



Systems and Interfaces Configuration Guide, Cisco SD-WAN Releases 19.1, 19.2, and 19.3

First Published: 2019-08-15

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883



CONTENTS

CHAPTER 1

What's New for Cisco SD-WAN 1

What's New for Cisco SD-WAN Release 19.2.x 1

CHAPTER 2

System and Interfaces Overview 3

Basic Settings for Cisco vManage 7

Configure Organization Name 8

Configure Cisco vBond DNS Name or IP Address 9

Configure Controller Certificate Authorization Settings 9

Enforce Software Version on Devices 11

Banner 11

Create a Custom Banner 13

Collect Device Statistics 13

Configure or Cancel vManage Server Maintenance Window 14

Configure Basic System Parameters 14

Configure Global Parameters 20

Create Global Settings Feature Template 20

CLI Equivalent 22

Configure NTP using Cisco vManage 22

Configure NTP using CLI 25

Configuring Time Using CLI on Cisco vEdge Device 27

Configure GPS Using Cisco vManage 27

Configure GPS Using CLI on Cisco vEdge Device 29

Configure System Logging 29

Syslog Message Format, Syslog Message Levels, and System Log Files 30

Configure Logging Using Cisco vManage 32

Export Cisco vManage NMS Audit Log to Syslog Server 35

Configure System Logging Using CLI	36
View System Logging Information	37
SSH Terminal	38
Tenant Management	38

CHAPTER 3**Configuring User Access and Authentication 41**

Manage Users using vManage	41
Configure User Using CLI	43
Manage a User Group	44
Creating Groups Using CLI	45
Configuring RADIUS Authentication Using CLI	45
Configure SSH Authentication	46
SSH Authentication using vManage on Cisco vEdge Devices	47
Configure SSH Authentication using CLI on Cisco vEdge Devices	47
Configure the Authentication Order	48
Configure NAS Attributes using CLI	50
Role-Based Access with AAA	51
Configuring AAA using vManage Template	60
Configuring IEEE 802.1X and IEEE 802.11i Authentication	66

CHAPTER 4**Create a Device Template from Feature Templates 79**

Configure Devices	82
Create a Device CLI Template	82
Manage Device Templates	83
View Device Templates	84
Attach and Detach a Device Template	84
Change the Device Rollback Timer	86
Preview Device Configuration and View Configuration Differences	87
Change Variable Values for a Device	87
Configuring Devices using vManage	88
Change Configuration Modes	88
Upload WAN Edge Router Authorized Serial Number File	89
Upload WAN Edge Router Serial Numbers from Cisco Smart Account	89
Generate Bootstrap Configuration for a vEdge Cloud Router	90

Export Device Data in CSV Format	90
View and Copy Device Configuration	91
Delete a WAN Edge Router	91
Decommission a vEdge Cloud router	92
View Template Log and Device Bringup	92
Add a Cisco vBond Orchestrator	92
Configure Cisco vSmart Controllers	93
Create a UCS-E Template	94

CHAPTER 5**Configure Network Interfaces 99**

Configure VPN	100
VPN	100
Create a VPN Template	100
Changing the Scope for a Parameter Value	101
Configure Basic VPN Parameters	102
Configure DNS and Static Hostname Mapping	103
Configure Interfaces in the WAN Transport VPN (VPN 0)	104
Extend the WAN Transport VPN	107
Configure GRE Interfaces and Advertise Services to Them	110
Configure the System Interface	114
Configure Control Plane High Availability	115
Configure Other Interfaces	115
Role-Based Access Control by VPN	117
VPN Dashboard Overview	117
Configure and Manage VPN Segments	118
Configure and Manage VPN Groups	119
Configure User with User group	119
Configure Interface Properties	120
Set the Interface Speed	120
Set the Interface MTU	121
Monitoring Bandwidth on a Transport Circuit	122
Enable DHCP Server using Cisco vManage	122
Configure DHCP Using CLI	126
Configuring PPPoE	127

Configure PPPoE from vManage Templates	128
Configure PPPoE from the CLI	131
Configuring VRRP	133
Network Interface Configuration Examples for Cisco vEdge Devices	136
Configure VPN Interfaces Using vManage	153
Configure VPN Ethernet Interface	153
Configure VPN Ethernet Interface	153
Configure Basic Interface Functionality	154
Create a Tunnel Interface	155
Configure an Interface as a NAT Device	160
Apply Access Lists and QoS Parameters	166
Add ARP Table Entries	167
VPN Interface Bridge	167
VPN Interface Ethernet PPPoE	169
VPN Interface GRE	176
VPN Interface IPsec (for Cisco vEdge Devices)	180
VPN Interface PPP	186
VPN Interface PPP Ethernet	194
Cellular Interfaces	199
Configure Cellular Interfaces Using vManage	200
Configuring Cellular Interfaces Using CLI	208
Interface CLI Reference	212



CHAPTER 1

What's New for Cisco SD-WAN



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

This chapter describes what's new in Cisco SD-WAN for each release.

- [What's New for Cisco SD-WAN Release 19.2.x, on page 1](#)

What's New for Cisco SD-WAN Release 19.2.x

This section applies to Cisco vEdge devices.

Cisco is constantly enhancing the SD-WAN solution with every release and we try and keep the content in line with the latest enhancements. The following table lists new and modified features we documented in the Configuration, Command Reference, and Hardware Installation guides. For information on additional features and fixes that were committed to the SD-WAN solution, see the *Resolved and Open Bugs* section in the Release Notes.

Table 1: What's New for Cisco vEdge Device

Feature	Description
Getting Started	
API Cross-Site Request Forgery Prevention	This feature adds protection against Cross-Site Request Forgery (CSRF) that occurs when using Cisco SD-WAN REST APIs. This protection is provided by including a CSRF token with API requests. You can put requests on an allowed list so that they do not require protection if needed. See Cross-Site Request Forgery Prevention .
Systems and Interfaces	

Feature	Description
Secure Shell Authentication Using RSA Keys	This feature helps configure RSA keys by securing communication between a client and a Cisco SD-WAN server. See SSH Authentication using vManage on Cisco XE SD-WAN Devices. See Configure SSH Authentication .
Policies	
Packet Duplication for Noisy Channels	This feature helps mitigate packet loss over noisy channels, thereby maintaining high application QoE for voice and video. See Configure and Monitor Packet Duplication .
Control Traffic Flow Using Class of Service Values	This feature lets you control the flow of traffic into and out of a Cisco device's interface based on the conditions defined in the quality of service (QoS) map. A priority field and a layer 2 class of service (CoS) were added for configuring the re-write rule. See Configure Localized Data Policy for IPv4 Using Cisco vManage .
Security	
Secure Communication Using Pairwise IPsec Keys	This feature allows private pairwise IPsec session keys to be created and installed for secure communication between IPsec devices and its peers. For related information, see IPsec Pairwise Keys Overview .
Configure IKE-Enabled IPsec Tunnels	The pre-shared key needs to be at least 16 bytes in length. The IPsec tunnel establishment fails if the key size is less than 16 characters when the router is upgraded to version 19.2. See Configure IKE-Enabled IPsec Tunnels .
Network Optimization and High Availability	
Disaster Recovery for vManage	This feature helps you configure Cisco vManage in an active or standby mode to counteract hardware or software failures that may occur due to unforeseen circumstances. See Configure Disaster Recovery .
Share VNF Devices Across Service Chains	This feature lets you share Virtual Network Function (VNF) devices across service chains to improve resource utilisation and reduce resource fragmentation. See Share VNF Devices Across Service Chains .
Monitor Service Chain Health	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. See Monitor Service Chain Health .
Manage PNF Devices in Service Chains	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. See Manage PNF Devices in Service Chains .



CHAPTER 2

System and Interfaces Overview

Setting up the basic system-wide functionality of network devices is a simple and straightforward process. These basic parameters include defining host properties, such as name and IP address; setting time properties, including NTP; setting up user access to the devices; defining system log (syslog) parameters; and creating network interfaces.

In addition, the Cisco SD-WAN software provides a number of management interfaces for accessing the Cisco SD-WAN devices in the overlay network.

Host Properties

All devices have basic system-wide properties that specify information that the Cisco SD-WAN software uses to construct a view of the network topology. Each device has a system IP address, which provides a fixed location of the device in the overlay network. This address, whose function is similar to that of a router ID on a router, is independent of any of the interfaces and interface IP addresses on the device. The system IP address is one of the four components of each device's TLOC property.

A second host property that must be set on all devices is the IP address of the vBond orchestrator for the network domain, or a DNS name that resolves to one or more IP addresses for vBond orchestrators. The vBond orchestrator automatically orchestrates the bringup of the overlay network, admitting a new device into the overlay and providing the introductions that allow device and vSmart controllers to locate each other.

Two other system-wide host properties are required on all devices, except for the vBond orchestrators, to allow the Cisco SD-WAN software to construct a view of the topology: the domain identifier and the site identifier.

To configure the host properties, see *Cisco SD-WAN Overlay Network Bringup* .

Time and NTP

The Cisco SD-WAN software implements the Network Time Protocol (NTP) to synchronize and coordinate time distribution across the Cisco SD-WAN overlay network. NTP uses an intersection algorithm to select applicable time servers and avoid issues caused due to network latency. The servers also can redistribute reference time using local routing algorithms and time daemons. NTP is defined in RFC 5905, *Network Time Protocol Version 4: Protocol and Algorithms Specification* .

User Authentication and Access with AAA, RADIUS, and TACACS+

The Cisco SD-WAN software uses Authentication, Authorization, and Accounting (AAA) to provide security for devices on the network. AAA, in combination with RADIUS and TACACS+ user authentication, controls

which users are allowed access to devices and what operations they are authorized to perform once they are logged in or connected to the devices.

Authentication refers to the process by which the user trying to access the device is authenticated. To access devices, users log in with name and a password. The local device can authenticate users, or authentication can be performed by a remote device, either by a Remote Authentication Dial-In User Service (RADIUS) server or by a Terminal Access Controller Access-Control System (TACACS+), or by both in sequence.

Authorization determines whether the user is authorized to perform a given activity on the device. In the Cisco SD-WAN software, authorization is implemented using role-based access. Access is based on groups that are configured on the devices. A user can be a member of one or more groups. External groups are also considered when performing authorization; that is, the Cisco SD-WAN software retrieves group names from RADIUS or TACACS+ servers. Each group is assigned privileges that authorize the group members to perform specific functions on the device. These privileges correspond to specific hierarchies of the configuration commands and the corresponding hierarchies of operational commands that members of the group are allowed to view or modify.

The Cisco SD-WAN software does not implement AAA accounting.

For more information, see *Role-Based Access with AAA*.

Authentication for WANs and WLANs

For wired networks (WANs), Cisco SD-WAN devices can run IEEE 802.1X software to prevent unauthorized network devices from gaining access to the WAN. IEEE 802.1X is a port-based network access control (PNAC) protocol that uses a client-server mechanism to provide authentication for devices wishing to connect to the network. You can enable 802.1X on vEdge router interfaces to have the router act as an 802.1X authenticator, responsible for authorizing or denying access to network devices.

IEEE 802.1X authentication requires three components:

- **Supplicant**—Client device, such as a laptop, that requests access to the WAN. In the Cisco SD-WAN overlay network, a supplicant is any service-side device that is running 802.1X-compliant software. These devices send network access requests to the router.
- **Authenticator**— A network device that provides a barrier to the WAN. In the overlay network, you can configure an interface device to act as an 802.1X authenticator. The device supports both controlled and uncontrolled ports. For controlled ports, the Cisco SD-WAN device acts as an 802.1X port access entity (PAE), allowing authorized network traffic and preventing unauthorized network traffic ingressing to and egressing from the controlled port. For uncontrolled ports, the Cisco SD-WAN, acting as an 802.1X PAE, transmits and receives Extensible Authentication Protocol over IEEE 802 (EAP over LAN, or EAPOL) frames.
- **Authentication server**—Host running authentication software that validates and authenticates supplicants that want to connect to the WAN. In the overlay network, this host is an external RADIUS server. This RADIUS server authenticates each client connected to the 802.1X port interface Cisco SD-WAN router and assigns the interface to a VLAN before the client is allowed to access any of the services offered by the router or by the LAN.

For wireless LANs (WLANs), routers can run IEEE 802.11i prevents unauthorized network devices from gaining access to the WLANs. IEEE 802.11i implements Wi-Fi Protected Access (WPA) and Wi-Fi Protected Access II (WPA2) to provide authentication and encryption for devices that want to connect to a WLAN. WPA authenticates individual users on the WLAN using a username and password. WPA uses the Temporal Key Integrity Protocol (TKIP), which is based on the RC4 cipher. WPA2 implements the NIST FIPS 140-2-compliant AES encryption algorithm along with IEEE 802.1X-based authentication, to enhance user

access security over WPA. WPA2 uses the Counter Mode Cipher Block Chaining Message Authentication Code Protocol (CCMP), which is based on the AES cipher. Authentication is done either using preshared keys or through RADIUS authentication.

Network Segmentation

The Layer 3 network segmentation in Cisco SD-WAN is achieved through VPNs on Cisco vEdge devices .

Network Interfaces

In the Cisco SD-WAN overlay network design, interfaces are associated with VPNs that translate to VRFs. The interfaces that participate in a VPN are configured and enabled in that VPN. Each interface can be present only in a single VPN.

The overlay network has the following types of VPNs/VRFs:



Note

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

- **VPN 0—Transport VPN**, which carries control traffic via the configured WAN transport interfaces. Initially, VPN 0 contains all of a device's interfaces except for the management interface, and all interfaces are disabled.
- **VPN 512—Management VPN**, which carries out-of-band network management traffic among the Cisco SD-WAN devices in the overlay network. The interface used for management traffic resides in VPN 512. By default, VPN 512 is configured and enabled on all Cisco SD-WAN devices except for Cisco vEdge 100. For controller devices, by default, VPN 512 is not configured.
- **Service VPNs**—VPNs 1 through VPN 65535 except for VPN 0 and VPN 512. All service-side interfaces activated in these VPNs connect to a local or branch network that is generally located at the same site as the Cisco SD-WAN router. These interfaces carry data traffic throughout the overlay network.

For each network interface, you can configure a number of interface-specific properties, such as DHCP clients and servers, VRRP, interface MTU and speed, and PPPoE. At a high level, for an interface to be operational, you must configure an IP address for the interface and mark it as operational (no shutdown). In practice, you always configure additional parameters for each interface.

Management and Monitoring Options

There are various ways you can manage and monitor a router. Management interfaces provide access to devices in Cisco SD-WAN overlay network, allowing you to collect information from the devices in an out-of-band fashion and to perform operations on the devices, such as configuring and rebooting them.

The following management interfaces are available:

- Command-line interface (CLI)
- IP Flow Information Export (IPFIX)
- RESTful API
- SNMP
- System logging (syslog) messages

- vManage web server

CLI

You can access a command-line interface (CLI) on each device, and from the CLI you configure overlay network features on the local device and gather operational status and information regarding that device. While a CLI is available, it is strongly recommended that you configure and monitor all Cisco SD-WAN network devices from a Cisco vManage web server, which provides visual views of network-wide operations and device status, including drill-downs that display details operation and status data. In addition, the vManage web server provides straightforward tools for bringing up and configuring overlay network devices, including bulk operations for setting up multiple devices simultaneously.

You access the CLI by establishing an SSH session to a Cisco SD-WAN device. For a hardware vEdge router, you can also connect to the device's console port.

For a Cisco SD-WAN device that is being managed by a vManage NMS, if you create or modify the configuration from the CLI, those changes are overwritten by the configuration that is stored in the vManage configuration database.

IPFIX

The IP Flow Information Export (IPFIX) protocol, also called cflowd, is a tool for monitoring the traffic flowing through Cisco SD-WAN routers in the overlay network and exporting information about the traffic to a flow collector. The exported information is sent in template reports, which contain both information about the flow and data extracted from the IP headers of the packets in the flow.

The Cisco SD-WAN cflowd performs 1:1 traffic sampling. Information about all flows is aggregated in the cflowd records; flows are not sampled. Cisco SD-WAN routers do not cache any of the records that are exported to a collector.

The Cisco SD-WAN cflowd software implements cflowd version 10, as specified in RFC 7011 and RFC 7012.

For a list of elements exported by IPFIX, see [Traffic Flow Monitoring with cflowd](#).

To enable the collection of traffic flow information, you create data policies that identify the traffic of interest and then direct that traffic to a cflowd collector. For more information, see [Traffic Flow Monitoring with Cflowd](#).

You can also enable cflowd visibility directly on Cisco SD-WAN routers without configuring data policy so that you can perform traffic flow monitoring on traffic coming to the router from all VPNs in the LAN. You then monitor the traffic from the vManage GUI or from the router's CLI.

RESTful API

The Cisco SD-WAN software provides a RESTful API, which is a programmatic interface for controlling, configuring, and monitoring the Cisco SD-WAN devices in an overlay network. You access the RESTful API through the vManage web server.

The Cisco SD-WAN RESTful API calls expose the functionality of Cisco SD-WAN software and hardware features and of the normal operations you perform to maintain the devices and the overlay network itself.

SNMP

The Simple Network Management Protocol (SNMP) allows you to manage all Cisco SD-WAN devices in the overlay network. The Cisco SD-WAN software supports SNMP v2c.

You can configure basic SNMP properties—device name, location, contact, and community—that allow the device to be monitored by an SNMP NMS.

You can configure trap groups and SNMP servers to receive traps.

The object identifier (OID) for the Internet port of the SNMP MIB is 1.3.6.1.

SNMP traps are asynchronous notifications that a Cisco SD-WAN device sends to an SNMP management server. Traps notify the management server of events, whether normal or significant, that occur on the Cisco SD-WAN device. By default, SNMP traps are not sent to an SNMP server. Note that for SNMPv3, the PDU type for notifications is either SNMPv2c inform (InformRequest-PDU) or trap (Trapv2-PDU).

Syslog Messages

System logging operations use a mechanism similar to the UNIX syslog command to record system-wide, high-level operations that occur on the Cisco SD-WAN devices in the overlay network. The log levels (priorities) of the messages are the same as those in standard UNIX commands, and you can configure which priority of syslog messages are logged. Messages can be logged to files on the Cisco SD-WAN device or to a remote host.

vManage NMS

The vManage NMS is a centralized network management system that allows configuration and management of all Cisco SD-WAN devices in the overlay network and provides a dashboard into the operations of the entire network and of individual devices in the network. Each vManage NMS runs on a web server in the network. Three or more vManage web servers are consolidated into a vManage cluster to provide scalability and management support for up to 6,000 Cisco SD-WAN devices, to distribute vManage functions across multiple devices, and to provide redundancy of network management operations.

- [Basic Settings for Cisco vManage, on page 7](#)
- [Configure Basic System Parameters, on page 14](#)
- [Configure Global Parameters, on page 20](#)
- [Configure NTP using Cisco vManage, on page 22](#)
- [Configure NTP using CLI, on page 25](#)
- [Configuring Time Using CLI on Cisco vEdge Device, on page 27](#)
- [Configure GPS Using Cisco vManage, on page 27](#)
- [Configure GPS Using CLI on Cisco vEdge Device, on page 29](#)
- [Configure System Logging, on page 29](#)
- [SSH Terminal, on page 38](#)
- [Tenant Management, on page 38](#)

Basic Settings for Cisco vManage

The System template is used to configure system-level Cisco vManage workflows.

Use the Settings screen to view the current settings and configure the setting for Cisco vManage parameters, including the organization name, vBond orchestrator's DNS name or IP address, certificate settings, and statistics collection.

The current setting for each item is displayed in the bar for each item, immediately following the name.

Setting	Value	Actions
Organization Name	vIPtela Inc Regression	View
vBond	10.0.12.26 : 12346	View Edit
Email Notifications	Disabled	View Edit
Controller Certificate Authorization	Manual	View Edit
vEdge Cloud Certificate Authorization	Automated	View Edit
Web Server Certificate	04 Nov 2019 9:07:40 AM	CSR Certificate
Enforce Software Version (ZTP)		View Edit
Banner	Disabled	View Edit
Reverse Proxy	Disabled	View Edit
Statistics Setting		View Edit
CloudExpress	Enabled	View Edit
vAnalytics	Disabled	View Edit
Client Session Timeout	Disabled	View Edit
Data Stream	Disabled	View Edit
Tenancy Mode	Single Tenant	View Edit
Statistics Configuration	Collection Interval: 30 minutes	View Edit
Maintenance Window	Not Configured	Edit
Identity Provider Settings	Disabled	View Edit
Statistics Database Configuration	Maximum Available Space: 17.7176 GB	View Edit
Google Map API Key	Maps API Key: AlzaSyA1PwZsBFTR4-PLCEResl6qMfEiqnRV898	View Edit
Software Install Timeout	Collection Interval: 60 minutes	View Edit

368729

Configure Organization Name

Before you can generate a Certificate Signing Request (CSR), you must configure the name of your organization. The organization name is included in the CSR.

In public key infrastructure (PKI) systems, a CSR is sent to a certificate authority to apply for a digital identity certificate.

To configure the organization name:

1. Click the **Edit** button to the right of the **Organization Name** bar.
2. In the **Organization Name** field, enter the name of your organization. The organization name must be identical to the name that is configured on the vBond orchestrator.
3. In the **Confirm Organization Name** field, re-enter and confirm your organization name.
4. Click **Save**.

Note that once the control connections are up and running, the organization name bar is no longer editable.

Configure Cisco vBond DNS Name or IP Address

1. Click the **Edit** button to the right of the vBond bar.
2. In the vBond **DNS/IP Address: Port** field, enter the DNS name that points to the vBond orchestrator or the IP address of the Cisco vBond orchestrator and the port number to use to connect to it.
3. Click **Save**.

Configure Controller Certificate Authorization Settings

Signed certificates are used to authenticate devices in the overlay network. Once authenticated, devices can establish secure sessions between each other. It is from the Cisco vManage that you generate these certificates and install them on the controller devices—Cisco vBond orchestrators, Cisco vManage, and Cisco vSmart controllers. You can use certificates signed by Symantec, or you can use enterprise root certificates.

The controller certification authorization settings establish how the certification generation for all controller devices will be done. They do not generate the certificates.

You need to select the certificate-generation method only once. The method you select is automatically used each time you add a device to the overlay network.

To have the Symantec signing server automatically generate, sign, and install certificates on each controller device:

1. Click the **Edit** button to the right of the **Controller Certificate Authorization** bar.
2. Click **Symantec Automated** (Recommended). This is the recommended method for handling controller signed certificates.
3. In the **Confirm Certificate Authorization Change** popup, click **Proceed** to confirm that you wish to have the Symantec signing server automatically generate, sign, and install certificates on each controller device.
4. Enter the first and last name of the requestor of the certificate.
5. Enter the email address of the requestor of the certificate. This address is required because the signed certificate and a confirmation email are sent to the requestor via email; they are also made available through the customer portal.
6. Specify the validity period for the certificate. It can be 1, 2, or 3 years.
7. Enter a challenge phrase. The challenge phrase is your certificate password and is required when you renew or revoke a certificate.
8. Confirm your challenge phrase.
9. In the Certificate **Retrieve Interval** field, specify how often the Cisco vManage server checks if the Symantec signing server has sent the certificate.
10. Click **Save**.

To manually install certificates that the Symantec signing server has generated and signed:

1. Click the **Edit** button to the right of the **Controller Certificate Authorization** bar.
2. Click **Symantec Manual**.

3. In the **Confirm Certificate Authorization Change** popup, click **Proceed** to manually install certificates that the Symantec signing server has generated and signed.
4. Click **Save**.

To use enterprise root certificates:

1. Click the **Edit** button to the right of the **Controller Certificate Authorization** bar.
2. Click **Enterprise Root Certificate**.
3. In the **Confirm Certificate Authorization Change** popup, click **Proceed** to confirm that you wish to use enterprise root certificates.
4. In the **Certificate** box, either paste the certificate, or click **Select a file** and upload a file that contains the enterprise root certificate.
5. By default, the enterprise root certificate has the following properties: To view this information, issue the **show certificate signing-request decoded** command on a controller device, and check the output in the Subject line. For example:
 - Country: United States
 - State: California
 - City: San Jose
 - Organizational unit: ENB
 - Organization: CISCO
 - Domain Name: cisco.com
 - Email: cisco-cloudops-sdwan@cisco.com

```
vSmart# show certificate signing-request decoded
...
Subject: C=US, ST=California, L=San Jose, OU=ENB, O=CISCO, CN=vsmart-uuid
.cisco.com/emailAddress=cisco-cloudops-sdwan@cisco.com
...
```

To change one or more of the default CSR properties:

- a. Click **Set CSR Properties**.
 - b. Enter the domain name to include in the CSR. This domain name is appended to the certificate number (CN).
 - c. Enter the organizational unit (OU) to include in the CSR.
 - d. Enter the organization (O) to include in the CSR.
 - e. Enter the city (L), state (ST), and two-letter country code (C) to include in the CSR.
 - f. Enter the email address (emailAddress) of the certificate requestor.
 - g. Specify the validity period for the certificate. It can be 1, 2, or 3 years.
6. Click **Import & Save**.

Enforce Software Version on Devices

If you are using the Cisco SD-WAN hosted service, you can enforce a version of the Cisco SD-WAN software to run on a router when it first joins the overlay network. To do so:

1. Ensure that the software image for the desired device software version is present in the vManage software image repository:
 - a. In Cisco vManage, select the **Maintenance > Software Repository** screen.
The Software Repository screen opens and displays a table of software images. If the desired software image is present in the repository, continue with Step 2.
 - b. If you need to add a software image, click **Add New Software**.
 - c. Select the location from which to download the software images, either Cisco vManage, Remote Server, or Remote Server - vManage.
 - d. Select an x86-based or a MIPS-based software image.
 - e. Click **Add** to play the image in the repository.
2. In the **Administration > Settings** screen, click the **Edit** button to the right of the Enforce Software Version (ZTP) bar.
3. In the **Enforce Software Version** field, click **Enabled**.
4. From the **Version** drop-down, select the version of the software to enforce on the device when they join the network.
5. Click **Save**.

If you enable this feature on the Cisco vManage, any device joining the network is configured with the version of the software specified in the **Enforce Software Version** field regardless of whether the device was running a higher or lower version of Cisco SD-WAN software.

Banner

Use the Banner template for Cisco vBond Orchestrators, Cisco vManages, Cisco vSmart Controllers, Cisco vEdge devices, and s.

You can configure two different banner text strings, one to be displayed before the CLI login prompt on a Cisco SD-WAN device and the other to be displayed after a successful login to the device.

- To configure the banner text for login screens using Cisco vManage templates, create a Banner feature template to configure PIM parameters, as described in this topic.
- To configure a login banner for the Cisco vManage system, go to **Administration > Settings**.

Configure a Banner

1. In Cisco vManage, select the **Configuration > Templates** screen.
2. In the **Device** tab, click **Create Template**.
3. From the **Create Template** drop-down, select **From Feature Template**.

4. From the **Device Model** drop-down, select the type of device for which you are creating the template.
5. Click the **Additional Templates** tab located directly beneath the Description field, or scroll to the **Additional Templates** section.
6. From the **Banner** drop-down, click **Create Template**. The **Banner** template form is displayed. The top of the form contains fields for naming the template, and the bottom contains fields for defining Banner parameters.

The screenshot shows the Cisco vManage interface for creating a Banner template. The left sidebar contains a navigation menu with options: Dashboard, Monitor, Configuration (selected), Devices, Certificates, Templates, Policies, Security, CloudExpress, Cloud onRamp, Tools, Maintenance, Administration, and vAnalytics. The main content area is titled 'CONFIGURATION | TEMPLATES' and has tabs for 'Device' and 'Feature'. The breadcrumb trail is 'Feature Template > Add Template > Banner'. The form includes the following fields:

- Device Type:** vEdge Cloud
- Template Name:** (empty text input field)
- Description:** (empty text input field)

Below these fields is a section titled 'BASIC CONFIGURATION' with two rows:

- Login Banner:** A dropdown menu with a checkmark icon and a text input field.
- MOTD Banner:** A dropdown menu with a checkmark icon and a text input field.

A 'Save' button is visible in the bottom right corner of the form area.

7. In the **Template Name** field, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
8. In the **Template Description** field, enter a description of the template. The description can be up to 2048 characters and can contain only alphanumeric characters.

When you first open a feature template, for each parameter that has a default value, the scope is set to Default (indicated by a check mark), and the default setting or value is shown. To change the default or to enter a value, click the **Scope** drop-down to the left of the parameter field.

- To set a banner, configure the following parameters:

Table 2: Parameters to be configured while setting a banner:

Parameter Name	Description
MOTD Banner	On a Cisco vEdge device enter message-of-the-day text to display after a successful login. The string can be up to 2048 characters long. To insert a line break, type \n.
Login Banner	Enter text to display before the login prompt. The string can be up to 2048 characters long. To insert a line break, type \n.

- To save the feature template, click **Save**.

CLI equivalent:

```
banner(login text | motd text)
```

Release Information

Introduced in Cisco vManage NMS in Release 15.2.

Create a Custom Banner

To create a custom banner that is displayed after you log in to the Cisco vManage:

- Click the **Edit** button to the right of the Banner bar.
- In the **Enable Banner** field, click **Enabled**.
- In the **Banner Info** text box, enter the text string for the login banner or click **Select a File** to download a file that contains the text string.
- Click **Save**.

Collect Device Statistics

To enable or disable the collection of statistics for devices in the overlay network:

- Click the **Edit** button to the right of the **Statistics Settings** bar. By default, all statistics collection settings are enabled for all Cisco SD-WAN devices.
- To set statistics collection parameters for all devices in the network, click **Disable All** for the parameter you wish to disable statistics collection for. To return to the saved settings during an edit operation, click **Reset**. To return the saved settings to the factory-default settings, click **Restore Factory Default**.
- To set statistics collection parameters for individual devices in the network, click **Custom** to select devices on which to enable or disable statistics collection. The **Select Devices** popup screen opens listing the hostname and device IP of all devices in the network. Select one or more devices from the **Enabled**

Devices column on the left and click the arrow pointing right to move the device to the **Disabled Devices** column on the right. To move devices from the **Disabled Devices** to the **Enabled Devices** column, select one or more devices and click the arrow pointing left. To select all devices in the **Select Devices** popup screen, click the **Select All** checkbox in either window. Click **Done** when all selections are made.

4. Click **Save**.

Set the Time Interval to Collect Device Statistics

To set the time interval at which vManage NMS should collect statistics for devices in the overlay network, use the **Administration > Settings** screen.

1. Click the **Edit** button to the right of the Statistics Configuration bar. By default, statistics is collected for all Viptela devices every 30 minutes.
2. Click the up or down arrow in the **Collection Interval** drop-down to change the frequency at which to collect device statistics. The minimum time you can specify is 5 minutes and the maximum is 180 minutes.
3. Click **Save**.

Configure or Cancel vManage Server Maintenance Window

You can set or cancel the start and end times and the duration of the maintenance window for the vManage server.

1. In vManage NMS, select the **Administration > Settings** screen.
2. Click the **Edit** button to the right of the Maintenance Window bar.
To cancel the maintenance window, click **Cancel**.
3. Click the **Start date and time** drop-down, and select the date and time when the maintenance window will start.
4. Click the **End date and time** drop-down, and select the date and time when the maintenance window will end.
5. Click **Save**. The start and end times and the duration of the maintenance window are displayed in the Maintenance Window bar.

Two days before the start of the window, the vManage Dashboard displays a maintenance window alert notification.

Configure Basic System Parameters

Use the System template for all Cisco SD-WAN devices.

To configure system-wide parameters using vManage templates:

1. Create a **System** feature template to configure system parameters.
2. Create an **NTP** feature template to configure NTP servers and authentication.

3. Configure the organization name and Cisco vBond Orchestrator IP address on the vManage NMS. These settings are appended to the device templates when the templates are pushed to devices.

Create System Template

1. In vManage NMS, select the **Configuration ► Templates** screen.
2. In the **Device** tab, click **Create Template**.
3. From the **Create Template** drop-down, select **From Feature Template**.
4. From the **Device Model** drop-down, select the type of device for which you are creating the template.
5. To create a custom template for System, select the **Factory_Default_System_Template** and click **Create Template**. The System template form is displayed. The top of the form contains fields for naming the template, and the bottom contains fields for defining System parameters.
6. In the **Template Name** field, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
7. In the **Template Description** field, enter a description of the template. The description can be up to 2048 characters and can contain only alphanumeric characters.

When you first open a feature template, for each parameter that has a default value, the scope is set to Default (indicated by a check mark), and the default setting or value is shown. To change the default or to enter a value, click the scope drop-down to the left of the parameter field and select one of the following:

Table 3:

Parameter Scope	Scope Description
Device Specific (indicated by a host icon)	<p>Use a device-specific value for the parameter. For device-specific parameters, you cannot enter a value in the feature template. You enter the value when you attach a Cisco SD-WAN device to a device template .</p> <p>When you click Device Specific, the Enter Key box opens. This box displays a key, which is a unique string that identifies the parameter in a CSV file that you create. This file is an Excel spreadsheet that contains one column for each key. The header row contains the key names (one key per column), and each row after that corresponds to a device and defines the values of the keys for that device. You upload the CSV file when you attach a Cisco SD-WAN device to a device template.</p> <p>To change the default key, type a new string and move the cursor out of the Enter Key box.</p> <p>Examples of device-specific parameters are system IP address, hostname, GPS location, and site ID.</p>
Global (indicated by a globe icon)	<p>Enter a value for the parameter, and apply that value to all devices.</p> <p>Examples of parameters that you might apply globally to a group of devices are DNS server, syslog server, and interface MTUs.</p>

Basic System-Wide Configuration

To set up system-wide functionality on a Cisco SD-WAN device, select the **Basic Configuration** tab and then configure the following parameters. Parameters marked with an asterisk are required.

Table 4:

Parameter Field	Description
Site ID* (on routers, vManage NMSs, and vSmart controllers)	Enter the identifier of the site in the Cisco SD-WAN overlay network domain in which the device resides, such as a branch, campus, or data center. The site ID must be the same for all Cisco SD-WAN devices that reside in the same site. <i>Range:</i> 1 through 4294967295 ($2^{32} - 1$)
System IP*	Enter the system IP address for the Cisco SD-WAN device, in decimal four-part dotted notation. The system IP address provides a fixed location of the device in the overlay network and is a component of the device's TLOC address. It is used as the device's loopback address in the transport VPN (VPN 0). You cannot use this same address for another interface in VPN 0.
Timezone*	Select the timezone to use on the device.
Hostname	Enter a name for the Cisco SD-WAN device. It can be up to 32 characters.
Location	Enter a description of the location of the device. It can be up to 128 characters.
Device Groups	Enter the names of one or more groups to which the device belongs, separated by commas.
Controller Groups	List the vSmart controller groups to which the router belongs.
Description	Enter any additional descriptive information about the device.
Console Baud Rate	Select the baud rate of the console connection on the router. <i>Values:</i> 1200, 2400, 4800, 9600, 19200, 38400, 57600, 115200 baud or bits per second (bps) <i>Default:</i> 115200 bps
Maximum OMP Sessions	Set the maximum number of OMP sessions that a router can establish to a vSmart controller. <i>Range:</i> 0 through 100 <i>Default:</i> 2
Dedicated Core for TCP Optimization (optional, on vEdge 1000 and 2000 routers only)	Click on to carve out a separate CPU core to use for performing TCP optimization.

To save the feature template, click **Save**.

CLI equivalent:

```

system
clock timezone timezone
console-baud-rate rate
controller-group-list numbers
description text
device-groups group-name
host-name string

```

```
location string
max-omp-sessions number
site-id site-id
system-ip ip-address
tcp-optimization-enabled
```

To configure the DNS name or IP address of the vBond orchestrator in your overlay network, go to the **Administration** > **Settings** screen and click **vBond**.

Configure the GPS Location

To configure a device location, select the **GPS** tab and configure the following parameters. This location is used to place the device on the Cisco vManage network map. Setting the location also allows Cisco vManage to send a notification if the device is moved to another location.

Table 5:

Parameter Field	Description
Latitude	Enter the latitude of the device, in the format <i>decimal-degrees</i> .
Longitude	Enter the longitude of the device, in the format <i>decimal-degrees</i> .

To save the feature template, click **Save**.

feature-id="sdwan-vedge"> *CLI equivalent:*

```
system gps-location (latitude decimal-degrees | longitude decimal-degrees)
```

Configure Interface Trackers

To track the status of transport interfaces that connect to the internet, click the **Tracker** tab. Then click **Add New Tracker** and configure the following parameters:

Table 6:

Parameter Field	Description
Name	Name of the tracker. The name can be up to 128 alphanumeric characters. You can configure up to eight trackers.
Threshold	How long to wait for the probe to return a response before declaring that the transport interface is down. <i>Range:</i> 100 through 1000 milliseconds <i>Default:</i> 300 milliseconds
Interval	How often probes are sent to determine the status of the transport interface. <i>Range:</i> 10 through 600 seconds <i>Default:</i> 60 seconds (1 minute)
Multiplier	Number of times to resend probes before declaring that the transport interface is down. <i>Range:</i> 1 through 10 <i>Default:</i> 3
End Point Type: IP Address	IP address of the end point of the tunnel interface. This is the destination in the internet to which the router sends probes to determine the status of the transport interface.

Parameter Field	Description
End Point Type: DNS Name	DNS name of the end point of the tunnel interface. This is the destination in the internet to which the router sends probes to determine the status of the transport interface.

To save a tracker, click **Add**.

To save the feature template, click **Save**.

Configuration examples using CLI:

```
system tracker tracker-name
  endpoint-dns-name dns-name
  endpoint-ip ip-address
  interval seconds
  multiplier number
  threshold milliseconds
```

To apply a tracker to an interface, configure it in the VPN Interface Cellular, VPN Interface Ethernet, VPN Interface NAT Pool, or VPN Interface PPP configuration templates. You can apply only one tracker to an interface.

Configure Advanced Options

To configure additional system parameters, click the **Advanced** tab:

Table 7:

Parameter Name	Description
Control Session Policer Rate	Specify a maximum rate of DTLS control session traffic, to police the flow of control traffic. <i>Range:</i> 1 through 65535 pps <i>Default:</i> 300 pps
MTU of DTLS Tunnel	Specify the MTU size to use on the DTLS tunnels that send control traffic between Cisco SD-WAN devices. <i>Range:</i> 500 through 2000 bytes <i>Default:</i> 1024 bytes
Port Hopping	Click On to enable port hopping, or click Off to disable it. When a Cisco SD-WAN device is behind a NAT, port hopping rotates through a pool of preselected OMP port numbers (called base ports) to establish DTLS connections with other Cisco SD-WAN devices when a connection attempt is unsuccessful. The default base ports are 12346, 12366, 12386, 12406, and 12426. To modify the base ports, set a port offset value. To disable port hopping on an individual TLOC (tunnel interface), use the VPN Interface Ethernet configuration template. <i>Default:</i> Enabled (on routers); disabled (on vManage NMSs and vSmart controllers)
Port Offset	Enter a number by which to offset the base port number. Configure this option when multiple Cisco SD-WAN devices are behind a single NAT device, to ensure that each device uses a unique base port for DTLS connections. <i>Values:</i> 0 through 19
DNS Cache Timeout	Specify when to time out the vBond orchestrator addresses that have been cached by the device. <i>Range:</i> 1 through 30 minutes <i>Default:</i> 30 minutes

Parameter Name	Description
Track Transport	Click On to regularly check whether the DTLS connection between the device and a vBond orchestrator is up. Click Off to disable checking. By default, transport checking is enabled
Local vBond (only on routers acting as vBond orchestrators)	Click On to configure the router to act as a vBond orchestrator. Then specify the DNS name for the vBond orchestrator or its IP address, in decimal four-part dotted notation.
Track Interface	Set the tag string to include in routes associated with a network that is connected to a non-operational interface. <i>Range:</i> 1 through 4294967295
Multicast Buffer	Specify the percentage of interface bandwidth that multicast traffic can use. <i>Range:</i> 5% through 100% <i>Default:</i> 20%
USB Controller (on vEdge 1000 and 2000 series routers only)	Click On to enable or click Off to disable the USB controller, which drives the external USB ports. If you enable the USB controller, the vEdge router reboots when you attach the device template to the device. <i>Default:</i> Disabled
Gateway Tracking	Click On to enable or click Off to Disable tracking of default gateway. Gateway tracking determines, for static routes, whether the next hop is reachable before adding that route to the device's route table. <i>Default:</i> Enabled
Host Policer Rate (on vEdge routers only)	Specify the maximum rate at which a policer delivers packets to the control plane. <i>Range:</i> 1000 through 20000 pps <i>Default:</i> 5000 pps
ICMP Error Rate (on vEdge routers only)	Specify how many ICMP error messages a policer can generate or receive. <i>Range:</i> 1 through 200 pps <i>Default:</i> 100 pps
Allow Same-Site Tunnel (on vEdge routers only)	Click On to allow tunnels to be formed between vEdge routers in the same site. Note that no BFD sessions are established between the two collocated vEdge routers. <i>Default:</i> Off
Route Consistency Check (on vEdge routers only)	Click On to check whether the IPv4 routes in the device's route and forwarding table are consistent.
Collect Admin Tech on Reboot	Click On to collect admin-tech information when the device reboots.
Idle Timeout	Set how long the CLI is inactive on a device before the user is logged out. If a user is connected to the device via an SSH connection, the SSH connection is closed after this time expires. <i>Range:</i> 0 through 300 seconds <i>Default:</i> CLI session does not time out
Eco-Friendly Mode (on vEdge Cloud routers only)	Click On to configure a Cloud router not to use its CPU minimally or not at all when the router is not processing any packets.

To save the feature template, click **Save**.

CLI equivalent:

```
system
  admin-tech-on-failure allow-same-site-tunnels
  control-session-pps rate eco-friendly-mode
  host-policer-pps rate
```

```

icmp-error-pps rate

idle-timeout seconds multicast-buffer-percent percentage

port-hop port-offset number route-consistency-check
system-tunnel-mtu bytes timer
dns-cache-timeout minutes track-default-gateway
track-interface-tag number

track-transport upgrade-confirm minutes [no] usb-controller (on Cisco vEdge 1000 and
Cisco vEdge2000 routers only)
vbond (dns-name | ip-address) local (on Cisco vEdge routers acting as Cisco vBond
controllers)

```

Release Information

Introduced in vManage NMS in Release 15.2. In Releases 15.3.8 and 15.4.3, add Track Interface field. In Release 17.1.0, add Route Consistency Check and Collect Admin Tech on Reboot fields. In Release 17.2.0, add support for CLI idle timeout and ecofriendly mode. In Release 17.2.2, add support for interface status tracking.

Configure Global Parameters

Use the Global Settings template to configure global parameters for all Cisco SD-WAN devices.

To configure global settings using vManage:

1. Create a feature template to configure global settings.
2. Create a device template and include the Global Settings feature template.
3. (Recommended) Before applying the device template to a device, use the View Configuration Differences feature to review the differences between the configuration currently on the device and the configuration to be sent to the device (overwriting its existing configuration).

Limitations

SD-WAN can apply the global settings feature template only to devices running Cisco IOS XE Gibraltar 17.2 or later.

Create Global Settings Feature Template

1. In vManage, select **Configuration** (gear icon) ► **Templates**.
2. Click the **Feature** tab.
3. Click **Add Template**.
4. In the left pane, select a device type.
5. In the right pane, select the **Global Settings** template.
6. Provide a name and description for the template.

7. For each of the parameters, use the default or set custom values as desired.

Parameter	Description
Services	
HTTP Server	Enable/disable HTTP server.
HTTPS Server	Enable/disable secure HTTPS server.
Passive FTP	Enable/disable passive FTP.
IP Domain-Lookup	Enable/disable domain name server (DNS) lookup.
Arp Proxy	Enable/disable proxy ARP.
RSH/RCP	Enable/disable remote shell (RSH) and remote copy (RCP) on the device.
Telnet (Outbound)	Enable/disable outbound telnet.
CDP	Enable/disable Cisco Discovery Protocol (CDP).
Other Settings	
TCP Keepalives (In)	Enable/disable generating keepalives on idle incoming network connections.
TCP Keepalives (Out)	Enable/disable generating keepalives on idle outgoing network connections.
TCP Small Servers	Enable/disable small TCP servers (for example, ECHO).
UDP Small Servers	Enable/disable small UDP servers (for example, ECHO).
Console Logging	Enable/disable console logging. By default, the router sends all log messages to its console port.
IP Source Routing	Enable/disable the originator of a packet to determine which path to use to get to the destination.
VTY Line Logging	Enable/disable the device to display log messages to a VTY session in real time.
SNMP IFINDEX Persist	Enable/disable SNMP IFINDEX persistence, which provides an interface index (ifIndex) value that is retained and used when the device reboots.
Ignore BOOTP	Enable/disable BOOTP server. This enables the device to listen for the bootp packet that comes in sourced from 0.0.0.0. When disabled, the device ignores these packets.
NAT 64	
UDP Timeout	NAT64 translation timeout for UDP Range: 1 to 65536 (seconds)
TCP Timeout	NAT64 translation timeout for TCP Range: 1 to 65536 (seconds)

Parameter	Description
HTTP Authentication	
HTTP Authentication	HTTP authentication mode Possible values: Local, AAA

- Enter a name for the template and click **Save**.

CLI Equivalent

Services:

```
[no] ip http server
[no] ip http secure-server
[no] ip ftp passive
[no] ip domain lookup
[no] ip arp proxy disable
[no] ip rcmd rsh-enable
[no] ip rcmd rcp-enable
(Telnet outbound enable) line vty 0 4, transport input telnet ssh
(Telnet outbound disable) line vty 0 4, transport input ssh
[no] cdp run enable
```

Other settings:

```
[no] service tcp-keepalives-in
[no] service tcp-keepalives-out
[no] service tcp-small-servers
[no] service udp-small-server
[no] logging console
[no] ip source-route
[no] logging monitor
[no] snmp-server ifindex persist
[no] ip bootp server
```

NAT 64:

```
nat64 translation timeout udp timeout
nat64 translation timeout tcp timeout
```

HTTP Authentication:

```
ip http authentication {local | aaa}
```

Configure NTP using Cisco vManage

Configure network time protocol (NTP) servers on your devices in order to synchronize time across all devices in the Cisco Overlay Network. You can configure up to four NTP servers, and they must all be located or reachable in the same VPN.

Other devices are allowed to ask a Cisco SD-WAN device for the time, but no devices are allowed to use the Cisco SD-WAN device as an NTP server.

To configure NTP using Cisco vManage templates:

- Create an NTP feature template to configure NTP parameters, as described in this article.

2. Configure the timezone in the System template.

Navigate to the Template Screen and Name the Template

1. In Cisco vManage NMS, select the **Configuration > Templates** screen.
2. In the **Device** tab, click **Create Template**.
3. From the **Create Template** drop-down, select **From Feature Template**.
4. From the **Device Model** drop-down, select the type of device for which you are creating the template.
5. Select the **Basic Information** tab.
6. Under **Additional System Templates**, located to the right of the screen, click **NTP**.
7. From the **NTP** drop-down, click **Create Template**. The NTP template form is displayed. The top of the form contains fields for naming the template, and the bottom contains fields for defining NTP parameters.
8. In the **Template Name** field, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
9. In the **Template Description** field, enter a description of the template. The description can be up to 2048 characters and can contain only alphanumeric characters.

When you first open a feature template, for each parameter that has a default value, the scope is set to Default (indicated by a check mark), and the default setting or value is shown. To change the default or to enter a value, click the scope drop-down to the left of the parameter field and select one of the following:

Table 8:

Parameter Scope	Scope Description
Device Specific (indicated by a host icon)	<p>Use a device-specific value for the parameter. For device-specific parameters, you cannot enter a value in the feature template. You enter the value when you attach a Viptela device to a device template .</p> <p>When you click Device Specific, the Enter Key box opens. This box displays a key, which is a unique string that identifies the parameter in a CSV file that you create. This file is an Excel spreadsheet that contains one column for each key. The header row contains the key names (one key per column), and each row after that corresponds to a device and defines the values of the keys for that device. You upload the CSV file when you attach a Viptela device to a device template. For more information, see Create a Template Variables Spreadsheet .</p> <p>To change the default key, type a new string and move the cursor out of the Enter Key box.</p> <p>Examples of device-specific parameters are system IP address, hostname, GPS location, and site ID.</p>
Global (indicated by a globe icon)	<p>Enter a value for the parameter, and apply that value to all devices.</p> <p>Examples of parameters that you might apply globally to a group of devices are DNS server, syslog server, and interface MTUs.</p>

Configure NTP Servers

To configure NTP servers, select the Server tab and click **Add New Server**. Then configure the following parameters. Parameters marked with an asterisk are required to configure NTP.

Table 9:

Parameter Name	Description
Hostname/IP Address*	Enter the IP address of an NTP server or of a DNS server that knows how to reach the NTP server.
Authentication Key*	Specify the MD5 key associated with the NTP server, to enable MD5 authentication. For the key to work, you must mark it as trusted in the Trusted Keys field, under the Authentication tab (discussed below).
VPN ID*	Enter the number of the VPN to use to reach the NTP server or the VPN in which the NTP server is located. If you configure multiple NTP servers, they must all be located or reachable in the same VPN. <i>Range: 0 through 65530</i>
Version*	Enter the version number of the NTP protocol software. <i>Range: 1 through 4</i> <i>Default: 4</i>
Source Interface	Enter the name of a specific interface to use for outgoing NTP packets. The interface must be located in the same VPN as the NTP server. If it is not, the configuration is ignored.
Prefer	Click On if multiple NTP servers are at the same stratum level and you want one to be preferred. For servers at different stratum levels, the software chooses the one with the highest stratum level.

To add the NTP server, click **Add**.

To add another NTP server, click **Add New Server**. You can configure up to four NTP servers. The Cisco SD-WAN software uses the server at the highest stratum level.

To edit an NTP server, click the pencil icon to the right of the entry.

To delete an NTP server, click the trash icon to the right of the entry.

To save the feature template, click **Save**.

CLI equivalent:

```
system ntp
  server (dns-server-address | ip-address)
    key key-id
  prefer
  source-interface interface-name
  version number
  vpn vpn-id
```

Configure NTP Authentication

To configure authentication keys used to authenticate NTP servers, in the **Authentication** tab, click the **Authentication Key** tab. Then click Add New Authentication Key, and configure the following parameters. Parameters marked with an asterisk are required to configure NTP.

Table 10:

Parameter Name	Description
Authentication Key*	Select the following values: <ul style="list-style-type: none"> • Authentication Key—Enter an MD5 key ID. It can be a number from 1 through 65535. • Authentication Value—Enter either a cleartext key or an AES-encrypted key.
Authentication Value*	Enter an MD5 authentication key. For the key to be used, you must designate it as trusted. To associate a key with a server, enter the same value as you use for the the Authentication Key field on the Server tab.

To configure trusted keys used to authenticate NTP servers, in the Authentication tab, click the **Trusted Keys** tab and configure the following parameters;

Table 11:

Parameter Name	Description
Trusted Keys*	Enter the MD5 authentication key to designate the key as trustworthy. To associate this key with a server, enter the same value as you use for the the Authentication Key field on the Server tab.

CLI equivalent:

```

system
 ntp
  keys
    authentication key-id md5 md5-key
    trusted key-id
  
```

Configure NTP using CLI

Configure Network-Wide Time with NTP

To coordinate and synchronize time across all devices in the Cisco SD-WAN overlay network, configure the IP address or DNS server address of an NTP server on each device. If necessary, specify the VPN through which the server is reachable.

```

vEdge(config)# system ntp server (dns-server-address | ipv4-address)
vEdge(config-system)# ntp server (dns-server-address | ipv4-address) vpnvpn-id
  
```

You can configure up to four NTP servers, and they must all be located or reachable in the same VPN. The software uses the server at the highest stratum level. If more than one server is at the same stratum level, you can configure the preference to use a specific server:

```
vEdge(config-ntp)# ntp
server (dns-server-address | ipv4-address) prefer
```

You can configure an MD5 authentication key to use as a password to access an NTP server:

```
vEdge(config-system)# ntp keys
vEdge(config-keys)# authentication key-id md5 md5-key
```

key-id is a number that identifies the MD5 authentication key. It can be a number from 1 through 65535.

md5-key is the MD5 authentication key. You can enter it as cleartext or as an AES-encrypted key.

To use an MD5 authentication key for an NTP server, the key must be configured to be trusted:

```
vEdge(config-system)# ntp keys trusted key-id
```

Finally, associate the MD5 authentication key with the NTP time server:

```
vEdge(config-system)# ntp server (dns-server-address | ipv4-address) key key-id
```

You can configure NTP packets to exit from a specific interface on the router. The interface must be located in the same VPN as the NTP server. If it is not, the configuration is ignored.

```
vEdge(config-system)# ntp server (dns-server-address | ipv4-address) source-interface
interface-name
```

The following example displays configuration three NTP servers. One of the NTP servers is at the NTP pool project at the Network Time Foundation and uses no authentication. The other two are internal servers and are configured with MD5 authentication:

```
vEdge# show running-config system ntp
system
ntp
keys
authentication 1001 md5 $4$KXLzYT9k6M8zj4BgLEFXKw==
authentication 1002 md5 $4$KXLzYTzk6M8zj4BgLEFXKw==
authentication 1003 md5 $4$KXLzYT1k6M8zj4BgLEFXKw==
trusted 1001 1002
!
server 192.168.15.243
key 1001
vpn 512
version 4
exit
server 192.168.15.242
key 1002
vpn 512
version 4
exit
server us.pool.ntp.org
vpn 512
version 4
exit
!
```

Configuring NTP on a Cisco SD-WAN device allows that device to contact NTP servers to synchronize time. Other devices are allowed to ask for the time, but no devices are allowed to use the Cisco SD-WAN as an NTP server.

Configuring Time Using CLI on Cisco vEdge Device

Configure the Timezone

The default timezone on all Cisco vEdge devices is UTC. If your devices are located in multiple timezones (and even if they are not), we recommend that you use the default timezone, which is UTC, on all devices so that the times in all logging and archive files are consistent.

To change the timezone on a device:

```
vEdge(config-system)# clock timezone timezone
```

Set the Time Locally

For Cisco vEdge devices that are part of a test or local network, you can set the time locally without using NTP because you do not need to ensure that time is synchronized across an entire network of devices. You can also set the time locally on any device as it is joining the network, in addition to configuring an NTP server, and this time will be overwritten by the official NTP time once the device contacts the NTP server.

To set the local time and date, issue the following operational commands:

```
vEdge# clock set time hh:mm:ss[.sss]  
vEdge# clock set date ccyy-mm-dd
```

You can also issue these commands as a single command:

```
vEdge# clock set date ccyy-mm-dd time hh:mm:ss[.sss]
```

or

```
vEdge# clock set time hh:mm:ss[.sss] date ccyy-mm-dd
```

To set the timezone, specify it in the configuration:

```
vEdge(config)# system clock timezone timezone
```

Configure GPS Using Cisco vManage

Use the GPS template for all Cisco cellular routers running Cisco SD-WAN software.

For Cisco devices running Cisco SD-WAN software, you can configure the GPS and National Marine Electronics Association (NMEA) streaming. You enable both these features to allow 4G LTE routers to obtain GPS coordinates.

Navigate to the Template Screen and Name the Template

1. In Cisco vManage NMS, select the **Configuration** > **Templates** screen.
2. In the Device tab, click **Create Template**.
3. From the **Create Template** drop-down, select **From Feature Template**.
4. From the **Device Model** drop-down, select the type of device for which you are creating the template.
5. Select the **Cellular** tab.
6. In **Additional Cellular Controller** Templates, click **GPS**.

7. To create a custom template for GPS, click the **GPS** drop-down and then click **Create Template**. The GPS template form is displayed. The top of the form contains fields for naming the template, and the bottom contains fields for defining GPS parameters.
8. In the **Template Name** field, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
9. In the **Template Description** field, enter a description of the template. The description can be up to 2048 characters and can contain only alphanumeric characters.

When you first open a feature template, for each parameter that has a default value, the scope is set to Default (indicated by a check mark), and the default setting or value is shown. To change the default or to enter a value, click the scope drop-down to the left of the parameter field and select either **Device Specific** or **Global**.

Configure GPS

To configure GPS parameters for the cellular router, configure the following parameters. Parameters marked with an asterisk are required to configure the GPS feature.

Table 12:

Parameter Name	Description
GPS	Click On to enable the GPS feature on the router.
GPS Mode	Select the GPS mode: <ul style="list-style-type: none"> • MS-based—Use mobile station–based assistance, also called assisted GPS mode, when determining position. In this mode, cell tower data is used to enhance the quality and precision in determining location, which is useful when satellite signals are poor. • Standalone—Use satellite information when determining position.
NMEA	Click On to enable the use of NMEA streams to help in determining position. NMEA streams data from the router's 4G LTE NIM to any marine device, such as a Windows-based PC, that is running a commercially available GPS-based application.
Source Address	Enter the IP address of the interface that connects to the router's NIM.
Destination Address	Enter the IP address of the marine NMEA server.
Destination Port	Enter the number of the port to use to send NMEA data to the server.

To save the feature template, click **Save**.

Release Information

Introduced in Cisco vManage Release 18.1.1.

Configure GPS Using CLI on Cisco vEdge Device

Configuring geographic location for a device by setting its latitude and longitude allows the device to be placed properly on the Cisco vManage network map.

To set a device's latitude and longitude:

```
vEdge(config-system)# gps-location latitude degrees.minutes-and-seconds longitude degrees.minutes-and-seconds
```

You can also set these values using two separate commands:

```
vEdge(config-system)# gps-location latitude degrees . minutes-and-seconds  
vEdge(config-system)# gps-location longitude degrees . minutes-and-seconds
```

For example:

```
vEdge(config-system)# gps-location latitude 37.0000 longitude 122.0600  
or  
vEdge(config-system)# gps-location latitude 37.000  
vEdge(config-system)# gps-location longitude 122.0600  
vEdge(config-system)# show full-configuration  
system  
  host-name          vEdge  
  gps-location latitude 36.972  
  gps-location longitude 122.0263  
  ...
```

You can also configure a text description of the device's location:

```
vEdge(config-system)# location "description of location"
```

For example:

```
vEdge(config-system)# location "UCSC in Santa Cruz, California"  
vEdge(config-system)# show full-configuration  
system  
  host-name          vEdge  
  location           "UCSC in Santa Cruz, California"  
  gps-location latitude 37.0000  
  gps-location longitude 122.0600  
  ...
```

Configure System Logging

System logging operations use a mechanism similar to the UNIX syslog command to record system-wide, high-level operations that occur on Cisco SD-WAN devices in the overlay network. The log levels (priorities) of the messages are the same as standard UNIX commands, and you can configure the priority of syslog messages. Cisco SD-WAN devices can send log messages to a UNIX-style syslog service.

Cisco vEdge devices send syslog messages to syslog servers using UDP. TCP is not supported.

The syslog service accepts messages and stores them in files on the Cisco SD-WAN device or to a remote host.

Syslog Message Format, Syslog Message Levels, and System Log Files

Syslog Message Format

Syslog messages begin with a percent sign (%) and following are the syslog message formats:

- Syslog message format

seq no:timestamp: %facility-severity-MENEMONIC:description (hostname-n)

The field descriptions of syslog messages are:

Table 13: Field Descriptions of Syslog Message Format

Field	Description
facility	Sets the logging facility to a value other than 20, which UNIX systems expect.
severity	The importance or severity of the message is categorized by the numerical code from 0 through 7. A lower number in this range indicates greater severity of the system condition.
description	A text string that describes the condition of syslog server. This portion of the syslog message sometimes includes IP addresses, interface names, port numbers, or usernames.

Usually, the syslog messages are preceded by extra text.

- The following is an example of a system logging message preceded by a priority value, sequence number, and time stamp:

*<45>10: polaris-user1: *Jun 21 10:76:84.100: %LINK-5-CHANGED: Interface GigabitEthernet0/0, changed state to administratively down*

Syslog Message Levels

All syslog messages are associated with priority levels that indicate the severity of syslog messages to save. The default priority value is "informational", so by default, all syslog messages are recorded. The priority level can be one of the following in order of decreasing severity:

- Emergency—System is unusable (corresponds to syslog severity 0).
- Alert—Ensure that you act immediately (corresponds to syslog severity 1).
- Critical—A serious condition (corresponds to syslog severity 2).
- Error—An error condition that does not fully impair system usability (corresponds to syslog severity 3).
- Warning—A minor error condition (corresponds to syslog severity 4).
- Notice—A normal, but significant condition (corresponds to syslog severity 5).
- Informational—Routine condition (the default) (corresponds to syslog severity 6).

System Log Files

All syslog messages that are at or above the default or configured priority value are recorded in a number of files in the `/var/log` directory on the local device of the syslog server. The following are the contents of the log files:

- `auth.log`—Login, logout, and superuser access events, and usage of authorization systems
- `kern.log`—Kernel messages
- `messages.log`—Consolidated log file that contains syslog messages from all sources.
- `vconfd.log`—All configuration-related syslog messages
- `vdebug.log`—All debug messages for modules whose debugging is turned on and all syslog messages that are above the default priority value. The debug log messages support various levels of logging based on the module. The different modules implement the logging levels differently. For example, the system manager (`sysmgr`) has two logging levels (on and off), while the chassis manager (`chmgr`) has four different logging levels (off, low, normal, and high). You cannot send debug messages to a remote host. Therefore, to enable debugging, use the **debug** operational command.
- `vsyslog.log`—All syslog messages from Cisco SD-WAN processes (daemons) that are above the configured priority value. The default priority value is "informational", so by default, all "notice", "warning", "error", "critical", "alert", and "emergency" syslog messages are saved.
- `vmanage-syslog.log`—Cisco vManage NMS Audit log messages

The following are the standard LINUX files that Cisco SD-WAN does not use and are available in the `/var/log` directory.

- `cron.log`
- `debug.log`
- `lpr.log`
- `mail.log`
- `syslog`

The messages sent to syslog files are not rate-limited and consequently:

- A storage limit of 10 log files with a capacity of up to 16 MB size is set for each syslog file.
 - When the storage capacity exceeds the 16 MB size limit, the log file is saved as a `.GZ` file along with the date appended to it.
 - When the storage limit exceeds 10 log files, the oldest log file is dropped.
- If many syslog messages are generated in a short span of time, the overflowing messages are buffered and queued to be stored in the syslog file.

For repeating syslog messages or identical messages that occur multiple times in succession, only one copy of the message is placed in the syslog file. The message is annotated to indicate the number of times the message occurred.

The maximum length of a log message is 1024 bytes. The longer messages are truncated.

The maximum length of a log message for Cisco vManage NMS audit logs is 1024 bytes. The longer messages are truncated into smaller fragments and each of these fragments are indicated by an identifier. The identifiers are, fragment 1/2, fragment 2/2, and so on. For example, a long audit log message when truncated into smaller fragments appears as:

```
local6.info: 18-Oct-2020 17:42:07 vm10 maintenance-fragment-1/2: {"logid":
"d9ed576a-43ae-49ce-921b-a51c1ed40698", "entry_time":
1576605512190, "statcycletime" 34542398334245, "logmodule":"maintenance", "logfeature":
"upgrade", "loguser": "admin", "logusersrcip":
"10.0.1.1", "logmessage": "Device validation Upgrade to version - Validation success",
"logdeviceid":"Validation", "auditdetails" :
["[18-Oct-2020 17:42:08 UTC] Published messages to vmanage(s)", "auditdetails":["[18-Oct-2020
17:42:07 UTC] Software image: vmanage-99.99.999-
x86_64.tar.gz", "Software image download may take up to 60}

local6.info: 18-Oct-2020 17:42:07 vm10 maintenance-fragment-2/2: { minutes", "logprocessid":
"software_install-7de0ec44-d290-4429-b24532435324", "tenant":, "default"}
```

The syslog messages related to AAA authentication and Netconf CLI access and usage are placed in the auth.log and messages.log files. Each time a Cisco vManage NMS logs into a router to retrieve statistics and status information and to push files to the router, the router generates AAA and Netconf log messages. So, over time, these messages can fill the log files. To prevent these messages from filling the log files, you can disable the logging of AAA and Netconf syslog messages by using the following commands from Cisco vManage NMS:

Disable logging of AAA and Netconf Syslog Messages

1. vManage# **config**
Enters the configuration mode terminal
2. vManage(config)# **system aaa logs**
Configures the logging of AAA and Netconf system logging (syslog) messages
3. vManage(config-logs)# **audit-disable**
Disable logging of AAA events
4. vManage(config-logs)# **netconf-disable**
Disable logging of Netconf events
5. vManage(config-logs)# **commit**
Commit complete.

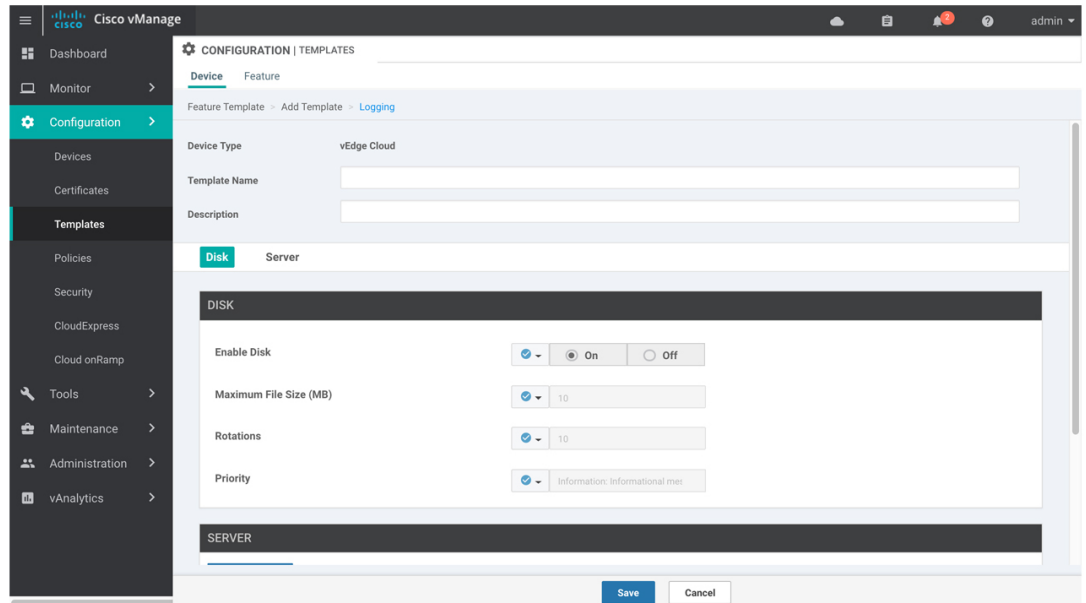
Configure Logging Using Cisco vManage

Use the Logging template for all Cisco SD-WANs to configure logging to either the local hard drive or a remote host.

Navigate to the Template Screen and Name the Template

1. In vManage NMS, select the **Configuration ► Templates** screen.
2. In the **Device** tab, click **Create Template**.
3. From the **Create Template** drop-down, select **From Feature Template**.

4. From the **Device Model** drop-down, select the type of device for which you are creating the template.
5. To create a custom template for Logging, select the **Factory Default Logging Template** and click **Create Template**. The Logging template form is displayed. The top of the form contains fields for naming the template, and the bottom contains fields for defining Logging parameters. You may need to click a tab or the plus sign (+) to display additional fields.



369421

6. In the **Template Name** field, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
7. In the **Template Description** field, enter a description of the template. The description can be up to 2048 characters and can contain only alphanumeric characters.

When you first open a feature template, for each parameter that has a default value, the scope is set to Default (indicated by a check mark), and the default setting or value is shown. To change the default or to enter a value, click the scope drop-down to the left of the parameter field.

Minimum Logging Configuration

The following logging parameters are configured by default:

- Log event notification system log (syslog) messages are logged to a file on the local device's hard disk, at a priority level of "information."
- Log files are placed in the directory /var/log on the local device.
- Log files are readable by the "admin" user.

Configure Logging to the Local Disk

To configure logging of event notification system log messages to the local device's hard disk, select the **Disk** tab and configure the following parameters:

Table 14:

Parameter Name	Description
Enable Disk	Click On to allow syslog messages to be saved in a file on the local hard disk, or click Off to disallow it. By default, logging to a local disk file is enabled on all Viptela devices.
Maximum File Size	Enter the maximum size of syslog files. Syslog files are rotated on an hourly basis based on the file's size. When the file size exceeds configured value, the file is rotated and the syslogd process is notified. <i>Range:</i> 1 through 20 MB <i>Default:</i> 10 MB
Rotations	Enter the number of syslog files to create before discarding the oldest files. <i>Range:</i> 1 through 10 <i>Default:</i> 10
Priority	Select the priority level of the syslog message to save to the log files. The severity indicates the seriousness of the event that generated the message. The default priority value is "informational", so, by default, all syslog messages are recorded. The priority level can be one of the following (in order of decreasing severity): <ul style="list-style-type: none"> • Emergency—System is unusable (corresponds to syslog severity 0). • Alert—Action must be taken immediately (corresponds to syslog severity 1). • Critical—Critical: A serious condition (corresponds to syslog severity 2). • Error—An error condition that does not fully impair system usability (corresponds to syslog severity 3). • Warning—A minor error condition (corresponds to syslog severity 4). • Notice—A normal, but significant condition (corresponds to syslog severity 5). • Informational—Routine condition (the default) (corresponds to syslog severity 6).

To save the feature template, click **Save**.

CLI equivalent:

```

system
 logging
  disk
  enable
  file
  rotate number size megabytes priority priority

```

Configure Logging to Remote Servers

To configure logging of event notification system log messages to a remote server, click the **Server** tab. Then click **Add New Server** and configure the following parameters:

Table 15:

Parameter Name	Description
Hostname/IP Address	Enter the DNS name, hostname, or IP address of the system on which to store syslog messages. To add another syslog server, click the plus sign (+). To delete a syslog server, click the trash icon to the right of the entry.
VPN ID	Enter the identifier of the VPN in which the syslog server is located or through which the syslog server can be reached. <i>Range:</i> 0 through 65530

Parameter Name	Description
Source Interface	Enter the specific interface to use for outgoing system log messages. The interface must be located in the same VPN as the syslog server. Otherwise, the configuration is ignored. If you configure multiple syslog servers, the source interface must be the same for all of them.
Priority	Select the severity of the syslog message to save. The severity indicates the seriousness of the event that generated the message. <i>priority</i> can be one of the following: <ul style="list-style-type: none"> Emergency—System is unusable (corresponds to syslog severity 0). Alert— Action must be taken immediately (corresponds to syslog severity 1). Critical—Critical: A serious condition (corresponds to syslog severity 2). Error—An error condition that does not fully impair system usability (corresponds to syslog severity 3). Warning—A minor error condition (corresponds to syslog severity 4). Notice—A normal, but significant condition (corresponds to syslog severity 5). Informational—Routine condition (the default) (corresponds to syslog severity 6). <p>Click Add to save the logging server.</p>

To edit a logging server, click the pencil icon to the right of the entry.

To remove a logging server, click the trash icon to the right of the entry.

To save the feature template, click **Save**.

CLI equivalent:

```

system
 logging
  server (dns-name | hostname | ip-address)
  priority priority
  source-interface interface-name
  vpn vpn-id

```

Release Information

Introduced in Cisco vManage NMS in Release 15.2.

Export Cisco vManage NMS Audit Log to Syslog Server

Table 16: Feature History

Feature Name	Release Information	Description
Export vManage Audit Log as Syslog	Cisco SD-WAN Release 20.3.1 Cisco vManage Release 20.3.1	The Cisco vManage NMS exports audit logs in syslog message format to a configured external syslog server. This feature allows you to consolidate and store network activity logs in a central location.

On Cisco XE SD-WAN devices and Cisco vEdge devices, you can log event notification system log (syslog) messages to files on a local device, or to files on a remote host using CLI. These event notification logs are

converted to system log files and exported to the syslog server. You can then retrieve system log information from the syslog server.

Configure System Logging Using CLI

Log Syslog Messages to a Local Device

By default, a priority level of “information” is enabled when you log syslog messages to a file on a local device. Use the following commands:

1. logging disk

Logs syslog messages on a hard disk

Example:

```
vm01(config-system)# logging disk
```

2. enable

Enables logging to a disk

Example:

```
vm01(config-logging-disk)# enable
```

3. file size *size*

Specifies the size of syslog files in megabytes (MB) By default, the syslog files are 10 MB. You can configure the size of syslog files to be 1–20 MB.

Example:

```
vm01(config-logging-disk)# file size 3
```

4. file rotate *number*

Rotates syslog files on an hourly basis based on the size of the file By default, 10 syslog files are created. You can configure the rotate command to be a number from 1 through 10.

Example:

```
vm01(config-logging-disk)# file rotate 3
```

For more information about logging disk commands, see the [logging disk](#) command.

Log Syslog Messages to a Remote Device

To log event notification system log (syslog) messages to a remote host, use the following commands:

1. logging server

Logs syslog messages to a remote host or syslog server You can configure the name of the server by DNS name, hostname, or IP address. You can configure up to four syslog servers.

Example:

```
vm01(config-system)# logging server 192.168.0.1
```

2. (Optional) vpn *vpn-id*

Specifies the VPN ID of the syslog server

3. (Optional) **source interface** *interface-name*

Specifies the source interface to reach the syslog server. The interface name can be a physical interface or a sub-interface (a VLAN-tagged interface). Ensure that the interface is located in the same VPN as the syslog server. Otherwise, the configuration is ignored. If you configure multiple syslog servers, the source interface must be the same for all of them.

Example:

```
vm01(config-server-192.168.0.1)# source interface eth0
```

4. **priority** *priority*

Specifies the severity of the syslog message to be saved. The default priority value is "informational" and by default, all syslog messages are recorded.

Example:

In the following example, set the syslog priority to log alert conditions.

```
vm01(config-server-192.168.0.1)# priority alert
```

If the syslog server is unreachable, the system suspends sending syslog messages for 180 seconds. When the server becomes reachable, logging resumes. For more information about logging server commands, see the [logging server](#) command.

View System Logging Information

To view system log settings after logging syslog messages to a remote host, use the **show logging** command. For example:

```
vm01(config-server-192.168.0.1)# show logging
```

```
System logging
  server 192.168.0.1
  source interface eth0
  exit
!
!
```

To view the contents of the syslog file, use the **show log** command. For example:

```
vm01(config-server-192.168.0.1)# show log nms/vmanage-syslog.log tail 10
```

To view the configured system logging settings from Cisco vManage, see [Audit Log](#).

To view device-specific syslog files from Cisco vManage, perform the following steps:

1. In the **Administration > Settings** screen, ensure that you enable **Data Stream**.
2. From the **Monitor > Network** screen, choose a Cisco vEdge device.
3. Click **Troubleshooting** from the left pane.
4. From the Logs pane, click **Debug Log**.
5. In the **Log Files** field, select a name of the log file. The lower part of the screen displays the log information.

SSH Terminal

Use the SSH Terminal screen to establish an SSH session to a Cisco vEdge device. From an SSH session, you can issue CLI commands on a Cisco vEdge device.

Establish an SSH Session to a Device

To establish an SSH session to a device:

1. From the left pane, select the device on which to collect statistics:
 - a. Select the device group to which the device belongs.
 - b. If needed, sort the device list by its status, hostname, system IP, site ID, or device type.
 - c. Click on the device to select it.
2. Enter the username and password to log in to the device.

You can now issue CLI commands to monitor or configure the device.

Tenant Management

Use the Tenant Management screen to add tenants to a Cisco vManage server that is operating in multitenant mode.

Add a Tenant

1. In the left pane, click the **Add Tenant** button.
2. In the **Add Tenant** window:
 - a. Enter a name for the tenant. It can be up to 128 characters and can contain only alphanumeric characters.
 - b. Enter a description for the tenant. It can be up to 256 characters and can contain only alphanumeric characters.
 - c. Enter the name of the organization. The name is case-sensitive. It is the name in the certificates for all Cisco SD-WAN network devices, and it must be identical on all devices in the overlay network.
 - d. In the URL subdomain field, enter the domain name for the tenant. The domain name must include the provider's domain name. You must also configure this same domain name when you enable multitenancy mode, in **vManage Administration > Settings > Tenancy Mode**
 - e. Click **Save**.
3. The Create Tenant screen is displayed, and the Status column shows In progress. To view status messages related to the creation of the tenant, click the > to the left of the status column. After about 1 minute, the Status column changes to Success, and the tenant table shows the tenant's system IP address.

View All Tenants

To view a summary of information about all tenants, in the center of the top bar, click the provider name.

View a Single Tenant

To view a summary of information about a single tenant:

1. In the center of the top bar, click the provider name.
2. In the table of tenants, click the tenant name. The summary information displays to the right of the name.
3. To hide the summary information, click the tenant name a second time.

To view the Cisco vManage dashboard for a single tenant:

1. In the center of the top bar, click **Select Tenant** to the right of the provider name.
2. Select the tenant name from the drop-down.

Edit a Tenant

1. In the left pane, click the name of the tenant.
2. In the right pane, click the Pencil icon to the right of the tenant's name.
3. In the **Edit Tenant** popup, modify the tenant's name, description, or domain name.
4. Click **Save**.

Remove a Tenant

1. In the left pane, click the name of the tenant.

2. In the right pane, click the **Trash** icon to the right of the tenant's name.
3. In the **Delete Tenant** popup, enter your Cisco vManage password and click **Save**.



CHAPTER 3

Configuring User Access and Authentication

Use the Manage Users screen to add, edit, or delete users and user groups from the vManage NMS.

Only a user logged in as the **admin** user or a user who has Manage Users write permission can add, edit, or delete users and user groups from the vManage NMS.

- [Manage Users using vManage, on page 41](#)
- [Configure User Using CLI, on page 43](#)
- [Manage a User Group, on page 44](#)
- [Creating Groups Using CLI, on page 45](#)
- [Configuring RADIUS Authentication Using CLI, on page 45](#)
- [Configure SSH Authentication, on page 46](#)
- [Configure the Authentication Order, on page 48](#)
- [Configure NAS Attributes using CLI, on page 50](#)
- [Role-Based Access with AAA, on page 51](#)
- [Configuring AAA using vManage Template, on page 60](#)
- [Configuring IEEE 802.1X and IEEE 802.11i Authentication, on page 66](#)

Manage Users using vManage

Use the Manage Users screen to add, edit, or delete users and user groups from the vManage NMS.

Only a user logged in as the **admin** user or a user who has Manage Users write permission can add, edit, or delete users and user groups from the vManage NMS.

Add a User

To perform operations on a device, you configure usernames and passwords for users who are allowed to access the device. The Cisco SD-WAN software provides one standard username, **admin**, and you can create custom usernames, as needed. We recommend that you configure strong passwords for users.

To add a user:

1. In the Users tab, click Add User.
2. In the Add User popup window, enter the full name, username, and password for the user. Note that uppercase characters are not allowed in usernames.
3. From the User Groups drop-down list, select the groups that the user will be a member of.
4. Click Add. The user is then listed in the user table.

Delete a User

If a user no longer needs access to devices, you can delete the user. When you delete a user, that user no longer has access to the device. Deleting a user does not force log out the user if the user is logged in.

To delete a user:

1. In the Users tab, select the user you wish to delete.
2. Click the More Actions icon to the right of the column and click Delete.
3. Click OK to confirm deletion of the user.

Edit User Details

Editing user details lets you update login information for a user, and add or remove a user from a user group. If you edit details for a user who is logged in, the changes take effect after the user logs out.

To edit user details:

1. In the Users tab, select the user whose details you wish to edit.
2. Click the More Actions icon to the right of the column and click Edit.
3. Edit login details, and add or remove the user from user groups.
4. Click Update.

Change User Password

You can update passwords for users as needed. We recommend that you use strong passwords.

To change a password for a user:

1. In the Users tab, select the user whose password you wish to change.
2. Click the More Actions icon to the right of the column and click Change Password.
3. Enter, and then confirm, the new password. Note that the user, if logged in, is logged out.
4. Click Done.

Configure User Using CLI

You can use the CLI to configure user credentials on each edge device. In this way, you can create additional users to give them access specific devices. The credentials that you create for a user by using the CLI can be different than the vManage credentials for the user, and you can create different credentials for a user on each device. Any user with the netadmin privilege can create a new user.

To create a user account, configure the username and password, and place the user into a group:

```
vEdge(config)# system aaa
vEdge(config)# user username password password
vEdge(config-aaa)# group group-name
```

username can be 1 to 128 characters long, and it must start with a letter. The name can contain only lowercase letters, the digits 0 through 9, hyphens (-), underscores (_), and periods (.). The name cannot contain any uppercase letters. Some usernames are reserved, so you cannot configure them. For a list of them, see the **aaa** configuration command.

password is the password for the user. Each username must have a password, and each user is allowed to change their own password. The CLI immediately encrypts the string and never displays a readable version of the password. When a user is logging in to the Cisco vEdge device, they have five chances to enter the correct password. After the fifth incorrect attempt, the user is locked out of the device, and they must wait 15 minutes before attempting to log in again.

group-name is the name of one of the standard Cisco SD-WAN groups (**basic**, **netadmin**, or **operator**) or of a group configured with the **usergroup** command (discussed below). If an **admin** user changes the permission of a user by changing their group, and if that user is currently logged in to the device, the user is logged out and must log back in again.

The factory-default password for the **admin** username is **admin**. It is strongly recommended that you modify this password the first time you configure a Cisco vEdge device.

```
vEdge(config)# system aaa admin password password
```

Configure the password as an ASCII string. The CLI immediately encrypts the string and never displays a readable version of the password. For example:

```
vEdge(config-user-admin)# show config
system
aaa
  user admin
    password $1$xULc8yYH$k71cTjvKESmeIGgImNDaC.
  !
  user eve
    password $1$8z3q4qoU$F6DMBr9vPBF0s/s145ax5.
    group basic
  !
  !
  !
```

If you are using RADIUS to perform AAA authentication, you can configure a specific RADIUS server to use to verify the password:

```
vEdge(config)# system aaa radius-servers tag
```

tag is a string that you defined with the **radius server tag** command, as described below.

Manage a User Group

Users are placed in groups, which define the specific configuration and operational commands that the users are authorized to view and modify. A single user can be in one or more groups. The Cisco SD-WAN software provides three standard user groups, and you can create custom user groups, as needed:

- **basic**—Includes users who have permission to view interface and system information.
- **netadmin**—Includes the admin user, by default, who can perform all operations on the vManage NMS. You can add other users to this group.
- **operator**—Includes users who have permission only to view information.

To add a user group:

1. In the User Groups tab, click Add User Group.
2. In the Add User Group popup window, enter the user group name and select the desired read and write permissions for each feature. Note that uppercase characters are not allowed in user group names.
3. Click OK. The user group is then listed in the left pane.

Each user group can have read or write permission for the features listed below. Write permission includes read permission.

Note: All user groups, regardless of the read or write permissions selected, can view the information displayed in the vManage Dashboard screen.

Delete a User Group

You can delete a user group when it is no longer needed. For example, you might delete a user group that you created for a specific project when that project ends.

1. In the User Groups tab, click the name of the user group you wish to delete. Note that you cannot delete any of the three standard user groups—basic, netadmin, and operator.
2. Click the Trash icon.
3. Click OK to confirm deletion of the user group.

Edit User Group Privileges

You can edit group privileges for an existing user group. This procedure lets you change configured feature read and write permissions for the user group needed.

1. In the User Groups tab, select the name of the user group whose privileges you wish to edit. Note that you cannot edit privileges for the three standard user groups—basic, netadmin, and operator.
2. Click the Edit button located directly above the privilege level table, and edit privileges as needed.
3. Click Save.

If an **admin** user changes the privileges of a user by changing their group, and if that user is currently logged in to the device, the user is logged out and must log back in again.

Creating Groups Using CLI

The Cisco SD-WAN software provides three fixed group names: **basic**, **netadmin**, and **operator**. The username **admin** is automatically placed in the **netadmin** usergroup.

If needed, you can create additional custom groups and configure privilege roles that the group members have. To create a custom group with specific authorization, configure the group name and privileges:

```
vEdge(config)# system aaa usergroup group-name task privilege
```

group-name can be 1 to 128 characters long, and it must start with a letter. The name can contain only lowercase letters, the digits 0 through 9, hyphens (-), underscores (_), and periods (.). The name cannot contain any uppercase letters. Some group names are reserved, so you cannot configure them. For a list of them, see the `aaa` configuration command.

If a remote RADIUS or TACACS+ server validates authentication but does not specify a user group, the user is placed into the user group **basic**. If a remote server validates authentication and specifies a user group (say, X) using VSA Cisco SD-WAN-Group-Name, the user is placed into that user group only. However, if that user is also configured locally and belongs to a user group (say, Y), the user is placed into both the groups (X and Y).

In the **task** option, list the privilege roles that the group members have. The role can be one or more of the following: **interface**, **policy**, **routing**, **security**, and **system**.

In the following example, the **basic** user group has full access to the **system** and **interface** portions of the configuration and operational commands, and the **operator** user group can use all operational commands but can make no modifications to the configuration:

```
vEdge# show running-config system aaa
system
aaa
  usergroup basic
    task system read write
    task interface read write
  !
  usergroup operator
    task system read
    task interface read
    task policy read
    task routing read
    task security read
  !
  user admin
    password $1$tokPB7tf$VchR2JI9Sw1/dqgkq9S.
  !
!
```

Configuring RADIUS Authentication Using CLI

The Remote Authentication Dial-In User Service (RADIUS) is a distributed client/server system that secures networks against unauthorized access. RADIUS clients run on supported Cisco devices and send authentication requests to a central RADIUS server, which contains all user authentication and network service access information.

To have a Cisco vEdge device use RADIUS servers for user authentication, configure one or up to 8 servers:

```

vEdge(config)# system radius
vEdge(config-radius)# server ip-address
vEdge(config-server)# secret-key password
vEdge(config-server)# priority number
vEdge(config-server)# auth-port port-number
vEdge(config-server)# acct-port port-number
vEdge(config-server)# source-interface interface-name
vEdge(config-server)# tag tag
vEdge(config-server)# vpn vpn-id

```

For each RADIUS server, you must configure, at a minimum, its IP address and a password, or key. You can specify the key as a clear text string up to 32 characters long or as an AES 128-bit encrypted key. The local device passes the key to the RADIUS server. The password must match the one used on the server. To configure more than one RADIUS server, include the **server** and **secret-key** commands for each server.

The remaining RADIUS configuration parameters are optional.

To set the priority of a RADIUS server, as a means of choosing or load balancing among multiple RADIUS servers, set a priority value for the server. The priority can be a value from 0 through 7. A server with a lower priority number is given priority over one with a higher number.

By default, the Cisco vEdge device uses port 1812 for authentication connections to the RADIUS server and port 1813 for accounting connections. To change these port numbers, use the **auth-port** and **acct-port** commands.

If the RADIUS server is reachable via a specific interface, configure that interface with the **source-interface** command.

You can tag RADIUS servers so that a specific server or servers can be used for AAA, IEEE 802.1X, and IEEE 802.11i authentication and accounting. Define the tag here, with a string from 4 to 16 characters long. Then associate the tag with the **radius-servers** command when you configure AAA, and when you configure interfaces for 802.1X and 802.11i.

If the RADIUS server is located in a different VPN from the Cisco vEdge device, configure the server's VPN number so that the Cisco vEdge device can locate it. If you configure multiple RADIUS servers, they must all be in the same VPN.

When a Cisco vEdge device is trying to locate a RADIUS server, it goes through the list of servers three times. To change this behavior, use the **retransmit** command, setting the number to a value from 1 to 1000:

```
vEdge(config-radius)# retransmit number
```

When waiting for a reply from the RADIUS server, a Cisco vEdge device waits 3 seconds before retransmitting its request. To change this time interval, use the **timeout** command, setting a value from 1 to 1000 seconds:

```
vEdge(config-radius)# timeout seconds
```

Configure SSH Authentication

Table 17: Feature History

Feature Name	Release Information	Description
Secure Shell Authentication Using RSA Keys	Cisco SD-WAN Release 19.2.1	This feature helps configure RSA keys by securing communication between a client and a Cisco SD-WAN server.

The Secure Shell (SSH) protocol provides secure remote access connection to network devices.

SSH supports user authentication using public and private keys. To enable SSH authentication, public keys of the users are stored in the home directory of authenticating user in the following location:

```
~<user>/.ssh/authorized_keys
```

A new key is generated on the client machine which owns the private-key. Any message encrypted using the public key of the SSH server is decrypted using the private key of the client.



Note By default, the SSH service on Cisco vEdge devices is always listening on both ports 22 and 830 on LAN. Cisco vManage uses these ports and the SSH service to perform device management. Due to this, any client machine that uses the Cisco vEdge device for internet access can attempt to SSH to the device. For each of the listening ports, we recommend that you create an ACL to block and/or allow access to Cisco vEdge devices and SSH connections for the listening ports.

Restrictions for SSH Authentication on Cisco SD-WAN

- The range of SSH RSA key size supported by Cisco vEdge devices is from 2048 to 4096. SSH RSA key size of 1024 and 8192 are not supported.
- A maximum of 10 keys are required on Cisco vEdge devices.

SSH Authentication using vManage on Cisco vEdge Devices

1. In vManage NMS, select the **Configuration ► Templates** screen.
2. In the **Feature** tab, click **Create Template**.
3. From the **Device Model** check box, select the type of device for which you are creating the template.
4. From the **Basic Information** tab, choose **AAA** template.
5. From the **Local** section, **New User** section, enter the **SSH RSA Key**. You must enter the complete public key from the `id_rsa.pub` file in the SSH RSA Key text box.

Configure SSH Authentication using CLI on Cisco vEdge Devices

When a user is created in the `/home/<user>` directory, SSH authentication configures the following parameters:

- Create the `.ssh` directory with permissions 700
- Create the `authorized_keys` files in the directory with permission 600

When the public-key is copied and pasted in the key-string, the public key is validated using the `ssh-keygen` utility. The `key-string` and `key-type` fields can be added, updated, or deleted based on your requirement. Similarly, the key-type can be changed.

When a user associated with an SSH directory gets deleted, the `.ssh` directory gets deleted.

Types of Public Keys Supported on Cisco vEdge devices:

- SSH-RSA

- SSH-DSS
- ecdsa-sha2-nistp256
- ecdsa-sha2-nistp384
- ecdsa-sha2-nistp521
- ssh-ed25519

SSH Authentication using CLI

```
vm5(config)# system aaa user ssh-user password vip group tenantadmin
vm5(config-user-ssh-user)# pubkey-chain ssh-usertag key-string
AAAAB3NzaClyc2EAAAADAQABAAQDAve2mZGPKveIgzHw6cjqsfTyIUogfPlkgsBJDuJfMhU1hrWZlh03sLvkci29Og2NNSJYM3OCy0TA7pFWvpDDXQw/gD4/
Eb2TH09CBNEChdV0zA6K2fMwOZfW2PvNRBlOzVlijjQaitd5Dqe7Ar5HGTAfLWzVmkU9HLQUDZSfeDt8cl/ftgn8sKQOxuTccIpwFnyZkth978Bqm029v8/O5R
BdQOVt3VBr9NNeC4egutS0yBNZeWBPfwecd4/aot38plF6jOo1DvUjn60CUUOu9TQIaSFg/dFUB0twEBOIUfMBeimRexIT+cI3z8vM1D9toqFRDAI8EUJegjU7BP
vm5(config-pubkey-chain-ssh-usertag)# commit
Commit complete.
```

Configure the Authentication Order

The authentication order dictates the order in which authentication methods are tried when verifying user access to a Cisco vEdge device through an SSH session or a console port. The default authentication order is **local**, then **radius**, and then **tacacs**. With the default authentication order, the authentication process occurs in the following sequence:

- The authentication process first checks whether a username and matching password are present in the running configuration on the local device.
- If local authentication fails, and if you have not configured authentication fallback (with the **auth-fallback** command), the authentication process stops. However, if you have configured authentication fallback, the authentication process next checks the RADIUS server. For this method to work, you must configure one or more RADIUS servers with the **system radius server** command. If a RADIUS server is reachable, the user is authenticated or denied access based on that server's RADIUS database. If a RADIUS server is unreachable and if you have configured multiple RADIUS servers, the authentication process checks each server sequentially, stopping when it is able to reach one of them. The user is then authenticated or denied access based on that server's RADIUS database.
- If the RADIUS server is unreachable (or all the servers are unreachable), the authentication process checks the TACACS+ server. For this method to work, you must configure one or more TACACS+ servers with the **system tacacs server** command. If a TACACS+ server is reachable, the user is authenticated or denied access based on that server's TACACS+ database. If a TACACS+ server is unreachable and if you have configured multiple TACACS+ servers, the authentication process checks each server sequentially, stopping when it is able to reach one of them. The user is then authenticated or denied access based on that server's TACACS+ database.
- If the TACACS+ server is unreachable (or all TACACS+ servers are unreachable), user access to the local Cisco vEdge device is denied.

To modify the default order, use the **auth-order** command:

```
vEdge(config-system-aaa)# auth-order (local | radius | tacacs)
```

Specify one, two, or three authentication methods in the preferred order, starting with the one to be tried first. If you configure only one authentication method, it must be **local**.

To have the "admin" user use the authentication order configured in the **auth-order** command, use the following command:

```
vEdge(config-system-aaa)# admin-auth-order
```

If you do not include this command, the "admin" user is always authenticated locally.

You can configure authentication to fall back to a secondary or tertiary authentication mechanism when the higher-priority authentication method fails to authenticate a user, either because the user has entered invalid credentials or because the authentication server is unreachable (or all the servers are unreachable):

```
vEdge(config-system-aaa)# auth-fallback
```

Fallback to a secondary or tertiary authentication mechanism happens when the higher-priority authentication server fails to authenticate a user, either because the credentials provided by the user are invalid or because the server is unreachable.

The following examples illustrate the default authentication behavior and the behavior when authentication fallback is enabled:

- If the authentication order is configured as **radius local**:
 - With the default authentication, local authentication is used only when all RADIUS servers are unreachable. If an authentication attempt via a RADIUS server fails, the user is not allowed to log in even if they have provided the correct credentials for local authentication.
 - With authentication fallback enabled, local authentication is used when all RADIUS servers are unreachable or when a RADIUS server denies access to a user.
- If the authentication order is configured as **local radius**:
 - With the default authentication, RADIUS authentication is tried when a username and matching password are not present in the running configuration on the local device.
 - With authentication fallback enabled, RADIUS authentication is tried when a username and matching password are not present in the running configuration on the local device. In this case, the behavior of two authentication methods is identical.
- If the authentication order is configured as **radius tacacs local**:
 - With the default authentication, TACACS+ is tried only when all RADIUS servers are unreachable, and local authentication is tried only when all TACACS+ servers are unreachable. If an authentication attempt via a RADIUS server fails, the user is not allowed to log in even if they have provided the correct credentials for the TACACS+ server. Similarly, if a TACACS+ server denies access, the user cannot log via local authentication.
 - With authentication fallback enabled, TACACS+ authentication is used when all RADIUS servers are unreachable or when a RADIUS server denies access a user. Local authentication is used next, when all TACACS+ servers are unreachable or when a TACACS+ server denies access to a user.

If a remote server validates authentication but does not specify a user group, the user is placed into the user group **basic**.

If a remote server validates authentication and specifies a user group (say, X), the user is placed into that user group only. However, if that user is also configured locally and belongs to a user group (say, Y), the user is placed into both the groups (X and Y).

If a remote server validates authentication and that user is not configured locally, the user is logged in to the vshell as the user **basic**, with a home directory of /home/basic.

If a remote server validates authentication and that user is configured locally, the user is logged in to the vshell under their local username (say, eve) with a home direction of /home/username (so, /home/eve).

Configure NAS Attributes using CLI

For RADIUS and TACACS+, you can configure Network Access Server (NAS) attributes for user authentication and authorization. To do this, you create a vendor-specific attributes (VSA) file, also called a RADIUS dictionary or a TACACS+ dictionary, on the RADIUS or TACACS+ server that contains the desired permit and deny commands for each user. The Cisco vEdge device retrieves this information from the RADIUS or TACACS+ server.

The VSA file must be named `dictionary.viptela`, and it must contain text in the following format:

```
localhost$ more dictionary.viptela
# -*- text -*-
#
# dictionary.viptela
#
#
# Version:      $Id$
#
VENDOR          Viptela                      41916
BEGIN-VENDOR    Viptela
ATTRIBUTE       Viptela-Group-Name          1      string
```

The Cisco SD-WAN software has three predefined user groups, as described above: **basic**, **netadmin**, and **operator**. These groups have the following permissions:

```
vEdge# show aaa usergroup
GROUP      USERS  TASK      PERMISSION
-----
basic      -      system    read
           -      interface read
netadmin   admin  system    read write
           -      interface read write
           -      policy     read write
           -      routing    read write
           -      security   read write
operator   -      system    read
           -      interface read
           -      policy     read
           -      routing    read
           -      security   read
```

To create new user groups, use this command:

```
vEdge(config)# system aaa usergroup
group-name task privilege
```

Here is a sample user configuration on a RADIUS server, which for FreeRADIUS would be in the file "users":


```

user1 Cleartext-password := "user123"
      Service-Type = NAS-Prompt-User,
      Viptela-Group-Name = operator,

user1 Cleartext-password := "user123"           Service-Type = NAS-Prompt-User,
      Viptela-Group-Name = operator,

```

Then in the dictionary on the RADIUS server, add a pointer to the VSA file:

```
$INCLUDE /usr/share/freeradius/dictionary.viptela
```

For TACACS+, here is a sample configuration, which would be in the file `tac_plus.conf`:

```

group = test_group {
    default service = permit
    service = ppp protocol = ip {
        Viptela-Group-Name = operator
    }
}
user = user1 {
    pap = cleartext "user123"
    member = test_group
}

```

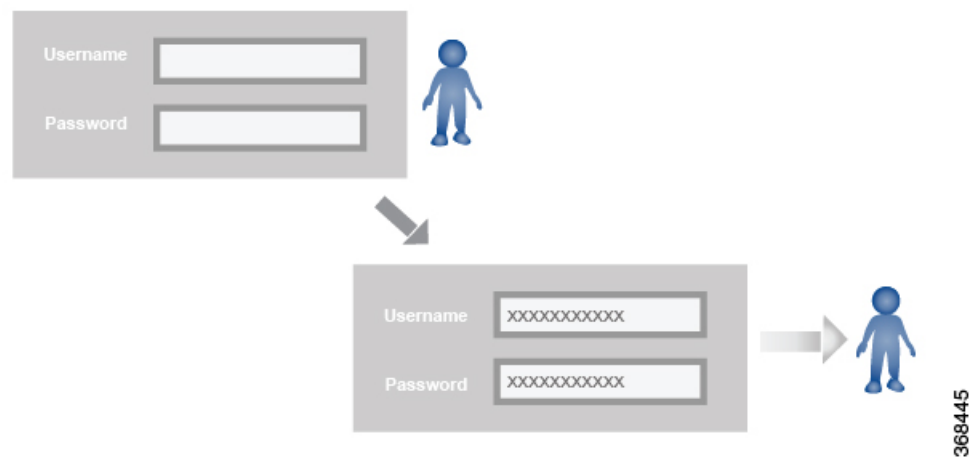
Role-Based Access with AAA

The Cisco SD-WAN AAA software implements role-based access to control the authorization permissions for users on Cisco vEdge devices. Role-based access consists of three components:

- Users are those who are allowed to log in to a Cisco vEdge device.
- User groups are collections of users.
- Privileges are associated with each group. They define the commands that the group's users are authorized to issue.


Users and User Groups

All users who are permitted to perform operations on a Cisco vEdge device must have a login account. For the login account, you configure a username and a password on the device itself. These allow the user to log in to that device. A username and password must be configured on each device that a user is allowed to access.




The Cisco SD-WAN software provides one standard username, **admin**, which is a user who has full administrative privileges, similar to a UNIX superuser. By default, the **admin** username password is **admin**. You cannot delete or modify this username, but you can and should change the default password.

User groups pool together users who have common roles, or privileges, on the Cisco vEdge device. As part of configuring the login account information, you specify which user group or groups that user is a member of. You do not need to specify a group for the **admin** user, because this user is automatically in the user group **netadmin** and is permitted to perform all operations on the Cisco vEdge device.



The user group itself is where you configure the privileges associated with that group. These privileges correspond to the specific commands that the user is permitted to execute, effectively defining the role-based access to the Cisco SD-WAN software elements.



The Cisco SD-WAN software provides three standard user groups. The two groups **basic** and **operator** are configurable. While you can use these two groups for any users and privilege levels, the **basic** group is designed to include users who have permission to both view and modify information on the device, while the **operator** group is designed to include users who have permission only to view information. The third group is **net admin**, which is non-configurable. By default, it includes the **admin** user. You can add other users to this group. Users in this group are permitted to perform all operations on the device.



Note Only admin users can view running and local configuration. Users associated with predefined operator user group do not have access to the running and local configurations. The predefined user group operator has only read access for the template configuration. If you need only a subset of admin user privileges, then you need to create a new user group with the selected features from the features list with both read and write access and associate the group with the custom user.

Privileges for Role-Based Access

Role-based access privileges are arranged into five categories, which are called *tasks*:

- Interface—Privileges for controlling the interfaces on the Cisco vEdge device.
- Policy—Privileges for controlling control plane policy, OMP, and data plane policy.
- Routing—Privileges for controlling the routing protocols, including BFD, BGP, OMP, and OSPF.
- Security—Privileges for controlling the security of the device, including installing software and certificates. Only users belonging to the **netadmin** group can install software on the system.
- System—General systemwide privileges.

The tables in the following sections detail the AAA authorization rules for users and user groups. These authorization rules apply to commands issued from the CLI and to those issued from Netconf.

User Authorization Rules for Operational Commands

The user authorization rules for operational commands are based simply on the username. Any user who is allowed to log in to the Cisco vEdge device can execute most operational commands. However, only the **admin** user can issue commands that affect the fundamental operation of the device, such as installing and upgrading the software and shutting down the device.

Note that any user can issue the **config** command to enter configuration mode, and once in configuration mode, they are allowed to issue any general configuration command. Also, any user is allowed to configure their password by issuing the **system aaa user self password password** command and then committing that configuration change. For the actual commands that configure device operation, authorization is defined according to user group membership. See User Group Authorization Rules for Configuration Commands.

The following table lists the AAA authorization rules for general CLI commands. All the commands are operational commands except as noted. Also, some commands available to the "admin" user are available only if that user is in the "netadmin" user group.

CLI Command	Any User	Admin User
clear history	X	X
commit confirm	X	X
complete-on-space	X	X
config	X	X
exit	X	X
file	X	X
help	X	X
[no] history	X	X
idle-timeout	X	X
job	X	X
logout	—	X (users in netadmin group only)

CLI Command	Any User	Admin User
monitor	X	X
nslookup	X	X
paginate	X	X
ping	X	X
poweroff	—	X(users in netadmin group only)
prompt1	X	X
prompt2	X	X
quit	X	X
reboot	—	X (users in netadmin group only)
request aaa request admin-tech request firmware request interface-reset request nms request reset request software	—	X (users in netadmin group only)
request execute request download request upload	X	X
request (everything else)	—	X
rollback (configuration mode command)	—	X (users in netadmin group only)
screen-length	X	X
screen-width	X	X
show cli	X	X
show configuration commit list	X	X
show history	X	X
show jobs	X	X
show parser dump	X	X
show running-config	X	X
show users	X	X
system aaa user <i>self</i> password <i>password</i> (configuration mode command) (Note: A user cannot delete themselves)		

CLI Command	Any User	Admin User
tcpdump	X	X
timestamp	X	X
tools ip-route	X	X
tools netstat	X	X
tools nping	X	X
traceroute	X	X
vshell	X	X (users in netadmin group only)

User Group Authorization Rules for Operational Commands

The following table lists the user group authorization roles for operational commands.

Operational Command	Interface	Policy	Routing	Security	System
clear app		X			
clear app-route		X			
clear arp	X				
clear bfd			X		X
clear bgp			X		X
clear bridge	X				
clear cellular	X				
clear control				X	
clear crash					X
clear dhcp					X
clear dns					X
clear igmp			X		
clear installed-certificates				X	
clear interface	X				
clear ip			X		
clear notification					X
clear omp			X		

Operational Command	Interface	Policy	Routing	Security	System
clear orchestrator				X	
clear ospf			X		
clear pim			X		
clear policy		X			
clear pppoe	X				
clear system					X
clear tunnel				X	
clear wlan	X				
clear ztp				X	X
clock					X
debug bgp			X		
debug cellular	X				
debug cflowd		X			
debug chmgr					X
debug config-mgr					X
debug dhcp-client					X
debug dhcp-helper					X
debug dhcp-server					X
debug fpm		X			
debug ftm					X
debug igmp			X		
debug netconf					X
debug omp			X		
debug ospf			X		
debug pim			X		
debug resolver			X		
debug snmp					X
debug sysmgr					X

Operational Command	Interface	Policy	Routing	Security	System
debug transport					X
debug ttm					X
debug vdaemon				X	X
debug vrrp				X	
debug wlan	X				
request certificate				X	
request control-tunnel				X	
request controller				X	
request controller-upload				X	
request csr				X	
request device				X	
request device-upload				X	
request on-vbond-controller				X	
request port-hop				X	
request root-cert-chain				X	
request security				X	
request vedge				X	
request vedge-upload				X	
request vsmart-upload				X	
show aaa					X
show app		X			
show app-route		X			
show arp	X				
show bfd			X		X
show bgp			X		
show boot-partition					X
show bridge	X				

Operational Command	Interface	Policy	Routing	Security	System
show cellular	X				
show certificate				X	
show clock					X
show control				X	X
show crash					X
show debugs—same as debug commands					
show dhcp					X
show external-nat				X	X
show hardware					X
show igmp			X		
show interface	X				
show ip			X		X
show ipsec				X	
show licenses					X
show logging					X
show multicast			X		
show nms-server					X
show notification					X
show ntp					X
show omp		X	X		X
show orchestrator				X	
show ospf			X		
show pim			X		
show policer		X			
show policy		X			
show ppp	X				
show pppoe	X				

Operational Command	Interface	Policy	Routing	Security	System
show reboot					X
show security-info				X	
show software					X
show system					X
show transport					X
show tunnel				X	
show uptime					X
show users					X
show version					X
show vrrp	X				
show wlan	X				
show ztp				X	

User Group Authorization Rules for Configuration Commands

The following table lists the user group authorization rules for configuration commands.

Configuration Command	Interface	Policy	Routing	Security	System
apply-policy		X			
banner					X
bfd			X		X
bridge	X				
omp		X	X		X
policy		X			
security				X	X
snmp					X
system					X
vpn interface	X				
vpn ip			X		
vpn router			X		

Configuration Command	Interface	Policy	Routing	Security	System
vpn service			X		
vpn (everything else, including creating, deleting, and naming)					X
wlan	X				

Configuring AAA using vManage Template

Configuring AAA by using the vManage template lets you make configuration setting in vManage and then push the configuration to selected devices of the same type. This procedure is a convenient way to configure several of the same type of devices at one time.

Use the AAA template for Cisco vBond Orchestrators, vManage NMSs, Cisco vSmart Controllers, and Cisco vEdge device s.

Cisco vEdge device s support configuration of authentication, authorization, and accounting (AAA) in combination with RADIUS and TACACS+.



Note You must configure a local user with a secret key via the template if you are using PPP or using MLPPP with CHAP.

Navigate to the Template Screen and Name the Template

1. In vManage NMS, select the **Configuration** ► **Templates** screen.
2. In the Device tab, click Create Template.
3. From the Create Template drop-down, select From Feature Template.
4. From the **Device Model** drop-down, select the type of device for which you are creating the template.
5. Select the **Basic Information** tab.
6. To create a custom template for AAA, select the Factory_Default_AAA_Template and click Create Template. The AAA template form is displayed. The top of the form contains fields for naming the template, and the bottom contains fields for defining AAA parameters.
7. In the Template Name field, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
8. In the Template Description field, enter a description of the template. The description can be up to 2048 characters and can contain only alphanumeric characters.

When you first open a feature template, for each parameter that has a default value, the scope is set to Default (indicated by a check mark), and the default setting or value is shown. To change the default or to enter a value, click the scope drop-down to the left of the parameter field and select one of the following:

Table 18:

Parameter Scope	Scope Description
Device Specific (indicated by a host icon)	<p>Use a device-specific value for the parameter. For device-specific parameters, you cannot enter a value in the feature template. You enter the value when you attach a Cisco vEdge device to a device template .</p> <p>When you click Device Specific, the Enter Key box opens. This box displays a key, which is a unique string that identifies the parameter in a CSV file that you create. This file is an Excel spreadsheet that contains one column for each key. The header row contains the key names (one key per column), and each row after that corresponds to a device and defines the values of the keys for that device. You upload the CSV file when you attach a Cisco vEdge device to a device template. For more information, see Create a Template Variables Spreadsheet .</p> <p>To change the default key, type a new string and move the cursor out of the Enter Key box.</p> <p>Examples of device-specific parameters are system IP address, hostname, GPS location, and site ID.</p>
Global (indicated by a globe icon)	<p>Enter a value for the parameter, and apply that value to all devices.</p> <p>Examples of parameters that you might apply globally to a group of devices are DNS server, syslog server, and interface MTUs.</p>

Configure Authentication Order and Fallback

You can configure the authentication order and authentication fallback for device. The authentication order specifies the order in which the system attempts to authenticate user, and provides a way to proceed with authentication if the current authentication method is unavailable. Fallback provides a mechanism for authentication if the user cannot be authenticated or if a RADIUS or TACACS+ server is unreachable.

To configure AAA authentication order and authentication fallback on a Cisco vEdge device , select the Authentication tab and configure the following parameters:

Table 19:

Parameter Name	Description
Authentication Order	<p>The default order is local, then radius, and then tacacs.</p> <p>To change the default order of authentication methods that the software tries when verifying user access to a Cisco vEdge device :</p> <ol style="list-style-type: none"> 1. Click the dropdown arrow to display the list of authentication methods. 2. In the list, click the up arrows to change the order of the authentication methods and click the boxes to select or deselect a method. <p>If you select only one authentication method, it must be local.</p>

Parameter Name	Description
Authentication Fallback	Click On to configure authentication to fall back from RADIUS or TACACS+ to the next priority authentication method if the user cannot be authenticated or if the RADIUS or TACACS+ servers are unreachable. With the default configuration (Off), authentication falls back only if the RADIUS or TACACS+ servers are unreachable.
Admin Authentication Order	Have the "admin" user use the authentication order configured in the Authentication Order parameter. If you do not configure the admin authentication order, the "admin" user is always authenticated locally.
Disable Netconf Logs	Click On to disable the logging of Netconf events. By default, these events are logged to the auth.info and messages log files.
Disable Audit Logs	Click On to disable the logging of AAA events. By default, these events are logged to the auth.info and messages log files.
RADIUS Server List	List the tags for one or two RADIUS servers. Separate the tags with commas. You set the tag under the RADIUS tab.

CLI equivalent:

```

system
aaa
admin-auth-order auth-fallback auth-order (local | radius | tacacs)
logs
  [no] audit-disable
  [no] netconf-disable
radius-servers tag

```

Configure Local Access for Users and User Groups

You can configure local access to a device for users and user groups. Local access provides access to a device if RADIUS or TACACS+ authentication fails.

To configure local access for individual users, select the Local tab. To add a new user, select the User tab, click Add New User, and configure the following parameters:

Table 20:

Parameter Name	Description
Name	<p>Enter a name for the user. It can be 1 to 128 characters long, and it must start with a letter. The name can contain only lowercase letters, the digits 0 through 9, hyphens (-), underscores (_), and periods (.). The name cannot contain any uppercase letters.</p> <p>The following usernames are reserved, so you cannot configure them: backup, basic, bin, daemon, games, gnats, irc, list, lp, mail, man, news, nobody, proxy, quagga, root, sshd, sync, sys, uucp, and www-data. Also, names that start with viptela-reserved are reserved.</p>

Parameter Name	Description
Password	<p>Enter a password for the user. The password is an MD5 digest string, and it can contain any characters, including tabs, carriage returns, and linefeeds. For more information, see Section 9.4 in RFC 7950, <i>The YANG 1.1 Data Modeling Language</i>.</p> <p>Each username must have a password. Users are allowed to change their own passwords.</p> <p>The default password for the admin user is admin. We strongly recommended that you change this password.</p>
Description	Enter a description for the user.
User Groups	Select from the list of configured groups. You must assign the user to at least one group. The admin user is automatically placed in the netadmin group and is the only member of this group.

Click Add to add the new user. Click Add New User again to add additional users.

To configure local access for user groups, you first place the user into either the basic or operator group. The admin is automatically placed in the netadmin group. Then you configure user groups. To make this configuration, select the Local tab, select the User Group tab, click Add New User Group, and configure the following parameters:

Table 21:

Parameter Name	Description
Name	<p>Name of an authentication group. It can be 1 to 128 characters long, and it must start with a letter. The name can contain only lowercase letters, the digits 0 through 9, hyphens (-), underscores (_), and periods (.). The name cannot contain any uppercase letters. The Cisco SD-WAN software provides three standard user groups, basic, netadmin, and operator. The user admin is automatically placed in the group netadmin and is the only user in this group. All users learned from a RADIUS or TACACS+ server are placed in the group basic. All users in the basic group have the same permissions to perform tasks, as do all users in the operator group. The following groups names are reserved, so you cannot configure them: adm, audio, backup, bin, cdrom, dialout, dip, disk, fax, floppy, games, gnats, input, irc, kmem, list, lp, mail, man, news, nogroup, plugdev, proxy, quagga, quaggavty, root, sasl, shadow, src, sshd, staff, sudo, sync, sys, tape, tty, uucp, users, utmp, video, voice, and www-data. Also, group names that start with the string viptela-reserved are reserved.</p>
Feature	<p>The feature table lists the roles for the user group. These roles are Interface, Policy, Routing, Security, and System. Each role allows the user group to read or write specific portions of the device's configuration and to execute specific types of operational commands. Click the appropriate boxes for Read, Write, and None to assign privileges to the group for each role.</p>

Click Add to add the new user group.

To add another user group, click Add New User Group again.

To delete a user group, click the trash icon at the right side of the entry. You cannot delete the three standard user groups, basic, netadmin, and operator.

CLI equivalent:

```

system
aaa
  user username
  group group-name
  password password usergroup group-name
  task (interface | policy | routing | security | system) (read | write)

```

Configure RADIUS Authentication

Configure RADIUS authentication if you are using RADIUS in your deployment.

To configure RADIUS authentication, select the RADIUS tab and configure the following parameters:

Table 22:

Parameter Name	Description
Retransmit Count	Specify how many times to search through the list of RADIUS servers while attempting to locate a server. <i>Range:</i> 1 through 1000 <i>Default:</i> 3
Timeout	Specify how long to wait to receive a reply from the RADIUS server before retransmitting a request. <i>Range:</i> 1 through 1000 <i>Default:</i> 5 seconds

To configure a connection to a RADIUS server, select the RADIUS tab, click Add New Radius Server, and configure the following parameters:

Table 23:

Parameter Name	Description
Address	Enter the IP address of the RADIUS server host.
Tag	Enter a text string to identify the RADIUS server. The tag can be 4 to 16 characters long. The tag allows you to configure authentication for AAA, IEEE 802.1X, and IEEE 802.11i to use a specific RADIUS server or servers. For Cisco vEdge devices running Cisco SD-WAN software, this field is ignored.
Authentication Port	Enter the UDP destination port to use for authentication requests to the RADIUS server. If the server is not used for authentication, configure the port number to be 0. <i>Default:</i> Port 1812
Accounting Port	Enter the UDP port to use to send 802.1X and 802.11i accounting information to the RADIUS server. <i>Range:</i> 0 through 65535 <i>Default:</i> 1813
Key (Deprecated)	This field is deprecated. Use the Secret Key field instead.
Secret Key	Enter the key the Cisco vEdge device passes to the RADIUS server for authentication and encryption. You can type the key as a text string from 1 to 32 characters long, and it is immediately encrypted, or you can type an AES 128-bit encrypted key. The key must match the AES encryption key used on the RADIUS server.

Parameter Name	Description
Source Interface	Enter the name of the interface on the local device to use to reach the RADIUS server.
VPN ID	Enter the number of the VPN in which the RADIUS server is located or through which the server can be reached. If you configure multiple RADIUS servers, they must all be in the same VPN.
Priority	Enter the priority of a RADIUS server. A server with a lower number is given priority. <i>Range: 0 through 7Default: 0</i>

Click Add to add the new RADIUS server.

To add another RADIUS server, click Add New RADIUS Server again.

To remove a server, click the trash icon on the right side of the line.

CLI equivalent:

```

system radius
  retransmit number
  server ip-address
    acct-port port-number
    auth-port port-number
    priority number
    secret-key key
    source-interface interface-name
    tag tag
    vpn vpn-id
  timeout seconds

```

Configure TACACS+ Authentication

Configure TACACS+ authentication if you are using TACACS+ in your deployment.

To configure the device to use TACACS+ authentication, select the TACACS tab and configure the following parameters:

Table 24:

Parameter Name	Description
Timeout	Enter how long to wait to receive a reply from the TACACS+ server before retransmitting a request. <i>Range: 1 through 1000Default: 5 seconds</i>
Authentication	Set the type of authentication to use for the server password. The default authentication type is PAP. You can change it to ASCII.

To configure a connection to a TACACS+ server, select the TACACS tab, click Add New TACSCS Server, and configure the following parameters:

Table 25:

Parameter Name	Description
Address	Enter the IP address of the TACACS+ server host.
Authentication Port	Enter the UDP destination port to use for authentication requests to the TACACS+ server. If the server is not used for authentication, configure the port number to be 0. <i>Default:</i> Port 49
Key (Deprecated)	This field is deprecated. Use the Secret Key field instead.
Secret Key	Enter the key the Cisco vEdge device passes to the TACACS+ server for authentication and encryption. You can type the key as a text string from 1 to 32 characters long, and it is immediately encrypted, or you can type an AES 128-bit encrypted key. The key must match the AES encryption key used on the TACACS+ server.
Source Interface	Enter the name of the interface on the local device to use to reach the TACACS+ server.
VPN ID	VPN in which the TACACS+ server is located or through which the server can be reached. If you configure multiple TACACS+ servers, they must all be in the same VPN.
Priority	Set the priority of a TACACS+ server. A server with lower priority number is given priority over one with a higher number. <i>Range:</i> 0 through 7 <i>Default:</i> 0

Click Add to add the new TACACS server.

To add another TACACS server, click Add New TACACS Server again.

To remove a server, click the trash icon on the right side of the line.

CLI equivalent:

```

system tacacs
  authentication password-authentication
  server ip-address
    auth-port port-number
    priority number
    key key
    source-interface interface-name
  vpn vpn-id
  timeout seconds

```

Configuring IEEE 802.1X and IEEE 802.11i Authentication

IEEE 802.1X is a port-based network access control (PNAC) protocol that prevents unauthorized network devices from gaining access to wired networks (WANs), by providing authentication for devices that want to connect to a WAN.

IEEE 802.11i prevents unauthorized network devices from gaining access to wireless networks (WLANs). 802.11i implements WiFi Protected Access II (WPA2) to provide authentication for devices that want to connect to a WLAN on a Cisco vEdge 100wm device.

A RADIUS authentication server must authenticate each client connected to a port before that client can access any services offered by network.

This section describes how to configure RADIUS servers to use for 802.1X and 802.11i authentication. It describes how to enable 802.1X on Cisco vEdge device interfaces to have the router act as an 802.1X authenticator, responsible for authorizing or denying access to network devices on a WAN.

It also describes how to enable 802.11i on Cisco vEdge 100wm device routers to control access to WLANs.

It describes how to enable IEEE 802.1X and AAA on a port, and how to enable IEEE 802.1X RADIUS accounting.

Configure RADIUS Authentication Servers

Authentication services for IEEE 802.1X and IEEE 802.11i are provided by RADIUS authentication servers. You configure the RADIUS servers to use for 802.1X and 802.11i authentication on a system-wide basis:

```
vEdge(config)# system radius  
vEdge(config-radius)# server ip-address
```

Specify the IP address of the RADIUS server. You can configure one or two RADIUS servers to perform 802.1X and 802.11i authentication. (Note that for AAA authentication, you can configure up to eight RADIUS servers.)

For each RADIUS server, you can configure a number of optional parameters.

You can configure the VPN through which the RADIUS server is reachable and the router interface to use to reach the server:

```
vEdge(config-server)# vpn vpn-id  
vEdge(config-server)# source-interface interface-name
```

If you configure two RADIUS servers, they must both be in the same VPN, and they must both be reachable using the same source interface.

You must configure a tag to identify the RADIUS server:

```
vEdge(config-server)# tag tag
```

The tag can be from 4 through 16 characters. You use this tag when configuring the RADIUS servers to use with IEEE 802.1X authentication and with IEEE 802.11i WPA enterprise authentication.

For authentication between the router and the RADIUS server, you can authenticate and encrypt packets sent between the Cisco vEdge device and the RADIUS server, and you can configure a destination port for authentication requests. To authenticate and encrypt packets, configure a key:

```
vEdge(config-server)# secret-key password
```

Enter the password as clear text, which is immediately encrypted, or as an AES 128-bit encrypted key. The key must match the AES encryption key used on the RADIUS server.

By default, UDP port 1812 is used as the destination port on the RADIUS server to use for authentication requests. You can change the port number to a number from 1 through 65535. To disable authentication, set the port number to 0.

```
vEdge(config-server)# auth-port number
```

You can set the priority of a RADIUS server, to choose which one to use first when performing 802.1X authentication:

```
vEdge(config-server)# priority number
```

The priority can be a value from 0 through 7. The server with the lower priority number is given priority. If you do not include this command in the RADIUS server configuration, the priority is determined by the order in which you enter the IP addresses in the **system radius server** command.

By default, accounting is enabled for 802.1X and 802.11i interfaces. Accounting information is sent to UDP port 1813 on the RADIUS server. To change this port:

```
vEdge(config-server)# acct-port number
```

The port number can be from 1 through 65535.

Configure IEEE 802.1X Port Security

To enable basic 802.1X port security on an interface, configure it and at least one RADIUS server to use for 802.1X authentication. The 802.1X interface must be in VPN 0.

```
vEdge(config)# vpn 0
interface interface-name
vEdge(config-interface)# dot1x
vEdge(config-dot1x)# radius-servers tag
```

For 802.1X authentication to work, you must also configure the same interface under an untagged bridge:

```
vEdge(config)# bridge number
vEdge(config)# interface interface-name
```

The interface name in the **vpn 0 interface** and **bridge interface** commands must be the same. Do not configure a VLAN ID for this bridge so that it remains untagged.

You can enable 802.1X on a maximum of four wired physical interfaces. The interface cannot also be configured as a tunnel interface.

Configure the tags associated with one or two RADIUS servers to use for 802.1X client authentication and accounting. (You configure the tags with the **system radius server tag** command.) If you specify tags for two RADIUS servers, they must both be reachable in the same VPN. If you do not configure a priority value when you configure the RADIUS server with the **system radius server priority** command, the order in which you list the IP addresses is the order in which the RADIUS servers are tried.

Enable RADIUS Accounting

By default, the Cisco vEdge device never sends interim accounting updates to the 802.1X RADIUS accounting server. Accounting updates are sent only when the 802.1X session ends.

To enable the sending of interim accounting updates, configure the interval at which to send the updates:

```
vEdge(config-dot1x)# accounting-interval seconds
```

The time can be from 0 through 7200 seconds.

Enable MAC Authentication Bypass

IEEE 802.1X authentication is accomplished through an exchange of Extensible Authentication Protocol (EAP) packets. After 802.1X-compliant clients respond to the EAP packets, they can be authenticated and granted access to the network. Enabling MAC authentication bypass (MAB) provides a mechanism to allow non-802.1X-compliant clients to be authenticated and granted access to the network.

The Cisco vEdge device determines that a device is non-802.1X-compliant clients when the 802.1X authentication process times out while waiting for an EAPOL response from the client.

To enable MAC authentication bypass for an 802.1X interface on the Cisco vEdge device :

```
vEdge(config)# vpn 0 interface interface-name dot1x
vEdge(config-dot1x)# mac-authentication-bypass
```

With this configuration, the Cisco vEdge device authenticates non-802.1X-compliant clients using the configured RADIUS servers. The RADIUS server must be configured with the MAC addresses of non-802.1X-compliant clients that are allowed to access the network.

To enable MAB on the RADIUS server:

```
vEdge(config-dot1x) # mac-authentication-bypass server
```

To allow authentication to be performed for one or more non-802.1X-compliant clients before performing an authentication check with the RADIUS server, list their MAC addresses in the following command:

```
vEdge(config-dot1x) # mac-authentication-bypass allow mac-addresses
```

You can configure up to eight MAC addresses for MAC authentication bypass. For these devices, the Cisco vEdge device grants immediate network access based on their MAC addresses, and then sends a request to the RADIUS server to authenticate the devices.

Configure VLANs for Authenticated and Unauthenticated Clients

For clients that cannot be authenticated but that you want to provide limited network services to, you create VLANs to handle network access for these clients. You also create VLANs to handle authenticated clients.

You can create the following kinds of VLAN:

- Guest VLAN—Provide limited services to non-802.1X-compliant clients.
- Authentication Reject VLAN—Provide limited services to 802.1X-compliant clients that failed RADIUS authentication. An authentication-reject VLAN is similar to a restricted VLAN.
- Authentication Fail VLAN—Provide network access when RADIUS authentication or the RADIUS server fails. An authentication-fail VLAN is similar to a critical VLAN.
- Default VLAN—Provide network access to 802.1X-compliant clients that are successfully authenticated by the RADIUS server. If you do not configure a default VLAN on the Cisco vEdge device, successfully authenticated clients are placed into VLAN 0, which is the VLAN associated with an untagged bridge.

To configure the VLANs for authenticated and unauthenticated clients, first create the VLAN in a bridging domain, and then create the 802.1X VLANs for the unauthenticated clients by associating the bridging domain VLAN with an 802.1X VLAN.

To create the VLAN, configure a bridging domain to contain the VLAN:

```
vEdge(config) # bridge bridge-id
vEdge(config-bridge) # name text
vEdge(config-bridge) # vlan vlan-id
vEdge(config-bridge) # interface interface-name
vEdge(config-interface) # no shutdown
```

The bridging domain identifier is a number from 1 through 63. A best practice is to have the bridge domain ID be the same as the VLAN number.

The name is optional, but it is recommended that you configure a name that identifies the 802.1X VLAN type, such as Guest-VLAN and Default-VLAN.

The VLAN number can be from 1 through 4095. This is the number that you associate with an 802.1X VLAN.

The interface name is the interface that is running 802.1X.

Then configure the 802.1X VLANs to handle unauthenticated clients.

A guest VLAN provides limited services to non-802.1X-compliant clients, and it can be used to allow clients to download 802.1X client software. An interface running 802.1X assigns clients to a guest VLAN when the interface does not receive a response to EAP request/identity packets that it has sent to the client, or when the client does not send EAPOL packets and MAC authentication bypass is not enabled. To configure a guest VLAN:

```
vEdge(config)# vpn 0 interface interface-name interface dot1x
vEdge(config-dot1x)# guest-vlan vlan-id
```

The VLAN number must match one of the VLANs you configured in a bridging domain. A best practice is to have the VLAN number be the same as the bridge domain ID.

An authentication-reject VLAN provides limited services to 802.1X-compliant clients that have failed RADIUS authentication. To configure an authentication-reject VLAN:

```
vEdge(config-dot1x)# auth-reject-vlan vlan-id
```

The VLAN number must match one of the VLANs you configure in a bridging domain. A best practice is to have the VLAN number be the same as the bridge domain ID.

When the RADIUS authentication server is not available, 802.1X-compliant clients attempting to authenticate are placed in an authentication-fail VLAN if it is configured. If this VLAN is not configured, the authentication request is eventually dropped. To configure the authentication-fail VLAN:

```
vEdge(config-dot1x)# auth-fail-vlan vlan-id
```

The VLAN number must match one of the VLANs you configure in a bridging domain. A best practice is to have the VLAN number be the same as the bridge domain ID.

The following configuration snippet illustrates the interrelationship between the 802.1X configuration and the bridging domain configuration. This snippet shows that the bridging domain numbers match the VLAN numbers, which is a recommended best practice. Also, the bridging domain name identifies the type of 802.1X VLAN.

```
system
...
radius
server 10.1.15.150
  tag freerad1
  source-interface ge0/0
  secret-key $4$L3rwZmsIic8zj4BgLEFXKw==
  priority 1
exit
server 10.20.24.150
  auth-port 2000
  acct-port 2001
  tag freerad2
  source-interface ge0/4
  secret-key $4$L3rwZmsIic8zj4BgLEFXKw==
  priority 2
exit
!
!
bridge 1
name Untagged_bridge
interface ge0/5
  no native-vlan
  no shutdown
!
!
bridge 10
name Authorize_VLAN
vlan 10
```

```
interface ge0/5
  no native-vlan
  no shutdown
!
!
bridge 20
  name Guest_VLAN
  vlan 20
  interface ge0/5
    no native-vlan
    no shutdown
  !
!
bridge 30
  name Critical_VLAN
  vlan 30
  interface ge0/5
    no native-vlan
    no shutdown
  !
!
bridge 40
  name Restricted_VLAN
  vlan 40
  interface ge0/5
    no native-vlan
    no shutdown
  !
!
vpn 0
  interface ge0/0
    ip address 10.1.15.15/24
    tunnel-interface
    encapsulation ipsec
    ...
  !
  no shutdown
  !
  interface ge0/1
    ip address 60.0.1.16/24
    no shutdown
  !
  interface ge0/2
    ip address 10.1.19.15/24
    no shutdown
  !
  interface ge0/4
    ip address 10.20.24.15/24
    no shutdown
  !
  interface ge0/5
    dot1x
    auth-reject-vlan 40
    auth-fail-vlan 30
    guest-vlan 20
    default-vlan 10
    radius-servers freerad1
  !
  no shutdown
  !
  interface ge0/7
    ip address 10.0.100.15/24
    no shutdown
  !
```

```

!
vpn 1
interface ge0/2.1
 ip address 10.2.19.15/24
 mtu      1496
 no shutdown
!
interface irb1
 ip address 56.0.1.15/24
 mac-address 00:00:00:00:aa:01
 no shutdown
 dhcp-server
  address-pool 56.0.1.0/25
  offer-time 600
  lease-time 86400
  admin-state up
  options
   default-gateway 56.0.1.15
!
!
!
vpn 10
interface ge0/2.10
 ip address 10.10.19.15/24
 mtu      1496
 no shutdown
!
interface irb10
 ip address 56.0.10.15/24
 mac-address 00:00:00:00:aa:10
 no shutdown
 dhcp-server
  address-pool 56.0.10.0/25
  offer-time 600
  lease-time 86400
  admin-state up
  options
   default-gateway 56.0.10.15
!
!
!
vpn 20
interface ge0/2.20
 ip address 10.20.19.15/24
 mtu      1496
 no shutdown
!
interface irb20
 ip address 56.0.20.15/24
 mac-address 00:00:00:00:aa:20
 no shutdown
!
!
!
vpn 30
interface ge0/2.30
 ip address 10.30.19.15/24
 mtu      1496
 no shutdown
!
interface irb30
 ip address 56.0.30.15/24
 mac-address 00:00:00:00:aa:30

```

```

    no shutdown
    !
!
vpn 40
interface ge0/2.40
 ip address 10.40.19.15/24
 mtu      1496
 no shutdown
 !
interface irb40
 ip address 56.0.40.15/24
 mac-address 00:00:00:00:aa:40
 no shutdown
 !
!
vpn 512
interface eth0
 ip dhcp-client
 no shutdown
 !
!

```

Configure Control Direction

To configure how the 802.1X interface handles traffic when the client is unauthorized, set the control direction:

```
vEdge(config-dot1x)# control-direction (in-and-out | in-only)
```

The direction can be one of the following:

- **in-and-out**—The 802.1X interface can both send packets to and receive packets from the authorized client. Bidirectional control is the default behavior.
- **in-only**—The 802.1X interface can send packets to the unauthorized client, but cannot receive packets from that client.

Configure Authentication with Wake on LAN

IEEE 802.1X authentication wake on LAN (WoL) allows dormant clients to be powered up when the Cisco vEdge device receives a type of Ethernet frame called the magic packet. Administrators can use wake on LAN when to connect to systems that have been powered down.

When a client that uses wake on LAN and that attaches through an 802.1X port powers off, the 802.1X port becomes unauthorized. The port can only receive and send EAPOL packets, and wake-on-LAN magic packets cannot reach the client. When the device is powered off, it is not authorized, and the switch port is not opened.

Without wake on LAN, when an 802.1X port is unauthorized, the router's 802.1X interface block traffic other than EAPOL packets coming from unauthorized clients.

When you enable wake on LAN on an 802.1X port, the Cisco vEdge device is able to send magic packets even if the 802.1X port is unauthorized.

To enable wake on LAN on an 802.1X interface, use the following command:

```
vEdge(config)# vpn 0 interface interface-name dot1x
vEdge(config-dot1x)# wake-on-lan
```

Configure 802.1X Host Mode

The host mode of an 802.1X interfaces determines whether the interface grants access to a single client or to multiple clients. Three host modes are available:

- Single-host mode—The 802.1X interface grants access only to the first authenticated client. All other clients attempting access are denied and dropped.
- Multiple-host mode—A single 802.1X interface grants access to multiple clients. In this mode, only one of the attached clients must be authorized for the interface to grant access to all clients. If the interface becomes unauthorized, the Cisco vEdge device denies network access to all the attached clients.
- Multiple-authentication mode—A single 802.1X interface grants access to multiple authenticated clients on data VLANs.

To configure the host mode of the 802.1X interface, use the following command:

```
vEdge(config)# vpn 0 interface interface-name dot1x
vEdge(config-dot1x)# host-mode (multi-auth | multi-host | single-host)
```

Set the Timeout for Inactive Clients

By default, when a client has been inactive on the network for 1 hour, its authentication is revoked, and the client is timed out. To change the timeout interval, use the following command:

```
vEdge(config)# vpn 0 interface interface-name dot1x
vEdge(config-dot1x)# timeout inactivity minutes
```

The timeout interval can be from 0 through 1440 minutes (24 hours).

Enable Periodic Client Reauthentication

By default, once a client session is authenticated, that session remains functional indefinitely. To enable the periodic reauthentication of 802.1X clients, configure the number of minutes between reauthentication attempts:

```
vEdge(config)# vpn 0 interface interface-name dot1x
vEdge(config-dot1x)# reauthentication minutes
```

The time can be from 0 through 1440 minutes (24 hours)

Configure Dynamic Authorization Service for RADIUS Change of Authorization

Dynamic authorization service (DAS) allows an 802.1X interface on a Cisco vEdge device to accept change of authorization (CoA) requests from a RADIUS or other authentication server and to act on the requests. The Cisco SD-WAN implementation of DAS supports disconnect packets, which immediately terminate user sessions, and reauthentication CoA requests, which modify session authorization attributes.

DAS, defined in RFC 5176, is an extension to RADIUS that allows the RADIUS server to dynamically change 802.1X session information without requiring the Cisco vEdge device to initiate the change request. When you enable DAS on the Cisco vEdge device, the router opens a socket to listen for CoA requests from the RADIUS server. If the network administrator of a RADIUS server modifies the authentication of an 802.1X client, the RADIUS server sends a CoA request to inform the router about the change of authorization. When the router receives the CoA request, it processes the requested change.

To enable DAS for an 802.1X interface, you configure information about the RADIUS server from which the interface can accept CoA requests. In the context of configuring DAS, the Cisco vEdge device is the server and the RADIUS server (or other authentication server) is the client.

To configure the RADIUS server from which to accept CoA requests, configure the server's IP address and the password that the RADIUS server uses to access the router's 802.1X interface:

```
vEdge(config)# vpn 0 interface interface-name dot1x
vEdge(config-dot1x)# das
```



```
vEdge(config-das)# client ip-address
vEdge(config-das)# secret-key password
```

You can configure the VPN through which the RADIUS server is reachable:

```
vEdge(config-das)# vpn vpn-id
```

By default, the 802.1X interface uses UDP port 3799 to listen for CoA request from the RADIUS server. You can change the port number:

```
vEdge(config-das)# port port-number
```

The port number can be a value from 1 through 65535. If you configure DAS on multiple 802.1X interfaces on a Cisco vEdge device, you must configure each interface to use a different UDP port.

By default, the CoA requests that the Cisco vEdge device receives from the DAS client are all honored, regardless of when the router receives them. To have the router handle CoA within a specified time, you require that the DAS client timestamp all CoA requests:

```
vEdge(config-das)# require-timestamp
```

With this configuration, the Cisco vEdge device processes only CoA requests that include an event timestamp. Non-timestamped CoA requests are dropped immediately.

When timestamping is configured, both the Cisco vEdge device and the RADIUS server check that the timestamp in the CoA request is current and within a specific time window. The default time window is 300 seconds (5 minutes). This behavior means that if the DAS timestamps a CoA at 15:00 and the router receives it at 15:04, the router honors the request. However, if the router receives the request at 15:10, the router drops the CoA request. You can change the time window to a time from 0 through 1000 seconds:

```
vEdge(config-das)# time-window seconds
```

Configure RADIUS Authentication and Accounting Attributes

For IEEE 802.1X authentication and accounting, the Cisco vEdge device, acting as a network access server (NAS), sends RADIUS attribute-value (AV) pairs to the RADIUS server. These AV pairs are defined in RFC 2865, RADIUS, RFC 2866, RADIUS Accounting, and RFC 2869, RADIUS Extensions. The AV pairs are placed in the Attributes field of the RADIUS packet.

By default, when you enable IEEE 802.1X port security, the following authentication attributes are included in messages sent to the RADIUS server:

Attribute Number	Attribute Name	Description
1	User-Name	Name of the user to be authenticated.
5	NAS-Port	Physical port number on the Cisco vEdge device that is authenticating the user.
12	Framed-MTU	Maximum MTU configured for the user.
30	Called-Station-Id	Phone number that the user called, using dialed number identification (DNIS) or similar technology used to access the RADIUS server.
31	Calling-Station-Id	Phone number that the call came in to the server, using automatic number identification (ANI) or similar technology.
44	Acct-Session-Id	Unique session identifier.

Attribute Number	Attribute Name	Description
61	NAS-Port-Type	Type of physical port on the Cisco vEdge device that is authenticating the user.
77	Connect-Info	Nature of the user's connection.
79	EAP-Message	Encapsulate Extended Access Protocol (EAP) packets, to allow the Cisco vEdge device to authenticate dial-in users via EAP without having to run EAP.
80	Message-Authenticator	Sign RADIUS Access-Requests to prevent these requests from being spoofed by ARAP, CHAP, or EAP.

When you enable RADIUS accounting, the following accounting attributes are included, by default, in messages sent to the RADIUS server:

Attribute Number	Attribute Name	Description
1	User-Name	Name of the user to be authenticated.
5	NAS-Port	Physical port number on the Cisco vEdge device that is authenticating the user.
30	Called-Station-Id	Phone number that the user called, using dialed number identification (DNIS) or similar technology used to access the RADIUS server.
31	Calling-Station-Id	Phone number that the call came in to the server, using automatic number identification (ANI) or similar technology.
40	Acct-Status-Type	Mark the beginning and end of an accounting request.
44	Acct-Session-Id	Unique accounting identifier used to match the start and stop records in a log file.
45	Acct-Authentic	How the user was authenticated.
61	NAS-Port-Type	Type of physical port on the Cisco vEdge device that is authenticating the user.
77	Connect-Info	Nature of the user's connection.

Several configuration commands allow you to add additional attribute information to RADIUS packets.

To include the NAS-IP-Address (attribute 4) in messages sent to the RADIUS server to indicate the IP address of the Cisco vEdge device that is acting as a NAS server:

```
vEdge(config-dot1x) nas-ip-address ip-address
```

To include the NAS-Identifier (attribute 32) in messages sent to the RADIUS server, use the following command:

```
vEdge(config-dot1x)# nas-identifier string
```

The NAS identifier is a unique string from 1 through 255 characters long that identifies the Cisco vEdge device that is acting as a NAS server.

To include a RADIUS authentication or accounting attribute of your choice in messages sent to the RADIUS server, use the following commands:

```
vEdge(config-dot1x)# auth-req-attr attribute-number (integer integer | octet
octet | string string)
vEdge(config-dot1x)# acct-req-attr attribute-number (integer integer | octet
octet | string
string)
```

Specify the desired value of the attribute as an integer, octet value, or string, depending on the attribute. For example, to set the Service-Type attribute to be authenticate-only:

```
vEdge(config-dot1x)# auth-req-attr 6 integer 8
```

Configure IEEE 802.11i Authentication

For Cisco vEdge device that support wireless LANs (WLANs), you can configure the router to support either a 2.4-GHz or 5-GHz radio frequency. Then, you segment the WLAN into multiple broadcast domains, which are called virtual access points, or VAPs. Users who connect to a VAP can be unauthenticated, or you can configure IEEE 802.11i authentication for each VAP.

For information about configuring the WLAN interface itself, see *Configuring WLAN Interfaces*.

To enable user authentication on the WLAN, you create a VAP on the desired radio frequency and then you configure Wi-Fi protected access (WPA) or WPA2 data protection and network access control for the VAP. WPA authenticates individual users on the WLAN using a username and password. WPA uses the Temporal Key Integrity Protocol (TKIP), which is based on the RC4 cipher. WPA2 implements the NIST FIPS 140-2-compliant AES encryption algorithm along with IEEE 802.1X-based authentication, to enhance user access security over WPA. WPA2 uses the Counter Mode Cipher Block Chaining Message Authentication Code Protocol (CCMP), which is based on the AES cipher. Authentication is done either using preshared keys or through RADIUS authentication.

To enable personal authentication, which requires users to enter a password to connect to the WLAN, configure the authentication and password:

```
vEdge(config)# wlan frequency
vEdge(config-wlan)# interface vap number
vEdge(config-vap)# no shutdown
vEdge(config-vap)# data-security (wpa-personal | wpa/wpa2-personal | wpa2-personal)
vEdge(config-vap)# wpa-personal-key password
```

For the security, configure either WPA, WPA2, or both (WPA/WPA2). Enter the password either as clear text or an AES-encrypted key.

For each VAP, you can customize the security mode to control wireless client access.

To enable enterprise WPA security, configure the authentication and the RADIUS server to perform the authentication:

```
vEdge(config-vap)# data-security (wpa-enterprise | wpa/wpa2-enterprise | wpa2-enterprise)
vEdge(config-vap)# radius-servers tag
```

For the security, configure either WPA, WPA2, or both (WPA/WPA2). Enter the password either as clear text or an AES-encrypted key.

In the **radius-servers** command, enter the tags associated with one or two RADIUS servers to use for 802.11i authentication. (You configure the tags with the **system radius server tag** command.) If you specify tags for two RADIUS servers, they must both be reachable in the same VPN. If you do not configure a priority value

when you configure the RADIUS server with the **system radius server priority** command, the order in which you list the IP addresses is the order in which the RADIUS servers are tried.

By default, management frames sent on the WLAN are not encrypted. For each VAP, you can configure the encryption to be optional or required:

```
vEdge(config-vap)# mgmt-security (none | optional | required)
```



CHAPTER 4

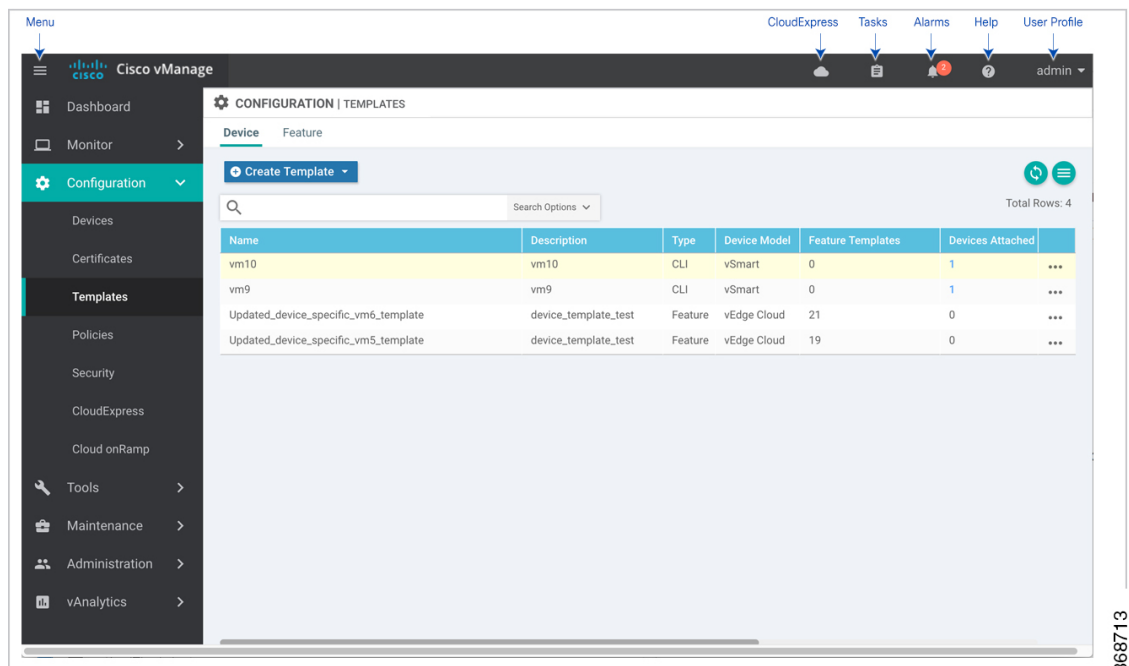
Create a Device Template from Feature Templates

Device templates define a device's complete operational configuration. A device template consists of a number of feature templates. Each feature template defines the configuration for a particular Cisco SD-WAN software feature. Some feature templates are mandatory, indicated with an asterisk (*), and some are optional. Each mandatory feature template, and some of the optional ones, have a factory-default template. For software features that have a factory-default template, you can use either the factory-default template (named `Factory_Default_feature-name_Template`) or you can create a custom feature template.

Create a Device Template from Feature Templates

To create a device template:

Figure 1: Create a Device Template in Cisco vManage



1. In the Device tab, click the Create Template drop-down and select From Feature Template.

2. From the Device Model drop-down, select the type of device for which you are creating the template. vManage NMS displays all the feature templates for that device type. The required feature templates are indicated with an asterisk (*), and the remaining templates are optional. The factory-default template for each feature is selected by default.
3. In the Template Name field, enter a name for the device template. This field is mandatory and can contain only uppercase and lowercase letters, the digits 0 through 9, hyphens (-), and underscores (_). It cannot contain spaces or any other characters.
4. In the Description field, enter a description for the device template. This field is mandatory, and it can contain any characters and spaces.
5. To view the factory-default configuration for a feature template, select the desired feature template and click View Template. Click Cancel to return to the Configuration Template screen.
6. To create a custom template for a feature, select the desired factory-default feature template and click Create Template. The template form is displayed. The top of the form contains fields for naming the template, and the bottom contains fields for defining feature parameters.
7. In the Template Name field, enter a name for the feature template. This field is mandatory and can contain only uppercase and lowercase letters, the digits 0 through 9, hyphens (-), and underscores (_). It cannot contain spaces or any other characters.
8. In the Description field, enter a description for the feature template. This field is mandatory, and it can contain any characters and spaces.
9. For each field, enter the desired value. You may need to click a tab or the plus sign (+) to display additional fields.
10. When you first open a feature template, for each parameter that has a default value, the scope is set to Default (indicated by a check mark), and the default setting or value is shown. To change the default or to enter a value, click the scope drop-down to the left of the parameter field and select one of the following:

Table 26:

Parameter Scope	Scope Description
Device Specific (indicated by a host icon)	<p>Use a device-specific value for the parameter. For device-specific parameters, you cannot enter a value in the feature template. You enter the value when you attach a device to a device template .</p> <p>When you click Device Specific, the Enter Key box opens. This box displays a key, which is a unique string that identifies the parameter in a CSV file that you create. This file is an Excel spreadsheet that contains one column for each key. The header row contains the key names (one key per column), and each row after that corresponds to a device and defines the values of the keys for that device. You upload the CSV file when you attach a device to a device template. For more information, see Use Variable Values in Configuration Templates .</p> <p>To change the default key, type a new string and move the cursor out of the Enter Key box.</p> <p>Examples of device-specific parameters are system IP address, hostname, GPS location, and site ID.</p>

Parameter Scope	Scope Description
Global (indicated by a globe icon)	Enter a value for the parameter, and apply that value to all devices. Examples of parameters that you might apply globally to a group of devices are DNS server, syslog server, and interface MTUs.

- For some groups of parameters, you can mark the entire group as device-specific. To do this, click the Mark as Optional Row box. These parameters are then grayed out so that you cannot enter a value for them in the feature template. You enter the value or values when you attach a device to a device template.
- Click Save.
- Repeat Steps 7 through 13 to create a custom template for each additional software feature. For details on creating specific feature templates, see the templates listed in Available Feature Templates.
- Click Create. The new configuration template is displayed in the Device Template table. The Feature Templates column shows the number of feature templates that are included in the device template, and the Type column shows "Feature" to indicate that the device template was created from a collection of feature templates.

Another way to create device templates from feature templates is to first create one or more custom feature templates and then create device templates. You can create multiple feature templates for the same feature. For a list of feature templates, see Available Feature Templates .

- From the Templates title bar, select Feature.
- Click the Add Template button.
- In the left pane, from Select Devices, select the type of device for which you are creating a template. You can create a single feature template for features that are available on multiple device types. You must, however, create separate feature templates for software features that are available only on the device type you are configuring.
- In the right pane, select the feature template. The template form is displayed. The top of the form contains fields for naming the template, and the bottom contains fields for defining required parameters. If the feature has optional parameters, the bottom of the template form shows a plus sign (+) after the required parameters.
- In the Template Name field, enter a name for the feature template. This field is mandatory and can contain only uppercase and lowercase letters, the digits 0 through 9, hyphens (-), and underscores (_). It cannot contain spaces or any other characters.
- In the Description field, enter a description for the feature template. This field is mandatory, and it can contain any characters and spaces.
- For each required parameter, choose the desired value, and if applicable, select the scope of the parameter. Select the scope from the drop-down menu to the left of each parameter's value box
- Click the plus sign (+) below the required parameters to set the values of optional parameters.
- Click Save.
- Repeat Steps 2 to 9 for each additional feature template you wish to create.
- From the Templates title bar, select Device.

12. Click the Create Template drop-down and select From Feature Template.
13. From the Device Model drop-down, select the type of device for which you are creating the device template. vManage NMS displays the feature templates for the device type you selected. The required feature templates are indicated with an asterisk (*). The remaining templates are optional.
14. In the Template Name field, enter a name for the device template. This field is mandatory and can contain only uppercase and lowercase letters, the digits 0 through 9, hyphens (-), and underscores (_). It cannot contain spaces or any other characters.
15. In the Description field, enter a description for the device template. This field is mandatory, and it can contain any characters and spaces.
16. To view the factory-default configuration for a feature template, select the desired feature template and click View Template. Click Cancel to return to the Configuration Template screen.
17. To use the factory-default configuration, click Create to create the device template. The new device template is displayed in the Device Template table. The Feature Templates column shows the number of feature templates that are included in the device template, and the Type column shows "Feature" to indicate that the device template was created from a collection of feature templates.
18. To modify the factory-default configuration, select the feature template for which you do not wish to use the factory-default template. From the drop-down list of available feature templates, select a feature template that you created.
19. Repeat Step 18 for each factory-default feature template you wish to modify.
20. Click Create. The new configuration template is displayed in the Device Template table. The Feature Templates column shows the number of feature templates that are included in the device template, and the Type column shows "Feature" to indicate that the device template was created from a collection of feature templates.

- [Configure Devices, on page 82](#)

Configure Devices

Create a Device CLI Template

To create a device template by entering a CLI text-style configuration directly on the vManage NMS:

1. In the Device tab, click the Create Template drop-down and select CLI Template.
2. From the Device Type drop-down, select the type of device for which you are creating the template.
3. In the Template Name field, enter a name for the device template. This field is mandatory and can contain only uppercase and lowercase letters, the digits 0 through 9, hyphens (-), and underscores (_). It cannot contain spaces or any other characters.
4. In the Description field, enter a description for the device template. This field is mandatory, and it can contain any characters and spaces.
5. In the CLI Configuration box, enter the configuration either by typing it, cutting and pasting it, or uploading a file.

6. To convert an actual configuration value to a variable, select the value and click Create Variable. Enter the variable name, and click Create Variable. You can also type the variable name directly, in the format `{{variable-name}}`; for example, `{{hostname}}`.
7. Click Add. The new device template is displayed in the Device Template table. The Feature Templates column shows the number of feature templates that are included in the device template, and the Type column shows "CLI" to indicate that the device template was created from CLI text.

Manage Device Templates

Edit a Device Template

1. In the Device or Feature tab, select a template.
2. Click the More Actions icon to the right of the row and click Edit.

You cannot change the name of a device or feature template when that template is attached to a device.

Note that you can edit templates simultaneously from one or more vManage servers. For simultaneous template edit operations, the following rules apply:

- You cannot edit the same device or feature template simultaneously.
- When you are editing a device template, all other feature templates attached to that device template are locked and you cannot perform any edit operations on them.
- When you are editing a feature template that is attached to a device template, that device template as well as all other feature templates attached to it are locked and you cannot perform any edit operations on them.

Delete a Template

Deleting a template does not remove the associated configuration from devices.

1. In the Device or Feature tab, select a template.
2. Click the More Actions icon to the right of the row and click Delete.
3. Click OK to confirm deletion of the template.

Copy a Template

1. In the Device or Feature tab, select a template.
2. Click the More Actions icon to the right of the row and click Copy.
3. Enter a new template name and description.
4. Click Copy.

Edit a CLI Device Template

1. In the Device tab, select a template.

2. Click the More Actions icon to the right of the row and click Edit.
3. In the Device CLI Template window, edit the template.
4. Click Update.

View Device Templates

•

View a Template

1. In the Device or Feature tab, select a template.
2. Click the More Actions icon to the right of the row and click View.

View Device Templates Attached to a Feature Template

1. In the Feature tab, select a template.
2. Click the More Actions icon to the right of the row and click Show Attached Device Templates. The View Attached Device Templates popup window opens, displaying the names of the device templates to which the feature template is attached.

View Devices Attached to a Device Template

For a device template that you created from feature templates:

1. In the Device tab, select a template.
2. Click the More Actions icon to the right of the row and click Attach Devices.
3. In the Attach Devices window, click the Attached Devices tab.

For a device template that you created from a CLI template:

1. In the Device tab, select a template.
2. Click the More Actions icon to the right of the row and click Show Attached Devices.

Attach and Detach a Device Template

On Cisco vEdge devices in the overlay network, you can perform the same operations, in parallel, from one or more vManage servers. You can perform the following template operations in parallel:

- Attach devices to a device template
- Detach devices from a device template
- Change the variable values for a device template that has devices attached to it

For template operations, the following rules apply:

- When a device template is already attached to a device, you can modify one of its feature templates. Then when you click Update ► Configure Devices, all other template operations—including attach devices, detach devices, and edit device values—are locked on all vManage servers until the update operation completes. This means that a user on another vManage server cannot perform any template operations until the update completes.
- You can perform the attach and detach device template operations on different devices, from one or more vManage servers, at the same time. However, if any one of these operations is in progress on one vManage server, you cannot edit any feature templates on any of the servers until the attach or detach operation completes.

Attach Devices to a Device Template

To attach one or more devices to a device template:

1. In the Device tab, select a template.
 2. Click the More Actions icon to the right of the row and click Attach Devices. The Attach Devices dialog box opens with the Select Devices tab selected
 3. In the Available Devices column on the left, select a group and search for one or more devices, select a device from the list, or click Select All.
 4. Click the arrow pointing right to move the device to the Selected Devices column on the right.
 5. Click Attach.
 6. If the template contains variables, enter the missing variable values for each device you selected in one of the following ways:
 - Enter the values manually for each device either in the table column or by clicking the More Actions icon to the right of the row and clicking Edit Device Template. When you are using optional rows, if you do not want to include the parameter for the specific device, do not specify a value.
 - Click Import File in the upper right corner of the screen to upload a CSV file that lists all the variables and defines each variable's value for each device.
1. Click Update
 2. Click Next. If any devices have the same system IP address, a pop-up or an error message is displayed when you click Next. Modify the system IP addresses so that there are no duplicates, and click Save. Then click Next again.
 3. In the left pane, select the device, to preview the configuration that is ready to be pushed to the device. The right pane displays the device's configuration and the Config Preview tab in the upper right corner is selected. Click the Config Diff tab to view the differences between this configuration and the configuration currently running on the device, if applicable. Click the Back button to edit the variable values entered in the previous screen.
 4. If you are attaching a Cisco vEdge device, click Configure Device Rollback Timer located at the bottom of the left pane, to configure the time interval at which the device rolls back to its previous configuration if the router loses its control connection to the overlay network. The Configure Device Rollback Time dialog box is displayed.
 - a. From the Devices drop-down, select a device.

- b. To enable the rollback timer, in the Set Rollback slider beneath the Devices drop-down, drag the slider to the left to enable the rollback timer. When you do this, the slider changes in color from gray to green.
 - c. To disable the rollback timer, click the Enable Rollback slider. When you disable the timer, the Password field pops up. Enter the password that you used to log in to the vManage NMS.
 - d. In the Device Rollback Time slider, drag the slider to the desired value. The default time is 5 minutes. You can configure a time from 6 to 15 minutes.
 - e. To exclude a device from the rollback timer setting, click Add Exception and select the devices to exclude.
 - f. The table at the bottom of the Configure Device Rollback Time dialog box lists all the devices to which you are attaching the template and their rollback time. To delete a configured rollback time, click the Trash icon to right right of the device name.
 - g. Click Save.
5. Click Configure Devices to push the configuration to the devices. The Status column displays whether the configuration was successfully pushed. Click the right angle bracket to the left of the row to display details of the push operation.

Export a Variables Spreadsheet in CSV Format for a Template

1. In the Device tab, select a device template.
2. Click the More Actions icon to the right of the row and click Export CSV.

Change the Device Rollback Timer

By default, when you attach a Cisco vEdge device to a configuration template, if the router is unable to successfully start after 5 minutes, it returns to, or rolls back to, the previous configuration. For a configuration that you have created from the CLI, you can change the device's rollback timer:

1. In the Device tab, select a device template.
2. Click the More Actions icon to the right of the row and click Change Device Values. The right pane displays the device's configuration, and the Config Preview tab in the upper right corner is selected.
3. In the left pane, click the name of a device.
4. Click Configure Device Rollback Timer located at the bottom of the left pane. The Configure Device Rollback Time dialog box is displayed.
5. From the Devices drop-down, select a device.
6. To enable the rollback timer, in the Set Rollback slider beneath the Devices drop-down, drag the slider to the left to enable the rollback timer. When you do this, the slider changes in color from gray to green.
7. To disable the rollback timer, click the Enable Rollback slider. When you disable the timer, the Password field pops up. Enter the password that you used to log in to the vManage NMS.
8. In the Device Rollback Time slider, drag the slider to the desired value. The default time is 5 minutes. You can configure a time from 6 to 15 minutes.

9. To exclude a device from the rollback timer setting, click Add Exception and select the devices to exclude.
10. The table at the bottom of the Configure Device Rollback Time dialog box lists all the devices to which you are attaching the template and their rollback time. To delete a configured rollback time, click the Trash icon to right right of the device name.
11. Click Save.
12. Click Configure Devices to push the configuration to the devices. The Status column displays whether the configuration was successfully pushed. Click the right angle bracket to the left of the row to display details of the push operation.

Preview Device Configuration and View Configuration Differences

For a configuration that you have created from the CLI:

1. In the Device tab, select a device template.
2. Click the More Actions icon to the right of the row and click Change Device Values. The right pane displays the device's configuration, and the Config Preview tab in the upper right corner is selected.
3. In the left pane, click the name of a device.
4. Click the Config Diff tab to view the differences between this configuration and the configuration currently running on the device, if applicable. Click the Back button to edit the variable values entered in the previous screen.
5. Click Configure Devices to push the configuration to the devices. The Status column displays whether the configuration was successfully pushed. Click the right angle bracket to the left of the row to display details of the push operation.

Change Variable Values for a Device

For a configuration that you have created from device configuration templates, if the templates contain variables, the vManage NMS can automatically populate the variables with actual values when you attach the templates to the devices. To do this, you create an Excel file that lists the variable values for each device and save the file in CSV format. You can also enter values for these variables manually.

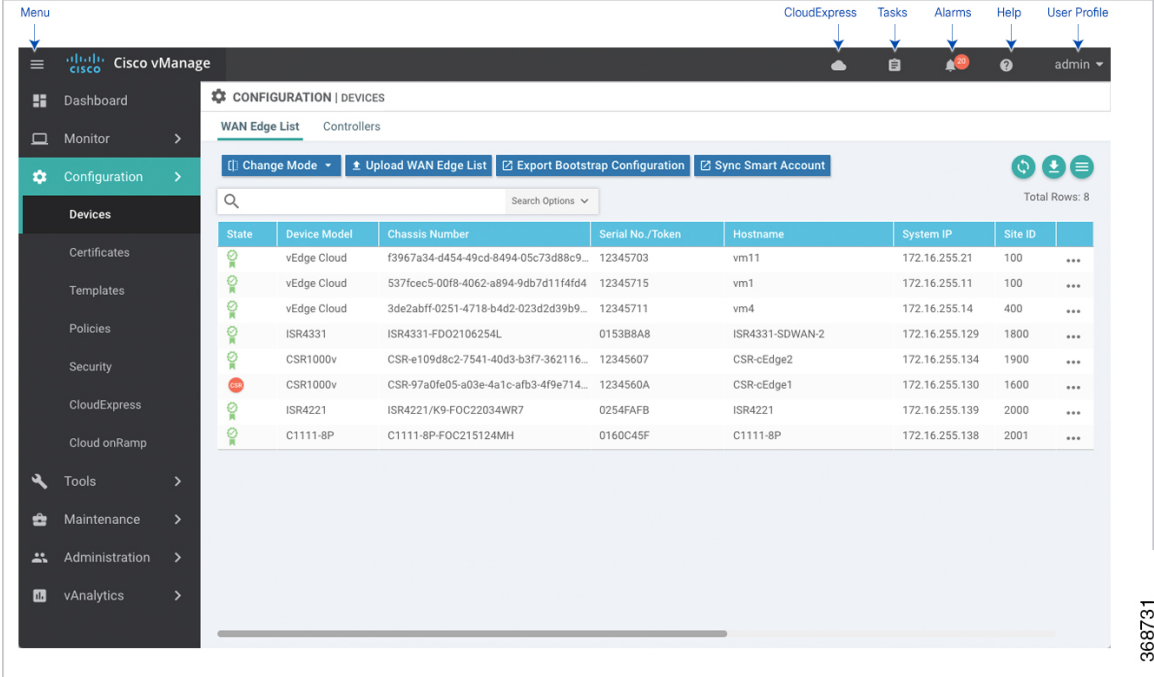
After you have pushed the configuration to a device, you can change the value assigned to any variable:

1. In the Device tab, select the device template.
2. Click the More Actions icon to the right of the row, and click Change Device Values. The screen displays a table of all the devices that are attached to that device template.
3. For the desired device, click the More Actions icon to the right of the row, and click Edit Device Template.
4. In the Update Device Template pop-up, enter values for the items in the variable list.
5. Click Update.
6. Click Next.

- Click Configure Devices to push the configuration to the device. The Status column displays if the configuration was successfully pushed or not. Click the right angle bracket to the left of the row to display details of the push operation.

Configuring Devices using vManage

Use the **Devices** screen to add and delete devices, toggle the mode of a device between CLI and vManage, upload the WAN Edge Serial number file, export bootstrap configuration and, and perform other device-related tasks.



The screenshot shows the Cisco vManage interface for configuring devices. The main content area is titled "CONFIGURATION | DEVICES" and has a "WAN Edge List" tab selected. Below the tab are several action buttons: "Change Mode", "Upload WAN Edge List", "Export Bootstrap Configuration", and "Sync Smart Account". A search bar is present above a table of devices. The table has 8 rows and 7 columns: State, Device Model, Chassis Number, Serial No./Token, Hostname, System IP, and Site ID. The status of each device is indicated by a green checkmark icon.

State	Device Model	Chassis Number	Serial No./Token	Hostname	System IP	Site ID
✓	vEdge Cloud	f3967a34-d454-49cd-8494-05c73d88c9...	12345703	vm11	172.16.255.21	100
✓	vEdge Cloud	537fcec5-00f8-4062-a894-9db7d11f4fd4	12345715	vm1	172.16.255.11	100
✓	vEdge Cloud	3de2abff-0251-4718-b4d2-023d2d39b9...	12345711	vm4	172.16.255.14	400
✓	ISR4331	ISR4331-FDO2106254L	0153B8A8	ISR4331-SDWAN-2	172.16.255.129	1800
✓	CSR1000v	CSR-e109d8c2-7541-40d3-b3f7-362116...	12345607	CSR-cEdge2	172.16.255.134	1900
✗	CSR1000v	CSR-97a0fe05-a03e-4a1c-afb3-4f9e714...	1234560A	CSR-cEdge1	172.16.255.130	1600
✓	ISR4221	ISR4221/K9-FOC22034WR7	0254F4FB	ISR4221	172.16.255.139	2000
✓	C1111-8P	C1111-8P-FOC215124MH	0160C45F	C1111-8P	172.16.255.138	2001

368731

Change Configuration Modes

A device can be in either of these configuration modes:

- vManage mode—A template is attached to the device and you cannot change the configuration on the device by using the CLI.
- CLI mode – No template is attached to the device and the device can be configured locally by using the CLI.

When you attach a template to a device from vManage, it puts the device in vManage mode. You can change the device back to CLI mode if needed to make local changes to its configuration.

To toggle a router from vManage mode to CLI mode:

- In WAN Edge List tab, select a device.
- Click the Change Mode drop-down and select CLI mode.

An SSH window opens. To log in to the device, enter a username and password. You can then issue CLI commands to configure or monitor the device.

To toggle a controller device from vManage mode to CLI mode:

1. In the Controllers tab, select a device.
2. Click the Change Mode drop-down.
3. Select CLI mode and then select the device type. The Change Mode CLI window opens.
4. From the vManage mode pane, select the device and click the right arrow to move the device to the CLI mode pane.
5. Click Update to CLI Mode.

An SSH window opens. To log in to the device, enter a username and password. You can then issue CLI commands to configure or monitor the device.

Upload WAN Edge Router Authorized Serial Number File

The WAN Edge router authorized serial number file contains the chassis and serial numbers of all valid Cisco vEdge devices in the overlay network. You retrieve a serial number file from the Cisco Plug-and-Play (PnP) portal and upload it to the vManage NMS. Then, from the vManage NMS, you send it to the controllers in the network. This file is required to allow the Cisco SD-WAN overlay network components to validate and authenticate each other and thus to allow the overlay network to become operational.

To upload the WAN edge router authorized serial number file to the vManage NMS and then download it to all the controllers in the overlay network:

1. In the WAN Edge List tab, click Upload WAN Edge List.
2. In the Upload WAN Edge List window:
 - a. Click Choose File and select the WAN edge router authorized serial number file you received from Cisco SD-WAN.
 - b. To automatically validate the routers and send their chassis and serial numbers to the controllers, ensure that the checkbox Validate the Uploaded WAN Edge List and Send to Controllers is selected. (It is selected by default.) If you do not select this option, you must individually validate each router in Configuration ► Certificates ► WAN Edge List.
 - c. Click Upload.

A list of routers in the network is displayed in the router table, with details about each router.

Upload WAN Edge Router Serial Numbers from Cisco Smart Account

Chassis and serial numbers of all valid Cisco vEdge devices in the overlay network are required to allow the Cisco SD-WAN overlay network components to validate and authenticate each other and thus to allow the overlay network to become operational.

To upload the WAN edge router authorized serial numbers from a Cisco Smart account to the vManage NMS and then download it to all the controllers in the overlay network:

1. In the WAN Edge List tab, click Sync Smart Account.
2. In the Sync Smart Account window:

- a. Enter the username and password for your Smart account..
- b. To automatically validate the routers and send their chassis and serial numbers to the controllers, ensure that the checkbox Validate the Uploaded WAN Edge List and Send to Controllers is selected. (It is selected by default.) If you do not select this option, you must individually validate each router in Configuration ► Certificates ► WAN Edge List.
- c. Click Sync.

A list of routers in the network is displayed in the router table, with details about each router.

Generate Bootstrap Configuration for a vEdge Cloud Router

For vEdge Cloud routers, you need to generate a bootstrap configuration file that you use when you create vEdge cloud VM instances.

To generate and download a bootstrap configuration for one or more vEdge Cloud routers:

1. In the WAN Edge List tab, click the Export Bootstrap Configuration button.
2. In the Export Bootstrap Configuration window, in the Bootstrap Configuration field, click Cloud-Init or Encoded String, depending the Hypervisor you are using to bring up the vEdge Cloud router.
3. Select the devices to configure from the Available Devices pane, or click Select All to select all devices.
4. Click the right arrow to move the devices to the Selected Devices pane.
5. Click Generate Configuration. The configurations are downloaded to the vManage NMS.
6. Provision the vEdge Cloud router instance in AWS, KVM, or ESXi with the bootstrap configuration. By default, ge0/0 is the device's tunnel interface and is a DHCP client. To use an interface other than ge0/0 as the tunnel interface or to use a static IP as the IP address, reconfigure the device through the CLI. For more information about configuring interfaces, see Configure Network Interfaces.

After you provision the vEdge Cloud router instance, vManage NMS installs a certificate on the device and the device's token changes to a serial number. After the device's control connections to vManage NMS come up, any templates attached to the device are automatically pushed to the device.

Export Device Data in CSV Format

In an overlay network, you might have multiple devices of the same type that have identical or effectively identical configurations. For example, in a network with redundant Cisco vSmart Controllers, each controller must be configured with identical policies. Another example is a network with Cisco vEdge devices at multiple sites, where each Cisco vEdge device is providing identical services at each site.

Because the configurations for these devices are essentially identical, you can create one set of feature templates, which you then consolidate into one device template that you use to configure all the devices. You can create an Excel file in CSV format that lists the variables and defines each device specific variable value for each device. Then you can load the file when you attach a device template to a device.

To export data for all devices to a file in CSV format, click the Export icon. This icon, which is a downward-pointing arrow, is located to the right of the filter criteria both in the WAN Edge List and in the Controllers tab.

vManage NMS downloads all data from the device table to an Excel file in CSV format.

View and Copy Device Configuration

View a Device's Running Configuration

Running configuration is configuration information that vManage obtains from the memory of a device. This information can be useful for troubleshooting.

To view a device's running configuration:

1. In the WAN Edge List or Controllers tab, select the device.
2. Click the More Actions icon to the right of the row and click Running Configuration.

View a Device's Local Configuration

Local configuration is configuration that vManage has stored for a device. This information can be useful for troubleshooting or for determining how to access a device if, for example, a device is not reachable from vManage.

To view a device's local configuration created using Configuration ► Templates:

1. In the WAN Edge List or Controllers tab, select the device.
2. Click the More Actions icon to the right of the row and click Local Configuration.

Copy Router Configuration

When you are replacing one router at a site with another router, you copy the old router's configuration to the new router. Then you remove the old router from the network and add the new one.

To copy the configuration from the old router to the new router:

1. In the Configuration ► Certificates screen, mark the new Cisco vEdge device as invalid.
2. In the Configuration ► Devices screen, in the WAN Edge List tab, select the old router.
3. Click the More Actions icon to the right of the row and click Copy Configuration.
4. In the Copy Configuration window, select the new router.
5. Click Update to confirm the copy of the configuration.

After you have copied the configuration to the new router, you can add the new router to the network. First, delete the old router from the network, as described below. Then add the new router to the network:

1. In the Configuration ► Certificates screen, mark the new router as valid.
2. Click Send to Controller.

Delete a WAN Edge Router

Deleting a router removes its serial and chassis numbers from the WAN edge router serial number list and permanently removes the router's configuration from the vManage NMS. Delete a router if you need to remove it from your deployment.

1. In the Configuration ► Certificates screen, mark the WAN Edge router as invalid.

2. In the Configuration ► Devices screen, in the WAN Edge List tab, select the router.
3. Click the More Actions icon to the right of the row and click Delete WAN Edge.
4. Click OK to confirm deletion of the device.
5. In the Configuration ► Certificates screen, click Send to Controller.

Decommission a vEdge Cloud router

Decommissioning a vEdge Cloud router removes the device's serial number from vManage NMS and generates a new token for the device. To do so:

1. In the WAN Edge List tab, select a vEdge Cloud router.
2. Click the More Actions icon to the right of the row and click Decommission WAN Edge.
3. Click OK to confirm the decommissioning of the router.

View Template Log and Device Bringup

View Log of Template Activities

A log of template activities contains information that relates to creating, editing, and deleting configuration templates, and the status of attaching configuration templates to devices. This information can be useful for troubleshooting.

To view a log of template activities:

1. In the WAN Edge List or Controllers tab, select the device.
2. Click the More Actions icon to the right of the row and click Template Log.

View Status of Device Bringup

You can view the status of the operations involved in bringing a router or controller up in the overlay network. This information can help you monitor these operations.

To view the status of a device bringup:

1. In the WAN Edge List or Controllers tab, select the device.
2. Click the More Actions icon to the right of the row and click Device Bring Up.

Add a Cisco vBond Orchestrator

A Cisco vBond Orchestrator automatically orchestrates connectivity between Cisco vEdge devices and vManage controllers. If any Cisco vEdge device or Cisco vSmart Controller is behind a NAT, the Cisco vBond Orchestrator also serves as an initial NAT-traversal orchestrator. To add a Cisco vBond Orchestrator:

1. In the Controllers tab, click the Add Controller drop-down and select vBond.
2. In the Add vBond window:
 - a. Enter the management IP address of the vBond controller.
 - b. Enter the username and password to access the vBond orchestrator.

- c. Select the Generate CSR checkbox to allow the certificate-generation process to occur automatically.
 - d. Click Add.
 3. Repeat Steps 1 and 2 to add additional Cisco vBond Orchestrators.

The new Cisco vBond Orchestrator is added to the list of controllers in the Controllers screen.

Configure Cisco vSmart Controllers

Add a vSmart Controller

After the Cisco vBond Orchestrator authenticates Cisco vEdge devices, the Cisco vBond Orchestrator provides Cisco vEdge devices information that they need to connect to the Cisco vSmart Controller. A Cisco vSmart Controller controls the flow of data traffic throughout the network via data and app-route policies. To configure Cisco vSmart Controllers:

1. In the Controllers tab, click the Add Controller drop-down and select vSmart.
2. In the Add vSmart window:
 - a. Enter the system IP address of the Cisco vSmart Controller.
 - b. Enter the username and password to access the Cisco vSmart Controller.
 - c. Select the protocol to use for control-plane connections. The default is DTLS. The DTLS (Datagram Transport Layer Security) protocol is designed to provide security for UDP communications.
 - d. If you select TLS, enter the port number to use for TLS connections. The default is 23456.
The TLS (Transport Socket Layer) protocol that provides communications security over a network.
 - e. Select the Generate CSR checkbox to allow the certificate-generation process to occur automatically.
 - f. Click Add.
3. Repeat Steps 1 and 2 to add additional Cisco vSmart Controllers. The vManage NMS can support up to 20 Cisco vSmart Controllers in the network.

The new Cisco vSmart Controller is added to the list of controllers in the Controllers screen.

Edit Controller Details

Editing controller details lets you update the IP address and login credentials of a controller device. To edit controller details:

1. In the Controllers tab, select the controller.
2. Click the More Actions icon to the right of the row and click Edit.
3. In the Edit window, edit the IP address and the login credentials.
4. Click Save.

Delete a Controller

Deleting a controller removes it from the overlay. Delete a controller if you are replacing it or if you no longer need it in your network.

To delete a controller:

1. In the Controllers tab, select the controller.
2. Click the More Actions icon to the right of the row and click Invalidate.
3. Click OK to confirm the removal of the device and all its control connections.

Configure Reverse Proxy on Controllers

To configure reverse proxy on an individual vManage NMS and Cisco vSmart Controller:

1. In the Controllers tab, select the device.
2. Click the More Actions icon to the right of the row, and click Add Reverse Proxy. The Add Reverse Proxy popup is displayed.
3. Click Add Reverse Proxy.
4. Configure the private IP address and port number for the device. The private IP address is the IP address of the transport interface in VPN 0. The default port number is 12346. This is the port used to establish the connections that handle control and traffic in the overlay network.
5. Configure the proxy IP address and port number for the device, to create the mapping between the private and public IP addresses and port numbers.
6. If the vManage NMS or Cisco vSmart Controller has multiple cores, repeat Steps 4 and 5 for each core.
7. Click Add.

To enable reverse proxy in the overlay network, in vManage NMS select Administration ► Settings. Then click Edit to the right of the Reverse Proxy bar, click Enabled, and click Save.

Create a UCS-E Template

Table 27: Feature History

Feature Name	Release Information	Feature Description
Create a UCS-E Template		This feature allows you to connect a UCS-E interface with a UCS-E server through the interface feature template.

For more information about the Cisco Unified Computing System (UCS) E-Series Servers, see the [Cisco UCS E-Series Servers and the Cisco UCS E-Series Network Compute Engine Hardware Installation Guide](#).

1. From the vManage menu, select Configuration ► Templates.
2. Click Feature.
3. Click Add Template.
4. Select a Cisco XE SD-WAN device from the list.

5. From the Other Templates section, click UCSE.

The UCSE Feature template opens. The top of the form contains fields for naming the template, and the bottom contains fields for configuring the Integrated Management Controller (IMC).

6. In the Template Name field, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
7. In the Description field, enter a description of the template. The description can be up to 2048 characters and can contain only alphanumeric characters.

Configure Bay and Slot for Template

Click the Basic Configuration tab to configure the bay and the slot for the template.

Parameter Name	Description
Bay	Specify the number for the SAS drive bays.
Slot	Specify the slot numbers for the mezzanine adapters.

IMC Configuration

Click the IMC tab to configure the IMC parameters for the template.

Parameter Name	Description
Access Port	<p>Configure the interface as an access port. You can configure only one VLAN on an access port, and the port can carry traffic for only one VLAN.</p> <p>Not all hardware models have a dedicated access port. See the Release Notes for your Cisco SD-WAN release for the supported hardware.</p> <p>Available options:</p> <ul style="list-style-type: none"> • Dedicated • Shared <p>The type of port, GE or TE, depends on the hardware model.</p> <p>For example:</p> <pre>Router(config-ucse)#imc access-port shared-lom ? GE1 GE1 TE2 TE2 TE3 TE3 console Console failover Failover</pre> <p>Some hardware models have GE ports whereas some have TE ports.</p> <p>Depending on the hardware module, the appropriate port (GE or TE) needs to be configured. Otherwise you will get an error.</p> <ul style="list-style-type: none"> • You can obtain the UCS-E module hardware model type by using the following commands: <ul style="list-style-type: none"> show inventory show platform • Failover - sub-option under Shared. <p>For example:</p> <pre>Router(config)#ucse subslot 1/0 Router(config-ucse)#imc access-port ? MGMT MGMT Interface shared-lom Shared LOM Router(config-ucse)#imc access-port shared-lom ? GE1 GE1 TE2 TE2 TE3 TE3 console Console failover Failover</pre>
IPv4 Address	Provide the UCS-E management port address.

Parameter Name	Description
Default Gateway	Gateway tracking determine, for static routes, whether the next hop is reachable before adding that route to the device's route table. Default: Enabled.
VLAN ID	Provide the VLAN number, which can be a value from 1 through 4094.
Assign Priority	Assign the priority.

Parameter Scope	Scope Description
Global (indicated by a globe icon)	Enter a value for the parameter and apply that value to all devices.
Device Specific (indicated by a host icon)	<p>Use a device-specific value for the parameter.</p> <p>For device-specific parameters, you cannot enter a value in the feature template. You enter the value when you attach a Cisco SD-WAN device to a device template.</p> <p>When you click Device Specific, the Enter Key box opens. This box displays a key, which is a unique string that identifies the parameter in a CSV file that you create. This file is an Excel spreadsheet that contains one column for each key. The header row contains the key names (one key per column), and each row after that corresponds to a device and defines the values of the keys for that device. You upload the CSV file when you attach a Cisco SD-WAN device to a device template.</p> <p>To change the default key, type a new string and move the cursor out of the Enter Key box.</p>
Default	When Default is selected, this field is not enabled.



CHAPTER 5

Configure Network Interfaces

In the Cisco SD-WAN overlay network design, interfaces are associated with VPNs. The interfaces that participate in a VPN are configured and enabled in that VPN. Each interface can be present only in a single VPN.

At a high level, for an interface to be operational, you must configure an IP address for the interface and mark it as operational (**no shutdown**). In practice, you always configure additional parameters for each interface.

You can configure up to 512 interfaces on a Cisco vEdge device. This number includes physical interfaces, loopback interfaces, and subinterfaces.



Note To maximize the efficiency of the load-balancing among Cisco vSmart Controllers, use sequential numbers when assigning system IP addresses to the Cisco vEdge devices in the domain. Example of a sequential numbering schemes is 172.16.1.1, 172.16.1.2, 172.16.1.3, and so on.



Note Ensure that any network interface configured on a device has a unique IP address. If the IP address of the interface conflicts with the system IP address of Cisco vManage instance, it can break the NETCONF session and lead Cisco vManage to read the device as offline.

- [Configure VPN, on page 100](#)
- [Configure Interfaces in the WAN Transport VPN \(VPN 0\), on page 104](#)
- [Extend the WAN Transport VPN, on page 107](#)
- [Configure GRE Interfaces and Advertise Services to Them, on page 110](#)
- [Configure the System Interface, on page 114](#)
- [Configure Control Plane High Availability, on page 115](#)
- [Configure Other Interfaces, on page 115](#)
- [Role-Based Access Control by VPN, on page 117](#)
- [Configure Interface Properties, on page 120](#)
- [Enable DHCP Server using Cisco vManage, on page 122](#)
- [Configuring PPPoE, on page 127](#)
- [Configuring VRRP , on page 133](#)
- [Network Interface Configuration Examples for Cisco vEdge Devices, on page 136](#)
- [Configure VPN Interfaces Using vManage, on page 153](#)

- [Interface CLI Reference](#), on page 212

Configure VPN

VPN

Use the VPN template for all Cisco SD-WAN devices running the Cisco SD-WAN software.

To configure VPNs using Cisco vManage templates, follow this general workflow:

1. Create VPN feature templates to configure VPN parameters. You create a separate VPN feature template for each VPN. For example, create one feature template for VPN 0, a second for VPN 1, and a third for VPN 512.

For Cisco vManage Network Management Systems and Cisco vSmart Controllers, you can configure only VPNs 0 and 512. Create templates for these VPNs only if you want to modify the default settings for the VPN. For Cisco vEdge devices, you can create templates for these two VPNs and for additional VPN feature templates to segment service-side user networks.

- **VPN 0—Transport VPN**, which carries control traffic via the configured WAN transport interfaces. Initially, VPN 0 contains all of a device's interfaces except for the management interface, and all interfaces are disabled.
 - **VPN 512—Management VPN**, which carries out-of-band network management traffic among the Cisco vEdge devices in the overlay network. The interface used for management traffic resides in VPN 512. By default, VPN 512 is configured and enabled on all Cisco vEdge devices except for Cisco vEdge 100. For controller devices, by default, VPN 512 is not configured.
 - **VPNs 1–511, 513–65530—Service VPNs**, for service-side data traffic on Cisco vEdge devices.
2. Create interface feature templates to configure the interfaces in the VPN. See [VPN-Interface-Ethernet](#).


Create a VPN Template



-
- Step 1** In Cisco vManage NMS, choose **Configuration > Templates**.
 - Step 2** In the Device tab, click **Create Template**.
 - Step 3** From the Create Template drop-down, select **From Feature Template**.
 - Step 4** From the **Device Model** drop-down, select the type of device for which you are creating the template.
 - Step 5** To create a template for VPN 0 or VPN 512:
 - a. Click the **Transport & Management VPN** tab located directly beneath the Description field, or scroll to the Transport & Management VPN section.
 - b. From the VPN 0 or VPN 512 drop-down, click **Create Template**. The VPN template form displays. The top of the form contains fields for naming the template, and the bottom contains fields for defining VPN parameters.
 - Step 6** To create a template for VPNs 1 through 511, and 513 through 65530:
 - a. Click the **Service VPN** tab located directly beneath the Description field, or scroll to the Service VPN section.

- b. Click the **Service VPN** drop-down.
- c. From the VPN drop-down, click **Create Template**. The VPN template form displays. The top of the form contains fields for naming the template, and the bottom contains fields for defining VPN parameters.

- Step 7** In the **Template Name** field, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
- Step 8** In the **Template Description** field, enter a description of the template. The description can be up to 2048 characters and can contain only alphanumeric characters.

Changing the Scope for a Parameter Value

When you first open a feature template, for each parameter that has a default value, the scope is set to Default (a ) and the default setting or value is shown. To change the default or to enter a value, click the scope drop-down to the left of the parameter field and select one of the following:

Parameter Name	Description
 Device Specific	<p>Use a device-specific value for the parameter. For device-specific parameters, you cannot enter a value in the feature template. You enter the value when you attach a device to a device template.</p> <p>When you click Device Specific, the Enter Key box opens. This box displays a key, which is a unique string that identifies the parameter in a CSV file that you create. This file is an Excel spreadsheet that contains one column for each key. The header row contains the key names (one key per column), and each row after that corresponds to a device and defines the values of the keys for that device. You upload the CSV file when you attach a device to a device template. For more information, see _Create a Template Variables Spreadsheet</p> <p>To change the default key, type a new string and move the cursor out of the Enter Key box.</p> <p>Examples of device-specific parameters are system IP address, hostname, GPS location, and site ID.</p>
 Global	<p>Enter a value for the parameter, and apply that value to all devices.</p> <p>Examples of parameters that you might apply globally to a group of devices are DNS server, syslog server, and interface MTUs.</p>

Once you have created and named the template, enter the following values. Parameters marked with an asterisk are required.

Configure Basic VPN Parameters

To configure basic VPN parameters, choose the Basic Configuration tab and then configure the following parameters. Parameters marked with an asterisk are required to configure a VPN.

Parameter Name	Description
VPN*	<p>Enter the numeric identifier of the VPN.</p> <p>Range for Cisco vEdge devices: 0 through 65530</p> <p>Values for Cisco vSmart Controller and Cisco vManage devices: 0, 512</p>
Name	Enter a name for the VPN.
Enhance ECMP keying (Cisco vEdge devices only)	<p>Click On to enable the use in the ECMP hash key of Layer 4 source and destination ports, in addition to the combination of the source IP address, destination IP address, protocol, and DSCP field, as the ECMP hash key. ECMP keying is Off by default.</p>

Parameter Name	Description
Enable TCP Optimization Cisco vEdge devices only	Click On to enable TCP optimization for a service-side VPN (a VPN other than VPN 0 and VPN 512). TCP optimization fine-tunes TCP to decrease round-trip latency and improve throughput for TCP traffic.



Note To complete the configuration of the transport VPN on a router, you must configure at least one interface in VPN 0.

To save the feature template, click **Save**.

Configure DNS and Static Hostname Mapping

To configure DNS addresses and static hostname mapping, click the **DNS** tab and configure the following parameters:

Parameter Name	Options	Description
Primary DNS Address	Select either IPv4 or IPv6 , and enter the IP address of the primary DNS server in this VPN.	
New DNS Address	Click New DNS Address and enter the IP address of a secondary DNS server in this VPN. This field appears only if you have specified a primary DNS address.	
	Mark as Optional Row	Check Mark as Optional Row to mark this configuration as device-specific. To include this configuration for a device, enter the requested variable values when you attach a device template to a device, or create a template variables spreadsheet to apply the variables.
	Hostname	Enter the hostname of the DNS server. The name can be up to 128 characters.
	List of IP Addresses	Enter up to eight IP addresses to associate with the hostname. Separate the entries with commas.
To save the DNS server configuration, click Add .		

To save the feature template, click **Save**.

CLI Equivalent

```
vpn vpn-id
  dns ip-address (primary | secondary)
  host hostname ip ip-address
```

Configure Interfaces in the WAN Transport VPN (VPN 0)

This topic describes how to configure the general properties of WAN transport and service-side network interfaces. For information about how to configure specific interface types and properties—including cellular interfaces, DHCP, PPPoE, VRRP, and WLAN interfaces.

VPN 0 is the WAN transport VPN. This VPN handles all control plane traffic, which is carried over OMP sessions, in the overlay network. For a Cisco vEdge device to participate in the overlay network, at least one interface must be configured in VPN 0, and at least one interface must connect to a WAN transport network, such as the Internet or an MPLS or a metro Ethernet network. This WAN transport interface is referred to as a tunnel interface. At a minimum, for this interface, you must configure an IP address, enable the interface, and set it to be a tunnel interface.

To configure a tunnel interface on a Cisco vSmart Controller or a Cisco vManage NMS, you create an interface in VPN 0, assign an IP address or configure the interface to receive an IP address from DHCP, and mark it as a tunnel interface. The IP address can be either an IPv4 or IPv6 address. To enable dual stack, configure both address types. You can optionally associate a color with the tunnel.



Note You can configure IPv6 addresses only on transport interfaces, that is, only in VPN 0.

```
vSmart/vManage(config)# vpn 0
vSmart/vManage(config-vpn-0)# interface interface-name
vSmart/vManage(config-interface)# [ip address prefix / length | ip dhcp-client [dhcp-distance
number]
vSmart/vManage(config-interface)# [ipv6 address prefix / length | ipv6 dhcp-client
[dhcp-distance number] [dhcp-rapid-commit]
vSmart/vManage(config-interface)# no shutdown
vSmart/vManage(config-interface)# tunnel-interface
vSmart/vManage(config-tunnel-interface)# color color
vSmart/vManage(config-tunnel-interface)# [no] allow-service service
```

Tunnel interfaces on Cisco vEdge devices must have an IP address, a color, and an encapsulation type. The IP address can be either an IPv4 or IPv6 address. To enable dual stack, configure both address types.

```
vEdge(config)# vpn 0
vEdge(config-vpn-0)# interface interface-name
vEdge(config-interface)# [ip address prefix / length | ip dhcp-client [dhcp-distance number]
vEdge(config-interface)# [ipv6 address prefix / length | ipv6 dhcp-client [dhcp-distance
number] [dhcp-rapid-commit]
vEdge(config-interface)# no shutdown
vEdge(config-interface)# tunnel-interface
vEdge(config-tunnel-interface)# color color [restrict]
vEdge(config-tunnel-interface)# encapsulation (gre | ipsec)
vEdge(config-tunnel-interface)# [no] allow-service service
```

On Cisco vSmart Controllers and Cisco vSmart Controller NMSs, *interface-name* can be either **eth number** or **loopback number**. Because Cisco vSmart Controllers and Cisco vSmart Controller NMSs participate only in the overlay network's control plane, the VPNs that you can configure on these devices are VPN 0 and VPN 512. Hence, all interfaces are present only on these VPNs.

On Cisco vEdge devices, *interface-name* can be **ge slot/port**, **gre number**, **ipsec number**, **loopback string**, **natpool number**, or **ppp number**.

To enable the interface, include the **no shutdown** command.

For the tunnel interface, you can configure a static IPv4 or IPv6 address, or you can configure the interface to receive its address from a DHCP server. To enable dual stack, configure both an IPv4 and an IPv6 address on the tunnel interface.

Color is a Cisco SD-WAN software construct that identifies the transport tunnel. It can be **3g**, **biz-internet**, **blue**, **bronze**, **custom1**, **custom2**, **custom3**, **default**, **gold**, **green**, **lte**, **metro-ethernet**, **mpls**, **private1** through **private6**, **public-internet**, **red**, and **silver**. The colors **metro-ethernet**, **mpls**, and **private1** through **private6** are referred to as *private colors*, because they use private addresses to connect to the remote side Cisco vEdge device in a private network. You can use these colors in a public network provided that there is no NAT device between the local and remote Cisco vEdge devices.

To limit the remote TLOCs that the local TLOC can establish BFD sessions with, mark the TLOC with the **restrict** option. When a TLOC is marked as restricted, a TLOC on the local router establishes tunnel connections with a remote TLOC only if the remote TLOC has the same color.

On a Cisco vSmart Controller or Cisco vSmart Controller NMS, you can configure one tunnel interface. On a Cisco vEdge device, you can configure up to eight tunnel interfaces.

This means that each Cisco vEdge device can have up to eight TLOCs.

On Cisco vEdge devices, you must configure the tunnel encapsulation. The encapsulation can be either IPsec or GRE. For IPsec encapsulation, the default MTU is 1442 bytes, and for GRE it is 1468 bytes, These values are a function of overhead required for BFD path MTU discovery, which is enabled by default on all TLOCs. (For more information, see *Configuring Control Plane and Data Plane High Availability Parameters* .) You can configure both IPsec and GRE encapsulation by including two **encapsulation** commands under the same **tunnel-interface** command. On the remote Cisco vEdge device, you must configure the same tunnel encapsulation type or types so that the two routers can exchange data traffic. Data transmitted out an IPsec tunnel can be received only by an IPsec tunnel, and data sent on a GRE tunnel can be received only by a GRE tunnel. The Cisco SD-WAN software automatically selects the correct tunnel on the destination Cisco vEdge device.

A tunnel interface allows only DTLS, TLS, and, for Cisco vEdge devices, IPsec traffic to pass through the tunnel. To allow additional traffic to pass without having to create explicit policies or access lists, enable them by including one **allow-service** command for each service. You can also explicitly disallow services by including the **no allow-service** command. Note that services affect only physical interfaces. You can allow or disallow these services on a tunnel interface:

Service	Cisco vEdge device	Cisco vSmart Controller	Cisco vSmart Controller
all (Overrides any commands that allow or disallow individual services)	X	X	X
bgp	X	—	—
dhcp (for DHCPv4 and DHCPv6)	X	—	—
dns	X	—	—
https	—	X	—
icmp	X	X	X
netconf	—	X	—
ntp	X	—	—

Service	Cisco vEdge device	Cisco vSmart Controller	Cisco vSmart Controller
ospf	X	—	—
sshd	X	X	X
stun	X	X	X

The **allow-service stun** command pertains to allowing or disallowing a Cisco vEdge device to generate requests to a generic STUN server so that the device can determine whether it is behind a NAT and, if so, what kind of NAT it is and what the device's public IP address and public port number are. On a Cisco vEdge device that is behind a NAT, you can also have tunnel interface to discover its public IP address and port number from the Cisco vBond Orchestrator.

```
vEdge (config-tunnel-interface) # vbond-as-stun-server
```

With this configuration, the Cisco vEdge device uses the Cisco vBond Orchestrator as a STUN server, so the router can determine its public IP address and public port number. (With this configuration, the router cannot learn the type of NAT that it is behind.) No overlay network control traffic is sent and no keys are exchanged over tunnel interface configured to the the Cisco vBond Orchestrator as a STUN server. However, BFD does come up on the tunnel, and data traffic can be sent on it. Because no control traffic is sent over a tunnel interface that is configured to use the Cisco vBond Orchestrator as a STUN server, you must configure at least one other tunnel interface on the Cisco vEdge device so that it can exchange control traffic with the Cisco vSmart Controller and the Cisco vSmart Controller NMS.

You can log the headers of all packets that are dropped because they do not match a service configured with an **allow-service** command. You can use these logs for security purposes, for example, to monitor the flows that are being directed to a WAN interface and to determine, in the case of a DDoS attack, which IP addresses to block.

```
vEdge (config) # policy implicit-acl-logging
```

When you enable implicit ACL logging, by default, the headers of all dropped packets are logged. It is recommended that you configure a limit to the number of packets logged with the **policy log-frequency** configuration command.

On a Cisco vEdge device, services that you configure on a tunnel interface act as implicit access lists (ACLs). If you apply a localized data policy on a tunnel interface by configuring an ACL with the **policy access-list** command, this ACL is an explicit ACL. For information about how packets matching both implicit and explicit ACLs are handled, see Configuring Localized Data Policy for IPv4 or Configuring Localized Data Policy for IPv6 .

For each transport tunnel on a vEdge router and for each encapsulation type on a single transport tunnel, the Cisco SD-WAN software creates a TLOC, which consists of the router's system IP address, the color, and the encapsulation. The OMP session running on the tunnel sends the TLOC, as a TLOC route, to the Cisco vSmart Controller, which uses it to determine the overlay network topology and to determine the best paths for data traffic across the overlay network.

To display information about interfaces in the WAN transport VPN that are configured with IPv4 addresses, use the **show interface** command. For example:

```
vEdge# show interface vpn 0
```

VPN	INTERFACE	IP ADDRESS	IF ADMIN STATUS	IF OPER STATUS	ENCAP TYPE	PORT TYPE	MTU	HWADDR	SPEED MBPS	DUPLEX	TCP MSS ADJUST	UPTIME	RX PACKETS	TX PACKETS
0	ge0/1	10.0.5.21/24	Up	Up	null	transport	1500	00:0c:29:6c:30:c1	10	full	0	0:04:03:41	260025	260145
0	ge0/2	-	Down	Up	null	service	1500	00:0c:29:6c:30:cb	10	full	0	0:04:03:41	3506	1
0	ge0/3	-	Down	Up	null	service	1500	00:0c:29:6c:30:d5	10	full	0	0:04:03:41	260	1


```

0 ge0/4 - Down Up null service 1500 00:0c:29:6c:30:df 10 full 0 0:04:03:41 260 1
0 ge0/5 - Down Up null service 1500 00:0c:29:6c:30:e9 10 full 0 0:04:03:41 260 1
0 ge0/6 10.0.7.21/24 Up Up null service 1500 00:0c:29:6c:30:f3 10 full 0 0:04:03:41 265 2
0 ge0/7 10.0.100.21/24 Up Up null service 1500 00:0c:29:6c:30:fd 10 full 0 0:04:03:41 278 2
0 system 172.16.255.21/32 Up Up null loopback 1500 00:00:00:00:00:00 10 full 0 0:04:03:37 0 0

```

To display information for interfaces configured with IPv6 addresses, use the **show ipv6 interface** command. For example:

```
vEdge# show ipv6 interface vpn 0
```

VPN	INTERFACE	AF	TYPE	IPV6 ADDRESS	IF		ENCAP	PORT	TYPE	MTU	HWADDR	SPEED		MSS		RX	TX	LINK	LOCAL	ADDRESS
					ADMIN	OPER						MBPS	DUPLEX	ADJUST	UPTIME					
0	ge0/1	ipv6		2001::a00:1a0b/120	Up	Up	null	service	1500	00:0c:29:ab:b7:62	1000	full	1420	0:01:30:00	2	6	fe80::20c:29ff:feab:b762/64			
0	ge0/2	ipv6		2001::a00:50b/120	Up	Up	null	service	1500	00:0c:29:ab:b7:6c	1000	full	1420	0:01:30:00	21	5	fe80::20c:29ff:feab:b76c/64			
0	ge0/3	ipv6		fd00:1234::/16	Up	Up	null	service	1500	00:0c:29:ab:b7:76	1000	full	1420	0:01:08:33	0	8	fe80::20c:29ff:feab:b776/64			
0	ge0/4	ipv6		-	Up	Up	null	service	1500	00:0c:29:ab:b7:80	1000	full	1420	0:01:30:00	18	5	fe80::20c:29ff:feab:b780/64			
0	ge0/5	ipv6		-	Down	Up	null	service	1500	00:0c:29:ab:b7:8a	1000	full	1420	0:01:44:19	1	1	fe80::20c:29ff:feab:b78a/64			
0	ge0/6	ipv6		-	Down	Up	null	service	1500	00:0c:29:ab:b7:94	1000	full	1420	0:01:44:19	0	1	fe80::20c:29ff:feab:b794/64			
0	ge0/7	ipv6		-	Up	Up	null	service	1500	00:0c:29:ab:b7:9e	1000	full	1420	0:01:43:02	55	5	fe80::20c:29ff:feab:b79e/64			
0	system	ipv6		-	Up	Up	null	loopback	1500	00:00:00:00:00:00	10	full	1420	0:01:29:31	0	0	-			
0	loopback1	ipv6		2001::a00:6501/128	Up	Up	null	transport	1500	00:00:00:00:00:00	10	full	1420	0:03:49:09	0	0	-			
0	loopback2	ipv6		2001::a00:6502/128	Up	Up	null	transport	1500	00:00:00:00:00:00	10	full	1420	0:03:49:05	0	0	-			
0	loopback3	ipv6		2001::a00:6503/128	Up	Up	null	transport	1500	00:00:00:00:00:00	10	full	1420	0:03:49:01	0	0	-			
0	loopback4	ipv6		2001::a00:6504/128	Up	Up	null	transport	1500	00:00:00:00:00:00	10	full	1420	0:03:48:54	0	0	-			

In the command output, a port type of "transport" indicates that the interface is configured as a tunnel interface, and a port type of "service" indicates that the interface is not configured as a tunnel interface and can be used for data plane traffic. The port type for the system IP address interface is "loopback".


Configure Other WAN Interface Properties

You can modify the distribution of data traffic across transport tunnels by applying a data policy in which the action sets TLOC attributes (IP address, color, and encapsulation) to apply to matching data packets. For more information, see the Configuring Centralized Data Policy .

Extend the WAN Transport VPN

When two Cisco vEdge devices are collocated at a physical site that has only one WAN circuit, you can configure the Cisco vEdge device that is not connected to the circuit to be able to establish WAN transport tunnels through the other router's TLOCs. In this way, you extend the WAN transport VPN so that both routers can establish tunnel interfaces, and hence can establish independent TLOCs, in the overlay network. (Note that you can configure the two routers themselves with different site identifiers.)

The following figure illustrates a site with two Cisco vEdge devices. Cisco vEdge device-1 terminates one WAN circuit from the Internet and the second Cisco vEdge device-2 terminates the private MPLS network. Each router has one TLOC. You can configure Cisco vEdge device-2 to extend its WAN transport VPN to Cisco vEdge device1 so that Cisco vEdge device-1 can participate independently in the overlay network. You can also make a similar configuration for vEdge1 so that the WAN transport can be extended from Cisco vEdge device1 to Cisco vEdge device2.



When you extend the WAN transport VPN, no BFD sessions are established between the two collocated vEdge routers.

You cannot configure TLOC extensions on cellular (LTE) interfaces.

To extend the WAN transport VPN, you configure the interface between the two routers:

- For the router that is not connected to the circuit, you configure a standard tunnel interface in VPN 0.
- For the router that is physically connected to the WAN or private transport, you associate the physical interface that connects to the circuit, configuring this in VPN 0 but not in a tunnel interface.

To configure the non-connected router (Cisco vEdge device-1 in the figure above), create a tunnel interface in VPN 0 on the physical interface to the connected router.

```
vEdge-1(config-vpn-0)# interface ge slot/ port
vEdge-1(config-interface)# ip address prefix / length
vEdge-1(config-interface)# no shutdown
vEdge-1(config-interface)# mtu number
vEdge-1(config-interface)# tunnel-interface
vEdge-1(config-tunnel-interface)# color color
```

For the router connected to the WAN or private transport (Cisco vEdge device-2 in the figure above), configure the interface that connects to the non-connected router, again in VPN 0:

```
vEdge-2(config-vpn-0)# interface ge slot/port
vEdge-2(config-interface)# ip address prefix / length
vEdge-2(config-interface)# tloc-extension ge slot / port
vEdge-2(config-interface)# no shutdown
vEdge-2(config-interface)# mtu number
```

The physical interface in the **interface** command is the one that connects to the other router.

The **tloc-extension** command creates the binding between the non-connected router and the WAN or private network. In this command, you specify the physical interface that connects to the WAN or private network circuit.

If the circuit connects to a public network:

- Configure a NAT on the public-network-facing interface on the Cisco vEdge device. The NAT configuration is required because the two Cisco vEdge devices are sharing the same transport tunnel.
- Configure a static route on the non-connected router to the TLOC-extended interface on the router connected to the public network.

If the circuit connects to a private network, such as an MPLS network:

- Enable routing on the non-connected router so that the interface on the non-connected router is advertised into the private network.

- Depending on the routing protocol you are using, enable either OSPF or BGP service on the non-connected router interface so that routing between the non-connected and the connected routers comes up. To do this, use the **allow-service** command.

You cannot extend a TLOC configured on a loopback interface, that is, when you use a loopback interface to connect to the public or private network. You can extend a TLOC only on a physical interface.

If one of the routers is connected to two WAN transports (such as the Internet and an MPLS network), create subinterfaces between the two routers, creating the tunnel on the subinterface. The subinterfaces on the two routers must be in the same subnet. Because you are using a subinterface, the interface's MTU must be at least 4 bytes less than the physical MTU.

Here is a sample configuration that corresponds to the figure shown above. Because the router Cisco vEdge device-2 connects to two transports, we create subinterfaces between the Cisco vEdge device-1 and Cisco vEdge device-2 routers. One subinterface binds to the Internet circuit, and the second one binds to the MPLS connection.

```
vEdge-1# show running-config vpn 0
interface ge0/2.101
  ip address 101.1.19.15/24
  mtu 1496
  tunnel-interface
    color lte
  ...
!
no shutdown
!
interface ge0/2.102
  ip address 102.1.19.15/24
  mtu 1496
  tunnel-interface
    color mpls
  ...
!
no shutdown
!
ip route 0.0.0.0/0 101.1.19.16
vEdge-2# show running-config vpn 0
interface ge0/0
  ip address 172.16.255.2
  tunnel-interface
    color lte
  ...
!
no shutdown
!
interface ge0/3
  ip address 172.16.255.16
  tunnel-interface
    color mpls
  ...
!
no shutdown
!
interface ge0/2.101
  ip address 101.1.19.16/24
  mtu 1496
  tloc-extension ge0/0
  no shutdown
!
interface ge0/2.102
  ip address 102.1.19.16/24
```

Configure GRE Interfaces and Advertise Services to Them

```

mtu 1496
tloc-extension ge0/3
no shutdown
!

```

For this example configuration, Cisco vEdge device-1 establishes two control connections to each Cisco vSmart Controller in the overlay network—one connection for the LTE tunnel and the second for the MPLS tunnel. These control connections are separate and independent from those established on Cisco vEdge device-2. The following output shows the control connections on vEdge-1 in a network with two Cisco vSmart Controllers:

```
vEdge-1# show control connections
```

PEER TYPE	PEER PROTOCOL	PEER SYSTEM IP	SITE ID	DOMAIN ID	PEER PRIVATE IP	PEER PRIVATE PORT	PEER PUBLIC IP	PEER PUBLIC PORT	LOCAL COLOR	STATE	UPTIME	CONTROLLER GROUP NAME
vsmart	dtls	172.16.255.19	100	1	10.0.5.19	12346	10.0.5.19	12346	lte	up	0:00:18:43	default
vsmart	dtls	172.16.255.19	100	1	10.0.5.19	12346	10.0.5.19	12346	mpls	up	0:00:18:32	default
vsmart	dtls	172.16.255.20	200	1	10.0.12.20	12346	10.0.12.20	12346	lte	up	0:00:18:38	default
vsmart	dtls	172.16.255.20	200	1	10.0.12.20	12346	10.0.12.20	12346	mpls	up	0:00:18:27	default

You can verify that the two Cisco vEdge devices have established no BFD sessions between them. On Cisco vEdge device-1, we see no BFD sessions to Cisco vEdge device-2 (system IP address 172.16.255.16):

```
vEdge-1# show bfd sessions
```

SYSTEM IP	SITE ID	STATE	SOURCE TLOC COLOR	REMOTE TLOC COLOR	SOURCE IP	DST PUBLIC IP	DST PUBLIC PORT	ENCAP	DETECT MULTIPLIER	TX INTERVAL(msec)	UPTIME	TRANSI-TIONS
172.16.255.11	100	up	lte	lte	101.1.19.15	10.0.101.1	12346	ipsec	20	1000	0:00:20:26	0
172.16.255.11	100	up	lte	3g	101.1.19.15	10.0.101.2	12346	ipsec	20	1000	0:00:20:26	0
172.16.255.11	100	up	lte	gold	101.1.19.15	10.0.101.3	12346	ipsec	20	1000	0:00:20:26	0
172.16.255.11	100	up	lte	red	101.1.19.15	10.0.101.4	12346	ipsec	20	1000	0:00:20:26	0
172.16.255.11	100	up	mpls	lte	102.1.19.15	10.0.101.1	12346	ipsec	20	1000	0:00:20:26	0
172.16.255.11	100	up	mpls	3g	102.1.19.15	10.0.101.2	12346	ipsec	20	1000	0:00:20:26	0
172.16.255.11	100	up	mpls	gold	102.1.19.15	10.0.101.3	12346	ipsec	20	1000	0:00:20:26	0
172.16.255.11	100	up	mpls	red	102.1.19.15	10.0.101.4	12346	ipsec	20	1000	0:00:20:26	0
172.16.255.14	400	up	lte	lte	101.1.19.15	10.1.14.14	12360	ipsec	20	1000	0:00:20:26	0
172.16.255.14	400	up	mpls	lte	102.1.19.15	10.1.14.14	12360	ipsec	20	1000	0:00:20:26	0
172.16.255.14	400	up	lte	lte	101.1.19.15	10.0.111.1	12346	ipsec	20	1000	0:00:20:26	0
172.16.255.21	100	up	lte	3g	101.1.19.15	10.0.111.2	12346	ipsec	20	1000	0:00:20:26	0
172.16.255.21	100	up	mpls	lte	102.1.19.15	10.0.111.1	12346	ipsec	20	1000	0:00:20:26	0
172.16.255.21	100	up	mpls	3g	102.1.19.15	10.0.111.2	12346	ipsec	20	1000	0:00:20:26	0

Configure GRE Interfaces and Advertise Services to Them

When a service, such as a firewall, is available on a device that supports only GRE tunnels, you can configure a GRE tunnel on the vEdge router to connect to the remote device. You then advertise that the service is available via a GRE tunnel, and you direct the appropriate traffic to the tunnel either by creating centralized data policy or by configuring GRE-specific static routes.

You create a GRE tunnel by configuring a GRE interface. GRE interfaces are logical interfaces, and you configure them just like any other physical interface. Because a GRE interface is a logical interface, you must bind it to a physical interface, as described below.

To configure a GRE tunnel interface to a remote device that is reachable through a transport network, configure the tunnel in VPN 0:

```

vEdge (config) # vpn 0 interface gre number
vEdge (config-interface-gre) # (tunnel-source ip-address | tunnel-source-interface
interface-name)
vEdge (config-interface-gre) # tunnel-destination ip-address
vEdge (config-interface-gre) # no shutdown

```

The GRE interface has a name in the format **gre number**, where *number* can be from 1 through 255.

To configure the source of the GRE tunnel on the local device, you can specify either the IP address of the physical interface (in the **tunnel-source** command) or the name of the physical interface (in the **tunnel-source-interface** command). Ensure that the physical interface is configured in the same VPN in which the GRE interface is located.

To configure the destination of the GRE tunnel, specify the IP address of the remote device in the **tunnel-destination** command.

The combination of a source address (or source interface name) and a destination address defines a single GRE tunnel. Only one GRE tunnel can exist that uses a specific source address (or interface name) and destination address pair.

You can optionally configure an IP address for the GRE tunnel itself:

```
vEdge(config-interface-gre)# ip address ip-address
```

Because GRE tunnels are stateless, the only way for the local router to determine whether the remote end of the tunnel is up, is to periodically send keepalive messages over the tunnel. The keepalive packets are looped back to the sender, and receipt of these packets by the local router indicates that the remote GRE device is up. By default, the GRE interface sends keepalive packets every 10 seconds, and if it receives no response, retries 3 times before declaring the remote device to be down. You can modify these default values with the **keepalive** command:

```
vEdge(config-interface-gre)# keepalive seconds retries
```

The keepalive interval can be from 0 through 65535 seconds, and the number of retries can be from 0 through 255. If you configure an IP address for the GRE interface, that IP address generates the keepalive messages.

If the vEdge router sits behind a NAT and you have configured GRE encapsulation, you must disable keepalives, with a **keepalive 0 0** command. (Note that you cannot disable keepalives by issuing a **no keepalive** command. This command returns the keepalive to its default settings of sending a keepalive packet every 10 seconds and retrying 3 times before declaring the remote device down.)

For GRE interfaces, you can configure only the following additional interface properties:

```
vEdge(config-interface-gre)# access-list acl-name
vEdge(config-interface-gre)# block-non-source-ip
vEdge(config-interface-gre)# clear-dont-fragment
vEdge(config-interface-gre)# description text
vEdge(config-interface-gre)# mtu bytes
vEdge(config-interface-gre)# policer policer-name
vEdge(config-interface-gre)# rewrite-rule rule-name
vEdge(config-interface-gre)# tcp-mss-adjust
```

GRE interfaces do not support cFlowd traffic monitoring.

You can configure one or two GRE interfaces per service. When you configure two, the first interface is the primary GRE tunnel, and the second is the backup tunnel. All packets are sent only to the primary tunnel. If that tunnel fails, all packets are then sent to the secondary tunnel. If the primary tunnel comes back up, all traffic is moved back to the primary GRE tunnel.

You direct data traffic from the service VPN to the GRE tunnel in one of two ways: either with a GRE-specific static route or with a centralized data policy.

To create a GRE-specific static route in the service VPN (a VPN other than VPN 0 or VPN 512), use the **ip gre-route** command:

```
vEdge(config-vpn)# ip gre-route prefix vpn 0 interface gre number [gre number2]
```

This GRE-specific static route directs traffic from the specified prefix to the primary GRE interface, and optionally to the secondary GRE interface, in VPN 0. The OMP administrative distance of a GRE-specific static route is 5, and the admin distance for a regular static route (configured with the **ip route** command) is 1. For more information, see *Unicast Overlay Routing Overview*.

To direct the data traffic to the GRE tunnel using a centralized data policy is a two-part process: you advertise the service in the service VPN, and then you create a centralized data policy on the Cisco vSmart Controller to forward matching traffic to that service.

To advertise the service, include the **service** command in the service VPN (a VPN other than VPN 0 or VPN 512):

```
vEdge(config-vpn)# service service-name interface gre number [gre number2]
```

The service name can be **FW**, **IDP**, **IDS**, or **TE**, or a custom service name **netsvc1** through **netsvc4**. For more information on service-names, refer to Service Chaining. The interface is the GRE interface in VPN 0 that is used to reach the service. If you have configured a primary and a backup GRE tunnel, list the two GRE interfaces (**gre number1 gre number2**) in the **service** command. Once you have configured a service as reachable at the GRE interface, you cannot delete the GRE interface from the configuration. To delete the GRE interface, you must first delete the service. You can, however, reconfigure the service itself, by modifying the **service** command.

Then, create a data policy on the Cisco vSmart Controller that applies to the service VPN. In the action portion of the data policy, you must explicitly configure the policy to service the packets destined for the GRE tunnel. To do this, include the **local** option in the **set service** command:

```
vSmart(config-policy-data-policy-vpn-list-vpn-sequence)# action accept
vSmart(config-action-accept)# set service service-name local
```


If the GRE tunnel used to reach the service is down, packet routing falls back to using standard routing. To drop packets when a GRE tunnel to the service is unreachable, add the **restrict** option:

```
vSmart(config-policy-data-policy-vpn-list-vpn-sequence)# action accept
vSmart(config-action-accept)# set service service-name local restrict
```

To monitor GRE tunnels and their traffic, use the following commands:

- **show interface** —List data traffic transmitted and received on GRE tunnels.
- **show tunnel gre-keepalives** —List GRE keepalive traffic transmitted and received on GRE tunnels.
- **show tunnel statistics** —List both data and keepalive traffic transmitted and received on GRE tunnels.

The following figure illustrates an example of configuring a GRE tunnel in VPN 0, to allow traffic to be redirected to a service that is not located at the same site as the vEdge router. In this example, local traffic is directed to the GRE tunnel using a centralized data policy, which is configured on the Cisco vSmart Controller.



The configuration looks like this:

```
vEdge# show running-config vpn 0
vpn 0
  interface gre1
    ip address 172.16.111.11/24
    keepalive 60 10
    tunnel-source 172.16.255.11
    tunnel-destination 10.1.2.27
```

```

        no shutdown
    !
!
vEdge# show running-config vpn 1 service
vpn 1
    service FW interface gre1

vSmart# show running-config policy
policy
    lists
        prefix-list for-firewall
            ip-prefix 58.0.1.0/24
        site-list my-site
        site-id 100
        vpn-list for-vpn-1
        vpn 1
    data-policy to-gre-tunnel
        vpn-list for-vpn-1
        sequence 10
        match
            source-data-prefix-list for-firewall
        action accept
        set service FW local
    apply-policy site-list my-site
    data-policy to-gre-tunnel from-service

```

Here is an example of the same configuring using a GRE-specific static route to direct data traffic from VPN 1 into the GRE tunnels:

```

vEdge# show running-config
vpn 0
    interface gre1
        ip address 172.16.111.11/24
        keepalive 60 10
        tunnel-source 172.16.255.11
        tunnel-destination 10.1.2.27
        no shutdown
    !
!
vpn 1
    ip gre-route 58.0.1.0/24 vpn 0 interface gre1

```

The **show interface** command displays the GRE interface in VPN 0:

```
vEdge# show interface vpn 0
```

VPN	INTERFACE	IP ADDRESS	IF ADMIN STATUS	IF OPER STATUS	ENCAP TYPE	PORT TYPE	MTU	HWADDR	SPEED MBPS	DUPLEX	TCP MSS ADJUST	UPTIME	RX PACKETS	TX PACKETS
0	gre1	172.16.111.11/24	Up	Down	null	service	1500	0a:00:05:0b:00:00	-	-	1420	-	0	0
0	ge0/1	10.0.26.11/24	Up	Up	null	service	1500	00:0c:29:ab:b7:62	10	full	1420	0:03:35:14	89	5
0	ge0/2	10.0.5.11/24	Up	Up	null	transport	1500	00:0c:29:ab:b7:6c	10	full	1420	0:03:35:14	9353	18563
0	ge0/3	-	Down	Up	null	service	1500	00:0c:29:ab:b7:76	10	full	1420	0:03:57:52	99	0
0	ge0/4	10.0.7.11/24	Up	Up	null	service	1500	00:0c:29:ab:b7:80	10	full	1420	0:03:35:14	89	5
0	ge0/5	-	Down	Up	null	service	1500	00:0c:29:ab:b7:8a	10	full	1420	0:03:57:52	97	0
0	ge0/6	-	Down	Up	null	service	1500	00:0c:29:ab:b7:94	10	full	1420	0:03:57:52	85	0
0	ge0/7	10.0.100.11/24	Up	Up	null	service	1500	00:0c:29:ab:b7:9e	10	full	1420	0:03:56:30	3146	2402
0	system	172.16.255.11/32	Up	Up	null	loopback	1500	00:00:00:00:00:00	10	full	1420	0:03:34:15	0	0

You can also view the GRE tunnel information:

```
vEdge# show tunnel gre-keepalives
```

VPN	IF NAME	SOURCE IP	DEST IP	ADMIN STATE	OPER STATE	KA ENABLED	REMOTE TX PACKETS	REMOTE RX PACKETS	TX PACKETS	RX PACKETS	TX ERRORS	RX ERRORS	TRANSITIONS
0	gre1	10.0.5.11	10.1.2.27	up	down	true	0	0	442	0	0	0	0

```
vEdge# show tunnel statistics
tunnel statistics gre 10.0.5.11 10.1.2.27 0 0
tunnel-mtu      1460
tx_pkts        451
tx_octets      54120
rx_pkts        0
rx_octets      0
tcp-mss-adjust 1380
```

Configure the System Interface

For each Cisco vEdge device, you configure a system interface with the **system system-ip** command. The system interface's IP address is a persistent address that identifies the Cisco vEdge device. It is similar to a router ID on a regular router, which is the address used to identify the router from which packets originated.

```
vEdge (config)# system system-ip ipv4-address
```

Specify the system IP address as an IPv4 address in decimal four-part dotted notation. Specify just the address; the prefix length (/32) is implicit.

The system IP address can be any IPv4 address except for 0.0.0.0/8, 127.0.0.0/8, and 224.0.0.0/4, and 240.0.0.0/4 and later. Each device in the overlay network must have a unique system IP address. You cannot use this same address for another interface in VPN 0.

The system interface is placed in VPN 0, as a loopback interface named **system**. Note that this is not the same as a loopback address that you configure for an interface.

To display information about the system interface, use the **show interface** command. For example:

```
vEdge# show running-config system system-ip
system
 system-ip 172.16.255.11
!
```

```
vEdge# show interface vpn 0
```

VPN	INTERFACE	IP ADDRESS	IF ADMIN STATUS	IF OPER STATUS	ENCAP TYPE	PORT TYPE	MTU	HWADDR	SPEED MBPS	DUPLEX	TCP MSS ADJUST	UPTIME	RX PACKETS	TX PACKETS
0	ge0/1	10.0.26.11/24	Up	Up	null	service	1500	00:0c:29:ab:b7:62	1000	full	1420	0:10:32:16	1606	8
0	ge0/2	10.0.5.11/24	Up	Up	null	transport	1500	00:0c:29:ab:b7:6c	1000	full	1420	0:10:32:16	307113	303457
0	ge0/3	-	Down	Up	null	service	1500	00:0c:29:ab:b7:76	1000	full	1420	0:10:47:49	1608	0
0	ge0/4	10.0.7.11/24	Up	Up	null	service	1500	00:0c:29:ab:b7:80	1000	full	1420	0:10:32:16	1612	8
0	ge0/5	-	Down	Up	null	service	1500	00:0c:29:ab:b7:8a	1000	full	1420	0:10:47:49	1621	0
0	ge0/6	-	Down	Up	null	service	1500	00:0c:29:ab:b7:94	1000	full	1420	0:10:47:49	1600	0
0	ge0/7	10.0.100.11/24	Up	Up	null	service	1500	00:0c:29:ab:b7:9e	1000	full	1420	0:10:47:31	3128	1165
0	system	172.16.255.11/32	Up	Up	null	loopback	1500	00:00:00:00:00:00	10	full	1420	0:10:31:58	0	0

The system IP address is used as one of the attributes of the OMP TLOC. Each TLOC is uniquely identified by a 3-tuple comprising the system IP address, a color, and an encapsulation. To display TLOC information, use the **show omp tlocs** command.

For device management purposes, it is recommended as a best practice that you also configure the same system IP address on a loopback interface that is located in a service-side VPN that is an appropriate VPN for management purposes. You use a loopback interface because it is always reachable when the router is operational and when the overlay network is up. If you were to configure the system IP address on a physical interface, both the router and the interface would have to be up for the router to be reachable. You use a service-side VPN because it is reachable from the data center. Service-side VPNs are VPNs other than VPN 0 (the WAN transport VPN) and VPN 512 (the management VPN), and they are used to route data traffic.

Here is an example of configuring the system IP address on a loopback interface in VPN 1:

```
vEdge# config
Entering configuration mode terminal
vEdge (config)# vpn 1
vEdge (config-vpn-1)# interface loopback0 ip address 172.16.255.11/32
vEdge (config-vpn-1)# no shutdown
```



```
vEdge(config-interface-loopback0)# commit and-quit
Commit complete.
vEdge# show interface
```

VPN	INTERFACE	IP ADDRESS	IF ADMIN STATUS	IF OPER STATUS	ENCAP TYPE	PORT TYPE	MTU	HWADDR	SPEED MBPS	DUPLEX	TCP MSS ADJUST	UPTIME	RX PACKETS	TX PACKETS
0	ge0/1	10.0.26.11/24	Up	Up	null	service	1500	00:0c:29:ab:b7:62	1000	full	1420	0:10:27:33	1597	8
0	ge0/2	10.0.5.11/24	Up	Up	null	transport	1500	00:0c:29:ab:b7:6c	1000	full	1420	0:10:27:33	304819	301173
0	ge0/3	-	Down	Up	null	service	1500	00:0c:29:ab:b7:76	1000	full	1420	0:10:43:07	1599	0
0	ge0/4	10.0.7.11/24	Up	Up	null	service	1500	00:0c:29:ab:b7:80	1000	full	1420	0:10:27:33	1603	8
0	ge0/5	-	Down	Up	null	service	1500	00:0c:29:ab:b7:8a	1000	full	1420	0:10:43:07	1612	0
0	ge0/6	-	Down	Up	null	service	1500	00:0c:29:ab:b7:94	1000	full	1420	0:10:43:07	1591	0
0	ge0/7	10.0.100.11/24	Up	Up	null	service	1500	00:0c:29:ab:b7:9e	1000	full	1420	0:10:42:48	3118	1164
0	system	172.16.255.11/32	Up	Up	null	loopback	1500	00:00:00:00:00:00	10	full	1420	0:10:27:15	0	0
1	ge0/0	10.2.2.11/24	Up	Up	null	service	1500	00:0c:29:ab:b7:58	1000	full	1420	0:10:27:30	5734	4204
1	loopback0	172.16.255.11/32	Up	Up	null	service	1500	00:00:00:00:00:00	10	full	1420	0:00:00:28	0	0
512	eth0	10.0.1.11/24	Up	Up	null	service	1500	00:50:56:00:01:0b	1000	full	0	0:10:43:03	20801	14368

Configure Control Plane High Availability

A highly available Cisco SD-WAN network contains two or more Cisco vSmart Controllers in each domain. A Cisco SD-WAN domain can have up to eight Cisco vSmart Controllers, and each Cisco vEdge device, by default, connects to two of them. You change this value on a per-tunnel basis:

```
vEdge(config-tunnel-interface)# max-controllers number
```

When the number of Cisco vSmart Controllers in a domain is greater than the maximum number of controllers that a domain's Cisco vEdge devices are allowed to connect to, the Cisco SD-WAN software load-balances the connections among the available Cisco vSmart Controllers.

Configure Other Interfaces

Configure Interfaces in the Management (VPN 512)

On all Cisco SD-WAN devices, VPN 512 is used for out-of-band management, by default as part of the factory-default configuration. On Cisco vEdge devices the interface type for management interfaces is **mgmt**, and the initial address for the interface is 192.168.1.1.

```
vEdge# show running-config vpn 512
vpn 512
  interface mgmt0
    ip dhcp-client
    no shutdown
  !
!
```

To display information about the configured management interfaces, use the **show interface** command. For example:

```
vEdge# show interface vpn 512
```

VPN	INTERFACE	IP ADDRESS	IF ADMIN STATUS	IF OPER STATUS	ENCAP TYPE	PORT TYPE	MTU	HWADDR	SPEED MBPS	DUPLEX	TCP MSS ADJUST	UPTIME	RX PACKETS	TX PACKETS
512	mgmt0	192.168.1.1/24	Up	Up	null	service	1500	00:50:56:00:01:1f	1000	full	0	0:04:08:01	1131	608



Note VPN 512 is not advertised in the overlay. It is local to the device. If you need a management VPN that is reachable through the overlay, create a VPN with a number other than 512.

Configure Service-Side Interfaces for Carrying Data Traffic

On Cisco vEdge devices, the VPNs other than 0 and 512 are service-side VPNs, and the interfaces in these VPNs connect the router to service-side LANs and WLANs. These interfaces are the interfaces that carry data traffic between vEdge routers and sites across the overlay network. At a minimum, for these interfaces, you must configure an IPv4 address, and you must enable the interface:

```
vEdge(config)# vpn vpn-id
vEdge(config-vpn)# interface ge slot / port
vEdge(config-interface)# ip address prefix/length
vEdge(config-interface)# no shutdown
```

For service-side interfaces, you can configure up to four secondary IP addresses.

```
vEdge(config)# vpn vpn-id
vEdge(config-vpn)# interface ge slot/port
vEdge(config-interface)# ip secondary-address ipv4-address
```

To display information about the configured data traffic interfaces, use the **show interface** command.

```
vEdge# show interface vpn 1
```

VPN	INTERFACE	IP ADDRESS	IF ADMIN STATUS	IF OPER STATUS	ENCAP TYPE	PORT TYPE	MTU	HWADDR	SPEED MBPS	DUPLEX	TCP MSS ADJUST	UPTIME	RX PACKETS	TX PACKETS
1	ge0/1	10.192.1.1/28	Up	Up	null	service	1500	00:0c:bd:05:f0:84	100	full	0	1:05:44:07	399	331
1	loopback1	10.255.1.1/32	Up	Up	null	service	1500	00:00:00:00:00:00	10	full	0	1:05:44:07	0	0

For some protocols, you specify an interface as part of the protocol's configuration. In these cases, the interface used by the protocol must be the same as one of the interfaces configured in the VPN. As example is OSPF, where you place interfaces in OSPF areas. In this example, the interface **ge0/0** is configured in VPN 1, and this interface is configured to be in the OSPF backbone area:

```
vEdge# show running-config vpn 1
vpn 1
router
  ospf
    router-id 172.16.255.21
    timers spf 200 1000 10000
    redistribute static
    redistribute omp
    area 0
      interface ge0/0
        exit
      exit
    !
  !
interface ge0/0
  ip address 10.2.3.21/24
  no shutdown
!
!
```

Configure Loopback Interfaces

Use the interface name format **loopback string**, where *string* can be any alphanumeric value and can include underscores (_) and hyphens (-). The total interface name, including the string "loopback", can be a maximum of 16 characters long. (Note that because of the flexibility of interface naming in the CLI, the interfaces **lo0** and **loopback0** are parsed as different strings and as such are not interchangeable. For the CLI to recognize as interface as a loopback interface, its name must start with the full string **loopback**.)

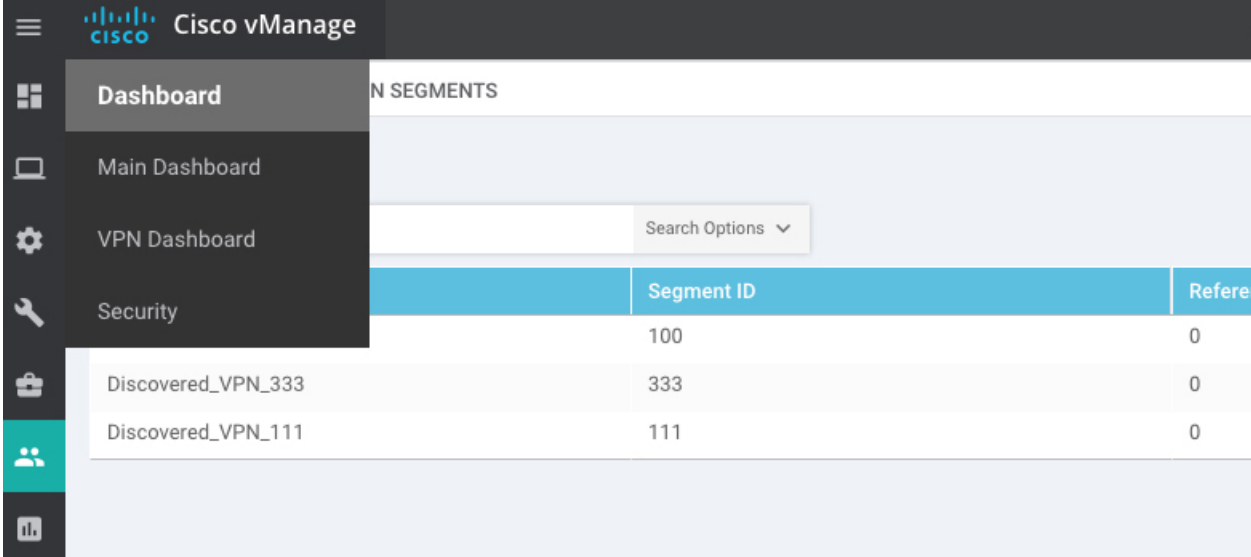
One special use of loopback interfaces is to configure data traffic exchange across private WANs, such as MPLS or metro Ethernet networks. To allow a router that is behind a private network to communicate directly

over the private WAN with other edge routers, you direct data traffic to a loopback interface that is configured as a tunnel interface rather than to an actual physical WAN interface.

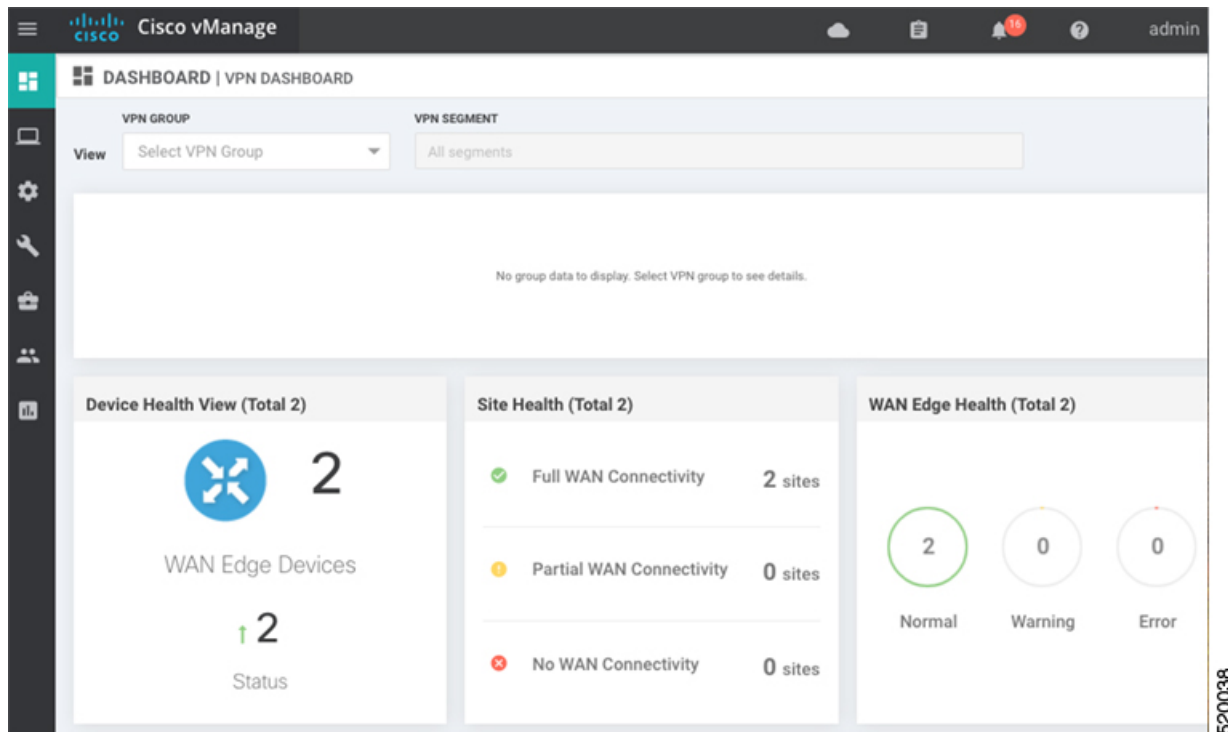
Role-Based Access Control by VPN

VPN Dashboard Overview

Users configured with VPN group can access only the VPN Dashboard, and it is read-only access. User with Admin access can create the VPN groups and has access to both Admin Dashboard and VPN Dashboard(s). Admin user can view these dashboards in the left panel as shown in the following figures:



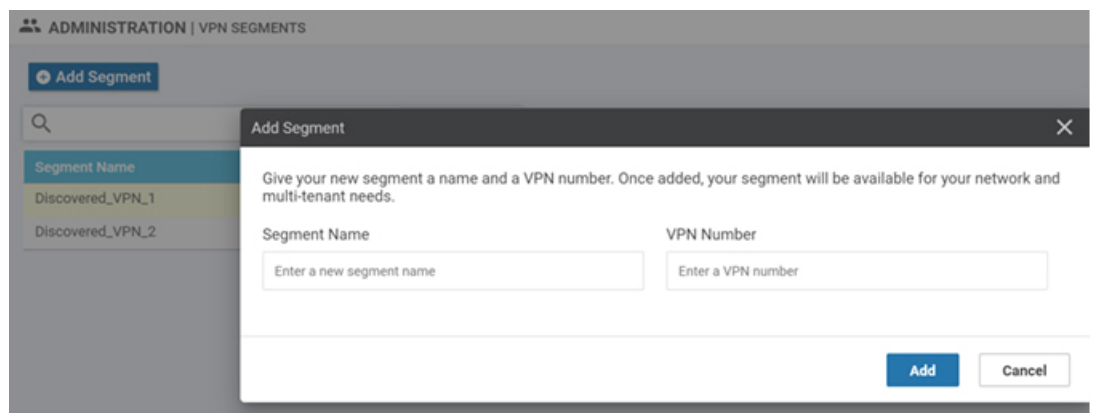
Segment ID	Reference
100	0
Discovered_VPN_333	0
Discovered_VPN_111	0



Configure and Manage VPN Segments

To configure VPN Segments:

1. Navigate to **Administration > VPN Segments** in Cisco vManage. The following web page displays with the list of segments that are configured.
2. To edit or delete an existing segment, click the **Edit or Delete** in the More Info (...) column on the right side.
3. To add new segment, click **Add Segment**. Add Segment window appears.



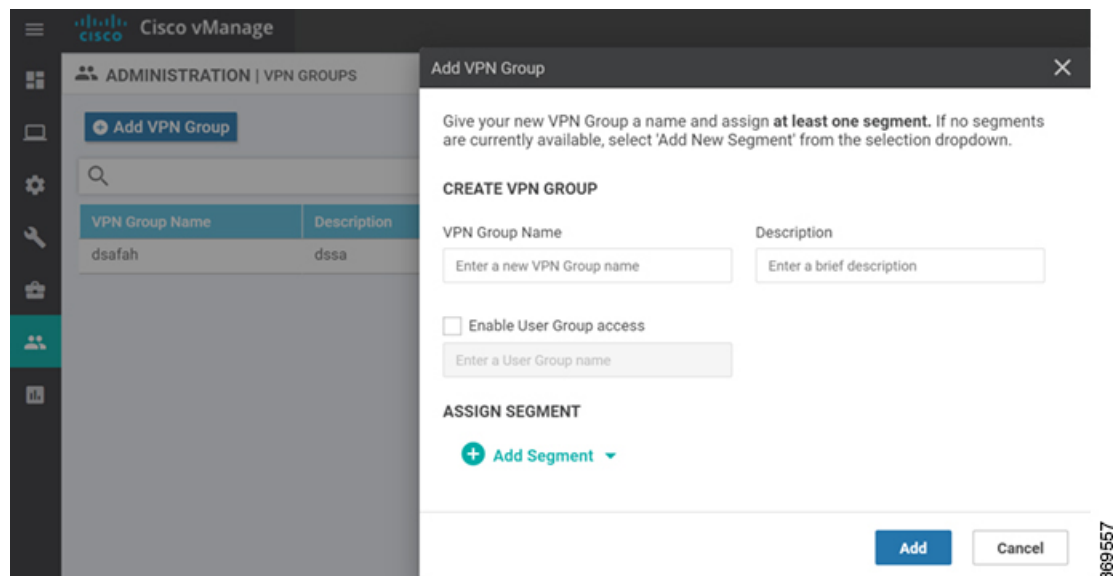
4. Enter the name of the segment in the **Segment Name** field.
5. Enter the number of VPNs you want to configure in VPN Number field.

6. Click **Add** to add a new segment.

Configure and Manage VPN Groups

To configure VPN Groups:

1. Navigate to **Administration > VPN Groups** in Cisco vManage. The following web page displays with the list of segments that are configured.
2. To edit or delete an VPN group, click the **Edit or Delete** in the More Info (...) column on the right side.
3. To view the existing VPN in the dashboard, click on **View Dashboard** in the More Info column. The VPN Dashboard displays the device details of the VPN device configured.
4. To add new VPN group, click **Add Group**. Add VPN Group window appears.



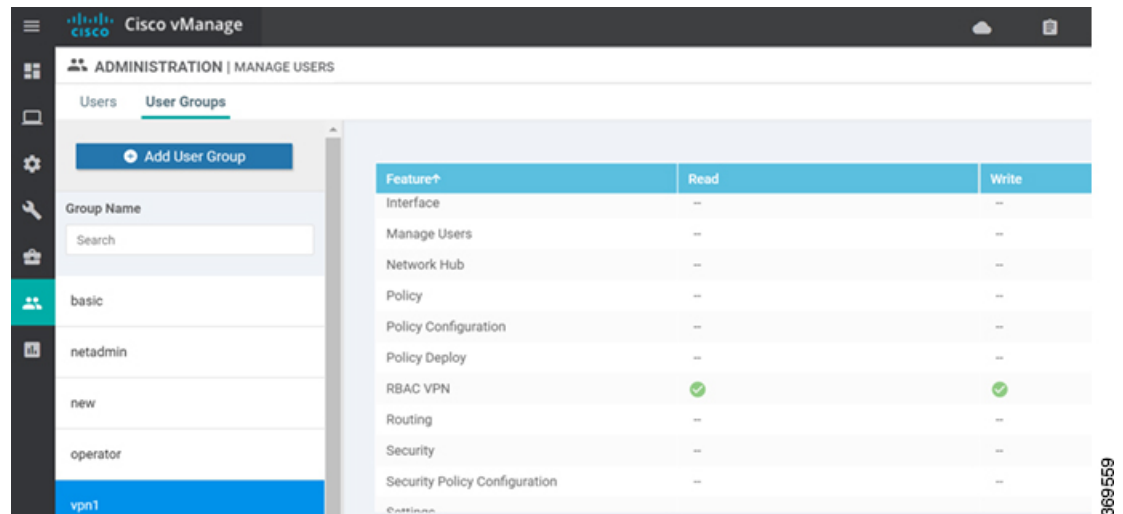
5. In the Create VPN Group pane, Enter VPN group name in the **VPN Group Name** field.
6. Enter a brief description of the VPN in the **Description** field.
7. Enable the user group access checkbox and enter the User Group Name.
8. In the Assign Segment pane, click on Add Segment drop-down to add new or existing segment to the VPN group.
9. Enter the Segment Name and VPN Number in the respective fields.
10. Click **Add** to add the configure VPN group to a device.

Configure User with User group

To create users with user group that is associated with the VPN group:

1. Navigate to **Administration > Manage Users** from Cisco vManage. The manage Users window appears.

- To edit, delete, or change password for an existing user, click the **Edit, Delete, or Change Password** in the More Info (...) column on the right side.
- Click on **Add User** to add a new user.
- In the Add New User page, add **Full Name, Username, Password, and Confirm Password** details.
- In the User Group drop-down, select the user group where you want to add a user.
- If you want to add a User Group, click on **Add User Group** button.



- Enter the user group name in the **Group Name** field.
- Select the Read or Write checkbox that you want to assign to a user group as shown in the figure.

Configure Interface Properties

Set the Interface Speed

When a Cisco vEdge device comes up, the Cisco SD-WAN software autodetects the SFPs present in the router and sets the interface speed accordingly. The software then negotiates the interface speed with the device at the remote end of the connection to establish the actual speed of the interface. To display the hardware present in the router, use the **show hardware inventory** command:

```
vEdge# show hardware inventory
```

HW TYPE	HW DEV INDEX	VERSION	PART NUMBER	SERIAL NUMBER	DESCRIPTION
Chassis	0	3.1	vEdge-1000	110D145130001	vEdge-1000
CPU	0	None	None	None	Quad-Core Octeon-II
DRAM	0	None	None	None	2048 MB DDR3
Flash	0	None	None	None	nor Flash - 16.00 MB
eMMC	0	None	None	None	eMMC - 7.31 GB
PIM	0	None	ge-fixed-8	None	8x 1GE Fixed Module
Transceiver	0	A	FCLF-8521-3	PQD3FHL	Port 0/0, Type 0x8 (Copper), Vendor FINISAR CORP.
Transceiver	1	PB	1GBT-SFP05	0000000687	Port 0/1, Type 0x8 (Copper), Vendor BEL-FUSE
FanTray	0	None	None	None	Fixed Fan Tray - 2 Fans

To display the actual speed of each interface, use the **show interface** command. Here, interface **ge0/0**, which connects to the WAN cloud, is running at 1000 Mbps (1Gbps; it is the 1GE P1M highlighted in the output above), and interface **ge0/1**, which connects to a device at the local site, has negotiated a speed of 100 Mbps.

```
vEdge# show interface
```

VPN	INTERFACE	IP ADDRESS	IF ADMIN STATUS	IF OPER STATUS	ENCAP TYPE	PORT TYPE	MTU	HWADDR	SPEED MBPS	DUPLEX	TCP MSS ADJUST	UPTIME	RX PACKETS	TX PACKETS
0	ge0/0	192.168.1.4/24	Up	Up	null	transport	1500	00:0c:bd:05:f0:83	1000	full	1300	0:06:10:59	2176305	2168760
0	ge0/2	-	Down	Down	null	service	1500	00:0c:bd:05:f0:81	-	-	0	-	0	0
0	ge0/3	-	Down	Down	null	service	1500	00:0c:bd:05:f0:82	-	-	0	-	0	0
0	ge0/4	-	Down	Down	null	service	1500	00:0c:bd:05:f0:87	-	-	0	-	0	0
0	ge0/5	-	Down	Down	null	service	1500	00:0c:bd:05:f0:88	-	-	0	-	0	0
0	ge0/6	-	Down	Down	null	service	1500	00:0c:bd:05:f0:85	-	-	0	-	0	0
0	ge0/7	-	Down	Down	null	service	1500	00:0c:bd:05:f0:86	-	-	0	-	0	0
0	system	1.1.1.1/32	Up	Up	null	loopback	1500	00:00:00:00:00:00	10	full	0	0:06:11:15	0	0
1	ge0/1	10.192.1.1/28	Up	Up	null	service	1500	00:0c:bd:05:f0:84	100	full	0	0:06:10:59	87	67
1	loopback1	1.1.1.1/32	Up	Up	null	service	1500	00:00:00:00:00:00	10	full	0	0:06:10:59	0	0
2	loopback0	10.192.1.2/32	Up	Up	null	service	1500	00:00:00:00:00:00	10	full	0	0:06:10:59	0	0
512	mgmt0	-	Up	Down	null	mgmt	1500	00:0c:bd:05:f0:80	-	-	0	-	0	0

For non-physical interfaces, such as those for the system IP address and loopback interfaces, the interface speed is set by default to 10 Mbps.

To override the speed negotiated by the two devices on the interface, disable autonegotiation and configure the desired speed:

```
vEdge(config-vpn) # interface interface-name no autonegotiate
vEdge(config-vpn) # interface interface-name speed (10 | 100)
```

For Cisco vSmart Controllers and Cisco vManage NMS systems, the initial interface speeds are 1000 Mbps, and the operating speed is negotiated with the device at the remote end of the interface. The controller interface speed may vary depending upon the virtualization platform, the NIC used, and the drivers that are present in the software.

Set the Interface MTU

By default, all interfaces have an MTU of 1500 bytes. You can modify this on an interface:

```
vEdge(config-vpn) # interface interface-name mtu bytes
```

The MTU can range from 576 through 2000 bytes.

To display an interface's MTU, use the **show interface** command.

For Cisco vBond Orchestrator, Cisco vManage, and Cisco vSmart Controller devices, you can configure interfaces to use ICMP to perform path MTU (PMTU) discovery. When PMTU discovery is enabled, the device to automatically negotiates the largest MTU size that the interface supports in an attempt to minimize or eliminate packet fragmentation:

```
vEdge(config-vpn) # interface interface-name pmtu
```

On Cisco vEdge device, the Cisco SD-WAN BFD software automatically performs PMTU discovery on each transport connection (that is, for each TLOC, or color). BFD PMTU discovery is enabled by default, and it is recommended that you use it and not disable it. To explicitly configure BFD to perform PMTU discovery, use the **bfd color pmtu-discovery** configuration command. However, you can choose to instead use ICMP to perform PMTU discovery:

```
vEdge(config-vpn) # interface interface-name pmtu
```

BFD is a data plane protocol and so does not run on Cisco vBond Orchestrator, Cisco vManage, and Cisco vSmart Controller devices.

Monitoring Bandwidth on a Transport Circuit

You can monitor the bandwidth usage on a transport circuit, to determine how the bandwidth usage is trending. If the bandwidth usage starts approaching a maximum value, you can configure the software to send a notification. Notifications are sent as Netconf notifications, which are sent to the Cisco vManage NMS, SNMP traps, and syslog messages. You might want to enable this feature for bandwidth monitoring, such as when you are doing capacity planning for a circuit or when you are gathering trending information about bandwidth utilization. You might also enable this feature to receive alerts regarding bandwidth usage, such as if you need to determine when a transport interface is becoming so saturated with traffic that a customer's traffic is impacted, or when customers have a pay-per-use plan, as might be the case with LTE transport.

To monitor interface bandwidth, you configure the maximum bandwidth for traffic received and transmitted on a transport circuit. The maximum bandwidth is typically the bandwidth that has been negotiated with the circuit provider. When bandwidth usage exceeds 85 percent of the configured value for either received or transmitted traffic, a notification, in the form of an SNMP trap, is generated. Specifically, interface traffic is sampled every 10 seconds. If the received or transmitted bandwidth exceeds 85 percent of the configured value in 85 percent of the sampled intervals in a continuous 5-minute period, an SNMP trap is generated. After the first trap is generated, sampling continues at the same frequency, but notifications are rate-limited to once per hour. A second trap is sent (and subsequent traps are sent) if the bandwidth exceeds 85 percent of the value in 85 percent of the 10-second sampling intervals over the next 1-hour period. If, after 1 hour, another trap is not sent, the notification interval reverts to 5 minutes.

You can monitor transport circuit bandwidth on Cisco vEdge devices and on Cisco vManage NMSs.

To generate notifications when the bandwidth of traffic received on a physical interface exceeds 85 percent of a specific bandwidth, configure the downstream bandwidth:

```
vEdge/vManage(config)# vpn vpn-id interface interface-name bandwidth-downstream kbps
```

To generate notifications when the bandwidth of traffic transmitted on a physical interface exceeds 85 percent of a specific bandwidth, configure the upstream bandwidth:

```
vEdge/vManage(config)# vpn vpn-id interface interface-name bandwidth-upstream kbps
```

In both configuration commands, the bandwidth can be from 1 through $2147483647 (2^{32} / 2) - 1$ kbps.

To display the configured bandwidths, look at the bandwidth-downstream and bandwidth-upstream fields in the output of the **show interface detail** command. The rx-kbps and tx-kbps fields in this command shows the current bandwidth usage on the interface.

Enable DHCP Server using Cisco vManage

Use the DHCP-Server template for all Cisco SD-WANs

You enable DHCP server functionality on a Cisco SD-WAN device interface so it can assign IP addresses to hosts in the service-side network.

To configure a Cisco SD-WAN device to act as a DHCP server using Cisco vManage templates:

1. Create a DHCP-Server feature template to configure DHCP server parameters, as described in this topic.
2. Create one or more interface feature templates, as described in the VPN-Interface-Ethernet and the VPN-Interface-PPP-Ethernet help topics.
3. Create a VPN feature template to configure VPN parameters. See the VPN help topic.

To configure a Cisco vEdge device interface to be a DHCP helper so that it forwards broadcast DHCP requests that it receives from DHCP servers, in the DHCP Helper field of the applicable interfaces template, enter the addresses of the DHCP servers.

Navigate to the Template Screen and Name the Template

1. In Cisco vManage NMS, select the Configuration ► Templates screen.
2. In the Device tab, click Create Template.
3. From the Create Template drop-down, select From Feature Template.
4. From the Device Model drop-down, select the type of device for which you are creating the template.
5. Click the Service VPN tab located directly beneath the Description field, or scroll to the Service VPN section.
6. Click the Service VPN drop-down.
7. Under Additional VPN Templates, located to the right of the screen, click VPN Interface.
8. From the Sub-Templates drop-down, select DHCP Server.
9. From the DHCP Server drop-down, click Create Template. The DHCP-Server template form is displayed. The top of the form contains fields for naming the template, and the bottom contains fields for defining DHCP Server parameters.

The screenshot shows the Cisco vManage interface for configuring a DHCP Server template. The left sidebar contains navigation options: Dashboard, Monitor, Configuration (selected), Devices, Certificates, Templates, Policies, Security, CloudExpress, Cloud onRamp, Tools, Maintenance, Administration, and vAnalytics. The main content area is titled 'CONFIGURATION | TEMPLATES' and shows the 'Device Feature' section. The 'Device Type' is set to 'vEdge Cloud'. Below this, there are fields for 'Template Name' and 'Description'. The 'Basic Configuration' tab is active, showing fields for 'Address Pool', 'Exclude Addresses', 'Maximum Leases', 'Lease Time (seconds)' (set to 86400), 'Offer Time (seconds)' (set to 600), and 'Administrative State'. 'Save' and 'Cancel' buttons are at the bottom right.

369419

10. In the Template Name field, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
11. In the Template Description field, enter a description of the template. The description can be up to 2048 characters and can contain only alphanumeric characters.

When you first open a feature template, for each parameter that has a default value, the scope is set to Default (indicated by a check mark), and the default setting or value is shown. To change the default or to enter a value, click the scope drop-down to the left of the parameter field.

Minimum DHCP Server Configuration

To configure DHCP server functionality, select the **Basic Configuration** tab and configure the following parameters. Parameters marked with an asterisk as required to configure DHCP servers.

Table 28:

Parameter Name	Description
Address Pool*	Enter the IPv4 prefix range, in the format <i>prefix/length</i> , for the pool of addresses in the service-side network for which the router interface acts as DHCP server.
Exclude Addresses	Enter one or more IP addresses to exclude from the DHCP address pool. To specify multiple individual addresses, list them separated by a comma. To specify a range of addresses, separate them with a hyphen.
Maximum Leases	Specify the number of IP addresses that can be assigned on this interface. <i>Range:</i> 0 through 4294967295
Lease Time	Specify how long a DHCP-assigned IP address is valid. <i>Range:</i> 0 through 4294967295 seconds
Offer Time	Specify how long the IP address offered to a DHCP client is reserved for that client. By default, an offered IP address is reserved indefinitely, until the DHCP server runs out of addresses. At that point, the address is offered to another client. <i>Range:</i> 0 through 4294967295 seconds <i>Default:</i> 600 seconds
Administrative State	Select Up to enable or Down to disable the DHCP functionality on the interface. By default, DHCP server functionality is disabled on an interface.

To save the feature template, click **Save**.

```
vpn vpn-id
interface geslot/port
dhcp-server address-pool prefix/length admin-state (down | up)
    exclude ip-address
    lease-time seconds
    max-leases number
    offer-time minutes
```

Configure Static Leases

To configure a static lease to assign a static IP address to a client device on the service-side network, click the Static Lease tab. Then click Add New Static Lease and configure the following parameters:

Table 29:

Parameter Name	Description
MAC Address	Enter the MAC address of the client to which the static IP address is being assigned.

Parameter Name	Description
IP Address	Enter the static IP address to assign to the client.
Hostname	Enter the hostname of the client device.

To edit a static lease, click the pencil icon to the right of the entry.

To remove a static lease, click the trash icon to the right of the entry.

To save the feature template, click **Save**.

CLI equivalent:

```
vpn vpn-id
interface geslot/port
dhcp-server static-lease mac-address ip ip-address host-name hostname
```

Configure Advanced Options

To configure a advanced DHCP server options, click the Advanced tab and then configure the following parameters:

Table 30:

Parameter Name	Description
Interface MTU	Specify the maximum MTU size of packets on the interface. <i>Range:</i> 68 to 65535 bytes
Domain Name	Specify the domain name that the DHCP client uses to resolve hostnames.
Default Gateway	Enter the IP address of a default gateway in the service-side network.
DNS Servers	Enter one or more IP address for a DNS server in the service-side network. Separate multiple entries with a comma. You can specify up to eight addresses.
TFTP Servers	Enter the IP address of a TFTP server in the service-side network. You can specify one or two addresses. If two, separate them with a comma.

To save the feature template, click **Save**.

CLI equivalent:

```
vpn vpn-id
interface geslot/port
dhcp-server options
    default-gateway ip-address
    dns-servers ip-address
    domain-name domain-name
    interface-mtu mtu
    tftp-servers ip-address
```

Release Information

Introduced in Cisco vManage NMS in Release 15.2.

Configure DHCP Using CLI

When you configure a tunnel interface on a Cisco vEdge device, a number of services are enabled by default on that interface, including DHCP.

A Cisco vEdge device can act as a DHCP server for the service-side network to which it is connected, and it can also act as a DHCP helper, forwarding requests for IP addresses from devices in the service-side network to a DHCP server that is in a different subnet on the service side of the Cisco vEdge device.

Enable DHCP on the WAN Interface

On a Cisco vEdge device's WAN interface—the interface configured as a tunnel interface in VPN 0, the transport VPN—DHCP is enabled by default. You can see this by using the **details** filter with the **show running-config** command. This command also shows that the DNS and ICMP services are enabled by default.

```
vm1# show running-config vpn 0 interface ge0/2 tunnel-interface | details
vpn 0
  interface ge0/2
    tunnel-interface
      encapsulation ipsec weight 1
      color lte
      control-connections
      carrier default
      no allow-service all
      no allow-service bgp
      allow-service dhcp
      allow-service dns
      allow-service icmp
      no allow-service ospf
      no allow-service sshd
      no allow-service ntp
      no allow-service stun
    !
  !
!
```

Enabling DHCP on the router's WAN interface allows the device that actually connects the router to the transport network (such as a DSL router) to dynamically assign a DHCP address to the Cisco vEdge device. The DHCP service in VPN 0 affects the transport-side network.

Configure Cisco vEdge Device as a DHCP Server

One or more service-side interfaces on Cisco vEdge device can act as a DHCP server, assigning IP addresses to hosts in the service-side network. To do this, configure this function on the interface that connects the Cisco vEdge device to the local site's network. At a minimum, you must configure the pool of IP addresses available for assigning to hosts:

```
vEdge(config-vpn)# interface ge slot / port dhcp-serveraddress-pool ip-address / prefix
vEdge(config-dhcp-server)#
```

You can exclude IP addresses that fall within the range of the DHCP address pool:

```
vEdge(config-dhcp-server)#exclude ip-address
```

To specify multiple individual addresses, list them in a single **exclude** command, separated by a space (for example, **exclude 10.1.1.1 10.2.2.2 10.3.3.3**). To specify a range of addresses, separate them with a hyphen (for example, **exclude 1.1.1.1-1.1.1.10**).

You can also statically assign IP addresses to a host:

```
vEdge(config-dhcp-server)# static-lease mac-address ip ip-address
```

By default, the DHCP server on a single interface can assign 254 DHCP leases, and each lease is valid for 24 hours. The offer of an IP address is valid indefinitely, until that DHCP server runs out of addresses to offer. You can modify these values:

```
vEdge(config-dhcp-server)# max-leases number
vEdge(config-dhcp-server)# lease-time seconds
vEdge(config-dhcp-server)# offer-time seconds
```

These values can range from 0 through $(2^{32} - 1)$.

The Cisco SD-WAN software supports DHCP server options that allow you to configure the IP addresses of a default gateway, DNS server, and TFTP server in the service-side network and the network mask of the service-side network:

```
vEdge(config-dhcp-server)# options default-gateway ip-address
vEdge(config-dhcp-server)# options dns-servers ip-address
vEdge(config-dhcp-server)# options domain-name domain-name
vEdge(config-dhcp-server)# options interface-mtu mtu
vEdge(config-dhcp-server)# options tftp-servers ip-address
vEdge(config-dhcp-server)# options option-code 43 ascii | hex
vEdge(config-dhcp-server)# options option-code 191 ascii
```

Configure a Cisco vEdge Device as a DHCP Helper

One or more service-side interfaces on a Cisco vEdge device can be a DHCP helper. With this configuration, the interface forwards any broadcast BOOTP DHCP requests that it receives from hosts on the service-side network to the DHCP server or servers specified by the configured IP helper address (or addresses) and returns the assigned IP address to the requester.

When the DHCP server at the Cisco vEdge device's local site is on a different segment than the devices connected to the Cisco vEdge device or than the Cisco vEdge device itself. When configured as a DHCP helper, the Cisco vEdge device interface forwards any broadcast BOOTP DHCP requests that it receives to the DHCP server specified by the configured IP helper address.

To configure an interface as a DHCP helper, configure the IP address of the DHCP server on the interface that connects to the local site's network:

```
vEdge(config-vpn)# interface ge slot/port dhcp-helper ip-address
```

You can configure up to four IP addresses, and you must enter the addresses in a single **dhcp-helper** command.

In Releases 17.2.2 and later, you can configure up to eight IP address. You must enter all the addresses in a single **dhcp-helper** command.

Configuring PPPoE

The Point-to-Point Protocol over Ethernet (PPPoE) connects multiple users over an Ethernet local area network to a remote site through common customer premises equipment. PPPoE is commonly used in a broadband aggregation, such as by digital subscriber line (DSL). PPPoE provides authentication with the CHAP or PAP protocol. In the Cisco SD-WAN overlay network, Cisco SD-WAN devices can run the PPPoE client. The PPPoE server component is not supported.

To configure PPPoE client on a Cisco SD-WAN device, you create a PPP logical interface and link it to a physical interface. The PPPoE connection comes up when the physical interface comes up. You can link a PPP interface to only one physical interface on a Cisco SD-WAN device, and you can link a physical interface to only one PPP interface. To enable more than one PPPoE interfaces on a Cisco SD-WAN device, configure multiple PPP interfaces.

It is recommended that you configure quality of service (QoS) and shaping rate on a PPPoE-enabled physical interface, and not on the PPP interface.

PPPoE-enabled physical interfaces do not support:

- 802.1Q
- Subinterfaces
- NAT, PMTU, and tunnel interfaces. These are configured on the PPP interface and therefore not available on PPPoE-enabled interfaces.

The Cisco SD-WAN implementation of PPPoE does not support the Compression Control Protocol (CCP) options, as defined in RFC 1962.

Configure PPPoE from vManage Templates

To use vManage templates to configure PPPoE on Cisco vEdge device, you create three feature templates and one device template:

- Create a VPN-Interface-PPP feature template to configure PPP parameters for the PPP virtual interface.
- Create a VPN-Interface-PPP-Ethernet feature template to configure a PPPoE-enabled interface.
- Optionally, create a VPN feature template to modify the default configuration of VPN 0.
- Create a device template that incorporates the VPN-Interface-PPP, VPN-Interface-PPP-Ethernet, and VPN feature templates.

To create a VPN-Interface-PPP feature template to configure PPP parameters for the PPP virtual interface:

Table 31:

Parameter Field	Procedure
Template Name	Enter a name for the template. It can be up to 128 alphanumeric characters.
Description	Enter a description for the template. It can be up to 2048 alphanumeric characters.
Shutdown	Click No to enable the PPP virtual interface.
Interface Name	Enter the number of the PPP interface. It can be from 1 through 31.
Description (optional)	Enter a description for the PPP virtual interface.
Authentication Protocol	Select either CHAP or PAP to configure one authentication protocol, or select PAP and CHAP to configure both. For CHAP, enter the hostname and password provided by your ISP. For PAP, enter the username and password provided by your ISP. If you are configuring both PAP and CHAP, to use the same username and password for both, click Same Credentials for PAP and CHAP.
AC Name (optional)	Select the PPP tab, and in the AC Name field, enter the name of the the name of the access concentrator used by PPPoE to route connections to the Internet.

Parameter Field	Procedure
IP MTU	Click the Advanced tab, and In the IP MTU field, ensure that the IP MTU is at least 8 bytes less than the MTU on the physical interface. The maximum MTU for a PPP interface is 1492 bytes. If the PPPoE server does not specify a maximum receive unit (MRU), the MTU value for the PPP interface is used as the MRU.
Save	Click Save to save the feature template.

1. In vManage NMS, select the Configuration ► Templates screen.
2. From the Templates title bar, select Feature.
3. Click Add Template.
4. In the left pane, select Cisco vEdge device Cloud or a router model.
5. In the right pane, select the VPN-Interface-PPP template.
6. In the template, configure the following parameters:

To create a VPN-Interface-PPP-Ethernet feature template to enable the PPPoE client on the physical interfaces:

1. In the vManage NMS, select the Configuration ► Templates screen.
2. From the Templates title bar, select Feature.
3. Click Add Template.
4. In the left pane, select Cisco vEdge device Cloud or a router model.
5. In the right pane, select the VPN-Interface-PPP-Ethernet template.
6. In the template, configure the following parameters:

Parameter Field	Procedure
Template Name	Enter a name for the template. It can be up to 128 alphanumeric characters.
Description	Enter a description for the template. It can be up to 2048 alphanumeric characters.
Shutdown	Click No to enable the PPPoE-enabled interface.
Interface Name	Enter the name of the physical interface in VPN 0 to associate with the PPP interface.
Description (optional)	Enter a description for the PPPoE-enabled interface.
IP Configuration	Assign an IP address to the physical interface: <ul style="list-style-type: none"> • To use DHCP, select Dynamic. The default administrative distance of routes learned from DHCP is 1. • To configure the IP address directly, enter of the IPv4 address of the interface.
DHCP Helper (optional)	Enter up to four IP addresses for DHCP servers in the network.

Parameter Field	Procedure
Save	Click Save to save the feature template.

To create a VPN feature template to configure the PPPoE-enabled interface in VPN 0, the transport VPN:

1. In the vManage NMS, select the Configuration ► Templates screen.
2. From the Templates title bar, select Feature.
3. Click Add Template.
4. In the left pane, select Cisco vEdge device Cloud or a router model.
5. In the right pane, select the VPN template.
6. In the template, configure the following parameters:

Parameter Field	Procedure
Template Name	Enter a name for the template. It can be up to 128 alphanumeric characters.
Description	Enter a description for the template. It can be up to 2048 alphanumeric characters.
VPN Identifier	Enter VPN identifier 0.
Name	Enter a name for the VPN.
Other interface parameters	Configure the desired interface properties.
Save	Click Save to save the feature template.

To create a device template that incorporates the VPN-Interface-PPP, VPN-Interface-PPP-Ethernet, and VPN feature templates:

1. In the vManage NMS, select the Configuration ► Templates screen.
2. From the Templates title bar, select Device.
3. Click Create Template, and from the drop-down list select From Feature Template.
4. From the Device Model drop-down, select the type of device for which you are creating the device template. vManage NMS displays the feature templates for the device type you selected. Required templates are indicated with an asterisk (*).
5. Enter a name and description for the device template. These fields are mandatory. The template name cannot contain special characters.
6. In the Transport & Management VPN section, under VPN 0, from the drop-down list of available templates, select the desired feature template. The list of available templates are the ones that you have previously created.
7. In the Additional VPN 0 Templates section to the right of VPN 0, click the plus sign (+) next to VPN Interface PPP.
8. In the VPN-Interface-PPP and VPN-Interface-PPP-Ethernet fields, select the feature templates to use.

9. To configure multiple PPPoE-enabled interfaces in VPN 0, click the plus sign (+) next to Sub-Templates.
10. To include additional feature templates in the device template, in the remaining sections, select the feature templates in turn, and from the drop-down list of available templates, select the desired template. The list of available templates are the ones that you have previously created. Ensure that you select templates for all mandatory feature templates and for any desired optional feature templates.
11. Click Create to create the device template.

To attach a device template to a device:

1. In the vManage NMS, select the Configuration ► Templates screen.
2. From the Templates title bar, select Device.
3. Select a template.
4. Click the More Actions icon to the right of the row and click Attach Device.
5. In the Attach Device window, either search for a device or select a device from the Available Device(s) column to the left.
6. Click the arrow pointing right to move the device to the Selected Device(s) column on the right.
7. Click Attach.

Configure PPPoE from the CLI

To use the CLI to configure PPPoE on Cisco vEdge devices:

1. Create a PPP interface. The interface number can be from 1 through 31.

```
vEdge (config-vpn) # interface ppp number
```

2. Configure an authentication method for PPPoE and authentication credentials. You can configure both CHAP and PAP authentication on the same PPP interface. The software tries both methods and uses the first one that succeeds.

```
vEdge (config-interface-ppp) # ppp authentication chap hostname name password password
vEdge (config-interface-ppp) # ppp authentication pap password password sent-username
username
```

- 3.
4. Enable the PPP interface to be operationally up:

```
vEdge (config-interface-ppp) # no shutdown
```

5. Configure the MTU of the PPP interface. The maximum MTU for a PPP interface is 1492 bytes. If maximum receive unit (MRU) is not specified by the PPPoE server, the MTU value for the PPP interface is used as the MRU.

```
vEdge (config-interface-ppp) # mtu bytes
```

6. Configure a tunnel interface for the PPP interface:

```
vEdge (config-interface-ppp) # tunnel-interface color color
```

7. Optionally, configure the name of the access concentrator used by PPPoE to route connections to the internet:

```
vEdge(config-interface-ppp)# ac-name name
```

8. Link a physical Gigabit Ethernet interface in VPN 0 to the PPP interface:

```
vEdge(config-interface-ge)# pppoe-client ppp-interface ppp number
```

9. Enable the physical Gigabit Ethernet interface to be operationally up:

```
vEdge(config-vpn-interface-ge)# no shutdown
```

Here is an example of a PPPoE configuration:

```
vEdge# show running-config vpn 0
vpn 0
interface ge0/1
 pppoe-client ppp-interface ppp10
 no shutdown
!
interface ppp10
 ppp authentication chap
 hostname branch100@corp.bank.myisp.net
 password $4$0HHjdmsC6M8zj4BgLEFXKw==
!
tunnel-interface
 encapsulation ipsec
 color gold
 no allow-service all
 no allow-service bgp
 allow-service dhcp
 allow-service dns
 allow-service icmp
 no allow-service ospf
 no allow-service sshd
 no allow-service ntp
 no allow-service stun
!
mtu 1492
 no shutdown
!
!
```

To view existing PPP interfaces, use the **show ppp interface** command. For example:

```
vEdge# show ppp interface
```

VPN	IFNAME	PPPOE INTERFACE	INTERFACE IP	GATEWAY IP	PRIMARY DNS	SECONDARY DNS	MTU
0	ppp10	ge0/1	11.1.1.1	115.0.1.100	8.8.8.8	8.8.4.4	1150

To view PPPoE session information, use the **show pppoe session** command. For example:

```
vEdge# show pppoe session
```

VPN	IFNAME	SESSION			PPP		SERVICE
		ID	SERVER MAC	LOCAL MAC	INTERFACE	AC NAME	NAME
0	ge0/1	1	00:0c:29:2e:20:1a	00:0c:29:be:27:f5	ppp1	branch100	-
0	ge0/3	1	00:0c:29:2e:20:24	00:0c:29:be:27:13	ppp2	branch100	-

Configuring VRRP

The Virtual Router Redundancy Protocol (VRRP) provides redundant gateway service for switches and other IP end stations. In the Cisco SD-WAN software, you configure VRRP on an interface, and typically on a subinterface, within a VPN .

For a VRRP interface to operate, its physical interface must be configured in VPN 0:

```
vEdge(config-vpn-0)# interface ge- slot / port
vEdge(config-interface-ge)# no shutdown
```

For each VRRP interface (or subinterface), you assign an IP address and you place that interface in a VRRP group.

```
vEdge(config-vpn)# interface ge- slot / port . subinterface
vEdge(config-interface-ge)# ip address prefix / length
vEdge(config-interface-ge)# vrrp group-number
```

The group number identifies the virtual router. You can configure a maximum of 24 groups on a router. In a typical VRRP topology, two physical routers are configured to act as a single virtual router, so you configure the same group number on interfaces on both these routers.

For each virtual router ID, you must configure an IP address.

```
vEdge(config-vrrp)# ipv4 ip-address
```

Within each VRRP group, the router with the higher priority value is elected as primary VRRP. By default, each virtual router IP address has a default primary election priority of 100, so the router with the higher IP address is elected as primary. You can modify the priority value, setting it to a value from 1 through 254.

```
vEdge(config-vrrp)# priority number
```

The primary VRRP periodically sends advertisement messages, indicating that it is still operating. If backup routers miss three consecutive VRRP advertisements, they assume that the primary VRRP is down and elect a new primary VRRP. By default, these messages are sent every second. You can change the VRRP advertisement time to be a value from 1 through 3600 seconds.

```
vEdge(config-vrrp)# timer seconds
```

By default, VRRP uses the state of the interface on which it is running, to determine which router is the primary virtual router. This interface is on the service (LAN) side of the router. When the interface for the primary VRRP goes down, a new primary VRRP virtual router is elected based on the VRRP priority value. Because VRRP runs on a LAN interface, if a router loses all its WAN control connections, the LAN interface still indicates that it is up even though the router is functionally unable to participate in VRRP. To take WAN side connectivity into account for VRRP, you can configure one of the following:

- Track the Overlay Management Protocol (OMP) session running on the WAN connection when determining the primary VRRP virtual router.

```
vEdge(config-vrrp)# track-omp
```

If all OMP sessions are lost on the primary VRRP router, VRRP elects a new default gateway from among all the gateways that have one or more active OMP sessions even if the gateway chosen has a lower VRRP priority than the current primary VRRP router. With this option, VRRP failover occurs once the OMP state changes from up to down, which occurs when the OMP hold timer expires. (The default OMP hold timer interval is 60 seconds.) Until the hold timer expires and a new primary VRRP is elected, all overlay traffic is dropped. When the OMP session recovers, the local VRRP interface claims itself as primary VRRP even

before it learns and installs OMP routes from the Cisco vSmart Controllers. Until the routers are learned, traffic is also dropped.

- Track both the OMP session and a list of remote prefixes. *list-name* is the name of a prefix list configured with the **policy lists prefix-list** command on the Cisco vEdge device :

```
vEdge(config-vrrp)# track-prefix-list list-name
```

If all OMP sessions are lost, VRRP failover occurs as described for the **track-omp** option. In addition, if reachability to all the prefixes in the list is lost, VRRP failover occurs immediately, without waiting for the OMP hold timer to expire, thus minimizing the amount of overlay traffic is dropped while the router determines the primary VRRP.

As discussed above, the IEEE 802.1Q protocol adds 4 bytes to each packet's length. Hence, for packets to be transmitted, either increase the MTU size on the physical interface in VPN 0 (the default MTU is 1500 bytes) or decrease the MTU size on the VRRP interface.

Here is an example of configuring VRRP on redundant physical interfaces. For subinterface 2, vEdge1 is configured to act as the primary VRRP, and for subinterface 3, vEdge2 acts as the primary VRRP.

```
vEdge1# show running-config vpn 1
vpn 1
interface ge0/6.2
 ip address 10.2.2.3/24
 mtu 1496
 no shutdown
 vrrp 2
  ipv4 10.2.2.1
  track-prefix-list vrrp-prefix-list1
!
!
interface ge0/6.3
 ip address 10.2.3.5/24
 mtu 1496
 shutdown
 vrrp 3
  ipv4 10.2.3.11
  track-prefix-list vrrp-prefix-list1
!
!
```

```
vEdge2# show running-config vpn 1
vpn 1
interface ge0/1.2
 ip address 10.2.2.4/24
 mtu 1496
 no shutdown
 vrrp 2
  ipv4 10.2.2.1
  track-prefix-list vrrp-prefix-list2
!
!
interface ge0/1.3
 ip address 10.2.3.6/24
 mtu 1496
 no shutdown
 vrrp 3
  ipv4 10.2.3.11
  track-prefix-list vrrp-prefix-list2
!
!
```

```
vEdge1# show interface vpn 1
```

IF	IF				TCP
ADMIN	OPER	ENCAP	PORT	SPEED	MSS

VPN	INTERFACE	RX		TX		STATUS	STATUS	TYPE	TYPE	MTU	HWADDR	MBPS	DUPLEX	
		IP ADDRESS	PACKETS	IP ADDRESS	PACKETS									
1	ge0/6.2	10.2.2.3/24	0	10.2.3.5/24	357	Up	Up	vlan	service	1496	00:0c:29:ab:b7:94	10	full	0
1	ge0/6.3	10.2.3.5/24	0	10.2.2.3/24	0	Down	Down	vlan	service	1496	00:0c:29:ab:b7:94	-	-	0

vEdge1# show vrrp interfaces

VPN	IF NAME	ID	GROUP	TRACK PREFIX		VRRP	OMP	ADVERTISEMENT	DOWN	TIMER	LAST
				VIRTUAL IP	VIRTUAL MAC						
1	ge0/6.2	2	10.2.2.1	00:0c:29:ab:b7:94	100	master	down	1	3		
			2015-05-01T20:09:37+00:00	-	-						
	ge0/6.3	3	10.2.3.11	00:00:00:00:00:00	100	init	down	1	3		
			0000-00-00T00:00:00+00:00	-	-						

In the following example, Router-1 is the primary VRRP, because it has a higher priority value than Router 2:

```
Router-1# show running-config vpn 1
vpn 1
!
interface ge0/1.15
 ip address 10.10.1.2/24
 mtu 1496
 no shutdown
 vrrp 15
  priority 110
  track-omp
  ipv4 10.20.23.1
!
!
```

Router-1# show vrrp vpn 1

VPN	IF NAME	ID	GROUP	TRACK PREFIX		VRRP	OMP	ADVERTISEMENT	DOWN	TIMER	LAST
				VIRTUAL IP	VIRTUAL MAC						
1	ge0/1.1	1	10.20.22.1	00:0c:bd:08:79:a4	100	backup	up	1	3		
			2016-01-13T03:10:55+00:00	-	-						
	ge0/1.5	5	10.20.22.193	00:0c:bd:08:79:a4	100	backup	up	1	3		
			2016-01-13T03:10:55+00:00	-	-						
	ge0/1.10	10	10.20.22.225	00:0c:bd:08:79:a4	100	backup	up	1	3		
			2016-01-13T03:10:55+00:00	-	-						
	ge0/1.15	15	10.20.23.1	00:0c:bd:08:79:a4	110	master	up	1	3		
			2016-01-13T03:10:56+00:00	-	-						
	ge0/1.20	20	10.20.24.1	00:0c:bd:08:79:a4	100	backup	up	1	3		
			2016-01-13T03:10:56+00:00	-	-						
	ge0/1.25	25	10.20.25.1	00:0c:bd:08:79:a4	110	master	up	1	3		
			2016-01-13T03:10:56+00:00	-	-						
	ge0/1.30	30	10.20.25.129	00:0c:bd:08:79:a4	100	backup	up	1	3		
			2016-01-13T03:10:56+00:00	-	-						

Router-1# show vrrp vpn 1 interfaces ge0/1.15 groups 15

GROUP	ID	IF NAME	ID	TRACK PREFIX		VRRP	OMP	ADVERTISEMENT	DOWN	TIMER	LAST STATE CHANGE
				VIRTUAL IP	VIRTUAL MAC						
				10.20.23.1	00:0c:bd:08:79:a4	110	master	up	1	3	

```
-----
1      10.20.33.1 00:0c:bd:08:79:a4 110      master up      1          3
2016-01-13T03:10:56+00:00 -      -
```

```
Router-2# show running-config vpn 1
vpn 1
!
interface ge0/1.15
ip address 10.10.1.3/24
mtu      1496
no shutdown
vrrp 15
 track-omp
  ipv4 10.20.23.1
!
!
```

```
Router-2# show vrrp vpn 1 interfaces groups
```

IF NAME	STATE	ID	GROUP	VIRTUAL IP	VIRTUAL MAC	PRIORITY	STATE	OMP	ADVERTISEMENT	MASTER	
										TIMER	DOWN
ge0/1.1	2016-01-13T00:22:15+00:00	1	10.20.32.1	00:0c:bd:08:2b:a5	110	master	up	1	3		
ge0/1.5	2016-01-13T00:22:15+00:00	5	10.20.32.193	00:0c:bd:08:2b:a5	110	master	up	1	3		
ge0/1.10	2016-01-13T00:22:15+00:00	10	10.20.32.225	00:0c:bd:08:2b:a5	110	master	up	1	3		
ge0/1.15	2016-01-13T03:10:56+00:00	15	10.20.33.1	00:0c:bd:08:2b:a5	100	backup	up	1	3		
ge0/1.20	2016-01-13T00:22:16+00:00	20	10.20.34.1	00:0c:bd:08:2b:a5	110	master	up	1	3		
ge0/1.25	2016-01-13T03:10:56+00:00	25	10.20.35.1	00:0c:bd:08:2b:a5	100	backup	up	1	3		
ge0/1.30	2016-01-13T00:22:16+00:00	30	10.20.35.129	00:0c:bd:08:2b:a5	100	master	up	1	3		

```
Router-2# show vrrp vpn 100 interfaces groups 15
```


IF NAME	STATE	ID	GROUP	VIRTUAL IP	VIRTUAL MAC	PRIORITY	STATE	OMP	ADVERTISEMENT	MASTER	
										TIMER	DOWN
ge0/0.15	2016-01-13T03:10:56+00:00	15	10.20.33.1	00:0c:bd:08:2b:a5	100	backup	up	1	3		

Network Interface Configuration Examples for Cisco vEdge Devices

This topic provides examples of configuring interfaces on Cisco vEdge devices to allow the flow of data traffic across both public and private WAN transport networks.

Connect to a Public WAN

This example shows a basic configuration for two connected to the same public WAN network (such as the Internet). The Cisco vSmart Controller and Cisco vBond Orchestrator are also connected to the public WAN network, and the Cisco vSmart Controller is able to reach all destinations on the public WAN.



For Cisco vEdge device-1, the interface ge0/1 connects to the public WAN, so it is the interface that is configured as a tunnel interface. The tunnel has a color of biz-internet, and the encapsulation used for data traffic is IPsec. The Cisco SD-WAN software creates a single TLOC for this interface, comprising the interface's IP address, color, and encapsulation, and the TLOC is sent to the Cisco vSmart Controller over the OMP session running on the tunnel. The configuration also includes a default route to ensure that the router can reach the Cisco vBond Orchestrator and Cisco vSmart Controller.

```
vpn 0
 interface ge0/1
   ip address 172.16.13.3/24
   tunnel-interface
     encapsulation ipsec
     color biz-internet
     allow-service dhcp
     allow-service dns
     allow-service icmp
     no allow-service sshd
     no allow-service ntp
     no allow-service stun
   !
   no shutdown
 !
 ip route 0.0.0.0/0 172.16.13.1
 !
```

The configuration for Cisco vEdge device-2 is similar to that for Cisco vEdge device-1:

```
vpn 0
 interface ge0/1
   ip address 172.16.15.5/24
   tunnel-interface
     encapsulation ipsec
     color biz-internet
     allow-service dhcp
     allow-service dns
     allow-service icmp
     no allow-service sshd
     no allow-service ntp
     no allow-service stun
   !
```

```

    no shutdown
    !
    ip route 0.0.0.0/0 172.16.15.1
    !

```

On the Cisco vSmart Controller and Cisco vBond Orchestrator, you configure a tunnel interface and default IP route to reach the WAN transport. For the tunnel, color has no meaning because these devices have no TLOCs.

```

vpn 0
  interface eth1
    ip address 172.16.8.9/24
    tunnel-interface
    !
    no shutdown
    !
    ip route 0.0.0.0/0 172.16.8.1
    !
!

vpn 0
  interface ge0/1
    ip address 172.16.16.6/24
    tunnel-interface
    !
    no shutdown
    !
    ip route 0.0.0.0/0 172.16.16.1
    !
!

```

Use the **show interface** command to check that the interfaces are operational and that the tunnel connections have been established. In the Port Type column, tunnel connections are marked as "transport."

```
vEdge-1# show interface vpn 0
```

RX VPN	TX INTERFACE	IP ADDRESS	IF		ENCAP	PORT TYPE	MTU	HWADDR	TCP			
			ADMIN	OPER					SPEED	MSS	ADJUST	UPTIME
PACKETS	PACKETS		STATUS	STATUS	TYPE			MBPS	DUPLEX			
0	ge0/0	172.16.13.3/24	Up	Up	null	transport	1500	00:0c:29:7d:1e:fe	10	full	0	0:02:26:20
88358	88202											
0	ge0/1	10.1.17.15/24	Up	Up	null	service	1500	00:0c:29:7d:1e:08	10	full	0	0:02:26:20
217	1											
0	ge0/2	-	Down	Up	null	service	1500	00:0c:29:7d:1e:12	10	full	0	0:02:26:20
217	0											
0	ge0/3	10.0.20.15/24	Up	Up	null	service	1500	00:0c:29:7d:1e:1c	10	full	0	0:02:26:20
218	1											
0	ge0/6	57.0.1.15/24	Up	Up	null	service	1500	00:0c:29:7d:1e:3a	10	full	0	0:02:26:20
217	1											
0	ge0/7	10.0.100.15/24	Up	Up	null	service	1500	00:0c:29:7d:1e:44	10	full	0	0:02:25:02
850	550											
0	system	172.16.255.3/32	Up	Up	null	loopback	1500	00:00:00:00:00:00	10	full	0	0:02:13:31
0	0											

Use the **show control connections** command to check that the Cisco vEdge device has a DTLS or TLS session established to the Cisco vSmart Controller.

```
vEdge-1# show control connections
```

PEER	PEER	PEER	SITE	DOMAIN	PEER	PEER		PEER	
						PRIVATE	PEER	PUBLIC	PEER
TYPE	PROTOCOL	SYSTEM IP	ID	ID	PRIVATE IP	PORT	PUBLIC IP	PORT	LOCAL COLOR
STATE		UPTIME							
vsmart	dtls	172.16.255.19	100	1	10.0.5.19	12346	10.0.5.19	12346	biz-internet
up		0:02:13:13							


```
vsmart dtls 172.16.255.20 200 1 10.0.12.20 12346 10.0.12.20 12346 biz-internet
up 0:02:13:13
```

Use the **show bfd sessions** command to display information about the BFD sessions that have been established between the local Cisco vEdge device and remote routers:

```
vEdge-1# show bfd sessions
```

DST PUBLIC	DETECT	TX	SOURCE TLOC	REMOTE TLOC	DST PUBLIC
SYSTEM IP	SITE ID STATE	COLOR	COLOR	SOURCE IP	IP
PORT	ENCAP MULTIPLIER	INTERVAL (msec)	UPTIME	TRANSITIONS	
172.16.255.11	100 up	biz-internet	biz-internet	10.1.15.15	10.0.5.11
12346	ipsec 20	1000	0:02:24:59	1	
172.16.255.14	400 up	biz-internet	biz-internet	10.1.15.15	10.1.14.14
12360	ipsec 20	1000	0:02:24:59	1	
172.16.255.16	600 up	biz-internet	biz-internet	10.1.15.15	10.1.16.16
12346	ipsec 20	1000	0:02:24:59	1	
172.16.255.21	100 up	biz-internet	biz-internet	10.1.15.15	10.0.5.21
12346	ipsec 20	1000	0:02:24:59	1	

Use the **show omp tlocs** command to list the TLOCs that the local router has learned from the Cisco vSmart Controller:

```
vEdge-1# show omp tlocs
```

```
C -> chosen
I -> installed
Red -> redistributed
Rej -> rejected
L -> looped
R -> resolved
S -> stale
Ext -> extranet
Inv -> invalid
```

ADDRESS	PRIVATE	BFD	ENCAP	FROM PEER	STATUS	PUBLIC IP	PUBLIC
FAMILY	TLOC IP	COLOR					PORT
PRIVATE IP	PORT	STATUS					
ipv4	172.16.255.11	biz-internet	ipsec	172.16.255.19	C,I,R	10.0.5.11	12346
10.0.5.11	12346	up		172.16.255.20	C,R	10.0.5.11	12346
10.0.5.11	12346	up		172.16.255.20	C,R	10.1.14.14	12360
10.1.14.14	172.16.255.14	biz-internet	ipsec	172.16.255.19	C,I,R	10.1.14.14	12360
10.1.14.14	12360	up		172.16.255.20	C,R	10.1.14.14	12360
10.1.14.14	12360	up		172.16.255.20	C,R	10.1.14.14	12360
10.1.16.16	172.16.255.16	biz-internet	ipsec	172.16.255.19	C,I,R	10.1.16.16	12346
10.1.16.16	12346	up		172.16.255.20	C,R	10.1.16.16	12346
10.1.16.16	12346	up		172.16.255.20	C,R	10.1.16.16	12346
10.1.16.16	12346	up		172.16.255.20	C,R	10.1.16.16	12346
10.0.5.21	172.16.255.21	biz-internet	ipsec	172.16.255.19	C,I,R	10.0.5.21	12346
10.0.5.21	12346	up		172.16.255.20	C,R	10.0.5.21	12346
10.0.5.21	12346	up		172.16.255.20	C,R	10.0.5.21	12346
10.0.5.21	12346	up	<				

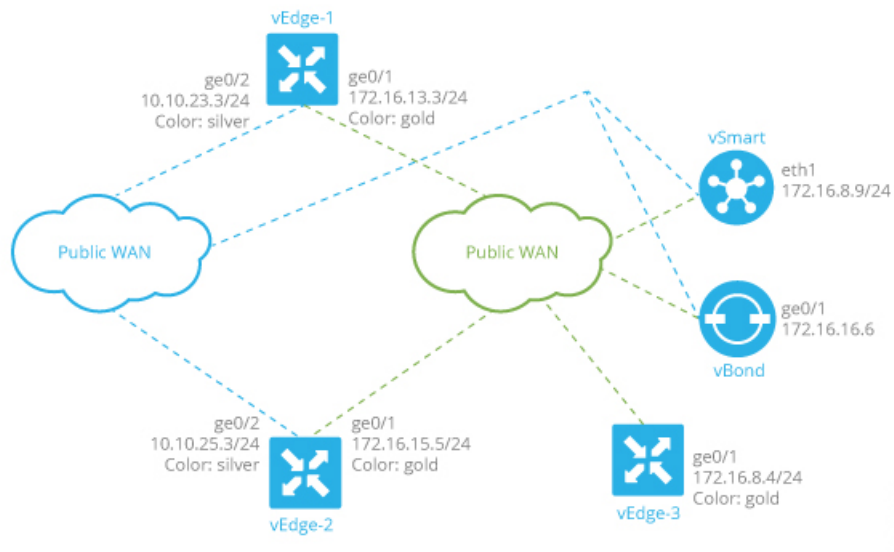
Connect to Two Public WANs

In this example, two Cisco vEdge devices at two different sites connect to two public WANs, and hence each router has two tunnel connections. To direct traffic to the two different WANs, each tunnel interface is assigned a different color (here, **silver** and **gold**). Because each router has two tunnels, each router has two TLOCs.

A third router at a third site, vEdge-3, connects only to one of the public WANs.

The Cisco vSmart Controller and Cisco vBond Orchestrator are connected to one of the public WAN networks. (In reality, it does not matter which of the two networks they are connected to, nor does it matter whether the two devices are connected to the same network.) The Cisco vSmart Controller is able to reach all destinations on the public WAN. To ensure that the Cisco vBond Orchestrator is accessible via each transport tunnel on

the routers, a default route is configured for each interface. In our example, we configure a static default route, but you can also use DHCP.



368481

The configurations for vEdge-1 and vEdge-2 are similar. We configure two tunnel interfaces, one with color **silver** and the other with color **gold**, and we configure static default routes for both tunnel interfaces. Here is the configuration for vEdge-1:

```
vpn 0
 interface ge0/1
   ip address 172.16.13.3/24
   tunnel-interface
     encapsulation ipsec
     color silver
   !
   no shutdown
 !
 interface ge0/2
   ip address 10.10.23.3/24
   tunnel-interface
     encapsulation ipsec
     color gold
   !
   no shutdown
 !
 ip route 0.0.0.0/0 172.16.13.1
 ip route 0.0.0.0/0 10.10.23.1
```

The configuration for vEdge-2 is similar:

```
vpn 0
 interface ge0/1
   ip address 172.16.15.5/24
   tunnel-interface
     encapsulation ipsec
     color silver
   !
   no shutdown
 !
 interface ge0/2
   ip address 10.10.25.3/24
   tunnel-interface
```

```

        encapsulation ipsec
        color gold
    !
    no shutdown
    !
ip route 0.0.0.0/0 172.16.15.1
ip route 0.0.0.0/0 10.10.25.1

```

The third router, vEdge-3, connects only to one of the public WAN networks, and its tunnel interface is assigned the color "gold":

```

vpn 0
  interface ge0/1
    ip address 172.16.8.4/24
    tunnel-interface
      encapsulation ipsec
      color gold
    !
    no shutdown
  !
  ip route 0.0.0.0/0 172.16.8.1

```

On the Cisco vSmart Controller and Cisco vBond Orchestrator, you configure a tunnel interface and default IP route to reach the WAN transport. For the tunnel, color has no meaning because these devices have no TLOCs.

```

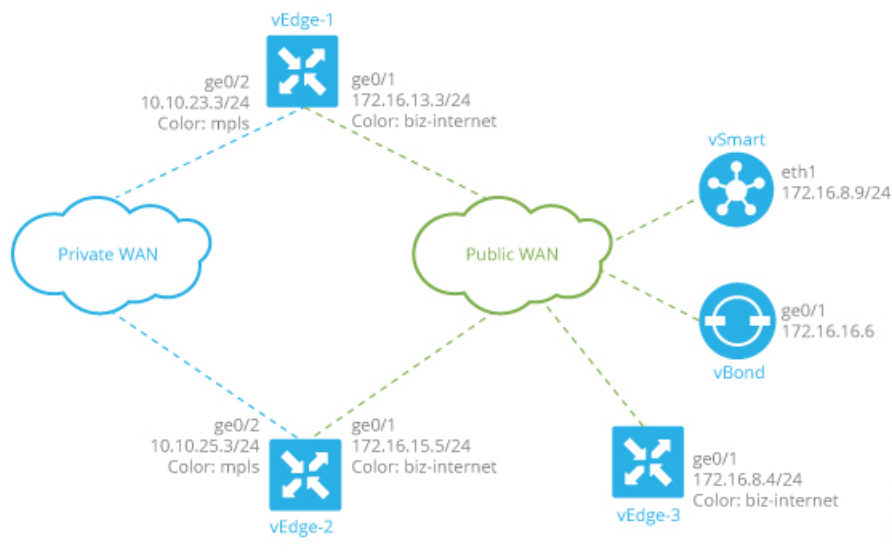
vpn 0
  interface eth1
    ip address 172.16.8.9/24
    tunnel-interface
    !
    no shutdown
  ip route 0.0.0.0/0 172.16.8.1

vpn 0
  interface ge0/1
    ip address 172.16.16.6/24
    tunnel-interface
    !
    no shutdown
  ip route 0.0.0.0/0 172.16.16.1

```

Connect to Public and Private WANs, with Separation of Network Traffic

In this example, two Cisco vEdge devices at two different sites each connect to the same public WAN (here, the Internet) and the same private WAN (here, an MPLS network). We want to separate the MPLS network completely so that it is not reachable by the Internet. The Cisco vSmart Controller and Cisco vBond Orchestrator are hosted in the provider's cloud, which is reachable only via the Internet. A third Cisco vEdge device at a third site connects only to the public WAN (Internet).



In this example topology, we need to ensure the following:

- Complete traffic separation exists between private-WAN (MPLS) traffic and public-WAN (Internet) traffic.
- Each site (that is, each Cisco vEdge device) must have a connection to the Internet, because this is the only way that the overlay network can come up.

To maintain complete separation between the public and private networks so that all MPLS traffic stays within the MPLS network, and so that only public traffic passes over the Internet, we create two overlays, one for the private MPLS WAN and the second for the public Internet. For the private overlay, we want to create data traffic tunnels (which run IPsec and BFD sessions) between private-WAN TLOCs, and for the public overlay we want to create these tunnel connections between Internet TLOCs. To make sure that no data traffic tunnels are established between private-WAN TLOCs and Internet TLOCs, or vice versa, we associate the **restrict** attribute with the color on the private-WAN TLOCs. When a TLOC is marked as restricted, a TLOC on the local router establishes tunnel connections with a remote TLOC only if the remote TLOC has the same color. Put another way, BFD sessions come up between two private-WAN TLOCs and they come up between two public-WAN TLOCs, but they do not come up between an MPLS TLOC and an Internet TLOC.

Each site must have a connection to the public (Internet) WAN so that the overlay network can come up. In this topology, the Cisco vSmart Controller and Cisco vBond Orchestrator are reachable only via the Internet, but the MPLS network is completely isolated from the Internet. This means that if a Cisco vEdge device were to connect just to the MPLS network, it would never be able to discover the Cisco vSmart Controller and Cisco vBond Orchestrators and so would never be able to establish control connections in the overlay network. In order for a Cisco vEdge device in the MPLS network to participate in overlay routing, it must have at least one tunnel connection, or more specifically, one TLOC, to the Internet WAN. (Up to seven TLOCs can be configured on each Cisco vEdge device.) The overlay network routes that the router learns over the public-WAN tunnel connection populate the routing table on the Cisco vEdge device and allow the router and all its interfaces and TLOCs to participate in the overlay network.

By default, all tunnel connections attempt to establish control connections in the overlay network. Because the MPLS tunnel connections are never going to be able to establish these connections to the Cisco vSmart Controller or Cisco vBond Orchestrators, we include the **max-control-connections 0** command in the configuration. While there is no harm in having the MPLS tunnels attempt to establish control connections, these attempts will never succeed, so disabling them saves resources on the Cisco vEdge device. Note that

max-control-connections 0 command works only when there is no NAT device between the Cisco vEdge device and the PE router in the private WAN.

Connectivity to sites in the private MPLS WAN is possible only by enabling service-side routing.

Here is the configuration for the tunnel interfaces on vEdge-1. This snippet does not include the service-side routing configuration.

```
vpn 0
  interface ge0/1
    ip address 172.16.13.3/24
    tunnel-interface
      encapsulation ipsec
      color biz-internet
    !
    no shutdown
  !
  interface ge0/2
    ip address 10.10.23.3/24
    tunnel-interface
      encapsulation ipsec
      color mpls restrict
      max-control-connections 0
    !
    no shutdown
  !
  ip route 0.0.0.0/0 172.16.13.1
```

The configuration on vEdge-2 is quite similar:

```
vpn 0
  interface ge0/1
    ip address 172.16.15.5/24
    tunnel-interface
      encapsulation ipsec
      color biz-internet
    !
    no shutdown
  !
  interface ge0/2
    ip address 10.10.25.3/24
    tunnel-interface
      encapsulation ipsec
      color mpls restrict
      max-control-connections 0
    !
    no shutdown
  !
  ip route 0.0.0.0/0 172.16.15.1
!
```

The vEdge-3 router connects only to the public Internet WAN:

```
vpn 0
  interface ge0/1
    ip address 172.16.8.4/24
    tunnel-interface
      encapsulation ipsec
      color biz-internet
    !
    no shutdown
  !
  ip route 0.0.0.0/0 172.16.8.1
!
```

On the Cisco vSmart Controller and Cisco vBond Orchestrator, you configure a tunnel interface and default IP route to reach the WAN transport. For the tunnel, color has no meaning because these devices have no TLOCs.

```

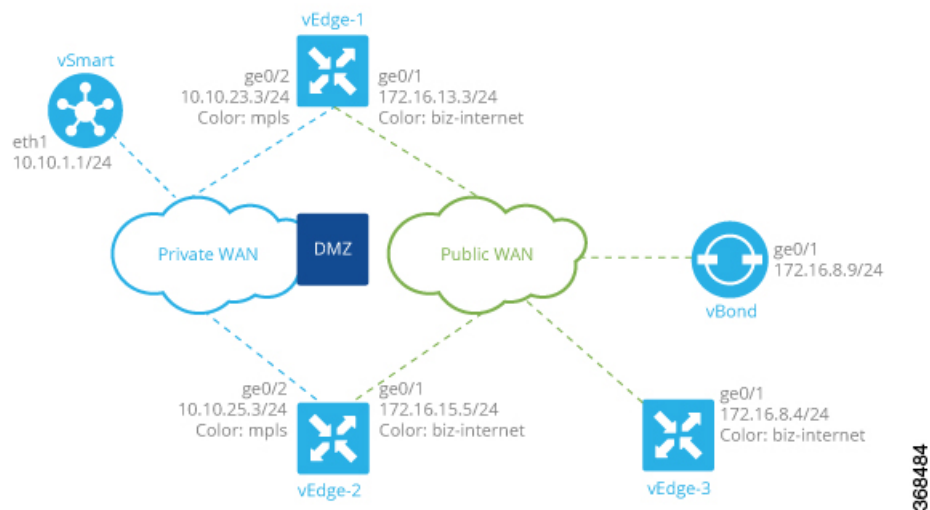
vpn 0
 interface eth1
   ip address 172.16.8.9/24
   tunnel-interface
   !
   no shutdown
   !
 ip route 0.0.0.0/0 172.16.8.1
 !

vpn 0
 interface ge0/1
   ip address 172.16.16.6/24
   tunnel-interface
   !
   no shutdown
   !
 ip route 0.0.0.0/0 172.16.16.1
 !

```

Connect to Public and Private WANs, with Ubiquitous Connectivity to Both WANs

This example is a variant of the previous example. We still have two Cisco vEdge devices at two different sites each connect to the same public WAN (here, the Internet) and the same private WAN (here, an MPLS network). However, now we want sites on the MPLS network and the Internet to be able to exchange data traffic. This topology requires a single overlay over both the public and private WANs. Control connections are present over both transports, and we want IPsec tunnel connections running BFD sessions to exist from private-WAN TLOCs to private-WAN TLOCs, from Internet TLOCs to Internet TLOCs, from private-WAN TLOCs to Internet TLOCs, and from Internet TLOCs to private-WAN TLOCs. This full possibility of TLOCs allows the establishment of a ubiquitous data plane in the overlay network.



For this configuration to work, the Cisco vBond Orchestrator must be reachable over both WAN transports. Because it is on the public WAN (that is, on the Internet), there needs to be connectivity from the private WAN to the Internet. This could be provided via a DMZ, as shown in the figure above. The Cisco vSmart Controller can be either on the public or the private LAN. If there are multiple controllers, some can be on public LAN and others on private LAN.

On each Cisco vEdge device, you configure private-WAN TLOCs, assigning a private color (**metro-ethernet**, **mpls**, or **private1** through **private6**) to the tunnel interface. You also configure public TLOCs, assigning any other color (or you can leave the color as **default**). Each Cisco vEdge device needs two routes to reach the Cisco vBond Orchestrator, one via the private WAN and one via the public WAN.

With such a configuration:

- Control connections are established over each WAN transport.
- BFD/IPsec comes up between all TLOCs (if no policy is configured to change this).
- A given site can be dual-homed to both WAN transports or single-homed to either one.

Here is an example of the configuration on one of the Cisco vEdge devices, vEdge-1:

```
vpn 0
 interface ge0/1
   description "Connection to public WAN"
   ip address 172.16.31.3/24
   tunnel-interface
     encapsulation ipsec
     color biz-internet
   !
   no shutdown
 !
 interface ge0/2
   description "Connection to private WAN"
   ip address 10.10.23.3/24
   tunnel-interface
     encapsulation ipsec
     color mpls
   !
   no shutdown
 !
 ip route 0.0.0.0/0 10.10.23.1
 ip route 0.0.0.0/0 172.16.13.1
 !
```

The **show control connections** command lists two DTLS sessions to the Cisco vSmart Controller, one from the public tunnel (color of **biz-internet**) and one from the private tunnel (color of **mpls**):

```
vEdge-1# show control connections
```

								PEER	
PEER	PEER	PEER		SITE	DOMAIN	PEER	PRIVATE	PEER	
PUBLIC									
TYPE	PROTOCOL	SYSTEM	IP	ID	ID	PRIVATE IP	PORT	PUBLIC IP	
PORT	LOCAL	COLOR	STATE		UPTIME				
vsmart	dtls	1.1.1.9		900	1	172.16.8.2	12346	172.16.8.2	
12346	mpls		up		0:01:41:17				
vsmart	dtls	1.1.1.9		900	1	172.16.8.2	12346	172.16.8.2	
12346	biz-internet		up		0:01:41:33				

The **show bfd sessions** command output shows that vEdge-1 has separate tunnel connections that are running separate BFD sessions for each color:

```
vEdge-1# show bfd sessions
```

			DETECT		SOURCE TLOC		REMOTE TLOC		DST PUBLIC		
DST PUBLIC	SYSTEM IP	PORT	SITE	ID	STATE	COLOR	INTERVAL (msec)	COLOR	UPTIME	SOURCE IP	IP
			ENCAP	MULTIPLIER						TRANSITIONS	
1.1.1.5			500		up	mpls		biz-internet		10.10.23.3	172.16.51.5
12346			ipsec	3			1000		0:06:07:19	1	
1.1.1.5			500		up	biz-internet		biz-internet		172.16.31.3	172.16.51.5
12360			ipsec	3			1000		0:06:07:19	1	

```

1.1.1.6      600    up      mpls      biz-internet  10.10.23.3    172.16.16.6
 12346      ipsec  3      1000     0:06:07:19  1
1.1.1.6      600    up      biz-internet  biz-internet  172.16.31.3   172.16.16.6
 12346      ipsec  3      1000     0:06:07:19  1

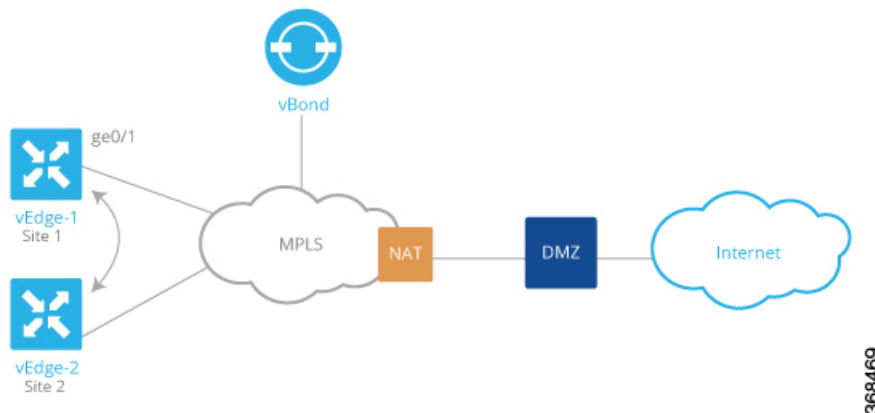
```

Exchange Data Traffic within a Single Private WAN

When the Cisco vEdge device is connected to a private WAN network, such as an MPLS or a metro Ethernet network, and when the carrier hosting the private network does not advertise the router's IP address, remote Cisco vEdge devices on the same private network but at different sites can never learn how to reach that router and hence are not able to exchange data traffic with it by going only through the private network. Instead, the remote routers must route data traffic through a local NAT and over the Internet to a Cisco vBond Orchestrator, which then provides routing information to direct the traffic to its destination. This process can add significant overhead to data traffic exchange, because the Cisco vBond Orchestrator may physically be located at a different site or a long distance from the two Cisco vEdge devices and because it may be situated behind a DMZ.

To allow Cisco vEdge devices at different overlay network sites on the private network to exchange data traffic directly using their private IP addresses, you configure their WAN interfaces to have one of eight private colors, **metro-ethernet**, **mpls**, and **private1** through **private6**. Of these four colors, the WAN interfaces on the Cisco vEdge devices must be marked with the same color so that they can exchange data traffic.

To illustrate the exchange of data traffic across private WANs, let's look at a simple topology in which two Cisco vEdge devices are both connected to the same private WAN. The following figure shows that these two Cisco vEdge devices are connected to the same private MPLS network. The vEdge-1 router is located at Site 1, and vEdge-2 is at Site 2. Both routers are directly connected to PE routers in the carrier's MPLS cloud, and you want both routers to be able to communicate using their private IP addresses.



This topology requires a special configuration to allow traffic exchange using private IP addresses because:

- The Cisco vEdge devices are in different sites; that is, they are configured with different site IDs.
- The Cisco vEdge devices are directly connected to the PE routers in the carrier's MPLS cloud.
- The MPLS carrier does not advertise the link between the Cisco vEdge device and its PE router.

To be clear, if the situation were one of the following, no special configuration would be required:

- vEdge-1 and vEdge-2 are configured with the same site ID.
- vEdge-1 and vEdge-2 are in different sites, and the Cisco vEdge device connects to a CE router that, in turn, connects to the MPLS cloud.

- vEdge-1 and vEdge-2 are in different sites, the Cisco vEdge device connects to the PE router in the MPLS cloud, and the private network carrier advertises the link between the Cisco vEdge device and the PE router in the MPLS cloud.
- vEdge-1 and vEdge-2 are in different sites, and you want them to communicate using their public IP addresses.

In this topology, because the MPLS carrier does not advertise the link between the Cisco vEdge device and the PE router, you use a loopback interface on the each Cisco vEdge device to handle the data traffic instead of using the physical interface that connects to the WAN. Even though the loopback interface is a virtual interface, when you configure it on the Cisco vEdge device, it is treated like a physical interface: the loopback interface is a terminus for both a DTLS tunnel connection and an IPsec tunnel connection, and a TLOC is created for it.

This loopback interface acts as a transport interface, so you must configure it in VPN 0.

For the vEdge-1 and vEdge-2 routers to be able to communicate using their private IP addresses over the MPLS cloud, you set the color of their loopback interfaces to be the same and to one of eight special colors—**metro-ethernet**, **mpls**, and **private1** through **private6**.

Here is the configuration on vEdge-1:

```
vedge-1(config)# vpn 0
vedge-1(config-vpn-0)# interface loopback1
vedge-1(config-interface-loopback1)# ip address 172.16.255.25/32
vedge-1(config-interface-loopback1)# tunnel-interface
vedge-1(config-tunnel-interface)# color mpls
vedge-1(config-interface-tunnel-interface)# exit
vedge-1(config-tunnel-interface)# no shutdown
vedge-1(config-tunnel-interface)# commit and-quit
vedge-1# show running-config vpn 0
...
interface loopback1
 ip-address 172.16.255.25/32
 tunnel-interface
  color mpls
 !
 no shutdown
 !
```

On vEdge-2, you configure a loopback interface with the same tunnel interface color that you used for vEdge-1:

```
vedge-2# show running-config vpn 0
vpn 0
 interface loopback2
 ip address 172.17.255.26/32
 tunnel-interface
  color mpls
 no shutdown
 !
```

Use the **show interface** command to verify that the loopback interface is up and running. The output shows that the loopback interface is operating as a transport interface, so this is how you know that it is sending and receiving data traffic over the private network.

```
vedge-1# show interface
```

TCP		IF	IF					SPEED		
MSS	RX	ADMIN	OPER	ENCAP						
VPN	INTERFACE	IP ADDRESS	STATUS	STATUS	TYPE	PORT	TYPE	MTU	HWADDR	MBPS
DUPLICATE	ADJUST	UPTIME	PACKETS	PACKETS						

```

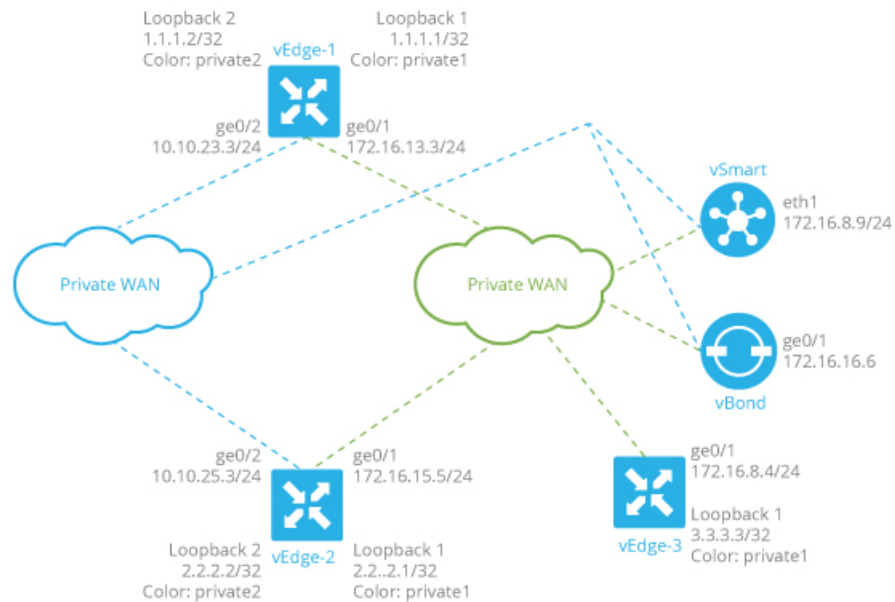
0   ge0/0      10.1.15.15/24   Up    Up    null  transport  1500  00:0c:29:7d:1e:fe  10   full
0   0          0:07:38:49 213199 243908
0   ge0/1      10.1.17.15/24   Up    Up    null  service    1500  00:0c:29:7d:1e:08  10   full
0   0          0:07:38:49 197     3
0   ge0/2      -              Down  Down  null  service    1500  00:0c:29:7d:1e:12  -    -
0   0          -            1
0   ge0/3      10.0.20.15/24   Up    Up    null  service    1500  00:0c:29:7d:1e:1c  10   full
0   0          0:07:38:49 221     27
0   ge0/6      57.0.1.15/24    Up    Up    null  service    1500  00:0c:29:7d:1e:3a  10   full
0   0          0:07:38:49 196     3
0   ge0/7      10.0.100.15/24  Up    Up    null  service    1500  00:0c:29:7d:1e:44  10   full
0   0          0:07:44:47 783     497
0   loopback1  172.16.255.25/32 Up    Up    null  transport  1500  00:00:00:00:00:00  10   full
0   0          0:00:00:20 0        0
0   system     172.16.255.15/32 Up    Up    null  loopback   1500  00:00:00:00:00:00  10   full
0   0          0:07:38:25 0        0
1   ge0/4      10.20.24.15/24   Up    Up    null  service    1500  00:0c:29:7d:1e:26  10   full
0   0          0:07:38:46 27594   27405
1   ge0/5      56.0.1.15/24    Up    Up    null  service    1500  00:0c:29:7d:1e:30  10   full
0   0          0:07:38:46 196     2
512 eth0        10.0.1.15/24    Up    Up    null  service    1500  00:50:56:00:01:05  1000 full
0   0          0:07:45:55 15053   10333

```

To allow Cisco vEdge device at different overlay network sites on the private network to exchange data traffic directly, you use a loopback interface on the each Cisco vEdge device to handle the data traffic instead of using the physical interface that connects to the WAN. You associate the same tag, called a carrier tag, with each loopback interface so that all the routers learn that they are on the same private WAN. Because the loopback interfaces are advertised across the overlay network, the vEdge routers are able to learn reachability information, and they can exchange data traffic over the private network. To allow the data traffic to actually be transmitted out the WAN interface, you bind the loopback interface to a physical WAN interface, specifically to the interface that connects to the private network. Remember that this is the interface that the private network does not advertise. However, it is still capable of transmitting data traffic.

Exchange Data Traffic between Two Private WANs

This example shows a topology with two different private networks, possibly the networks of two different network providers, and all the Cisco SD-WAN devices are located somewhere on one or both of the private networks. Two Cisco vEdge devices are located at two different sites, and they both connect to both private networks. A third Cisco vEdge device connects to only one of the private WANs. The Cisco vBond Orchestrator and Cisco vSmart Controller both sit in one of the private WANs, perhaps in a data center, and they are reachable over both private WANs. For the Cisco vEdge devices to be able to establish control connections, the subnetworks where the Cisco vBond Orchestrator and Cisco vSmart Controller devices reside must be advertised into each private WAN. Each private WAN CPE router then advertises these subnets in its VRF, and each Cisco vEdge device learns those prefixes from each PE router that it is connected to.



368483

Because both WANs are private, we need only a single overlay. In this overlay network, without policy, IPsec tunnels running BFD sessions exist from any TLOC connected to either transport network to any TLOC in the other transport as well as to any TLOC in the same WAN transport network.

As with the previous examples in this topic, it is possible to configure the tunnel interfaces on the routers' physical interfaces. If you do this, you also need to configure a routing protocol between the Cisco vEdge device at its peer PE router, and you need to configure access lists on the Cisco vEdge device to advertise all the routes in both private networks.

A simpler configuration option that avoids the need for access lists is to use loopback interfaces as the tunnel interfaces, and then bind each loopback interface to the physical interface that connects to the private network. Here, the loopback interfaces become the end points of the tunnel, and the TLOC connections in the overlay network run between loopback interfaces, not between physical interfaces. So in the figure shown above, on router vEdge-1, the tunnel connections originate at the Loopback1 and Loopback2 interfaces. This router has two TLOCs: {1.1.1.1, private2, ipsec} and {1.1.1.2, private1, ipsec}.

The WAN interfaces on the Cisco vEdge devices must run a routing protocol with their peer PE routers. The routing protocol must advertise the Cisco vEdge device's loopback addresses to both PE routers so that all Cisco vEdge devices on the two private networks can learn routes to each other. A simple way to advertise the loopback addresses is to redistribute routes learned from other (connected) interfaces on the same router. (You do this instead of creating access lists.) If, for example, you are using OSPF, you can advertise the loopback addresses by including the **redistribute connected** command in the OSPF configuration. Looking at the figure above, the **ge0/2** interface on vEdge-1 needs to advertise both the Loopback1 and Loopback2 interfaces to the blue private WAN, and **ge0/1** must also advertise both these loopback interfaces to the green private WAN.

With this configuration:

- The Cisco vEdge devices learn the routes to the Cisco vBond Orchestrator and Cisco vSmart Controller over each private WAN transport.
- The Cisco vEdge devices learn every other Cisco vEdge device's loopback address over each WAN transport network.

- The end points of the tunnel connections between each pair of Cisco vEdge devices are the loopback interfaces, not the physical (**ge**) interfaces.
- The overlay network has data plane connectivity between any TLOCs and has a control plane over both transport networks.

Here is the interface configuration for VPN 0 on vEdge-1. Highlighted are the commands that bind the loopback interfaces to their physical interfaces. Notice that the tunnel interfaces, and the basic tunnel interface properties (encapsulation and color), are configured on the loopback interfaces, not on the Gigabit Ethernet interfaces.

```
vpn 0
 interface loopback1
   ip address 1.1.1.2/32
   tunnel-interface
     encapsulation ipsec
     color private1
     bind ge0/1
   !
   no shutdown
 !
 interface loopback2
   ip address 1.1.1.1/32
   tunnel-interface
     encapsulation ipsec
     color private2
     bind ge0/2
   !
   no shutdown
 !
 interface ge0/1
   ip address 172.16.13.3/24
   no shutdown
 !
 interface ge0/2
   ip address 10.10.23.3/24
   no shutdown
 !
 ip route 0.0.0.0/0 10.10.23.1
 ip route 0.0.0.0/0 172.16.13.1
```

The configuration for vEdge-2 is similar:

```
vpn 0
 interface loopback1
   ip address 2.2.2.1/32
   tunnel-interface
     encapsulation ipsec
     color private1
     bind ge0/1
   !
   no shutdown
 !
 interface loopback2
   ip address 2.2.2.2/32
   tunnel-interface
     encapsulation ipsec
     color private2
     bind ge0/2
   !
   no shutdown
 !
 interface ge0/1
   ip address 172.16.15.5/24
```

```

    no shutdown
  !
  interface ge0/2
    ip address 10.10.25.3/24
    no shutdown
  !
  ip route 0.0.0.0/0 10.10.25.1
  ip route 0.0.0.0/0 172.16.15.1
  !

```

The vEdge-3 router connects only to the green private WAN:

```

vpn 0
  interface loopback1
    ip address 3.3.3.3/32
    tunnel-interface
      encapsulation ipsec
      color private1
      bind ge0/1
  !
  no shutdown
  !
  interface ge0/1
    ip address 172.16.8.4/24
    no shutdown
  !
  ip route 0.0.0.0/0 172.16.8.1
  !

```

On the Cisco vSmart Controller and Cisco vBond Orchestrator, you configure a tunnel interface and default IP route to reach the WAN transport. For the tunnel, color has no meaning because these devices have no TLOCs.

```


vpn 0
  interface eth1
    ip address 172.16.8.9/24
    tunnel-interface
  !
  no shutdown
  !
  ip route 0.0.0.0/0 172.16.8.1
  !

vpn 0
  interface ge0/1
    ip address 172.16.16.6/24
    tunnel-interface
  !
  no shutdown
  !
  ip route 0.0.0.0/0 172.16.16.1
  !

```

Connect to a WAN Using PPPoE

This example shows a Cisco vEdge device with a TLOC tunnel interface and an interface enabled for Point-to-Point Protocol over Ethernet (PPPoE). The PPP interface defines the authentication method and credentials and is linked to the PPPoE-enabled interface.



368512

Here is the interface configuration for VPN 0:

```
vpn 0
interface ge0/1
  no shutdown
  !
  tunnel-interface
    encapsulation ipsec
    color biz-internet
    allow-service dhcp
    allow-service dns
    allow-service icmp
    no allow-service sshd
    no allow-service ntp
    no allow-service stun
  !
  no shutdown
  !
interface ge0/3
  pppoe-client ppp-interface ppp10
  no shutdown
  !
interface ppp10
  ppp authentication chap
  hostname branch100@corp.bank.myisp.net
  password $4$OHHjdmsC6M8zj4BgLEFXKw==
  !
  tunnel-interface
    encapsulation ipsec
    color gold
    allow-service dhcp
    allow-service dns
    allow-service icmp
    no allow-service sshd
    no allow-service ntp
    no allow-service stun
  !
  no shutdown
  !
```

Use the **show ppp interface** command to view existing PPP interfaces:

```
vEdge# show ppp interface
```

VPN	IFNAME	PPPOE INTERFACE	INTERFACE IP	GATEWAY IP	PRIMARY DNS	SECONDARY DNS	MTU
0	ppp10	ge0/3	11.1.1.1	115.0.1.100	8.8.8.8	8.8.4.4	1150

Use the **show pppoe session** and **show pppoe statistics** commands to view information about PPPoE sessions:

```
vEdge# show pppoe session
```

VPN	SESSION			PPP		SERVICE NAME
	IFNAME	ID	SERVER MAC	LOCAL MAC	INTERFACE	
0	ge0/1	1	00:0c:29:2e:20:1a	00:0c:29:be:27:f5	ppp1	branch100 -
0	ge0/3	1	00:0c:29:2e:20:24	00:0c:29:be:27:13	ppp2	branch100 -

```
vEdge# show pppoe statistics
```

```
pppoe_tx_pkts           : 73
pppoe_rx_pkts           : 39
pppoe_tx_session_drops  : 0
pppoe_rx_session_drops  : 0
pppoe_inv_discovery_pkts : 0
pppoe_ccp_pkts          : 12
pppoe_ipcp_pkts         : 16
pppoe_lcp_pkts          : 35
pppoe_padi_pkts         : 4
pppoe_pado_pkts         : 2
pppoe_padr_pkts         : 2
pppoe_pads_pkts         : 2
pppoe_padt_pkts         : 2
```

Configure VPN Interfaces Using vManage

Configure VPN Ethernet Interface

Configure VPN Ethernet Interface

-
- Step 1** In Cisco vManage, select the **Configuration > Templates** screen.
- Step 2** In the **Device** tab, click **Create Template**.
- Step 3** From the Create Template drop-down, select **From Feature Template**.
- Step 4** From the **Device Model** drop-down, select the type of device for which you are creating the template.
- Step 5** To create a template for VPN 0 or VPN 512:
- Click the **Transport & Management VPN** tab located directly beneath the Description field, or scroll to the Transport & Management VPN section.
 - Under **Additional VPN 0 Templates**, located to the right of the screen, click **VPN Interface**.
 - From the VPN Interface drop-down, click **Create Template**. The **VPN Interface Ethernet** template form displays. The top of the form contains fields for naming the template, and the bottom contains fields for defining VPN Interface Ethernet parameters.
- Step 6** To create a template for VPNs 1 through 511, and 513 through 65530:
- Click the Service VPN tab located directly beneath the Description field, or scroll to the Service VPN section.

- b. Click the **Service VPN** drop-down.
- c. Under Additional VPN templates, located to the right of the screen, click **VPN Interface**.
- d. From the **VPN Interface** drop-down, click **Create Template**. The VPN Interface Ethernet template form displays. The top of the form contains fields for naming the template, and the bottom contains fields for defining VPN Interface Ethernet parameters.

Step 7 In the **Template Name** field, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.

Step 8 In the **Template Description** field, enter a description of the template. The description can be up to 2048 characters and can contain only alphanumeric characters.

Configure Basic Interface Functionality

To configure basic interface functionality in a VPN, choose the **Basic Configuration** tab and configure the following parameters:



Note Parameters marked with an asterisk are required to configure an interface.

Parameter Name	IPv4 or IPv6	Options	Description
Shutdown*			Click No to enable the interface.
Interface name*			Enter a name for the interface.
Description			Enter a description for the interface.
IPv4 / IPv6			Click IPv4 to configure an IPv4 VPN interface. Click IPv6 to configure an IPv6 interface.
Dynamic			Click Dynamic to set the interface as a Dynamic Host Configuration Protocol (DHCP) client, so that the interface receives its IP address from a DHCP server.
	Both	DHCP Distance	Optionally, enter an administrative distance value for routes learned from a DHCP server. Default is 1.
	IPv6	DHCP Rapid Commit	Optionally, configure the DHCP IPv6 local server to support DHCP Rapid Commit, to enable faster client configuration and confirmation in busy environments. Click On to enable DHCP rapid commit Click Off to continue using the regular commit process.
Static			Click Static to enter an IP address that doesn't change.
	IPv4	IPv4 Address	Enter a static IPv4 address.
	IPv6	IPv6 Address	Enter a static IPv6 address.

Parameter Name	IPv4 or IPv6	Options	Description
Secondary IP Address	IPv4		Click Add to enter up to four secondary IPv4 addresses for a service-side interface.
IPv6 Address	IPv6		Click Add to enter up to two secondary IPv6 addresses for a service-side interface.
DHCP Helper	Both		To designate the interface as a DHCP helper on a router, enter up to eight IP addresses, separated by commas, for DHCP servers in the network. A DHCP helper interface forwards BootP (broadcast) DHCP requests that it receives from the specified DHCP servers.
Block Non-Source IP	Yes / No		Click Yes to have the interface forward traffic only if the source IP address of the traffic matches the interface's IP prefix range. Click No to allow other traffic.
Bandwidth Upstream			For Cisco vEdge devices and vManage: For transmitted traffic, set the bandwidth above which to generate notifications. Range: 1 through (232 / 2) – 1 kbps
Bandwidth Downstream			For Cisco vEdge devices and vManage: For received traffic, set the bandwidth above which to generate notifications. Range: 1 through (232 / 2) – 1 kbps

To save the feature template, click **Save**.

CLI Equivalent

```

vpn vpn-id
  interface interface-name
    bandwidth-downstream kbps
    bandwidth-upstream kbps
    block-non-source-ip
    description text
    dhcp-helper ip-address
    (ip address ipv4-prefix/length | ip dhcp-client [dhcp-distance number])
    (ipv6 address ipv6-prefix/length | ipv6 dhcp-client [dhcp-distance number])
  [dhcp-rapid-commit]
  secondary-address ipv4-address
  [no] shutdown

```

Create a Tunnel Interface

On Cisco vEdge devices, you can configure up to four tunnel interfaces. This means that each Cisco vEdge device router can have up to four TLOCs. On Cisco vSmart Controllers and Cisco vManage, you can configure one tunnel interface.

For the control plane to establish itself so that the overlay network can function, you must configure WAN transport interfaces in VPN 0. The WAN interface will enable the flow of tunnel traffic to the overlay. You can add other parameters shown in the table below only after you configure the WAN interface as a tunnel interface.

To configure a tunnel interface, select the **Interface Tunnel** tab and configure the following parameters:

Parameter Name	Cisco vEdge Devices Only	Description
Tunnel Interface	No	Click On to create a tunnel interface.
Color	No	Select a color for the TLOC.
Control Connection	Yes	If the Cisco vEdge device has multiple TLOCs, click No to have the tunnel not establish a TLOC. The default is On , which establishes a control connection for the TLOC.
Maximum Control Connections	Yes	Specify the maximum number of Cisco vSmart Controllers that the WAN tunnel interface can connect to. To have the tunnel establish no control connections, set the number to 0. Range: 0 through 8 Default: 2
Cisco vBond Orchestrator As Stun Server	Yes	Click On to enable Session Traversal Utilities for NAT (STUN) to allow the tunnel interface to discover its public IP address and port number when the Cisco vEdge device router is located behind a NAT.
Exclude Controller Group List	Yes	Set the Cisco vSmart Controllers that the tunnel interface is not allowed to connect to. Range: 0 through 100
vManage Connection Preference	Yes	Set the preference for using a tunnel interface to exchange control traffic with the vManage NMS. Range: 0 through 8 Default: 5
Port Hop	No	Click On to enable port hopping, or click Off to disable it. If port hopping is enabled globally, you can disable it on an individual TLOC (tunnel interface). To control port hopping on a global level, use the System configuration template. Default: Enabled vManage NMS and Cisco vSmart Controller default: Disabled
Low-Bandwidth Link	Yes	Select to characterize the tunnel interface as a low-bandwidth link.
Allow Service	No	Select On or Off for each service to allow or disallow the service on the interface.

To configure additional tunnel interface parameters, click **Advanced Options**:

Parameter Name	Cisco vEdge devices Only	Description
GRE	Yes	Use GRE encapsulation on the tunnel interface. By default, GRE is disabled. If you select both IPsec and GRE encapsulations, two TLOCs are created for the tunnel interface that have the same IP addresses and colors, but that differ by their encapsulation.
IPsec	Yes	Use IPsec encapsulation on the tunnel interface. By default, IPsec is enabled. If you select both IPsec and GRE encapsulations, two TLOCs are created for the tunnel interface that have the same IP addresses and colors, but that differ by their encapsulation.
IPsec Preference	Yes	Specify a preference value for directing traffic to the tunnel. A higher value is preferred over a lower value. Range: 0 through 4294967295 Default: 0
IPsec Weight	Yes	Enter a weight to use to balance traffic across multiple TLOCs. A higher value sends more traffic to the tunnel. Range: 1 through 255 Default: 1
Carrier	No	Select the carrier name or private network identifier to associate with the tunnel. Values: carrier1, carrier2, carrier3, carrier4, carrier5, carrier6, carrier7, carrier8, default Default: default
Bind Loopback Tunnel	Yes	Enter the name of a physical interface to bind to a loopback interface.
Last-Resort Circuit	Yes	Select to use the tunnel interface as the circuit of last resort.
NAT Refresh Interval	No	Enter the interval between NAT refresh packets sent on a DTLS or TLS WAN transport connection. Range: 1 through 60 seconds Default: 5 seconds
Hello Interval	No	Enter the interval between Hello packets sent on a DTLS or TLS WAN transport connection. Range: 100 through 10000 milliseconds Default: 1000 milliseconds (1 second)

Parameter Name	Cisco vEdge devices Only	Description
Hello Tolerance	No	Enter the time to wait for a Hello packet on a DTLS or TLS WAN transport connection before declaring that transport tunnel to be down. Range: 12 through 60 seconds Default: 12 seconds

To save the feature template, click **Save**.

Configure Tunnel Interface CLI on vEdge Devices

```

vpn 0
  interface interface-name
    tunnel-interface
      allow-service service-name
      bind interface-name (on vEdge routers only)
      carrier carrier-name
      color color
      encapsulation (gre | ipsec) (on vEdge routers only)
        preference number
        weight number
      exclude-controller-group-list number (on vEdge routers only)
      hello-interval milliseconds
      hello-tolerance seconds
      last-resort-circuit (on vEdge routers only)
      low-bandwidth-link
      max-control-connections number (on vEdge routers only)
      nat-refresh-interval seconds
      vbond-as-stun-server
      vmanage-connection-preference number (on vEdge routers only)

```

Associate a Carrier Name with a Tunnel Interface

To associate a carrier name or private network identifier with a tunnel interface, use the **carrier** command. *carrier-name* can be **default** and **carrier1** through **carrier8**:

```

vEdge(config)# vpn 0
vEdge(config-vpn-0)# interface interface-name
vEdge(config-interface)# tunnel-interface
vEdge(config-tunnel-interface)# carrier carrier-name

```

Limit Keepalive Traffic on a Tunnel Interface

By default, Cisco vEdge devices send a Hello packet once per second to determine whether the tunnel interface between two devices is still operational and to keep the tunnel alive. The combination of a hello interval and a hello tolerance determines how long to wait before declaring a DTLS or TLS tunnel to be down. The default hello interval is 1 second, and the default tolerance is 12 seconds. With these default values, if no Hello packet is received within 11 seconds, the tunnel is declared down at 12 seconds.

If the hello interval or the hello tolerance, or both, are different at the two ends of a DTLS or TLS tunnel, the tunnel chooses the interval and tolerance as follows:

- For a tunnel connection between two controller devices, the tunnel uses the lower hello interval and the higher tolerance interval for the connection between the two devices. (Controller devices are vBond

controllers, vManage NMSs, and vSmart controllers.) This choice is made in case one of the controllers has a slower WAN connection. The hello interval and tolerance times are chosen separately for each pair of controller devices.

- For a tunnel connection between a Cisco vEdge device and any controller device, the tunnel uses the hello interval and tolerance times configured on the router. This choice is made to minimize the amount of traffic sent over the tunnel, to allow for situations where the cost of a link is a function of the amount of traffic traversing the link. The hello interval and tolerance times are chosen separately for each tunnel between a Cisco vEdge device and a controller device.

To minimize the amount of keepalive traffic on a tunnel interface, increase the Hello packet interval and tolerance on the tunnel interface:

```
vEdge(config-tunnel-interface)# hello-interval milliseconds
vEdge(config-tunnel-interface)# hello-tolerance seconds
```

The default hello interval is 1000 milliseconds, and it can be a time in the range 100 through 600000 milliseconds (10 minutes). The default hello tolerance is 12 seconds, and it can be a time in the range 12 through 600 seconds (10 minutes). The hello tolerance interval must be at most one-half the OMP hold time. The default OMP hold time is 60 seconds, and you configure it with the **omp timers holdtime** command.

Configure Multiple Tunnel Interfaces on a vEdge Router

On a Cisco vEdge device, you can configure up to eight tunnel interfaces in the transport interface (VPN 0). This means that each Cisco vEdge device can have up to eight TLOCs.

When a Cisco vEdge device has multiple TLOCs, each TLOC is preferred equally and traffic to each TLOC is weighted equally, resulting in ECMP routing. ECMP routing is performed regardless of the encapsulation used on the transport tunnel, so if, for example, a router has one IPsec and one GRE tunnel, with ECMP traffic is forwarded equally between the two tunnels. You can change the traffic distribution by modifying the preference or the weight, or both, associated with a TLOC. (Note that you can also affect or change the traffic distribution by applying a policy on the interface that affects traffic flow.)

```
vEdge(config)# vpn 0
vEdge(config-vpn-0)# interface interface-name
vEdge(config-tunnel-interface) encapsulation (gre | ipsec)
vEdge(config-encapsulation)# preference number
vEdge(config-encapsulation)# weight number
```

The **preference** command controls the preference for directing inbound and outbound traffic to a tunnel. The preference can be a value from 0 through 4294967295 ($2^{32} - 1$), and the default value is 0. A higher value is preferred over a lower value.

When a Cisco vEdge device has two or more tunnels, if all the TLOCs have the same preference and no policy is applied that affects traffic flow, all the TLOCs are advertised into OMP. When the router transmits or receives traffic, it distributes traffic flows evenly among the tunnels, using ECMP.

When a Cisco vEdge device has two or more tunnels, if the TLOCs have different preferences and a policy that affects traffic flow is not applied, all the TLOCs are advertised to Cisco vSmart Controller via OMP for further processing based on the control policy applied on Cisco vSmart Controller for the corresponding vEdge site-id. When the router transmits or receives traffic, it sends traffic to or receives traffic from only the TLOC with the highest preference. When there are three or more tunnels and two of them have the same preference, traffic flows are distributed evenly between these two tunnels.

A remote Cisco vEdge device trying to reach one of these prefixes selects which TLOC to use from the set of TLOCs that have been advertised. So, for example, if a remote router selects a GRE TLOC on the local router, the remote router must have its own GRE TLOC to be able to reach the prefix. If the remote router

has no GRE TLOC, it is unable to reach the prefix. If the remote router has a single GRE TLOC, it selects that tunnel even if there is an IPsec TLOC with a higher preference. If the remote router has multiple GRE TLOCs, it selects from among them, choosing the one with the highest preference or using ECMP among GRE TLOCs with equal preference, regardless of whether there is an IPsec TLOC with a higher preference.

The **weight** command controls how traffic is balanced across multiple TLOCs that have equal preference values. The weight can be a value from 1 through 255, and the default is 1. When the weight value is higher, the router sends more traffic to the TLOC. You typically set the weight based on the bandwidth of the TLOC. When a router has two or more TLOCs, all with the highest equal preference value, traffic distribution is weighted according to the configured weight value. For example, if TLOC A has weight 10, and TLOC B has weight 1, and both TLOCs have the same preference value, then roughly 10 flows are sent out TLOC A for every 1 flow sent out TLOC B.

Configure an Interface as a NAT Device

You can configure IPv4 and IPv6 interfaces to act as a network address translation (NAT) device for applications such as port forwarding. To configure a NAT device:

1. In the **VPN Interface Ethernet Template**, click the **NAT** tab, and select either **IPv4** or **IPv6**.
2. Change the scope from Default (blue check) to **Global** (green globe).
3. Click **On** to enable NAT (IPv4) or NAT64 (IPv6). The correct set of parameters will display.
4. Enter the parameter values.
5. To save the feature template, click **Save**.



Note Optionally, click either **Port Forward** or **Static NAT** to enable those parameters.

IPv4 NAT Parameter Values

Table 32: IPv4 NAT Parameter Values

Parameter Name	Description
Refresh Mode	Select how NAT mappings are refreshed, either outbound or bidirectional (outbound and inbound). Default: Outbound
Log NAT flow creations or deletions	Enable logging when NAT flows are created or deleted. Default: Off 1. Change the scope from Default to Global . 2. Click On .
UDP Timeout	Specify when NAT translations over UDP sessions time out. Range: 1 through 65536 minutes Default: 1 minutes

Parameter Name	Description
TCP Timeout	Specify when NAT translations over TCP sessions time out. Range: 1 through 65536 minutes Default: 60 minutes (1 hour)
Block ICMP	Select On to block inbound ICMP error messages. By default, an acting as a NAT device receives these error messages. Default: Off
Respond to Ping	Select On to have the device respond to ping requests to the NAT interface's IP address that are received from the public side of the connection.
NAT Pool Range Start	Enter a starting IP address for the NAT pool. 1. Change the scope from Default to Global to enable the field. 2. Enter the starting IP address for the NAT pool.
NAT Pool Range End	Enter a closing IP address for the NAT pool. 1. Change the scope from Default to Global to enable the field. 2. Enter the last IP address for the NAT pool.

Configure Static NAT

To configure a static NAT of service-side source IP addresses:

1. In the **VPN Interface Ethernet Template**, click the **NAT** tab, and select either **IPv4** or **IPv6**.

Click the **Static NAT** tab. Click **New Static NAT** and configure the following parameters to add a static NAT mapping:

Table 33:

Parameter Name	Description
Mark as Optional Row	Check Mark as Optional Row to mark this configuration as device-specific. To include this configuration for a device, enter the requested variable values when you attach a device template to a device, or create a template variables spreadsheet to apply the variables.
Source IP	Enters the NAT private source IP address.
Translated Source IP Address	Maps a public IP address to a private source address, enter the public IP address.
Static NAT Direction	Selects the direction in which to perform network address translation.
inside	Translates the IP address of packets that are coming from the service side of the device and that are destined for the transport side of the router.

Parameter Name	Description
outside	Translates the IP address of packets that are coming to the device from the transport side device and that are destined for a service-side device.
Source VPN ID	Configures Source VPN ID

- To save the NAT mapping, click **Add**.
- To save the feature template, click **Save**.

Configure IPv4 NAT CLI Equivalent on vEdge

CLI Equivalent

```
vpn vpn-id
interface interface-name
  nat
    block-icmp-error
    refresh (bi-directional | outbound)
    respond-to-ping
    tcp-timeout minutes
    udp-timeout minutes
```

IPv6 NAT Parameter Values

Table 34: IPv4 NAT Parameter Values

Parameter Name	Description
UDP Timeout	Enter the timeout value for User Datagram Protocol (UDP) traffic <ol style="list-style-type: none"> Change the scope from Default to Global. Enter a timeout value. Range: 1–536870 seconds Default: 1 second
TCP Timeout	Enter the timeout value for Transmission Control Protocol (TCP) traffic. <ol style="list-style-type: none"> Change the scope from Default to Global. Enter a timeout value. Enter a timeout value. Default: 60 seconds

Configure NAT64 CLI Equivalent on Cisco vEdge Device

CLI Equivalent

```
interface interface-name
nat64 enable
  tcp-timeout minutes
  udp-timeout minutes
```


VPN Interface NAT Pool using Cisco vManage

Create NAT Pool Interfaces in a VPN

Use the **VPN Interface NATPool** template for Cisco vEdge devices and , to create Network Address Translation (NAT) pools of IP addresses in virtual private networks (VPNs). To configure NAT pool interfaces in a VPN using Cisco vManage templates:

1. Create a **VPN Interface NATPool** template for Cisco vEdge devices to configure Ethernet interface parameters, as described in this article.
2. Create a VPN feature template to configure parameters for a service-side VPN.
3. Optionally, create a data policy to direct data traffic to a service-side NAT.

Create a VPN Interface NAT Pool Template

You can open a new **VPN Interface NATPool** template for Cisco vEdge devices from the VPN section of a device template.

1. From the vManage menu, select **Configuration > Templates**
2. Click **Feature**.
3. Click **Add Template**.
4. Select a device from the list.
5. From the VPN section, click **VPN Interface NATPool**.

The VPN Interface Ethernet template form displays. The top of the form contains fields for naming the template, and the bottom contains fields for defining VPN Interface NAT Pool parameters.

1. In the required **Template Name** field, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
2. In the optional **Template Description** field, enter a description of the template. The description can be up to 2048 characters and can contain only alphanumeric characters.

Parameter Menus and Options

Parameter Menus and Options

When you first open a feature template, for each parameter that has a default value, the scope is set to Default (indicated by a



), and the default setting or value is shown. To change the default or to enter a value, click the scope drop-down to the left of the parameter field and select the appropriate option.

Configure a NAT Pool Interface

To configure a NAT pool interface, configure the following parameters. Parameters marked with an asterisk are required to configure the interface.

Basic Configuration

Enter the following basic configuration parameters:

Table 35:

Parameter Name	Description
Shutdown*	Yes Click No to enable the interface. No
Interface Name (1...31)*	Enter a number for the NAT pool interface to use for service-side NAT. For example, <i>natpool22</i> . Range: 1-31
Description	Enter a description for the interface.
IPv4 Address*	Enter the IPv4 address of the interface. The address length determines the number of NAT addresses that the router use at the same time. A Cisco vEdge device router can support a maximum of 250 NAT IP addresses.
Refresh Mode	Select how NAT mappings are refreshed:
bi-directional	Keep active the NAT mappings for inbound and outbound traffic.
outbound	Keep active the NAT mappings for outbound traffic. This is the default.
UDP Timeout	Enter the time when NAT translations over UDP sessions time out. <i>Default: 1 minute</i> Range: 1-65536 minutes
TCP Timeout	Enter the time when NAT translations over TCP sessions time out. <i>Default: 60 minutes (1 hour)</i> Range: 1-65536 minutes
Block ICMP	Select whether a Cisco vEdge device that is acting as a NAT device should receive inbound ICMP error messages. By default, the router blocks these error messages. Click Off to receive the ICMP error messages.
Direction	Select the direction in which the NAT interface performs address translation:
inside	Translate the source IP address of packets that are coming from the service side of the Cisco vEdge device and that are destined to transport side of the router. This is the default.
outside	Translate the source IP address of packets that are coming to the Cisco vEdge device from the transport side of the Cisco vEdge device and that are destined to a service-side device.
Overload	Click No to disable dynamic NAT. By default, dynamic NAT is enabled.

Configure a Tracker Interface

1. To create one or more tracker interfaces, select the **Tracker** tab and click **New Tracker**.
2. Select one or more interfaces to track the status of service interfaces.
3. To save the tracker interfaces, click **Add**. To save the feature template, click **Save**.

NAT Pool Interface CLI Equivalent Commands on Cisco vEdge Devices

Use the following commands to configure NAT Pool interfaces on Cisco vEdge devices.

```
vpn vpn-id
  interface natpoolnumber
    ip address prefix/length
    nat
      tracker tracker-name1
        tracker-name2, tracker-name3
    direction (inside | outside)
    [no] overload
    refresh (bi-directional | outbound)
    static source-ip ip-address1 translate-ip ip-address2 (inside | outside)
    tcp-timeout minutes
    udp-timeout minutes
    [no] shutdown
```

Configure Port-Forwarding Rules

To create port-forwarding rules to allow requests from an external network to reach devices on the internal network:

1. Select the **Port Forward** tab.
2. Click **New Port Forwarding Rule**, and configure the parameters. You can create up to 128 rules.
3. To save the rule, click **Add**.
4. To save the feature template, click **Save**.

Table 36:

Parameter Name	Values	Description
Port Start Range	Enter the starting port number. This number must be less than or equal to the ending port number.	
Port End Range	Enter the ending port number. To apply port forwarding to a single port, specify the same port number for the starting and ending numbers. When applying port forwarding to a range of ports, the range includes the two port numbers that you specify.	
Protocol	TCP UDP	

Parameter Name	Values	Description
VPN	0-65535	Private VPN in which the internal server resides.
Private IP	Enter an IP address to use within the firewall. A best practice is to specify the IP address of a service-side VPN.	

Port Forwarding CLI Equivalent for vEdge

```
vpn vpn-id
interface natpoolnumber
nat
port-forward port-start port-number1 port-end port-number2 proto (tcp | udp)
private-ip-address ip address private-vpn vpn-id
```

Static NAT CLI Equivalent Commands on Cisco vEdge Device

```
vpn vpn-id
interface natpoolnumber
nat
port-forward port-start port-number1 port-end port-number2 proto (tcp | udp)
private-ip-address ip address private-vpn vpn-id
```

Release Information

Introduced in Cisco vManage NMS Release 16.3. In Release 17.2.2, add support for tracker interface status. In Release 18.4, updated images; add support for multiple tracker interfaces.

Apply Access Lists and QoS Parameters

Quality of service (QoS) helps determine how a service will perform. By configuring QoS, enhance the performance of an application on the WAN. To configure a shaping rate for an interface and to apply a QoS map, a rewrite rule, access lists, and policers to a interface, select the ACL/QoS tab and configure the following parameters:

Parameter Name	Description
Shaping rate	Configure the aggregate traffic transmission rate on the interface to be less than line rate, in kilobits per second (kbps).
QoS Map	Specify the name of the QoS map to apply to packets being transmitted out the interface.
Rewrite Rule	Click On , and specify the name of the rewrite rule to apply on the interface.
Ingress ACL – IPv4	Click On , and specify the name of the access list to apply to IPv4 packets being received on the interface.
Egress ACL – IPv4	Click On , and specify the name of the access list to apply to IPv4 packets being transmitted on the interface.

Parameter Name	Description
Ingress ACL – IPv6	Click On , and specify the name of the access list to apply to IPv6 packets being received on the interface.
Egress ACL – IPv6	Click On , and specify the name of the access list to apply to IPv6 packets being transmitted on the interface.
Ingress Policer	Click On , and specify the name of the policer to apply to packets received on the interface.
Egress Policer	Click On , and specify the name of the policer to apply to packets being transmitted on the interface.

To save the feature template, click **Save**.

CLI Equivalent

```
vpn vpn-id
  interface interface-name
    access-list acl-list (in | out)
    policer policer-name (in |out)
    qos-map name
    rewrite-rule name
    shaping-rate name
```

Add ARP Table Entries

The Address Resolution Protocol (ARP) helps associate a link layer address (such as the MAC address of a device) to its assigned internet layer address. Configure a static ARP address when dynamic mapping is not functional. To configure static ARP table entries on the interface, select the ARP tab. Then click **Add New ARP** and configure the following parameters:

Parameter Name	Description
IP Address	Enter the IP address for the ARP entry in dotted decimal notation or as a fully qualified host name.
MAC Address	Enter the MAC address in colon-separated hexadecimal notation.

To save the ARP configuration, click **Add**.

To save the feature template, click **Save**.

CLI Equivalent

```
vpn vpn-id
  interface interface-name arp ip ip-address mac mac-address
```

VPN Interface Bridge

Use the VPN Interface Bridge template for all Cisco vEdge device Cloud and Cisco vEdge devices.

Integrated routing and bridging (IRB) allows Cisco vEdge devices in different bridge domains to communicate with each other. To enable IRB, create logical IRB interfaces to connect a bridge domain to a VPN. The VPN provides the Layer 3 routing services necessary so that traffic can be exchanged between different VLANs.

Each bridge domain can have a single IRB interface and can connect to a single VPN, and a single VPN can connect to multiple bridge domains on a Cisco vEdge device.

To configure a bridge interface using Cisco vManage templates:

1. Create a VPN Interface Bridge feature template to configure parameters for logical IRB interfaces, as described in this article.
2. Create a Bridge feature template for each bridging domain, to configure the bridging domain parameters. See the Bridge help topic.

Navigate to the Template Screen and Name the Template

1. In Cisco vManage NMS, select the **Configuration > Templates** screen.
2. In the Device tab, click **Create Template**.
3. From the Create Template drop-down, select **From Feature Template**.
4. From the **Device Model** drop-down, select the type of device for which you are creating the template.
5. Click the **Service VPN** tab located directly beneath the Description field, or scroll to the **Service VPN** section.
6. Click the Service VPN drop-down.
7. Under Additional VPN Templates, located to the right of the screen, click VPN Interface Bridge.
8. From the VPN Interface Bridge drop-down, click Create Template. The VPN Interface Bridge template form is displayed. The top of the form contains fields for naming the template, and the bottom contains fields for defining VPN Interface Bridge parameters.
9. In the Template Name field, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
10. In the Template Description field, enter a description of the template. The description can be up to 2048 characters and can contain only alphanumeric characters.

When you first open a feature template, for each parameter that has a default value, the scope is set to Default (indicated by a check mark), and the default setting or value is shown. To change the default or to enter a value, click the scope drop-down to the left of the parameter field and select one of the following:

Table 37:

Parameter Scope	Scope Description
Device Specific (indicated by a host icon)	<p>Use a device-specific value for the parameter. For device-specific parameters, you cannot enter a value in the feature template. You enter the value when you attach a Viptela device to a device template .</p> <p>When you click Device Specific, the Enter Key box opens. This box displays a key, which is a unique string that identifies the parameter in a CSV file that you create. This file is an Excel spreadsheet that contains one column for each key. The header row contains the key names (one key per column), and each row after that corresponds to a device and defines the values of the keys for that device. You upload the CSV file when you attach a Viptela device to a device template. For more information, see Create a Template Variables Spreadsheet .</p> <p>To change the default key, type a new string and move the cursor out of the Enter Key box.</p> <p>Examples of device-specific parameters are system IP address, hostname, GPS location, and site ID.</p>
Global (indicated by a globe icon)	<p>Enter a value for the parameter, and apply that value to all devices.</p> <p>Examples of parameters that you might apply globally to a group of devices are DNS server, syslog server, and interface MTUs.</p>

Release Information

Introduced in Cisco vManage NMS in Release 15.3. In Release 18.2, add support for disabling ICMP redirect messages.

VPN Interface Ethernet PPPoE

Use the PPPoE template for Cisco XE SD-WAN devices.

You configure PPPoE over GigabitEthernet interfaces on Cisco IOS XE routers, to provide PPPoE client support.

To configure interfaces on Cisco routers using Cisco vManage templates:

1. Create a VPN Interface Ethernet PPPoE feature template to configure Ethernet PPPoE interface parameters, as described in this article.
2. Create a VPN feature template to configure VPN parameters. See the VPN help topic.

Navigate to the Template Screen and Name the Template

1. In Cisco vManage NMS, select the Configuration ► Templates screen.
2. In the Device tab, click Create Template.
3. From the Create Template drop-down, select "From Feature Template."
4. From the Device Model drop-down, select the type of device for which you are creating the template.

5. Click the Transport & Management VPN tab located directly beneath the Description field, or scroll to the Transport & Management VPN section.
6. Under Additional VPN 0 Templates, located to the right of the screen, click VPN Interface Ethernet PPPoE.
7. From the VPN Interface Ethernet PPPoE drop-down, click Create Template. The VPN Interface Ethernet PPPoE template form is displayed. The top of the form contains fields for naming the template, and the bottom contains fields for defining Ethernet PPPoE parameters.



8. In the Template Name field, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
9. In the Template Description field, enter a description of the template. The description can be up to 2048 characters and can contain only alphanumeric characters.

When you first open a feature template, for each parameter that has a default value, the scope is set to Default (indicated by a check mark), and the default setting or value is shown. To change the default or to enter a value, click the scope drop-down to the left of the parameter field and select one of the following:

Table 38:

Parameter Scope	Scope Description
Device Specific (indicated by a host icon)	<p>Use a device-specific value for the parameter. For device-specific parameters, you cannot enter a value in the feature template. You enter the value when you attach a Cisco SD-WAN device to a device template .</p> <p>When you click Device Specific, the Enter Key box opens. This box displays a key, which is a unique string that identifies the parameter in a CSV file that you create. This file is an Excel spreadsheet that contains one column for each key. The header row contains the key names (one key per column), and each row after that corresponds to a device and defines the values of the keys for that device. You upload the CSV file when you attach a Cisco SD-WAN device to a device template. For more information, see Create a Template Variables Spreadsheet .</p> <p>To change the default key, type a new string and move the cursor out of the Enter Key box.</p> <p>Examples of device-specific parameters are system IP address, hostname, GPS location, and site ID.</p>
Global (indicated by a globe icon)	<p>Enter a value for the parameter, and apply that value to all devices.</p> <p>Examples of parameters that you might apply globally to a group of devices are DNS server, syslog server, and interface MTUs.</p>

Configure PPPoE Functionality

To configure basic PPPoE functionality, select the Basic Configuration tab and configure the following parameters. Required parameters are indicated with an asterisk.

Table 39:

Parameter Name	Description
Shutdown*	Click No to enable the GigabitEthernet interface.
Ethernet Interface Name	Enter the name of a GigabitEthernet interface. For IOS XE routers, you must spell out the interface names completely (for example, GigabitEthernet0/0/0).
VLAN ID	VLAN tag of the sub-interface.
Description	Enter a description of the Ethernet-PPPoE-enabled interface.
Dialer Pool Member	Enter the number of the dialer pool to which the interface belongs. <i>Range:</i> 100 to 255.
PPP Maximum Payload	Enter the maximum receive unit (MRU) value to be negotiated during PPP Link Control Protocol (LCP) negotiation. <i>Range:</i> 64 through 1792 bytes

To save the feature template, click Save.

Configure the PPP Authentication Protocol

To configure the PPP Authentication Protocol, select the PPP tab and configure the following parameters. Required parameters are indicated with an asterisk.

Table 40:

Parameter Name	Description
PPP Authentication Protocol	Select the authentication protocol used by the MLP: <ul style="list-style-type: none"> • CHAP—Enter the hostname and password provided by your Internet Service Provider (ISP). <i>hostname</i> can be up to 255 characters. • PAP—Enter the username and password provided by your ISP. <i>username</i> can be up to 255 characters. • PAP and CHAP—Configure both authentication protocols. Enter the login credentials for each protocol. To use the same username and password for both, click Same Credentials for PAP and CHAP.

To save the feature template, click Save.

Create a Tunnel Interface

On IOS XE routers, you can configure up to four tunnel interfaces. This means that each router can have up to four TLOCs.

For the control plane to establish itself so that the overlay network can function, you must configure WAN transport interfaces in VPN 0.

To configure a tunnel interface for the multilink interface, select the Tunnel Interface tab and configure the following parameters:

Table 41:

Parameter Name	Description
Tunnel Interface	Click On to create a tunnel interface.
Color	Select a color for the TLOC.
Control Connection	If the router has multiple TLOCs, click No to have the tunnel not establish a TLOC. The default is On, which establishes a control connection for the TLOC.
Maximum Control Connections	Specify the maximum number of Cisco vSmart Controllers that the WAN tunnel interface can connect to. To have the tunnel establish no control connections, set the number to 0. <i>Range: 0 through 8Default: 2</i>
Cisco vBond Orchestrator As STUN Server	Click On to enable Session Traversal Utilities for NAT (STUN) to allow the tunnel interface to discover its public IP address and port number when the router is located behind a NAT.
Exclude Controller Group List	Set the Cisco vSmart Controllers that the tunnel interface is not allowed to connect to. <i>Range: 0 through 100</i>
Cisco vManage Connection Preference	Set the preference for using a tunnel interface to exchange control traffic with the Cisco vManage NMS. <i>Range: 0 through 8Default: 5</i>
Port Hop	Click On to enable port hopping, or click Off to disable it. When a router is behind a NAT, port hopping rotates through a pool of preselected OMP port numbers (called base ports) to establish DTLS connections with other routers when a connection attempt is unsuccessful. The default base ports are 12346, 12366, 12386, 12406, and 12426. To modify the base ports, set a port offset value. <i>Default: Enabled</i>
Low-Bandwidth Link	Select to characterize the tunnel interface as a low-bandwidth link.
Allow Service	Select On or Off for each service to allow or disallow the service on the interface.

To configure additional tunnel interface parameters, click Advanced Options and configure the following parameters:

Table 42:

Parameter Name	Description
GRE	Use GRE encapsulation on the tunnel interface. By default, GRE is disabled. If you select both IPsec and GRE encapsulations, two TLOCs are created for the tunnel interface that have the same IP addresses and colors, but that differ by their encapsulation.

Parameter Name	Description
IPsec	Use IPsec encapsulation on the tunnel interface. By default, IPsec is enabled. If you select both IPsec and GRE encapsulations, two TLOCs are created for the tunnel interface that have the same IP addresses and colors, but that differ by their encapsulation.
IPsec Preference	Specify a preference value for directing traffic to the tunnel. A higher value is preferred over a lower value. <i>Range:</i> 0 through 4294967295 <i>Default:</i> 0
IPsec Weight	Enter a weight to use to balance traffic across multiple TLOCs. A higher value sends more traffic to the tunnel. <i>Range:</i> 1 through 255 <i>Default:</i> 1
Carrier	Select the carrier name or private network identifier to associate with the tunnel. <i>Values:</i> carrier1, carrier2, carrier3, carrier4, carrier5, carrier6, carrier7, carrier8, default <i>Default:</i> default
Bind Loopback Tunnel	Enter the name of a physical interface to bind to a loopback interface.
Last-Resort Circuit	Select to use the tunnel interface as the circuit of last resort.
NAT Refresh Interval	Enter the interval between NAT refresh packets sent on a DTLS or TLS WAN transport connection. <i>Range:</i> 1 through 60 seconds <i>Default:</i> 5 seconds
Hello Interval	Enter the interval between Hello packets sent on a DTLS or TLS WAN transport connection. <i>Range:</i> 100 through 10000 milliseconds <i>Default:</i> 1000 milliseconds (1 second)
Hello Tolerance	Enter the time to wait for a Hello packet on a DTLS or TLS WAN transport connection before declaring that transport tunnel to be down. <i>Range:</i> 12 through 60 seconds <i>Default:</i> 12 seconds

Configure the Interface as a NAT Device

To configure an interface to act as a NAT device for applications such as port forwarding, select the NAT tab, click On and configure the following parameters:

Table 43:

Parameter Name	Description
NAT	Click On to have the interface act as a NAT device.
Refresh Mode	Select how NAT mappings are refreshed, either outbound or bidirectional (outbound and inbound). <i>Default:</i> Outbound

UDP Timeout	Specify when NAT translations over UDP sessions time out. <i>Range:</i> 1 through 65536 minutes <i>Default:</i> 1 minutes
TCP Timeout	Specify when NAT translations over TCP sessions time out. <i>Range:</i> 1 through 65536 minutes <i>Default:</i> 60 minutes (1 hour)
Block ICMP	Select On to block inbound ICMP error messages. By default, a router acting as a NAT device receives these error messages. <i>Default:</i> Off
Respond to Ping	Select On to have the router respond to ping requests to the NAT interface's IP address that are received from the public side of the connection.

To create a port forwarding rule, click Add New Port Forwarding Rule and configure the following parameters. You can define up to 128 port-forwarding rules to allow requests from an external network to reach devices on the internal network.

Table 44:

Parameter Name	Description
Port Start Range	Enter a port number to define the port or first port in the range of interest. <i>Range:</i> 0 through 65535
Port End Range	Enter the same port number to apply port forwarding to a single port, or enter a larger number to apply it to a range of ports. <i>Range:</i> 0 through 65535
Protocol	Select the protocol to which to apply the port-forwarding rule, either TCP or UDP. To match the same ports for both TCP and UDP traffic, configure two rules.
VPN	Specify the private VPN in which the internal server resides. This VPN is one of the VPN identifiers in the overlay network. <i>Range:</i> 0 through 65530
Private IP	Specify the IP address of the internal server to which to direct traffic that matches the port-forwarding rule.

To save a port forwarding rule, click Add.

To save the feature template, click Save.

Apply Access Lists

To apply a rewrite rule, access lists, and policers to a router interface, select the ACL tab and configure the following parameters:

Table 45:

Parameter Name	Description
Shaping rate	Configure the aggregate traffic transmission rate on the interface to be less than line rate, in kilobits per second (kbps).
QoS map	Specify the name of the QoS map to apply to packets being transmitted out the interface.

Parameter Name	Description
Rewrite Rule	Click On, and specify the name of the rewrite rule to apply on the interface.
Ingress ACL – IPv4	Click On, and specify the name of the access list to apply to IPv4 packets being received on the interface.
Egress ACL – IPv4	Click On, and specify the name of the access list to apply to IPv4 packets being transmitted on the interface.
Ingress ACL – IPv6	Click On, and specify the name of the access list to apply to IPv6 packets being received on the interface.
Egress ACL – IPv6	Click On, and specify the name of the access list to apply to IPv6 packets being transmitted on the interface.
Ingress Policer	Click On, and specify the name of the policer to apply to packets being received on the interface.
Egress Policer	Click On, and specify the name of the policer to apply to packets being transmitted on the interface.

To save the feature template, click Save.

Configure Other Interface Properties

To configure other interface properties, select the Advanced tab and configure the following properties:

Table 46:

Parameter Name	Description
Bandwidth Upstream	For transmitted traffic, set the bandwidth above which to generate notifications. <i>Range:</i> 1 through $(2^{32} / 2) - 1$ kbps
Bandwidth Downstream	For received traffic, set the bandwidth above which to generate notifications. <i>Range:</i> 1 through $(2^{32} / 2) - 1$ kbps
IP MTU	Specify the maximum MTU size of packets on the interface. <i>Range:</i> 576 through 1804 <i>Default:</i> 1500 bytes
TCP MSS	Specify the maximum segment size (MSS) of TCP SYN packets passing through the router. By default, the MSS is dynamically adjusted based on the interface or tunnel MTU such that TCP SYN packets are never fragmented. <i>Range:</i> 552 to 1460 bytes <i>Default:</i> None
TLOC Extension	Enter the name of the physical interface on the same router that connects to the WAN transport circuit. This configuration then binds this service-side interface to the WAN transport. A second router at the same site that itself has no direct connection to the WAN (generally because the site has only a single WAN connection) and that connects to this service-side interface is then provided with a connection to the WAN.

Parameter Name	Description
Tracker	Enter the name of a tracker to track the status of transport interfaces that connect to the internet.
IP Directed-Broadcast	Enables translation of a directed broadcast to physical broadcasts. An IP directed broadcast is an IP packet whose destination address is a valid broadcast address for some IP subnet but which originates from a node that is not itself part of that destination subnet.

To save the feature template, click Save.

Release Information

Introduced in Cisco vManage NMS in Release 18.4.1.

VPN Interface GRE

Use the VPN Interface GRE template for all vEdge Cloud router and Cisco vEdge devices.

When a service, such as a firewall, is available on a device that supports only GRE tunnels, you can configure a GRE tunnel on the Cisco vEdge device to connect to the remote device by configuring a logical GRE interface. You then advertise that the service is available via a GRE tunnel, and you create data policies to direct the appropriate traffic to the tunnel. GRE interfaces come up as soon as they are configured, and they stay up as long as the physical tunnel interface is up.

To configure GRE interfaces using Cisco vManage templates:

1. Create a VPN Interface GRE feature template to configure a GRE interface, as described in this article.
2. Create a VPN feature template to advertise a service that is reachable via a GRE tunnel, to configure GRE-specific static routes, and to configure other VPN parameters. See the VPN help topic.
3. Create a data policy on the Cisco vSmart Controller controller that applies to the service VPN, including a **set service *service-name* local** command. See the Policies help topic.

Navigate to the Template Screen and Name the Template

1. In Cisco vManage NMS, select the Configuration ► Templates screen.
2. In the Device tab, click Create Template.
3. From the Create Template drop-down, select From Feature Template.
4. From the Device Model drop-down, select the type of device for which you are creating the template.
5. To create a template for VPN 0 or VPN 512:
 - a. Click the Transport & Management VPN tab located directly beneath the Description field, or scroll to the Transport & Management VPN section.
 - b. Under Additional VPN 0 Templates, located to the right of the screen, click VPN Interface GRE.
 - c. From the VPN Interface GRE drop-down, click Create Template. The VPN Interface GRE template form is displayed. The top of the form contains fields for naming the template, and the bottom contains fields for defining VPN Interface GRE parameters.

6. To create a template for VPNs 1 through 511, and 513 through 65530:
 - a. Click the Service VPN tab located directly beneath the Description field, or scroll to the Service VPN section.
 - b. Click the Service VPN drop-down.
 - c. Under Additional VPN templates, located to the right of the screen, click VPN Interface GRE.
 - d. From the VPN Interface GRE drop-down, click Create Template. The VPN Interface GRE template form is displayed. The top of the form contains fields for naming the template, and the bottom contains fields for defining VPN Interface GRE parameters.

The screenshot displays the Cisco vManage configuration page for creating a feature template. The breadcrumb trail is 'Feature Template > Add Template > VPN Interface GRE'. The 'Device Type' is 'vEdge Cloud'. The 'Template Name' and 'Description' fields are empty. The 'Basic Configuration' section is expanded, showing 'Shutdown' (checked), 'Interface Name (1..255)' (gre), and 'Description' (checked). The 'Source' section has 'IP Address' selected. 'Save' and 'Cancel' buttons are at the bottom.

7. In the Template Name field, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
8. In the Template Description field, enter a description of the template. The description can be up to 2048 characters and can contain only alphanumeric characters.

When you first open a feature template, for each parameter that has a default value, the scope is set to Default (indicated by a check mark), and the default setting or value is shown. To change the default or to enter a value, click the scope drop-down to the left of the parameter field and select one of the following:

Table 47:

Parameter Scope	Scope Description
Device Specific (indicated by a host icon)	<p>Use a device-specific value for the parameter. For device-specific parameters, you cannot enter a value in the feature template. You enter the value when you attach a Viptela device to a device template .</p> <p>When you click Device Specific, the Enter Key box opens. This box displays a key, which is a unique string that identifies the parameter in a CSV file that you create. This file is an Excel spreadsheet that contains one column for each key. The header row contains the key names (one key per column), and each row after that corresponds to a device and defines the values of the keys for that device. You upload the CSV file when you attach a Viptela device to a device template. For more information, see Create a Template Variables Spreadsheet .</p> <p>To change the default key, type a new string and move the cursor out of the Enter Key box.</p> <p>Examples of device-specific parameters are system IP address, hostname, GPS location, and site ID.</p>
Global (indicated by a globe icon)	<p>Enter a value for the parameter, and apply that value to all devices.</p> <p>Examples of parameters that you might apply globally to a group of devices are DNS server, syslog server, and interface MTUs.</p>

Configuring a Basic GRE Interface

To configure a basic GRE interface, select the Basic Configuration and then configure the following parameters. Parameters marked with an asterisk are required to configure a GRE interface.

Table 48:

Parameter Name	Description
Shutdown*	Click Off to enable the interface.
Interface Name*	Enter the name of the GRE interface, in the format gre number . <i>number</i> can be from 1 through 255.
Description	Enter a description of the GRE interface.
Source*	<p>Enter the source of the GRE interface:</p> <ul style="list-style-type: none"> GRE Source IP Address—Enter the source IP address of the GRE tunnel interface. This address is on the local router. Tunnel Source Interface—Enter the physical interface that is the source of the GRE tunnel.
Destination*	Enter the destination IP address of the GRE tunnel interface. This address is on a remote device

Parameter Name	Description
GRE Destination IP Address*	Enter the destination IP address of the GRE tunnel interface. This address is on a remote device
IPv4 Address	Enter an IPv4 address for the GRE tunnel.
IP MTU	Specify the maximum MTU size of packets on the interface. <i>Range: 576 through 1804</i> <i>Default: 1500 bytes</i>
Clear-Dont-Fragment	Click On to clear the Don't Fragment bit in the IPv4 packet header for packets being transmitted out the interface.
TCP MSS	Specify the maximum segment size (MSS) of TPC SYN packets passing through the Cisco vEdge device. By default, the MSS is dynamically adjusted based on the interface or tunnel MTU such that TCP SYN packets are never fragmented. <i>Range: 552 to 1460 bytes</i> <i>Default: None</i>
Keepalive Interval	Specify how often the GRE interface sends keepalive packets on the GRE tunnel. Because GRE tunnels are stateless, sending of keepalive packets is the only way to determine whether the remote end of the tunnel is up. The keepalive packets are looped back to the sender. Receipt of these packets by the sender indicates that the remote end of the GRE tunnel is up. <i>Range: 0 through 65535 seconds</i> <i>Default: 10 seconds</i>
Keepalive Retries	Specify how many times the GRE interface tries to resend keepalive packets before declaring the remote end of the GRE tunnel to be down. <i>Range: 0 through 255</i> <i>Default: 3</i>

To save the feature template, click Save.

CLI equivalent:

```
vpn vpn-id interface grenumber clear-dont-fragment description text
ip address ipv4-prefix/length keepalive seconds retries mtu bytes
policer policer-name (in |out)
    qos-map name rewrite-rule name shaping-rate name
    [no] shutdown tcp-mss-adjust bytes tunnel-destination ip-address
    ( tunnel-source ip-address | tunnel-source-interface interface-name)
```

Configure Interface Access Lists

To configure access lists on a GRE interface, select the ACL tab and configure the following parameters:

Table 49:

Parameter Name	Description
Rewrite Rule	Click On, and specify the name of the rewrite rule to apply on the interface.
Ingress ACL – IPv4	Click On, and specify the name of the access list to apply to IPv4 packets being received on the interface.
Egress ACL – IPv4	Click On, and specify the name of the access list to apply to IPv4 packets being transmitted on the interface.

Parameter Name	Description
Ingress Policer	Click On, and specify the name of the policer to apply to packets being received on the interface.
Egress Policer	Click On, and specify the name of the policer to apply to packets being transmitted on the interface.

CLI equivalent:

```
vpn vpn-id interface gnumber access-list acl-list (in | out)
  policer policer-name (in |out)
  qos-map name rewrite-rule name shaping-rate name
```

Release Information

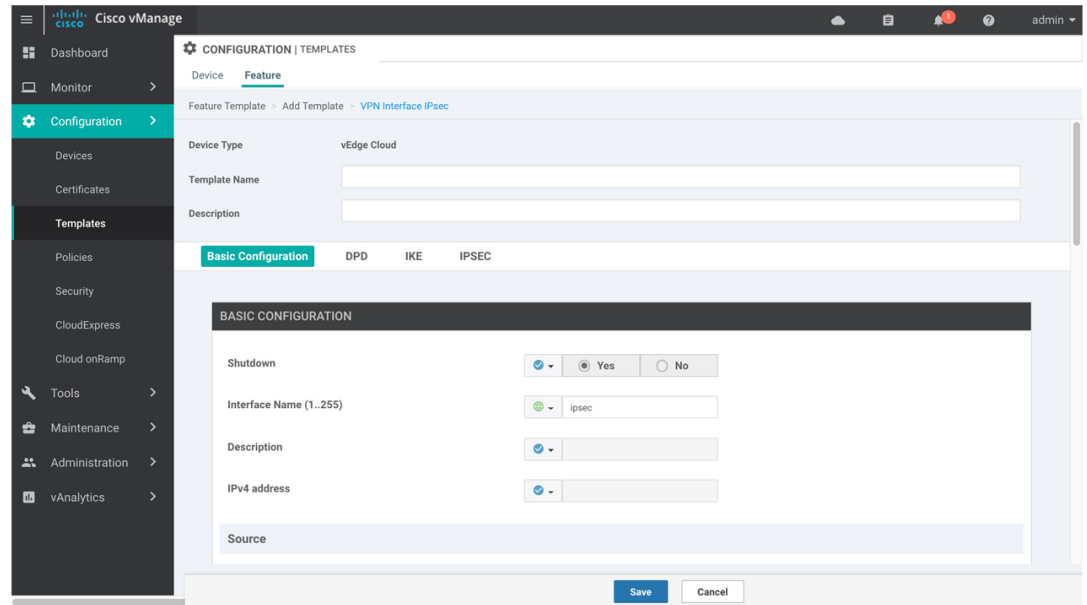
Introduced in Cisco vManage NMS Release 15.4.1.

VPN Interface IPsec (for Cisco vEdge Devices)

Use the VPN Interface IPsec feature template to configure IPsec tunnels on Cisco vEdge devices that are being used for Internet Key Exchange (IKE) sessions. You can configure IPsec on tunnels in the transport VPN (VPN 0) and in service VPNs (VPN 1 through 65530, except for 512).

Navigate to the Template Screen and Name the Template

1. In Cisco vManage NMS, select the Configuration ► Templates screen.
2. In the Device tab, click Create Template.
3. From the Create Template drop-down, select From Feature Template.
4. From the Device Model drop-down, select the type of device for which you are creating the template.
5. Click the Service VPN tab located directly beneath the Description field, or scroll to the Service VPN section.
6. Click the Service VPN drop-down.



369436

7. Under Additional VPN Templates, located to the right of the screen, click VPN Interface IPsec.
8. From the VPN Interface IPsec drop-down, click Create Template. The VPN Interface IPsec template form is displayed. The top of the form contains fields for naming the template, and the bottom contains fields for defining VPN Interface IPsec parameters.
9. In the Template Name field, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
10. In the Template Description field, enter a description of the template. The description can be up to 2048 characters and can contain only alphanumeric characters.

When you first open a feature template, for each parameter that has a default value, the scope is set to Default (indicated by a check mark), and the default setting or value is shown. To change the default or to enter a value, click the scope drop-down to the left of the parameter field and select one of the following:

Table 50:

Parameter Scope	Scope Description
Device Specific (indicated by a host icon)	<p>Use a device-specific value for the parameter. For device-specific parameters, you cannot enter a value in the feature template. You enter the value when you attach a Viptela device to a device template .</p> <p>When you click Device Specific, the Enter Key box opens. This box displays a key, which is a unique string that identifies the parameter in a CSV file that you create. This file is an Excel spreadsheet that contains one column for each key. The header row contains the key names (one key per column), and each row after that corresponds to a device and defines the values of the keys for that device. You upload the CSV file when you attach a Viptela device to a device template. For more information, see Create a Template Variables Spreadsheet .</p> <p>To change the default key, type a new string and move the cursor out of the Enter Key box.</p> <p>Examples of device-specific parameters are system IP address, hostname, GPS location, and site ID.</p>
Global (indicated by a globe icon)	<p>Enter a value for the parameter, and apply that value to all devices.</p> <p>Examples of parameters that you might apply globally to a group of devices are DNS server, syslog server, and interface MTUs.</p>

Configure a Basic IPsec Tunnel Interface

To configure an IPsec tunnel to use for IKE sessions, select the Basic Configuration tab and configure the following parameters. Parameters marked with an asterisk are required to configure an IPsec tunnel.

Table 51:

Parameter Name	Description
Shutdown*	Click No to enable the interface.
Interface Name*	Enter the name of the IPsec interface, in the format ipsec number . <i>number</i> can be from 1 through 256.
Description	Enter a description of the IPsec interface.
IPv4 Address*	Enter the IPv4 address of the IPsec interface, in the format <i>ipv4-prefix/length</i> . The address must be a /30.
Source*	<p>Set the source of the IPsec tunnel that is being used for IKE key exchange:</p> <ul style="list-style-type: none"> • Click IP Address—Enter the IPv4 address that is the source tunnel interface. This address must be configured in VPN 0. • Click Interface—Enter the name of the physical interface that is the source of the IPsec tunnel. This interface must be configured in VPN 0.

Parameter Name	Description
Destination: IPsec Destination IP Address/FQDN*	Set the destination of the IPsec tunnel that is being used for IKE key exchange. Enter either an IPv4 address or the fully qualified DNS name that points to the destination.
TCP MSS	Specify the maximum segment size (MSS) of TPC SYN packets passing through the Cisco vEdge device. By default, the MSS is dynamically adjusted based on the interface or tunnel MTU such that TCP SYN packets are never fragmented. <i>Range: 552 to 1460 bytesDefault: None</i>
IP MTU	Specify the maximum MTU size of packets on the interface. <i>Range: 576 through 1804Default: 1500 bytes</i>

To save the feature template, click Save.

CLI equivalent:

```
vpn vpn-id
 interface ipsec number ip address ipv4-prefix/length mtu bytes
   no shutdown
   tcp-mss-adjust bytes tunnel-destination ipv4-address
   ( tunnel-source ip-address | tunnel-source-interface interface-name)
```

Configure Dead-Peer Detection

To configure IKE dead-peer detection to determine whether the connection to an IKE peer is functional and reachable, select the DPD tab and configure the following parameters:

Table 52:

Parameter Name	Description
DPD Interval	Specify the interval for IKE to send Hello packets on the connection. <i>Range: 0 through 65535 seconds (1 hour through 14 days)Default: 10 seconds</i>
DPD Retries	Specify how many unacknowledged packets to accept before declaring an IKE peer to be dead and then tearing down the tunnel to the peer. <i>Range: 0 through 255Default: 3</i>

To save the feature template, click Save.

CLI equivalent:

```
vpn vpn-id interface ipsec number dead-peer-detection seconds retries number
```

Configure IKE

To configure IKE, select the IKE tab and configure the parameters discussed below.

When you create an IPsec tunnel on a Cisco vEdge device, IKE Version 1 is enabled by default on the tunnel interface. The following properties are also enabled by default for IKEv1:

- Authentication and encryption—AES-256 advanced encryption standard CBC encryption with the HMAC-SHA1 keyed-hash message authentication code algorithm for integrity

- Diffie-Hellman group number—16
- Rekeying time interval—4 hours
- SA establishment mode—Main

To modify IKEv1 parameters, configure the following:

Table 53:

Parameter Name	Description
IKE Version	Enter 1 to select IKEv1.
IKE Mode	Specify the IKE SA establishment mode. <i>Values:</i> Aggressive mode, Main mode <i>Default:</i> Main mode
IPsec Rekey Interval	Specify the interval for refreshing IKE keys. <i>Range:</i> 3600 through 1209600 seconds (1 hour through 14 days) <i>Default:</i> 14400 seconds (4 hours)
IKE Cipher Suite	Specify the type of authentication and encryption to use during IKE key exchange. <i>Values:</i> aes128-cbc-sha1, aes256-cbc-sha1 <i>Default:</i> aes256-cbc-sha1
IKE Diffie-Hellman Group	Specify the Diffie-Hellman group to use in IKE key exchange. <i>Values:</i> 1024-bit modulus, 2048-bit modulus, 3072-bit modulus, 4096-bit modulus <i>Default:</i> 4096-bit modulus
IKE Authentication: Preshared Key	To use preshared key (PSK) authentication, enter the password to use with the preshared key.
IKE ID for Local End Point	If the remote IKE peer requires a local end point identifier, specify it. <i>Range:</i> <i>Default:</i> Tunnel's source IP address
IKE ID for Remote End Point	If the remote IKE peer requires a remote end point identifier, specify it. <i>Range:</i> 1 through 64 characters <i>Default:</i> Tunnel's destination IP address

To save the feature template, click Save.

CLI equivalent:

```
vpn vpn-id interface ipsec number ike authentication-type type
  local-id id
  pre-shared-secret password
  remote-id id cipher-suite suite group number mode mode rekey-interval seconds
  version 1
```

To configure IKEv2, configure the following parameters:

Table 54:

Parameter Name	Description
IKE Version	Enter 2 to select IKEv2.

IPsec Rekey Interval	Specify the interval for refreshing IKE keys. <i>Range:</i> 3600 through 1209600 seconds (1 hour through 14 days) <i>Default:</i> 14400 seconds (4 hours)
IKE Cipher Suite	Specify the type of authentication and encryption to use during IKE key exchange. <i>Values:</i> aes128-cbc-sha1, aes256-cbc-sha1 <i>Default:</i> aes256-cbc-sha1
IKE Diffie-Hellman Group	Specify the Diffie-Hellman group to use in IKE key exchange. <i>Values:</i> 1024-bit modulus, 2048-bit modulus, 3072-bit modulus, 4096-bit modulus <i>Default:</i> 4096-bit modulus
IKE Authentication: Preshared Key	To use preshared key (PSK) authentication, enter the password to use with the preshared key.
IKE ID for Local End Point	If the remote IKE peer requires a local end point identifier, specify it. <i>Range:</i> <i>Default:</i> Tunnel's source IP address
IKE ID for Remote End Point	If the remote IKE peer requires a remote end point identifier, specify it. <i>Range:</i> 1 through 64 characters <i>Default:</i> Tunnel's destination IP address

To save the feature template, click Save.

CLI equivalent:

```
vpn vpn-id interface ipsec number ike authentication-type type
    local-id id
    pre-shared-secret password
    remote-id id cipher-suite suite group number rekey-interval seconds
version 2
```

Configure IPsec Tunnel Parameters

To configure the IPsec tunnel that carries IKE traffic, select the IPsec tab and configure the following parameters:

Table 55:

Parameter Name	Description
IPsec Rekey Interval	Specify the interval for refreshing IKE keys. <i>Range:</i> 3600 through 1209600 seconds (1 hour through 14 days) <i>Default:</i> 14400 seconds (4 hours)
IKE Replay Window	Specify the replay window size for the IPsec tunnel. <i>Values:</i> 64, 128, 256, 512, 1024, 2048, 4096, 8192 bytes <i>Default:</i> 32 bytes
IPsec Cipher Suite	Specify the authentication and encryption to use on the IPsec tunnel. <i>Values:</i> aes256-cbc-sha1, aes256-gcm, null-sha1 <i>Default:</i> aes256-gcm
Perfect Forward Secrecy	Specify the PFS settings to use on the IPsec tunnel. <i>Values:</i> • group-2 —Use the 1024-bit Diffie-Hellman prime modulus group. • group-14 —Use the 2048-bit Diffie-Hellman prime modulus group. • group-15 —Use the 3072-bit Diffie-Hellman prime modulus group. • group-16 —Use the 4096-bit Diffie-Hellman prime modulus group. • none —Disable PFS. <i>Default:</i> group-16

To save the feature template, click Save.

CLI equivalent:

```
vpn vpn-id interface ipsec number ipsec cipher-suite suite perfect-forward-secrecy
pfs-setting rekey-interval seconds replay-window number
```

Release Information

Introduced in Cisco vManage NMS in Release 17.2. In Release 17.2.3, add support for PFS. In Release 18.2, support support for IPsec tunnels in VPN 0. In Release 18.4, standard IPsec support for IOS XE routers.

VPN Interface PPP

Point-to-Point Protocol (PPP) is a data link protocol used to establish a direct connection between two nodes. PPP properties are associated with a PPPoE-enabled interface on Cisco SD-WAN devices to connect multiple users over an Ethernet link.

To configure PPPoE on Cisco vEdge devices using Cisco vManage templates:

1. Create a VPN Interface PPP feature template to configure PPP parameters for the PPP virtual interface, as described in this article.
2. Create a VPN Interface PPP Ethernet feature template to configure a PPPoE-enabled interface. See the VPN Interface PPP Ethernet help topic.
3. Optionally, create a VPN feature template to modify the default configuration of VPN 0. See the VPN help topic.

Navigate to the Template Screen and Name the Template

1. In vManage NMS, select the **Configuration** > **Templates** screen.
2. In the Device tab, click **Create Template**.
3. From the **Create Template** drop-down, select **From Feature Template**.
4. From the **Device Model** drop-down, select the type of device for which you are creating the template.
5. Click the **Transport & Management VPN** tab located directly beneath the **Description** field, or scroll to the **Transport & Management VPN** section.
6. Under **Additional VPN 0 Templates**, located to the right of the screen, click **VPN Interface PPP**.

369437

7. From the **VPN Interface PPP** drop-down, click **Create Template**. The VPN Interface PPP template form is displayed. The top of the form contains fields for naming the template, and the bottom contains fields for defining VPN Interface PPP parameters.
8. In the **Template Name** field, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
9. In the **Template Description** field, enter a description of the template. The description can be up to 2048 characters and can contain only alphanumeric characters.

When you first open a feature template, for each parameter that has a default value, the scope is set to Default (indicated by a check mark), and the default setting or value is shown. To change the default or to enter a value, click the scope drop-down to the left of the parameter field and select one of the following:

Table 56:

Parameter Scope	Scope Description
Device Specific (indicated by a host icon)	<p>Use a device-specific value for the parameter. For device-specific parameters, you cannot enter a value in the feature template. You enter the value when you attach a Cisco vEdge device to a device template .</p> <p>When you click Device Specific, the Enter Key box opens. This box displays a key, which is a unique string that identifies the parameter in a CSV file that you create. This file is an Excel spreadsheet that contains one column for each key. The header row contains the key names (one key per column), and each row after that corresponds to a device and defines the values of the keys for that device. You upload the CSV file when you attach a Viptela device to a device template. For more information, see Create a Template Variables Spreadsheet .</p> <p>To change the default key, type a new string and move the cursor out of the Enter Key box.</p> <p>Examples of device-specific parameters are system IP address, hostname, GPS location, and site ID.</p>
Global (indicated by a globe icon)	<p>Enter a value for the parameter, and apply that value to all devices.</p> <p>Examples of parameters that you might apply globally to a group of devices are DNS server, syslog server, and interface MTUs.</p>

Configure a PPP Virtual Interface

To configure a PPP virtual interface, select the **Basic Configuration** tab and configure the following parameters. Parameters marked with an asterisk are required to configure the interface. You must also configure an authentication protocol and a tunnel interface for the PPP interface, and you must ensure that the maximum MTU for the PPP interface is 1492 bytes.

Table 57:

Parameter Name	Description
Shutdown*	Click No to enable the PPP virtual interface.
PPP Interface Name*	Enter the number of the PPP interface. It can be a number from 1 through 31.
Description	Enter a description for the PPP virtual interface.
Bandwidth Upstream	For transmitted traffic, set the bandwidth above which to generate notifications. <i>Range:</i> 1 through $(2^{32} / 2) - 1$ kbps
Bandwidth Downstream	For received traffic, set the bandwidth above which to generate notifications. <i>Range:</i> 1 through $(2^{32} / 2) - 1$ kbps
Block Non-Source IP	Click Yes to have the interface forward traffic only if the source IP address of the traffic matches the interface's IP prefix range.

To save the feature template, click **Save**.

CLI equivalent:

```

vpn 0
  interface pppnumber bandwidth-downstream kbps bandwidth-upstream kbps block-non-source-ip
  ppp
    no shutdown

```

Configure the Access Concentrator Name and Authentication Protocol

To configure the access concentrator name, select the PPP tab and configure the following parameters:

Table 58:

Parameter Name	Description
AC Name	Name of the access concentrator used by PPPoE to route connections to the Internet.
Authentication Protocol	Select the authentication protocol used by PPPoE: <ul style="list-style-type: none"> • CHAP—Enter the hostname and password provided by your Internet Service Provider (ISP). <i>hostname</i> can be up to 255 characters. • PAP—Enter the username and password provided by your ISP. <i>username</i> can be up to 255 characters. • PAP and CHAP—Configure both authentication protocols. Enter the login credentials for each protocol. To use the same username and password for both, click Same Credentials for PAP and CHAP.

To save the feature template, click **Save**.

CLI equivalent:

```

vpn 0
  interface pppnumber ppp
    ac-name name
    authentication
      chap hostname name password password
      pap password password sent-username name

```

Create a Tunnel Interface

On Cisco vEdge devices, you can configure up to four tunnel interfaces. This means that each Cisco vEdge device can have up to four TLOCs.

For the control plane to establish itself so that the overlay network can function, you must configure WAN transport interfaces in VPN 0.

To configure a tunnel interface for the PPP interface, select the **Tunnel Interface** tab and configure the following parameters:

Table 59:

Parameter Name	Description
Tunnel Interface	Click On to create a tunnel interface.
Color	Select a color for the TLOC.

Parameter Name	Description
Control Connection	If the Cisco vEdge device has multiple TLOCs, click No to have the tunnel not establish a TLOC. The default is On, which establishes a control connection for the TLOC.
Maximum Control Connections	Specify the maximum number of Cisco vSmart Controller that the WAN tunnel interface can connect to. To have the tunnel establish no control connections, set the number to 0. <i>Range: 0 through 8Default: 2</i>
vBond As STUN Server	Click On to enable Session Traversal Utilities for NAT (STUN) to allow the tunnel interface to discover its public IP address and port number when the Cisco vEdge device is located behind a NAT.
Exclude Controller Group List	Set the Cisco vSmart Controller that the tunnel interface is not allowed to connect to. <i>Range: 0 through 100</i>
vManage Connection Preference	Set the preference for using a tunnel interface to exchange control traffic with the vManage NMS. <i>Range: 0 through 8Default: 5</i>
Low-Bandwidth Link	Select to characterize the tunnel interface as a low-bandwidth link.
Allow Service	Select On or Off for each service to allow or disallow the service on the interface.

To configure additional tunnel interface parameters, click **Advanced Options** and configure the following parameters:

Table 60:

Parameter Name	Description
GRE	Use GRE encapsulation on the tunnel interface. By default, GRE is disabled. If you select both IPsec and GRE encapsulations, two TLOCs are created for the tunnel interface that have the same IP addresses and colors, but that differ by their encapsulation.
IPsec	Use IPsec encapsulation on the tunnel interface. By default, IPsec is enabled. If you select both IPsec and GRE encapsulations, two TLOCs are created for the tunnel interface that have the same IP addresses and colors, but that differ by their encapsulation.
IPsec Preference	Specify a preference value for directing traffic to the tunnel. A higher value is preferred over a lower value. <i>Range: 0 through 4294967295Default: 0</i>
IPsec Weight	Enter a weight to use to balance traffic across multiple TLOCs. A higher value sends more traffic to the tunnel. <i>Range: 1 through 255Default: 1</i>

Parameter Name	Description
Carrier	Select the carrier name or private network identifier to associate with the tunnel. <i>Values:</i> carrier1, carrier2, carrier3, carrier4, carrier5, carrier6, carrier7, carrier8, default <i>Default:</i> default
Bind Loopback Tunnel	Enter the name of a physical interface to bind to a loopback interface.
Last-Resort Circuit	Select to use the tunnel interface as the circuit of last resort.
NAT Refresh Interval	Enter the interval between NAT refresh packets sent on a DTLS or TLS WAN transport connection. <i>Range:</i> 1 through 60 seconds <i>Default:</i> 5 seconds
Hello Interval	Enter the interval between Hello packets sent on a DTLS or TLS WAN transport connection. <i>Range:</i> 100 through 10000 milliseconds <i>Default:</i> 1000 milliseconds (1 second)
Hello Tolerance	Enter the time to wait for a Hello packet on a DTLS or TLS WAN transport connection before declaring that transport tunnel to be down. <i>Range:</i> 12 through 60 seconds <i>Default:</i> 12 seconds

CLI equivalent:

```
vpn 0
  interface interface-name tunnel-interface allow-service service-name
bind interface-name
  carrier carrier-name
  color color encapsulation (gre | ipsec)
  preference number
  weight number hello-interval milliseconds hello-tolerance seconds
last-resort-circuit max-control-connections number nat-refresh-interval seconds
vbond-as-stun-server
```

Configure the Interface as a NAT Device

To configure an interface to act as a NAT device, select the **NAT** tab and configure the following parameters:

Table 61:

Parameter Name	Description
NAT	Click On to have the interface act as a NAT device.
Refresh Mode	Select how NAT mappings are refreshed, either outbound or bidirectional (outbound and inbound). <i>Default:</i> Outbound
UDP Timeout	Specify when NAT translations over UDP sessions time out. <i>Range:</i> 1 through 65536 minutes <i>Default:</i> 1 minutes

TCP Timeout	Specify when NAT translations over TCP sessions time out. <i>Range:</i> 1 through 65536 minutes <i>Default:</i> 60 minutes (1 hour)
Block ICMP	Select On to block inbound ICMP error messages. By default, a Cisco vEdge device acting as a NAT device receives these error messages. <i>Default:</i> Off
Respond to Ping	Select On to have the Cisco vEdge device respond to ping requests to the NAT interface's IP address that are received from the public side of the connection.

To create a port forwarding rule, click **Add New Port Forwarding Rule** and configure the following parameters. You can define up to 128 port-forwarding rules to allow requests from an external network to reach devices on the internal network.

Table 62:

Parameter Name	Description
Port Start Range	Enter a port number to define the port or first port in the range of interest. <i>Range:</i> 0 through 65535
Port End Range	Enter the same port number to apply port forwarding to a single port, or enter the larger number to apply it to a range or ports. <i>Range:</i> 0 through 65535
Protocol	Select the protocol to which to apply the port-forwarding rule, either TCP or UDP. To match the same ports for both TCP and UDP traffic, configure two rules.
VPN	Specify the private VPN in which the internal server resides. This VPN is one of the VPN identifiers in the overlay network. <i>Range:</i> 0 through 65535
Private IP	Specify the IP address of the internal server to which to direct traffic that matches the port-forwarding rule.

To save a port forwarding rule, click **Add**.

To save the feature template, click **Save**.

CLI equivalent:

```
vpn vpn-id
interface interface-name nat block-icmp-error port-forward port-start port-number1
port-end port-number2 proto (tcp | udp)
    private-ip-address ip-address private-vpn vpn-id refresh (bi-directional | outbound)

    respond-to-ping tcp-timeout minutes
    udp-timeout minutes
```

Apply Access Lists

To apply a rewrite rule, access lists, and policers to a router interface, select the **ACL** tab and configure the following parameters:

Table 63:

Parameter Name	Description
Rewrite Rule	Click On , and specify the name of the rewrite rule to apply on the interface.
Ingress ACL – IPv4	Click On , and specify the name of the access list to apply to IPv4 packets being received on the interface.
Egress ACL – IPv4	Click On , and specify the name of the access list to apply to IPv4 packets being transmitted on the interface.
Ingress ACL – IPv6	Click On , and specify the name of the access list to apply to IPv6 packets being received on the interface.
Egress ACL – IPv6	Click On , and specify the name of the access list to apply to IPv6 packets being transmitted on the interface.
Ingress Policer	Click On , and specify the name of the policer to apply to packets being received on the interface.
Egress Policer	Click On , and specify the name of the policer to apply to packets being transmitted on the interface.

To save the feature template, click **Save**.

CLI equivalent:

```
vpn 0
 interface pppnumber access-list acl-name (in | out)
   ipv6 access-list acl-name (in | out)
   policer policer-name (in |out)
   rewrite-rule name
```

Configure Other Interface Properties

To configure other interface properties, select the **Advanced** tab and configure the following properties:

Table 64:

Parameter Name	Description
MAC Address	Specify a MAC address to associate with the interface, in colon-separated hexadecimal notation.
IP MTU	Specify the maximum MTU size of packets on the interface. <i>Range:</i> 576 through 1804 <i>Default:</i> 1500 bytes
TCP MSS	Specify the maximum segment size (MSS) of TCP SYN packets passing through the Cisco vEdge device. By default, the MSS is dynamically adjusted based on the interface or tunnel MTU such that TCP SYN packets are never fragmented. <i>Range:</i> 552 to 1460 bytes <i>Default:</i> None

Parameter Name	Description
Clear Dont Fragment	Click On to clear the Don't Fragment bit in the IPv4 packet header for packets being transmitted out the interface. When the DF bit is cleared, packets larger than that interface's MTU are fragmented before being sent.
TLOC Extension	Enter the name of the physical interface on the same router that connects to the WAN transport circuit. This configuration then binds this service-side interface to the WAN transport. A second Cisco vEdge device at the same site that itself has no direct connection to the WAN (generally because the site has only a single WAN connection) and that connects to this service-side interface is then provided with a connection to the WAN.
Tracker	Enter the name of a tracker to track the status of transport interfaces that connect to the internet.
ICMP Redirect	Click Disable to disable ICMP redirect messages on the interface. By default, an interface allows ICMP redirect messages.

To save the feature template, click **Save**.

CLI equivalent:

```
vpn vpn-id interface interface-name clear-dont-fragment icmp-redirect-disable
mac-address mac-address mtu bytes tcp-mss-adjust bytes tloc-extension
interface-name tracker tracker-name
```

Release Information

Introduced in vManage NMS in Release 15.3. In Release 16.3, add support for IPv6. In Release 17.1, support ability to configure both CHAP and PAP authentication on a PPP interface. In Release 17.2.2, add support for interface status tracking. In Release 18.2, add support for disabling ICMP redirect messages.

VPN Interface PPP Ethernet

Use the VPN Interface PPP Ethernet template for Cisco vEdge devices.

Point-to-Point Protocol (PPP) is a data link protocol used to establish a direct connection between two nodes. PPP properties are associated with a PPPoE-enabled interface on Cisco vEdge devices to connect multiple users over an Ethernet link.

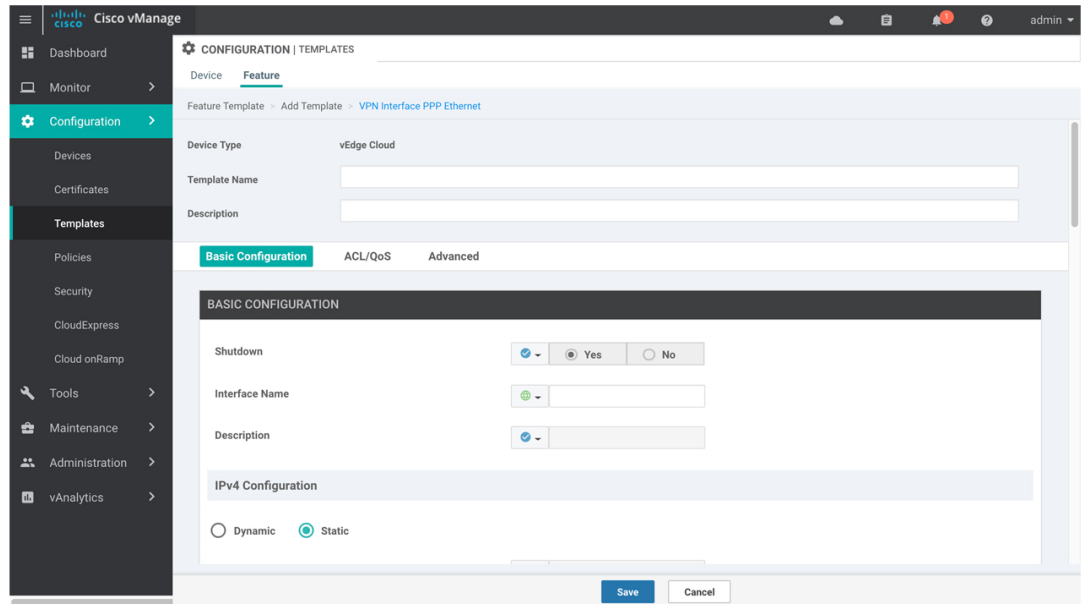
To configure PPPoE on Cisco vEdge device using Cisco vManage templates:

1. Create a VPN Interface PPP Ethernet feature template to configure a PPPoE-enabled interface as described in this article.
2. Create a VPN Interface PPP feature template to configure PPP parameters for the PPP virtual interface. See the VPN Interface PPP help topic.
3. Optionally, create a VPN feature template to modify the default configuration of VPN 0. See the VPN help topic.

Navigate to the Template Screen and Name the Template

1. In Cisco vManage, select the **Configuration > Templates** screen.

2. In the **Device** tab, click **Create Template**.
3. From the **Create Template** drop-down, select **From Feature Template**.
4. From the **Device Model** drop-down, select the type of device for which you are creating the template.
5. Click the **Transport & Management VPN** tab located directly beneath the **Description** field, or scroll to the **Transport & Management VPN** section.
6. Under **Additional VPN 0 Templates**, located to the right of the screen, click **VPN Interface PPP**.



7. From the **VPN Interface PPP Ethernet** drop-down, click **Create Template**. The **VPN Interface PPP Ethernet** template form is displayed. The top of the form contains fields for naming the template, and the bottom contains fields for defining VPN Interface PPP parameters.
8. In the **Template Name** field, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
9. In the **Template Description** field, enter a description of the template. The description can be up to 2048 characters and can contain only alphanumeric characters.

When you first open a feature template, for each parameter that has a default value, the scope is set to Default (indicated by a check mark), and the default setting or value is shown. To change the default or to enter a value, click the scope drop-down to the left of the parameter field and select one of the following:

Table 65:

Parameter Scope	Scope Description
Device Specific (indicated by a host icon)	<p>Use a device-specific value for the parameter. For device-specific parameters, you cannot enter a value in the feature template. You enter the value when you attach a Viptela device to a device template .</p> <p>When you click Device Specific, the Enter Key box opens. This box displays a key, which is a unique string that identifies the parameter in a CSV file that you create. This file is an Excel spreadsheet that contains one column for each key. The header row contains the key names (one key per column), and each row after that corresponds to a device and defines the values of the keys for that device. You upload the CSV file when you attach a Viptela device to a device template. For more information, see Create a Template Variables Spreadsheet .</p> <p>To change the default key, type a new string and move the cursor out of the Enter Key box.</p> <p>Examples of device-specific parameters are system IP address, hostname, GPS location, and site ID.</p>
Global (indicated by a globe icon)	<p>Enter a value for the parameter, and apply that value to all devices.</p> <p>Examples of parameters that you might apply globally to a group of devices are DNS server, syslog server, and interface MTUs.</p>

Configure a Basic PPPoE-Enabled Interface

To create a PPPoE-enabled interface on a Cisco vEdge device, select the **Basic Configuration** tab and configure the following parameters. Parameters marked with an asterisk are required to configure the interface.

Table 66:

Parameter Name	Description
Shutdown*	Click No to enable the PPPoE-enabled interface.
Interface Name*	<p>Enter the name of the physical interface in VPN 0 to associate with the PPP interface.</p> <p>For Cisco XE SD-WAN devices, you must spell out the interface names completely (for example, GigabitEthernet0/0/0), and you must configure all the router's interfaces even if you are not using them so that they are configured in the shutdown state and so that all default values for them are configured.</p>
Description	Enter a description of the PPPoE-enabled interface.
IPv4 Configuration*	<p>To configure a static address, click Static and enter an IPv4 address.</p> <p>To set the interface as a DHCP client so that the interface to receive its IP address from a DHCP server, click Dynamic. You can optionally set the DHCP distance to specify the administrative distance of routes learned from a DHCP server. The default DHCP distance is 1.</p>

Parameter Name	Description
IPv6 Configuration*	To configure a static address for an interface in VPN 0, click Static and enter an IPv6 address. To set the interface as a DHCP client so that the interface to receive its IP address from a DHCP server, click Dynamic . You can optionally set the DHCP distance to specify the administrative distance of routes learned from a DHCP server. The default DHCP distance is 1. You can optionally enable DHCP rapid commit, to speed up the assignment of IP addresses.
DHCP Helper	Enter up to eight IP addresses for DHCP servers in the network, separated by commas, to have the interface be a DHCP helper. A DHCP helper interface forwards BOOTP (Broadcast) DHCP requests that it receives from the specified DHCP servers.
Bandwidth Upstream	For transmitted traffic, set the bandwidth above which to generate notifications. <i>Range:</i> 1 through $(2^{32} / 2) - 1$ kbps
Bandwidth Downstream	For received traffic, set the bandwidth above which to generate notifications. <i>Range:</i> 1 through $(2^{32} / 2) - 1$ kbps

To save the feature template, click **Save**.

CLI equivalent:

```
vpn 0
  interface pppnumber bandwidth-downstream kbps bandwidth-upstream kbps description text
  dhcp-helper ip-address
    ( ip address ipv4-prefix/length | ip-dhcp-client [dhcp-distance number])
    ( ipv6 address ipv6-prefix/length | ipv6 dhcp-client [dhcp-distance number] [
dhcp-rapid-commit]
    pppoe-client ppp-interface pppnumber
    [no] shutdown
```

Apply Access Lists

To configure a shaping rate to a PPPoE-enabled interface and to apply a QoS map, a rewrite rule, access lists, and policers to the interface, select the ACL/QoS tab and configure the following parameters:

Table 67:

Parameter Name	Description
Shaping Rate	Configure the aggregate traffic transmission rate on the interface to be less than line rate, in kilobits per second (kbps).
QoS Map	Specify the name of the QoS map to apply to packets being transmitted out the interface.
Rewrite Rule	Click On , and specify the name of the rewrite rule to apply on the interface.
Ingress ACL – IPv4	Click On , and specify the name of the access list to apply to IPv4 packets being received on the interface.
Egress ACL – IPv4	Click On , and specify the name of the access list to apply to IPv4 packets being transmitted on the interface.

Parameter Name	Description
Ingress ACL – IPv6	Egress ACL – IPv6
Egress ACL – IPv6	Egress ACL – IPv6
Ingress Policer	Click On and specify the name of the policer to apply to packets being received on the interface.
Egress Policer	Click On , and specify the name of the policer to apply to packets being transmitted on the interface.

To save the feature temp

CLI equivalent:

```
vpn 0
 interface pppnumber access-list acl-list (in | out)
   policer policer-name (in |out)
   qos-map name rewrite-rule name shaping-rate name
```

Configure Other Interface Properties

To configure other interface properties, select the **Advanced** tab and configure the following properties:

Table 68:

Parameter Name	Description
Duplex	Choose full or half to specify whether the interface runs in full-duplex or half-duplex mode. <i>Default:</i> Full
MAC Address	Specify a MAC address to associate with the interface, in colon-separated hexadecimal notation.
IP MTU	Specify the maximum MTU size of packets on the interface. <i>Range:</i> 576 through 1804 <i>Default:</i> 1500 bytes
PMTU Discovery	Click On to enable path MTU discovery on the interface. PMTU determines the largest MTU size that the interface supports so that packet fragmentation does not occur.
Flow Control	Select a setting for bidirectional flow control, which is a mechanism for temporarily stopping the transmission of data on the interface. <i>Values:</i> autonet, both, egress, ingress, none <i>Default:</i> autoneg
TCP MSS	Specify the maximum segment size (MSS) of TPC SYN packets passing through the Cisco vEdge device. By default, the MSS is dynamically adjusted based on the interface or tunnel MTU such that TCP SYN packets are never fragmented. <i>Range:</i> 552 to 1460 bytes <i>Default:</i> None

Parameter Name	Description
Speed	Specify the speed of the interface, for use when the remote end of the connection does not support autonegotiation. <i>Values:</i> 10, 100, or 1000 Mbps <i>Default:</i> Autonegotiate (10/100/1000 Mbps)
Static Ingress QoS	Specify a queue number to use for incoming traffic. <i>Range:</i> 0 through 7
ARP Timeout	Specify how long it takes for a dynamically learned ARP entry to time out. <i>Range:</i> 0 through 2678400 seconds (744 hours) <i>Default:</i> 1200 seconds (20 minutes)
Autonegotiation	Click Off to turn off autonegotiation. By default, an interface runs in autonegotiation mode.
TLOC Extension	Enter the name of a physical interface on the same router that connects to the WAN transport. This configuration then binds this service-side interface to the WAN transport. A second Cisco vEdge device at the same site that itself has no direct connection to the WAN (generally because the site has only a single WAN connection) and that connects to this service-side interface is then provided with a connection to the WAN.
Power over Ethernet (on Cisco vEdge 100m and Cisco vEdge 100wm routers)	Click On to enable PoE on the interface.
ICMP Redirect	Click Disable to disable ICMP redirect messages on the interface. By default, an interface allows ICMP redirect messages.

To save the feature template, click **Save**.

CLI equivalent:

```
vpn 0
  interface pppnumber arp-timeout seconds
    [no] autonegotiate duplex (full | half)
    flow-control control icmp-redirect-disable mac-address mac-address mtu bytes pmtu
pppoe-client
  ppp-interface pppnumber speed speed
  static-ingress-qos number tcp-mss-adjust bytes tloc-extension interface-name
```

Release Information

Introduced in vManage NMS Release 15.3. In Release 16.3, add support for IPv6. In Release 18.2, add support for disabling ICMP redirect messages.

Cellular Interfaces

To enable LTE connectivity, configure cellular interfaces on a router that has a cellular module. The cellular module provides wireless connectivity over a service provider's cellular network. One use case is to provide wireless connectivity for branch offices.

A cellular network is commonly used as a backup WAN link, to provide network connectivity if all the wired WAN tunnel interfaces on the router become unavailable. You can also use a cellular network as the primary WAN link for a branch office, depending on usage patterns within the branch office and the data rates supported by the core of the service provider's cellular network.

When you configure a cellular interface on a device, you can connect the device to the Internet or another WAN by plugging in the power cable of the device. The device then automatically begins the process of joining the overlay network, by contacting and authenticating with Cisco vBond Orchestrators, Cisco vSmart Controllers, and Cisco vManage systems.

vEdge routers support LTE and CDMA radio access technology (RAT) types.

Configure Cellular Interfaces Using vManage

To configure cellular interfaces using vManage templates:

1. Create a VPN Interface Cellular feature template to configure cellular module parameters, as described in this article.
2. Create a Cellular Profile template to configure the profiles used by the cellular modem.
3. Create a VPN feature template to configure VPN parameters.

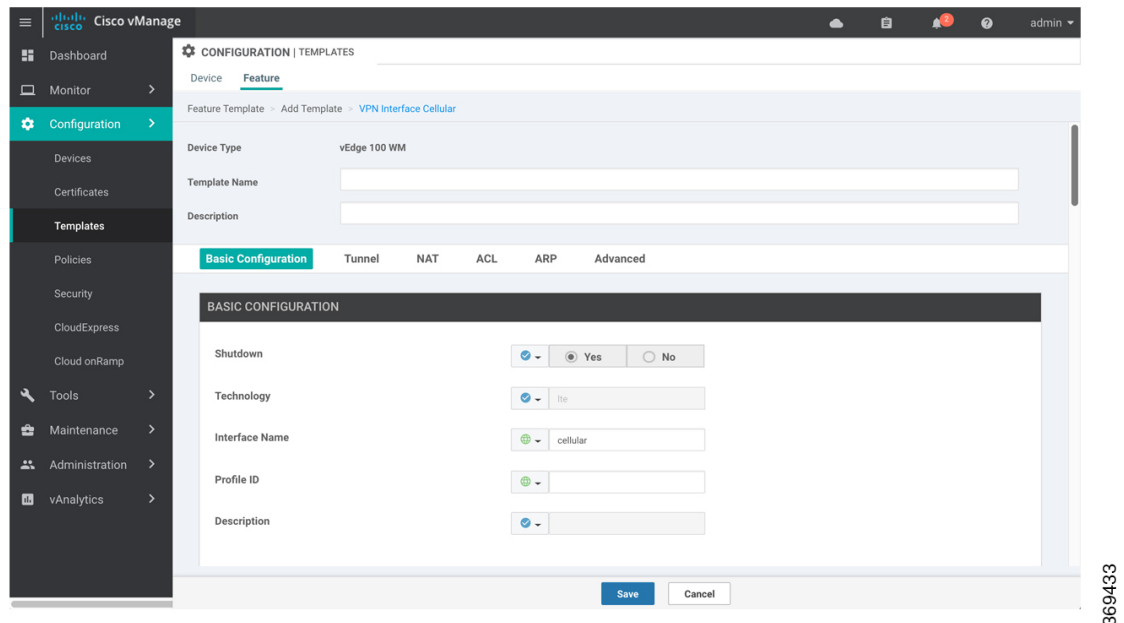


Note

If your deployment includes devices with cellular interface, you must include cellular controller templates in Cisco vManage, even if these templates are not used.

Create VPN Interface Cellular

1. In vManage NMS, select the **Configuration > Templates** screen.
2. In the **Device** tab, click **Create Template**.
3. From the **Create Template** drop-down, select **From Feature Template**.
4. From the **Device Model** drop-down, select the type of device for which you are creating the template.
5. Click the **Transport & Management VPN** tab or scroll to the Transport & Management VPN section.
6. Under Additional VPN 0 Templates, click **VPN Interface Cellular**.



369433

7. From the **VPN Interface Cellular** drop-down, click **Create Template**. The VPN Interface Cellular template form is displayed. The top of the form contains fields for naming the template, and the bottom contains fields for defining VPN Interface Cellular parameters.
8. In the **Template Name** field, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
9. In the **Template Description** field, enter a description of the template. The description can be up to 2048 characters and can contain only alphanumeric characters.

When you first open a feature template, for each parameter that has a default value, the scope is set to Default (indicated by a check mark), and the default setting or value is shown. To change the default or to enter a value, click the scope drop-down to the left of the parameter field.

Configure Basic Cellular Interface Functionality

To configure basic cellular interface functionality, select the **Basic Configuration** tab and configure the following parameters. Parameters marked with an asterisk are required to configure an interface. You must also configure a tunnel interface for the cellular interface.

Table 69:

Parameter Name	Description
Shutdown*	Click No to enable the interface.
Technology	Cellular technology. The default is lte . Other values are auto and cdma . For ZTP to work, the technology must be auto .
Interface Name*	Enter the name of the interface. It must be cellular0 .

Parameter Name	Description
Profile ID*	Enter the identification number of the cellular profile. This is the profile identifier that you configure in the Cellular-Profile template. <i>Range:</i> 1 through 15
Description	Enter a description of the cellular interface.
IPv4 Configuration	To configure a static address, click Static and enter an IPv4 address. To set the interface as a DHCP client so that the interface to receive its IP address from a DHCP server, click Dynamic . You can optionally set the DHCP distance to specify the administrative distance of routes learned from a DHCP server. The default DHCP distance is 1.
IPv6 Configuration	To configure a static address for an interface in VPN 0, click Static and enter an IPv6 address. To set the interface as a DHCP client so that the interface to receive its IP address from a DHCP server, click Dynamic . You can optionally set the DHCP distance to specify the administrative distance of routes learned from a DHCP server. The default DHCP distance is 1. You can optionally enable DHCP rapid commit, to speed up the assignment of IP addresses.
DHCP Helper	Enter up to four IP addresses for DHCP servers in the network, separated by commas, to have the interface be a DHCP helper. A DHCP helper interface forwards BOOTP (Broadcast) DHCP requests that it receives from the specified DHCP servers.
Block Non-Source IP	Click Yes to have the interface forward traffic only if the source IP address of the traffic matches the interface's IP prefix range.
Bandwidth Upstream	For transmitted traffic, set the bandwidth above which to generate notifications. <i>Range:</i> 1 through $(2^{32} / 2) - 1$ kbps
Bandwidth Downstream	For received traffic, set the bandwidth above which to generate notifications. <i>Range:</i> 1 through $(2^{32} / 2) - 1$ kbps
IP MTU*	Enter 1428 to set the MTU size, in bytes. This value must be 1428. You cannot use a different value.

To save the feature template, click **Save**.

CLI equivalent:

```

vpn 0
  interface cellular0
    bandwidth-downstream kbps bandwidth-upstream kbps block-non-source-ip ( ip address
ip-address/length | ip dhcp-client [dhcp-distance number])
    ( ipv6 address ipv6-prefix/length | ipv6 dhcp-client [dhcp-distance number]
[dhcp-rapid-comit])
    mtu 1428
    profile number
    no shutdown

```

Create a Tunnel Interface

To configure an interface in VPN 0 to be a WAN transport connection, you must configure a tunnel interface on the cellular interface. The tunnel, which provides security from attacks, is used to send the phone number.

At a minimum, select On and select a color for the interface, as described in the previous section. You can generally accept the system defaults for the remainder of the tunnel interface settings.

To configure a tunnel interface, select the Tunnel tab, set Tunnel Interface to On, and configure the following parameters. Parameters marked with an asterisk are required to configure a cellular interface.

Table 70:

Parameter Name	Description
Tunnel Interface*	Click On to create a tunnel interface.
Color*	Select a color for the TLOC. The color typically used for cellular interface tunnels is Ite .
Control Connection	The default is On, which establishes a control connection for the TLOC. If the router has multiple TLOCs, click No to have a tunnel not establish a TLOC.
Maximum Control Connections	Set the maximum number of vSmart controllers that the WAN tunnel interface can connect to. To have the tunnel establish no control connections, set the number to 0. <i>Range:</i> 0 through 8 Default: 2
vBond As STUN Server	Click On to enable Session Traversal Utilities for NAT (STUN) to allow the tunnel interface to discover its public IP address and port number when the router is located behind a NAT.
Exclude Control Group List	Set the identifiers of one or more vSmart controller groups that this tunnel is not allowed to establish control connections with. <i>Range:</i> 0 through 100
vManage Connection Preference	Set the preference for using the tunnel to exchange control traffic with the vManage NMS. <i>Range:</i> 0 through 9 Default: 5
Low-Bandwidth Link	Click On to set the tunnel interface as a low-bandwidth link. Default: Off
Allow Service	Click On or Off for each service to allow or disallow the service on the cellular interface.

To configure additional tunnel interface parameters, click Advanced Options and configure the following parameters:

Table 71:

Parameter Name	Description
GRE	Use GRE encapsulation on the tunnel interface. By default, GRE is disabled. If you select both IPsec and GRE encapsulations, two TLOCs are created for the tunnel interface that have the same IP addresses and colors, but that differ by their encapsulation.
IPsec	Use IPsec encapsulation on the tunnel interface. By default, IPsec is enabled. If you select both IPsec and GRE encapsulations, two TLOCs are created for the tunnel interface that have the same IP addresses and colors, but that differ by their encapsulation.
IPsec Preference	Enter a value to set the preference for directing traffic to the tunnel. A higher value is preferred over a lower value. <i>Range:</i> 0 through 4294967295 <i>Default:</i> 0
IPsec Weight	Enter a weight to use to balance traffic across multiple TLOCs. A higher value sends more traffic to the tunnel. <i>Range:</i> 1 through 255 <i>Default:</i> 1
Carrier	Select the carrier name or private network identifier to associate with the tunnel. <i>Values:</i> carrier1, carrier2, carrier3, carrier4, carrier5, carrier6, carrier7, carrier8, default <i>Default:</i> default
Bind Loopback Tunnel	Enter the name of a physical interface to bind to a loopback interface. The interface name has the format ge slot/port .
Last-Resort Circuit	Use the tunnel interface as the circuit of last resort
NAT Refresh Interval	Set the interval between NAT refresh packets sent on a DTLS or TLS WAN transport connection. <i>Range:</i> 1 through 60 seconds <i>Default:</i> 5 seconds
Hello Interval	Enter the interval between Hello packets sent on a DTLS or TLS WAN transport connection. <i>Range:</i> 100 through 10000 milliseconds <i>Default:</i> 1000 milliseconds (1 second)
Hello Tolerance	Enter the time to wait for a Hello packet on a DTLS or TLS WAN transport connection before declaring that transport tunnel to be down. <i>Range:</i> 12 through 60 seconds <i>Default:</i> 12 seconds

To save the feature template, click **Save**.

CLI equivalent:

```

vpn 0
  interface cellular0
    tunnel-interface allow-service service-name
  bind interface-name carrier carrier-name
    color color encapsulation (gre | ipsec)
    preference number
    weight number exclude-controller-group-list number hello-interval milliseconds
    hello-tolerance seconds hold-time milliseconds low-bandwidth-link
max-control-connections number last-resort-circuit nat-refresh-interval seconds
vbond-as-stun-server vmanage-connection-preference number

```

Configure the Cellular Interface as a NAT Device

To configure a cellular interface to act as a NAT device for applications such as port forwarding, select the NAT tab, click **On** and configure the following parameters:

Table 72: Configure the Cellular Interface as a NAT Device

Parameter Name	Description
NAT	Click On to have the interface act as a NAT device.
Refresh Mode	Select how NAT mappings are refreshed, either outbound or bidirectional (outbound and inbound). <i>Default: Outbound</i>
UDP Timeout	Specify when NAT translations over UDP sessions time out. <i>Range: 1 through 65536 minutesDefault: 1 minute</i>
TCP Timeout	Specify when NAT translations over TCP sessions time out. <i>Range: 1 through 65536 minutesDefault: 60 minutes (1 hour)</i>
Block ICMP	Select On to block inbound ICMP error messages. By default, a router acting as a NAT device receives these error messages. <i>Default: Off</i>
Respond to Ping	Select On to have the router respond to ping requests to the NAT interface's IP address that are received from the public side of the connection.

To create a port forwarding rule, click **Add New Port Forwarding Rule** and configure the following parameters. You can define up to 128 port-forwarding rules to allow requests from an external network to reach devices on the internal network.

Table 73:

Parameter Name	Description
Port Start Range	Enter a port number to define the port or first port in the range of interest. <i>Range: 0 through 65535</i>
Port End Range	Enter the same port number to apply port forwarding to a single port, or enter a larger number to apply it to a range of ports. <i>Range: 0 through 65535</i>
Protocol	Select the protocol to which to apply the port-forwarding rule, either TCP or UDP. To match the same ports for both TCP and UDP traffic, configure two rules.
VPN	Specify the private VPN in which the internal server resides. This VPN is one of the VPN identifiers in the overlay network. <i>Range: 0 through 65530</i>
Private IP	Specify the IP address of the internal server to which to direct traffic that matches the port-forwarding rule.

To save a port forwarding rule, click **Add**.

To save the feature template, click **Save**.

CLI equivalent:

```

vpn 0
interface cellular0
  nat block-icmp-error port-forward port-start port-number1 port-end port-number2
    proto (tcp | udp) private-ip-address ip address private-vpn vpn-id refresh
  (bi-directional | outbound)
  respond-to-ping tcp-timeout minutes
  udp-timeout minutes

```

Apply Access Lists

To configure a shaping rate to a cellular interface and to apply a QoS map, a rewrite rule, access lists, and policers to a router interface, select the ACL/QoS tab and configure the following parameters:

Table 74: Access Lists Parameters

Parameter Name	Description
Shaping rate	Configure the aggregate traffic transmission rate on the interface to be less than line rate, in kilobits per second (kbps).
QoS map	Specify the name of the QoS map to apply to packets being transmitted out the interface.
Rewrite rule	Click On , and specify the name of the rewrite rule to apply on the interface.
Ingress ACL – IPv4	Click On , and specify the name of an IPv4 access list to packets being received on the interface.
Egress ACL–IPv4	Click On , and specify the name of an IPv4 access list to packets being transmitted on the interface.
Ingress ACL – IPv6	Click On , and specify the name of an IPv6 access list to packets being received on the interface.
Egress ACL–IPv6	Click On , and specify the name of an IPv6 access list to packets being transmitted on the interface.
Ingress policer	Click On , and specify the name of the policer to apply to packets being received on the interface.
Egress policer	Click On , and specify the name of the policer to apply to packets being transmitted on the interface.

To save the feature template, click **Save**.

CLI equivalent:

```

vpn 0
interface cellular0
  access-list acl-name (in | out)
  ipv6 access-list acl-name (in | out)
  policer policer-name (in |out)
  qos-map name rewrite-rule name shaping-rate name

```

Add ARP Table Entries

To configure static Address Resolution Protocol (ARP) table entries on the interface, select the **ARP** tab. Then click **Add New ARP** and configure the following parameters:

Table 75:

Parameter Name	Description
IP Address	Enter the IP address for the ARP entry in dotted decimal notation or as a fully qualified host name.
MAC Address	Enter the MAC address in colon-separated hexadecimal notation.

To save the ARP configuration, click **Add**.

To save the feature template, click **Save**.

CLI equivalent:

```
vpn vpn-id interface irbnumber arp
    ip address ip-address mac mac-address
```

Configure Other Interface Properties

To configure other interface properties, select the **Advanced** tab and configure the following parameters.

Table 76: Cellular Interfaces Advanced Parameters

Parameter Name	Description
PMTU Discovery	Click On to enable path MTU discovery on the interface, to allow the router to determine the largest MTU size supported without requiring packet fragmentation.
TCP MSS	Specify the maximum segment size (MSS) of TCP SYN packets passing through the router. By default, the MSS is dynamically adjusted based on the interface or tunnel MTU such that TCP SYN packets are never fragmented. <i>Range:</i> 552 to 1460 bytes <i>Default:</i> None
Clear-Don't-Fragment	Click On to clear the Don't Fragment (DF) bit in the IPv4 packet header for packets being transmitted out the interface. When the DF bit is cleared, packets larger than that interface's MTU are fragmented before being sent.
Static Ingress QoS	Select a queue number to use for incoming traffic. <i>Range:</i> 0 through 7
ARP Timeout	Specify how long it takes for a dynamically learned ARP entry to time out. <i>Range:</i> 0 through 2678400 seconds (744 hours) <i>Default:</i> 1200 seconds (20 minutes)
Autonegotiate	Click Off to turn off autonegotiation. By default, an interface runs in autonegotiation mode.

Parameter Name	Description
TLOC Extension	Enter the name of a physical interface on the same router that connects to the WAN transport. This configuration then binds this service-side interface to the WAN transport. A second router at the same site that itself has no direct connection to the WAN (generally because the site has only a single WAN connection) and that connects to this service-side interface is then provided with a connection to the WAN.
Tracker	Enter the name of a tracker to track the status of transport interfaces that connect to the internet.
ICMP Redirect	Click Disable to disable ICMP redirect messages on the interface. By default, an interface allows ICMP redirect messages.

To save the feature template, click **Save**.

CLI equivalent:

```
vpn 0
 interface cellular0
 arp-timeout seconds
 [no] autonegotiate clear-dont-fragment icmp-redirect-disable mtu 1428
 pmtu static-ingress-qos number tcp-mss-adjust bytes
 tloc-extension interface-name tracker tracker-name
```

Release Information

Introduced in vManage NMS in Release 16.1. In Release 16.2, add circuit of last resort and its associated hold time. In Release 16.3, add support for IPv6. In Release 17.2.2, add support for tracker interface status. In Release 18.2, add support for disabling ICMP redirect messages.

Configuring Cellular Interfaces Using CLI

To configure a cellular interface on a Cisco vEdge device that has a cellular module:

1. Create a cellular profile:

```
vEdge(config)# cellular cellular number
vEdge(config-cellular)# profile profile-id
```

Each Cisco vEdge device has only one LTE module, so *number* must be 0. The profile identifier can be a value from 1 through 15.

2. If your ISP requires that you configure profile properties, configure one or more of the following:

```
vEdge(config-profile)# apn
 name
vEdge(config-profile)# auth auth-method
vEdge(config-profile)# ip-addr ip-address
vEdge(config-profile)# name name
vEdge(config-profile)# pdn-type type
vEdge(config-profile)# primary-dns ip-address
vEdge(config-profile)# secondary-dns ip-address
vEdge(config-profile)# user-name username
vEdge(config-profile)# user-pass password
```

1. Create the cellular interface:

```
vEdge(config)# vpn 0 interface cellular0
```

2. Enable the cellular interface:

```
vEdge(config-interface)# no shutdown
```

3. For cellular interfaces, you must use a DHCP client to dynamically configure the IP address. This is the default option. To explicitly configure this:

```
vEdge(config-interface)# ip dhcp-client [dhcp-distance number]
```

number is the administrative distance of routes learned from a DHCP server. You can configure it to a value from 1 through 255.

4. Associate the cellular profile with the cellular interface:

```
vEdge(config-interface)# profile profile-id
```

The profile identifier is the number you configured in Step 1.

5. Set the interface MTU:

```
vEdge(config-interface)# mtu bytes
```

The MTU can be 1428 bytes or smaller.

6. By default, the radio access technology (RAT) type is LTE. For 2G/3G networks, change it to CDMA:

```
vEdge(config-interface)# technology cdma
```

If you are using the interface for ZTP, change the technology to **auto**:

```
vEdge(config-interface)# technology auto
```

7. Configure any other desired interface properties.

8. Create a tunnel interface on the cellular interface:

```
vEdge(config-interface)# tunnel-interface  
vEdge(config-tunnel-interface)# color color  
vEdge(config-tunnel-interface)# encapsulation (gre | ipsec)
```

9. By default, the tunnel interface associated with a cellular interface is not considered to be the circuit of last resort. To allow the tunnel to be the circuit of last resort:

```
vEdge(config-tunnel-interface)# last-resort-circuit
```

When the interface is configured as a circuit of last resort, the cellular modem becomes dormant and no traffic is sent over the circuit. However, the cellular modem is kept in online mode so that the modem radio can be monitored at all times and to allow for faster switchover in the case the tunnel interface needs to be used as the last resort. By default, there is a delay of 7 seconds before switching back to the primary tunnel interface from a circuit of last resort. This delay is to ensure that the primary interface is once again fully operational and is not still flapping.

10. To minimize the amount of control plane keepalive traffic on the cellular interface, increase the Hello packet interval and tolerance on the tunnel interface:

```
vEdge(config-tunnel-interface)# hello-interval milliseconds  
vEdge(config-tunnel-interface)# hello-tolerance seconds
```

The default hello interval is 1000 milliseconds, and it can be a time in the range 100 through 600000 milliseconds (10 minutes). The default hello tolerance is 12 seconds, and it can be a time in the range 12 through 600 seconds (10 minutes). To reduce outgoing control packets on a TLOC, it is recommended that on the tunnel interface you set the hello interval to 60000 milliseconds (10 minutes) and the hello tolerance to 600 seconds (10 minutes) and include the **no track-transport** to disable regular checking of the DTLS connection between the Cisco vEdge device and the vBond orchestrator. For a tunnel

connection between a Cisco vEdge device and any controller device, the tunnel uses the hello interval and tolerance times configured on the Cisco vEdge device. This choice is made to minimize the amount of traffic sent over the tunnel, to allow for situations where the cost of a link is a function of the amount of traffic traversing the link. The hello interval and tolerance times are chosen separately for each tunnel between a Cisco vEdge device and a controller device. Another step taken to minimize the amount of control plane traffic is to not send or receive OMP control traffic over a cellular interface when other interfaces are available. This behavior is inherent in the software and is not configurable.

11. If the Cisco vEdge device has two or more cellular interfaces, you can minimize the amount of traffic between the vManage NMS and the cellular interfaces by setting one of the interfaces to be the preferred one to use when sending updates to the vManage NMS and receiving configurations from the vManage NMS:

```
vEdge(config-tunnel-interface) # vmanage-connection-preference number
```

The preference can be a value from 0 through 8. The default preference is 5. To have a tunnel interface never connect to the vManage NMS, set the number to 0. At least one tunnel interface on the Cisco vEdge device must have a nonzero vManage connection preference.

12. Configure any other desired tunnel interface properties.
13. To minimize the amount of data plane keepalive traffic on the cellular interface, increase the BFD Hello packet interval:

```
vEdge(bfd-color-lte) # hello-interval milliseconds
```

The default hello interval is 1000 milliseconds (1 second), and it can be a time in the range 100 through 300000 milliseconds (5 minutes).

To determine the status of the cellular hardware, use the **show cellular status** command.

To determine whether a Cisco vEdge device has a cellular module, use the **show hardware inventory** command.

To determine whether a cellular interface is configured as a last-resort circuit, use the **show control affinity config** and **show control local-properties** commands.



Note If you want to remove a property from the cellular profile, delete the profile entirely from the configuration, and create it again with only the required parameters.



Note When you activate the configuration on a Cisco vEdge device with cellular interfaces, the primary interfaces (that is, those interfaces not configured as circuits of last resort) and the circuit of last resort come up. In this process, all the interfaces begin the process of establishing control and BFD connections. When one or more of the primary interfaces establishes a TLOC connection, the circuit of last resort shuts itself down because it is not needed. During this shutdown process, the circuit of last resort triggers a BFD TLOC Down alarm and a Control TLOC Down alarm on the Cisco vEdge device. These two alarms are cleared only when all the primary interfaces lose their BFD connections to remote nodes and the circuit of last resort activates itself. This generation and clearing of alarms is expected behavior.

Best Practices for Configuring Cellular Interfaces

Cellular technology on Cisco vEdge devices can be used in a number of ways:

- **Circuit of last resort**—You can use a cellular interface as a backup circuit on a Cisco vEdge device. Such a circuit is activated only if all transport links on the Cisco vEdge device fail. In this mode the radio interface is turned off, and no control or data connections exist over the cellular interface. To configure an cellular interface to be a circuit of last resort, include the **last-resort-circuit** command when you configure the cellular interface's tunnel interface.
- **Active circuit**—You can choose to use a cellular interface as an active circuit, perhaps because it is the only last-mile circuit or to always keep the cellular interface active so that you can measure the performance of the circuit. In this scenario the amount of bandwidth utilized to maintain control and data connections over the cellular interface can become a concern. Here are some best practices to minimize bandwidth usage over a cellular interface:
 - **When a device with cellular interface is deployed as a spoke, and data tunnels are established in a hub-and-spoke manner, you can configure the cellular interface as a low-bandwidth interface.** To do this, include the **low-bandwidth-link** command when you configure the cellular interface's tunnel interface. When the cellular interface is operating as a low-bandwidth interface, the device spoke site is able to synchronize all outgoing control packets. The spoke site can also proactively ensure that no control traffic, except for routing updates, is generated from one of the remote hub nodes. Routing updates continue to be sent, because they are considered to be critical updates.
 - **Increase control packet timers**—To minimize control traffic on a cellular interface, you can decrease how often protocol update messages are sent on the interface. OMP sends Update packets every second, by default. You can increase this interval to a maximum of 65535 seconds (about 18 hours) by including the **omp timers advertisement-interval** configuration command. BFD sends Hello packets every second, by default. You can increase this interval to a maximum of 5 minutes (30000 milliseconds) by including the **bfd color hello-interval** configuration command. (Note that you specify the OMP Update packet interval in seconds and the BFD Hello packet interval in milliseconds.)
 - **Prioritize vManage control traffic over a non-cellular interface**—When a Cisco vEdge device has both cellular and non-cellular transport interfaces, by default, the Cisco vEdge device chooses one of the interfaces to use to exchange control traffic with the vManage NMS. You can configure the Cisco vEdge device to never use the cellular interface to exchange traffic with the NMS, or you can configure a lower preference for using the cellular interface for this traffic. You configure the preference by including the **vmanage-connection-preference** command when configuring the tunnel interface. By default, all tunnel interface have a vManage connection preference value of 5. The value can range from 0 through 8, where a higher value is more preferred. A tunnel with a preference value of 0 can never exchange control traffic with the vManage NMS.



Note At least one tunnel interface on the Cisco vEdge device must have a non-0 vManage connection preference value. Otherwise, the device has no control connections.

Interface CLI Reference

CLI commands for configuring and monitoring system-wide parameters, interfaces, and SNMP on vEdge routers and vSmart controllers.

Interface Configuration Commands

Use the following commands to configure interfaces and interface properties in the Cisco SD-WAN overlay network. Interfaces must be configured on a per-VPN basis.

```

vpn vpn-id
  interface interface-name
    access-list acl-list (on vEdge routers only)
    arp
      ip ip-address mac mac-address
    arp-timeout seconds (on vEdge routers only)
    autonegotiate (on vEdge routers only)
    block-non-source-ip (on vEdge routers only)
    clear-dont-fragment
    dead-peer-detection interval seconds retries number (on vEdge routers only)
    description text
    dhcp-helper ip-address (on vEdge routers only)
    dhcp-server (on vEdge routers only)
      address-pool prefix/length
      exclude ip-address
      lease-time seconds
      max-leases number
      offer-time minutes
    options
      default-gateway ip-address
      dns-servers ip-address
      domain-name domain-name
      interface-mtu mtu
      tftp-servers ip-address
    static-lease mac-address ip ip-address host-name hostname
dot1x
  accounting-interval seconds
  acct-req-attr attribute-number (integer integer | octet octet | string string)
  auth-fail-vlan vlan-id
  auth-order (mab | radius)
  auth-reject-vlan vlan-id
  auth-req-attr attribute-number (integer integer | octet octet | string string)
  control-direction direction
  das
    client ip-address
    port port-number
    require-timestamp
    secret-key password
    time-window seconds
    vpn vpn-id
  default-vlan vlan-id
  guest-vlan vlan-id
  host-mode (multi-auth | multi-host | single-host)
  mac-authentication-bypass
    allow mac-addresses
    server
  nas-identifier string
  nas-ip-address ip-address
  radius-servers tag
  reauthentication minutes
  timeout

```

```

    inactivity minutes
    wake-on-lan
duplex (full | half)
flow-control (bidirectional | egress | ingress)
ike (on vEdge routers only)
    authentication-type type
        local-id id
        pre-shared-secret password
        remote-id id
    cipher-suite suite
    group number
    mode mode
    rekey seconds
    version number
(ip address prefix/length | ip dhcp-client [dhcp-distance number])
(ipv6 address prefix/length | ipv6 dhcp-client [dhcp-distance number] [dhcp-rapid-commit])

ip address-list prefix/length (on vSmart controller containers only)
ip secondary-address ipv4-address (on vEdge routers only)
ipsec (on vEdge routers only)
    cipher-suite suite
    perfect-forward-secrecy pfs-setting
    rekey seconds
    replay-window number
keepalive seconds retries (on vEdge routers only)
mac-address mac-address
mtu bytes
nat (on vEdge routers only)
    block-icmp-error
    block-icmp-error
    direction (inside | outside)
    log-translations
    [no] overload
    port-forward port-start port-number1 port-end port-number2
        proto (tcp | udp) private-ip-address ip address private-vpn vpn-id
    refresh (bi-directional | outbound)
    respond-to-ping
    static source-ip ip-address1 translate-ip ip-address2 (inside | outside)
    static source-ip ip-address1 translate-ip ip-address2 source-vpn vpn-id protocol (tcp
| udp) source-port number translate-port number
    tcp-timeout minutes
    udp-timeout minutes
pmtu (on vEdge routers only)
policer policer-name (on vEdge routers only)
ppp (on vEdge routers only)
    ac-name name
    authentication (chap | pap) hostname name password password
pppoe-client (on vEdge routers only)
    ppp-interface name
profile profile-id (on vEdge routers only)
qos-map name (on vEdge routers only)
rewrite-rule name (on vEdge routers only)
shaping-rate name (on vEdge routers only)
shutdown
speed speed
static-ingress-qos number (on vEdge routers only)
tcp-mss-adjust bytes
technology technology (on vEdge routers only)
tloc-extension interface-name (on vEdge routers only)
tracker tracker-name (on vEdge routers only)
tunnel-interface
    allow-service service-name
    bind geslot/port (on vEdge routers only)
    carrier carrier-name

```

```

color color [restrict]
connections-limit number
encapsulation (gre | ipsec) (on vEdge routers only)
  preference number
  weight number
hello-interval milliseconds
hello-tolerance seconds
low-bandwidth-link (on vEdge routers only)
max-control-connections number (on vEdge routers only)
nat-refresh-interval seconds
port-hop
vbond-as-stun-server (on vEdge routers only)
vmanage-connection-preference number (on vEdge routers only)
tunnel-destination ip-address (GRE interfaces; on vEdge routers only)
(tunnel-source ip-address | tunnel-source-interface interface-name) (GRE interfaces;
on vEdge routers only)
(tunnel-source ip-address | tunnel-source-interface interface-name) (IPsec interfaces;
on vEdge routers only)
upgrade-confirm minutes
vrrp group-name (on vEdge routers only)
  priority number
  timer seconds
track-omp

```

Interface Monitoring Commands

Use the following commands to monitor interfaces:

show dhcp interface

show dhcp server

show interface

show interface arp-stats

show interface errors

show interface packet-sizes

show interface port-stats

show interface queue

show interface statistics

show vrrp

System Configuration Commands

Use the following commands to configure system-wide parameters:

```

banner
  login "text"
  motd "text"
system
  aaa
  admin-auth-order (local | radius | tacacs)
  auth-fallback
  auth-order (local | radius | tacacs)
  logs
  audit-disable
  netconf-disable

```

```

radius-servers tag
user user-name
  group group-name
  password password
  usergroup group-name
  task (interface | policy | routing | security | system) (read | write)
admin-tech-on-failure
archive
  interval minutes
  path file-path/filename
  ssh-id-file file-path/filename
  vpn vpn-id
clock
  timezone timezone
console-baud-rate rate
control-session-pps rate
description text
device-groups group-name
domain-id domain-id
eco-friendly-mode (on vEdge Cloud routers only)
gps-location (latitude decimal-degrees | longitude decimal-degrees)
host-name string
host-policer-pps rate (on vEdge routers only)
icmp-error-pps rate
idle-timeout minutes
iptables-enable
location string
logging
  disk
    enable
    file
      name filename
      rotate number
      size megabytes
      priority priority
  host
    name (name | ip-address)
    port udp-port-number
    priority priority
    rate-limit number interval seconds
multicast-buffer-percent percentage (on vEdge routers only)
ntp
  keys
    authentication key-id md5 md5-key
    trusted key-id
  server (dns-server-address | ipv4-address)
    key key-id
    prefer
    source-interface interface-name
    version number
    vpn vpn-id
organization-name string
port-hop
port-offset number
radius
  retransmit number
  server ip-address
    auth-port port-number
    priority number
    secret-key key
    source-interface interface-name
    tag tag
    vpn vpn-id
  timeout seconds

```

```

route-consistency-check (on vEdge routers only)
site-id site-id
sp-organization-name name (on vBond orchestrators and vSmart controllers only)
system-ip ip-address
system-tunnel-mtu bytes
tacacs
  authentication authentication-type
  server ip-address
    auth-port port-number
    priority number
    secret-key key
    source-interface interface-name
  vpn vpn-id
  timeout seconds
tcp-optimization-enabled
timer
  dns-cache-timeout minutes
track-default-gateway
track-interface-tag number (on vEdge routers only)
track-transport
tracker tracker-name
  endpoint-dns-name dns-name
  endpoint-ip ip-address
  interval seconds
  multiplier number
  threshold milliseconds
upgrade-confirm minutes
[no] usb-controller (on vEdge 1000 and vEdge 2000 routers only)
vbond (dns-name | ip-address) [local] [port number] [ztp-server]

```

System Monitoring Commands on a Cisco vEdge device

Use the following commands to monitor system-wide parameters:

show aaa usergroup

show control local-properties

show logging

show ntp associations

show ntp peer

show orchestrator local-properties

show running-config system

show system status

show uptime

show users