamp for SaaS Dashboard, select **Configuration** > **Cloud onRamp for SaaS**.

To edit Cloud OnRamp for SaaS interfaces on gateway sites:

1. Select the sites you want to edit and click **Edit Gateways**.
2. In the **Edit Interfaces** of Selected Sites screen, select a site to edit.
   - To add interfaces, click the **Interfaces** field to select available interfaces.
   - To remove an interface, click the **X** beside its name.
3. Click **Save Changes** to push the new template to the Cisco CSR 1000V routers.

Configure DIA Sites

1. In vManage NMS, select the **Configuration** > **Cloud onRamp for SaaS** screen. The Cloud OnRamp for SaaS Dashboard screen opens.

   In the title bar, choose **Manage Cloud OnRamp for SaaS** > **DIA**. The screen changes and displays the following elements:

   - Attach DIA Sites—Attach DIA sites.
   - Detach DIA Sites—Remove DIA sites.
   - Edit DIA Sites—Edit interfaces on DIA sites.
   - Sites table—Display sites configured for Cloud onRamp service.

2. From the **Manage Cloud OnRamp for SaaS** drop-down, located to the right of the title bar, choose **Direct Internet Access (DIA) Sites**.
3. In the Manage DIA screen, click **Attach DIA Sites**. The Attach DIA Sites popup window displays all sites in your overlay network with available sites highlighted. For a site to be available, all devices at that site must be running in vManage mode.
4. In the Available Sites pane, select a site to attach and click the right arrow. To remove a site, select it in the Selected Sites pane and click the left arrow.
5. If you do not specify interfaces for Cloud OnRamp for SaaS to use, the system will select a NAT-enabled physical interface from VPN 0. If you would like to specify GRE interfaces for Cloud OnRamp for SaaS to use:
   a. Click the link, **Add interfaces** to selected sites (optional), located in the bottom right corner of the window.
   b. In the **Select Interfaces** drop-down, choose GRE interfaces to add.
   c. Click **Save Changes**.
6. Click **Attach**. vManage NMS pushes the feature template configuration to the devices. The Task View window displays a Validation Success message.
7. To return to the Cloud OnRamp for SaaS Dashboard, choose **Configuration** > **Cloud onRamp for SaaS**.

To edit Cloud onRamp interfaces on DIA sites:

1. Select the sites you want to edit and click Edit DIA Sites.
2. In the Edit Interfaces of Selected Sites screen, select a site to edit.
   - To add interfaces, click the **Interfaces** field to select available interfaces.
   - To remove an interface, click the **X** beside its name.
3. Click **Save Changes** to push the new template to the Cisco IOS XE Catalyst SD-WAN device s.

You have now completed configuring the Cloud OnRamp for SaaS. To return to the Cloud OnRamp for SaaS Dashboard, choose the **Configuration** > **Cloud onRamp for SaaS** screen.

Monitor Performance of Cloud OnRamp for SaaS

View Application Performance

In vManage NMS, select the **Configuration** > **Cloud OnRamp for SaaS** screen. The Cloud OnRamp for SaaS Dashboard displays the performance of each cloud application in a separate pane.

Each application pane displays the number of Cisco IOS XE Catalyst SD-WAN devices accessing the application and the quality of the connection:

- The bottom status bar displays green for devices experiencing good quality.
- The middle status bar displays yellow for devices experiencing average quality.
- The top status bar displays red for devices experiencing bad quality.

The number to the right of each status bar indicates how many devices are experiencing that quality of connection.

View Application Details

1. In vManage NMS, choose the **Configuration** > **Cloud OnRamp for SaaS** screen. The Cloud OnRamp for SaaS Dashboard displays each cloud application in a separate pane.
2. Click in an application's pane. vManage NMS displays a list of sites accessing the application.
3. Click a graph icon in the vQoE Score column to display vQoE history for that site:
    - Click a predefined or custom time period for which to display data.
    - Hover over a point on the chart to display vQoE details for that point in time.

Cloud OnRamp for Colocation Solution Overview

As more applications move to the cloud, the traditional approach of backhauling traffic over expensive WAN circuits to a data center is no longer relevant. The conventional WAN infrastructure was not designed for accessing applications in the cloud. The infrastructure is expensive and introduces unnecessary latency that degrades the experience.

Network architects are reevaluating the design of the WANs to achieve the following:

- Support a cloud transition.
- Reduce network costs.
- Increase the visibility and manageability of the cloud traffic.

The architects are turning to Software-Defined WAN (SD-WAN) fabric to take advantage of inexpensive broadband Internet services and to route intelligently a trusted SaaS cloud-bound traffic directly from remote branches.

With colocation solution built specifically for colocation facilities, the traffic is routed to the best-permissible path from branches and remote workers to where those applications are hosted. The solution also allows

distributed enterprises to have an alternative to enabling direct internet access at the branch and enhance their connectivity to infrastructure-as-a-service (IaaS) and software-as-a-service (SaaS) providers.

The solution provides enterprises with multiple distributed branch offices that are clustered around major cities or spread over several countries the ability to regionalize the routing services in colocation facilities. Reason being, these facilities are physically closer to the branches and can host the cloud resources that the enterprise needs to access. So, essentially by distributing a virtual Cisco SD-WAN over a regional architecture of colocation centers, the processing power is brought to the cloud edge.

The following image shows how you can aggregate the access to the multicloud applications from multiple branches to regional colocation facilities.

Figure 1: Cisco SD-WAN Cloud onRamp for CoLocations

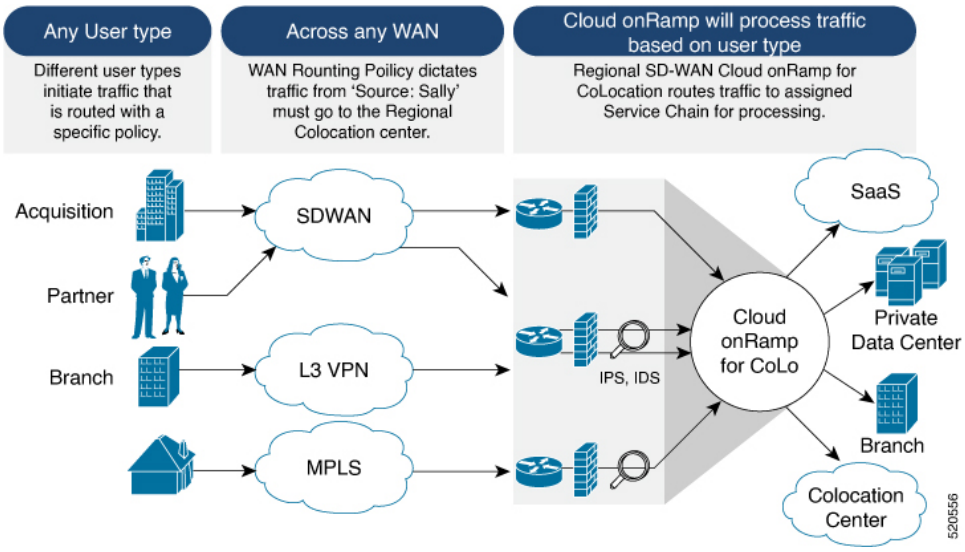

The solution can serve four specific types of enterprises:

- Multinational companies that cannot use direct internet connections to the cloud and SaaS platforms due to security restrictions and privacy regulations.
- Partners and vendors without Cisco SD-WAN but still need connectivity to their customers. They do not want to install SD-WAN routing appliances in their site.
- Global organizations with geographically distributed branch offices that require high bandwidth, optimum application performance, and granular security.
- Remote access that need secure VPN connections to an enterprise over inexpensive direct internet links.

The colocation solution can be hosted within certain colocation facilities by a colocation IaaS provider. You can select the colocation provider that meets your needs in a region on a regional basis as long as it supports the necessary components.

Manage Clusters

Use the Cloud OnRamp for Colocation screen to configure a colocation cluster and service groups that can be used with the cluster.

The three steps to configure are:

- Create a cluster.
- Create a service group.
- Attach a cluster with a service group.

A colocation cluster is a collection of two to eight CSP devices and two switches. The supported cluster templates are:

- Small cluster—2 Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C+2 CSP
- Medium Cluster—2 Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C+4 CSP
- Large Cluster—2 Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C+6 CSP
- X-Large Cluster—2 Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C+8 CSP

Note Ensure that you add a minimum of two CSP devices one-by-one to a cluster. You can keep adding three, four, and so on, up to a maximum of eight CSP devices. You can edit a Day-N configuration of any cluster, and add pairs of CSP devices to each site up to a maximum of eight CSP devices.

Ensure that all devices that you bring into a cluster have the same software version.

Note You can’t use the CSP-5444 and CSP-5456 devices in the same cluster.

Following are the cluster states:

- Incomplete—When a cluster is created from the Cisco SD-WAN Manager interface without providing the minimum requirement of two CSP devices and two switches. Also, cluster activation is not yet triggered.
- Inactive—When a cluster is created from the Cisco SD-WAN Manager interface after providing the minimum requirement of two CSP devices and two Switches, and cluster activation is not yet triggered.
- Init—When the cluster activation is triggered from the Cisco SD-WAN Manager interface and Day-0 configuration push to the end devices is pending.
- Inprogress—When one of the CSP devices within a cluster comes up with control connections, the cluster moves to this state.
- Pending—When the Day-0 configuration push is pending or VNF install is pending.
- Active—When a cluster is activated successfully and NCS has pushed the configuration to the end device.
- Failure—If Cisco Colo Manager has not been brought up or if any of the CSP devices that failed to receive an UP event.

A cluster transitioning to an active state or failure state is as follows:

- **Inactive** > **Init** > **Inprogress** > **Pending** > **Active**—Success
- **Inactive** > **Init** > **Inprogress** > **Pending** > **Failure**—Failure

Provision and Configure Cluster

This topic describes about activating a cluster that enables deployment of service chains.

To provision and configure a cluster, perform the following:

1. Create a colocation cluster by adding two to eight CSP devices and two switches.

   CSP devices can be added to a cluster and configured using Cisco SD-WAN Manager before bringing them up. You can configure CSP devices and Catalyst 9K switches with the global features such as, AAA, default user (admin) password, NTP, syslog, and more.

2. Configure colocation cluster parameters including IP address pool input such as, service chain VLAN pool, VNF management IP address pool, management gateway, VNF data plane IP pool, and system IP address pool.

3. Configure a service group.

   A service group consists of one or more service chains.

Note You can add a service chain by selecting one of the predefined or validated service chain template, or create a custom one. For each service chain, configure input and output VLAN handoff and service chain throughput or bandwidth, as mentioned.

4. Configure each service chain by selecting each VNF from the service template. Choose a VNF image that is already uploaded to the VNF repository to bring up the VM along with required resources (CPU, memory, and disk). Provide the following information for each VNF in a service chain:

   - The specific VM instance behavior such as, HA, shared VM can be shared across service chains.
   - Day-0 configuration values for tokenized keys and not part of the VLAN pool, management IP address, or data HA IP address. The first and last VMs handoff-related information such as peering IP and autonomous system values must be provided. The internal parameters of a service chain are automatically updated by Cisco SD-WAN Validator from the VLAN, or Management, or Data Plane IP address pool provided.

5. Add the required number of service chains for each service group and create the required number of service groups for a cluster.

6. To attach a cluster to a site or location, activate the cluster after all configuration is complete.

   You can watch the cluster status change from In progress to active or error in the **Task View** window.

To edit a cluster:

1. Modify the activated cluster by adding or deleting service groups or service chains.

2. Modify the global features configuration such as, AAA, system setting, and more.

You can predesign a service group and service chain before creating a cluster. You can then attach the service group with a cluster after the cluster is active.

Create and Activate Clusters

This topic provides the steps on how you can form a cluster with CSP devices, Cisco Catalyst switches as a single unit, and provision the cluster with cluster-specific configuration.

Before you begin

- Ensure that you synchronize the clocks for Cisco SD-WAN Manager and CSP devices. To synchronize a clock for CSP devices, configure the NTP server for CSP devices when you enter information about cluster settings.
- Ensure that you configure the NTP server for Cisco SD-WAN Manager and Cisco SD-WAN Validator. To configure the NTP server.
- Ensure that you configure the OTP for the CSP devices to bring up the CSP devices.
- Ensure that you power on both the Catalyst 9500 switches and ensure that they are operational.

Procedure

Step 1 From the Cisco SD-WAN Manager menu, choose Cisco SD-WAN Manager, click **Configuration** > **Cloud OnRamp for Colocation**.

a) Click **Configure & Provision Cluster**.

b) Provide the following information:

Table 1: Cluster Information

Field	Description
Cluster Name	The cluster name can contain 128 alphanumeric characters.
Description	The description can contain 2048 alphanumeric characters.
Site ID	The overlay network site identifier. Ensure that the value you enter for Site ID is similar to the organizations Site ID structure for the other Cisco Catalyst SD-WAN overlay elements.
Location	The location can contain 128 alphanumeric characters.
Cluster Type	To configure a cluster in a multitenant mode so that it can be shared across multiple tenants, choose **Shared**. **Note** In the single-tenant mode, the cluster type **Non Shared** is selected by default.

c) To configure switches, click a switch icon in the **Switches** box. In the **Edit Switch** dialog box, enter a switch name and choose the switch serial number from the drop-down list. Click **Save**.

The switch name can contain128 alphanumeric characters.

The switch serial numbers that you view in the drop-down list are obtained and integrated with Cisco SD-WAN Manager using the PnP process. These serial numbers are assigned to switches when you order Cisco Catalyst SD-WAN Cloud OnRamp for Colocation solution PID on the CCW and procure the switch devices.

Note
You can keep the serial number field blank for switch devices and CSP devices, design your colocation cluster, and then edit the cluster later to add the serial number after you procure the devices. However, you can't activate a cluster with the CSP devices or switch devices without the serial numbers.

d) To configure another switch, repeat Step c.
e) To configure CSP devices, click a CSP icon in the **Appliances** box. The **Edit CSP** dialog box is displayed. Provide a CSP device name and choose the CSP serial number from the drop-down list. Click **Save**.

The CSP device name can contain 128 alphanumeric characters.

f) Configure OTP for the CSP devices to bring up the devices.
g) To add remaining CSP devices, repeat Step e.
h) Click **Save**.
After you create a cluster, on the cluster configuration window, an ellipsis enclosed in a yellow circle appears next to a device where the serial number isn't assigned for the device. You can edit a device to enter the serial numbers.
i) To edit a CSP device configuration, click a CSP icon, and perform the process mentioned in substep e.
j) To set the mandatory and optional global parameters for a cluster, on the cluster configuration page, enter the parameters for **Cluster Configuration**.
k) Click **Save**.

You can view the cluster that you created in a table on the cluster configuration page.

Step 2 To activate a cluster,

a) Click a cluster from the cluster table.
b) For the desired cluster, click **…** and choose **Activate**.

When you activate the cluster, Cisco SD-WAN Manager establishes a DTLS tunnel with the CSP devices in the cluster, where it connects with the switches through Cisco Colo Manager. When the DTLS tunnel connection is running, a CSP device in the cluster is chosen to host the Cisco Colo Manager. Cisco Colo Manager starts up and Cisco SD-WAN Manager sends global parameter configurations to the CSP devices and Cisco Catalyst 9500 switches. For information about cluster activation progress, see Progress of Cluster Activation. .

Note In Cisco vManage Release 20.7.1 and earlier releases, the Cisco Colo Manager (CCM) and CSP device configuration tasks time out 30 minutes after the tasks are created. In the case of long-running image installation operations, these configuration tasks may time out and fail, while the cluster activation state continues to be in a pending state.

From Cisco vManage Release 20.8.1, the CCM and CSP device configuration tasks time out 30 minutes after the last heartbeat status message that Cisco SD-WAN Manager received from the target devices. With this change, long-running image installation operations do not cause configuration tasks to fail after a predefined interval of time after task creation.

Cluster Configuration

The cluster configuration parameters are:

View Cluster

To view cluster configuration, perform the following steps:

Procedure

Step 1 From the Cisco SD-WAN Manager menu, choose **Configuration** > **Cloud OnRamp for Colocation**.

Step 2 For the desired cluster, click **...** and choose **View**.

The **Cluster** window displays the switch devices and CSP devices in the cluster and shows the cluster settings that are configured.

You can only view the global parameters of a cluster, configuration of switch devices and CSP devices.

Step 3 Click **Cancel** to return to the **Cluster** window.

Edit Cluster

To modify any existing cluster configuration such as global parameters, perform the following steps:

Procedure

Step 1 From the Cisco SD-WAN Manager menu, choose **Configuration** > **Cloud OnRamp for Colocation**

Step 2 For the desired cluster, click **...** and choose **Edit**.

The **Cluster** window displays the switch devices and CSP devices in the cluster and shows the cluster settings that are configured.

Step 3 In the cluster design window, you can modify some of the global parameters. Based on whether a cluster is in active or inactive state, you can perform the following operations on a cluster:

a. Inactive state:

- Edit all global parameters, and the Resource Pool parameter.
- Add more CSP devices (up to eight).
- Can’t edit the name or serial number of a switch or CSP device. Instead, delete the CSP or switch and add another switch or CSP with a different name and serial number.
- Delete an entire cluster configuration.

b. Active state:

- Edit all global parameters, except the Resource Pool parameter.

Note

You can't change the Resource pool parameter when the cluster is active. However, the only option to change the Resource Pool parameter is to delete the cluster and recreate it with the correct Resource Pool parameter.

- Can’t edit the name or serial number of a switch or CSP device.
- Can’t delete a cluster in an active state.
- Add more CSP devices (up to eight).

Step 4 Click **Save Cluster**.

Remove Cluster

To decommission an entire cluster from Cisco SD-WAN Manager, perform the following steps:

Procedure

Step 1 From the Cisco SD-WAN Manager menu, choose **Configuration** > **Certificates**.

Step 2 Verify the **Validate** column for the CSP devices that you wish to delete, and click **Invalid**.

Step 3 For the invalid devices, click **Send to Controllers**.

Step 4 From the Cisco SD-WAN Manager menu, choose **Configuration** > **Cloud OnRamp for Colocation**.

Step 5 For the cluster that has invalid CSP devices, click **...** and choose **Deactivate**.

If the cluster is attached to one or more service groups, a message appears that displays the service chains hosting the VMs that are running on the CSP device and whether you can continue with the cluster deletion. However, although you confirm the deletion of a cluster, you’re not allowed to remove the cluster without detaching the service groups that are hosted on this CSP device. If the cluster isn’t attached to any service group, a message appears that gets a confirmation from you about the cluster deletion.

Note
You can delete the cluster, if necessary, or can keep it in deactivated state.

Step 6 To delete the cluster, choose **Delete**.

Step 7 Click **Cancel** if you don't wish to delete the cluster.

Step 8 To decommission invalid devices, from the Cisco SD-WAN Manager menu, choose **Configuration** > **Devices**.

Step 9 For the devices that are in the deactivated cluster, click **...** and choose **Decommission WAN Edge**.

This action provides new tokens to your devices.

Step 10 Reset the devices to the factory default by using the command:

factory-default-reset all

Step 11 Log into Cisco NFVIS by using **admin** as the login name and **Admin123#** as the default password.

Step 12 Reset switch configuration and reboot switches. See the Troubleshooting chapter in Cisco Catalyst SD-WAN Cloud OnRamp for Colocation Solution Guide.

Reactivate Cluster

To add new CSP devices or when CSP devices are considered for RMA process, perform the following steps:

Procedure

Step 1 From the Cisco SD-WAN Manager menu, choose **Configuration** > **Devices**.

Step 2 Locate the devices that are in a deactivated cluster.

Step 3 Get new token from Cisco SD-WAN Manager for the devices.

Step 4 Log into Cisco NFVIS using **admin** as the login name and **Admin123#** as the default password.

Step 5 Use the **request activate chassis-number** *chassis-serial-number* **token** *token-number* command.

Step 6 Use Cisco SD-WAN Manager to configure the colocation devices and activate the cluster. See Create and Activate Clusters, on page 29.

If you've deleted the cluster, recreate and then activate it.

Step 7 From the Cisco SD-WAN Manager menu, choose **Configuration** > **Certificates**. Locate and verify status of the colocation devices.

Step 8 For the desired device that should be valid, click **Valid**.

Step 9 For the valid devices, click **Send to Controllers**.

Create Service Chain in a Service Group

A service group consists of one or more service chains.

Table 2: Feature History

Feature Name	Release Information	Feature Description
Monitor Service Chain Health	Cisco IOS XE Catalyst SD-WAN Release 16.12.1b	This feature lets you configure periodic checks on the service chain data path and reports the overall status. To enable service chain health monitoring, NFVIS version 3.12.1 or later should be installed on all CSP devices in a cluster.

Procedure

From the Cisco SD-WAN Manager menu, choose **Configuration** > **Cloud OnRamp for Colocation**

a) Click **Service Group** and click **Create Service Group**. Enter the service group name, description, and colocation group.

The service group name can contain 128 alphanumeric characters.

The service group description can contain 2048 alphanumeric characters.

For a multitenant cluster, choose a colocation group or a tenant from the drop-down list. For a single-tenant cluster, the colocation group **admin** is chosen by default.

b) Click **Add Service Chain**.

c) In the **Add Service Chain** dialog box, enter the following information:

Table 3: Add Service Chain Information

Field	Description
Name	The service chain name can contain 128 alphanumeric characters.
Description	The service chain description can contain alphanumeric 2048 characters.
Bandwidth	The service chain bandwidth is in Mbps. The default bandwidth is 10 Mbps and you can configure a maximum bandwidth of 5 Gbps.
Input Handoff VLANS and Output Handoff VLANS	The Input VLAN handoff and output VLAN handoff can be comma-separated values (10, 20), or a range from 10–20.
Monitoring	A toggle button that allows you to enable or disable service chain health monitoring. The service chain health monitoring is a periodic monitoring service that checks health of a service chain data path and reports the overall service chain health status. By default, the monitoring service is disabled. A service chain with subinterfaces such as, SCHM (Service Chain Health Monitoring Service) can only monitor the service chain including the first VLAN from the subinterface VLAN list. The service chain monitoring reports status based on end-to-end connectivity. Therefore, ensure that you take care of the routing and return traffic path, with attention to the Cisco Catalyst SD-WAN service chains for better results. **Note** • Ensure that you provide input and output monitoring IP addresses from input and output handoff subnets. However, if the first and last VNF devices are VPN terminated, you don't need to provide input and output monitoring IP addresses. For example, if the network function isn't VPN terminated, the input monitoring IP can be 192.0.2.1/24 from the inbound subnet, 192.0.2.0/24. The inbound subnet connects to the first network function and the output monitoring IP can be, 203.0.113.11/24 that comes from outbound subnet, 203.0.113.0/24 of the last network function of a service chain. • If the first or last VNF firewall in a service chain is in transparent mode, you can't monitor these service chains.
Service Chain	A topology to choose from the service chain drop-down list. For a service chain topology, you can choose any of the validated service chains such as, Router - Firewall - Router, Firewall, Firewall - Router. See the Validated Service Chains topic in Cisco Catalyst SD-WAN Cloud OnRamp Colocation Solution Guide. You can also create a customized service chain. See Create Custom Service Chain, on page 39.

d) In the **Add Service Chain** dialog box, click **Add**.
Based on the service chain configuration information, a graphical representation of the service group with all the service chains and the VNFs automatically appear in the design view window. A VNF or PNF appears with a "V" or "P" around the circumference for a virtual a physical network function. It shows all the configured service chains within each service group. A check mark next to the service chain indicates that the service chain configuration is complete.

After you activate a cluster, attach it with the service group and enable monitoring service for the service chain, when you bring up the CSP device where CCM is running. Cisco SD-WAN Manager chooses the same CSP device to start the monitoring service. The monitoring service monitors all service chains periodically in a round robin fashion by setting the monitoring interval to 30 minutes. See Monitor Cloud OnRamp Colocation Clusters, on page 55.

e) In the design view window, to configure a VNF, click a VNF in the service chain.
The **Configure VNF** dialog box appears.

f) Configure the VNF with the following information and perform the actions, as appropriate:

Note
The following fields are available from Cisco vManage Release 20.7.1:

- **Disk Image/Image Package (Select File)**
- **Disk Image/Image Package (Filter by Tag, Name and Version)**
- **Scaffold File (Select File)**
- **Scaffold File (Filter by Tag, Name and Version)**

Table 4: VNF Properties of Router and Firewall

Field	Description
Image Package	Choose a router, firewall package.
Disk Image/Image Package (Select File)	Choose a tar.gz package or a qcow2 image file.
Disk Image/Image Package (Filter by Tag, Name and Version)	(Optional) Filter an image or a package file based on the name, version, and tags that you specified when uploading a VNF image.
Scaffold File (Select File)	Choose a scaffold file. **Note** • This field is mandatory if a qcow2 image file has been chosen. It is optional if a tar.gz package has been chosen. • If you choose both a tar.gz package and a scaffold file, then all image properties and system properties from the scaffold file override the image properties and system properties, including the Day-0 configuration files, specified in the tar.gz package.

Field	Description
Scaffold File (Filter by Tag, Name and Version)	(Optional) Filter a scaffold file based on the name, version, and tags that you specified when uploading a VNF image.
Click **Fetch VNF Properties**. The available information for the image is displayed in the **Configure VNF** dialog box.	
Name	VNF image name
CPU	(Optional) Specifies the number of virtual CPUs that are required for a VNF. The default value is 1 vCPU.
Memory	(Optional) Specifies the maximum primary memory in MB that the VNF can use. The default value is 1024 MB.
Disk	(Optional) Specifies disk in GB required for the VM. The default value is 8 GB.
A dialog box with any custom tokenized variables from Day-0 that requires your input appears. Provide the values.	

In the following image, all IP addresses, VLAN, and autonomous system within the green box are system-specific information that is generated from the VLAN, IP pools provided for the cluster. The information is automatically added into the Day-0 configurations of VMs.

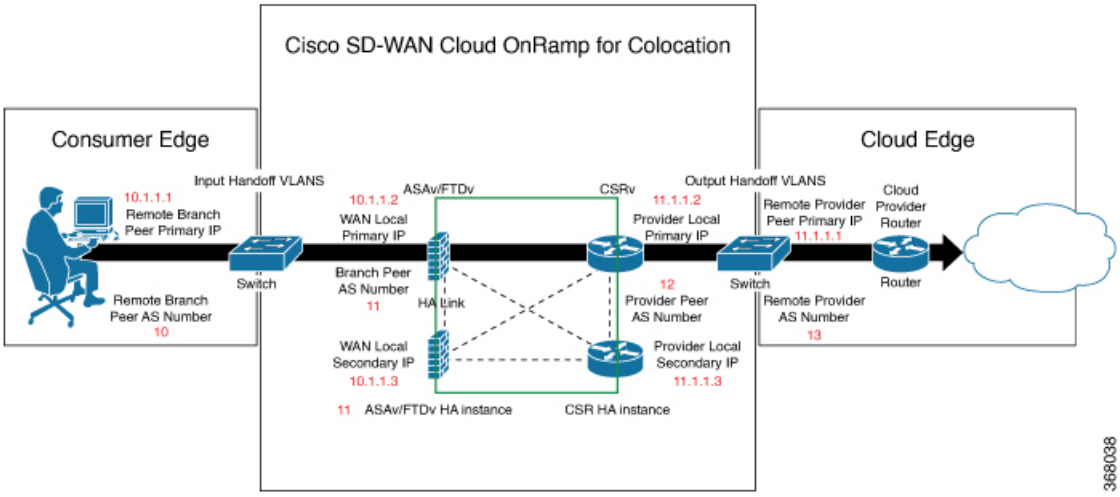

The following images are a sample configuration for VNF IP addresses and autonomous system numbers, in Cisco SD-WAN Manager.

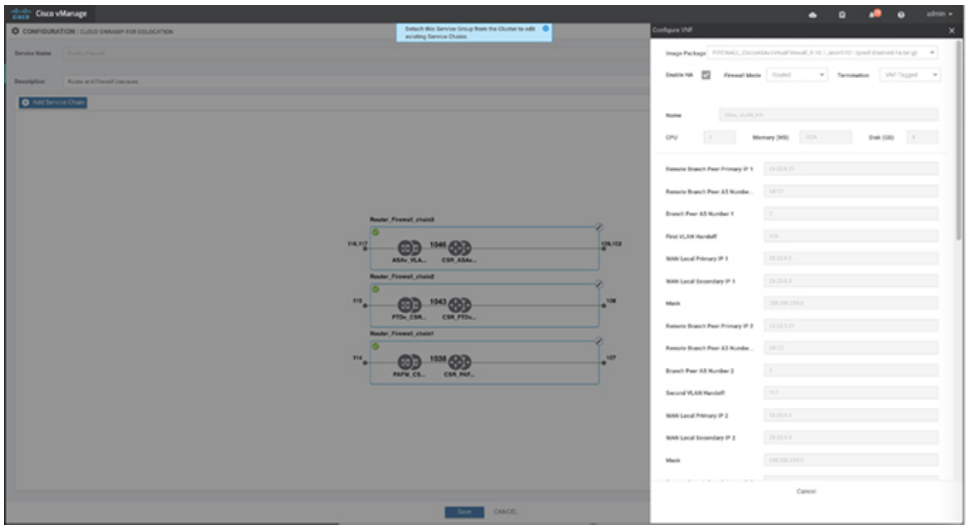

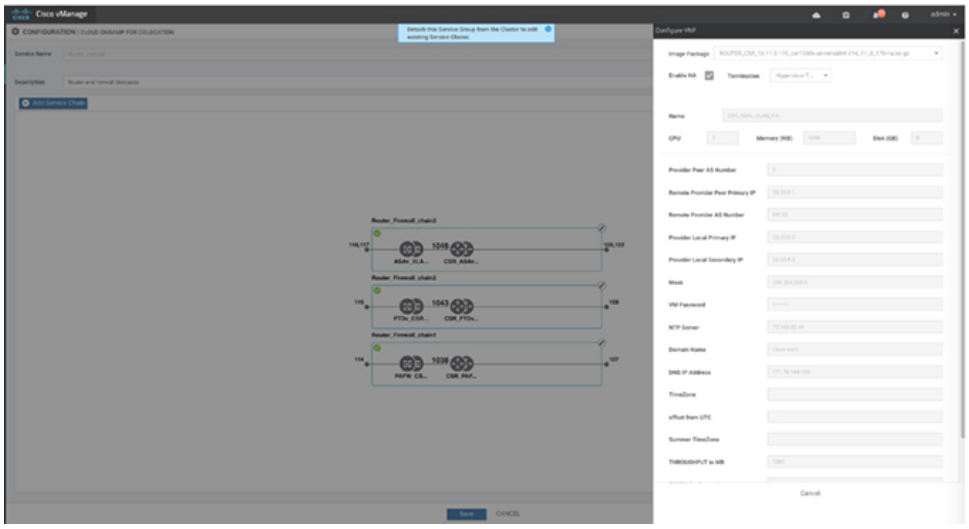

If you're using a multitenant cluster and a comanged scenario, configure theCisco Catalyst SD-WAN VM by entering the values for the following fields and the remaining fields, as required for the service chain design:

Note
To join the tenant overlay network, the provider should provide correct values for the following fields.

Field	Description
Serial Number	The authorized serial number of a Cisco Catalyst SD-WAN device. The service provider can get the device serial number from the tenant before creating the service chain.
OTP	The OTP of the Cisco Catalyst SD-WAN device that is available after authenticating it with Cisco SD-WAN Control Components. The service provider can get the OTP for the corresponding serial number from the tenant before creating the service chain.

Field	Description
Site Id	The identifier of the site in the tenant Cisco Catalyst SD-WAN overlay network domain in which the Cisco Catalyst SD-WAN device resides, such as a branch, campus, or data center. The service provider can get the site Id from the tenant before creating the service chain.
Tenant ORG Name	The tenant organization name that is included in the Certificate Signing Request (CSR). The service provider can get the organization name from the tenant before creating the service chain.
System IP connect to Tenant	The IP address to connect to the tenant overlay network. The service provider can get the IP address from the tenant before creating the service chain.
Tenant vBond IP	The IP address of the tenant Cisco SD-WAN Validator. The service provider can get the Cisco SD-WAN Validator IP address from the tenant before creating the service chain.

For edge VMs such as first and last VM in a service chain, you must provide the following addresses as they peer with a branch router and the provider router.

Table 5: VNF Options for First VM in Service Chain

Field	Mandatory or Optional	Description
Firewall Mode	Mandatory	Choose Routed or Transparent mode. **Note** Firewall mode is applicable to firewall VMs only.
Enable HA	Optional	Enable HA mode for the VNF.

Termination	Mandatory	Choose one of the following modes:
		• L3 mode selection with subinterfaces that are in trunk mode
		`<type>selection</type> <val help="L3 Mode With Sub-interfaces(Trunked)" display="VNF-Tagged">vlan</val>`
		• L3 mode with IPSEC termination from a consumer-side and rerouted to the provider gateway
		`<val help="L3 Mode With IPSEC Termination From Consumer and Routed to Provider GW" display="Tunneled">vpn</val>`
		• L3 mode with access mode (nontrunk mode)
		`<val help="L3 Mode In Access Mode (Non-Trunked)" display="Hypervisor-Tagged">routed</val>`

g) Click **Configure**. The service chain is configured with the VNF configuration.
h) To add another service chain, repeat the procedure from Steps b-g.
i) Click **Save**.

The new service group appears in a table under the **Service Group**. To view the status of the service chains that are monitored, use the **Task View** window, which displays a list of all running tasks along with the total number of successes and failures. To determine the service chain health status, use the **show system:system status** command on the CSP device that has service chain health monitoring enabled.

Create Custom Service Chain

You can customize service chains,

- By including extra VNFs or add other VNF types.
- By creating new VNF sequence that isn't part of the predefined service chains.

Procedure

Step 1 Create a service group and service chains within the service group. See Create Service Chain in a Service Group, on page 33.

Step 2 In the **Add Service Chain** dialog box, enter the service chain name, description, bandwidth, input VLAN handoff, output VLAN handoff, monitoring health information of a service chain, and service chain configuration. Click **Add**.

For the service chain configuration, choose **Create Custom** from the drop-down. An empty service chain in the design view window is available.

Step 3 To add a VNF such as a router, load balancer, firewall, and others, click a VNF icon and drag the icon to its proper location within the service group box. After adding all required VNFs and forming the VNF service chain, configure each of the VNFs. Click a VNF in the service group box. The **Configure VNF** dialog box appears. Enter the following parameters:

a) Choose the software image to load from the **Disk Image/Image Package** (**Select File**) drop-down list.

Note
You can select a qcow2 image file from Cisco vManage Release 20.7.1.

b) Choose a scaffold file from the **Scaffold File** (**Select File**) drop-down list if you have chosen a qcow2 image file.

Note
This option is available from Cisco vManage Release 20.7.1.

c) Optionally, filter an image, a package file, or a scaffold file based on the name, version, and tags that you specified when uploading a VNF image.

Note
This option is available from Cisco vManage Release 20.7.1.

d) Click **Fetch VNF Properties**.
e) In the **Name** field, enter a name of the VNF.
f) In the **CPU** field, enter the number of virtual CPUs required for the VNF.
g) In the **Memory** field, enter the amount of memory in megabytes to be allocated for the VNF.
h) In the **Disk** field, enter the amount of memory for storage in gigabytes to be allocated for the VNF.
i) Enter VNF-specific parameters, as required.

Note

These VNF details are the custom variables that are required for Day-0 operations of the VNF.

j) Click **Configure**.
k) To delete the VNF or cancel the VNF configuration, click **Delete** or **Cancel** respectively.

The customized service chains are added to a service group.

Note You can customize a VNF sequence with only up to four VNFs in a service chain.

Custom Service Chain with Shared PNF Devices

You can customize service chains by adding supported PNF devices.

Caution Ensure that you don't share PNF devices across colocation clusters. A PNF device can be shared across service chains, or across service groups. However, a PNF device can now be shared only across a single cluster.

Table 6: Feature History

Feature Name	Release Information	Feature Description
Manage PNF Devices in Service Chains	Cisco IOS XE Catalyst SD-WAN Release 16.12.1b	This feature lets you add Physical Network Function (PNF) devices to a network, in addition to the Virtual Network function (VNF) devices. These PNF devices can be added to service chains and shared across service chains, service groups, and a cluster. Inclusion of PNF devices in the service chain can overcome the performance and scaling issues caused by using only VNF devices in a service chain.

Before you begin

For more information about validated physical network functions, see the Validated Physical Network Functions topic in the Cisco Catalyst SD-WAN Cloud OnRamp for Colocation Solution Guide.

To create a customized service chain by adding a router or firewall to an existing service chain, ensure that you note the following points:

- If a PNF device needs to be managed byCisco SD-WAN Manager, ensure that the serial number is already available in Cisco SD-WAN Manager, which can then be available for selection during PNF configuration.
- The FTD device can be in any position in a service chain.
- An ASR 1000 Series Aggregation Services Routers can only be in the first and last position in a service chain.
- PNF devices can be added across service chains and service groups.
- PNF devices can be shared across service groups. They can be shared across service groups by entering the same serial numbers.

- PNF devices can be shared across a single colocation cluster, and can't be shared across multiple colocation clusters.

Procedure

Step 1 Create a service group and service chains within the service group. See Create Service Chain in a Service Group, on page 33.

Step 2 In the **Add Service Chain** dialog box, enter the service chain name, description, bandwidth, input VLAN handoff, output VLAN handoff, monitoring health information of a service chain, and service chain configuration. Click **Add**.

For the service chain configuration, choose **Create Custom** from the drop-down list. An empty service chain in the design view window is available. At the left, a set of VNF devices and PNF devices that you can add into the service chain appears. The 'V' in the circumference of VNF devices represents a VNF and 'P' in the circumference of PNF devices represent a PNF.

Note
Ensure that you choose the **Create Custom** option for creating service chains by sharing PNF devices.

Step 3 To add a PNF such as physical routers, physical firewalls in a service chain, click the required PNF icon, and drag the icon to the proper location within the service chain box.

After adding all required PNF devices, configure each of them.

a) Click a PNF device in the service chain box.

The **Configure PNF** dialog box appears. To configure a PNF, enter the following parameters:

b) Check **HA Enabled** if HA is enabled for the PNF device.

c) If the PNF is HA enabled, ensure that you add the HA serial number in **HA Serial**.

If the PNF device is FTD, enter the following information.

1. In the **Name** field, enter a name of the PNF.
2. Choose Routed or Transparent mode as the **Firewall Mode**.
3. In the **PNF Serial** field, enter the serial number of the PNF device.

If the PNF device is ASR 1000 Series Aggregation Services Routers, enter the following information.

1. Check the **vManaged** check box if the device is managed by Cisco SD-WAN Manager.
2. Click **Fetch Properties**.
3. In the **Name** field, enter a name of the PNF.
4. In the **PNF Serial** field, enter the serial number of the PNF device.

d) Click **Configure**.

Step 4 To add service chains and share PNF devices, repeat from Step 2.

Step 5 To edit an existing PNF configuration, click the PNF.

Step 6 In the **Share NF To** drop-down list, choose the service chains with which the PNF should be shared.

After a PNF is shared, if you hover over a PNF, the respective shared PNF devices are highlighted in blue color. However, the PNFs from different service groups aren't highlighted in blue color. After you choose an NF to be shared, a blue color rim appears. If the same PNF is shared across multiple service chains, it can be used in different positions by dragging and placing the PNF icons in a specific position.

Figure 2: Single PNF in a Service Chain

The following image shows a service chain that consists of a single PNF, Ftd_Pnf (not shared with other service chains).

Figure 3: Two PNF Devices in Service Chains

The following image shows service chains that consist of two PNFs, FTdv_PNF shared across service chain 1 (SC1) and service chain 2 (SC2) and ASR_PNF (non-shared).

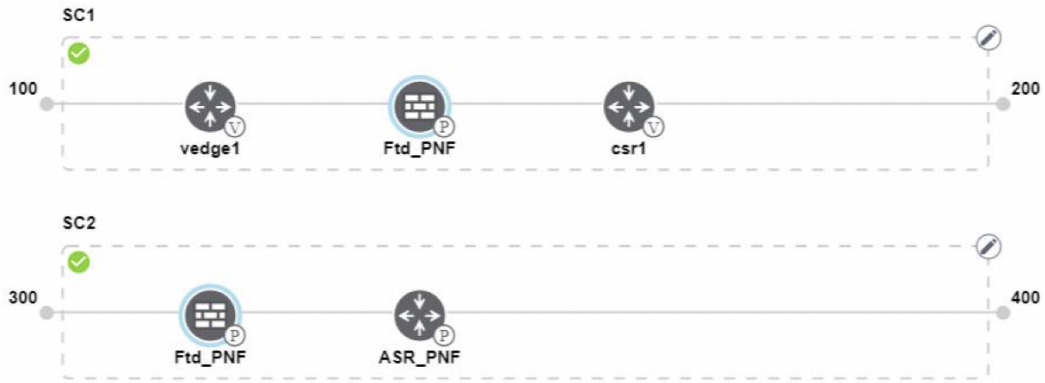

Figure 4: Three PNF Devices in Service Chains

The following image shows service chains that consist of three PNF devices in two different positions along with Cisco SD-WAN Manager configuration.

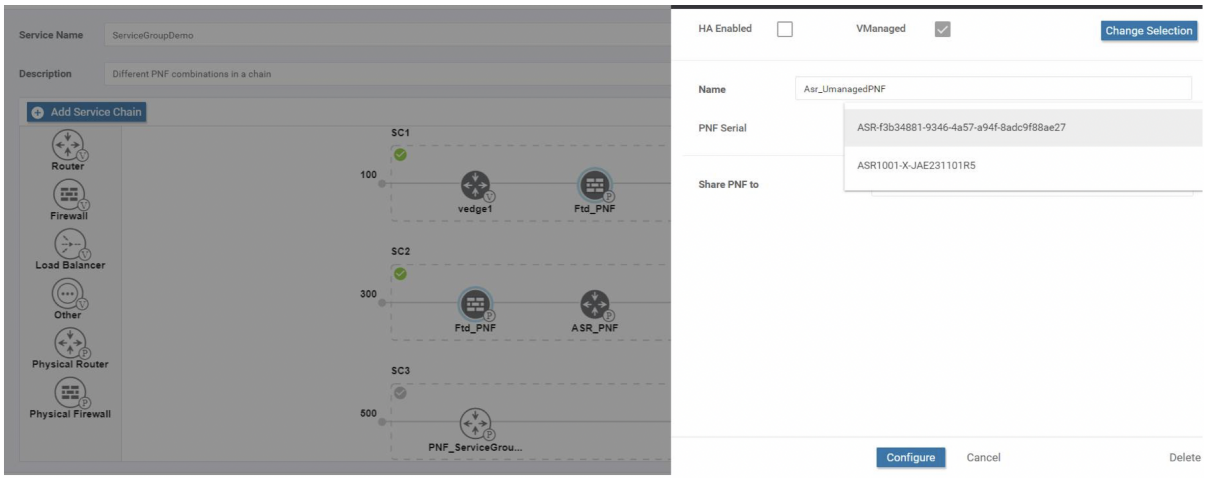

Step 7 To delete or cancel a Network Function configuration, click **Delete** or **Cancel** respectively.

You must attach the service groups to a colocation cluster. After attaching service groups that contain PNF devices, the PNF configuration isn't automatically pushed to the PNF devices unlike VNF devices. Instead, you must manually configure the PNF device by noting configuration that is generated on the Monitor window. The VLANs must be also configured on the Cisco Catalyst 9500-40X switch devices. See the ASR 1000 Series Aggregation Services Routers Configuration Guides and Cisco Firepower Threat Defense Configuration Guides for more information about the specific PNF configuration.

Configure PNF and Cisco Catalyst 9500 Switches

Procedure

Step 1 Identify ports from the switches where the PNF devices should be added, which are part of a service chain. To verify the availability of the ports, see "Service Chains and Port Connectivity Details" topic in Cisco SD-WAN Cloud OnRamp for Colocation Solution Guide.

Step 2 Connect with Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C by using either the terminal server of any of the Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switches or use the **vty session** command with the IP address of the active switch.

Step 3 Configure VLANs from the generated configuration parameters on Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switches with interfaces that are connected to the PNF. See the Monitor screen for the generated VLAN configuration.

Step 4 To configure an FTD or an ASR 1000 Series device, note the configuration from the **Monitor** window and then manually configure it on a device.

Custom Service Chain with Shared VNF Devices

You can customize service chains by including supported VNF devices.

Table 7: Feature History

Feature Name	Release Information	Feature Description
Share VNF Devices Across Service Chains	Cisco IOS XE Catalyst SD-WAN Release 16.12.1b	This feature lets you share Virtual Network Function (VNF) devices across service chains to improve resource utilisation and reduce resource fragmentation.

Before you begin

Ensure that you note the following points about sharing VNF devices:

- You can share only the first, last, or both first and last VNF devices in a service chain.
- You can share a VNF with a minimum of one more service chain and maximum up to five service chains.
- Each service chain can have a maximum of up to four VNF devices in a service chain.
- You can share VNF devices only in the same service group.

Procedure

Step 1 Create a service group and service chains within the service group. See Create Service Chain in a Service Group, on page 33.

Step 2 In the **Add Service Chain** dialog box, enter the service chain name, description, bandwidth, input VLAN handoff, output VLAN handoff, monitoring health information of a service chain, and service chain configuration. Click **Add**.

For the service chain configuration, choose **Create Custom** from the drop-down list. An empty service chain in the design view window is available. At the left, a set of VNF devices and PNF devices that you can add into the service chain appears. The 'V' in the circumference of VNF devices represents a VNF and 'P' in the circumference of PNF devices represent a PNF.

Note
Ensure that you choose the **Create Custom** option for creating a shared VNF package.

Step 3 To add a VNF such as a router, load balancer, firewall, and others, click a VNF icon from the left panel, and drag the icon to a proper location within the service chain box.

After adding all required VNF devices, configure each of them.

a) Click a VNF in the service chain box.

The **Configure VNF** dialog box appears. To configure VNF, enter the following parameters:

b) From the **Image Package** drop-down list, choose the software image to load.

To create a customized VNF package from Cisco SD-WAN Manager, see Create Customized VNF Image, on page 62.

c) Click **Fetch VNF Properties**.
d) In the **Name** field, enter a name of the VNF.
e) In the **CPU** field, enter the number of virtual CPUs required for the VNF.

f) In the **Memory** field, enter the amount of memory in megabytes to be allocated for the VNF.
g) In the **Disk** field, enter the amount of memory for storage in gigabytes to be allocated for the VNF.
h) Enter VNF-specific parameters, as required. See Create Service Chain in a Service Group, on page 33 for more information about VNF-specific properties.

These VNF-specific parameters are the custom user variables that are required for Day-0 operations of a VNF.

Note
Ensure that you enter the values of the user variables if they are defined as mandatory, and the system variables are automatically set by Cisco SD-WAN Manager.

i) Click **Configure**.

Step 4 To share VNF devices, repeat from Step 2.

Step 5 To edit an existing VNF configuration, click the VNF.

Step 6 Scroll down the VNF configuration to find the **Share NF To** field. From the **Share NF To** drop-down list, choose the service chains with which the VNF should be shared.

After a VNF is shared, if you hover over a VNF, the specific shared VNF devices are highlighted in blue color. After you choose an NF to be shared, a blue rim appears on it.

Step 7 To delete a VNF or cancel the VNF configuration, click **Delete** or **Cancel** respectively.

You must attach service groups to a cluster.

Shared VNF Use Cases

The following are the sample images for some of the shared VNF use cases and their predefined variable list:

Figure 5: Shared–Cisco vEdge Router VNF in First Position

The Cisco vEdge Router VNF in the first position is shared with the second service chain in the first position. The input to the first VNF is in access mode (hypervisor-tagged) and the neighbor (ASAv firewall) is in HA mode.

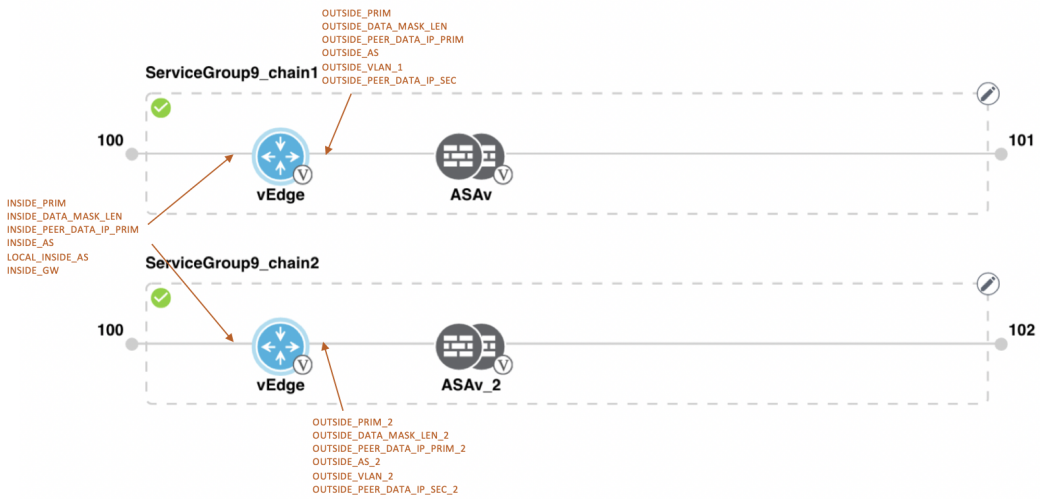

Figure 6: Shared–Cisco vEdge Router VNF in First Position

The Cisco vEdge Router VNF in the first position is shared with the second service chain in the first position. The input to the first VNF is in access mode (hypervisor-tagged) and the neighbor is in StandAlone mode.

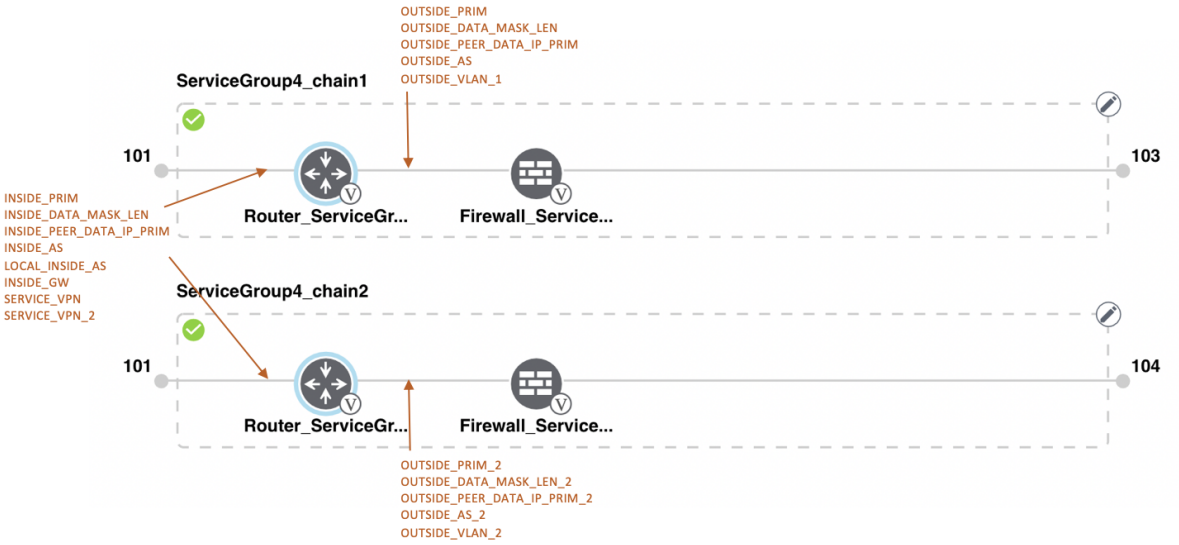

Figure 7: Shared–Cisco vEdge Router VNF in First Position

The Cisco vEdge Router VNF in the first position is shared with the second service chain in the first position. The input to the first VNF is in trunk mode (VNF-tagged) and the neighbor is in StandAlone mode.

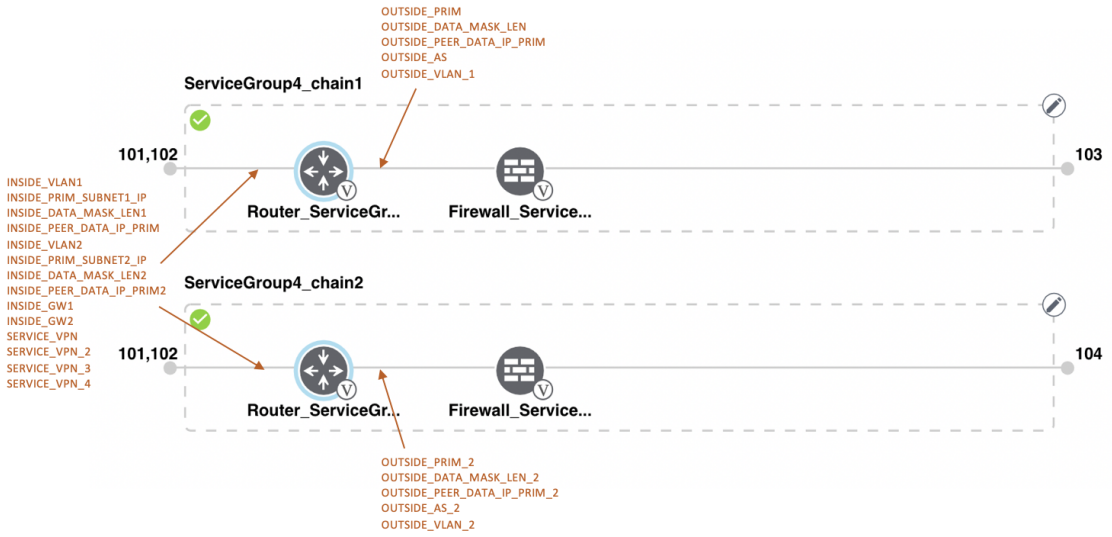

Figure 8: Shared–Cisco vEdge Router VNF in First Position

The Cisco vEdge Route VNF in the first position is shared with the second service chain in the first position. The input to the first VNF is in trunk mode (VNF-tagged) and the neighbor is in HA mode.

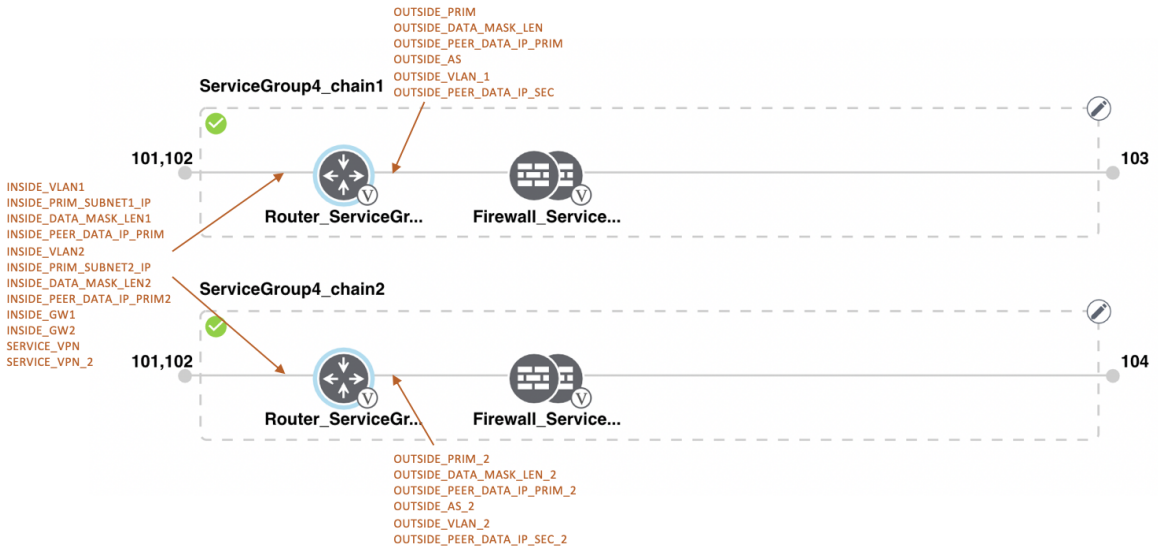

Figure 9: Shared-Cisco CSR1000V VNF in Last Position

The Cisco CSR1000V VNF in the last position is shared with the second service chain in the second position. The output from the last VNF is in access mode (hypervisor-tagged) and the neighbor (ASAv firewall) is in StandAlone mode.

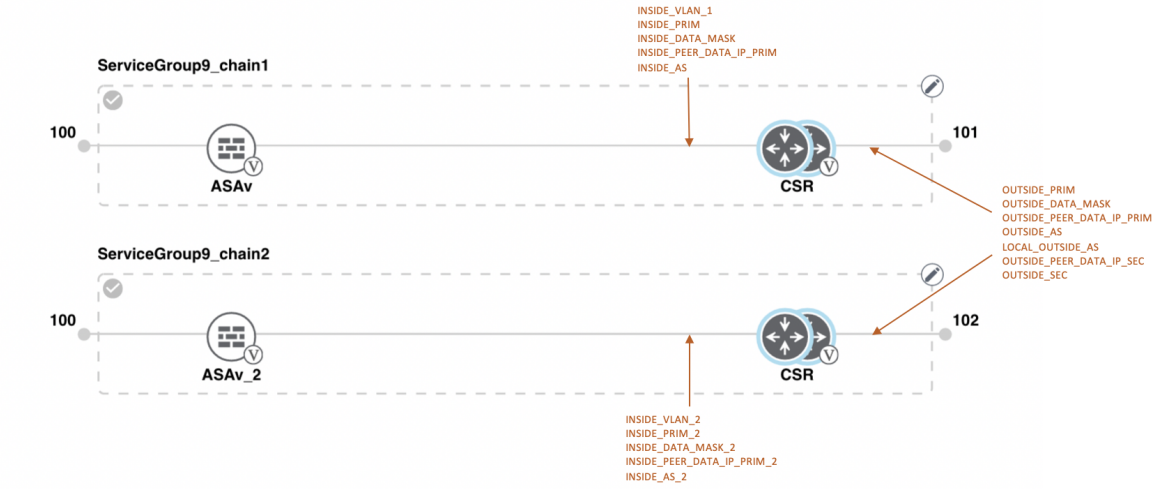

Figure 10: Shared–Cisco CSR1000V VNF in Last Position

The Cisco CSR1000V VNF in the last position is shared with the second service chain in the second position. The output from the last VNF is in access mode (hypervisor-tagged) and the neighbor is in StandAlone mode.

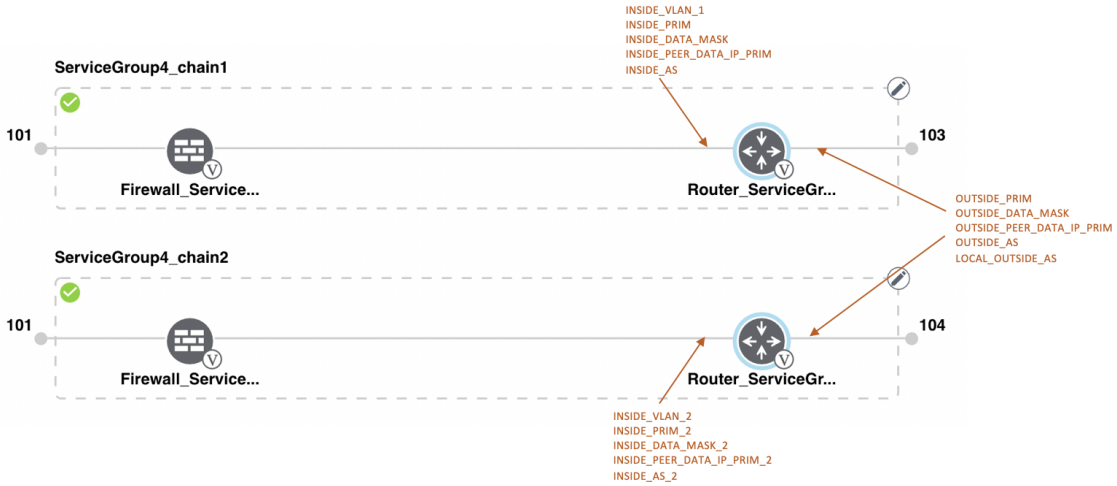

Figure 11: Shared–Cisco CSR1000V VNF in Last Position

The Cisco CSR1000V VNF in the last position is shared with the second service chain in the second position. The output from the last VNF is in access mode (hypervisor-tagged) and the neighbor (Firewall_Service) is in HA mode.

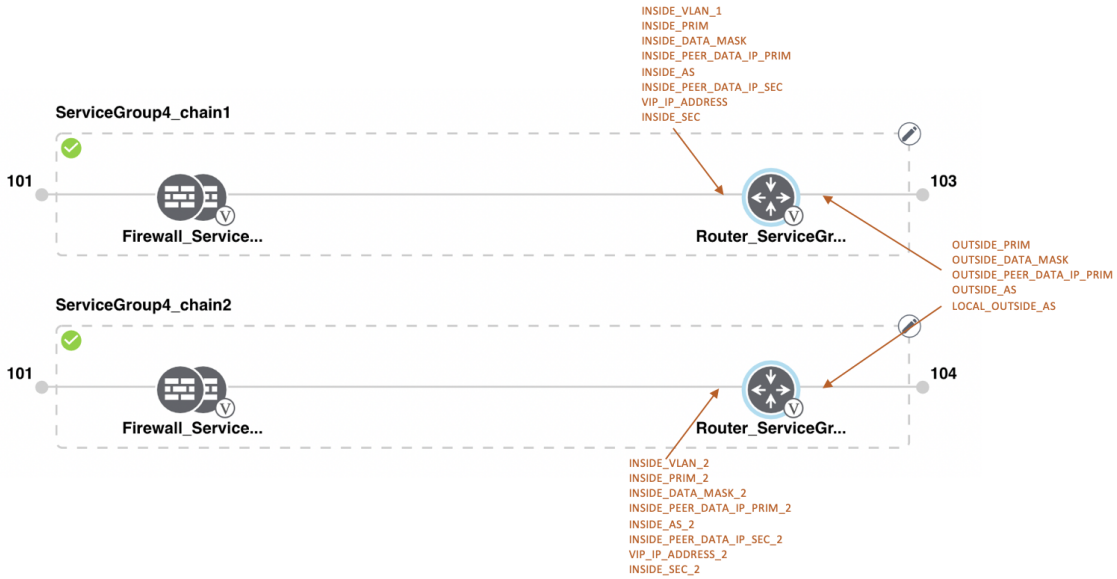

Figure 12: Shared–Cisco CSR1000V VNF in Last Position

The Cisco CSR1000V VNF in the last position is shared with the second service chain in the second position. The output from the last VNF is in access mode (hypervisor-tagged) and the neighbor (Firewall_Service) is in HA mode.

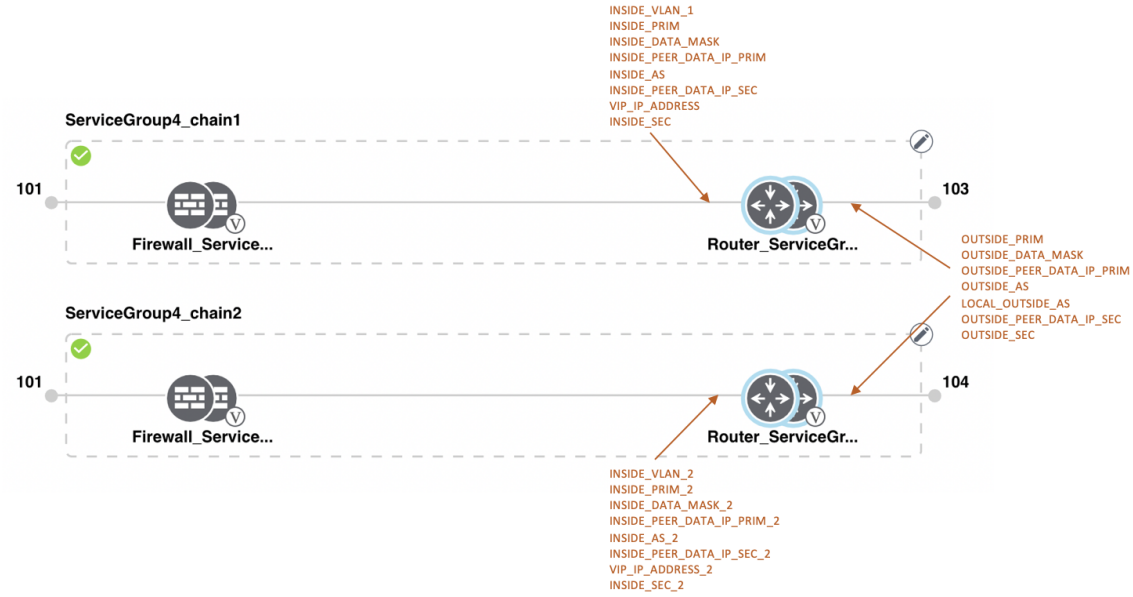

Figure 13: Shared–ASAv VNF in First Position

The ASAv VNF in the first position is shared with the second service chain in the first position. The input to the first VNF is in access mode (hypervisor-tagged) and the neighbor is in redundant mode.

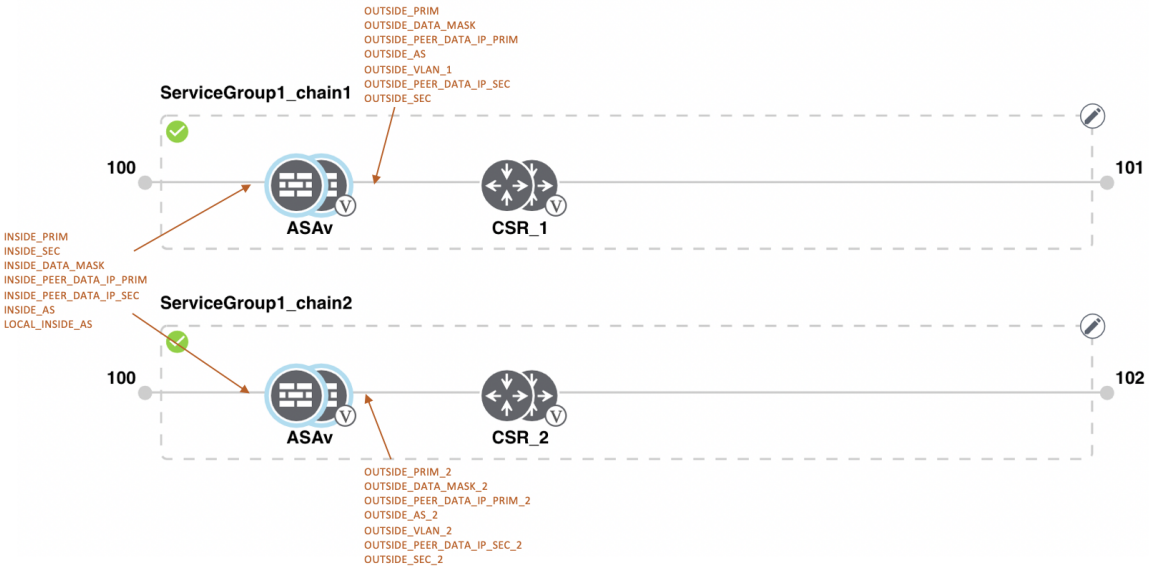

Figure 14: Shared–ASAv VNF in First Position

The ASAv (Firewall_Service) VNF in the first position is shared with the second service chain in the first position. The input to the first VNF is in access mode (hypervisor-tagged) and the neighbor is in StandAlone mode.

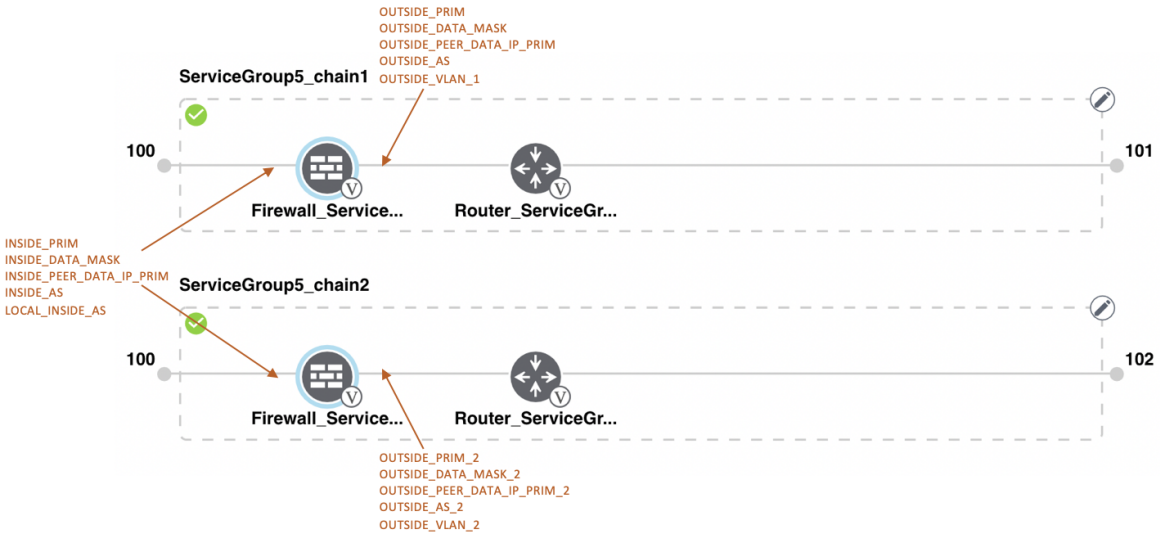

Figure 15: Shared–ASAv VNF in First Position

The ASAv (Firewall_Service) VNF in the first position is shared with the second service chain in the first position. The input to the first VNF is in access mode (hypervisor-tagged) and the neighbor, which is a router is in redundant mode.

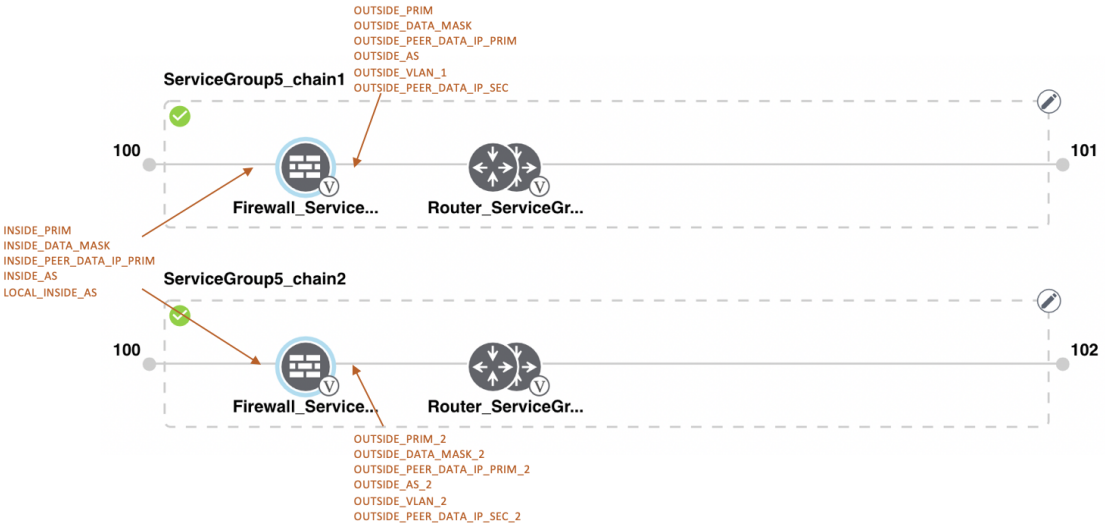

Figure 16: Shared–ASAv VNF in First Position

The ASAv VNF in the first position in HA mode is shared with the second service chain in the first position. The input to the first VNF is in trunk mode (vnf-tagged) and the neighbor is in redundant mode.

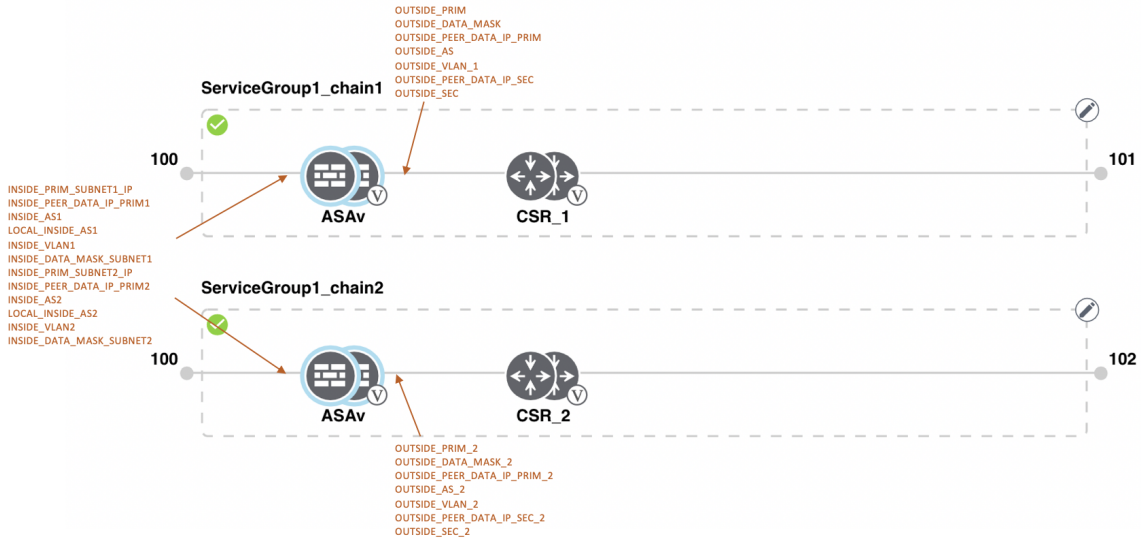

View Service Groups

To view service groups, perform the following steps:

Procedure

Step 1 From the Cisco SD-WAN Manager menu, choose **Configuration** > **Cloud OnRamp for Colocation**

Step 2 Click **Service Group**.

Step 3 For the desired service group, click **...** and choose **View**.

You can view the service chains in the design window.

Edit Service Groups

Before attaching a service group with a cluster, you can edit all parameters. After attaching a service group with a cluster, you can only edit monitoring configuration parameters. Also, after attaching a service group, you can only add new service chains but not edit or attach a service chain. Hence, ensure that you detach a service group from a cluster before editing an existing service chain. To edit and delete a service group, perform the following steps:

Procedure

Step 1 From the Cisco SD-WAN Manager menu, choose **Configuration** > **Cloud OnRamp for Colocation**.

Step 2 Click **Service Group**.

Step 3 For the desired service group, click **...** and choose **Edit**.

Step 4 To modify either service chain configuration or modify a VNF configuration, click a router or firewall VNF icon.

Step 5 To add new service chains, click **Add Service Chain**.

Attach or Detach a Service Group in a Cluster

To complete the Cisco Catalyst SD-WAN Cloud OnRamp for Colocation configuration, you must attach service groups to a cluster. To attach or detach a service group to and from a cluster, perform the following steps:

Procedure

Step 1 From the Cisco SD-WAN Manager menu, choose **Configuration** > **Cloud OnRamp for Colocation**.

Step 2 Click **...** adjacent to the corresponding cluster and choose **Attach Service Groups**.

Step 3 In the **Attach Service Groups** dialog box, choose one or more service groups in **Available Service Groups** and click **Add** to move the selected groups to **Selected Service Groups**.

Step 4 Click **Attach**.

Step 5 To detach a service group from a cluster, click **...** adjacent to the corresponding cluster and choose **Detach Service Groups**.

You can’t attach or detach a single service chain within a service group.

Step 6 In the **Config Preview** window that is displayed, click **Cancel** to cancel the attach or detach task.

Note

.

Step 7 To verify if service groups are attached or detached, you can view the status using Cisco SD-WAN Manager. Note the following points:

- If the status of the tasks in the **Task View** window is displayed as **FAILURE** or in **PENDING** for a long duration, see the "Troubleshoot Service Chain Issues" topic in the Cisco Catalyst SD-WAN Cloud OnRamp for Colocation solution guide.
- If a Cisco Colo Manager task fails, see the "Troubleshoot Cisco Colo Manager Issues" topic in the Cisco SD-WAN Cloud OnRamp for Colocation Solution Guide.

If a colocation cluster moves to **PENDING** state, for a cluster, click **...**, and choose **Sync**. This action moves the cluster back to **ACTIVE** state. The **Sync** option keeps Cisco SD-WAN Manager synchronized with the colocation devices.

View Information About VNFs

Table 8: Feature History

Feature Name	Release Information	Description
VNF States and Color Codes	Cisco IOS XE Catalyst SD-WAN Release 17.2.1r	This feature allows you to determine the state of a deployed VM using color codes, which you can view on the **Monitor** > **Devices** page. These color codes help you make decisions on creating service chains based on the state of the VM.

Table 9: Feature History

Feature Name	Release Information	Description
Network Utilization Charts for SR-IOV Enabled NICs and OVS Switch	Cisco IOS XE Catalyst SD-WAN Release 17.2.1r	This feature allows you to view network utilization charts of VM VNICs connected to both SR-IOV enabled NICs and OVS switch. These charts help you determine if the VM utilization is optimal to create service chains.

You can view performance specifications and required resources for each VNF. Reviewing this information can help you to determine which VNF to use when you're designing a network service. To view information about VNFs, perform the following steps:

Procedure

Step 1 From the Cisco SD-WAN Manager menu, choose **Monitor** > **Devices**.

Cisco vManage Release 20.6.x and earlier: From the Cisco SD-WAN Manager menu, choose **Monitor** > **Network**.

Cisco SD-WAN Manager displays the VNF information in a tabular format. The table includes information such as CPU use, memory consumption, and disk, and other core parameters that define performance of a network service.

Step 2 Click a CSP device from the table.

Step 3 From the left pane, click **VNF Status**.

Step 4 From the table, click the VNF name. Cisco SD-WAN Manager displays information about the specific VNF. You can click the network utilization, CPU utilization, memory utilization, and disk utilization to monitor the VNF resources utilization.

The following VNF information is displayed:

Table 10: VNF Information

Chart options bar	VNF information in graphical format	VNF information in color coded format
• Chart Options drop-down—Click Chart Options drop-down list to select the type of data to display. • Time periods—Click either a predefined time period, or a custom time period for which to display data.	Choose a VNF from the **Select Device** drop-down list to display information for the VNF.	The VNFs are shown in specific colors based on the following operational status of the VNF life cycle: • Green—VNF is healthy, deployed, and successfully booted up. • Red—VNF deployment or any other operation fails, or VNF stops. • Yellow—VNF is transitioning from one state to another.

The right pane displays the following:

- Filter criteria
- VNF table that lists information about all VNFs or VMs. By default, the first six VNFs are selected. The network utilization charts for VNICs connected to SR-IOV enabled NICs and OVS switch are displayed.

Figure 17: VNF Information

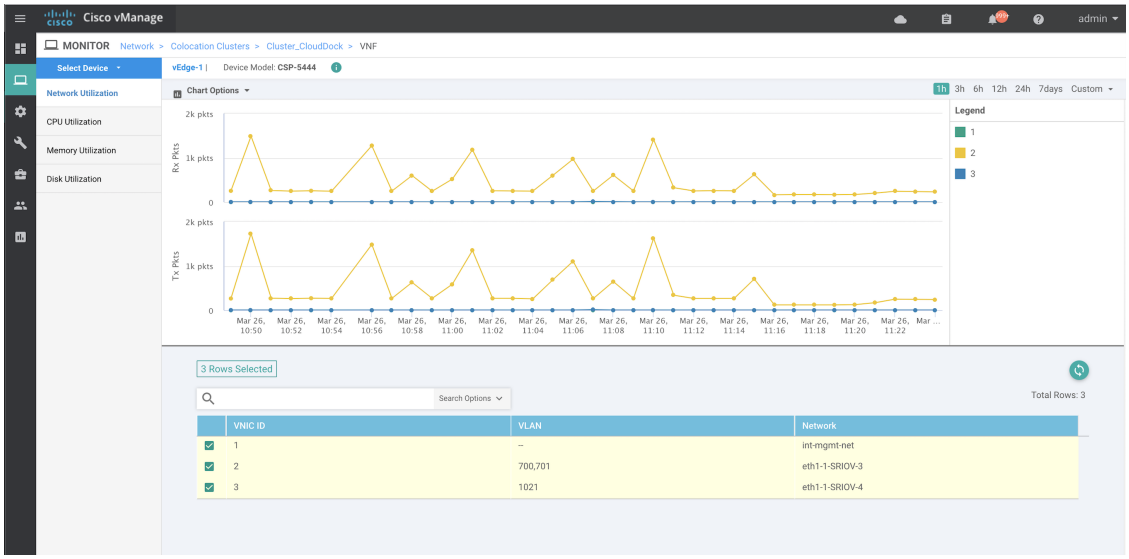

The graphical display plots information for the VNFs that you have selected by checking the check box.

- Click the check box at the left to select and deselect VNFs. You can select and display information for a maximum of six VNFs at a time.
- To change the sort order of a column, click the column title.

View Cisco Colo Manager Health

You can view Cisco Colo Manager (CCM) health for a device, CCM host system IP, CCM IP, and CCM state. Reviewing this information can help you to determine which VNF to use when you're designing a network service chain. To view information about VNFs, perform the following steps:

Procedure

Step 1 From the Cisco SD-WAN Manager menu, choose **Monitor** > **Devices**.

Cisco SD-WAN Manager Release 20.6.x and earlier: From the Cisco SD-WAN Manager menu, choose **Monitor** > **Network**.

The information of all devices is displayed in a tabular format.

Step 2 Click a CSP device from the table.

Step 3 From the left pane, click **Colo Manager**.

The right pane displays information about the memory usage, CPU usage, uptime, and so on, of the Cisco Colo Manager.

Monitor Cloud OnRamp Colocation Clusters

Table 11: Feature History

Feature Name	Release Information	Description
Network Assurance –VNFs: Stop/Start/Restart	Cisco IOS XE Catalyst SD-WAN Release 17.3.1a Cisco vManage Release 20.3.1	This feature provides the capability to stop, start, or restart VNFs on Cisco CSP devices from the **Colocation Cluster** tab. You can easily perform the operations on VNFs using Cisco SD-WAN Manager.

You can view the cluster information and their health states. Reviewing this information can help you to determine which Cisco CSP device is responsible for hosting each VNF in a service chain. To view information about a cluster, perform the following steps:

Procedure

Step 1 From the Cisco SD-WAN Manager menu, choose **Monitor** > **Devices**.

Cisco vManage Release 20.6.1 and earlier: From the Cisco SD-WAN Manager menu, choose **Monitor** > **Network**.

Step 2 To monitor clusters, click **Colocation Cluster**.

Cisco vManage Release 20.6.1 and earlier: Click **Colocation Clusters**.

All clusters with relevant information are displayed in a tabular format. Click a cluster name. You can monitor cluster by clicking **Config. View** and **Port Level View**.

- **Config. View**: The primary part of the window displays the CSP devices and switch devices that form the cluster. In the right pane, you can view the cluster information such as the available and total CPU resources, available and allocated memory, and so on, based on colocation size.

  The detail part of the window contains:

  - Search: To filter the search results, use the Filter option in the search bar.
  - A table that lists information about all devices in a cluster (Cisco CSP devices, PNFs, and switches).

    Click a Cisco CSP device. VNF information is displayed in a tabular format. The table includes information such as VNF name, service chains, number of CPUs, memory consumption, and other core parameters that define performance of a network service chain. See View Information About VNFs , on page 53 .

    To start, stop, or reboot a VNF, for the desired VNF, click **...** and choose one of the following operations:

    - **Start**.
    - **Stop**.
    - **Restart**.

Note
Ensure that service chain provisioning is complete and VMs are deployed, before issuing start, stop, restart operations on any of the VNFs in the service chain.

After you choose an operation on a VNF, wait until the operation is complete before you issue another operation. You can view the progress of an operation from the **Task View** window.

- **Port Level View**: After you activate the cluster, to view the port connectivity details, click **Port Level View**.

  You can view detailed port connectivity information for the switches and CSP devices in a color coded format based on the SR-IOV and OVS modes.

  To view the mapping of ports between the Catalyst 9500 switches and CSP devices, click or hover over a CSP device.

Figure 18: Monitor Port Connectivity Details of a Cluster

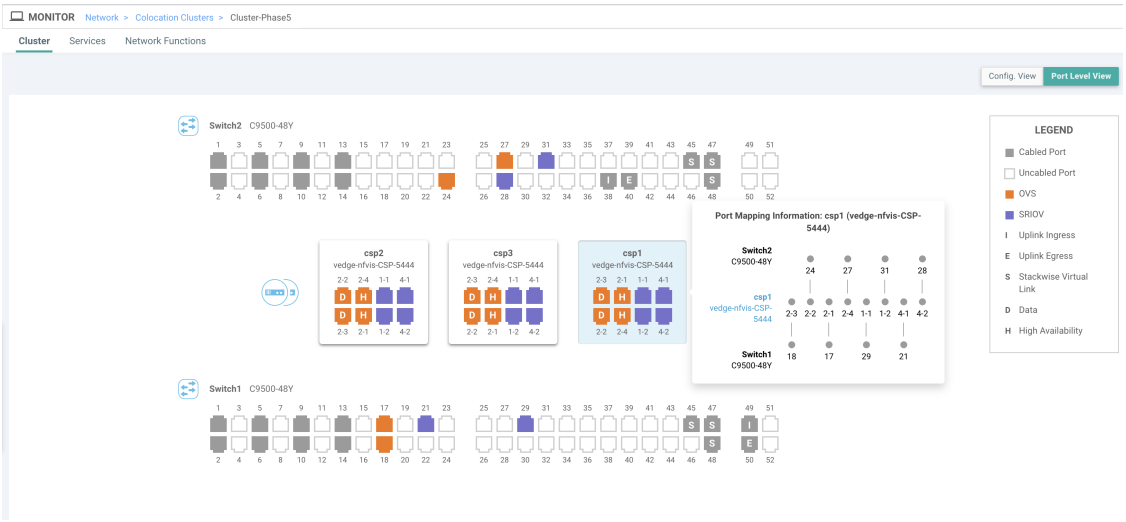

Step 3 Click **Services**.

Here, you can view the following:

- Complete information of a service chain. The first two columns display the name and description of the service chain in the service group and the remaining columns mention about the VNF, PNF statuses, monitoring service enablement, and the overall health of a service chain. You can also view the colocation user group associated with a service chain. The various health statuses and their representations are:
  - Healthy—An up arrow in green. A service chain is in 'Healthy' status when all the VNF, PNF devices are running and are in healthy state. Ensure that you configure the routing and policy correctly.
  - Unhealthy—A down arrow in red. If one of the VNFs or PNFs are in unhealthy state, the service chain is reported to be in 'Unhealthy' status. For example, after deploying a service chain, if one of the network function IP address changes on the WAN or LAN side, or the firewall policy isn't configured to let the traffic pass through, then unhealthy state is reported. This is because the network function or overall service chain is Unhealthy or both are in Unhealthy state.
  - Undetermined—Down arrow in yellow. This state is reported when the health of the service chain can't be determined. This state is also reported when there's no status such as healthy or unhealthy available for the monitored service chain over a time period. You can't query or search a service chain with undetermined status.

    If a service chain consists of a single PNF and PNF is outside the reachability of Cisco SD-WAN Manager, it can't be monitored. If a service chain consists of a single network function, the firewall that has VPN termination on both sides which can't be monitored, then it's reported as Undetermined.

    Note
    If the status of a service chain is undetermined, you can't choose the service chain to view the detailed monitoring information.

- If you had configured a service chain by enabling the monitoring field, then click a service group that is in Healthy or Unhealthy state. The primary part of the service chain monitoring window contains the following elements:

  Graphical display that plots the latency information of the service chain, VNFs, PNFs.

  The detail part of the service chain monitoring window contains:

- Search: To filter the search results, use the Filter option in the search bar.
- A table that lists information about all service chains, VNFs, PNFs, their health status, and types.
  - Check the service chain, VNF, PNF check boxes for the service chains, VNFs, PNFs you want to choose.
  - To change the sort order of a column, click the column title.

The status details column indicates the monitored data path and it provides the per hop analysis.

- Click **Diagram** and view the service group with all the service chains and VNFs in the design view window.
- Click a VNF. You can view CPU, memory, and disk allocated to the VNF in a dialog box.
- Choose a service group from the **Service Groups** drop-down. The design view displays the selected service group with all the service chains and VNFs.

Step 4 Click **Network Functions**.

Here, you can view the following:

- All the virtual or physical network functions in a tabular format. Use the **Show** button, and choose to display either a VNF or PNF.

  VNF information is displayed in a tabular format. The table includes information such as VNF name, service chains, colocation user groups, CPU use, memory consumption, and other core parameters that define performance of network service. To view more information about the VNF, click a VNF name. See View Information About VNFs , on page 53 .

- PNF information is displayed in tabular format. The table includes information such as the serial number and PNF type. To view and note configuration of a specific PNF, click the desired PNF serial number. Ensure that you manually note all the configuration of the PNFs and then configure the PNF devices. For example, the following are some of the PNF configuration where you position the PNF at various locations in the service chain. See the ASR 1000 Series Aggregation Services Routers Configuration Guides and Cisco Firepower Threat Defense Configuration Guides to configure the PNFs manually.

Figure 19: PNF in the First Position with Service Chain Side Parameters

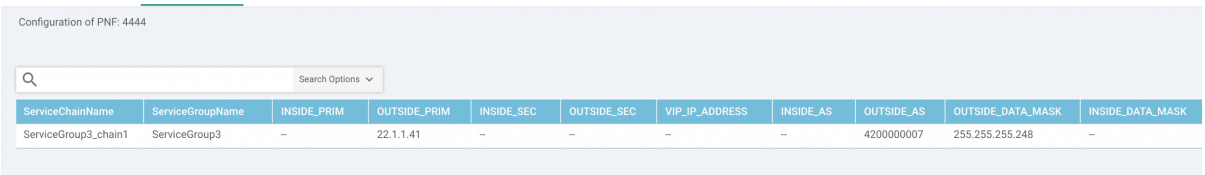

Figure 20: PNF in the First Position with Outside Neighbor Information

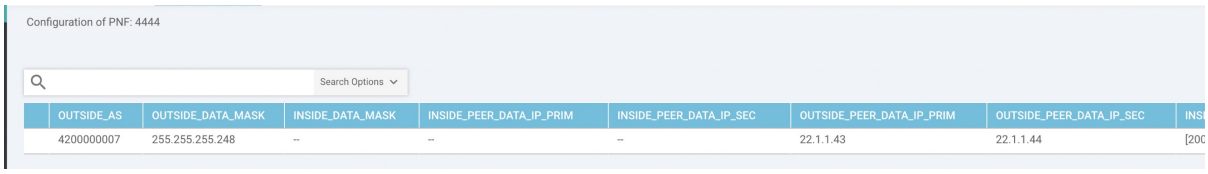

Figure 21: PNF Shared Across Two Service Chains

The ServiceGroup2_chain3 is a PNF-only service chain and therefore no configuration gets generated. The PNF is in the last position of the ServiceGroup2_chain1, so only INSIDE variables gets generated.

Figure 22: PNF Shared Across Two Service Chains with Outside Neighbor Information

Manage VM Catalog and Repository

Table 12: Feature History

Feature Name	Release Information	Description
Support for Cisco VM Image Upload in qcow2 Format	Cisco IOS XE Catalyst SD-WAN Release 17.7.1a Cisco vManage Release 20.7.1	This feature allows you to upload a virtual machine image to Cisco SD-WAN Manager in qcow2 format. Earlier, you could upload only a prepackaged image file in tar.gz format.

Cisco SD-WAN Manager supports uploading a prepackaged Cisco virtual machine image, tar.gz, or an image in qcow2 format. It is mandatory to upload a scaffold file if you choose a qcow2 image file. Similarly, you can now select either an image package file or a qcow2 image file with a scaffold file when configuring a Virtual Network Function (VNF) during service chain creation.

A scaffold file contains the following components:

- VNF metadata (image_properties.xml)
- System-generated variables from cluster resource pools for service chaining (system_generated_propeties.xml)
- Tokenized Day-0 configuration files
- Package manifest file (package.mf)

Alternatively, you can package the VM image by providing a root disk image in any of the supported formats (qcow2). Use the linux command-line NFVIS VM packaging tool, **nfvpt.py** to package the qcow2 or alternatively create a customized VM image using Cisco SD-WAN Manager. See Create Customized VNF Image, on page 62.

A VM is SR-IOV capable means sriov_supported is set to true in image_properties.xml in the vm package *.tar.gz. Also, the service chain network is automatically connected to SR-IOV network. If sriov_supported is set to false, an OVS network is created on the data port channel. It's attached to VM VNICs for service chaining by using the OVS network. For the Cloud OnRamp for Colocation solution, a VM uses homogeneous type of network in service chains. This type of network means it's either OVS or SR-IOV, and not a combination of SR-IOV and OVS.

Only two data VNICs are attached to any VM–one for inbound traffic and the other for outbound traffic. If more than two data interfaces are required, use subinterfaces configuration within the VM. The VM packages are stored in the VM catalog.

Note Each VM type such as firewall can have multiple VM images that are uploaded to Cisco SD-WAN Manager from same or different vendors and added to a catalog. Also, different versions that are based on the release of the same VM can be added to a catalog. However, ensure that the VM name is unique.

The Cisco VM image format can be bundled as *.tar.gz and can include:

- Root disk images to boot the VM.
- Package manifest for checksum validation of the file listing in the package.
- Image properties file in XML format that lists the VM meta data.
- (Optional) Day-0 configuration, other files that are required to bootstrap the VM.
- (Optional) HA Day-0 configuration if VM supports stateful HA.
- System-generated properties file in XML format that lists the VM system properties.

VM images can be hosted on both HTTP server local repository that Cisco SD-WAN Manager hosts or on the remote server.

If VM is in Cisco NFVIS supported VM package format such as, tar.gz, Cisco SD-WAN Manager performs all the processing and you can provide variable key and values during VNF provisioning.

Note Cisco SD-WAN Manager manages the Cisco VNFs, and the Day-1 and Day-N configurations within VNF aren't supported for other VNFs. See the Cisco NFVIS Configuration Guide, VM Image Packaging for more information about VM package format and content, and samples on image_properties.xml and manifest (package.mf).

To upload multiple packages for the same VM, same version, communication manager (CM) type, ensure that one of the three values (name, version, VNF type) are different. Then, you can repackage the VM *.tar.gz to be uploaded.

Upload VNF Images

The VNF images are stored in the Cisco SD-WAN Manager software repository. These VNF images are referenced during service chain deployment, and then they are pushed to Cisco NFVIS during service chain attachment.

Procedure

Step 1 From the Cisco SD-WAN Manager menu, choose **Maintenance** > **Software Repository**.

Step 2 To add a prepackaged VNF image, click **Virtual Images**, and then click **Upload Virtual Image**.

Step 3 Choose the location to store the virtual image.

- To store the virtual image on the local Cisco SD-WAN Manager server and download it to CSP devices over a control plane connection, click **Manager**. The **Upload VNF's Package to Manager** dialog box appears.
  - **a.** Drag and drop the virtual image file or the qcow2 image file to the dialog box or click **Browse** to choose the virtual image from the local Cisco SD-WAN Manager server. For example, CSR.tar.gz, ASAv.tar.gz, or ABC.qcow2
  - **b.** If you upload a file, specify the type of the uploaded file: **Image Package** or **Scaffold**. Optionally, specify a description of the file and add custom tags to the file. The tags can be used to filter images and scaffold files when creating a service chain.
  - **c.** If you upload a qcow2 image file, specify the service or VNF type: **FIREWALL** or **ROUTER**. Optionally, specify the following:
    - Description of the image
    - Version number of the image
    - Checksum
    - Hash algorithm

    You can also add custom tags to the file that can be used to filter images and scaffold files when creating a service chain.

    Note
    - It is mandatory to upload a scaffold file if you choose a qcow2 image file.
    - The option to select a qcow2 image file is available from Cisco vManage Release 20.7.1. In Cisco vManage Release 20.6.1 and earlier releases, you can select only a tar.gz file.

  - **d.** Click **Upload** to add the image to the virtual image repository. The virtual image repository table displays the added virtual image, and it available for installing on the CSP devices.

- To store the image on a remote Cisco SD-WAN Manager server and then download it to CSP devices, click **Remote Server - Manager**. The **Upload VNF's Package to Remote Server-Manager** dialog box appears.
  - **a.** In the **Manager Hostname/IP Address** field, enter the IP address of an interface on Cisco SD-WAN Manager server that is in the management VPN (typically, VPN 512).
  - **b.** Drag and drop the virtual image file or the qcow2 image file to the dialog box, or click **Browse** to choose the virtual image from the local Cisco SD-WAN Manager server.
  - **c.** If you upload a file, specify the type of the uploaded file: **Image Package** or **Scaffold**. Optionally, specify a description of the file and add custom tags to the file. The tags can be used to filter images and scaffold files when creating a service chain.

d. If you upload a qcow2 image file, specify the service or VNF type: **FIREWALL** or **ROUTER**. Optionally, specify the following:

- Description of the image
- Version number of the image
- Checksum
- Hash algorithm

You can also add custom tags to the file that can be used to filter images and scaffold files when creating a service chain.

Note
- It is mandatory to upload a scaffold file if you choose a qcow2 image file.
- The option to select a qcow2 image file is available from Cisco vManage Release 20.7.1. In Cisco vManage Release 20.6.1 and earlier releases, you can select only a tar.gz file.

e. Click **Upload** to add the image to the virtual image repository. The virtual image repository table displays the added virtual image, and it is available for installing on the CSP devices.

You can have multiple VNF entries such as a firewall from same or from different vendors. Also, you can add different versions of VNF that are based on the release of the same VNF. However, ensure that the VNF name is unique.

Create Customized VNF Image

Before you begin

You can upload one or more qcow2 images in addition to a root disk image as an input file along with VM-specific properties, bootstrap configuration files (if any), and generate a compressed TAR file. Through custom packaging, you can:

- Create a custom VM package along with image properties and bootstrap files (if needed) into a TAR archive file.
- Tokenize custom variables and apply system variables that are passed with the bootstrap configuration files.

Ensure that the following custom packaging requirements are met:

- Root disk image for a VNF–qcow2
- Day-0 configuration files–system and tokenized custom variables
- VM configuration–CPU, memory, disk, NICs
- HA mode–If a VNF supports HA, specify Day-0 primary and secondary files, NICs for a HA link.
- Additional Storage–If more storage is required, specify predefined disks (qcow2), storage volumes (NFVIS layer)

Procedure

Step 1 From the Cisco SD-WAN Manager menu, choose **Maintenance** > **Software Repository** .

Step 2 Click **Virtual Images** > **Add Custom VNF Package**.

Step 3 Configure the VNF with the following VNF package properties and click **Save**.

Table 13: VNF Package Properties

Field	Mandatory or Optional	Description
Package Name	Mandatory	The filename of the target VNF package. It's the Cisco NFVIS image name with .tar or .gz extensions.
App Vendor	Mandatory	Cisco VNFs or third-party VNFs.
Name	Mandatory	Name of the VNF image.
Version	Optional	Version number of a program.
Type	Mandatory	Type of VNF to choose. Supported VNF types are: Router, Firewall, Load Balancer, and Other.

Step 4 To package a VM qcow2 image, click **File Upload**, and browse to choose a qcow2 image file.

Step 5 To choose a bootstrap configuration file for VNF, if any, click **Day 0 Configuration** and click **File Upload** to browse and choose the file.

Include the following Day-0 configuration properties:

Table 14: Day-0 Configuration

Field	Mandatory or Optional	Description
Mount	Mandatory	The path where the bootstrap file gets mounted.
Parseable	Mandatory	A Day-0 configuration file can be parsed or not. Options are: **Enable** or **Disable**. By default, **Enable** is chosen.
High Availability	Mandatory	High availability for a Day-0 configuration file to choose. Supported values are: Standalone, HA Primary, HA Secondary.

Note
If any bootstrap configuration is required for a VNF, create a *bootstrap-config* or a *day0-config* file.

Step 6 To add a Day-0 configuration, click **Add**, and then click **Save**. The Day-0 configuration appears in the **Day 0 Config File** table. You can tokenize the bootstrap configuration variables with system and custom variables. To tokenize variables of a Day-0 configuration file, click **View Configuration File** next to the desired Day-0 configuration file. In the **Day 0 configuration file** dialog box, perform the following tasks:

Note
The bootstrap configuration file is an XML or a text file, and contains properties specific to a VNF and the environment. For a shared VNF, see the Custom Packaging Details for Shared VNF topic and additional references in Cisco SD-WAN Cloud OnRamp for Colocation Solution Guide for the list of system variables that must be added for different VNF types..

a) To add a system variable, in the **CLI configuration** dialog box, select, and highlight a property from the text fields. Click **System Variable**. The **Create System Variable** dialog box appears.
b) Choose a system variable from the **Variable Name** drop-down list, and click **Done**. The highlighted property is replaced by the system variable name.
c) To add a custom variable, in the **CLI configuration** dialog box, choose and highlight a custom variable attribute from the text fields. Click **Custom Variable**. The **Create Custom Variable** dialog box appears.
d) Enter the custom variable name and choose a type from **Type** drop-down list.
e) To set the custom variable attribute, do the following:
   - To ensure that the custom variable is mandatory when creating a service chain, click **Type** next to **Mandatory**.
   - To ensure that a VNF includes both primary and secondary day-0 files, click **Type** next to **Common**.
f) Click **Done**, and then click **Save**. The highlighted custom variable attribute is replaced by the custom variable name.

Step 7 To upload extra VM images, expand **Advance Options**, click **Upload Image**, and then browse to choose an extra qcow2 image file. Choose the root disk, Ephemeral disk 1, or Ephemeral disk 2, and click **Add**. The newly added VM image appears in the **Upload Image** table.

Note
Ensure that you don't combine ephemeral disks and storage volumes when uploading extra VM images.

Step 8 To add the storage information, expand **Add Storage**, and click **Add volume**. Provide the following storage information and click **Add**. The added storage details appear in the **Add Storage** table.

Table 15: Storage Properties

Field	Mandatory or Optional	Description
Size	Mandatory	The disk size that is required for the VM operation. If the size unit is GiB, the maximum disk size can be 256 GiB.
Size Unit	Mandatory	Choose size unit. The supported units are: MIB, GiB, TiB.
Device Type	Optional	Choose a disk or CD-ROM. By default, disk is chosen.

Field	Mandatory or Optional	Description
Location	Optional	The location of the disk or CD-ROM. By default, it's local.
Format	Optional	Choose a disk image format. The supported formats are: qcow2, raw, and vmdk. By default, it's raw.
Bus	Optional	Choose a value from the drop-down list. The supported values for a bus are: virtio, scsi, and ide. By default, it's virtio.

Step 9 To add VNF image properties, expand **Image Properties** and enter the following image information.

Table 16: VNF Image Properties

Field	Mandatory or Optional	Description
SR-IOV Mode	Mandatory	Enable or disable SR-IOV support. By default, it's enabled.
Monitored	Mandatory	VM health monitoring for those VMs that you can bootstrap. The options are: enable or disable. By default, it's enabled.
Bootup Time	Mandatory	The monitoring timeout period for a monitored VM. By default, it's 600 seconds.
Serial Console	Optional	The serial console that is supported or not. The options are: enable or disable. By default, it's disabled.
Privileged Mode	Optional	Allows special features like promiscuous mode and snooping. The options are: enable or disable. By default, it's disabled.
Dedicate Cores	Mandatory	Facilitates allocation of a dedicated resource (CPU) to supplement a VM's low latency (for example, router and firewall). Otherwise, shared resources are used. The options are: enable or disable. By default, it's enabled.

Step 10 To add VM resource requirements, expand **Resource Requirements** and enter the following information.

Table 17: VM Resource Requirements

Field	Mandatory or Optional	Description
Default CPU	Mandatory	The CPUs supported by a VM. The maximum numbers of CPUs supported are 8.
Default RAM	Mandatory	The RAM supported by a VM. The RAM can range 2–32.
Disk Size	Mandatory	The disk size in GB supported by a VM. The disk size can range 4–256.
Max number of VNICs	Optional	The maximum number of VNICs allowed for a VM. The number of VNICs can from range 8–32 and by default, the value is 8.
Management VNIC ID	Mandatory	The management VNIC ID corresponding to the management interface. The valid range is from 0 to maximum number of VNICs.
Number of Management VNICs ID	Mandatory	The number of VNICs.
High Availability VNIC ID	Mandatory	The VNIC IDs where high availability is enabled. The valid range is from 0–maximum number of VNICs. It shouldn't conflict with management VNIC Id. By default, the value is 1.
Number of High Availability VNICs ID	Mandatory	The maximum number of VNIC IDs where high availability is enabled. The valid range is 0–(maximum number of VNICs-number of management VNICs-2) and by default, the value is 1.

Step 11 To add day-0 configuration drive options, expand **Day 0 Configuration Drive options** and enter the following information.

Table 18: Day-0 Configuration Drive Options

Field	Mandatory or Optional	Description
Volume Label	Mandatory	The volume label of the Day-0 configuration drive. The options are: V1 or V2. By default, the option is V2. V2 is the config-drive label config-2. V1 is config-drive label cidata.
Init Drive	Optional	The Day-0 configuration file as a disk when mounted. The default drive is CD-ROM.
Init Bus	Optional	Choose an init bus. The supported values for a bus are: virtio, scsi, and ide. By default, it's ide.

The Software Repository table displays the customized VNF image, and image is available for choosing when creating a custom service chain.

View VNF Images

Procedure

Step 1 From the Cisco SD-WAN Manager menu, choose **Maintenance** > **Software Repository**.

Step 2 Click **Virtual Images**.

Step 3 To filter the search results, use the filter option in the search bar.

The Software Version column provides the version of the software image.

The Software Location column indicates where the software images are stored. Software images can be stored either in the repository on the Cisco SD-WAN Manager server or in a repository in a remote location.

The **Version Type Name** column provides the type of firewall.

The **Available Files** column lists the names of the VNF image files.

The **Update On** column displays when the software image was added to the repository.

Step 4 For the desired VNF image, click **...** and choose **Show Info**.

Delete VNF Images

Procedure

Step 1 From the Cisco SD-WAN Manager menu, choose **Maintenance** > **Software Repository**.

Step 2 Click **Virtual Images**. The images in the repository are displayed in a table.

Step 3 For the desired image, click **...** and choose **Delete**.

Note If you're downloading a VNF image to a device, you can't delete the VNF image until the download process completes.

Note If the VNF image is referenced by a service chain, it can't be deleted.

Upgrade Cisco NFVIS Using Cisco SD-WAN Manager

To upload and upgrade Cisco NFVIS, the upgrade image must be available as an archive file that can be uploaded to the Cisco SD-WAN Manager repository using Cisco SD-WAN Manager. After you upload the Cisco NFVIS image, the upgraded image can be applied to a CSP device by using the **Software Upgrade** window in Cisco SD-WAN Manager. You can perform the following tasks when upgrading Cisco NFVIS software using Cisco SD-WAN Manager:

- Upload Cisco NFVIS upgrade image. See Upload NFVIS Upgrade Image, on page 68.
- Upgrade a CSP device with the uploaded image. See Upgrade a CSP Device with a Cisco NFVIS Upgrade Image, on page 69.
- View the upgrade status for the CSP device by clicking the **Tasks** icon located in the Cisco SD-WAN Manager toolbar.

Upload NFVIS Upgrade Image

Procedure

Step 1 Download the Cisco NFVIS upgrade image from a prescribed location to your local system. You can also download the software image to an FTP server in your network.

Step 2 From the Cisco SD-WAN Manager menu, choose **Maintenance** > **Software Repository** .

Step 3 Click **Add New Software** > **Remote Server/Remote Server - Manager**.

You can either store the software image on a remote file server, on a remote Cisco SD-WAN Manager server, or on a Cisco SD-WAN Manager server.

Cisco SD-WAN Manager server: Saves software images on a local Cisco SD-WAN Manager server.

Remote server: Saves the URL pointing to the location of the software image and can be accessed using an FTP or HTTP URL.

Remote Cisco SD-WAN Manager server: Saves software images on a remote Cisco SD-WAN Manager server and location of the remote Cisco SD-WAN Manager server is stored in the local Cisco SD-WAN Manager server.

Step 4 To add the image to the software repository, browse and choose the Cisco NFVIS upgrade image that you had downloaded in Step1.

Step 5 Click **Add|Upload**.

The Software Repository table displays the added NFVIS upgrade image, and it's available for installing on the CSP devices. See the Manage Software Upgrade and Repository topic in the Cisco Catalyst SD-WAN Monitor and Maintain Configuration Guide.

Upgrade a CSP Device with a Cisco NFVIS Upgrade Image

Before you begin

Ensure that the Cisco NFVIS software versions are the files that have `.nfvispkg` extension.

Procedure

Step 1 From theCisco SD-WAN Manager menu, choose **Maintenance** > **Software Upgrade** > **WAN Edge**.

Step 2 Check one or more CSP device check boxes for the devices you want to choose.

Step 3 Click **Upgrade**. The **Software Upgrade** dialog box appears.

Step 4 Choose the Cisco NFVIS software version to install on the CSP device. If software is located on a remote server, choose the appropriate remote version.

Step 5 To automatically upgrade and activate with the new Cisco NFVIS software version and reboot the CSP device, check the **Activate and Reboot** check box.

If you don't check the **Activate and Reboot** check box, the CSP device downloads and verifies the software image. However, the CSP device continues to run the old or current version of the software image. To enable the CSP device to run the new software image, you must manually activate the new Cisco NFVIS software version by choosing the device again and clicking the **Activate** button in the **Software Upgrade** window.

Step 6 Click **Upgrade**.

The **Task View** window displays a list of all running tasks along with total number of successes and failures. The window periodically refreshes and displays messages to indicate the progress or status of the upgrade. You can easily access the software upgrade status window by clicking the **Task View** icon located in the Cisco SD-WAN Manager toolbar.

Note
If two or more CSP devices belonging to the same cluster are upgraded, the software upgrade for the CSP devices happens in a sequence.

Note
The **Set the Default Software Version** option isn't available for the Cisco NFVIS images.

The CSP device reboots and the new NFVIS version is activated on the device. This reboot happens during the **Activate** phase. The activation can either happen immediately after upgrade if you check the **Activate and Reboot** check box, or by manually clicking **Activate** after choosing the CSP device again.

To verify if CSP device has rebooted and is running, use the task view window. Cisco SD-WAN Manager polls your entire network every 90 seconds up to 30 times and shows the status on th task view window.

Note You can delete a Cisco NFVIS software image from a CSP device if the image version isn't the active version that is running on the device.