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 2
- Configure Interfaces in the WAN Transport VPN (VPN 0), on page 6
- Extend the WAN Transport VPN, on page 9
- Configure GRE Interfaces and Advertise Services to Them, on page 12
- Configure the System Interface, on page 16
- Configure Control Plane High Availability, on page 17
- Configure Other Interfaces, on page 17
- Role-Based Access Control by VPN, on page 19
- Configure Interface Properties, on page 22
- Enable DHCP Server using Cisco vManage, on page 24
- Configuring PPPoE, on page 29
- Configuring VRRP , on page 35
- Network Interface Configuration Examples for Cisco vEdge Devices, on page 38
- Configure VPN Interfaces Using vManage, on page 55

• Interface CLI Reference, on page 114

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.

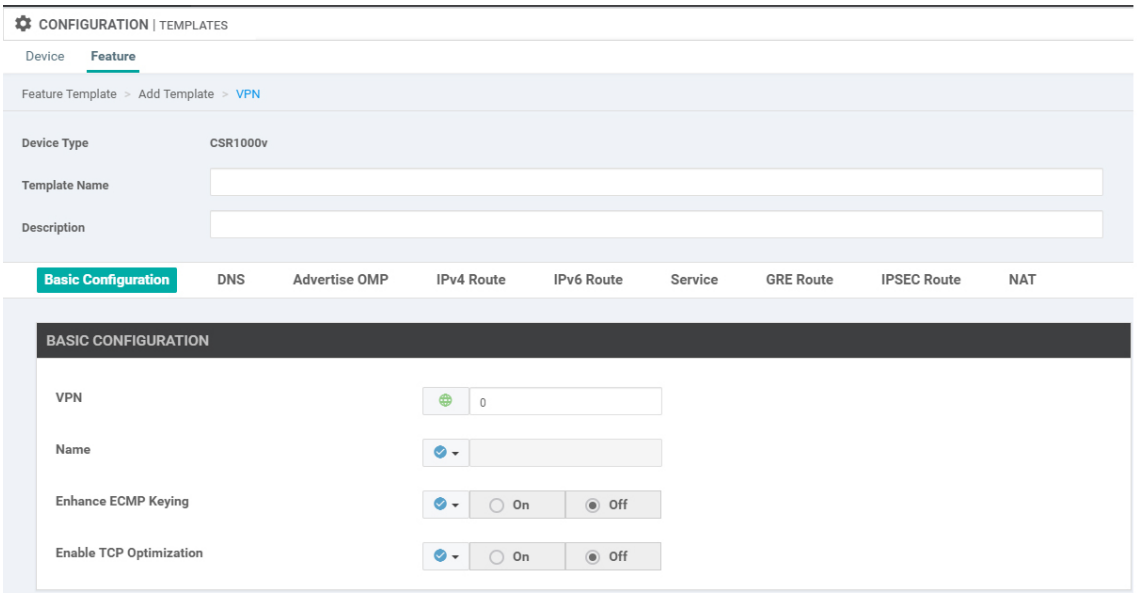

Step 7 In the **Template Name** field, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.

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	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. 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 To change the default key, type a new string and move the cursor out of the Enter Key box. Examples of device-specific parameters are system IP address, hostname, GPS location, and site ID.
Global	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.

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*	Enter the numeric identifier of the VPN. Range for Cisco vEdge devices: 0 through 65530 Values for Cisco vSmart Controller and Cisco vManage devices: 0, 512
Name	Enter a name for the VPN.
Enhance ECMP keying (Cisco vEdge devices only)	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.

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 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 packets matching both implicit and explict 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' 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
                                IF      IF                                                                TCP
                                ADMIN   OPER    ENCAP                                       SPEED         MSS                  RX       TX
VPN  INTERFACE  IP ADDRESS      STATUS  STATUS  TYPE   PORT TYPE  MTU   HWADDR             MBPS   DUPLEX  ADJUST  UPTIME       PACKETS  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

                                      IF      IF                                                                  TCP
                AF                    ADMIN   OPER    ENCAP                                       SPEED           MSS                 RX       TX
VPN  INTERFACE  TYPE  IPV6 ADDRESS      STATUS  STATUS  TYPE   PORT TYPE  MTU   HWADDR             MBPS    DUPLEX  ADJUST  UPTIME      PACKETS  PACKETS  LINK LOCAL ADDRESS
-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------
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.

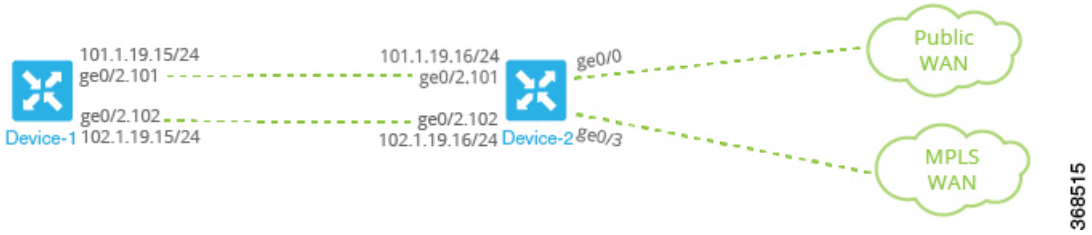

When you extend the WAN transport VPN, no BFD sessions are established between the two collocated vEdge routers.

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 geslot / 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 TLOConly 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
```

```
  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                    PEER                                                  CONTROLLER
PEER     PEER     PEER            SITE       DOMAIN     PEER     PRIVATE  PEER           PUBLIC                                                GROUP
TYPE     PROTOCOL SYSTEM IP       ID         ID         PRIVATE IP PORT     PUBLIC IP      PORT   LOCAL COLOR     STATE      UPTIME          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
                                       SOURCE TLOC   REMOTE TLOC                    DST PUBLIC       DST PUBLIC         DETECT      TX                          TRANSI-
SYSTEM IP         SITE ID  STATE       COLOR         COLOR         SOURCE IP        IP               PORT        ENCAP  MULTIPLIER  INTERVAL(msec)  UPTIME      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
172.16.255.21     100      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 a 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-namelocal
```

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.

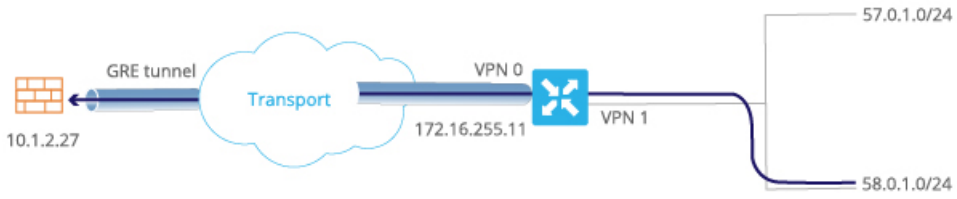

368516

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

                                  IF      IF                                                                      TCP
                                  ADMIN   OPER    ENCAP                                           SPEED           MSS                  RX       TX
VPN  INTERFACE  IP ADDRESS        STATUS  STATUS  TYPE   PORT TYPE  MTU   HWADDR                 MBPS   DUPLEX   ADJUST  UPTIME       PACKETS  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
                                                                 REMOTE   REMOTE
     IF                                  ADMIN  OPER   KA       TX       RX       TX       RX       TX      RX
VPN  NAME  SOURCE IP  DEST IP           STATE  STATE  ENABLED  PACKETS  PACKETS  PACKETS  PACKETS  ERRORS  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
                                IF      IF                                                              TCP
                                ADMIN   OPER    ENCAP                                           SPEED   MSS                     RX      TX
VPN  INTERFACE  IP ADDRESS      STATUS  STATUS  TYPE   PORT TYPE  MTU   HWADDR              MBPS    DUPLEX  ADJUST  UPTIME      PACKETS 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
                                  IF      IF                                                                  TCP
                                  ADMIN   OPER    ENCAP                                               SPEED   MSS                   RX       TX
VPN  INTERFACE  IP ADDRESS        STATUS  STATUS  TYPE   PORT TYPE  MTU   HWADDR             MBPS    DUPLEX  ADJUST  UPTIME      PACKETS  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
                              IF      IF                                                          TCP
                              ADMIN   OPER    ENCAP  PORT                                 SPEED   MSS                   RX       TX
VPN  INTERFACE  IP ADDRESS    STATUS  STATUS  TYPE   TYPE     MTU   HWADDR              MBPS    DUPLEX  ADJUST  UPTIME      PACKETS  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 device s, 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

                             IF      IF                                                                    TCP
                             ADMIN   OPER    ENCAP  PORT                                    SPEED          MSS                  RX       TX
VPN  INTERFACE  IP ADDRESS   STATUS  STATUS  TYPE   TYPE     MTU   HWADDR                   MBPS   DUPLEX  ADJUST  UPTIME      PACKETS  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:

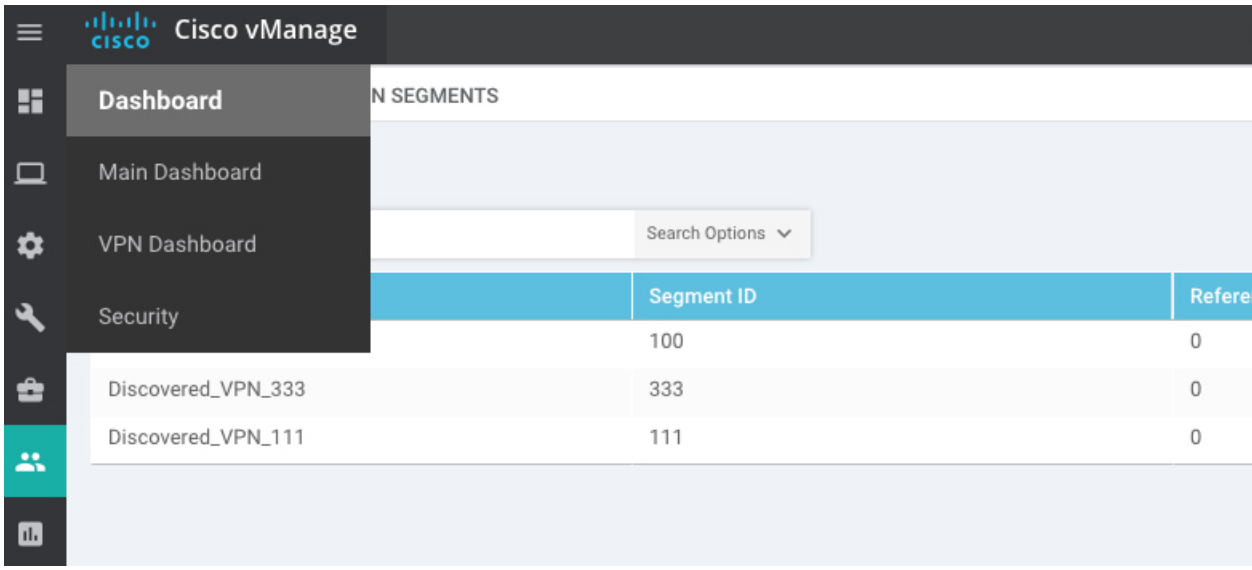

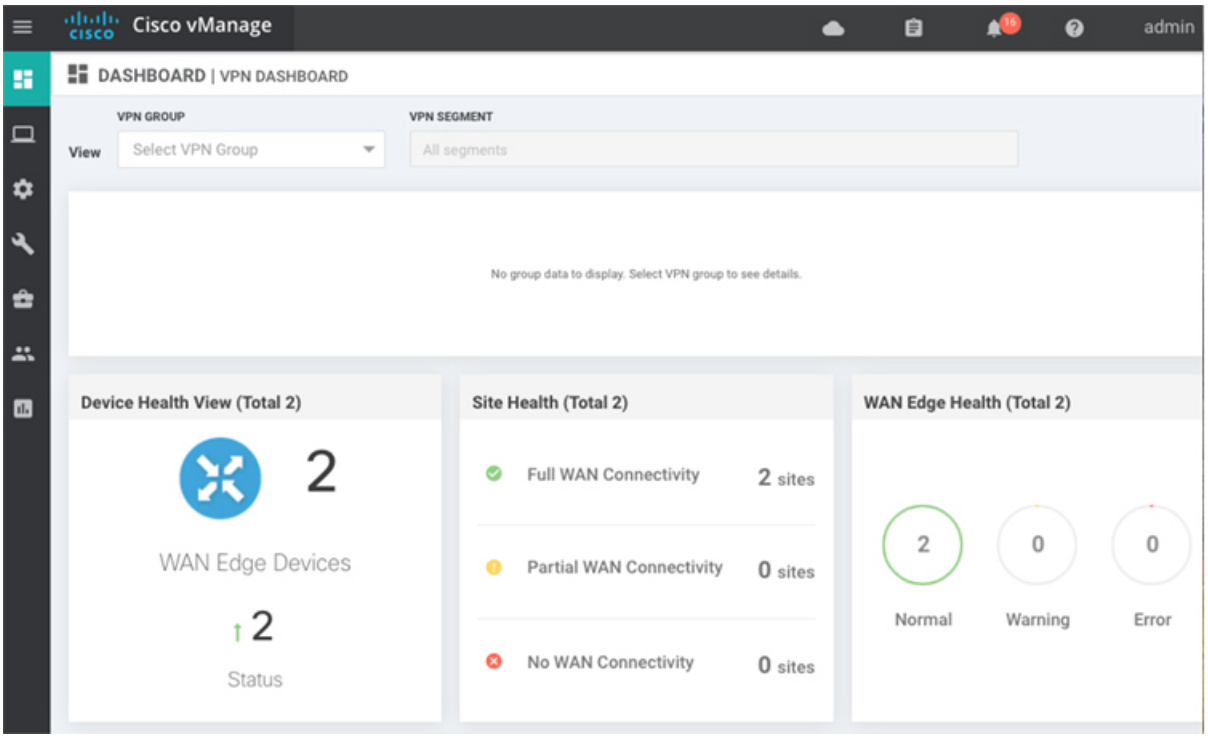

Configure and Manage VPN Segments

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.

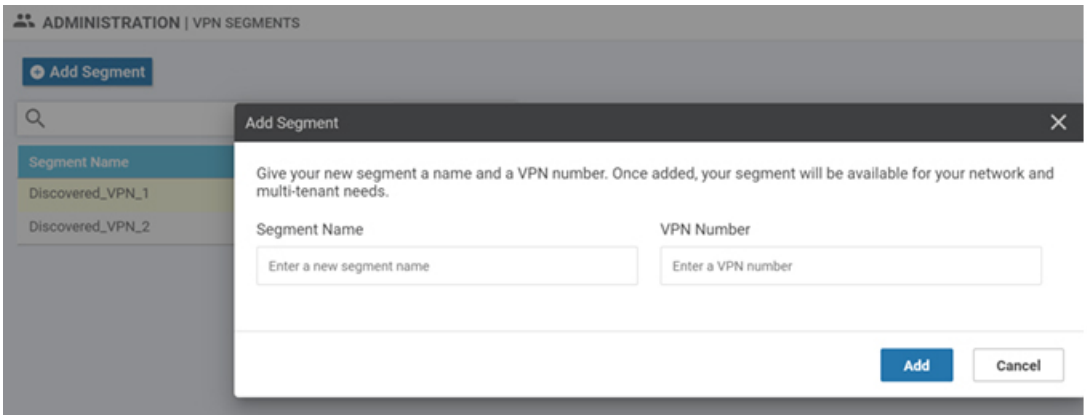

4. Enter the name of the segment in the **Segment Name** field.
5. Enter the number of VPNs you want to configure in VPN Number field.

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.

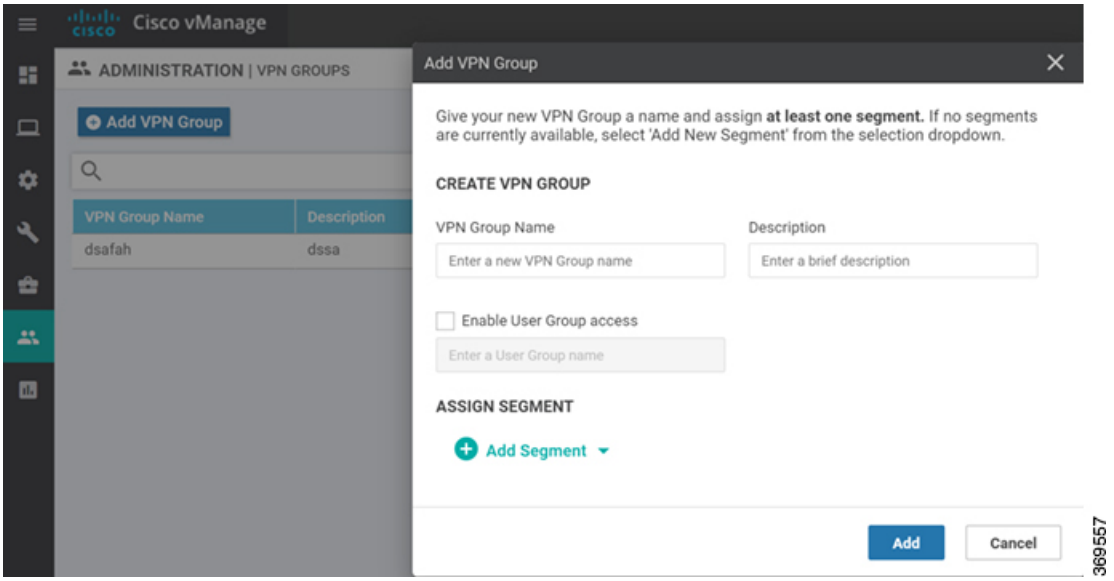

5. In the Create VPN Group pane, Enter VPN group name in the **VPN Group Name** field.
6. Enter a brief description of the VPN in the **Description** field.
7. Enable the user group access checkbox and enter the User Group Name.
8. In the Assign Segment pane, click on Add Segment drop-down to add new 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.

2. 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.
3. Click on **Add User** to add a new user.
4. In the Add New User page, add **Full Name, Username, Password,** and **Confirm Password details**.
5. In the User Group drop-down, select the user group where you want to add a user.
6. If you want to add a User Group, click on **Add User Group** button.

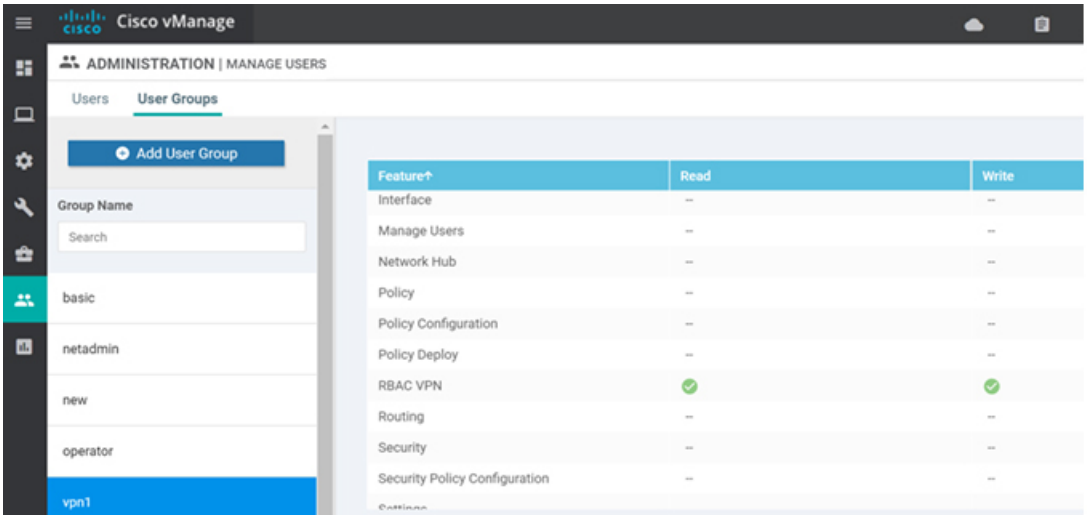

7. Enter the user group name in the **Group Name** field.
8. Select the Read or Write checkbox that you want to assign to a user group as shown in the figure.

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
              DEV
HW TYPE       INDEX  VERSION  PART NUMBER       SERIAL NUMBER     DESCRIPTION
-------------------------------------------------------------------------------------------------------------------
Chassis       0      3.1      vEdge-1000        11OD145130001     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 PIM 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

                                IF      IF                                                                      TCP
                                ADMIN   OPER    ENCAP                                                   SPEED           MSS                     RX       TX
VPN  INTERFACE  IP ADDRESS      STATUS  STATUS  TYPE    PORT TYPE  MTU   HWADDR             SPEED  MBPS   DUPLEX  ADJUST  UPTIME          PACKETS  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.

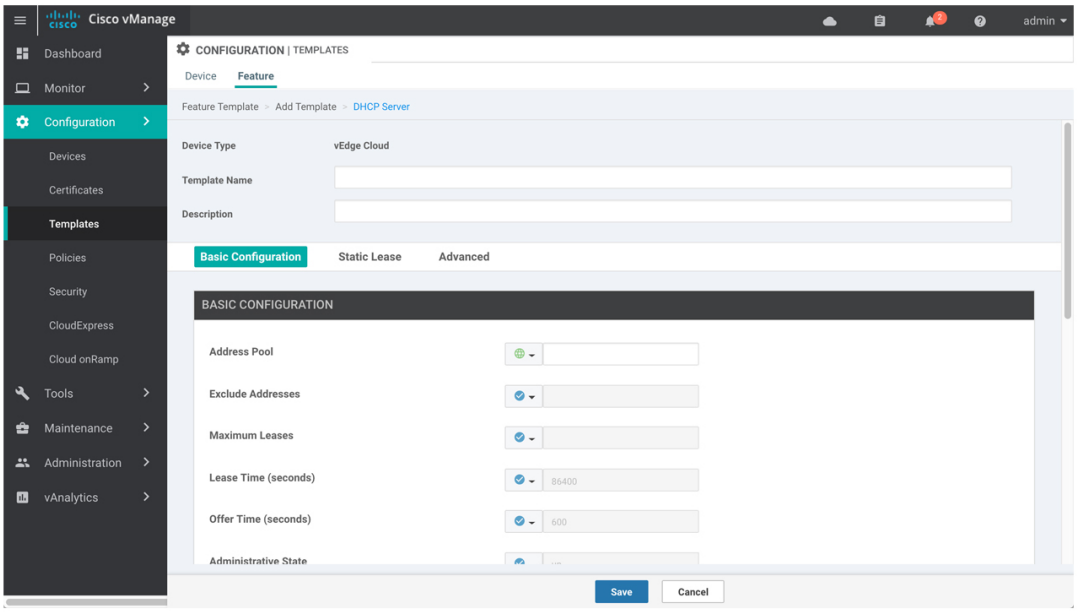

369419

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 1:

Parameter Name	Description
Address Pool*	Enter the IPv4 prefix range, in the format *prefix/length*, 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.*Range:* 0 through 4294967295
Lease Time	Specify how long a DHCP-assigned IP address is valid.*Range:* 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.*Range:* 0 through 4294967295 seconds*Default:* 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 2:

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 3:

Parameter Name	Description
Interface MTU	Specify the maximum MTU size of packets on the interface.*Range:* 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 4:

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 Confguration	Assign an IP address to the physical interface: • 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 aname 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 $4$OHHjdmsC6M8zj4BgLEFXKw==
  !
  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

             PPPOE      INTERFACE               PRIMARY  SECONDARY
VPN  IFNAME  INTERFACE  IP         GATEWAY IP   DNS      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

             SESSION                                          PPP                    SERVICE

VPN  IFNAME  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
```

```
                  RX       TX
VPN  INTERFACE  IP ADDRESS   STATUS  STATUS  TYPE   TYPE     MTU   HWADDR             MBPS   DUPLEX
ADJUST  UPTIME       PACKETS  PACKETS
------------------------------------------------------------------------------------------------------
1    ge0/6.2    10.2.2.3/24  Up      Up      vlan   service  1496  00:0c:29:ab:b7:94  10     full    0
      0:00:05:52  0        357
1    ge0/6.3    10.2.3.5/24  Down    Down    vlan   service  1496  00:0c:29:ab:b7:94  -      -       0
      -           0        0

vEdge1# show vrrp interfaces

                                                                                                  MASTER
                      TRACK   PREFIX
              GROUP  VIRTUAL                                    VRRP    OMP    ADVERTISEMENT  DOWN
                      PREFIX  LIST
VPN  IF NAME  ID     IP         VIRTUAL MAC        PRIORITY  STATE   STATE  TIMER          TIMER   LAST
 STATE CHANGE TIME     LIST    STATE
------------------------------------------------------------------------------------------------------
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
                                                                                                  MASTER
                           TRACK   PREFIX
               GROUP                                        VRRP    OMP    ADVERTISEMENT  DOWN
                           PREFIX  LIST
VPN  IF NAME   ID     VIRTUAL IP     VIRTUAL MAC        PRIORITY  STATE   STATE  TIMER          TIMER
 LAST STATE CHANGE TIME     LIST    STATE
------------------------------------------------------------------------------------------------------
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

                                                                                  MASTER
           TRACK   PREFIX
GROUP                                          VRRP    OMP    ADVERTISEMENT  DOWN
           PREFIX  LIST
ID     VIRTUAL IP   VIRTUAL MAC        PRIORITY  STATE   STATE  TIMER          TIMER   LAST STATE CHANGE
 TIME     LIST    STATE
```

```
-------------------------------------------------------------------------------------------------------------------------------
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

                                                                                                                MASTER
                      TRACK   PREFIX
          GROUP                                                 VRRP    OMP    ADVERTISEMENT  DOWN
                      PREFIX  LIST
IF NAME   ID     VIRTUAL IP     VIRTUAL MAC        PRIORITY  STATE   STATE  TIMER          TIMER   LAST
 STATE CHANGE TIME     LIST    STATE
-------------------------------------------------------------------------------------------------------------------------------
ge0/1.1   1      10.20.32.1     00:0c:bd:08:2b:a5  110       master  up     1              3
2016-01-13T00:22:15+00:00  -       -
ge0/1.5   5      10.20.32.193   00:0c:bd:08:2b:a5  110       master  up     1              3
2016-01-13T00:22:15+00:00  -       -
ge0/1.10  10     10.20.32.225   00:0c:bd:08:2b:a5  110       master  up     1              3
2016-01-13T00:22:15+00:00  -       -
ge0/1.15  15     10.20.33.1     00:0c:bd:08:2b:a5  100       backup  up     1              3
2016-01-13T03:10:56+00:00  -       -
ge0/1.20  20     10.20.34.1     00:0c:bd:08:2b:a5  110       master  up     1              3
2016-01-13T00:22:16+00:00  -       -
ge0/1.25  25     10.20.35.1     00:0c:bd:08:2b:a5  100       backup  up     1              3
2016-01-13T03:10:56+00:00  -       -
ge0/1.30  30     10.20.35.129   00:0c:bd:08:2b:a5  100       master  up     1              3
2016-01-13T00:22:16+00:00  -       -

Router-2# show vrrp vpn 100 interfaces groups 15

                                                                                              MASTER
                    TRACK   PREFIX
          GROUP                                             VRRP    OMP    ADVERTISEMENT  DOWN
                    PREFIX  LIST
IF NAME   ID     VIRTUAL IP   VIRTUAL MAC        PRIORITY  STATE   STATE  TIMER          TIMER   LAST
STATE CHANGE TIME     LIST    STATE
-------------------------------------------------------------------------------------------------------------------------------
ge0/0.15  15     10.20.33.1   00:0c:bd:08:2b:a5  100       backup  up     1              3
2016-01-13T03:10:56+00:00  -       -
```

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). TheCisco 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.

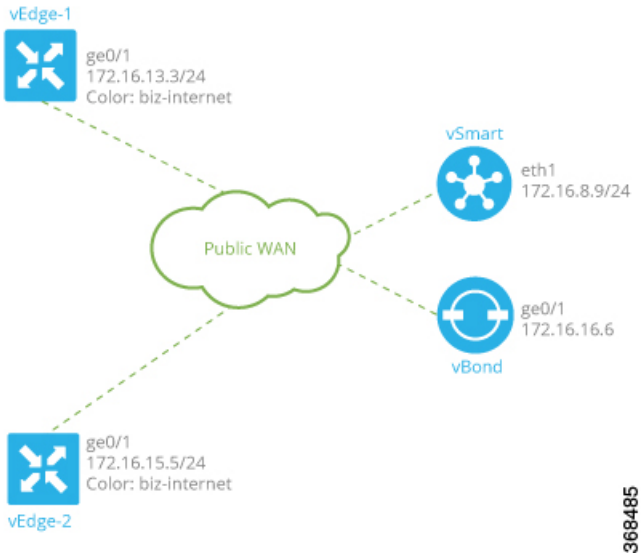

For Cisco vEdge device-1, the interface ge0/1 connects to the public WAN, so it is the interface that is configured as a tunnel interface. The tunnel has a color of biz-internet, and the encapsulation used for data traffic is IPsec. The Cisco SD-WAN software creates a single TLOC for this interface, 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

                                IF      IF                                                                      TCP
                                ADMIN   OPER    ENCAP                                           SPEED           MSS
   RX       TX
VPN  INTERFACE  IP ADDRESS      STATUS  STATUS  TYPE   PORT TYPE  MTU   HWADDR            MBPS   DUPLEX  ADJUST  UPTIME
   PACKETS  PACKETS
-------------------------------------------------------------------------------------------------------------------------------
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     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     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
                                          SOURCE TLOC      REMOTE TLOC                      DST PUBLIC
   DST PUBLIC         DETECT      TX
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                                                                                                 PUBLIC
                  PRIVATE  BFD
FAMILY   TLOC IP          COLOR            ENCAP  FROM PEER        STATUS    PUBLIC IP        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
         172.16.255.14    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
         172.16.255.16    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
         172.16.255.21    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   <
```

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.

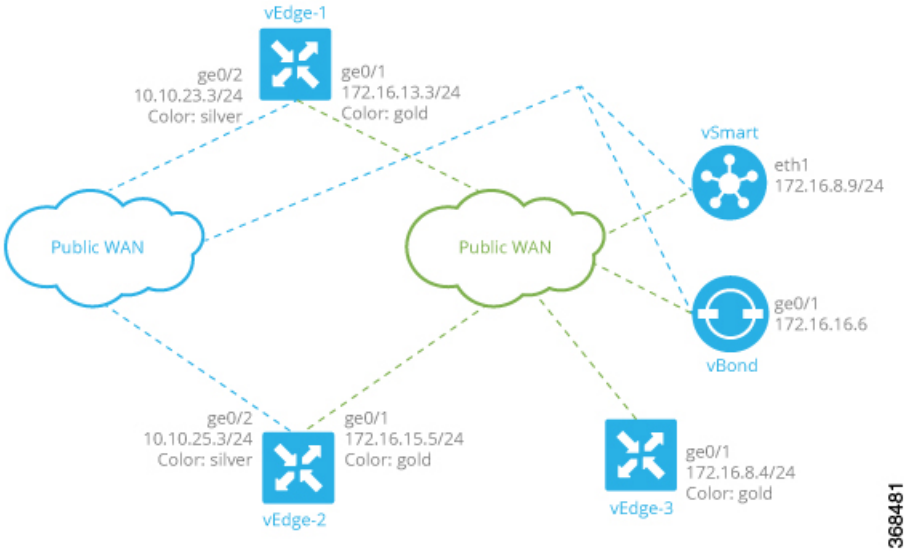

The configurations for vEdge-1 and vEdge-2 are similar. We configure two tunnel interfaces, one with color **silver** and the other with color **gold**, and we configure static default routes for both tunnel interfaces. Here is the configuration for vEdge-1:

```
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).

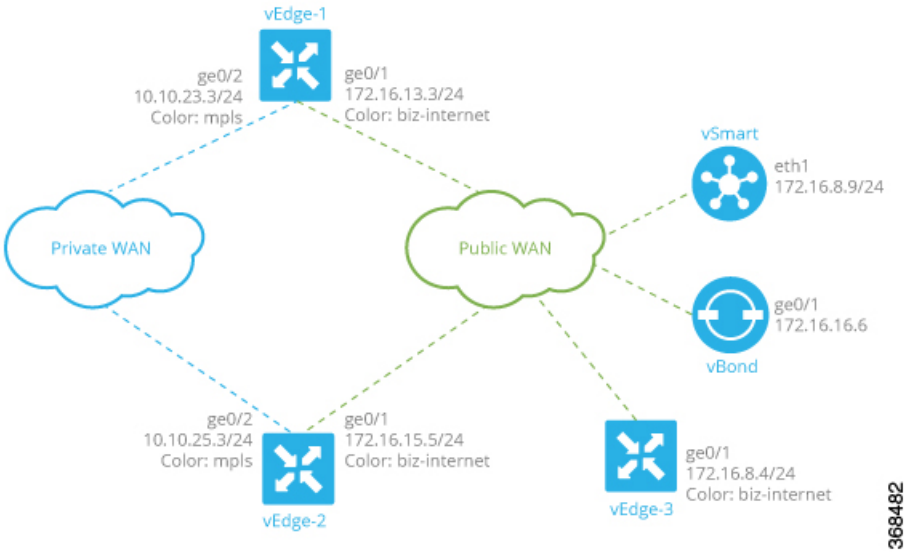

In this example topology, we need to ensure the following:

- 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 never 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 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.

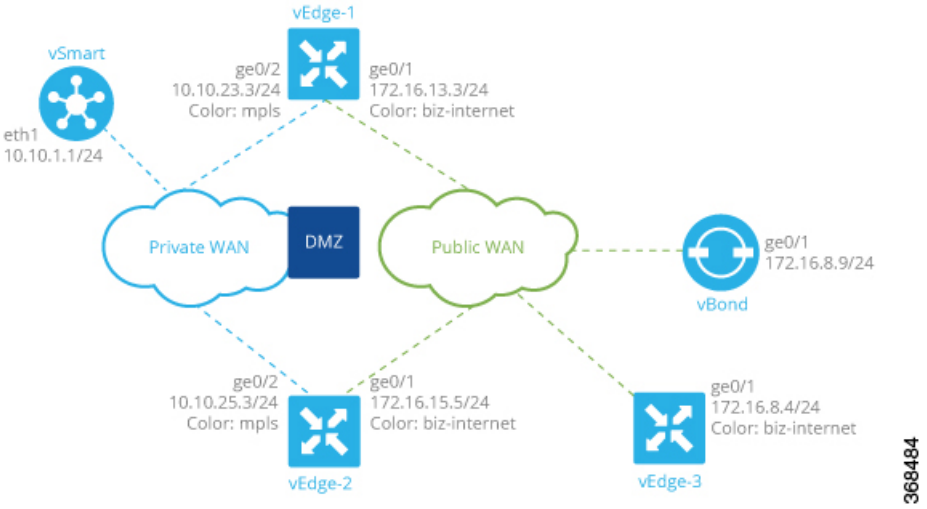

For this configuration to work, the Cisco vBond Orchestrator must be reachable over both WAN transports. Because it is on the public WAN (that is, on the Internet), there needs to be connectivity from the private WAN to the Internet. This could be provided via a DMZ, as shown in the figure above. The 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 theCisco 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     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
                                       SOURCE TLOC      REMOTE TLOC                       DST PUBLIC
   DST PUBLIC         DETECT      TX
SYSTEM IP        SITE ID  STATE       COLOR            COLOR           SOURCE IP         IP
   PORT        ENCAP  MULTIPLIER  INTERVAL(msec)  UPTIME         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 is 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.

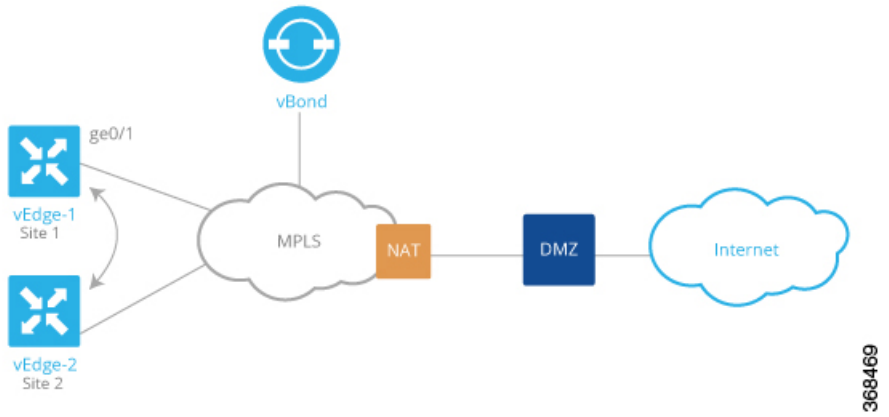

This topology requires a special configuration to allow traffic exchange using private IP addresses because:

- 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
                                 IF      IF
     TCP
                                 ADMIN   OPER    ENCAP                                       SPEED
     MSS                 RX      TX
VPN  INTERFACE  IP ADDRESS       STATUS  STATUS  TYPE   PORT TYPE  MTU   HWADDR             MBPS
DUPLEX  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       0:07:38:49  213199   243908
0    ge0/1      10.1.17.15/24     Up      Up      null   service    1500  00:0c:29:7d:1e:08  10     full
     0       0:07:38:49  197      3
0    ge0/2      -                 Down    Down    null   service    1500  00:0c:29:7d:1e:12  -      -
     0       -           1        1
0    ge0/3      10.0.20.15/24     Up      Up      null   service    1500  00:0c:29:7d:1e:1c  10     full
     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       0:07:38:49  196      3
0    ge0/7      10.0.100.15/24    Up      Up      null   service    1500  00:0c:29:7d:1e:44  10     full
     0       0:07:44:47  783      497
0    loopback1  172.16.255.25/32  Up      Up      null   transport  1500  00:00:00:00:00:00  10     full
     0       0:00:00:20  0        0
0    system     172.16.255.15/32  Up      Up      null   loopback   1500  00:00:00:00:00:00  10     full
     0       0:07:38:25  0        0
1    ge0/4      10.20.24.15/24    Up      Up      null   service    1500  00:0c:29:7d:1e:26  10     full
     0       0:07:38:46  27594    27405
1    ge0/5      56.0.1.15/24      Up      Up      null   service    1500  00:0c:29:7d:1e:30  10     full
     0       0:07:38:46  196      2
512  eth0       10.0.1.15/24      Up      Up      null   service    1500  00:50:56:00:01:05  1000   full
     0       0:07:45:55  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.

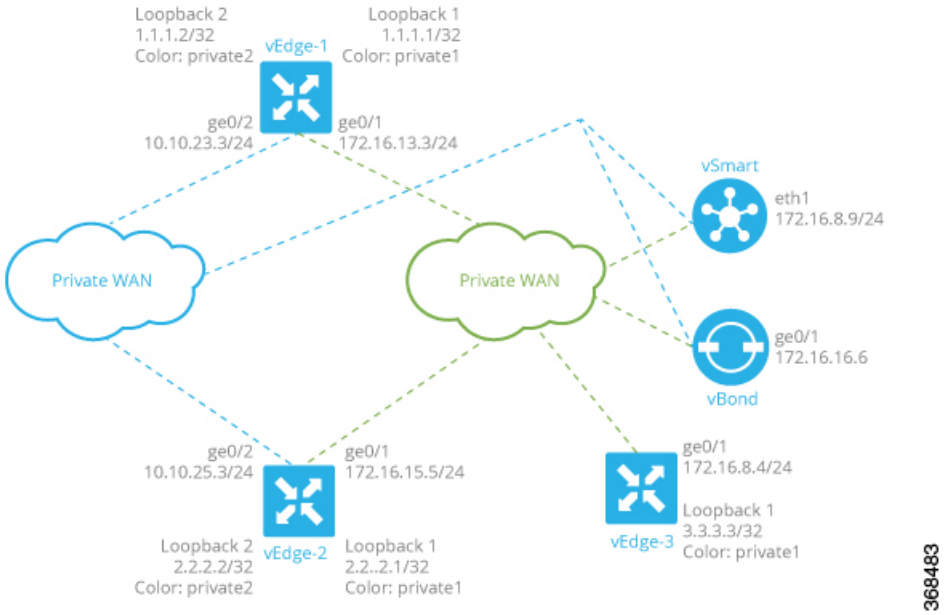

Because both WANs are private, we need only a single overlay. In this overlay network, without policy, IPsec tunnels running BFD sessions exist from any TLOC connected to either transport network to any TLOC in the other transport as well as to any TLOC in the same WAN transport network.

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 advertise 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.

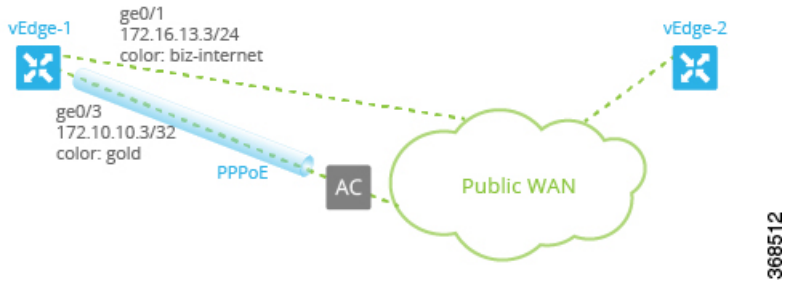

Here is the interface configuration for VPN 0:

```
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

              PPPOE       INTERFACE                PRIMARY   SECONDARY
VPN   IFNAME  INTERFACE   IP          GATEWAY IP   DNS       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 ppppoe session** and **show pppoe statistics** commands to view information about PPPoE sessions:

```
vEdge# show pppoe session

             SESSION                                          PPP                    SERVICE

VPN  IFNAME  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   -


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:

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** .

c. 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:

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**.

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 device s, 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 theCisco 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 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 is 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 preferences 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 5: 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 6:

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

2. To save the NAT mapping, click **Add**.

3. 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 7: IPv4 NAT Parameter Values

Parameter Name	Description
UDP Timeout	Enter the timeout value for User Datagram Protocol (UDP) traffic **1.** Change the scope from Default to **Global**. **2.** Enter a timeout value. Range: 1–536870 seconds Default: 1 second
TCP Timeout	Enter the timeout value for Transmission Control Protocol (TCP) traffic. **1.** Change the scope from Default to **Global**. **2.** 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 usingCisco 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 8:

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, *natpool22*. 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.*Default:* 1 minute Range: 1-65536 minutes
TCP Timeout	Enter the time when NAT translations over TCP sessions time out.*Default:* 60 minutes (1 hour) 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 9:

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	Select the protocol to apply the port-forwarding rule to. To match the same ports for both TCP and UDP traffic, configure two rules.

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 10:

Parameter Scope	Scope Description
Device Specific (indicated by a host icon)	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 . 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 . To change the default key, type a new string and move the cursor out of the Enter Key box. Examples of device-specific parameters are system IP address, hostname, GPS location, and site ID.
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.

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 11:

Parameter Scope	Scope Description
Device Specific (indicated by a host icon)	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 . 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 . To change the default key, type a new string and move the cursor out of the Enter Key box. Examples of device-specific parameters are system IP address, hostname, GPS location, and site ID.
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.

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 12:

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. *Range*: 100 to 255.
PPP Maximum Payload	Enter the maximum receive unit (MRU) value to be negotiated during PPP Link Control Protocol (LCP) negotiation.*Range*: 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 13:

Parameter Name	Description
PPP Authentication Protocol	Select the authentication protocol used by the MLP: • CHAP—Enter the hostname and password provided by your Internet Service Provider (ISP). *hostname* can be up to 255 characters. • PAP—Enter the username and password provided by your ISP. *username* 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 14:

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. *Range:* 0 through 8*Default:* 2
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.*Range:* 0 through 100
Cisco vManage Connection Preference	Set the preference for using a tunnel interface to exchange control traffic with the Cisco vManage NMS.*Range:* 0 through 8*Default:* 5
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.*Default:* Enabled
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 15:

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. *Range:* 0 through 4294967295*Default:* 0
IPsec Weight	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	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	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.*Range:* 1 through 60 seconds*Default:* 5 seconds
Hello Interval	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)
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. *Range:* 12 through 60 seconds*Default:* 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 16:

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).*Default*: Outbound

UDP Timeout	Specify when NAT translations over UDP sessions time out.*Range*: 1 through 65536 minutes*Default*: 1 minutes
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, a router acting as a NAT device receives these error messages.*Default*: 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 17:

Parameter Name	Description
Port Start Range	Enter a port number to define the port or first port in the range of interest.*Range:* 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.*Range:* 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.*Range:* 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 18:

Parameter Name	Description
Shaping rate	Configure the aggreate 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 19:

Parameter Name	Description
Bandwidth Upstream	For transmitted traffic, set the bandwidth above which to generate notifications.*Range:* 1 through (2^{32} / 2) – 1 kbps
Bandwidth Downstream	For received traffic, set the bandwidth above which to generate notifications.*Range:* 1 through (2^{32} / 2) – 1 kbps
IP MTU	Specify the maximum MTU size of packets on the interface.*Range:* 576 through 1804*Default:* 1500 bytes
TCP MSS	Specify the maximum segment size (MSS) of TPC 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.*Range:* 552 to 1460 bytes*Default:* 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.

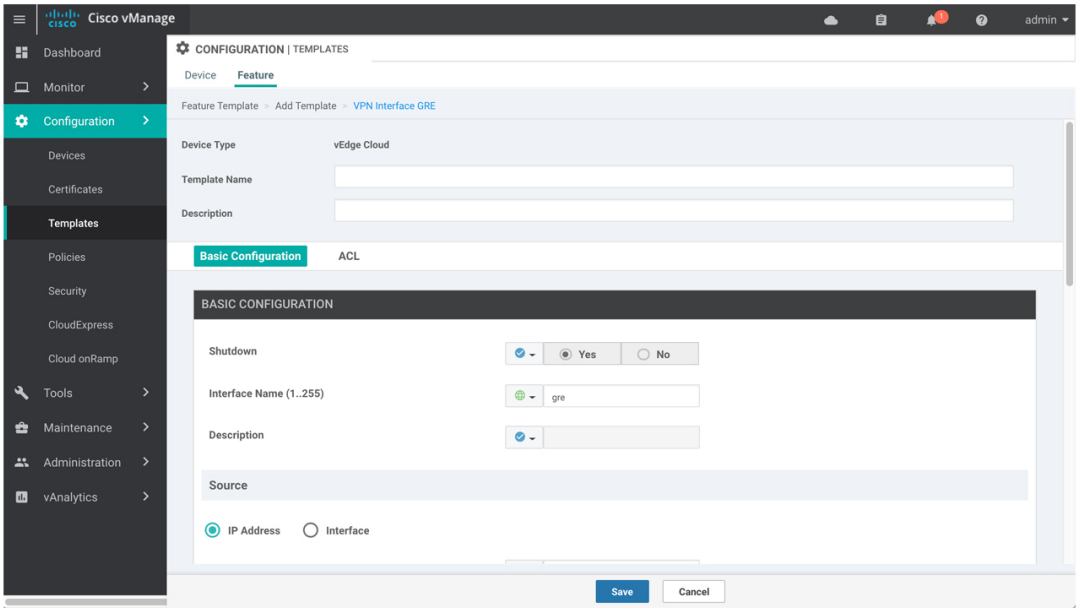

369435

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 20:

Parameter Scope	Scope Description
Device Specific (indicated by a host icon)	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 . 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 . To change the default key, type a new string and move the cursor out of the Enter Key box. Examples of device-specific parameters are system IP address, hostname, GPS location, and site ID.
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.

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 21:

Parameter Name	Description
Shutdown*	Click Off to enable the interface.
Interface Name*	Enter the name of the GRE interface, in the format **gre** *number*. *number* can be from 1 through 255.
Description	Enter a description of the GRE interface.
Source*	Enter the source of the GRE interface: • 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.*Range:* 576 through 1804*Default:* 1500 bytes
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.*Range:* 552 to 1460 bytes*Default:* None
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.*Range:* 0 through 65535 seconds*Default:* 10 seconds
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.*Range:* 0 through 255*Default:* 3

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 22:

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  grenumber 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.

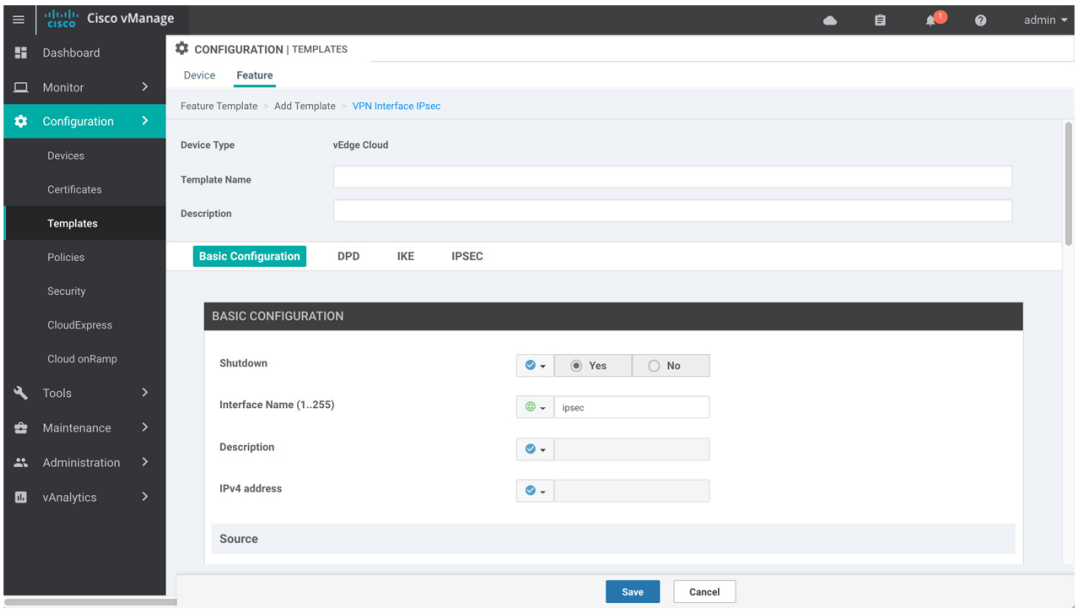

369436

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 23:

Parameter Scope	Scope Description
Device Specific (indicated by a host icon)	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 . 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 . To change the default key, type a new string and move the cursor out of the Enter Key box. Examples of device-specific parameters are system IP address, hostname, GPS location, and site ID.
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.

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 24:

Parameter Name	Description
Shutdown*	Click No to enable the interface.
Interface Name*	Enter the name of the IPsec interface, in the format **ipsec** *number*. *number* 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 *ipv4-prefix*/*length*. The address must be a /30.
Source*	Set the source of the IPsec tunnel that is being used for IKE key exchange: • 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.*Range:* 552 to 1460 bytes*Default:* None
IP MTU	Specify the maximum MTU size of packets on the interface.*Range:* 576 through 1804*Default:* 1500 bytes

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 25:

Parameter Name	Description
DPD Interval	Specify the interval for IKE to send Hello packets on the connection.*Range:* 0 through 65535 seconds (1 hour through 14 days)*Default:* 10 seconds
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.*Range:* 0 through 255*Default:* 3

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 26:

Parameter Name	Description
IKE Version	Enter 1 to select IKEv1.
IKE Mode	Specify the IKE SA establishment mode.*Values:* Aggressive mode, Main mode*Default:* Main mode
IPsec Rekey Interval	Specify the interval for refreshing IKE keys.*Range:* 3600 through 1209600 seconds (1 hour through 14 days)*Default:* 14400 seconds (4 hours)
IKE Cipher Suite	Specify the type of authentication and encryption to use during IKE key exchange.*Values:* aes128-cbc-sha1, aes256-cbc-sha1*Default:* aes256-cbc-sha1
IKE Diffie-Hellman Group	Specify the Diffie-Hellman group to use in IKE key exchange.*Values:* 1024-bit modulus, 2048-bit modulus, 3072-bit modulus, 4096-bit modulus*Default:* 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.*Range: Default:* Tunnel's source IP address
IKE ID for Remote End Point	If the remote IKE peer requires a remote end point identifier, specify it.*Range:* 1 through 64 characters*Default:* 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 27:

Parameter Name	Description
IKE Version	Enter 2 to select IKEv2.

IPsec Rekey Interval	Specify the interval for refreshing IKE keys.*Range:* 3600 through 1209600 seconds (1 hour through 14 days)*Default:* 14400 seconds (4 hours)
IKE Cipher Suite	Specify the type of authentication and encryption to use during IKE key exchange.*Values:* aes128-cbc-sha1, aes256-cbc-sha1*Default:* aes256-cbc-sha1
IKE Diffie-Hellman Group	Specify the Diffie-Hellman group to use in IKE key exchange.*Values:* 1024-bit modulus, 2048-bit modulus, 3072-bit modulus, 4096-bit modulus*Default:* 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.*Range: Default:* Tunnel's source IP address
IKE ID for Remote End Point	If the remote IKE peer requires a remote end point identifier, specify it.*Range:* 1 through 64 characters*Default:* 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 28:

Parameter Name	Description
IPsec Rekey Interval	Specify the interval for refreshing IKE keys.*Range:* 3600 through 1209600 seconds (1 hour through 14 days)*Default:* 14400 seconds (4 hours)
IKE Replay Window	Specify the replay window size for the IPsec tunnel.*Values:* 64, 128, 256, 512, 1024, 2048, 4096, 8192 bytes*Default:* 32 bytes
IPsec Cipher Suite	Specify the authentication and encryption to use on the IPsec tunnel.*Values:* **aes256-cbc-sha1**, **aes256-gcm**, **null-sha1** *Default:* **aes256-gcm**
Perfect Forward Secrecy	Specify the PFS settings to use on the IPsec tunnel.*Values:* • **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.*Default:* **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**.

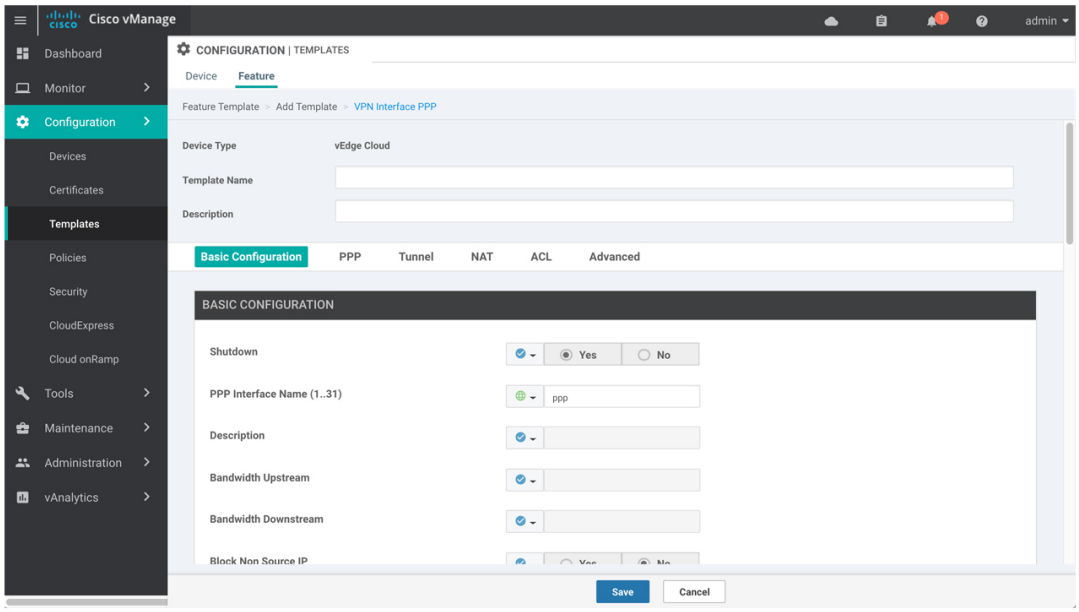

369437

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 29:

Parameter Scope	Scope Description
Device Specific (indicated by a host icon)	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 . 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 . To change the default key, type a new string and move the cursor out of the **Enter Key** box. Examples of device-specific parameters are system IP address, hostname, GPS location, and site ID.
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.

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 30:

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.*Range:* 1 through $(2^{32} / 2) – 1$ kbps
Bandwidth Downstream	For received traffic, set the bandwidth above which to generate notifications.*Range:* 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  ppp number banddwidth-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 31:

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: • CHAP—Enter the hostname and password provided by your Internet Service Provider (ISP). *hostname* can be up to 255 characters. • PAP—Enter the username and password provided by your ISP. *username* 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  ppp number 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 eachCisco 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 32:

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. *Range:* 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 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.*Range:* 0 through 100
vManage Connection Preference	Set the preference for using a tunnel interface to exchange control traffic with the vManage NMS.*Range:* 0 through 8*Default:* 5
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 33:

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. *Range:* 0 through 4294967295*Default:* 0
IPsec Weight	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

Parameter Name	Description
Carrier	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	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.*Range:* 1 through 60 seconds*Default:* 5 seconds
Hello Interval	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)
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. *Range:* 12 through 60 seconds*Default:* 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 34:

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). *Default*: Outbound
UDP Timeout	Specify when NAT translations over UDP sessions time out. *Range*: 1 through 65536 minutes *Default*: 1 minutes

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, a Cisco vEdge devicer acting as a NAT device receives these error messages. *Default*: 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 35:

Parameter Name	Description
Port Start Range	Enter a port number to define the port or first port in the range of interest. *Range:* 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.*Range:* 0 through 65535
Protocol	Select the protocol to whcih 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.*Range:* 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 36:

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 ppp number 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 37:

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.*Range:* 576 through 1804*Default:* 1500 bytes
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.*Range:* 552 to 1460 bytes*Default:* 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 secondCisco 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**.

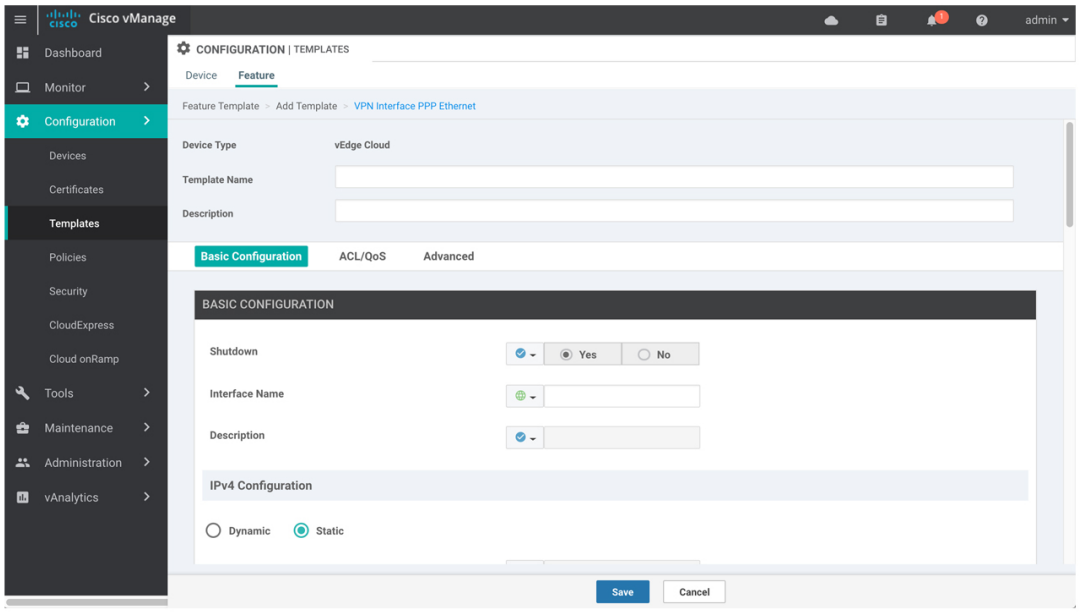

369438

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 38:

Parameter Scope	Scope Description
Device Specific (indicated by a host icon)	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 . 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 . To change the default key, type a new string and move the cursor out of the Enter Key box. Examples of device-specific parameters are system IP address, hostname, GPS location, and site ID.
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.

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 39:

Parameter Name	Description
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. 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.
Description	Enter a description of the PPPoE-enabled 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.

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.*Range:* 1 through (2^{32} / 2) – 1 kbps
Bandwidth Downstream	For received traffic, set the bandwidth above which to generate notifications.*Range:* 1 through (2^{32} / 2) – 1 kbps

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

CLI equivalent:

```
vpn  0
  interface  ppp number 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 ppp number
   [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 40:

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 ppp number 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 41:

Parameter Name	Description
Duplex	Choose full or half to specify whether the interface runs in full-duplex or half-duplex mode.*Default:* 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.*Range:* 576 through 1804*Default:* 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.*Values:* autonet, both, egress, ingress, none*Default:* 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.*Range:* 552 to 1460 bytes*Default:* None

Parameter Name	Description
Speed	Specify the speed of the interface, for use when the remote end of the connection does not support autonegotiation.*Values:* 10, 100, or 1000 Mbps*Default:* Autonegotiate (10/100/1000 Mbps)
Static Ingress QoS	Specify a queue number to use for incoming traffic.*Range:* 0 through 7
ARP Timeout	Specify how long it takes for a dynamically learned ARP entry to time out.*Range:* 0 through 2678400 seconds (744 hours)*Default:* 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 ppp number 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 ppp number 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**.

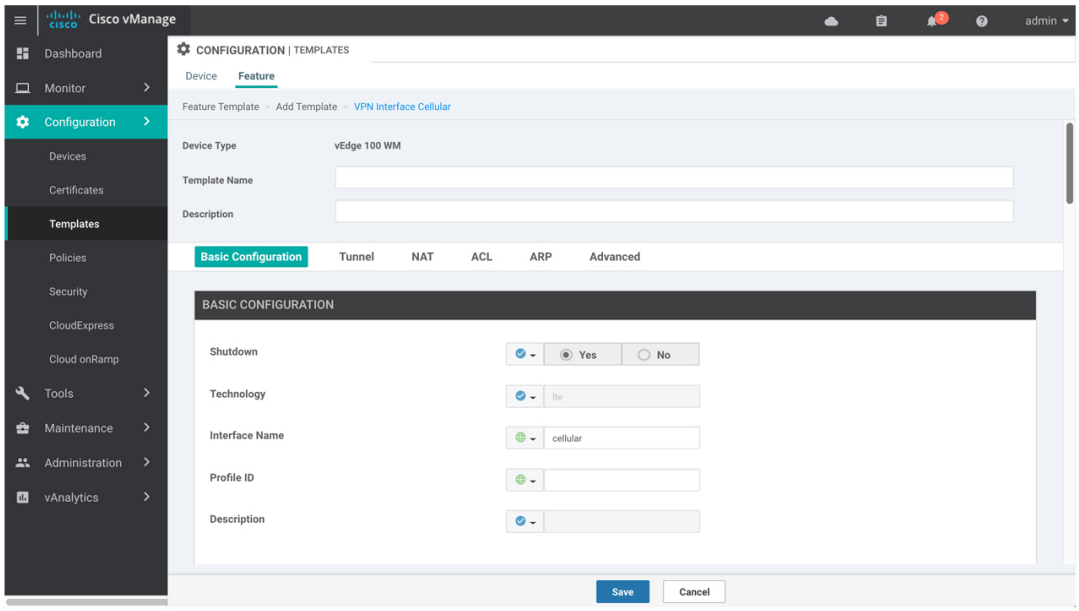

369433

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 42:

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.*Range:* 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.*Range:* 1 through $(2^{32} / 2) – 1$ kbps
Bandwidth Downstream	For received traffic, set the bandwidth above which to generate notifications.*Range:* 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 43:

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 **lte**.
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.*Range:* 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 allows to establish control connections with. Range: 0 through 100
vManage Connection Preference	Set the preference for using the tunnel to exchange control traffic with the vManage NMS. Range: 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 44:

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.*Range:* 0 through 4294967295*Default:* 0
IPsec Weight	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	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	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.*Range:* 1 through 60 seconds*Default:* 5 seconds
Hello Interval	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)
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. *Range:* 12 through 60 seconds*Default:* 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 45: 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).*Default*: Outbound
UDP Timeout	Specify when NAT translations over UDP sessions time out.*Range*: 1 through 65536 minutes*Default*: 1 minute
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, a router acting as a NAT device receives these error messages.*Default*: 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 46:

Parameter Name	Description
Port Start Range	Enter a port number to define the port or first port in the range of interest.*Range:* 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.*Range:* 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.*Range:* 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**.

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 47: Access Lists Parameters

Parameter Name	Description
Shaping rate	Configure the aggreate 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 48:

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 49: 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 TPC 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.*Range:* 552 to 1460 bytes*Default:* None
Clear-Dont-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.*Range:* 0 through 7
ARP Timeout	Specify how long it takes for a dynamically learned ARP entry to time out.*Range:* 0 through 2678400 seconds (744 hours)*Default:* 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(confg-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 interface to 60000 milliseconds (10 minutes) and the hello tolerance to 600 seconds (10 minutes) and include the **no track-transport** 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 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, increate 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 (300000 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-celluar 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-destination (dns-name | ipv4-address) (IPsec 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