



Systems and Interfaces Configuration Guide, Cisco SD-WAN Release 20.x

First Published: 2019-08-15

Last Modified: 2022-08-25

Americas Headquarters

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

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

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

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2019–2022 Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1	Read Me First	1
------------------	----------------------	----------

CHAPTER 2	What's New in Cisco SD-WAN	3
------------------	-----------------------------------	----------

CHAPTER 3	System and Interfaces Overview	5
	Basic Settings for Cisco vManage	9
	Configure Organization Name	10
	Configure Cisco vBond DNS Name or IP Address	10
	Configure Controller Certificate Authorization Settings	10
	Enforce Software Version on Devices	12
	Banner	13
	Create a Custom Banner	14
	Collect Device Statistics	15
	Configure or Cancel vManage Server Maintenance Window	16
	Configure Basic System Parameters	16
	Configure Global Parameters	25
	Create Global Settings Feature Template	26
	CLI Equivalent	28
	Configure NTP Servers Using Cisco vManage	29
	Configure NTP using CLI	32
	Configuring Time Using CLI on Cisco vEdge Device	34
	Configure GPS Using CLI on Cisco vEdge Device	34
	Configure System Logging	35
	Syslog Message Format, Syslog Message Levels, and System Log Files	35
	Configure Logging Using Cisco vManage	38
	Export Cisco vManage NMS Audit Log to Syslog Server	41

Configure System Logging Using CLI	42
View System Logging Information	43
SSH Terminal	44
HTTP/HTTPS Proxy Server for Cisco vManage Communication with External Servers	44
Configure HTTP/HTTPS Proxy Server	45
Bulk API Rate Limit for a Cisco vManage Cluster	46
Configure Bulk API Rate Limit	46
View Bulk API Rate Limit	47

CHAPTER 4

Configure User Access and Authentication	49
Configure Hardened Passwords	50
Enforce Strong Passwords	50
Password Requirements	51
Password Attempts Allowed	52
Password Change Policy	52
Reset a Locked User	53
Reset a Locked User Using the CLI	53
Manage Users	53
Configure Users Using CLI	54
Manage a User Group	55
Creating Groups Using CLI	56
Ciscotac User Access	57
Configure Sessions in Cisco vManage	58
Set a Client Session Timeout in Cisco vManage	58
Set a Session Lifetime in Cisco vManage	58
Set the Server Session Timeout in Cisco vManage	59
Enable Maximum Sessions Per User	59
Configuring RADIUS Authentication Using CLI	60
Configure SSH Authentication	61
SSH Authentication using vManage on Cisco vEdge Devices	61
Configure SSH Authentication using CLI on Cisco vEdge Devices	62
Configure the Authentication Order	62
Configure NAS Attributes using CLI	64
Role-Based Access with AAA	66

Configuring AAA using Cisco vManage Template	75
Navigating to the Template Screen and Naming the Template	75
Configuring Authentication Order and Fallback	76
Configuring Local Access for Users and User Groups	77
Configuring RADIUS Authentication	79
Configuring TACACS+ Authentication	81
Configure Authorization and Accounting	82
Configuring Authorization	82
Configuring Accounting	84
Configuring Password Policy for AAA on Devices	85
Configure Password Policies Using Cisco vManage	86
Configuring IEEE 802.1X and IEEE 802.11i Authentication	87

CHAPTER 5
Role-Based Access Control 99

Information About RBAC	101
Role-Based Access Control by VPN	101
RBAC by VPN	101
VPN Dashboard Overview	101
Role-Based Access with AAA	101
RBAC By Resource Group Overview	110
RBAC for Policies Overview	112
Information About Granular RBAC for Templates	112
Information About Granular Configuration Task Permissions	113
Information About Assigning Roles Locally to a User Defined by an Identity Provider	113
Benefits of RBAC	114
Benefits of Granular RBAC for Feature Templates	114
Restrictions for RBAC	114
Restrictions for Granular RBAC for Feature Templates	114
Use Cases for RBAC	115
Use Cases for Assigning Roles Locally to a User Defined by an Identity Provider	115
Configure RBAC	115
Manage Users	115
User Group Permissions: Cisco IOS XE SD-WAN Devices	116
User Group Permissions: Cisco Catalyst Wireless Gateway Devices	136

RBAC User Group in a Multitenant Environment	138
Add a User	138
Delete a User	139
Edit User Details	139
Change a User Password	139
Check Users Logged In to a Device Using SSH Sessions	140
Check Users Logged In to a Device Using HTTP Sessions	140
Manage a User Group	140
Create User Groups	141
Configure and Manage VPN Segments	142
Configure and Manage VPN Groups	142
Managing Resource Groups	143
Workflow to Configure RBAC for Policies	143
Modify Policy Configurations	144
Assign Users to Configure RBAC for Policies	144
Configure Granular RBAC for Feature Templates	145
Configure RBAC Using the CLI	145
Configure Users Using CLI	145
Creating Groups Using CLI	146
Verify RBAC	147
Verify Granular RBAC Permissions	147
Monitor RBAC	147
Monitor devices for VPN Groups	147

CHAPTER 6

Configure Devices	149
Device Configuration Workflow	149
Feature Templates	150
Device Templates	150
Template Variables	151
Configuration Prerequisites	151
Create a Device Template from Feature Templates	152
Create a Device CLI Template	155
Manage Device Templates	156
Use Variable Values in Configuration Templates	157

Use a File for Variable Parameters	158
Manually Enter Values for Device-Specific Variables and for Optional Rows	160
View Device Templates	162
Attach and Detach a Device Template	163
Determine Why a Device Rejects a Template	165
Change the Device Rollback Timer	165
Preview Device Configuration and View Configuration Differences	166
Change Variable Values for a Device	167
Default Device Templates	167
Configuring Devices using vManage	168
Change Configuration Modes	169
Upload WAN Edge Router Authorized Serial Number File	170
Upload WAN Edge Router Serial Numbers from Cisco Smart Account	171
Generate Bootstrap Configuration for a vEdge Cloud Router	171
Export Device Data in CSV Format	172
View and Copy Device Configuration	172
Delete a WAN Edge Router	173
Decommission a Cloud Router	174
View Template Log and Device Bringup	174
Add a Cisco vBond Orchestrator	174
Configure Cisco vSmart Controllers	175

CHAPTER 7
Network Hierarchy and Resource Management 177

Information About Network Hierarchy and Resource Management	178
Benefits of Network Hierarchy and Resource Management	179
Supported Devices for Network Hierarchy and Resource Management	179
Restrictions for Network Hierarchy and Resource Management	179
Manage a Network Hierarchy	179
Create a Region in a Network Hierarchy	179
Create an Area in a Network Hierarchy	180
Create a Site in a Network Hierarchy	180
Edit a Region	181
Delete a Region	181
Edit an Area	181

Delete an Area	181
Edit a Site	181
Delete a Site	182
Create a System IP Pool	182
Edit a System IP Pool	182
Create a Remote Access Pool	183
Edit a Remote Access Pool	183
Delete a Pool	183
Assign Resource IDs to Devices	184
Assign a Site ID to a Device	184
Use the Quick Connect Workflow	184
Use a Template	184
Use a Configuration Group	185
Assign a Region ID to a Device	185
Assign a System IP to a Device	186
Assign a Hostname to a Device	186

CHAPTER 8
Configure Network Interfaces 187

Configure VPN	188
VPN	188
Create a VPN Template	188
Changing the Scope for a Parameter Value	189
Configure Basic VPN Parameters	190
Configure DNS and Static Hostname Mapping	191
Configure Interfaces in the WAN Transport VPN (VPN 0)	191
Extend the WAN Transport VPN	195
Configure GRE Interfaces and Advertise Services to Them	198
Configure the System Interface	202
Configure Control Plane High Availability	203
Configure Other Interfaces	203
Configure Loopback Interfaces	204
Configure Interface Properties	205
Set the Interface Speed	205
Set the Interface MTU	206

Monitoring Bandwidth on a Transport Circuit	206
Enable DHCP Server using Cisco vManage	207
Configure DHCP Using CLI	210
Configuring PPPoE	212
Configure PPPoE from vManage Templates	212
Configure PPPoE from the CLI	216
Configure PPPoE Over ATM	218
Supported Platforms for PPPoE Over ATM	218
Configure PPPoE Over ATM using Cisco vManage	218
Configure PPPoE Over ATM on the CLI	219
Configuration Example for Configuring PPPoE Over ATM Interfaces	220
Configuring VRRP	220
Network Interface Configuration Examples for Cisco vEdge Devices	225
Configure VPN Ethernet Interface	240
Configure Basic Interface Functionality	241
Create a Tunnel Interface	243
Configure Tunnel Interface CLI on vEdge Devices	244
Associate a Carrier Name with a Tunnel Interface	245
Create Tunnel Groups	245
Configure Tunnel Groups on Cisco vEdge devices Using CLI	245
Limit Keepalive Traffic on a Tunnel Interface	245
Configure Multiple Tunnel Interfaces on a vEdge Router	246
Configure an Interface as a NAT Device	247
Configure IPv4 NAT CLI Equivalent on vEdge	247
Configure NAT64 CLI Equivalent on Cisco vEdge Device	247
VPN Interface NAT Pool using Cisco vManage	247
Apply Access Lists and QoS Parameters	251
Add ARP Table Entries	252
Configuring VRRP	252
Configure a Prefix List for VRRP	253
Configure a Prefix List for VRRP in the Device Template	254
Configure Advanced Properties	255
VPN Interface Bridge	256
Create a Bridging Interface	258

	Apply Access Lists	259
	Configure VRRP	259
	Add ARP Table Entries	261
	Configure Advanced Properties	261
	VPN Interface Ethernet PPPoE	262
	VPN Interface GRE	271
	VPN Interface IPsec (for Cisco vEdge Devices)	274
	VPN Interface PPP	279
	VPN Interface PPP Ethernet	287
	Cellular Interfaces	292
	Configure Cellular Interfaces Using Cisco vManage	292
	Configuring Cellular Interfaces Using the CLI	301
	Best Practices for Configuring Cellular Interfaces	303
	WiFi Radio	304
	WiFi SSID	306
	Interface CLI Reference	308
<hr/>		
CHAPTER 9	IPv6 Functionality	313
<hr/>		
CHAPTER 10	Configure a Cellular Gateway	319
<hr/>		
CHAPTER 11	Track Static Routes for Service VPNs	323
	Information About Static Route Tracking	323
	Restrictions for IPv4 Static Route Tracking	324
	Workflow to Configure IPv4 Static Route Tracking	324
	Create a Static Route Tracker	324
	Configure a Next Hop Static Route with Tracker	326
	Monitor Static Route Tracker Configuration	327
	Configure Static Routes Using CLI	328
	Configuration Examples for Static Route Tracking Using the CLI	330
	Verify Static Route Tracking Configuration Using CLI	331
<hr/>		
CHAPTER 12	VRRP Interface Tracking	333
	Information About VRRP Interface Tracking	333

Restrictions and Limitations	334
VRRP Tracking Use Cases	334
Workflow to Configure VRRP Tracking	335
Configure an Object Tracker	335
Configure VRRP for a VPN Interface Template and Associate Interface Object Tracker	336
Configure VRRP Tracking Using CLI Templates	337
VRRP Object Tracking Using CLI	337
SIG Container Tracking	337
Configuration Example for VRRP Object Tracking Using CLI	338
Configuration Examples for SIG Object Tracking	339
Verify VRRP Tracking	339

CHAPTER 13**Configure a Cisco vEdge Device as an NTP Parent 341**

Configure an NTP Parent	341
Configure Support for NTP in Symmetric Active Mode	342

CHAPTER 14**Dynamic On-Demand Tunnels 345**

On-Demand Tunnel Mechanism in Detail	346
Notes and Limitations	347
Configure On-Demand Tunnels	348
Prerequisites for On-Demand Tunnels	348
Prerequisites: Cisco vSmart Controller Centralized Control Policy	348
Prerequisites: OMP Settings	350
Prerequisites: Hub Device	350
Prerequisites: Spoke Devices	351
Configure On-Demand Tunnels Using Cisco vManage	352
Configure On-Demand Tunnels Using the CLI	353
View Current Status of On-Demand Tunnels in Cisco vManage	353
View Chart of On-Demand Tunnel Status Over Time in Cisco vManage	353

CHAPTER 15**Cisco SD-WAN Multitenancy 355**

Overview of Cisco SD-WAN Multitenancy	355
User Roles in Multitenant Environment	357
Supported Devices and Controller Specifications	359

Restrictions	360
Initial Setup for Multitenancy	361
Create a 3-Node Cisco vManage Cluster	362
Create a 6-Node Cisco vManage Cluster	364
Enable Multitenancy on Cisco vManage	367
Add Cisco vSmart Controller	367
Expand a Multitenant Deployment to Support More Tenants and Tenant Devices	369
Expand a 3-Node Cluster to a 6-node Cluster	369
Manage Tenants	371
Add a New Tenant	372
Modify Tenant Information	374
Delete a Tenant	375
Cisco vManage Dashboard for Multitenancy	376
View Tenant Activity, Device, and Network Information	376
View Detailed Information of a Tenant Setup	376
Manage Tenant WAN Edge Devices	380
Add a WAN Edge Device to a Tenant Network	380
Delete a WAN Edge Device from a Tenant Network	381
Tenant-Specific Policies on Cisco vSmart Controllers	381
Manage Tenant Data	382
Back Up Tenant Data	382
Create, Extract, and List Configuration Data Backup File	383
Restore and Delete Tenant Data Backup File	384
View OMP Statistics per Tenant on a Cisco vSmart Controller	385
View Tenants Associated with a Cisco vSmart Controller	386
Migrate Single-Tenant Cisco SD-WAN Overlay to Multitenant Cisco SD-WAN Deployment	386
Migrate Multitenant Cisco SD-WAN Overlay	390
Upgrade Cisco SD-WAN Controller and Edge Device Software	392
Multitenant Cisco vManage: Disaster Recovery	393
Multitenant Cisco vManage: Disaster Recovery in an Overlay Network with Virtual Routers	398
Multitenant Cisco vManage: Disaster Recovery After a Failed Data Center Becomes Operational	404
Replace Faulty Cisco vSmart Controller	408

Information About Flexible Tenant Placement on Multitenant Cisco vSmart Controllers	411
Benefits of Flexible Tenant Placement on Multitenant Cisco vSmart Controllers	412
Restrictions for Flexible Tenant Placement on Multitenant Cisco vSmart Controllers	413
Assign Cisco vSmart Controllers to Tenants During Onboarding	413
Update Cisco vSmart Controllers Placement For a Tenant	418

CHAPTER 17**Cisco SD-WAN Multitenancy (Cisco SD-WAN Releases 20.4.x and 20.5.x) 421**

Overview of Cisco SD-WAN Multitenancy	421
User Roles in Multitenant Environment	424
Hardware Supported and Specifications	425
Initial Setup for Multitenancy	426
Enable Multitenancy on Cisco vManage	428
Add Cisco vSmart Controller	428
Manage Tenants	430
Add a New Tenant	430
Modify Tenant Information	432
Delete a Tenant	433
Cisco vManage Dashboard for Multitenancy	433
View Tenant Activity, Device, and Network Information	433
View Detailed Information of a Tenant Setup	434
Manage Tenant WAN Edge Devices	437
Add a WAN Edge Device to a Tenant Network	437
Delete a WAN Edge Device from a Tenant Network	438
Tenant-Specific Policies on Cisco vSmart Controllers	438
Manage Tenant Data	439
Back Up Tenant Data	439
Create, Extract, and List Configuration Data Backup File	440
Restore and Delete Tenant Data Backup File	441
View OMP Statistics per Tenant on a Cisco vSmart Controller	442
View Tenants Associated with a Cisco vSmart Controller	443
Migrate Single-Tenant Cisco SD-WAN Overlay to Multitenant Cisco SD-WAN Deployment	443

CHAPTER 18**Appendix: vManage How-Tos 449**

How to Load a Custom vManage Application Server Logo	449
------------------------------------------------------	-----



CHAPTER 1

Read Me First

Related References

- [Cisco SD-WAN Controller Compatibility Matrix and Server Recommendations](#)
- [Cisco SD-WAN Device Compatibility](#)

User Documentation

- [User Documentation for Cisco SD-WAN Release 20](#)

Communications, Services, and Additional Information

- Sign up for Cisco email newsletters and other communications at: [Cisco Profile Manager](#).
- For information on the latest technical, advanced, and remote services to increase the operational reliability of your network visit [Cisco Services](#).
- To browse and discover secure, validated enterprise-class apps, products, solutions, and services, visit [Cisco Devnet](#).
- To obtain general networking, training, and certification titles from Cisco Press Publishers, visit [Cisco Press](#).
- To find warranty information for a specific product or product family, visit [Cisco Warranty Finder](#).
- To view open and resolved bugs for a release, access the [Cisco Bug Search Tool](#).
- To submit a service request, visit [Cisco Support](#).

Documentation Feedback

To provide feedback about Cisco technical documentation use the feedback form available in the right pane of every online document.



CHAPTER 2

What's New in Cisco SD-WAN

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

[What's New in Cisco SD-WAN \(vEdge\) Release 20.x](#)



CHAPTER 3

System and Interfaces Overview

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

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

Host Properties

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

A second host property that must be set on all devices is the IP address of the Cisco vBond Orchestrator for the network domain, or a Domain Name System (DNS) name that resolves to one or more IP addresses for Cisco vBond Orchestrators. A Cisco vBond Orchestrator automatically orchestrates the process of bringing up the overlay network, admitting a new device into the overlay, and providing the introductions that allow the device and Cisco vSmart Controllers to locate each other.

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

To configure the host properties, see [Cisco SD-WAN Overlay Network Bring-Up Process](#).

Time and NTP

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

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

The Cisco SD-WAN software uses Authentication, Authorization, and Accounting (AAA) to provide security for the devices on a network. AAA, in combination with RADIUS and Terminal Access Controller

Access-Control System (TACACS+) user authentication, controls which users are allowed access to devices, and what operations they are authorized to perform after they are logged in or connected to the devices.

Authentication refers to the process by which users trying to access the devices are authenticated. To access devices, users log in with a username and a password. The local device can authenticate users. Alternatively, authentication can be performed by a remote device, either a RADIUS server or a TACACS+ server, or both in a sequence.

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

Beginning in Cisco SD-WAN Release 20.5.1, accounting generates a record of commands that a user executes on a device. Accounting is performed by a TACACS+ server.

For more information, see [Role-Based Access with AAA](#).

Authentication for WANs and WLANs

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

IEEE 802.1X authentication requires three components:

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

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

140-2–compliant AES encryption algorithm along with IEEE 802.1X-based authentication, to enhance user access security over WPA. WPA2 uses the Counter Mode Cipher Block Chaining Message Authentication Code Protocol (CCMP), which is based on the AES cipher. Authentication is done by either using preshared keys or through RADIUS authentication.

Network Segmentation

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

Network Interfaces

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



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

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

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

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

Management and Monitoring Options

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

The following management interfaces are available:

- CLI
- IP Flow Information Export (IPFIX)
- RESTful API
- SNMP

- System logging (syslog) messages
- Cisco vManage

CLI

You can access a CLI on each device, and from the CLI, you configure overlay network features on the local device and gather operational status and information regarding that device. Using an available CLI, we strongly recommend that you configure and monitor all the Cisco SD-WAN network devices from Cisco vManage, which provides views of network-wide operations and device status, including detailed operational and status data. In addition, Cisco vManage provides straightforward tools for bringing up and configuring overlay network devices, including bulk operations for setting up multiple devices simultaneously.

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

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

IPFIX

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

Cisco SD-WAN cflowd performs 1:1 traffic sampling. Information about all the flows is aggregated in the cflowd records; flows are not sampled.



Note Cisco SD-WAN devices do not cache any of the records that are exported to a collector.

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

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

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

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

RESTful API

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

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

SNMP

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

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

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

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

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

Syslog Messages

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

Cisco vManage

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

- [Basic Settings for Cisco vManage, on page 9](#)
- [Configure Basic System Parameters, on page 16](#)
- [Configure Global Parameters, on page 25](#)
- [Configure NTP Servers Using Cisco vManage, on page 29](#)
- [Configure NTP using CLI, on page 32](#)
- [Configuring Time Using CLI on Cisco vEdge Device, on page 34](#)
- [Configure GPS Using CLI on Cisco vEdge Device, on page 34](#)
- [Configure System Logging, on page 35](#)
- [SSH Terminal, on page 44](#)
- [HTTP/HTTPS Proxy Server for Cisco vManage Communication with External Servers, on page 44](#)
- [Bulk API Rate Limit for a Cisco vManage Cluster, on page 46](#)

Basic Settings for Cisco vManage

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

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

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

Configure Organization Name

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

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

To configure the organization name:

1. From the Cisco vManage menu, choose **Administration > Settings**.
2. From **Organization Name**, click **Edit**.
3. In **Organization Name**, enter the name of your organization. The organization name must be identical to the name that is configured on the vBond orchestrator.
4. In **Confirm Organization Name**, re-enter and confirm your organization name.
5. Click **Save**.



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

Configure Cisco vBond DNS Name or IP Address

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



Note The DNS cache timeout should be proportional to the number of Cisco vBond Orchestrator IP addresses that DNS has to resolve, otherwise the control connection for Cisco vManage might not come up during a link failure. This is because, when there are more than six IP addresses (this is the recommended number since the default DNS cache timeout is currently two minutes) to check, the DNS cache timer expires even as the highest preferred interface tries all vBond IP addresses, before failing over to a different color. For instance, it takes about 20 seconds to attempt to connect to one IP address. So, if there are eight IP addresses to be resolved, the DNS cache timeout should be $20 \times 8 = 160$ seconds or three minutes.

Configure Controller Certificate Authorization Settings

Signed certificates are used to authenticate devices in the overlay network. Once authenticated, devices can establish secure sessions between each other. It is from the Cisco vManage that you generate these certificates

and install them on the controller devices—Cisco vBond orchestrators, Cisco vManage, and Cisco vSmart controllers. You can use certificates signed by Symantec, or you can use enterprise root certificates.

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

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

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

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

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

1. From **Controller Certificate Authorization**, click **Edit**.
2. Click **Symantec Manual**.
3. In the **Confirm Certificate Authorization Change** dialog box, click **Proceed** to manually install certificates that the Symantec signing server has generated and signed.
4. Click **Save**.

To use enterprise root certificates:

1. From **Controller Certificate Authorization**, click **Edit**.
2. Click **Enterprise Root Certificate**.
3. In the **Confirm Certificate Authorization Change** dialog box, click **Proceed** to confirm that you wish to use enterprise root certificates.

4. In the **Certificate** box, either paste the certificate, or click **Select a file** and upload a file that contains the enterprise root certificate.
5. By default, the enterprise root certificate has the following properties: To view this information, issue the **show certificate signing-request decoded** command on a controller device, and check the output in the Subject line. For example:
 - Country: United States
 - State: California
 - City: San Jose
 - Organizational unit: ENB
 - Organization: CISCO
 - Domain Name: cisco.com
 - Email: cisco-cloudops-sdwan@cisco.com

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

To change one or more of the default CSR properties:

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

Enforce Software Version on Devices

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

To ensure that templates are in sync after an upgrade that enforces a software version, make sure of the following before you perform the upgrade:

- The bootflash and flash on the router must have enough free space to support the upgrade
- The version of the SD-WAN image that is on the device before the upgrade must be a lower version than the enforced SD-WAN version you specify in the following procedure

To enforce a version of the Cisco SD-WAN software to run on a router when it first joins the overlay network, follow these steps:

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

Banner

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

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

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

Configure a Banner

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Device Templates**, and click **Create Template**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is called **Device**.

3. From the **Create Template** drop-down list, select **From Feature Template**.

4. From the **Device Model** drop-down list, select the type of device for which you are creating the template.
5. Click **Additional Templates** or scroll to the **Additional Templates** section.
6. From the **Banner** drop-down list, click **Create Template**. The **Banner** template form is displayed. This form contains fields for naming the template, and the fields for defining Banner parameters.
7. In **Template Name**, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
8. In **Template Description**, 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 list.

9. To set a banner, configure the following parameters:

Table 1: Parameters to be configured while setting a banner:

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

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

CLI equivalent:

```
banner{login text | motd text}
```

Release Information

Introduced in Cisco vManage NMS in Release 15.2.

Create a Custom Banner

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

1. From **Banner**, click **Edit**.
2. In **Enable Banner**, click **Enabled**.
3. In **Banner Info**, enter the text string for the login banner or click **Select a File** to download a file that contains the text string.
4. Click **Save**.

Collect Device Statistics

Enable or disable the collection of statistics for devices in the overlay network. By default, the collection of statistics is enabled for all the devices in the overlay network.

1. From the Cisco vManage menu, choose **Administration > Settings**.
2. To modify the settings for collecting device statistics, click **Statistics Setting**, and click **Edit**.



Tip To view the configured settings, click **View**.

By default, for every group of statistics (such as **Aggregated DPI** and **AppHosting**), collection of statistics is enabled for all devices.

3. To enable the collection of a group of statistics for all devices, click **Enable All** for the particular group.
4. To disable the collection of a group of statistics for all devices, click **Disable All** for the particular group.
5. To enable the collection of a group of statistics for all devices only for consumption by Cisco vAnalytics, click **vAnalytics only** for the particular group.
6. To enable or disable the collection of a group of statistics for specific devices in the overlay network, click **Custom** for the particular group.

In the **Select Devices** dialog box, depending on whether statistics collection is enabled or disabled for a device, the device is listed among **Enabled Devices** or **Disabled Devices** respectively.

- a. To enable statistics collection for one or more devices, choose the devices from **Disabled Devices** and move them to **Enabled Devices**.



Tip To choose all **Disabled Devices**, click **Select All**.

- b. To disable statistics collection for one or more devices, choose the devices from **Enabled Devices** and move them to **Disabled Devices**.



Tip To choose all **Enabled Devices**, click **Select All**.

- c. To save your selections, click **Done**.
To discard your selections, click **Cancel**.

7. To apply the modified settings, click **Save**.
To discard your changes, click **Cancel**.
To revert to the default settings, click **Restore Factory Default**.

Configure the Time Interval to Collect Device Statistics

1. From the Cisco vManage menu, choose **Administration > Settings**.

- To modify the time interval at which device statistics are collected, find **Statistics Configuration** and click **Edit**.



Tip To view the configured time interval, click **View**.

- Enter the desired **Collection Interval** in minutes.
 - Default value: 30 minutes
 - Minimum value: 5 minutes
 - Maximum value: 180 minutes
- To apply the modified settings, click **Save**.
To discard your changes, click **Cancel**.
To revert to the default settings, click **Restore Factory Default**.

Configure or Cancel vManage Server Maintenance Window

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

- From the Cisco vManage menu, choose **Administration > Settings**.
- From **Maintenance Window**, click **Edit**.
To cancel the maintenance window, click **Cancel**.
- Click the **Start date and time** drop-down list, and select the date and time when the **Maintenance Window** will start.
- Click the **End date and time** drop-down list, and select the date and time when the **Maintenance Window** will end.
- Click **Save**. The start and end times and the duration of the maintenance window are displayed in the **Maintenance Window** bar.

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

Configure Basic System Parameters

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

To configure system-wide parameters using Cisco vManage templates:

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

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

Create System Template

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Device Templates**, and click **Create Template**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is called **Device**.

3. From the **Create Template** drop-down list, select **From Feature Template**.
4. From the **Device Model** drop-down list, select the type of device for which you are creating the template.
5. To create a custom template for System, select the **Factory_Default_System_Template** and click **Create Template**.

The System template form is displayed. This form contains fields for naming the template, and fields for defining the System parameters.

6. In **Template Name**, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
7. In **Template Description**, 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 2:

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

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

Basic System-Wide Configuration

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

Table 3:

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

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

CLI equivalent:

```

system
clock timezone timezone
console-baud-rate rate
controller-group-list numbers
description text
device-groups group-name
host-name string
location string
max-omp-sessions number
site-id site-id
system-ip ip-address
tcp-optimization-enabled

```

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

Configure the GPS Location

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

Table 4:

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

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

CLI equivalent:

```

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

```

Configure Interface Trackers for NAT Direct Internet Access

Table 5: Feature History

Feature Name	Release Information	Description
Support for Interface Status Tracking on Cisco vEdge Devices	Cisco vManage Release 17.2.2	This feature supports interface tracking on Cisco vEdge devices.
Dual Endpoint Support for Interface Status Tracking on Cisco vEdge Devices	Cisco SD-WAN Release 20.6.1 Cisco vManage Release 20.6.1	This feature allows you to configure tracker groups with dual endpoints using the Cisco System template and associate each template group to an interface. The dual endpoints provide redundancy for tracking the status of transport interfaces to avoid false negatives.

The DIA tracker helps determine if the internet or external network becomes unavailable. This feature is useful when NAT is enabled on a transport interface in VPN 0 to allow data traffic from the router to exit directly to the internet.

If the internet or external network becomes unavailable, the router continues to forward traffic based on the NAT route in the service VPN. Traffic that is forwarded to the internet gets dropped. To prevent the internet-bound traffic from being dropped, configure the DIA tracker on the edge router to track the status of the transport interface. The tracker periodically probes the interface IP address of the end point of the tunnel interface to determine the status of the transport interface. The tracker determines the status of the internet and returns the data to the attach points that are associated with the tracker.

When the tracker is configured on the transport interface, the interface IP address is used as a source IP address for probe packets.

IP SLA monitors the status of probes and measures the round trip time of these probe packets and compares the values with the configured latency in the probe. When the latency exceeds the configured threshold value, the tracker considers the network as unavailable.

If the tracker determines that the local internet is unavailable, the router withdraws the NAT route and reroutes the traffic based on the local routing configuration to overlay.

The local router continues to periodically check the status of the path to the interface. When it detects that the path is functioning again, the router reinstalls the NAT route to the internet.

Starting Cisco SD-WAN Release 20.6.1 you can configure a tracker group with dual endpoints on Cisco vEdge devices (using two trackers) and associate this tracker group to an interface. Dual endpoints help in avoiding false negatives that might be introduced regarding unavailability of the internal or external network.

Restrictions for Configuring Tracker Groups for Dual Endpoints

A tracker group with dual endpoints can only be configured on the following types of interfaces:

- Ethernet Interfaces
- Subinterfaces
- PPPoE Interfaces

Configure NAT DIA Tracker

To track the status of transport interfaces that connect to the internet (Network Address Translation Direct Internet Access (NAT DIA)), click **Tracker > Add New Tracker** and configure the following parameters:

Table 6:

Parameter Field	Description
Name	Name of the tracker. The name can be up to 128 alphanumeric characters. You can configure up to eight trackers.

Parameter Field	Description
Tracker Type	<p>Choose an interface, static route, or a tracker group.</p> <p>Starting Cisco SD-WAN Release 20.6.1 you can configure a tracker group with dual endpoints on Cisco vEdge devices and associate this tracker group to an interface.</p> <p>Choose Tracker type as Interface for NAT DIA and dual endpoint tracker configuration.</p>
Tracker Type: Tracker Elements	This field is displayed only if you chose Tracker Type as a tracker-group. Add the existing interface tracker names (separated by a space). When you add this tracker to the template, the tracker group is associated with these individual trackers and you can then associate the tracker group to an interface.
Tracker Type: Tracker Boolean	<p>This field is displayed only if you chose Tracker Type as a tracker-group. Select AND or OR explicitly.</p> <p>An OR operation ensures that the transport interface status is reported as active if either one of the associated trackers of the tracker group report that the interface is active.</p> <p>If you select the AND operation, the transport interface status is reported as active if both the associated trackers of the tracker group report that the interface is active.</p>
Threshold	How long to wait for the probe to return a response before declaring that the transport interface is down. <i>Range:</i> 100 through 1000 milliseconds. <i>Default:</i> 300 milliseconds.
Interval	How often probes are sent to determine the status of the transport interface. <i>Range:</i> 10 through 600 seconds. <i>Default:</i> 60 seconds (1 minute)
Multiplier	Number of times to resend probes before declaring that the transport interface is down. <i>Range:</i> 1 through 10. <i>Default:</i> 3
End Point Type: IP Address	<p>IP address of the end point of the tunnel interface. This is the destination in the internet to which the router sends probes to determine the status of the transport interface.</p> <p>Note In Cisco SD-WAN Release 20.5.1 and later releases, if the tracker receives an HTTP response status code, which is less than 400, the endpoint is reachable.</p> <p>Prior to Cisco SD-WAN Release 20.5.1, the endpoint is reachable if the tracker receives an HTTP response status code of 200.</p>
End Point Type: DNS Name	DNS name of the end point of the tunnel interface. This is the destination in the internet to which the router sends probes to determine the status of the transport interface.

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

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

Configure NAT DIA Tracker Using the CLI

Configure NAT DIA tracker

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

Configure tracker group and assign it to an interface



Note You can configure only one endpoint per tracker.

```
system
  tracker nat-tracker1
    endpoint-ip 10.1.1.1
  !
  tracker nat-tracker2
    endpoint-ip 10.2.2.2
  !
  tracker nat-tracker3
    tracker-type tracker-group
    boolean or
    tracker-elements nat-tracker1 nat-tracker2
  !
  !
  vpn 0
  interface ge0/1
    nat
    tracker nat-tracker3
  !
  !
```

Verify dual endpoints configuration

```
vEdge1# show running-config system | begin tracker
```

```
tracker nat-tracker1
  endpoint-ip 10.1.1.1
  !
  tracker nat-tracker2
    endpoint-ip 10.2.2.2
  !
  tracker nat-tracker3
    boolean or
    tracker-type tracker-group
    tracker-elements nat-tracker1 nat-tracker2
  !
```

```
vEdge1# show tracker tracker-group
```

VPN	INTERFACE	TRACKER NAME	BOOLEAN	STATUS	TRACKER ELEMENT NAME	TRACKER ELEMENT STATUS	TRACKER ELEMENT RTT

```
0    ge0_1    nat-tracker3  or    DOWN    nat-tracker1  DOWN    Timeout
                                nat-tracker2  DOWN    Timeout
```

Apply Tracker to an Interface

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

A tracker group with dual endpoints can only be configured on the following types of interfaces:

- Ethernet Interfaces
- Subinterfaces
- PPPoE Interfaces

Monitor NAT DIA Endpoint Tracker Configuration

1. From the Cisco vManage menu, choose **Monitor > Devices**.

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

2. Choose a device from the list of devices.
3. Click **Real Time**.
4. From the **Device Options** drop-down list, choose **Dual Endpoint Tracker Info**.

Configure Advanced Options

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

Table 7:

Parameter Name	Description
Control Session Policer Rate	Specify a maximum rate of DTLS control session traffic, to police the flow of control traffic. <i>Range:</i> 1 through 65535 pps. <i>Default:</i> 300 pps
MTU of DTLS Tunnel	Specify the MTU size to use on the DTLS tunnels that send control traffic between Cisco SD-WAN devices. <i>Range:</i> 500 through 2000 bytes. <i>Default:</i> 1024 bytes
Port Hopping	Click On to enable port hopping, or click Off to disable it. When a Cisco SD-WAN device is behind a NAT, port hopping rotates through a pool of preselected OMP port numbers (called base ports) to establish DTLS connections with other Cisco SD-WAN devices when a connection attempt is unsuccessful. The default base ports are 12346, 12366, 12386, 12406, and 12426. To modify the base ports, set a port offset value. To disable port hopping on an individual TLOC (tunnel interface), use the VPN Interface Ethernet configuration template. <i>Default:</i> Enabled (on routers); disabled (on Cisco vManage devices and Cisco vSmart Controllers).

Parameter Name	Description
Port Offset	Enter a number by which to offset the base port number. Configure this option when multiple Cisco SD-WAN devices are behind a single NAT device, to ensure that each device uses a unique base port for DTLS connections. <i>Values:</i> 0 through 19
DNS Cache Timeout	Specify when to time out the Cisco vBond Orchestrator addresses that have been cached by the device. <i>Range:</i> 1 through 30 minutes. <i>Default:</i> 30 minutes
Track Transport	Click On to regularly check whether the DTLS connection between the device and a Cisco vBond Orchestrator is up. Click Off to disable checking. By default, transport checking is enabled.
Local vBond (only on routers acting as vBond orchestrators)	Click On to configure the router to act as a Cisco vBond Orchestrator. Then specify the DNS name for the Cisco vBond Orchestrator or its IP address, in decimal four-part dotted notation.
Track Interface	Set the tag string to include in routes associated with a network that is connected to a non-operational interface. <i>Range:</i> 1 through 4294967295
Multicast Buffer	Specify the percentage of interface bandwidth that multicast traffic can use. <i>Range:</i> 5% through 100% <i>Default:</i> 20%
USB Controller (on vEdge 1000 and 2000 series routers only)	Click On to enable or click Off to disable the USB controller, which drives the external USB ports. If you enable the USB controller, the vEdge router reboots when you attach the device template to the device. <i>Default:</i> Disabled
Gateway Tracking	Click On to enable or click Off to Disable tracking of default gateway. Gateway tracking determines, for static routes, whether the next hop is reachable before adding that route to the device's route table. <i>Default:</i> Enabled
Host Policer Rate (on vEdge routers only)	Specify the maximum rate at which a policer delivers packets to the control plane. <i>Range:</i> 1000 through 20000 pps. <i>Default:</i> 5000 pps
ICMP Error Rate (on vEdge routers only)	Specify how many ICMP error messages a policer can generate or receive. <i>Range:</i> 1 through 200 pps <i>Default:</i> 100 pps
Allow Same-Site Tunnel (on vEdge routers only)	Click On to allow tunnels to be formed between vEdge routers in the same site. Note that no BFD sessions are established between the two collocated vEdge routers. <i>Default:</i> Off
Route Consistency Check (on vEdge routers only)	Click On to check whether the IPv4 routes in the device's route and forwarding table are consistent.
Collect Admin Tech on Reboot	Click On to collect admin-tech information when the device reboots.
Idle Timeout	Set how long the CLI is inactive on a device before the user is logged out. If a user is connected to the device via an SSH connection, the SSH connection is closed after this time expires. <i>Range:</i> 0 through 300 seconds. <i>Default:</i> CLI session does not time out.

Parameter Name	Description
Eco-Friendly Mode (on vEdge Cloud routers only)	Click On to configure a Cloud router not to use its CPU minimally or not at all when the router is not processing any packets.

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

CLI equivalent:

```

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

  icmp-error-pps rate

  idle-timeout seconds multicast-buffer-percent percentage

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

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

```

Release Information

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

Configure Global Parameters

Use the Global Settings template to configure a variety of global parameters for all Cisco IOS XE SD-WAN devices, including:

- Various services, such as HTTP and Telnet
- NAT64 timeouts
- HTTP authentication mode
- TCP keepalive
- TCP and UDP small servers
- Console logging
- IP source routing
- VTY line logging

- SNMP IFINDEX persistence
- BOOTP server

Before applying the global parameters to a device, you can view the current configuration of the device and view the differences between the parameter values that you have set in the Global Settings template and the current values on a device.

To configure global settings using Cisco vManage:

1. Create a feature template to configure global settings.
2. Create a device template and include the Global Settings feature template.
3. (Recommended) Before applying the device template to a device, use the [Preview Device Configuration and View Configuration Differences, on page 166](#) feature to review the differences between the configuration currently on the device and the configuration to be sent to the device. This step is recommended because applying the device template overwrites the existing configuration on a device.

Limitations

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

Create Global Settings Feature Template

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Feature Templates**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

3. Click **Add Template**.
4. In the left pane, select a device type.
5. Select the **Global Settings** template.
6. Provide a name and description for the template.
7. For each of the parameters, use the default or set custom values as desired.

Parameter	Description
Services	
HTTP Server	Enable or disable HTTP server.
HTTPS Server	Enable or disable secure HTTPS server.
Passive FTP	Enable or disable passive FTP.
IP Domain-Lookup	Enable or disable domain name server (DNS) lookup.

Parameter	Description
Arp Proxy	Enable or disable proxy ARP.
RSH/RCP	Enable or disable remote shell (RSH) and remote copy (RCP) on the device.
Telnet (Outbound)	Enable or disable outbound telnet.
CDP	
Other Settings	
TCP Keepalives (In)	Enable or disable generation of keepalive timers when incoming network connections are idle.
TCP Keepalives (Out)	Enable or disable generation of keepalive timers when outgoing network connections are idle.
TCP Small Servers	Enable or disable small TCP servers (for example, ECHO).
UDP Small Servers	Enable or disable small UDP servers (for example, ECHO).
Console Logging	Enable or disable console logging. By default, the router sends all log messages to its console port.
IP Source Routing	Enable or disable IP source routing. IP source routing is a feature that enables the originator of a packet to specify the path for the packet to use to get to the destination.
VTY Line Logging	Enable or disable the device to display log messages to a VTY session in real time.
SNMP IFINDEX Persist	Enable or disable SNMP IFINDEX persistence, which provides an interface index (ifIndex) value that is retained and used when the device reboots.
Ignore BOOTP	Enable or disable BOOTP server. When enabled, the device listens for the bootp packet that comes in sourced from 0.0.0.0. When disabled, the device ignores these packets.
NAT64	
UDP Timeout	<p>NAT64 translation timeout for UDP</p> <p>Range: 1 to 65536 (seconds)</p> <p>Default: 300 seconds (5 minutes)</p> <p>Note Starting from Cisco IOS XE Release 17.6.1a and Cisco vManage Release 20.6.1, the default UDP Timeout value for NAT64 has been changed to 300 seconds (5 minutes).</p>

Parameter	Description
TCP Timeout	NAT64 translation timeout for TCP Range: 1 to 65536 (seconds) Default: 3600 seconds (1 hour) Note Starting from Cisco IOS XE Release 17.6.1a and Cisco vManage Release 20.6.1, the default TCP Timeout value for NAT64 has been changed to 3600 seconds (1 hour).
HTTP Authentication	
HTTP Authentication	HTTP authentication mode Accepted values: Local, AAA Default: Local
SSH Version	
SSH version	Specify an SSH version. Default value: Version 2

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

CLI Equivalent

Services (enable):

```
system
 ip http server
 ip http secure-server
 ip ftp passive
 ip domain lookup
 ip arp proxy disable
 ip rcmd rsh-enable
 ip rcmd rcp-enable
 cdp run enable
```

Telnet outbound enable:

```
system
 line vty 0 4
   transport input telnet ssh
```

Services (disable):

```
system
 no ip http server
 no ip http secure-server
 no ip ftp passive
 no ip domain lookup
 no ip arp proxy disable
 no ip rcmd rsh-enable
 no ip rcmd rcp-enable
 no cdp run enable
```

Telnet outbound disable:

```

system
  line vty 0 4
    transport input ssh

```

Other settings (enable):

```

system
  service tcp-keepalives-in
  service tcp-keepalives-out
  service tcp-small-servers
  service udp-small-server
  logging console
  ip source-route
  logging monitor
  snmp-server ifindex persist
  ip bootp server

```

Other settings (disable):

```

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

```

NAT 64:

```

system
  nat64 translation timeout udp timeout
  nat64 translation timeout tcp timeout

```

HTTP Authentication:

```

system
  ip http authentication {local | aaa}

```

Configure NTP Servers Using Cisco vManage

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

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



Note For the NTP to properly function when using VPN0 on the Cisco vEdge devices, you must configure **allow-service ntp** for the tunnel interface on the Cisco VPN Interface Ethernet template.

To configure an NTP server using Cisco vManage templates:

1. Create an NTP feature template to configure NTP parameters, as described in this section.
2. Configure the timezone in the System template.

Name the Template

1. From the Cisco vManage menu, choose **Configuration** > **Templates**.
2. Click **Device Templates**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. From the **Create Template** drop-down list, choose **From Feature Template**.
4. From the **Device Model** drop-down list, choose the type of device for which you wish to create the template.
5. Click **Basic Information**.
6. From **Additional Cisco System Templates**, click **NTP**.
7. From the **NTP** drop-down list, choose **Create Template**.

The **Cisco NTP** template form is displayed. This form contains fields for naming the template, and fields for defining NTP parameters.

8. In **Template Name**, enter a name for the template.
The name can be up to 128 characters and can contain only alphanumeric characters.
9. In **Template Description**, 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 value or to enter a value, click the scope drop-down list to the left of the parameter field and select one of the following:

Table 8: Setting Parameter Scope

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

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

Configure an NTP Server

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

Table 9: Parameters for Configuring an NTP Server

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

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

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

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

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

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

CLI equivalent:

```
system ntp
  server (dns-server-address | ip-address)
    key key-id
  prefer
```

```
source-interface interface-name
version number
vpn vpn-id
```

Configure NTP Authentication Keys

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

Table 10: Parameters for Configuring NTP Authentication Keys

Parameter Name	Description
Authentication Key ID*	Enter the following values: <ul style="list-style-type: none"> • Authentication Key: Enter an MD5 authentication key ID. Valid range is from 1 to 65535. • Authentication Value: Enter either a cleartext key or an AES-encrypted key.
Authentication Value*	Enter an MD5 authentication key. For this key to be used, you must designate it as trusted. To associate a key with a server, enter the same value that you entered in the Authentication Key ID field under Server .

To configure the trusted keys used to authenticate NTP servers, under **Authentication**, click **Trusted Key**, and configure the following parameters.

Table 11: Parameters for Configuring Trusted Keys

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

CLI equivalent:

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

Configure NTP using CLI

Configure Network-Wide Time with NTP

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Configuring Time Using CLI on Cisco vEdge Device

Configure the Timezone

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

To change the timezone on a device:

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

Set the Time Locally

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

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

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

You can also issue these commands as a single command:

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

or

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

To set the timezone, specify it in the configuration:

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

Configure GPS Using CLI on Cisco vEdge Device

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

To set a device's latitude and longitude:

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

You can also set these values using two separate commands:

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

For example:

```
vEdge(config-system)# gps-location latitude 37.0000 longitude 122.0600
or
vEdge(config-system)# gps-location latitude 37.000
vEdge(config-system)# gps-location longitude 122.0600
vEdge(config-system)# show full-configuration
```



```

system
 host-name          vEdge
 gps-location latitude 36.972
 gps-location longitude 122.0263
 ...

```

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

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

For example:

```

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

```

Configure System Logging

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

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

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

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

Syslog Message Format

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

- Syslog message format

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

The field descriptions of syslog messages are:

Table 12: Field Descriptions of Syslog Message Format

Field	Description
facility	Sets the logging facility to a value other than 20, which UNIX systems expect.
severity	The importance or severity of the message is categorized by the numerical code from 0 through 7. A lower number in this range indicates greater severity of the system condition.

Field	Description
description	A text string that describes the condition of syslog server. This portion of the syslog message sometimes includes IP addresses, interface names, port numbers, or usernames.

Usually, the syslog messages are preceded by extra text.

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

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

Syslog Message Levels

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

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

System Log Files

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

- `auth.log`—Login, logout, and superuser access events, and usage of authorization systems
- `kern.log`—Kernel messages
- `messages.log`—Consolidated log file that contains syslog messages from all sources.
- `vconfd.log`—All configuration-related syslog messages
- `vdebug.log`—All debug messages for modules whose debugging is turned on and all syslog messages that are above the default priority value. The debug log messages support various levels of logging based on the module. The different modules implement the logging levels differently. For example, the system manager (`sysmgr`) has two logging levels (on and off), while the chassis manager (`chmgr`) has four different logging levels (off, low, normal, and high). You cannot send debug messages to a remote host. Therefore, to enable debugging, use the **debug** operational command.

- vsyslog.log—All syslog messages from Cisco SD-WAN processes (daemons) that are above the configured priority value. The default priority value is "informational", so by default, all "notice", "warning", "error", "critical", "alert", and "emergency" syslog messages are saved.
- vmanage-syslog.log—Cisco vManage NMS Audit log messages

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

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

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

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

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

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

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

```
local6.info: 18-Oct-2020 17:42:07 vm10 maintenance-fragment-1/2: {"logid":
"d9ed576a-43ae-49ce-921b-a51c1ed40698", "entry_time":
1576605512190, "statcycletime" 34542398334245, "logmodule":"maintenance", "logfeature":
"upgrade", "loguser": "admin", "logusersrcip":
"10.0.1.1", "logmessage": "Device validation Upgrade to version - Validation success",
"logdeviceid":"Validation", "auditdetails" :
["[18-Oct-2020 17:42:08 UTC] Published messages to vmanage(s)", "auditdetails":["[18-Oct-2020
17:42:07 UTC] Software image: vmanage-99.99.999-
x86_64.tar.gz", "Software image download may take up to 60]
local6.info: 18-Oct-2020 17:42:07 vm10 maintenance-fragment-2/2: { minutes", "logprocessid":
"software_install-7de0ec44-d290-4429-b24532435324", "tenant":, "default"}
```

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

disable the logging of AAA and Netconf syslog messages by using the following commands from Cisco vManage NMS:

Disable logging of AAA and Netconf Syslog Messages


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

Configure Logging Using Cisco vManage

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

Navigate to the Template Screen and Name the Template

1. In vManage NMS, select the **Configuration ► Templates** screen.
2. In the **Device** tab, click **Create Template**.
3. From the **Create Template** drop-down, select **From Feature Template**.
4. From the **Device Model** drop-down, select the type of device for which you are creating the template.
5. To create a custom template for Logging, select the **Factory_Default_Logging_Template** and click **Create Template**. The Logging template form is displayed. The top of the form contains fields for naming the template, and the bottom contains fields for defining Logging parameters. You may need to click a tab or the plus sign (+) to display additional fields.



369421

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

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

Minimum Logging Configuration

The following logging parameters are configured by default:

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

Configure Logging to the Local Disk

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

Table 13:

Parameter Name	Description
Enable Disk	Click On to allow syslog messages to be saved in a file on the local hard disk, or click Off to disallow it. By default, logging to a local disk file is enabled on all Viptela devices.

Parameter Name	Description
Maximum File Size	Enter the maximum size of syslog files. Syslog files are rotated on an hourly basis based on the file's size. When the file size exceeds configured value, the file is rotated and the syslogd process is notified. <i>Range:</i> 1 through 20 MB <i>Default:</i> 10 MB
Rotations	Enter the number of syslog files to create before discarding the oldest files. <i>Range:</i> 1 through 10 <i>Default:</i> 10
Priority	Select the priority level of the syslog message to save to the log files. The severity indicates the seriousness of the event that generated the message. The default priority value is "informational", so, by default, all syslog messages are recorded. The priority level can be one of the following (in order of decreasing severity): • Emergency—System is unusable (corresponds to syslog severity 0). • Alert— Action must be taken immediately (corresponds to syslog severity 1). • Critical—Critical: A serious condition (corresponds to syslog severity 2). • Error—An error condition that does not fully impair system usability (corresponds to syslog severity 3). • Warning—A minor error condition (corresponds to syslog severity 4). • Notice—A normal, but significant condition (corresponds to syslog severity 5). • Informational—Routine condition (the default) (corresponds to syslog severity 6).

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

CLI equivalent:

```

system
 logging
  disk
    enable
    file
      rotate numbersize megabytes priority priority

```

Configure Logging to Remote Servers

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

Table 14:

Parameter Name	Description
Hostname/IP Address	Enter the DNS name, hostname, or IP address of the system on which to store syslog messages. To add another syslog server, click the plus sign (+). To delete a syslog server, click the trash icon to the right of the entry.
VPN ID	Enter the identifier of the VPN in which the syslog server is located or through which the syslog server can be reached. <i>Range:</i> 0 through 65530
Source Interface	Enter the specific interface to use for outgoing system log messages. The interface must be located in the same VPN as the syslog server. Otherwise, the configuration is ignored. If you configure multiple syslog servers, the source interface must be the same for all of them.

Parameter Name	Description
Priority	<p>Select the severity of the syslog message to save. The severity indicates the seriousness of the event that generated the message. <i>priority</i> can be one of the following:</p> <ul style="list-style-type: none"> • Emergency—System is unusable (corresponds to syslog severity 0). • Alert— Action must be taken immediately (corresponds to syslog severity 1). • Critical—Critical: A serious condition (corresponds to syslog severity 2). • Error—An error condition that does not fully impair system usability (corresponds to syslog severity 3). • Warning—A minor error condition (corresponds to syslog severity 4). • Notice—A normal, but significant condition (corresponds to syslog severity 5). • Informational—Routine condition (the default) (corresponds to syslog severity 6). <p>Click Add to save the logging server.</p>

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

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

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

CLI equivalent:

```

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

```

Release Information

Introduced in Cisco vManage NMS in Release 15.2.

Export Cisco vManage NMS Audit Log to Syslog Server

Table 15: Feature History

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

On Cisco IOS XE SD-WAN devices and Cisco vEdge devices, you can log event notification system log (syslog) messages to files on a local device, or to files on a remote host using CLI. These event notification logs are converted to system log files and exported to the syslog server. You can then retrieve system log information from the syslog server.

Configure System Logging Using CLI

Log Syslog Messages to a Local Device

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

1. logging disk

Logs syslog messages on a hard disk

Example:

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

2. enable

Enables logging to a disk

Example:

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

3. file size *size*

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

Example:

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

4. file rotate *number*

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

Example:

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

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

Log Syslog Messages to a Remote Device

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

1. logging server

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

Example:

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

2. (Optional) vpn *vpn-id*

Specifies the VPN ID of the syslog server

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

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

Example:

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

4. priority *priority*

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

Example:

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

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

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

View System Logging Information

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

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

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

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

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

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

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

1. From the Cisco vManage menu, choose **Administration > Settings**, and ensure that you enable **Data Stream**.
2. From the Cisco vManage menu, choose **Monitor > Devices**, and choose a Cisco vEdge device
Cisco vManage Release 20.6.x and earlier: From the Cisco vManage menu, choose **Monitor > Network**, and choose a Cisco vEdge device.
3. Click **Troubleshooting**.
4. From **Logs**, click **Debug Log**.
5. From **Log Files**, select a name of the log file to view the log information.

SSH Terminal

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

Establish an SSH Session to a Device

To establish an SSH session to a device:

1. From the Cisco vManage menu, choose **Tools > SSH Terminal**.
2. Select the device on which you wish to collect statistics:
 - a. Select the device group to which the device belongs.
 - b. If needed, sort the device list by its status, hostname, system IP, site ID, or device type.
 - c. Click the device to select it.
3. Enter the username and password to log in to the device.

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

HTTP/HTTPS Proxy Server for Cisco vManage Communication with External Servers

Table 16: Feature History

Feature Name	Release Information	Description
HTTP/HTTPS Proxy Server for Cisco vManage Communication with External Servers	Cisco SD-WAN Release 20.5.1 Cisco vManage Release 20.5.1	Cisco vManage uses HTTP/HTTPS to access some web services and for some REST API calls. With this feature, you can channel the HTTP/HTTPS communication through an HTTP/HTTPS proxy server.

The following are some instances in which Cisco vManage uses an HTTP/HTTPS connection to an external server:

- Certificate request or renewal
- Cisco Plug and Play integration
- Smart Licensing Using Policy
- Cloud OnRamp
- Software image download
- Data upload to Cisco SD-WAN vAnalytics

In Cisco vManage Release 20.4.1 and earlier releases, you must permit this HTTP/HTTPS communication in the firewall configured on your on-premises Cisco vManage instance. Beginning Cisco vManage 20.5.1, you can channel the HTTP/HTTPS communication via an HTTP/HTTPS proxy server. With the HTTP/HTTPS proxy server configured, you can restrict HTTP/HTTPS communication with external servers while configuring the firewall and secure the system further.

Traffic is directed through the HTTP/HTTPS proxy server in the following cases:

- HTTPS connection for Symantec or Cisco automated certificate request or renewal
- REST API calls to URLs of the following domains:
 - cisco.com
 - amazonaws.com
 - microsoft.com
 - office.com
 - microsoftonline.com

Once every 24 hours, Cisco vManage checks whether the configured HTTP/HTTPS proxy server is reachable. If the proxy server is unreachable, Cisco vManage raises the alarm `HTTPS proxy server {IP} not reachable`.

Restrictions

- When configured to communicate with external servers via an HTTP/HTTPS proxy server, Cisco vManage resolves FQDNs locally or through configured DNS servers, bypassing the proxy server. Cisco vManage then sends the HTTP/HTTPS connections resulting from the resolution to the proxy server. DNS queries for the resolution of external server FQDNs must be successful before Cisco vManage can send resulting HTTP/HTTPS connections to the HTTP/HTTPS proxy server.
- Use of the HTTP/HTTPS proxy server is not supported for communication between the SD-AVC container in Cisco vManage and external services.

Configure HTTP/HTTPS Proxy Server

1. From the Cisco vManage menu, choose **Administration > Settings**.
2. For the **HTTP/HTTPS Proxy** setting, click **Edit**.
3. For the **Enable HTTP/HTTPS Proxy** setting, click **Enabled**.
4. Enter the **HTTP/HTTPS Proxy IP Address** and **Port** number.
5. Click **Save**.



Note Cisco vManage uses TCP port 7 echo request to validate reachability of the proxy server. Ensure that you configure your firewall and proxy server to allow the echo requests to make the destination host ports accessible.

Cisco vManage verifies that the HTTP/HTTPS proxy server is reachable and saves the server details in the configuration database. HTTP/HTTPS connections and REST API calls to external servers are directed through the proxy server.

If the HTTP/HTTPS proxy server is not reachable, Cisco vManage displays an error message on the GUI indicating the reason for failure.

Bulk API Rate Limit for a Cisco vManage Cluster

Table 17: Feature History

Feature Name	Release Information	Description
Bulk API Rate Limit for a Cisco vManage Cluster	Cisco vManage Release 20.10.1	For a Cisco vManage cluster, the rate limit for bulk APIs equals (rate-limit per node) * (number of nodes in the cluster). Cisco vManage distributes bulk API requests among the nodes in the cluster. With these changes, you can retrieve data faster from a Cisco vManage cluster through bulk APIs.

In Cisco vManage Release 20.9.x and earlier releases, you send bulk API requests to a node in the Cisco vManage cluster. The bulk API throughput is constrained by the rate-limit per node. To increase the throughput, you must send separate bulk API requests to each node in the cluster and collate the API responses.

From Cisco vManage Release 20.10.1, send bulk API requests to the Cisco vManage cluster. Cisco vManage distributes the API requests among the clusters in the node. This distribution increases the rate limit to (rate-limit per node) * (number of nodes in the cluster), allowing you to retrieve more data in a shorter duration compared to a bulk API request addressed to a single node. With the distribution, you need not send separate bulk API requests to two or more nodes in the cluster or collate the API responses.

Configure Bulk API Rate Limit

1. Log in to one of the Cisco vManage nodes in the Cisco vManage cluster and configure the following command:

```
vManage# request nms server-proxy set ratelimit
```

2. The command-line displays the following prompt about the rate limit for non-bulk APIs:

```
Do you want to reconfigure rate limit for URL non bulk api [y/n] :
```

Enter **n**.

3. The command-line displays the following prompt about the rate limit for bulk APIs:

```
Do you want to reconfigure rate limit for URL bulk api /dataservice/data/device/statistics [y/n] :
```

Enter **y**.

4. Enter the per-node rate limit in response to a prompt similar to the following:

```
Enter the PER NODE rate limit for URL bulk api /dataservice/data/device/statistics [144
load balanced across all nodes at present] :
```

This prompt is from a three-node Cisco vManage cluster, with the bulk API rate limit configured to the default value of 48 requests per node. Across all the three nodes, the bulk API rate limit is (rate-limit/node) * 3, which is 144 requests.

Before you enter the rate limit, consider its effect on Cisco vManage resources.

5. Enter the unit time for which the rate limit applies in response to a prompt similar to the following.

You can apply a rate limit per second, minute, hour, or day. The default unit is minute.

```
Enter the rate limit unit (second, minute, hour, day) for URL bulk api
/dataservice/data/device/statistics [minute] :
```

Cisco vManage applies the rate limit on all the Cisco vManage instances in the cluster. The command line displays the following message:

```
Propagating rate limit update across all nodes. Please wait.
```

After the rate limit is applied, Cisco vManage prompts you to restart the server-proxy on all nodes and the command line returns to the privileged EXEC mode:

```
Done. Please restart server-proxy on all nodes using "request nms server-proxy restart"
command.
vManage#
```

6. Restart the server-proxy using the following command:

```
vManage# request nms server-proxy restart
```

7. Log in to the other Cisco vManage nodes in the cluster and restart the server-proxy using the **request nms server-proxy restart** command.

In the following example, the bulk API rate limit per node is set to 50 requests per minute.

```
vManage# request nms server-proxy set ratelimit
Do you want to reconfigure rate limit for URL non bulk api [y/n] : n
Do you want to reconfigure rate limit for URL bulk api /dataservice/data/device/statistics
[y/n] : y
Enter the PER NODE rate limit for URL bulk api /dataservice/data/device/statistics [144
load balanced across all nodes at present] : 50
Enter the rate limit unit (second, minute, hour, day) for URL bulk api
/dataservice/data/device/statistics [minute] : minute
Propagating rate limit update across all nodes. Please wait.
Done. Please restart server-proxy on all nodes using "request nms server-proxy restart"
command.
vManage# request nms server-proxy restart
```

View Bulk API Rate Limit

To view the bulk API rate limit, log in to any node in the Cisco vManage cluster and use the **show nms server-proxy ratelimit** command.

The following is a sample command output:

```
vManage# show nms server-proxy ratelimit
Non Bulk API: 100/second (per node)
Bulk API: 150/minute (across cluster)
```

This sample output is from three-node Cisco vManage cluster with the bulk API rate limit per node configured to 50 requests per minute. Therefore, the bulk API rate limit for the cluster is $50 * 3 = 150$ requests per minute.



CHAPTER 4

Configure User Access and Authentication

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

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

- [Configure Hardened Passwords](#) , on page 50
- [Manage Users](#), on page 53
- [Configure Users Using CLI](#), on page 54
- [Manage a User Group](#), on page 55
- [Creating Groups Using CLI](#), on page 56
- [CiscoTAC User Access](#), on page 57
- [Configure Sessions in Cisco vManage](#), on page 58
- [Configuring RADIUS Authentication Using CLI](#), on page 60
- [Configure SSH Authentication](#), on page 61
- [Configure the Authentication Order](#), on page 62
- [Configure NAS Attributes using CLI](#), on page 64
- [Role-Based Access with AAA](#), on page 66
- [Configuring AAA using Cisco vManage Template](#), on page 75
- [Configuring Password Policy for AAA on Devices](#), on page 85
- [Configuring IEEE 802.1X and IEEE 802.11i Authentication](#), on page 87

Configure Hardened Passwords

Table 18: Feature History

Feature Name	Release Information	Description
Hardened Passwords	Cisco vManage Release 20.3.1	This feature enables password policy rules in Cisco vManage. After password policy rules are enabled, Cisco vManage enforces the use of strong passwords.
	Cisco SD-WAN Release 20.9.1 Cisco vManage Release 20.9.1	This feature lets you configure Cisco vManage to enforce predefined-medium security or high-security password criteria.

Enforce Strong Passwords

We recommend the use of strong passwords. You must enable password policy rules in Cisco vManage to enforce use of strong passwords.

After you enable a password policy rule, the passwords that are created for new users must meet the requirements that the rule defines. In addition, for releases from Cisco vManage Release 20.9.1, you are prompted to change your password the next time you log in if your existing password does not meet the requirements that the rule defines.

1. From the Cisco vManage menu, choose **Administration > Settings**.
2. In **Password Policy**, choose **Edit**.
3. Perform one of these actions, based on your Cisco vManage release:
 - For releases before Cisco vManage Release 20.9.1, click **Enabled**.
 - For releases from Cisco vManage Release 20.9.1 click **Medium Security** or **High Security** to choose the password criteria.

By default, **Password Policy** is set to **Disabled**.

4. In the **Password Expiration Time (Days)** field, you can specify the number of days for when the password expires.

By default, password expiration is 90 days.

Before your password expires, a banner prompts you to change your password. If the password expiration time is 60 days or more, this banner first appears at 30 days before your password expires. If the password expiration time is less than 60 days, this banner first appears at half the number of days that are configured for the expiration time. If you do not change your password before it expires, you are blocked from logging in. In such a scenario, an admin user can change your password and restore your access.



Note The password expiration policy does not apply to the admin user.

5. Click **Save**.

Password Requirements

Cisco vManage enforces the following password requirements after you have enabled the password policy rules:

- The following password requirements apply to releases before Cisco vManage Release 20.9.1:
 - Must contain a minimum of eight characters, and a maximum of 32 characters.
 - Must contain at least one uppercase character.
 - Must contain at least one lowercase character.
 - Must contain at least one numeric character.
 - Must contain at least one of the following special characters: # ? ! @ \$ % ^ & * - .
 - Must not contain the full name or username of the user.
 - Must not reuse a previously used password.
 - Must contain different characters in at least four positions in the password.
- Minimum releases: Cisco SD-WAN Release 20.9.1, Cisco vManage Release 20.9.1:

Password Criteria	Requirements
Medium Security	<ul style="list-style-type: none"> • Must contain a minimum of 8 characters • Must contain no more than 32 characters • Must contain at least 1 lowercase character • Must contain at least 1 uppercase character • Must contain at least 1 numeric character • Must contain at least 1 of the following special characters: # ? ! @ \$ % ^ & * - . • Must not be identical to any of the last 5 passwords used • Must not contain the full name or username of the user

Password Criteria	Requirements
High Security	<ul style="list-style-type: none"> • Must contain a minimum of 15 characters • Must contain no more than 32 characters • Must contain at least 1 lowercase character • Must contain at least 1 uppercase character • Must contain at least 1 numeric character • Must contain at least 1 of the following special characters: # ? ! @ \$ % ^ & * - • Must not be identical to any of the last 5 passwords used • Must not contain the full name or username of the user • Must have at least eight characters that are not in the same position they were in the old password

Password Attempts Allowed

You are allowed five consecutive password attempts before your account is locked. After six failed password attempts, you are locked out for 15 minutes. If you enter an incorrect password on the seventh attempt, you are not allowed to log in, and the 15-minute lock timer starts again.

If your account is locked, wait for 15 minutes for the account to automatically be unlocked. Alternatively, reach out to an administrator to reset the password, or have an administrator unlock your account.



Note Your account gets locked even if no password is entered multiple times. When you do not enter anything in the password field, it is considered as invalid or wrong password.

Password Change Policy



Note You must have enabled password policy rules first for strong passwords to take effect. For more information, see [Enforce Strong Passwords, on page 50](#).

When resetting your password, you must set a new password. You cannot reset a password using an old password.



Note In Cisco vManage Release 20.6.4, Cisco vManage Release 20.9.1 and later releases, a user that is logged out, or a user whose password has been changed locally or on the remote TACACS server cannot log in using their old password. The user can log in only using their new password.

Reset a Locked User

If a user is locked out after multiple password attempts, an administrator with the required rights can update passwords for this user.

There are two ways to unlock a user account, by changing the password or by getting the user account unlocked.



Note Only a **netadmin** user or a user with the User Management Write role can perform this operation.

To reset the password of a user who has been locked out:

1. In **Users (Administration > Manage Users)**, choose the user in the list whose account you want to unlock.
2. Click **...** and choose **Reset Locked User**.
3. Click **OK** to confirm that you want to reset the password of the locked user. Note that this operation cannot be undone.

Alternatively, you can click **Cancel** to cancel the operation.

Reset a Locked User Using the CLI

You can reset a locked user using the CLI as follows:

1. Log in to the device as an `admin` user.
2. Run the following command:

```
Device# request aaa unlock-user username
```
3. When prompted, enter a new password for the user.

Manage Users

From the Cisco vManage menu, choose **Administration > Manage Users** to add, edit, view, or delete users and user groups.

Please note the following:

- Only a user logged in as the **admin** user or a user who has Manage Users write permission can add, edit, or delete users and user groups from Cisco vManage.
- Each user group can have read or write permission for the features listed in this section. Write permission includes Read permission.
- All user groups, regardless of the read or write permissions selected, can view the information displayed in the Cisco vManage Dashboard.

Table 19: User Group Permissions for Different Device Types

Permissions	See This Section
User group permissions related to Cisco IOS XE SD-WAN device configuration.	User Group Permissions: Cisco IOS XE SD-WAN Devices
User group permissions related to Cisco Catalyst Wireless Gateway device configuration.	User Group Permissions: Cisco Catalyst Wireless Gateway Devices

Configure Users Using CLI

You can use the CLI to configure user credentials on each device. This way, you can create additional users and give them access to specific devices. The credentials that you create for a user by using the CLI can be different from the Cisco vManage credentials for the user. In addition, you can create different credentials for a user on each device. All users with the **netadmin** privilege can create a new user.

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

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

The Username can be 1 to 128 characters long, and it must start with a letter. The name can contain only lowercase letters, the digits 0 through 9, hyphens (-), underscores (_), and periods (.). The name cannot contain any uppercase letters. Because some usernames are reserved, you cannot configure them. For a list of reserved usernames, see the **aaa** configuration command in the Cisco SD-WAN Command Reference Guide.

The Password is the password for a user. Each username must have a password, and users are allowed to change their own password. The CLI immediately encrypts the string and does not display a readable version of the password. When a user logs in to a Cisco vEdge device, they have five chances to enter the correct password. After the fifth incorrect attempt, the user is locked out of the device, and must wait for 15 minutes before attempting to log in again.



Note Enclose any user passwords that contain the special character ! in double quotation marks (“”). If a double quotation is not included for the entire password, the config database (?) treats the special character as a space and ignores the rest of the password.

For example, if the password is C!sc0, use “C!sc0”.

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

The factory-default password for the admin username is admin. We strongly recommend that you modify this password the first time you configure a Cisco vEdge device :

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

Configure the password as an ASCII string. The CLI immediately encrypts the string and does not display a readable version of the password, for example:

```
vEdge(config-user-admin)# show config
system
aaa
  user admin
    password $1$xULc8yYH$k71cTjvKESmeIGgImNDaC.
  !
  user eve
    password $1$8z3q4qoU$F6DMBr9vPBF0s/sl45ax5.
    group basic
  !
!
```

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

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

The tag is a string that you defined with the **radius server tag** command, as described in the Cisco SD-WAN Command Reference Guide.

Manage a User Group

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

- **basic**: Includes users who have permission to view interface and system information.
- **netadmin**: Includes the admin user, by default, who can perform all operations on the Cisco vManage. You can add other users to this group.
- **operator**: Includes users who have permission only to view information.
- Minimum supported release: Cisco vManage Release 20.9.1
- **network_operations**: Includes users who can perform non-security operations on Cisco vManage, such as viewing and modifying non-security policies, attaching and detaching device templates, and monitoring non-security data.
- Minimum supported release: Cisco vManage Release 20.9.1
- **security_operations**: Includes users who can perform security operations on Cisco vManage, such as viewing and modifying security policies, and monitoring security data.

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

Delete a User Group

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

1. From the Cisco vManage menu, choose **Administration > Manage Users**.
2. Click **User Groups**.
3. Click the name of the user group you wish to delete.



Note You cannot delete any of the default user groups—`basic`, `netadmin`, `operator`, `network_operations`, and `security_operations`.

4. Click **Trash** icon.
5. To confirm the deletion of the user group, click **OK**.

Edit User Group Privileges

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

1. From the Cisco vManage menu, choose **Administration > Manage Users**.
2. Click **User Groups**.
3. Select the name of the user group whose privileges you wish to edit.



Note You cannot edit privileges for any of the default user groups—`basic`, `netadmin`, `operator`, `network_operations`, and `security_operations`.

4. Click **Edit**, and edit privileges as needed.
5. Click **Save**.

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

Creating Groups Using CLI

The Cisco SD-WAN software provides default user groups: **basic**, **netadmin**, **operator**, **network_operations**, and **security_operations**. The username **admin** is automatically placed in the **netadmin** usergroup.

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

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

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

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

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

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

```
vEdge# show running-config system aaa
system
aaa
  usergroup basic
    task system read write
    task interface read write
  !
  usergroup operator
    task system read
    task interface read
    task policy read
    task routing read
    task security read
  !
user admin
  password $1$tokPB7tf$VchR2JI9Sw1/dqgkqup9S.
!
!
```

Ciscotac User Access

The Cisco SD-WAN software provides two users—**ciscotacro** and **ciscotacrw**—that are for use only by the Cisco Support team. These users are available for both cloud and on-premises installations. They operate on a consent-token challenge and token response authentication in which a new token is required for every new login session. The **ciscotacro** and **ciscotacrw** users can use this token to log in to Cisco vManage web server as well as the SSH Terminal on Cisco vManage. These users can also access Cisco vBond Orchestrators, Cisco vSmart Controllers, and Cisco vEdge devices using the SSH Terminal on Cisco vManage.

The default CLI templates include the **ciscotacro** and **ciscotacrw** user configuration. These users are enabled by default. However, a customer can disable these users, if needed.

- **ciscotacro User:** This user is part of the operator user group with only read-only privileges. This user can only monitor a configuration but cannot perform any operation that will modify the configuration of the network.
- **ciscotacrw User:** This user is part of the netadmin user group with read-write privileges. This user can modify a network configuration. In addition, only this user can access the root shell using a consent token.

For more information on managing these users, see [Manage Users, on page 53](#).

Limitations

- Only 16 concurrent sessions are supported for the **ciscotacro** and **ciscotacrw** users.
- The session duration is restricted to four hours. It is not configurable.

- The inactivity timer functionality closes user sessions that have been idle for a specified period of time. This feature is enabled by default and the timeout value is 30 minutes. However, the user configuration includes the option of extending the inactivity timer.
- A customer can remove these two users. If removed, the customer can open a case and share temporary login credentials or share the screen with the Cisco Support team for troubleshooting an issue.

Configure Sessions in Cisco vManage

Table 20: Feature History

Feature History	Release Information	Description
Configure Sessions in Cisco vManage	Cisco SD-WAN Release 20.3.1 Cisco vManage Release 20.3.1	This feature lets you see all the HTTP sessions that are open within Cisco vManage. It gives you details about the username, source IP address, domain of the user, and other information. A user with User Management Write access, or a netadmin user can trigger a log out of any suspicious user's session.

Set a Client Session Timeout in Cisco vManage

You can set a client session timeout in Cisco vManage. When a timeout is set, such as no keyboard or keystroke activity, the client is automatically logged out of the system.



Note You can edit Client Session Timeout in a multitenant environment only if you have a Provider access.

1. From the Cisco vManage menu, choose **Administration** > **Settings**.
2. Click **Client Session Timeout**.
3. Click **Edit**.
4. Click **Enabled**.
5. Specify the timeout value, in minutes.
6. Click **Save**.

Set a Session Lifetime in Cisco vManage

You can specify how long to keep your session active by setting the session lifetime, in minutes. A session lifetime indicates the amount of time for which a session can be active. If you keep a session active without

letting the session expire, you will be logged out of the session in 24 hours, which is the default session timeout value.

The default session lifetime is 1440 minutes or 24 hours.



Note You can edit Session Lifetime in a multitenant environment only if you have a Provider access.

1. From the Cisco vManage menu, choose **Administration** > **Settings**.
2. Click **Session Life Time**.
3. Click **Edit**.
4. In the **SessionLifeTime** field, specify the session timeout value, in minutes, from the drop-down list.
5. Click **Save**.

Set the Server Session Timeout in Cisco vManage

You can configure the server session timeout in Cisco vManage. The server session timeout indicates how long the server should keep a session running before it expires due to inactivity. The default server session timeout is 30 minutes.



Note Server Session Timeout is not available in a multitenant environment even if you have a Provider access or a Tenant access.

1. From the Cisco vManage menu, choose **Administration** > **Settings**.
2. Click **Server Session Timeout**.
3. Click **Edit**.
4. In the **Timeout(minutes)** field, specify the timeout value, in minutes.
5. Click **Save**.

Enable Maximum Sessions Per User

You can enable the maximum number of concurrent HTTP sessions allowed per username. If you enter 2 as the value, you can only open two concurrent HTTP sessions. If you try to open a third HTTP session with the same username, the third session is granted access, and the oldest session is logged out.



Note Maximum Session Per User is not available in a multitenant environment even if you have a Provider access or a Tenant access.

1. From the Cisco vManage menu, choose **Administration** > **Settings**.

2. Click **Max Sessions Per User**.
3. Click **Edit**.
4. Click **Enabled**.
By default, **Max Sessions Per User**, is set to **Disabled**.
5. In the **Max Sessions Per User** field, specify a value for the maximum number of user sessions.
6. Click **Save**.

Configuring RADIUS Authentication Using CLI

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

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

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

For each RADIUS server, you must configure, at a minimum, its IP address and a password, or key. You can specify the key as a clear text string up to 31 characters long or as an AES 128-bit encrypted key. The local device passes the key to the RADIUS server. The password must match the one used on the server. To configure more than one RADIUS server, include the **server** and **secret-key** commands for each server.

The remaining RADIUS configuration parameters are optional.

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

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

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

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

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

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

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

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

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

Configure SSH Authentication

Table 21: Feature History

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

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

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

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

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



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

Restrictions for SSH Authentication on Cisco SD-WAN

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

SSH Authentication using vManage on Cisco vEdge Devices

1. In vManage NMS, select the **Configuration ► Templates** screen.
2. In the **Feature Templates** tab, click **Create Template**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

3. From the **Device Model** check box, select the type of device for which you are creating the template.
4. From the **Basic Information** tab, choose **AAA** template.
5. From the **Local** section, **New User** section, enter the **SSH RSA Key**. You must enter the complete public key from the id_rsa.pub file in the SSH RSA Key text box.

Configure SSH Authentication using CLI on Cisco vEdge Devices

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

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

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

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

Types of Public Keys Supported on Cisco vEdge devices:

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

SSH Authentication using CLI

```
vm5(config)# system aaa user ssh-user password vip group tenantadmin
vm5(config-user-ssh-user)# pubkey-chain ssh-usertag key-string
AAAAB3NzaC1yc2EAAAADAQABAAQDAve2mZGFLkveIgzHm6cjqsFTyIUcgfPikgsBJDuJfMnUlhWZLh03sLvki29Og2NNSJYM3OCy0TA7pFWvpDDXQw/gD4/
Bb2TH09CBNEChdV0zA6K2fMbwOZfmw2PvNRElOzVlijjQaitd5Dqe7Ar5HGtafLwVnku9HLQUDZSfeDt8cl/ftgn8skQQXuifccTpwFhYZkth978Bqm029v8/05R
BdQOVtT3VBr9NNeC4egutS0yBNZeXWBPfrwecd4/aot38plF6jOo1DvUjn60CUUOu9TQIaSFg/dFFUB0twE0IUfMBeimRexIT+cI3z8vMLD9tqFRDAI8EUegjU7BP
vm5(config-pubkey-chain-ssh-usertag)# commit
Commit complete.
```

Configure the Authentication Order

The authentication order dictates the order in which authentication methods are tried when verifying user access to a Cisco vEdge device through an SSH session or a console port. The default authentication order is

local, then **radius**, and then **tacacs**. With the default authentication order, the authentication process occurs in the following sequence:

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

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

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

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

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

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

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

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

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

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

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

- If the authentication order is configured as **radius local**:
 - With the default authentication, local authentication is used only when all RADIUS servers are unreachable. If an authentication attempt via a RADIUS server fails, the user is not allowed to log in even if they have provided the correct credentials for local authentication.

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

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

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

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

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

Configure NAS Attributes using CLI

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

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

```
localhost$ more dictionary.viptela
# -*- text -*-
#
# dictionary.viptela
#
#
```

```
# Version:      $Id$
#
VENDOR          Viptela                      41916
BEGIN-VENDOR    Viptela
ATTRIBUTE       Viptela-Group-Name          1    string
```

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

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

To create new user groups, use this command:

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

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

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

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

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

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

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

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



Note Starting from Cisco vManage Release 20.8.1, the unknown mandatory attributes from TACACS are not allowed. The authorization fails, when a client receives the configurations with the arguments that are not supported. For information about configuring ISE for Cisco SDWAN devices, see [RADIUS and TACACS-Based User Authentication and Authorization](#).


Role-Based Access with AAA

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

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


Users and User Groups

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




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

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



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



The Cisco SD-WAN software provides the following standard user groups:

- **basic**: The basic group is a configurable group and can be used for any users and privilege levels. This group is designed to include users who have permission to both view and modify information on the device.

- **operator**: The operator group is also a configurable group and can be used for any users and privilege levels. This group is designed to include users who have permission only to view information.
 - **netadmin**: The netadmin group is a non-configurable group. By default, this group includes the **admin** user. You can add other users to this group. Users in this group are permitted to perform all operations on the device.
 - Minimum supported release: Cisco vManage Release 20.9.1
- network_operations**: The **network_operations** group is a non-configurable group. Users in this group can perform all non-security-policy operations on the device and only view security policy information. For example, users can create or modify template configurations, manage disaster recovery, and create non-security policies such as application aware routing policy or CFlowD policy.
- Minimum supported release: Cisco vManage Release 20.9.1
- security_operations**: The **security_operations** group is a non-configurable group. Users in this group can perform all security operations on the device and only view non-security-policy information. For example, users can manage umbrella keys, licensing, IPS signatures auto update, TLS/SSL proxy settings, and so on.

Users of the **network_operations** group are authorized to apply policies to a device, revoke applied policies, and edit device templates. Users of the **security_operations** group require **network_operations** users to intervene on day-0 to deploy security policy on a device and on day-N to remove a deployed security policy. However, after a security policy is deployed on a device, **security_operations** users can modify the security policy without needing the **network_operations** users to intervene.



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

Privileges for Role-Based Access

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

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

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

User Authorization Rules for Operational Commands

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

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

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

CLI Command	Any User	Admin User
clear history	X	X
commit confirm	X	X
complete-on-space	X	X
config	X	X
exit	X	X
file	X	X
help	X	X
[no] history	X	X
idle-timeout	X	X
job	X	X
logout	—	X (users in netadmin group only)
monitor	X	X
nslookup	X	X
paginate	X	X
ping	X (users in netadmin group only)	X (users in netadmin group only)
poweroff	—	X(users in netadmin group only)
prompt1	X	X
prompt2	X	X
quit	X	X

CLI Command	Any User	Admin User
reboot	—	X (users in netadmin group only)
request aaa request admin-tech request firmware request interface-reset request nms request reset request software	—	X (users in netadmin group only)
request execute request download request upload	X	X
request (everything else)	—	X
rollback (configuration mode command)	—	X (users in netadmin group only)
screen-length	X	X
screen-width	X	X
show cli	X	X
show configuration commit list	X	X
show history	X	X
show jobs	X	X
show parser dump	X	X
show running-config	X	X
show users	X	X
system aaa user <i>self</i> password <i>password</i> (configuration mode command) (Note: A user cannot delete themselves)		
tcpdump	X	X
timestamp	X	X
tools ip-route	X	X
tools netstat	X	X
tools nping	X	X
traceroute	X	X
vshell	X	X (users in netadmin group only)

User Group Authorization Rules for Operational Commands

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

Operational Command	Interface	Policy	Routing	Security	System
clear app		X			
clear app-route		X			
clear arp	X				
clear bfd			X		X
clear bgp			X		X
clear bridge	X				
clear cellular	X				
clear control				X	
clear crash					X
clear dhcp					X
clear dns					X
clear igmp			X		
clear installed-certificates				X	
clear interface	X				
clear ip			X		
clear notification					X
clear omp			X		
clear orchestrator				X	
clear ospf			X		
clear pim			X		
clear policy		X			
clear pppoe	X				
clear system					X
clear tunnel				X	
clear wlan	X				

Operational Command	Interface	Policy	Routing	Security	System
clear ztp				X	X
clock					X
debug bgp			X		
debug cellular	X				
debug cflowd		X			
debug chmgr					X
debug config-mgr					X
debug dhcp-client					X
debug dhcp-helper					X
debug dhcp-server					X
debug fpm		X			
debug ftm					X
debug igmp			X		
debug netconf					X
debug omp			X		
debug ospf			X		
debug pim			X		
debug resolver			X		
debug snmp					X
debug sysmgr					X
debug transport					X
debug ttm					X
debug vdaemon				X	X
debug vrrp				X	
debug wlan	X				
request certificate				X	
request control-tunnel				X	
request controller				X	

Operational Command	Interface	Policy	Routing	Security	System
request controller-upload				X	
request csr				X	
request device				X	
request device-upload				X	
request on-vbond-controller				X	
request port-hop				X	
request root-cert-chain				X	
request security				X	
request vedge				X	
request vedge-upload				X	
request vsmart-upload				X	
show aaa					X
show app		X			
show app-route		X			
show arp	X				
show bfd			X		X
show bgp			X		
show boot-partition					X
show bridge	X				
show cellular	X				
show certificate				X	
show clock					X
show control				X	X
show crash					X
show debugs—same as debug commands					
show dhcp					X

Operational Command	Interface	Policy	Routing	Security	System
show external-nat				X	X
show hardware					X
show igmp			X		
show interface	X				
show ip			X		X
show ipsec				X	
show licenses					X
show logging					X
show multicast			X		
show nms-server					X
show notification					X
show ntp					X
show omp		X	X		X
show orchestrator				X	
show ospf			X		
show pim			X		
show policer		X			
show policy		X			
show ppp	X				
show pppoe	X				
show reboot					X
show security-info				X	
show software					X
show system					X
show transport					X
show tunnel				X	
show uptime					X
show users					X

Operational Command	Interface	Policy	Routing	Security	System
show version					X
show vrrp	X				
show wlan	X				
show ztp				X	

User Group Authorization Rules for Configuration Commands

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

Configuration Command	Interface	Policy	Routing	Security	System
apply-policy		X			
banner					X
bfd			X		X
bridge	X				
omp		X	X		X
policy		X			
security				X	X
snmp					X
system					X
vpn interface	X				
vpn ip			X		
vpn router			X		
vpn service			X		
vpn (everything else, including creating, deleting, and naming)					X
wlan	X				

Configuring AAA using Cisco vManage Template

Table 22: Feature History

Feature Name	Release Information	Description
Authorization and Accounting	Cisco SD-WAN Release 20.5.1 Cisco vManage Release 20.5.1	This feature provides for the configuration of authorization, which authorizes commands that a user enters on a device before the commands can be executed, and accounting, which generates a record of commands that a user executes on a device.

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

Use the AAA template for Cisco vBond Orchestrators, Cisco vManage instances, Cisco vSmart Controllers, and Cisco vEdge devices.

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



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

Navigating to the Template Screen and Naming the Template

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Device Templates**, and click **Create Template**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

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

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

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

Configuring Authentication Order and Fallback

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

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

Table 24:

Parameter Name	Description
Authentication Order	<p>The default order is local, then radius, and then tacacs.</p> <p>To change the default order of authentication methods that the software tries when verifying user access to a Cisco vEdge device:</p> <ol style="list-style-type: none"> 1. Click the drop-down arrow to display the list of authentication methods. 2. In the list, click the up arrows to change the order of the authentication methods and click the boxes to select or deselect a method. <p>If you select only one authentication method, it must be local.</p>
Authentication Fallback	<p>Click On to configure authentication to fall back from RADIUS or TACACS+ to the next priority authentication method if the user cannot be authenticated or if the RADIUS or TACACS+ servers are unreachable. With the default configuration (Off), authentication falls back only if the RADIUS or TACACS+ servers are unreachable.</p>
Admin Authentication Order	<p>Have the "admin" user use the authentication order configured in the Authentication Order parameter. If you do not configure the admin authentication order, the "admin" user is always authenticated locally.</p>
Disable Netconf Logs	<p>Click On to disable the logging of Netconf events. By default, these events are logged to the auth.info and messages log files.</p>
Disable Audit Logs	<p>Click On to disable the logging of AAA events. By default, these events are logged to the auth.info and messages log files.</p>
RADIUS Server List	<p>List the tags for one or two RADIUS servers. Separate the tags with commas. You set the tag under the RADIUS tab.</p>

CLI equivalent:

```

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

```

Configuring Local Access for Users and User Groups

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

To configure local access for individual users, select **Local**.

To add a new user, from **Local** click + **New User**, and configure the following parameters:

Table 25:

Parameter Name	Description
Name	Enter a name for the user. It can be 1 to 128 characters long, and it must start with a letter. The name can contain only lowercase letters, the digits 0 through 9, hyphens (-), underscores (_), and periods (.). The name cannot contain any uppercase letters. The following usernames are reserved, so you cannot configure them: backup, basic, bin, daemon, games, gnats, irc, list, lp, mail, man, news, nobody, proxy, quagga, root, sshd, sync, sys, uucp, and www-data. Also, names that start with viptela-reserved are reserved.
Password	Enter a password for the user. Each username must have a password. Users are allowed to change their own passwords. The default password for the admin user is admin. We strongly recommended that you change this password.
Description	Enter a description for the user.
User Groups	Select from the list of configured groups. You must assign the user to at least one group. The admin user is automatically placed in the netadmin group and is the only member of this group.
SSH RSA Key(s)	Add SSH RSA Keys by clicking the + Add button. A new field is displayed in which you can paste your SSH RSA key. To remove a key, click the - button. Devices support a maximum of 10 SSH RSA keys.

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

To configure local access for user groups, you first place the user into either the basic or operator group. The admin is automatically placed in the netadmin group. Then you configure user groups. To make this configuration, from **Local** select **User Group**.

Click + **New User Group**, and configure the following parameters:

Table 26:

Parameter Name	Description
Name	Name of an authentication group. It can be 1 to 128 characters long, and it must start with a letter. The name can contain only lowercase letters, the digits 0 through 9, hyphens (-), underscores (_), and periods (.). The name cannot contain any uppercase letters. The Cisco SD-WAN software provides three standard user groups, basic, netadmin, and operator. The user admin is automatically placed in the group netadmin and is the only user in this group. All users learned from a RADIUS or TACACS+ server are placed in the group basic. All users in the basic group have the same permissions to perform tasks, as do all users in the operator group. The following groups names are reserved, so you cannot configure them: adm, audio, backup, bin, cdrom, dialout, dip, disk, fax, floppy, games, gnats, input, irc, kmem, list, lp, mail, man, news, nogroup, plugdev, proxy, quagga, quaggavty, root, sasl, shadow, src, sshd, staff, sudo, sync, sys, tape, tty, uucp, users, utmp, video, voice, and www-data. Also, group names that start with the string viptela-reserved are reserved.

Parameter Name	Description
Feature Type	Click Preset to display a list of preset roles for the user group. Click Custom to display a list of authorization tasks that have been configured.
Feature	<p>The Preset list in the feature table lists the roles for the user group. These roles are Interface, Policy, Routing, Security, and System. Each role allows the user group to read or write specific portions of the device's configuration and to execute specific types of operational commands. Click the appropriate boxes for Read, Write, and None to assign privileges to the group for each role.</p> <p>The Custom list in the feature table lists the authorization tasks that you have created (see "Configure Authorization). To associate a task with this user group, choose Read, Write, or both options. The Read option grants to users in this user group read authorization to XPath's as defined in the task. The Write option allows users in this user group write access to XPath's as defined in the task.</p>

Click **Add** to add the new user group.

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

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

CLI equivalent:

```

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

```

Configuring RADIUS Authentication

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

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

Table 27:

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

To configure a connection to a RADIUS server, from **RADIUS**, click + **New Radius Server**, and configure the following parameters:

Table 28:

Parameter Name	Description
Address	Enter the IP address of the RADIUS server host.
Tag	Enter a text string to identify the RADIUS server. The tag can be 4 to 16 characters long. The tag allows you to configure authentication for AAA, IEEE 802.1X, and IEEE 802.11i to use a specific RADIUS server or servers. For Cisco vEdge devices running Cisco SD-WAN software, this field is ignored.
Authentication Port	Enter the UDP destination port to use for authentication requests to the RADIUS server. If the server is not used for authentication, configure the port number to be 0. <i>Default:</i> Port 1812
Accounting Port	Enter the UDP port to use to send 802.1X and 802.11i accounting information to the RADIUS server. <i>Range:</i> 0 through 65535. <i>Default:</i> 1813.
Key	This field is deprecated. Use the Secret Key field instead.
Secret Key	Enter the key the Cisco vEdge device passes to the RADIUS server for authentication and encryption. You can type the key as a text string from 1 to 31 characters long, and it is immediately encrypted, or you can type an AES 128-bit encrypted key. The key must match the AES encryption key used on the RADIUS server.
Source Interface	Enter the name of the interface on the local device to use to reach the RADIUS server.
VPN ID	Enter the number of the VPN in which the RADIUS server is located or through which the server can be reached. If you configure multiple RADIUS servers, they must all be in the same VPN.
Priority	Enter the priority of a RADIUS server. A server with a lower number is given priority. <i>Range:</i> 0 through 7. <i>Default:</i> 0

Click **Add** to add the new RADIUS server.

To add another RADIUS server, click + **New RADIUS Server** again.

To remove a server, click the trash icon.

CLI equivalent:

```

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

```

Configuring TACACS+ Authentication

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

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

Table 29:

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

To configure a connection to a TACACS+ server, from **TACACS**, click + **New TACACS Server**, and configure the following parameters:

Table 30:

Parameter Name	Description
Address	Enter the IP address of the TACACS+ server host.
Authentication Port	Enter the UDP destination port to use for authentication requests to the TACACS+ server. If the server is not used for authentication, configure the port number to be 0. <i>Default: Port 49</i>
Key (Deprecated)	This field is deprecated. Use the Secret Key field instead.
Secret Key	Enter the key the Cisco vEdge device passes to the TACACS+ server for authentication and encryption. You can type the key as a text string from 1 to 31 characters long, and it is immediately encrypted, or you can type an AES 128-bit encrypted key. The key must match the AES encryption key used on the TACACS+ server.
Source Interface	Enter the name of the interface on the local device to use to reach the TACACS+ server.
VPN ID	VPN in which the TACACS+ server is located or through which the server can be reached. If you configure multiple TACACS+ servers, they must all be in the same VPN.
Priority	Set the priority of a TACACS+ server. A server with lower priority number is given priority over one with a higher number. <i>Range: 0 through 7Default: 0</i>

Click **Add** to add the new TACACS server.

To add another TACACS server, click + **New TACACS Server** again.

To remove a server, click the trash icon.

CLI equivalent:

```

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

```

Configure Authorization and Accounting

Table 31: Feature History

Feature Name	Release Information	Description
Authorization and Accounting	Cisco SD-WAN Release 20.5.1 Cisco vManage Release 20.5.1	This feature provides for the configuration of authorization, which authorizes commands that a user enters on a device before the commands can be executed, and accounting, which generates a record of commands that a user executes on a device.

Configuring Authorization

You can configure authorization, which causes the device to authorize commands that users enter on a device before the commands can be executed.

Configuring authorization involves creating one or more tasks. A task consists of a set of operational commands and a set of configuration commands. Operational commands are show commands and exec commands. Configuration commands are the XPath of configuration commands.

You define the default user authorization action for each command type. The default action can be accept or deny. You also can define user authorization accept or deny actions for individual commands or for XPath strings within a command type. In this way, you can override the default action for specific commands as needed.

A task is mapped to a user group, so all users in the user group are granted the authorizations that the command sets in the task define.

To configure authorization, choose the **Authorization** tab, click + **New Task**, and configure the following parameters:

Table 32:

Parameter Name	Description
Name	Enter a unique name for the task

Parameter Name	Description
+ Add Oper	<p>Click to add a set of operational commands. In the Add Oper window that pops up:</p> <ol style="list-style-type: none"> From the Default action drop-down list, choose the default authorization action for operational commands. Choose accept to grant user authorization by default, or choose deny to prevent user authorization by default. To designate specific operational commands for which user authorization is granted or denied authorization, click + Add Oper to expand the Add Oper area. In the Oper field that displays, click accept to grant user authorization for a command, or click deny to prevent user authorization for a command, and enter the command in the CLI field. Then click Add in the Add Oper area. <p>Do not include quotes or a command prompt when entering a command. For example, config terminal is a valid entry, but "config terminal" is not valid.</p> <p>Repeat this Step 2 as needed to designate other commands.</p> <p>The actions that you specify here override the default action. In this way, you can designate specific commands that are not authorized when the default action is accept, and designate specific commands that are authorized when the default action is deny.</p> <p>To remove a specific command, click the trash icon on the right side of its line in the table at the bottom of the Add Oper window.</p> Click Add at the bottom right of the Add Oper window.
+ Add Config	<p>Click to add a set of XPath strings for configuration commands. In the Add Config window that pops up:</p> <ol style="list-style-type: none"> From the Default action drop-down list, choose the default authorization action for configuration commands. Choose accept to grant user authorization by default, or choose deny to prevent user authorization by default. To designate specific configuration command XPath strings for which user is granted or denied authorization Click + Add Config to expand the Add Config area. In the Config field that displays, click accept to grant user authorization for an XPath, or click deny to prevent user authorization for an XPath, and enter the XPath string in the CLI field. Then click Add in the Add Config area. <p>To display the XPath for a device, enter the show running-config display xpath command on the device.</p> <p>Do not include quotes or a command prompt when entering an XPath string.</p> <p>Repeat this Step 2 as needed to designate other XPath strings.</p> <p>The actions that you specify here override the default action. In this way, you can designate specific XPath strings that are not authorized when the default action is accept, and designate specific XPath strings that are authorized when the default action is deny.</p> <p>To remove a specific command, click the trash icon on the right side of its line in the table at the bottom of the Add Config window.</p> Click Add at the bottom right of the Add Config window.

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

After you create a tasks, perform these actions:

- Create or update a user group. Use the Custom feature type to associate one or more tasks with the user group by assigning read, write, or both privileges to each task. See [Configure Local Access for Users and User Groups](#).



Note A user group can be associated with either a predefined task or with user-defined tasks. Associating a user group with a combination of both predefined and user-defined tasks is not supported.

- Add users to the user group. These users then receive the authorization for operational and configuration commands that the tasks that are associated with the user group define. See [Configure Local Access for Users and User Groups](#).

If a user is attached to multiple user groups, the user receives the authorization access that is configured for the last user group that was created.

CLI equivalent:

```
system aaa
  accounting
  task name
    config
      default-action {accept | deny}
      accept "xpath"
      deny "xpath"
    oper-exec
      default-action {accept | deny}
      accept "command"
      deny "command-id"
  usergroup group-name
    task authorization-task {read | write}
```

Configuring Accounting

You can configure accounting, which causes a TACACS+ server to generate a record of commands that a user executes on a device.



Note Accounting does not generate a record of CLI commands for Cisco vManage template configuration.

Prerequisites

- The TACACS+ server must be configured with a secret key on the **TACACS** tab
- The TACACS+ server must be configured as first in the authentication order on the **Authentication** tab

To configure accounting, choose the **Accounting** tab and configure the following parameter:

Table 33:

Parameter Name	Description
Enable/disable user accounting	Click On to enable the accounting feature. Click Off to disable this feature.

CLI equivalent:

```
system aaa
  accounting
```

Configuring Password Policy for AAA on Devices

In Cisco vManage Release 20.4.1, you can create password policies using Cisco AAA on Cisco vEdge devices. We recommend configuring a password policy to ensure that all users or users of a specific group are prompted to use strong passwords. You can customize the password policy to meet the requirements of your organization.



Note You can only configure password policies for Cisco AAA using device CLI templates.

You can configure the following parameters:

password-policy min-password-length <i>length</i>	The minimum allowed length of a password. You can specify between 8 to 32 characters.
password-policy num-lower-case-characters <i>number-of-lower-case-characters</i>	The minimum number of lower case characters. You can specify between 1 to 128 characters.
password-policy num-numeric-characters <i>number-of-numeric-characters</i>	The minimum number of numeric characters. You can specify between 1 to 128 characters.
password-policy num-special-characters <i>number-of-special-characters</i>	The minimum number of special characters. You can specify between 1 to 128 characters.
password-policy num-upper-case-characters <i>number-of-upper-case-characters</i>	The minimum number of upper case characters. You can specify between 1 to 128 characters.

Configure Password Policies Using Cisco vManage

Table 34: Feature History

Feature Name	Release Information	Description
Support for Password Policies using Cisco AAA	Cisco SD-WAN Release 20.4.1 Cisco vManage Release 20.4.1	This feature allows you to create password policies for Cisco AAA. Password policies ensure that your users use strong passwords and can be customized based on your requirements. To configure password policies, push the <code>password-policy</code> commands to your device using Cisco vManage device CLI templates. For more information on the <code>password-policy</code> commands, see the aaa command reference page .

Configure password policies for Cisco AAA by doing the following:

1. Navigate to **Configuration > Templates**.
2. Click **Device Templates**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. Click **Create Template**.
4. Click **CLI Template**.
5. From the **Device Model** drop-down list, choose your Cisco vEdge device.
6. Enter a **Template Name**.
7. Enter a **Description**.
8. (Optional) From the **Load Running config from reachable device:** drop-down list, choose a device from which to load the running configuration.
9. Enter or append the password policy configuration.
For more information on the `password-policy` commands, see the [aaa command reference page](#).
10. Click **Add**.
The device templates page appears.
11. Attach the templates to your devices as described in [Attach a Device Template to Devices](#).

Configuring IEEE 802.1X and IEEE 802.11i Authentication

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

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

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

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

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

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

Configure RADIUS Authentication Servers

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

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

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

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

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

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

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

You must configure a tag to identify the RADIUS server:

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

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

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

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

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

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

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

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

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

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

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

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

The port number can be from 1 through 65535.

Configure IEEE 802.1X Port Security

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

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

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

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

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

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

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

Enable RADIUS Accounting

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

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

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

The time can be from 0 through 7200 seconds.

Enable MAC Authentication Bypass

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

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

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

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

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

To enable MAB on the RADIUS server:

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

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

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

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

Configure VLANs for Authenticated and Unauthenticated Clients

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

You can create the following kinds of VLAN:

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

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

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

```
vEdge(config)# bridge bridge-id
vEdge(config-bridge)# name text
vEdge(config-bridge)# vlan vlan-id
```

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

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

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

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

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

Then configure the 802.1X VLANs to handle unauthenticated clients.

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

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

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

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

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

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

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

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

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

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

```
system
...
radius
server 10.1.15.150
  tag          freerad1
  source-interface ge0/0
  secret-key   $4$L3rwZmsIic8zj4BgLEFXKw==
  priority     1
exit
server 10.20.24.150
  auth-port    2000
  acct-port    2001
  tag          freerad2
  source-interface ge0/4
```



```
        secret-key      $4$L3rwZmsIic8zj4BgLEFXKw==
        priority        2
    exit
    !
    !
    bridge 1
    name Untagged_bridge
    interface ge0/5
        no native-vlan
        no shutdown
    !
    !
    bridge 10
    name Authorize_VLAN
    vlan 10
    interface ge0/5
        no native-vlan
        no shutdown
    !
    !
    bridge 20
    name Guest_VLAN
    vlan 20
    interface ge0/5
        no native-vlan
        no shutdown
    !
    !
    bridge 30
    name Critical_VLAN
    vlan 30
    interface ge0/5
        no native-vlan
        no shutdown
    !
    !
    bridge 40
    name Restricted_VLAN
    vlan 40
    interface ge0/5
        no native-vlan
        no shutdown
    !
    !
    vpn 0
    interface ge0/0
        ip address 10.1.15.15/24
        tunnel-interface
            encapsulation ipsec
        ...
    !
    no shutdown
    !
    interface ge0/1
        ip address 60.0.1.16/24
        no shutdown
    !
    interface ge0/2
        ip address 10.1.19.15/24
        no shutdown
    !
    interface ge0/4
        ip address 10.20.24.15/24
        no shutdown
```

```

!
interface ge0/5
 dot1x
  auth-reject-vlan 40
  auth-fail-vlan 30
  guest-vlan 20
  default-vlan 10
  radius-servers freerad1
!
no shutdown
!
interface ge0/7
 ip address 10.0.100.15/24
 no shutdown
!
!
vpn 1
 interface ge0/2.1
  ip address 10.2.19.15/24
  mtu 1496
  no shutdown
!
 interface irb1
  ip address 56.0.1.15/24
  mac-address 00:00:00:00:aa:01
  no shutdown
  dhcp-server
  address-pool 56.0.1.0/25
  offer-time 600
  lease-time 86400
  admin-state up
  options
  default-gateway 56.0.1.15
!
!
!
vpn 10
 interface ge0/2.10
  ip address 10.10.19.15/24
  mtu 1496
  no shutdown
!
 interface irb10
  ip address 56.0.10.15/24
  mac-address 00:00:00:00:aa:10
  no shutdown
  dhcp-server
  address-pool 56.0.10.0/25
  offer-time 600
  lease-time 86400
  admin-state up
  options
  default-gateway 56.0.10.15
!
!
!
vpn 20
 interface ge0/2.20
  ip address 10.20.19.15/24
  mtu 1496
  no shutdown
!

```

```

interface irb20
 ip address 56.0.20.15/24
 mac-address 00:00:00:00:aa:20
 no shutdown
!
!
vpn 30
 interface ge0/2.30
 ip address 10.30.19.15/24
 mtu 1496
 no shutdown
!
 interface irb30
 ip address 56.0.30.15/24
 mac-address 00:00:00:00:aa:30
 no shutdown
!
!
vpn 40
 interface ge0/2.40
 ip address 10.40.19.15/24
 mtu 1496
 no shutdown
!
 interface irb40
 ip address 56.0.40.15/24
 mac-address 00:00:00:00:aa:40
 no shutdown
!
!
vpn 512
 interface eth0
 ip dhcp-client
 no shutdown
!
!

```

Configure Control Direction

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

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

The direction can be one of the following:

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

Configure Authentication with Wake on LAN

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

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

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

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

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

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

Configure 802.1X Host Mode

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

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

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

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

Set the Timeout for Inactive Clients

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

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

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

Enable Periodic Client Reauthentication

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

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

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

Configure Dynamic Authorization Service for RADIUS Change of Authorization

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

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

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

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

```
vEdge(config)# vpn 0 interface interface-name dot1x
vEdge(config-dot1x)# das
vEdge(config-das)# client ip-address
vEdge(config-das)# secret-key password
```

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

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

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

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

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

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

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

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

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

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

Configure RADIUS Authentication and Accounting Attributes

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

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

Attribute Number	Attribute Name	Description
1	User-Name	Name of the user to be authenticated.
5	NAS-Port	Physical port number on the Cisco vEdge device that is authenticating the user.
12	Framed-MTU	Maximum MTU configured for the user.
30	Called-Station-Id	Phone number that the user called, using dialed number identification (DNIS) or similar technology used to access the RADIUS server.
31	Calling-Station-Id	Phone number that the call came in to the server, using automatic number identification (ANI) or similar technology.
44	Acct-Session-Id	Unique session identifier.
61	NAS-Port-Type	Type of physical port on the Cisco vEdge device that is authenticating the user.
77	Connect-Info	Nature of the user's connection.
79	EAP-Message	Encapsulate Extended Access Protocol (EAP) packets, to allow the Cisco vEdge device to authenticate dial-in users via EAP without having to run EAP.
80	Message-Authenticator	Sign RADIUS Access-Requests to prevent these requests from being spoofed by ARAP, CHAP, or EAP.

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

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

Attribute Number	Attribute Name	Description
61	NAS-Port-Type	Type of physical port on the Cisco vEdge device that is authenticating the user.
77	Connect-Info	Nature of the user's connection.

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

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

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

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

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

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

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

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

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

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

Configure IEEE 802.11i Authentication

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

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

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

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

```
vEdge(config)# wlan frequency
vEdge(config-wlan)# interface vap number
```

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

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

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

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

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

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

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

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

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




CHAPTER 5

Role-Based Access Control

Table 35: Feature History

Feature Name	Release Information	Description
Role-Based Access Control By Resource Group	Cisco SD-WAN Release 20.5.1 Cisco vManage Release 20.5.1	This feature introduces role-based access control (RBAC) based on sites or resource groups. It is a method of authorizing system access for users based on a combination of user groups and resource groups. For large Cisco SD-WAN deployments across multiple geographical locations, this feature helps you to split the network administration among different regional administrators.
RBAC for Policies	Cisco vManage Release 20.6.1 Cisco SD-WAN Release 20.6.1	This feature allows you to create users and user groups with required read and write permissions for Cisco vManage policies. RBAC for policies provides users with the access to all the details of policies to help maximize the operational efficiency. It makes it easier to meet configuration requirements and guarantees that authorized users on the system are only given access to what they need.
Co-Management: Granular Role-Based Access Control for Feature Templates	Cisco vManage Release 20.7.1	This feature introduces greater granularity in assigning RBAC permissions for template use. This enables you to give a tenant self-management of network configuration tasks. Network administrators and managed service providers can use this feature to assign permissions to their end customers.
Co-Management: Improved Granular Configuration Task Permissions	Cisco vManage Release 20.9.1	To enable a user to self-manage specific configuration tasks, you can assign the user permissions to perform specific configuration tasks while excluding other tasks. This feature introduces numerous new permission options, enabling fine granularity in determining which configuration task permissions to provide to a user. .

Feature Name	Release Information	Description
RBAC for Security Operations and Network Operations Default User Groups	Cisco vManage Release 20.9.1	<p>This feature provides the following default user groups:</p> <ul style="list-style-type: none"> • network_operations user group for non-security policies • security_operations user group for security policies <p>RBAC for policies allows you to create users and user groups with the required read and write permissions for security and non-security policies. Users can perform configuration and monitoring actions only for the authorized policy type.</p>
Co-Management: Improved Granular Configuration for Resource group features	Cisco vManage Release 20.11.1	<p>To enable a user to self-manage specific configuration tasks, you can assign the user permissions to perform specific configuration tasks while excluding other tasks.</p> <p>This feature introduces new permission options for the following configuration groups and feature profiles.</p> <ul style="list-style-type: none"> • AppQoS under other feature profile • GPS under transport feature profile • Cisco VPN Interface GRE under WAN/LAN profile. • Cisco VPN Interface IPsec under WAN profile. • Cisco Multicast under LAN profile. • UCSE under other feature profile. • IPv4 Tracker and Tracker Group under transport and service feature profiles. • IPv6 DIA Tracker and Tracker Group, under transport feature profile.
Assigning Roles Locally for SSO-Authenticated Users	Cisco vManage Release 20.11.1	<p>If you are using an identity provider, such as Okta, for security assertion markup language (SAML)-based single sign-on (SSO), then in most use cases, you define user roles through the identity provider. This feature enables you to assign user groups locally in Cisco vManage, in case no roles are defined for the user by the identity provider.</p>

- [Information About RBAC, on page 101](#)
- [Restrictions for RBAC, on page 114](#)

- [Use Cases for RBAC, on page 115](#)
- [Configure RBAC, on page 115](#)
- [Configure RBAC Using the CLI, on page 145](#)
- [Verify RBAC, on page 147](#)
- [Monitor RBAC, on page 147](#)

Information About RBAC

Role-Based Access Control by VPN

Role-based access control (RBAC) is the process of restricting user access to network configurations and resources. In RBAC, users are assigned roles depending on the resources they need access to. The RBAC by VPN feature helps you to manage and control access to your network based on the VPNs. It involves setting permissions and privileges to enable access to authorized users.

RBAC by VPN

Role-based access by VPN allows a network administrator to define VPN groups with one or more network segments. The network administrator can associate a user with a VPN group that restricts user access to devices in the network and features of Cisco vManage.

RBAC by VPN provides the following restricted access to users configured with a VPN group:

- Access to VPN Dashboard
- Monitor devices, network, and application status via VPN dashboard
- VPN dashboard information restricted to devices with segments in the VPN group
- Monitor option restricted to devices with segments in the VPN group
- Interface monitoring on each device restricted to interfaces of segments in the VPN group

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 access these dashboards by choosing **Dashboard** from the Cisco vManage menu.


Role-Based Access with AAA

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

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


Users and User Groups

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




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

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



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



The Cisco SD-WAN software provides the following standard user groups:

- **basic**: The basic group is a configurable group and can be used for any users and privilege levels. This group is designed to include users who have permission to both view and modify information on the device.
- **operator**: The operator group is also a configurable group and can be used for any users and privilege levels. This group is designed to include users who have permission only to view information.
- **netadmin**: The netadmin group is a non-configurable group. By default, this group includes the **admin** user. You can add other users to this group. Users in this group are permitted to perform all operations on the device.
- Minimum supported release: Cisco vManage Release 20.9.1

network_operations: The **network_operations** group is a non-configurable group. Users in this group can perform all non-security-policy operations on the device and only view security policy information. For example, users can create or modify template configurations, manage disaster recovery, and create non-security policies such as application aware routing policy or CFlowD policy.

- Minimum supported release: Cisco vManage Release 20.9.1

security_operations: The **security_operations** group is a non-configurable group. Users in this group can perform all security operations on the device and only view non-security-policy information. For example, users can manage umbrella keys, licensing, IPS signatures auto update, TLS/SSL proxy settings, and so on.

Users of the **network_operations** group are authorized to apply policies to a device, revoke applied policies, and edit device templates. Users of the **security_operations** group require **network_operations** users to intervene on day-0 to deploy security policy on a device and on day-N to remove a deployed security policy. However, after a security policy is deployed on a device, **security_operations** users can modify the security policy without needing the **network_operations** users to intervene.



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

Privileges for Role-Based Access

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

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

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

User Authorization Rules for Operational Commands

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

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

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

CLI Command	Any User	Admin User
clear history	X	X
commit confirm	X	X
complete-on-space	X	X
config	X	X
exit	X	X
file	X	X
help	X	X
[no] history	X	X
idle-timeout	X	X
job	X	X
logout	—	X (users in netadmin group only)
monitor	X	X
nslookup	X	X
paginate	X	X
ping	X (users in netadmin group only)	X (users in netadmin group only)
poweroff	—	X(users in netadmin group only)
prompt1	X	X
prompt2	X	X
quit	X	X
reboot	—	X (users in netadmin group only)
request aaa request admin-tech request firmware request interface-reset request nms request reset request software	—	X (users in netadmin group only)
request execute request download request upload	X	X
request (everything else)	—	X
rollback (configuration mode command)	—	X (users in netadmin group only)

CLI Command	Any User	Admin User
screen-length	X	X
screen-width	X	X
show cli	X	X
show configuration commit list	X	X
show history	X	X
show jobs	X	X
show parser dump	X	X
show running-config	X	X
show users	X	X
system aaa user <i>self</i> password <i>password</i> (configuration mode command) (Note: A user cannot delete themselves)		
tcpdump	X	X
timestamp	X	X
tools ip-route	X	X
tools netstat	X	X
tools nping	X	X
traceroute	X	X
vshell	X	X (users in netadmin group only)

User Group Authorization Rules for Operational Commands

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

Operational Command	Interface	Policy	Routing	Security	System
clear app		X			
clear app-route		X			
clear arp	X				
clear bfd			X		X
clear bgp			X		X
clear bridge	X				

Operational Command	Interface	Policy	Routing	Security	System
clear cellular	X				
clear control				X	
clear crash					X
clear dhcp					X
clear dns					X
clear igmp			X		
clear installed-certificates				X	
clear interface	X				
clear ip			X		
clear notification					X
clear omp			X		
clear orchestrator				X	
clear ospf			X		
clear pim			X		
clear policy		X			
clear pppoe	X				
clear system					X
clear tunnel				X	
clear wlan	X				
clear ztp				X	X
clock					X
debug bgp			X		
debug cellular	X				
debug cflowd		X			
debug chmgr					X
debug config-mgr					X
debug dhcp-client					X

Operational Command	Interface	Policy	Routing	Security	System
debug dhcp-helper					X
debug dhcp-server					X
debug fpm		X			
debug ftm					X
debug igmp			X		
debug netconf					X
debug omp			X		
debug ospf			X		
debug pim			X		
debug resolver			X		
debug snmp					X
debug sysmgr					X
debug transport					X
debug ttm					X
debug vdaemon				X	X
debug vrrp				X	
debug wlan	X				
request certificate				X	
request control-tunnel				X	
request controller				X	
request controller-upload				X	
request csr				X	
request device				X	
request device-upload				X	
request on-vbond-controller				X	
request port-hop				X	
request root-cert-chain				X	

Operational Command	Interface	Policy	Routing	Security	System
request security				X	
request vedge				X	
request vedge-upload				X	
request vsmart-upload				X	
show aaa					X
show app		X			
show app-route		X			
show arp	X				
show bfd			X		X
show bgp			X		
show boot-partition					X
show bridge	X				
show cellular	X				
show certificate				X	
show clock					X
show control				X	X
show crash					X
show debugs—same as debug commands					
show dhcp					X
show external-nat				X	X
show hardware					X
show igmp			X		
show interface	X				
show ip			X		X
show ipsec				X	
show licenses					X
show logging					X

Operational Command	Interface	Policy	Routing	Security	System
show multicast			X		
show nms-server					X
show notification					X
show ntp					X
show omp		X	X		X
show orchestrator				X	
show ospf			X		
show pim			X		
show policer		X			
show policy		X			
show ppp	X				
show pppoe	X				
show reboot					X
show security-info				X	
show software					X
show system					X
show transport					X
show tunnel				X	
show uptime					X
show users					X
show version					X
show vrrp	X				
show wlan	X				
show ztp				X	

User Group Authorization Rules for Configuration Commands

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

Configuration Command	Interface	Policy	Routing	Security	System
apply-policy		X			
banner					X
bfd			X		X
bridge	X				
omp		X	X		X
policy		X			
security				X	X
snmp					X
system					X
vpn interface	X				
vpn ip			X		
vpn router			X		
vpn service			X		
vpn (everything else, including creating, deleting, and naming)					X
wlan	X				

RBAC By Resource Group Overview

Minimum supported releases: Cisco IOS XE Release 17.5.1a and Cisco vManage Release 20.5.1

RBAC by resource groups is a method of restricting or authorizing system access for users based on user groups and resource groups. A user group defines the privileges of a user in the system and the resource group defines the organizations (domains) to which a user is allowed access. Because users are not directly assigned privileges, management of individual user privileges is simply a matter of assigning the appropriate user and resource groups.

For large Cisco SD-WAN deployments across multiple geographical locations, you can split the network administration among different regional administrators.

Based on the user groups and resources groups to which network administrators are assigned, we can broadly classify them as Global Administrators and Regional Administrators. Global administrators have access to resources in every resource group and have full read-write privileges for all the features. Regional Administrators group have full read-write privileges for all the features, but the resources they can access is controlled by the resource groups to which they are assigned.

Global Admin

User accounts in the global resource group have access to all resources. A global admin is responsible for overseeing the entire network, but not involved in the operations of the individual devices on a daily basis. The global admin can assign devices to their corresponding regions, assign the regional admin accounts, manage the controllers, maintain sharable and centralized configurations, and when necessary, operate on the individual devices.

Any user in a single tenant setup with netadmin privileges and also part of global resource group is considered as global admin. Default admin user on Cisco vManage is also a global-admin, and that user can assign more global-admins. Global resource group encompasses all the WAN edges, controllers in the single view.

Global admin can switch to view only a specific resource group and can create templates. Local resource group admins, also called regional admins can clone the global templates and reuse them within their resource groups.

Regional Admin

The regional admins are responsible for day-to-day operations (configuration, monitoring, onboarding, and so on) for devices in their corresponding regions. They should not have access to or visibility into devices outside of their region. The following user groups can be created:

- resource group admin – full read/write access to devices in the corresponding resource group, can troubleshoot, monitor, attach or detach templates for the WAN edges in their group
- resource group operator – read-only access to WAN edges within their resource group
- resource group basic – basic access

Resource group admins can create new templates and attach or detach to the WAN edges in their group. They can also copy global templates and re-use them.

Resource group decides which resources the user has access to. However, the level of access is controlled by the existing user group.

- If user is in **resource_group_a** and user group **resource_group_admin**, they have full read/write access to all resources in **resource_group_a**.
- If user is in **resource_group_a** and user group **resource_group_operator**, they have read only access to all resources in **resource_group_a**.
- If user is in **resource_group_a** and user group **resource_group_basic**, they have read only access to interface and system resources in **resource_group_a**.

Global Resource Group

Global group is a special system pre-defined resource group that has different access control rules.

- Users within this group are considered as global-admins, who can have full access to all resources (devices, templates and policies) in the system and they can manage the resource groups and assign resources and users to groups.
- All other users have read-only access to resources within this group.
- The system default admin account (or tenantadmin account in a multi-tenant setup) is always in this group. This privilege cannot be changed. However, the admin account may add/remove other user accounts to or from this group.

IdP (SSO)-Managed Group

An identity provider (IdP) is a service that stores and verifies user identity. IdPs typically work with single sign-on (SSO) providers to authenticate users. If a user is authenticated with a SSO service of an IdP, the group information is also provided and managed by the IDP. An IdP passes the information about the user, including the user name and all the group names, where the user belongs to. Cisco vManage matches the group names with the group names stored in the database to further distinguish if a particular group name passed from IdP is for user group or resource group or VPN group.

Multi-Tenancy Support

With Cisco SD-WAN multitenancy, a service provider can manage multiple customers, called tenants, from Cisco vManage. The tenants share Cisco vManage instances, Cisco vBond Orchestrators, and Cisco vSmart Controllers. The domain name of the service provider has subdomains for each tenant. Cisco vManage is deployed and configured by the service provider. The provider enables multitenancy and creates a Cisco vManage cluster to serve tenants. Only the provider can access a Cisco vManage instance through the SSH terminal.

Provider has the following features:

- resource group is not applicable as the provider manages only the controllers.
- when provider provisions a new tenant, the default user account for the tenant is tenantadmin.
- other user accounts created by the provider are included in the default global resource group.
- when a provider creates a template for a tenant, the template is included in to the global resource group.

RBAC for Policies Overview

Minimum supported releases: Cisco SD-WAN Release 20.6.1 and Cisco vManage Release 20.6.1

RBAC for policies allows a user or user group to have selective Read and Write (RW) access to Cisco vManage policies. For example,

- A user with RW access for Cflowd policy can only configure Cflowd policy, but cannot configure application-aware routing policy.
- A user with RW access for application aware routing policy can only configure application-aware routing policy, but cannot configure other policies.

This feature is only supported for centralized and localized policies, but not supported for security policies.

Information About Granular RBAC for Templates

Minimum supported release: Cisco vManage Release 20.7.1

When setting user group permissions, you can use the following template permissions to provide an RBAC user with a specific degree of access to different types of templates. This gives you control over the types of device configurations that an RBAC user can apply.

Permission	Description
CLI Add-On Template	Provides access to the CLI add-on feature template.

Permission	Description
Device CLI Template	Provides access to the device CLI template.
SIG Template	Provides access to the SIG feature template and SIG credential template.
Other Feature Templates	Provides access to all feature templates except the SIG feature template, SIG credential template, and CLI add-on feature template.
Feature Profile	Provides access to all feature profiles.
Config Group	Provides access to all the configuration groups.

You can specify granular RBAC for each feature profile by expanding it. With the set permissions to the usergroup, ensure that you are able to access required feature profiles from **Templates > Configuration Groups**.

Single-Tenant and Multi-Tenant Scenarios

You can use granular RBAC for feature templates in single-tenant and multi-tenant Cisco vManage scenarios.

You can create user groups to assign specific permissions to a tenant's various teams, enabling teams to manage only specific network services without granting permission to use device CLI templates. It might be undesirable to give a tenant permission to apply device CLI templates, as the device CLI template can override any other template or device configuration.

For example, you can create a user group for a tenant's security operations group, giving them read/write access only to the SIG Template option, which would enable the security operations group to work on security configuration.

Information About Granular Configuration Task Permissions

From Cisco vManage Release 20.9.1, numerous user permission options are available, providing you fine granularity when assigning a user with permissions to manage specific configuration tasks related to configuration groups and feature profiles.

Information About Assigning Roles Locally to a User Defined by an Identity Provider

Minimum release: Cisco vManage Release 20.11.1

When you define users in an identity provider, such as Okta, for SAML SSO, one attribute that you can define for each user is the role.

When a user logs in to a Cisco vManage instance, Cisco vManage retrieves information about the user from the identity provider, including the user's role or roles. The roles defined in the identity provider map to user group permissions in Cisco vManage. Based on the roles of the user, Cisco vManage provides the user with the permissions defined by the corresponding user group.

You can assign roles locally (not depending on the identify provider) for a user profile that does not have a role defined in the identity provider.

If you have defined roles for a user through the identity provider and have also assigned user groups locally for the same user, the roles defined through the identity provider take priority.

The following table summarizes the ways to provide a user with specific permissions:

Using or Not Using an Identity Provider for SAML SSO	Roles Defined in the Identity Provider	How User Permissions Are Defined
Not using an identity provider	Not applicable	In Cisco vManage, assign a user to one or more user groups locally. This provides the user with the corresponding user group permissions.
Using an identity provider	Identity provider has one or more roles defined for the user.	Define roles for the user through the identity provider. Cisco vManage provides the user with the user group permissions corresponding to the roles.
	Identity provider does not have a role defined for the user.	Use the Remote User option when adding a user (Administration > Manage Users > Add User). See Add a User, on page 138 . In Cisco vManage, assign a user to one or more user groups locally. This provides the user with the corresponding user group permissions.

Benefits of RBAC

Benefits of Granular RBAC for Feature Templates

Minimum supported release: Cisco vManage Release 20.7.1

The permissions that you add for co-management are useful for providing detailed control over access to network configuration. They are useful when using Cisco SD-WAN with tenants, enabling you to provide a tenant access to specific types of templates. This enables you to give the tenant self-management of network configuration tasks within the tenant's VPN.

For information about the permissions added for co-management, see [Information About Granular RBAC for Templates, on page 112](#).

Restrictions for RBAC

Restrictions for Granular RBAC for Feature Templates

Minimum supported release: Cisco vManage Release 20.7.1

- To use any of the template restriction options that are provided for RBAC for co-management, provide permissions for the **Template Configuration** option. If a specific user role does not have any permissions assigned in the **Template Configuration** option, the **Templates** menu does not appear for the user in Cisco vManage. See [Manage Users](#).

- To enable an RBAC user to apply templates to devices, provide **Write** permission to the **Template Deploy** option.

Use Cases for RBAC

Use Cases for Assigning Roles Locally to a User Defined by an Identity Provider

Minimum release: Cisco vManage Release 20.11.1

An organization uses the identity provider, Okta, to authenticate users logging in to Cisco vManage.

A user defined through the identity provider has not been assigned any roles. A network administrator with access to Cisco vManage, but no access to the identity provider, can locally assign the user to a specific user group to provide the user with specific permissions.

Configure RBAC

Manage Users

From the Cisco vManage menu, choose **Administration** > **Manage Users** to add, edit, view, or delete users and user groups.

Please note the following:

- Only a user logged in as the **admin** user or a user who has Manage Users write permission can add, edit, or delete users and user groups from Cisco vManage.
- Each user group can have read or write permission for the features listed in this section. Write permission includes Read permission.
- All user groups, regardless of the read or write permissions selected, can view the information displayed in the Cisco vManage Dashboard.

Table 36: User Group Permissions for Different Device Types

Permissions	See This Section
User group permissions related to Cisco IOS XE SD-WAN device configuration.	User Group Permissions: Cisco IOS XE SD-WAN Devices
User group permissions related to Cisco Catalyst Wireless Gateway device configuration.	User Group Permissions: Cisco Catalyst Wireless Gateway Devices

User Group Permissions: Cisco IOS XE SD-WAN Devices

Table 37: User Group Permissions: Cisco IOS XE SD-WAN devices

Feature	Read Permission	Write Permission
Alarms	<p>Set alarm filters and view the alarms generated on the devices on the Monitor > Logs > Alarms page.</p> <p>Cisco vManage Release 20.6.x and earlier: Set alarm filters and view the alarms generated on the devices on the Monitor > Alarms page.</p>	No additional permissions.
Audit Log	<p>Set audit log filters and view a log of all the activities on the devices on the Monitor > Logs > Alarms page and the Monitor > Logs > Audit Log page.</p> <p>Cisco vManage Release 20.6.x and earlier: Set audit log filters and view a log of all the activities on the devices on the Monitor > Alarms page and the Monitor > Audit Log page.</p>	No additional permissions.
Certificates	<p>View a list of the devices in the overlay network under Configuration > Certificates > WAN Edge List.</p> <p>View a certificate signing request (CSR) and certificate on the Configuration > Certificates > Controllers window.</p>	<p>Validate and invalidate a device, stage a device, and send the serial number of valid controller devices to the Cisco vBond Orchestrator on the Configuration > Certificates > WAN Edge List window.</p> <p>Generate a CSR, install a signed certificate, reset the RSA key pair, and invalidate a controller device on the Configuration > Certificates > Controllers window.</p>

Feature	Read Permission	Write Permission
CLI Add-On Template (Minimum supported release: Cisco vManage Release 20.7.1)	View the CLI add-on feature template on the Configuration > Templates window. Note This operation requires read permission for Template Configuration .	Create, edit, delete, and copy a CLI add-on feature template on the Configuration > Templates window. Note These operations require write permission for Template Configuration . Note For information about this option, see Information About Granular RBAC for Feature Templates
Cloud OnRamp	View the cloud applications on the Configuration > Cloud OnRamp for SaaS and Configuration > Cloud OnRamp for IaaS window.	No additional permissions.
Cluster	View information about the services running on Cisco vManage, a list of devices connected to a Cisco vManage server, and the services that are available and running on all the Cisco vManage servers in the cluster on the Administration > Cluster Management window.	Change the IP address of the current Cisco vManage, add a Cisco vManage server to the cluster, configure the statistics database, edit, and remove a Cisco vManage server from the cluster on the Administration > Cluster Management window.
Colocation	View the cloud applications on the Configuration > Cloud OnRamp for Colocation window.	No additional permissions.
Config Group > Device > Deploy (Minimum supported release: Cisco vManage Release 20.9.1)	This permission does not provide any functionality.	Deploy a configuration onto Cisco IOS XE SD-WAN devices. Note To edit an existing feature configuration requires write permission for Template Configuration . For more details on deploying devices, see Deploy Devices .

Feature	Read Permission	Write Permission
<p>Device CLI Template (Minimum supported release: Cisco vManage Release 20.7.1)</p>	<p>View the device CLI template on the Configuration > Templates window.</p> <p>Note This operation requires read permission for Template Configuration.</p>	<p>Create, edit, delete, and copy a device CLI template on the Configuration > Templates window.</p> <p>Note These operations require write permission for Template Configuration.</p> <p>Note For information about this option, see Information About Granular RBAC for Feature Templates</p>
<p>Device Inventory</p>	<p>View the running and local configuration of devices, a log of template activities, and the status of attaching configuration templates to devices on the Configuration > Devices > WAN Edge List window.</p> <p>View the running and local configuration of the devices and the status of attaching configuration templates to controller devices on the Configuration > Devices > Controllers window.</p>	<p>Upload a device's authorized serial number file to Cisco vManage, toggle a device from Cisco vManage configuration mode to CLI mode, copy a device configuration, and delete the device from the network on the Configuration > Devices > WAN Edge List window.</p> <p>Add and delete controller devices from the overlay network, and edit the IP address and login credentials of a controller device on the Configuration > Devices > Controllers window.</p>

Feature	Read Permission	Write Permission
Device Monitoring	<p>View the geographic location of the devices on the Monitor > Geography window.</p> <p>View events that have occurred on the devices on the Monitor > Logs > Events page.</p> <p>Cisco vManage Release 20.6.x and earlier: View events that have occurred on the devices on the Monitor > Events page.</p> <p>View a list of devices in the network, along with device status summary, SD-WAN Application Intelligence Engine (SAIE) and Cflowd flow information, transport location (TLOC) loss, latency, and jitter information, control and tunnel connections, system status, and events on the Monitor > Devices page (only when a device is selected).</p> <p>Note In Cisco vManage Release 20.7.x and earlier releases, the SAIE flow is called the deep packet inspection (DPI) flow.</p> <p>Cisco vManage Release 20.6.x and earlier: Device information is available in the Monitor > Network page.</p>	<p>Ping a device, run a traceroute, and analyze the traffic path for an IP packet on the Monitor > Devices page (only when a device is selected).</p> <p>Note These operations require read and write permissions for Device Monitoring.</p>
Device Reboot	View the list of devices on which the reboot operation can be performed on the Maintenance > Device Reboot window.	Reboot one or more devices on the Maintenance > Device Reboot window.
Disaster Recovery	View information about active and standby clusters running on Cisco vManage on the Administration > Disaster Recovery window.	No additional permissions.

Feature	Read Permission	Write Permission
Events	View the geographic location of the devices on the Monitor > Logs > Events page. View the geographic location of the devices on the Monitor > Events page.	Ping a device, run a traceroute, and analyze the traffic path for an IP packet on the Monitor > Logs > Events page (only when a device is selected).
Feature Profile > Other > Thousandeyes (Minimum supported release: Cisco vManage Release 20.9.1)	View the ThousandEyes settings on the Configuration > Templates > (View configuration group) page, in the Other Profile section. Note This operation requires read permission for Template Configuration .	Create, edit, and delete the ThousandEyes settings on the Configuration > Templates > (Add or edit configuration group) page, in the Other Profile section. Note These operations require write permission for Template Configuration .
Feature Profile > Service > Dhcp (Minimum supported release: Cisco vManage Release 20.9.1)	View the DHCP settings on the Configuration > Templates > (View configuration group) page, in the Service Profile section. Note This operation requires read permission for Template Configuration .	Create, edit, and delete the DHCP settings on the Configuration > Templates > (Add or edit configuration group) page, in the Service Profile section. Note These operations require write permission for Template Configuration .
Feature Profile > Service > Lan/Vpn (Minimum supported release: Cisco vManage Release 20.9.1)	View the LAN/VPN settings on the Configuration > Templates > (View configuration group) page, in the Service Profile section. Note This operation requires read permission for Template Configuration .	Create, edit, and delete the LAN/VPN settings on the Configuration > Templates > (Add or edit configuration group) page, in the Service Profile section. Note These operations require write permission for Template Configuration .

Feature	Read Permission	Write Permission
Feature Profile > Service > Lan/Vpn/Interface/Ethernet (Minimum supported release: Cisco vManage Release 20.9.1)	View the Ethernet Interface settings on the Configuration > Templates > (View configuration group) page, in the Service Profile section. Note This operation requires read permission for Template Configuration .	Create, edit, and delete the Ethernet Interface settings on the Configuration > Templates > (Add or edit configuration group) page, in the Service Profile section. Note These operations require write permission for Template Configuration .
Feature Profile > Service > Lan/Vpn/Interface/Svi (Minimum supported release: Cisco vManage Release 20.9.1)	View the SVI Interface settings on the Configuration > Templates > (View configuration group) page, in the Service Profile section. Note This operation requires read permission for Template Configuration .	Create, edit, and delete the SVI Interface settings on the Configuration > Templates > (Add or edit configuration group) page, in the Service Profile section. Note These operations require write permission for Template Configuration .
Feature Profile > Service > Routing/Bgp (Minimum supported release: Cisco vManage Release 20.9.1)	View the Routing/BGP settings on the Configuration > Templates > (View configuration group) page, in the Service Profile section. Note This operation requires read permission for Template Configuration .	Create, edit, and delete the Routing/BGP settings on the Configuration > Templates > (Add or edit configuration group) page, in the Service Profile section. Note These operations require write permission for Template Configuration .
Feature Profile > Service > Routing/Ospf (Minimum supported release: Cisco vManage Release 20.9.1)	View the Routing/OSPF settings on the Configuration > Templates > (View configuration group) page, in the Service Profile section. Note This operation requires read permission for Template Configuration .	Create, edit, and delete the Routing/OSPF settings on the Configuration > Templates > (Add or edit configuration group) page, in the Service Profile section. Note These operations require write permission for Template Configuration .

Feature	Read Permission	Write Permission
<p>Feature Profile > Service > Switchport</p> <p>(Minimum supported release: Cisco vManage Release 20.9.1)</p>	<p>View the Switchport settings on the Configuration > Templates > (View configuration group) page, in the Service Profile section.</p> <p>Note This operation requires read permission for Template Configuration.</p>	<p>Create, edit, and delete the Switchport settings on the Configuration > Templates > (Add or edit configuration group) page, in the Service Profile section.</p> <p>Note These operations require write permission for Template Configuration.</p>
<p>Feature Profile > Service > Wirelesslan</p> <p>(Minimum supported release: Cisco vManage Release 20.9.1)</p>	<p>View the Wireless LAN settings on the Configuration > Templates > (View configuration group) page, in the Service Profile section.</p> <p>Note This operation requires read permission for Template Configuration.</p>	<p>Create, edit, and delete the Wireless LAN settings on the Configuration > Templates > (Add or edit configuration group) page, in the Service Profile section.</p> <p>Note These operations require write permission for Template Configuration.</p>
<p>Feature Profile > System > Interface/Ethernet > Aaa</p> <p>(Minimum supported release: Cisco vManage Release 20.9.1)</p>	<p>View the AAA settings on the Configuration > Templates > (View configuration group) page, in the System Profile section.</p> <p>Note This operation requires read permission for Template Configuration.</p>	<p>Create, edit, and delete the AAA settings on the Configuration > Templates > (Add or edit configuration group) page, in the System Profile section.</p> <p>Note These operations require write permission for Template Configuration.</p>
<p>Feature Profile > System > Interface/Ethernet > Banner</p> <p>(Minimum supported release: Cisco vManage Release 20.9.1)</p>	<p>View the Banner settings on the Configuration > Templates > (View configuration group) page, in the System Profile section.</p> <p>Note This operation requires read permission for Template Configuration.</p>	<p>Create, edit, and delete the Banner settings on the Configuration > Templates > (Add or edit configuration group) page, in the System Profile section.</p> <p>Note These operations require write permission for Template Configuration.</p>

Feature	Read Permission	Write Permission
Feature Profile > System > Basic (Minimum supported release: Cisco vManage Release 20.9.1)	View the Basic settings on the Configuration > Templates > (View configuration group) page, in the System Profile section. Note This operation requires read permission for Template Configuration .	Create, edit, and delete the Basic settings on the Configuration > Templates > (Add or edit configuration group) page, in the System Profile section. Note These operations require write permission for Template Configuration .
Feature Profile > System > Bfd (Minimum supported release: Cisco vManage Release 20.9.1)	View the BFD settings on the Configuration > Templates > (View configuration group) page, in the System Profile section. Note This operation requires read permission for Template Configuration .	Create, edit, and delete the BFD settings on the Configuration > Templates > (Add or edit configuration group) page, in the System Profile section. Note These operations require write permission for Template Configuration .
Feature Profile > System > Global (Minimum supported release: Cisco vManage Release 20.9.1)	View the Global settings on the Configuration > Templates > (View configuration group) page, in the System Profile section. Note This operation requires read permission for Template Configuration .	Create, edit, and delete the Global settings on the Configuration > Templates > (Add or edit configuration group) page, in the System Profile section. Note These operations require write permission for Template Configuration .
Feature Profile > System > Logging (Minimum supported release: Cisco vManage Release 20.9.1)	View the Logging settings on the Configuration > Templates > (View configuration group) page, in the System Profile section. Note This operation requires read permission for Template Configuration .	Create, edit, and delete the Logging settings on the Configuration > Templates > (Add or edit configuration group) page, in the System Profile section. Note These operations require write permission for Template Configuration .

Feature	Read Permission	Write Permission
Feature Profile > System > Ntp (Minimum supported release: Cisco vManage Release 20.9.1)	View the NTP settings on the Configuration > Templates > (View configuration group) page, in the System Profile section. Note This operation requires read permission for Template Configuration .	Create, edit, and delete the NTP settings on the Configuration > Templates > (Add or edit configuration group) page, in the System Profile section. Note These operations require write permission for Template Configuration .
Feature Profile > System > Omp (Minimum supported release: Cisco vManage Release 20.9.1)	View the OMP settings on the Configuration > Templates > (View configuration group) page, in the System Profile section. Note This operation requires read permission for Template Configuration .	Create, edit, and delete the OMP settings on the Configuration > Templates > (Add or edit configuration group) page, in the System Profile section. Note These operations require write permission for Template Configuration .
Feature Profile > System > Snmp (Minimum supported release: Cisco vManage Release 20.9.1)	View the SNMP settings on the Configuration > Templates > (View configuration group) page, in the System Profile section. Note This operation requires read permission for Template Configuration .	Create, edit, and delete the SNMP settings on the Configuration > Templates > (Add or edit configuration group) page, in the System Profile section. Note These operations require write permission for Template Configuration .
Feature Profile > Transport > Cellular Controller (Minimum supported release: Cisco vManage Release 20.9.1)	View the Cellular Controller settings on the Configuration > Templates > (View a configuration group) page, in the Transport & Management Profile section. Note This operation requires read permission for Template Configuration .	Create, edit, and delete the Cellular Controller settings on the Configuration > Templates > (Add or edit a configuration group) page, in the Transport & Management Profile section. Note These operations require write permission for Template Configuration .

Feature	Read Permission	Write Permission
<p>Feature Profile > Transport > Cellular Profile</p> <p>(Minimum supported release: Cisco vManage Release 20.9.1)</p>	<p>View the Cellular Profile settings on the Configuration > Templates > (View a configuration group) page, in the Transport & Management Profile section.</p> <p>Note This operation requires read permission for Template Configuration.</p>	<p>Create, edit, and delete the Cellular Profile settings on the Configuration > Templates > (Add or edit a configuration group) page, in the Transport & Management Profile section.</p> <p>Note These operations require write permission for Template Configuration.</p>
<p>Feature Profile > Transport > Management/Vpn</p> <p>(Minimum supported release: Cisco vManage Release 20.9.1)</p>	<p>View the Management VPN settings on the Configuration > Templates > (View configuration group) page, in the Transport & Management Profile section.</p> <p>Note This operation requires read permission for Template Configuration.</p>	<p>Create, edit, and delete the Management VPN settings on the Configuration > Templates > (Add or edit a configuration group) page, in the Transport & Management Profile section.</p> <p>Note These operations require write permission for Template Configuration.</p>
<p>Feature Profile > Transport > Management/Vpn/Interface/Ethernet</p> <p>(Minimum supported release: Cisco vManage Release 20.9.1)</p>	<p>View the Management Ethernet Interface settings on the Configuration > Templates > (View configuration group) page, in the Transport & Management Profile section.</p> <p>Note This operation requires read permission for Template Configuration.</p>	<p>Create, edit, and delete the Management VPN and Management Internet Interface settings on the Configuration > Templates > (Add or edit a configuration group) page, in the Transport & Management Profile section.</p> <p>Note These operations require write permission for Template Configuration.</p>

Feature	Read Permission	Write Permission
<p>Feature Profile > Transport > Routing/Bgp</p> <p>(Minimum supported release: Cisco vManage Release 20.9.1)</p>	<p>View the BGP Routing settings on the Configuration > Templates > (View configuration group) page, in the Transport & Management Profile section.</p> <p>Note This operation requires read permission for Template Configuration.</p>	<p>Create, edit, and delete the BGP Routing settings on the Configuration > Templates > (Add or edit a configuration group) page, in the Transport & Management Profile section.</p> <p>Note These operations require write permission for Template Configuration.</p>
<p>Feature Profile > Transport > Tracker</p> <p>(Minimum supported release: Cisco vManage Release 20.9.1)</p>	<p>View the Tracker settings on the Configuration > Templates > (View configuration group) page, in the Transport & Management Profile section.</p> <p>Note This operation requires read permission for Template Configuration.</p>	<p>Create, edit, and delete the Tracker settings on the Configuration > Templates > (Add or edit a configuration group) page, in the Transport & Management Profile section.</p> <p>Note These operations require write permission for Template Configuration.</p>
<p>Feature Profile > Transport > Wan/Vpn</p> <p>(Minimum supported release: Cisco vManage Release 20.9.1)</p>	<p>View the Wan/Vpn settings on the Configuration > Templates > (View configuration group) page, in the Transport & Management Profile section.</p> <p>Note This operation requires read permission for Template Configuration.</p>	<p>Create, edit, and delete the Wan/Vpn settings on the Configuration > Templates > (Add or edit a configuration group) page, in the Transport & Management Profile section.</p> <p>Note These operations require write permission for Template Configuration.</p>

Feature	Read Permission	Write Permission
Feature Profile > Transport > Wan/Vpn/Interface/Cellular (Minimum supported release: Cisco vManage Release 20.9.1)	View the Wan/Vpn/Interface/Cellular settings on the Configuration > Templates > (View configuration group) page, in the Transport & Management Profile section. Note This operation requires read permission for Template Configuration .	Create, edit, and delete the Wan/Vpn/Interface/Cellular settings on the Configuration > Templates > (Add or edit a configuration group) page, in the Transport & Management Profile section. Note These operations require write permission for Template Configuration .
Feature Profile > Transport > Wan/Vpn/Interface/Ethernet (Minimum supported release: Cisco vManage Release 20.9.1)	View the Wan/Vpn/Interface/Ethernet settings on the Configuration > Templates > (View configuration group) page, in the Transport & Management Profile section. Note This operation requires read permission for Template Configuration .	Create, edit, and delete the Wan/Vpn/Interface/Ethernet settings on the Configuration > Templates > (Add or edit a configuration group) page, in the Transport & Management Profile section. Note These operations require write permission for Template Configuration .
Integration Management	View information about controllers running on Cisco vManage, on the Administration > Integration Management window.	No additional permissions.
License Management	View license information of devices running on Cisco vManage, on the Administration > License Management window.	On the Administration > License Management page, configure use of a Cisco Smart Account, choose licenses to manage, and synchronize license information between Cisco vManage and the license server.
Interface	View information about the interfaces on a device on the Monitor > Devices > Interface page. Cisco vManage Release 20.6.x and earlier: View information about the interfaces on a device on the Monitor > Network > Interface page	Edit Chart Options to select the type of data to display, and edit the time period for which to display data on the Monitor > Devices > Interface page.

Feature	Read Permission	Write Permission
Manage Users	View users and user groups on the Administration > Manage Users window.	Add, edit, and delete users and user groups from Cisco vManage, and edit user group privileges on the Administration > Manage Users window.
Other Feature Templates (Minimum supported release: Cisco vManage Release 20.7.1)	View all feature templates except the SIG feature template, SIG credential template, and CLI add-on feature template on the Configuration > Templates window. Note This operation requires read permission for Template Configuration .	Create, edit, delete, and copy all feature templates except the SIG feature template, SIG credential template, and CLI add-on feature template on the Configuration > Templates window. Note These operations require write permission for Template Configuration . Note For information about this option, see Information About Granular RBAC for Feature Templates
Policy	View the common policies for all Cisco vSmart Controllers or devices in the network on the Configuration > Policies window.	Create, edit, and delete the common policies for all Cisco vSmart Controllers or devices in the network on the Configuration > Policies window.
Policy Configuration	View the list of policies created and details about them on the Configuration > Policies window.	Create, edit, and delete the common policies for all the Cisco vSmart Controllers and devices in the network on the Configuration > Policies window.
Policy Deploy	View the current status of the Cisco vSmart Controllers to which a policy is being applied on the Configuration > Policies window.	Activate and deactivate the common policies for all Cisco vManage servers in the network on the Configuration > Policies window.

Feature	Read Permission	Write Permission
RBAC VPN	View the VPN groups and segments based on roles on the Monitor > VPN page. Cisco vManage Release 20.6.x and earlier: View the VPN groups and segments based on roles on the Dashboard > VPN Dashboard page.	Add, edit, and delete VPNs and VPN groups from Cisco vManage, and edit VPN group privileges on the Administration > VPN Groups window.
Routing	View real-time routing information for a device on the Monitor > Devices > Real-Time page. Cisco vManage Release 20.6.x and earlier: View real-time routing information for a device on the Monitor > Network > Real-Time page.	Add command filters to speed up the display of information on the Monitor > Devices > Real-Time page.
Security	View the current status of the Cisco vSmart Controllers to which a security policy is being applied on the Configuration > Security window.	Activate and deactivate the security policies for all Cisco vManage servers in the network on the Configuration > Security window.
Security Policy Configuration	Activate and deactivate the common policies for all Cisco vManage servers in the network on the Configuration > Security > Add Security Policy window.	Activate and deactivate the security policies for all Cisco vManage servers in the network on the Configuration > Security > Add Security Policy window.
Session Management	View user sessions on the Administration > Manage Users > User Sessions window.	Add, edit, and delete users and user groups from Cisco vManage, and edit user sessions on the Administration > Manage Users > User Sessions window.
Settings	View the organization name, Cisco vBond Orchestrator DNS or IP address, certificate authorization settings, software version enforced on a device, custom banner on the Cisco vManage login page, and the current settings for collecting statistics on the Administration > Settings window.	Edit the organization name, Cisco vBond Orchestrator DNS or IP address, certificate authorization settings, software version enforced on a device, custom banner on the Cisco vManage login page, current settings for collecting statistics, generate a certificate signing request (CSR) for a web server certificate, and install a certificate on the Administration > Settings window.

Feature	Read Permission	Write Permission
<p>SIG Template (Minimum supported release: Cisco vManage Release 20.7.1)</p>	<p>View the SIG feature template and SIG credential template on the Configuration > Templates window.</p> <p>Note This operation requires read permission for Template Configuration.</p>	<p>Create, edit, delete, and copy a SIG feature template and SIG credential template on the Configuration > Templates window.</p> <p>Note These operations require write permission for Template Configuration.</p> <p>Note For information about this option, see Information About Granular RBAC for Feature Templates</p>
<p>Software Upgrade</p>	<p>View a list of devices, the custom banner on Cisco vManage on which a software upgrade can be performed, and the current software version running on a device on the Maintenance > Software Upgrade window.</p>	<p>Upload new software images on devices, upgrade, activate, and delete a software image on a device, and set a software image to be the default image on devices on the Maintenance > Software Upgrade window.</p>
<p>System</p>	<p>View system-wide parameters configured using Cisco vManage templates on the Configuration > Templates > Device Templates window.</p> <p>Note In Cisco vManage Release 20.7.x and earlier releases, Device Templates is called Device.</p>	<p>Configure system-wide parameters using Cisco vManage templates on the Configuration > Templates > Device Templates window.</p> <p>Note In Cisco vManage Release 20.7.x and earlier releases, Device Templates is called Device.</p>

Feature	Read Permission	Write Permission
Template Configuration	View feature and device templates on the Configuration > Templates window.	Create, edit, delete, and copy a feature or device template on the Configuration > Templates window. Note Beginning with Cisco vManage Release 20.7.1, to create, edit, or delete a template that is already attached to a device, the user requires write permission for the Template Deploy option.
Template Deploy	View the devices attached to a device template on the Configuration > Templates window.	Attach a device to a device template on the Configuration > Templates window.
Tools	Use the admin tech command to collect the system status information for a device on the Tools > Operational Commands window.	Use the admin tech command to collect the system status information for a device, and use the interface reset command to shut down and then restart an interface on a device in a single operation on the Tools > Operational Commands window. Rediscover the network to locate new devices and synchronize them with Cisco vManage on the Tools > Operational Commands window. Establish an SSH session to the devices and issue CLI commands on the Tools > Operational Commands window.
vAnalytics	Launch vAnalytics on Cisco vManage > vAnalytics window.	No additional permissions.
Workflows	Launch workflow library from Cisco vManage > Workflows window.	No additional permissions.

Feature	Read Permission	Write Permission
<p>Config Group > Device > Deploy (Minimum supported release: Cisco vManage Release 20.11.1)</p>	<p>View the devices associated to a configuration group on the Configuration > Templates > Edit Configuration Group > Associated Devices window.</p>	<p>Deploy a configuration onto Cisco IOS XE SD-WAN devices.</p> <p>Note To edit an existing feature configuration requires write permission for Template Configuration.</p> <p>For more details on deploying devices, see Deploy Devices.</p>
<p>Feature Profile > Transport > IPv4 Tracker and Tracker Group (Minimum supported release: Cisco vManage Release 20.11.1)</p>	<p>View the IPv4 Tracker and Tracker Group settings on the Configuration > Templates > (View configuration group) page, in the Transport & Management Profile section.</p> <p>Note This operation requires read permission for Template Configuration.</p>	<p>Create, edit, and delete the IPv4 Tracker and Tracker Group settings on the Configuration > Templates > (Add or edit a configuration group) page, in the Transport & Management Profile section.</p> <p>Note These operations require write permission for Template Configuration.</p>
<p>Feature Profile > Transport > IPv6 Tracker and Tracker Group (Minimum supported release: Cisco vManage Release 20.11.1)</p>	<p>View the IPv6 Tracker and Tracker Group settings on the Configuration > Templates > (View configuration group) page, in the Transport & Management Profile section.</p> <p>Note This operation requires read permission for Template Configuration.</p>	<p>Create, edit, and delete the IPv6 Tracker and Tracker Group settings on the Configuration > Templates > (Add or edit a configuration group) page, in the Transport & Management Profile section.</p> <p>Note These operations require write permission for Template Configuration.</p>

Feature	Read Permission	Write Permission
<p>Feature Profile > Transport > Gps</p> <p>(Minimum supported release: Cisco vManage Release 20.11.1)</p>	<p>View the GPS settings on the Configuration > Templates > (View configuration group) page, in the Transport & Management Profile section.</p> <p>Note This operation requires read permission for Template Configuration.</p>	<p>Create, edit, and delete the Gps settings on the Configuration > Templates > (Add or edit a configuration group) page, in the Transport & Management Profile section.</p> <p>Note These operations require write permission for Template Configuration.</p>
<p>Feature Profile > Other > APPQoS</p> <p>(Minimum supported release: Cisco vManage Release 20.11.1)</p>	<p>View the APPQoS settings on the Configuration > Templates > (View configuration group) page, in the Other section.</p> <p>Note This operation requires read permission for Template Configuration.</p>	<p>Create, edit, and delete the APPQoS settings on the Configuration > Templates > (Add or edit a configuration group) page, in the Other section.</p> <p>Note These operations require write permission for Template Configuration.</p>
<p>Feature Profile > Other > UCSE</p> <p>(Minimum supported release: Cisco vManage Release 20.11.1)</p>	<p>View the UCSE settings on the Configuration > Templates > (View configuration group) page, in the Other section.</p> <p>Note This operation requires read permission for Template Configuration.</p>	<p>Create, edit, and delete the UCSE settings on the Configuration > Templates > (Add or edit a configuration group) page, in the Other section.</p> <p>Note These operations require write permission for Template Configuration.</p>
<p>Feature Profile > Wan Profile > Cisco VPN Interface IPSec</p> <p>(Minimum supported release: Cisco vManage Release 20.11.1)</p>	<p>View the Cisco VPN Interface IPSec settings on the Configuration > Templates > (View configuration group) page, in the Wan Profile section.</p> <p>Note This operation requires read permission for Template Configuration.</p>	<p>Create, edit, and delete the Cisco VPN Interface IPSec settings on the Configuration > Templates > (Add or edit a configuration group) page, in the Wan Profile section.</p> <p>Note These operations require write permission for Template Configuration.</p>

Feature	Read Permission	Write Permission
<p>Feature Profile > Wan/Lan Profile > Cisco VPN Interface GRE</p> <p>(Minimum supported release: Cisco vManage Release 20.11.1)</p>	<p>View the Cisco VPN Interface GRE settings on the Configuration > Templates > (View configuration group) page, in the Wan/Lan Profile section.</p> <p>Note This operation requires read permission for Template Configuration.</p>	<p>Create, edit, and delete the Cisco VPN Interface GRE settings on the Configuration > Templates > (Add or edit a configuration group) page, in the Wan/Lan Profile section.</p> <p>Note These operations require write permission for Template Configuration.</p>
<p>Feature Profile > Lan Profile > Cisco Multicast</p> <p>(Minimum supported release: Cisco vManage Release 20.11.1)</p>	<p>View the Cisco Multicast settings on the Configuration > Templates > (View configuration group) page, in the Lan Profile section.</p> <p>Note This operation requires read permission for Template Configuration.</p>	<p>Create, edit, and delete the Cisco Multicast settings on the Configuration > Templates > (Add or edit a configuration group) page, in the Lan Profile section.</p> <p>Note These operations require write permission for Template Configuration.</p>



Note To create Service, System and Transport feature profiles using configuration groups, you need to provide read and write permissions on the following features to access each configuration group.

- **Feature Profile > System**
- **Feature Profile > System > AAA**
- **Feature Profile > System > BFD**
- **Feature Profile > System > Banner**
- **Feature Profile > System > Basic**
- **Feature Profile > System > Logging**
- **Feature Profile > System > NTP**
- **Feature Profile > System > OMP**
- **Feature Profile > System > SNMP**
- **Feature Profile > Service**
- **Feature Profile > Service > BFD**
- **Feature Profile > Service > LAN/VPN**
- **Feature Profile > Service > LAN/VPN/Interface/Ethernet**
- **Feature Profile > Service > Routing/BGP**
- **Feature Profile > Service > Routing/OSPF**
- **Feature Profile > Service > Routing/DHCP**
- **Feature Profile > Service > Routing/Multicast**
- **Feature Profile > Transport**
- **Feature Profile > Transport > Routing/BGP**
- **Feature Profile > Transport > WAN/VPN**
- **Feature Profile > Transport > WAN/VPN/Interface/Ethernet**

For more details on configuring features using Configuration Groups, see [Feature Management](#).

User Group Permissions: Cisco Catalyst Wireless Gateway Devices

Table 38: User Group Permissions: Cisco Catalyst Wireless Gateway Devices

Feature	Read Permission	Write Permission
Feature Profile > Teleworker > Basic (Minimum supported release: Cisco vManage Release 20.9.1)	View the basic settings on the Configuration > Templates > (View mobility configuration group) page, in the Global Profile section. Note This operation requires read permission for Template Configuration .	Configure the basic settings on the Configuration > Templates > (Add or edit mobility configuration group) page, in the Global Profile section. Note This operation requires write permission for Template Configuration .
Feature Profile > Teleworker > Cellular (Minimum supported release: Cisco vManage Release 20.9.1)	View the cellular network settings on the Configuration > Templates > (View mobility configuration group) page, in the Global Profile section. Note This operation requires read permission for Template Configuration .	Configure the cellular network settings on the Configuration > Templates > (Add or edit mobility configuration group) page, in the Global Profile section. Note This operation requires write permission for Template Configuration .
Feature Profile > Teleworker > Ethernet (Minimum supported release: Cisco vManage Release 20.9.1)	View the ethernet settings on the Configuration > Templates > (View mobility configuration group) page, in the Global Profile section. Note This operation requires read permission for Template Configuration .	Configure the ethernet settings on the Configuration > Templates > (Add or edit mobility configuration group) page, in the Global Profile section. Note This operation requires write permission for Template Configuration .

Feature	Read Permission	Write Permission
<p>Feature Profile > Teleworker > NetworkProtocol</p> <p>(Minimum supported release: Cisco vManage Release 20.9.1)</p>	<p>View the network protocol settings on the Configuration > Templates > (View mobility configuration group) page, in the Global Profile section.</p> <p>Note This operation requires read permission for Template Configuration.</p>	<p>Configure the network protocol settings on the Configuration > Templates > (Add or edit mobility configuration group) page, in the Global Profile section.</p> <p>Note This operation requires write permission for Template Configuration.</p>
<p>Feature Profile > Teleworker > SecurityPolicy</p> <p>(Minimum supported release: Cisco vManage Release 20.9.1)</p>	<p>View the security policy settings on the Configuration > Templates > (View mobility configuration group) page, in the Global Profile section.</p> <p>Note This operation requires read permission for Template Configuration.</p>	<p>Configure the security policy settings on the Configuration > Templates > (Add or edit mobility configuration group) page, in the Global Profile section.</p> <p>Note This operation requires write permission for Template Configuration.</p>
<p>Feature Profile > Teleworker > Vpn</p> <p>(Minimum supported release: Cisco vManage Release 20.9.1)</p>	<p>View the VPN settings on the Configuration > Templates > (View mobility configuration group) page, in the Global Profile section.</p> <p>Note This operation requires read permission for Template Configuration.</p>	<p>Configure the VPN settings on the Configuration > Templates > (Add or edit mobility configuration group) page, in the Global Profile section.</p> <p>Note This operation requires write permission for Template Configuration.</p>
<p>Feature Profile > Teleworker > Wifi</p> <p>(Minimum supported release: Cisco vManage Release 20.9.1)</p>	<p>View the Wi-Fi settings on the Configuration > Templates > (View mobility configuration group) page, in the Global Profile section.</p> <p>Note This operation requires read permission for Template Configuration.</p>	<p>Configure the Wi-Fi settings on the Configuration > Templates > (Add or edit mobility configuration group) page, in the Global Profile section.</p> <p>Note This operation requires write permission for Template Configuration.</p>

RBAC User Group in a Multitenant Environment

The following is the list of user group permissions for role-based access control (RBAC) in a multitenant environment:

- R stands for read permission.
- W stands for write permission.

Table 39: RBAC User Group in Multitenant Environment

Feature	Provider Admin	Provider Operator	Tenant Admin	Tenant Operator
Cloud OnRamp	RW	R	RW	R
Colocation	RW	R	RW	R
RBAC VPN	RW	R	RW	R
Security	RW	R	RW	R
Security Policy Configuration	RW	R	RW	R
vAnalytics	RW	R	RW	R

Add a User

1. From the Cisco vManage menu, choose **Administration > Manage Users**.
2. By default **Users** is selected. The table displays the list of users configured in the device.
3. To edit, delete, or change password for an existing user, click **...** and click **Edit**, **Delete**, or **Change Password** respectively.
4. To add a new user, click **Add User**.
5. Add **Full Name**, **Username**, **Password**, and **Confirm Password** details.
6. If you are using an identity provider, such as Okta, for security assertion markup language (SAML)-based single sign-on (SSO), then in most use cases, you define user roles through the identity provider. If no roles are defined for the user through the identity provider, you can enable the **Remote User** option and assign user groups locally in Cisco vManage. Assigning user groups locally provides an alternate method for assigning the user with permissions.

If you enable this option, enter an email address for the user.

If you have defined roles for a user through the identity provider and have also assigned user groups locally for the same user, the roles defined through the identity provider take priority.



Note This option is available from Cisco vManage Release 20.11.1.

7. In the **User Groups** drop-down list, select the user group where you want to add a user.

8. In the **Resource Group** drop-down list, select the resource group.



Note This field is available from Cisco SD-WAN Release 20.5.1.

9. Click **Add**.

Delete a User

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

To delete a user:

1. From the Cisco vManage menu, choose **Administration > Manage Users**.
2. For the user you wish to delete, click **...**, and click **Delete**.
3. To confirm the deletion of the user, click **OK**.

Edit User Details

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

To edit user details:

1. From the Cisco vManage menu, choose **Administration > Manage Users**.
2. For the user you wish to edit, click **...**, and click **Edit**.
3. Edit the user details.

You can also add or remove the user from user groups.

4. Click **Update**.

Change a User Password

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

Before You Begin

If you are changing the password for an admin user, detach device templates from all Cisco vManage instances in the cluster before you perform this procedure. You can reattach the device templates after you complete this procedure.

To change a password for a user:

1. From the Cisco vManage menu, choose **Administration > Manage Users**.
2. For the user you wish to change the password, click **...** and click **Change Password**.
3. Enter the new password, and then confirm it.



Note Note that the user, if logged in, is logged out.

4. Click **Done**.

Check Users Logged In to a Device Using SSH Sessions

1. From the Cisco vManage menu, choose **Monitor** > **Devices**.
Cisco vManage Release 20.6.x and earlier: From the Cisco vManage menu, choose **Monitor** > **Network**.
2. Select the device you want to use under the **Hostname** column.
3. Click **Real Time**.
4. From **Device Options**, choose **AAA users** for Cisco IOS XE SD-WAN devices or **Users** for Cisco vEdge devices.

A list of users logged in to this device is displayed.

Check Users Logged In to a Device Using HTTP Sessions

1. From the Cisco vManage menu, choose **Administration** > **Manage Users**.
2. Click **User Sessions**.

A list of all the active HTTP sessions within Cisco vManage is displayed, including, username, domain, source IP address, and so on.

Manage a User Group

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

- **basic**: Includes users who have permission to view interface and system information.
- **netadmin**: Includes the admin user, by default, who can perform all operations on the Cisco vManage. You can add other users to this group.
- **operator**: Includes users who have permission only to view information.
- Minimum supported release: Cisco vManage Release 20.9.1
- **network_operations**: Includes users who can perform non-security operations on Cisco vManage, such as viewing and modifying non-security policies, attaching and detaching device templates, and monitoring non-security data.
- Minimum supported release: Cisco vManage Release 20.9.1
- **security_operations**: Includes users who can perform security operations on Cisco vManage, such as viewing and modifying security policies, and monitoring security data.

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

Delete a User Group

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

1. From the Cisco vManage menu, choose **Administration > Manage Users**.
2. Click **User Groups**.
3. Click the name of the user group you wish to delete.



Note You cannot delete any of the default user groups—basic, netadmin, operator, network_operations, and security_operations.

4. Click **Trash** icon.
5. To confirm the deletion of the user group, click **OK**.

Edit User Group Privileges

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

1. From the Cisco vManage menu, choose **Administration > Manage Users**.
2. Click **User Groups**.
3. Select the name of the user group whose privileges you wish to edit.



Note You cannot edit privileges for the any of the default user groups—basic, netadmin, operator, network_operations, and security_operations.

4. Click **Edit**, and edit privileges as needed.
5. Click **Save**.

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

Create User Groups

1. From the Cisco vManage menu, choose **Administration > Manage Users**.
2. Click **User Groups**.
3. Click **Add User Group**.
4. Enter **User Group Name**.

5. Select the **Read** or **Write** check box against feature that you want to assign to a user group.
6. Click **Add**.
7. You can view the new user group in the left navigation path. Click **Edit** to edit the existing read or write rules.
8. Click **Save**.

Configure and Manage VPN Segments

To configure VPN Segments:

1. From the Cisco vManage menu, choose **Administration > VPN Segments**. A web page displays the list of segments that are configured.
2. To edit or delete an existing segment, click **...**, and click **Edit** or **Delete**.
3. To add new segment, click **Add Segment**.
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. To add a new segment, click **Add**.

Configure and Manage VPN Groups

To configure VPN Groups:

1. From the Cisco vManage menu, choose **Administration > VPN Groups**. A web page displays the list of segments that are configured.
2. To edit or delete a VPN group, click **...**, and click **Edit** or **Delete**.
3. To view the existing VPN in the dashboard, click **...**, and click **View Dashboard**. The **VPN Dashboard** displays the device details of the VPN device configured.
4. To add new VPN group, click **Add Group**.
5. From **Create VPN Group**, enter VPN group name in the **VPN Group Name** field.
6. Enter a brief description of the VPN in the **Description** field.
7. Check **Enable User Group access** check box and enter the user group name.
8. From **Assign Segment**, click **Add Segment** drop-down list to add new or existing segment to the VPN group.
9. Enter the **Segment Name** and **VPN Number** in the respective fields.
10. To add the configure VPN group to a device, click **Add**.

Managing Resource Groups

Minimum supported releases: Cisco IOS XE Release 17.5.1a and Cisco vManage Release 20.5.1

To configure Resource Groups:

1. From the Cisco vManage menu, choose **Administration** > **Resource Groups**. The table displays a list of resource groups that are configured in Cisco vManage.
2. To edit or delete a resource group, click **...**, and click **Edit** or **Delete**.
3. To add new resource group, click **Add Resource Group**.
4. Enter **Resource Group Name** and the **Description**.
5. Under **Site ID**, enter **Range** or **Select ID(S)** from the drop-down list to include in the resource group.
6. To add the resource group to a device, click **Add**.

To add Users:

1. From the Cisco vManage menu, choose **Administration** > **Manage Users**. The Manage Users screen appears.
2. By default **Users** is selected. The table displays the list of users configured in the device.
3. To edit, delete, or change password for an existing user, click **...**, and click **Edit**, **Delete**, or **Change Password** respectively.
4. To add a new user, click **Add User**.
5. Add **Full Name**, **Username**, **Password**, and **Confirm Password** details.
6. From the **User Groups** drop-down list, select the user group where you want to add a user.
7. From the **Resource Group** drop-down list, select the resource group.



Note This field is available from Cisco SD-WAN Release 20.5.1.

8. Click **Add**.

Workflow to Configure RBAC for Policies

Minimum supported releases: Cisco SD-WAN Release 20.6.1 and Cisco vManage Release 20.6.1

To configure RBAC for policies, use the following workflow:

1. Create user groups with required Read or Write (R/W) access to selected control or data policies. For details on creating user groups, refer [Create User Groups](#).
2. Create users and assign them to required user groups. Refer [Create Users](#).
3. Create or modify or view policy configurations as required. For information about configuring policies, see [Configure Centralized Policies Using Cisco vManage](#).

Modify Policy Configurations

Minimum supported releases: Cisco SD-WAN Release 20.6.1 and Cisco vManage Release 20.6.1

1. Login to Cisco vManage with the new user details.
2. You can modify or update the configurations based on the requirement.

When you login to Cisco vManage with new user details, you can view only the user group components that are assigned to you. For more details on configuring policies, see [Policies Configuration Guide for vEdge Routers](#)

Assign Users to Configure RBAC for Policies

Minimum supported releases: Cisco SD-WAN Release 20.6.1 and Cisco vManage Release 20.6.1

To Assign User to Create or Modify a CFlowd Data Policy

To create a CFlowd user group:

1. From the Cisco vManage, choose **Administration > Manage Users**.
2. Click **User Groups** and **Add User Group**.
3. Enter **User Group Name**.
For example, cflowd-policy-only.
4. Check the Read or Write check box against the CFlowD Policy feature that you want to assign to a user group.
5. Click **Add**.
6. You can view the new user group in the left navigation path. Click **Edit** to edit the existing read or write rules.
7. Click **Save**.

To create a CFlowd user:

1. In Cisco vManage, choose **Administration > Manage Users**.
2. Click **Users**.
3. Click **Add User**.
4. In the Add New User page, enter **Full Name**, **Username**, **Password**, and **Confirm Password** details.
5. Choose **cflowd-policy-only** from the **User Groups** drop-down.
Allow the **Resource Group** to select the default resource group.
6. Click **Add**. You can view the new user in the Users window.
7. To edit the existing read or write rules for a user, click **Edit**.

To modify a Cflowd policy:

1. Login to Cisco vManage with the new user credentials.

You can view access only to CFlowd Policies as your login is assigned to **cflowd-policy-only** user group.

2. You can create, modify, or update the configurations based on the requirement.

Configure Granular RBAC for Feature Templates

Minimum supported release: Cisco vManage Release 20.7.1

To configure specific template access, create a user group and assign the read and write permissions using the permission types described in Information About RBAC for Co-Management. The permission options for limiting template access appear with the other permission options that you choose when adding a user group.

For information about granular RBAC for feature templates, see [Information About Granular RBAC for Templates, on page 112](#).

For information about adding a user group, see [Create User Groups](#).

For a list of permission types and descriptions, see [Manage Users](#).

Configure RBAC Using the CLI

Configure Users Using CLI

You can use the CLI to configure user credentials on each device. This way, you can create additional users and give them access to specific devices. The credentials that you create for a user by using the CLI can be different from the Cisco vManage credentials for the user. In addition, you can create different credentials for a user on each device. All users with the **netadmin** privilege can create a new user.

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

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

The Username can be 1 to 128 characters long, and it must start with a letter. The name can contain only lowercase letters, the digits 0 through 9, hyphens (-), underscores (_), and periods (.). The name cannot contain any uppercase letters. Because some usernames are reserved, you cannot configure them. For a list of reserved usernames, see the **aaa** configuration command in the Cisco SD-WAN Command Reference Guide.

The Password is the password for a user. Each username must have a password, and users are allowed to change their own password. The CLI immediately encrypts the string and does not display a readable version of the password. When a user logs in to a Cisco vEdge device, they have five chances to enter the correct password. After the fifth incorrect attempt, the user is locked out of the device, and must wait for 15 minutes before attempting to log in again.



Note Enclose any user passwords that contain the special character ! in double quotation marks (“”). If a double quotation is not included for the entire password, the config database (?) treats the special character as a space and ignores the rest of the password.

For example, if the password is C!sc0, use “C!sc0”.

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

The factory-default password for the admin username is admin. We strongly recommend that you modify this password the first time you configure a Cisco vEdge device :

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

Configure the password as an ASCII string. The CLI immediately encrypts the string and does not display a readable version of the password, for example:

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

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

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

The tag is a string that you defined with the **radius server tag** command, as described in the Cisco SD-WAN Command Reference Guide.

Creating Groups Using CLI

The Cisco SD-WAN software provides default user groups: **basic**, **netadmin**, **operator**, **network_operations**, and **security_operations**. The username **admin** is automatically placed in the **netadmin** usergroup.

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

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

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

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

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

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

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

Verify RBAC

Verify Granular RBAC Permissions

Minimum supported release: Cisco vManage Release 20.7.1

Use this procedure to verify the permissions that you have configured for a user group.

1. From the Cisco vManage menu, choose **Administration > Manage Users**.
2. Click **User Groups**.
3. In the pane that displays the user groups, select a user group to display the read and write permissions assigned to the user group.
4. Scroll to the permissions that control template access to verify your configuration for the user group.

Monitor RBAC

Monitor devices for VPN Groups

To monitor devices:

1. From the Cisco vManage menu, choose **Monitor > Network**.
2. Click **WAN - Edge**.
3. Select the **VPN Group** and **VPN Segment** for which you want to monitor the network.

A web page displays the list of VPN groups and segments that are configured to a device.



CHAPTER 6

Configure Devices

You can create and store configurations for all devices—the Cisco vManage systems themselves, Cisco vSmart Controllers, Cisco vBond Orchestrators, and routers— by using Cisco vManage. When the devices start up, they contact Cisco vManage, which then downloads the device configuration to the device. (A device that is starting up first contacts the Cisco vBond Orchestrator, which validates the device and then sends it the IP address of Cisco vManage.)

The general procedure for creating configuration for all devices is the same. This section provides a high-level description of the configuration procedure. It also describes the prerequisite steps that must be performed before you can create configurations and configure devices in the overlay network.

- [Device Configuration Workflow, on page 149](#)
- [Feature Templates, on page 150](#)
- [Device Templates, on page 150](#)
- [Template Variables, on page 151](#)
- [Configuration Prerequisites, on page 151](#)
- [Create a Device Template from Feature Templates, on page 152](#)
- [Default Device Templates, on page 167](#)
- [Configuring Devices using vManage, on page 168](#)

Device Configuration Workflow

Devices in the overlay network that are managed by Cisco vManage must be configured from Cisco vManage. The basic configuration procedure is straightforward:

1. Create feature templates.
 - a. From the Cisco vManage menu, choose **Configuration > Templates**.
 - b. Click **Feature Templates**, and click **Add Templates**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

2. Create device templates.
 - a. From the Cisco vManage menu, choose **Configuration > Templates**.

- b. Click **Device Templates**, and click **Create Templates**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. Attach device templates to individual devices.
 - a. From the Cisco vManage menu, choose **Configuration > Templates**.
 - b. Click **Device Templates**, and choose a template.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

- c. Click **...**, and select **Attach Devices**.

Feature Templates

Feature templates are the building blocks of complete configuration for a device. For each feature that you can enable on a device, Cisco vManage provides a template form that you fill out. The form allows you to set the values for all configurable parameters for that feature.

Because device configurations vary for different device types and the different types of routers, feature templates are specific to the type of device.

Some features are mandatory for device operation, so creating templates for these features is required. Also for the same feature, you can create multiple templates for the same device type.



Note In releases prior to Cisco SD-WAN Release 20.7.1, if you enter < or > special characters in a Cisco vManage feature template definition or description, Cisco vManage generates a 500 exception error while attempting to preview a Cisco vManage feature template.

Starting from Cisco SD-WAN Release 20.7.1, if you enter < or > special characters in a Cisco vManage feature template definition or description, the special characters are converted to their HTML equivalents, **<** and **>**. This applies to all feature templates. You no longer receive a 500 exception error when previewing a Cisco vManage feature template.

Device Templates

You create and store configurations for all devices—the Cisco vManage systems themselves, Cisco vSmart Controllers, Cisco vBond Orchestrators, and routers— by using Cisco vManage. When the devices start up, they contact Cisco vManage, which then downloads the device configuration to the device. (A device that is starting up first contacts the Cisco vBond Orchestrator, which validates the device and then sends it the IP address of Cisco vManage.)

Device templates contain complete operational configuration for a device. You create device templates by consolidating individual feature templates.

Each device template is specific for a type of device. For each device type, if multiple devices have the same configuration, you can use the same device template for them. For example, many of the routers in a network might have the same basic configuration, so you can configure them with the same templates. (You specify the differences in the templates using configuration variables, which are discussed below.) If the configurations for the same type of devices are different, you create separate device templates.

You can also create a device template by entering a CLI text-style configuration directly on Cisco vManage. Typically, you upload a text file containing the configuration text (or cut the configuration text from a text file and paste it into Cisco vManage). You can also directly type the configuration text into Cisco vManage.

From Cisco vManage Release 20.5.1, device variable page shows text area instead of text input field to configure CLI device template for the ease of configuration.

Template Variables

Within a feature template, some configuration commands and command options are identical across all device types. Others—such as a device system IP address, its geographic latitude and longitude, the timezone, and the overlay network site identifier—are variable, changing from device to device. When you attach the device template to a device, you are prompted to enter actual values for these command variables. You can do this either manually, by typing the values for each variable and for each device, or you can upload an Excel file in CSV format that contains the values for each device.

Configuration Prerequisites

Security Prerequisites

Before you can configure any device in the network, that device must be validated and authenticated so that Cisco vManage systems, Cisco vSmart Controllers, and Cisco vBond Orchestrators recognize it as being allowed in the overlay network.

To validate and authenticate the controllers in the overlay network—Cisco vManage systems, vSmart controllers, and Cisco vSmart Controllers, and Cisco vBond Orchestrators—a signed certificate must be installed on these devices.

To validate and authenticate the routers, you receive an authorized serial number file from Cisco, which lists the serial and chassis numbers for all the routers allowed in your network. Then, you upload the serial number file to Cisco vManage.

Variables Spreadsheet

The feature templates that you create most likely contain variables. To have Cisco vManage populate the variables with actual values when you attach a device template to a device, create an Excel file that lists the variable values for each device and save the file in CSV format.

In the spreadsheet, the header row contains the variable name and each row after that corresponds to a device, defining the values of the variables. The first three columns in the spreadsheet must be the following, in this order:

- `csv-deviceId`—Serial number of the device (used to uniquely identify the device). For routers, you receive the serial numbers in the authorized serial number file sent to you from Cisco. For other devices, the serial number is included in the signed certificate you receive from Symantec or from your root CA.
- `csv-deviceIP`—System IP address of the device (used to populate the **system ip address** command).
- `csv-host-name`—Hostname of the device (used to populate the **system hostname** command).

You can create a single spreadsheet for all devices in the overlay network—Cisco vSmart Controllers, Cisco vBond Orchestrators, and routers. You do not need to specify values for all variables for all devices.

Create a Device Template from Feature Templates

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

Create a Device Template from Feature Templates

To create a device template:

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Device Templates**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. Click the **Create Template** drop-down list, and select **From Feature Template**.
4. From the **Device Model** drop-down list, select the type of device for which you wish to create the template.

vManage NMS displays all the feature templates for that device type. The required feature templates are indicated with an asterisk (*), and the remaining templates are optional. The factory-default template for each feature is selected by default.
5. In the **Template Name** field, enter a name for the device template.

This field is mandatory and can contain only uppercase and lowercase letters, the digits 0 through 9, hyphens (-), and underscores (_). It cannot contain spaces or any other characters.
6. In the **Description** field, enter a description for the device template.

This field is mandatory, and it can contain any characters and spaces.
7. To view the factory-default configuration for a feature template, select the desired feature template and click **View Template**.
8. Click **Cancel** to return to the **Configuration Template** screen.

9. To create a custom template for a feature, select the desired factory-default feature template and click **Create Template**. The template form is displayed.
This form contains fields for naming the template and defining the feature parameters.
10. In the **Template Name** field, enter a name for the feature template.
This field is mandatory and can contain only uppercase and lowercase letters, the digits 0 through 9, hyphens (-), and underscores (_). It cannot contain spaces or any other characters.
11. In the **Description** field, enter a description for the feature template.
This field is mandatory, and it can contain any characters and spaces.
12. For each field, enter the desired value. You may need to click a tab or the plus sign (+) to display additional fields.
13. 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 list of the parameter field and select one of the following:

Table 40:

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

14. For some groups of parameters, you can mark the entire group as device-specific. To do this, check the **Mark as Optional Row** check box.
These parameters are then grayed out so that you cannot enter a value for them in the feature template. You enter the value or values when you attach a device to a device template.
15. Click **Save**.

16. Repeat Steps 6 through 13 to create a custom template for each additional software feature. For details on creating specific feature templates, see the templates listed in **Available Feature Templates**.

17. Click **Create**. The new configuration template is displayed in the Device Template table.

The Feature Templates column shows the number of feature templates that are included in the device template, and the Type column shows "Feature" to indicate that the device template was created from a collection of feature templates.

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

1. Click **Feature**.
2. Click **Add Template**.
3. From **Select Devices**, select the type of device for which you wish to create a template.

You can create a single feature template for features that are available on multiple device types. You must, however, create separate feature templates for software features that are available only on the device type you are configuring.

4. Select the feature template. The template form is displayed.

This form contains fields for naming the template and fields for defining the required parameters. If the feature has optional parameters, then the template form shows a plus sign (+) after the required parameters.

5. In the **Template Name** field, enter a name for the feature template.

This field is mandatory and can contain only uppercase and lowercase letters, the digits 0 through 9, hyphens (-), and underscores (_). It cannot contain spaces or any other characters.

6. In the **Description** field, enter a description for the feature template.

This field is mandatory, and it can contain any characters and spaces.

7. For each required parameter, choose the desired value, and if applicable, select the scope of the parameter. Select the scope from the drop-down list of each parameter's value box.

8. Click the plus sign (+) from the required parameters to set the values of optional parameters.

9. Click **Save**.

10. Repeat Steps 2 to 9 for each additional feature template you wish to create.

11. Click **Device**.

12. Click the **Create Template** drop-down list and select **From Feature Template**.

13. From the **Device Model** drop-down list, select the type of device for which you wish to create the device template.

vManage NMS displays the feature templates for the device type you selected. The required feature templates are indicated with an asterisk (*). The remaining templates are optional.

14. In the **Template Name** field, enter a name for the device template.

This field is mandatory and can contain only uppercase and lowercase letters, the digits 0 through 9, hyphens (-), and underscores (_). It cannot contain spaces or any other characters.

15. In the **Description** field, enter a description for the device template.
This field is mandatory, and it can contain any characters and spaces.
16. To view the factory-default configuration for a feature template, select the desired feature template and click **View Template**.
17. Click **Cancel** to return to the **Configuration Template** screen.
18. To use the factory-default configuration, click **Create** to create the device template. The new device template is displayed in the **Device Template** table. The Feature Templates column shows the number of feature templates that are included in the device template, and the Type column shows "Feature" to indicate that the device template was created from a collection of feature templates.
19. To modify the factory-default configuration, select the feature template for which you do not wish to use the factory-default template. From the drop-down list of available feature templates, select a feature template that you created.
20. Repeat Step 19 for each factory-default feature template you wish to modify.
21. Click **Create**. The new configuration template is displayed in the **Device Template** table.
The Feature Templates column shows the number of feature templates that are included in the device template, and the Type column shows "Feature" to indicate that the device template was created from a collection of feature templates.

Create a Device CLI Template

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

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Device Templates**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. Click the **Create Template** drop-down list and select **CLI Template**.
4. From the **Device Type** drop-down list, select the type of device for which you wish to create the template.
5. In the **Template Name** field, enter a name for the device template.
This field is mandatory and can contain only uppercase and lowercase letters, the digits 0 through 9, hyphens (-), and underscores (_). It cannot contain spaces or any other characters.
6. In the **Description** field, enter a description for the device template.
This field is mandatory, and it can contain any characters and spaces.
7. In the CLI Configuration box, enter the configuration either by typing it, cutting and pasting it, or uploading a file.
8. To convert an actual configuration value to a variable, select the value and click **Create Variable**. Enter the variable name, and click **Create Variable**. You can also type the variable name directly, in the format `{{variable-name}}`; for example, `{{hostname}}`.

- Click **Add**. The new device template is displayed in the Device Template table.

The **Feature Templates** column shows the number of feature templates that are included in the device template, and the Type column shows "CLI" to indicate that the device template was created from CLI text.

Manage Device Templates

Edit a Device Template

- From the Cisco vManage menu, choose **Configuration > Templates**.
- Click **Device Templates** or **Feature Templates**, and select a template.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**, and **Feature Templates** is titled **Feature**.

- Click ..., and click **Edit**.

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



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

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

Delete a Template

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

- From the Cisco vManage menu, choose **Configuration > Templates**.
- Click **Device Templates** or **Feature Templates**, and select a template.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**, and **Feature Templates** is titled **Feature**.

- Click ..., and click **Delete**.
- To confirm the deletion of the template, click **OK**.

Copy a Template

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Device Templates** or **Feature Templates**, and select a template.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**, and **Feature Templates** is titled **Feature**.

3. Click **...**, and click **Copy**.
4. Enter a new template name and description.
5. Click **Copy**.

Edit a CLI Device Template

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Device Templates**, and select a template.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. Click **...**, and click **Edit**.
4. Under **Device CLI Template**, edit the template.
5. Click **Update**.

Use Variable Values in Configuration Templates

An overlay network might have multiple devices of the same type that have nearly identical configurations. This situation most commonly occurs with routers when the routers that are located in multiple stores or branch locations provide identical services, but each individual router has its own hostname, IP address, GPS location, and other site-specific properties, such as BGP neighbors. This situation also occurs in a network with redundant controller devices, such as Cisco vSmart Controllers, which must all be configured with identical policies, and Cisco vManage systems. Again, each controller has its own individual parameters, such as hostname and IP address.

To simplify the configuration process for these devices, you can create a single configuration template that contains both static configuration values and variable values. The static values are common across all the devices, and the variable values apply only to an individual device. You provide the actual values for the variables when you attach the individual device to the device configuration template.

You can configure a variable value for a parameter in a feature configuration template in two ways:

- Select the parameter scope to be **Device Specific**—For an individual configuration parameter, select **Device Specific** to mark the parameter as a variable. Each variable must be identified by a unique text string, which is called a *key*. When you select **Device Specific**, an **Enter Key** box opens and displays the

default key. You can use the default key, or you can change it by typing a new string and then moving the cursor out of the Enter Key box.

- Mark a group of related parameters as optional—For some features in some feature configuration templates, you can mark the entire feature as optional. To mark the feature in this way, click Mark as Optional Row in a section of a feature configuration template. The variable parameters are then dimmed, and you cannot configure values for them in the feature configuration template.

You enter the device-specific values for the variables when you attach the device to the configuration, in one of the following ways:

- From a file—When you are attaching a template to a device, you load a file to the vManage NMS. This is an Excel file in CSV format that lists all the variables and defines the variable's value for each device.
- Manually—When you attach a device template to a device, the Cisco vManage prompts you for the values for each of device-specific parameters, and you type in the value for each parameter.



Note Cisco SD-WAN supports up to 500 variables in a template push operation.

Use a File for Variable Parameters

To load device-specific variable values from a file, you create a template variables file. This file is an Excel file in CSV format that lists all the variables in your the configurations of your devices and defines the values for each variable. You create this file offline and then import it into Cisco vManage server when you attach a device configuration to one or more devices in the overlay network.

We recommend that you create a template variables CSV file when your overlay network has more than a small number of Cisco vEdge devices.

CSV File Format

The CSV file is an Excel spreadsheet that contains one column for each variable that is required for the configuration of a device. The header row contains the variable names (one variable per column), and each row after that corresponds to a device and defines the values of the variables for that device.

You can create a single spreadsheet for all devices in the overlay network—Cisco vEdge devices, Cisco vManage systems, Cisco vSmart Controllers, and Cisco vBond Orchestrators—or you can create one spreadsheet for each device type. The system determines the device type from its serial number.

In the spreadsheet, for each device type and for each individual device, you specify values only for the required variables. When you do not need to specify a value for a variable, simply leave that cell blank.

The first three columns in the spreadsheet must be the following items and must be in the order shown:

Column	Column Heading	Description
1	csv-deviceId	Serial number of the device (used to uniquely identify the device). For Cisco vEdge devices, you receive the serial numbers in the authorized serial number file sent to you from Cisco. For other devices, the serial number is included in the signed certificate you receive from Symantec or from your root CA.
2	csv-deviceIP	System IP address of the device (used to populate the system ip address command).
3	csv-host-name	Hostname of the device (used to populate the system hostname command).

The headings for the remaining columns must be unique variable keys that are defined in the Enter Key box of a feature configuration template. These remaining columns can be in any order.

Generate a Skeleton CSV File

You can create a template variables CSV file manually, with the format described in the previous section, or you can have Cisco vManage generate a skeleton CSV file that contains all the required columns and column headings. This generated CSV file has one row for each Cisco device type, and it has the column headings for each of the variables that are required by all the feature templates included in the device configuration. The column heading text corresponds to the key string that identifies a device-specific parameter. Then you populate the rows with values for each variable.

To have Cisco vManage generate a skeleton CSV file:

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Feature Templates**, and click **Add Template**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

3. Create the required feature templates for one Cisco vEdge device router, one Cisco vSmart Controller, one Cisco vManage system, and one Cisco vBond Orchestrator.
In each feature template:
 - a. For fields that have default values, verify that you want to use that value for all devices. If you do not want to use the default, change the scope to **Global** or **Device-specific**.
 - b. For fields that apply to all devices, select the **Global** icon next to the field and set the desired global values.
 - c. For fields that are device specific, select the **Device-specific** icon next to the field and leave the field blank.
4. For each Cisco device type, create a device template.

5. From the Cisco vManage menu, choose **Configuration > Templates**.
6. Click **Device Templates**, and select the desired device template from the template list table.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

7. Click **...**, and click **Export CSV**.
8. Repeat Steps 7 and 8 for each device template.

Edit the exported CSV file, adding at a minimum the device serial number, device system IP address, and device hostname for each device in the overlay network. Then add values for desired device-specific variables for each device. Note that variable names cannot contain forward slashes (/), backwards slashes (\), or parentheses (()).

If desired, you can combine the CSV files into a single file.

Import a CSV File

To use the device-specific variable values in the CSV file, import the file when you are attaching a device template to the Viptela device:

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Device Templates**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. For the desired template, click **...**, and select **Attach Devices**.
4. In the **Attach Devices** dialog box, select the desired devices in **Available Devices** and click the arrow to move them to **Selected Devices**.
5. Click **Attach**.
6. Click the Up arrow. The Upload CSV File box displays.
7. Choose the CSV file to upload, and click **Upload**.

During the attachment process, click Import file to load the Excel file. If Cisco vManage detects duplicate system IP addresses for devices in the overlay network, it displays a warning message or a pop-up window. You must correct the system IP addresses to remove any duplicates before you can continue the process of attaching device templates to Viptela devices.

Manually Enter Values for Device-Specific Variables and for Optional Rows

For parameters in a feature template that you configure as device-specific, when you attach a device template to a device, Cisco vManage prompts you for the values to use for these parameters. Entering device-specific values in this manner is useful in test or POC networks, or if you are deploying a small network. This method generally does not scale well for larger networks.

For situations in which the configuration for many devices is identical except for a few parameters, in the feature configuration template, you can specify that the parameter be an optional row in the configuration. By selecting optional row, the feature template automatically marks the parameters as device-specific, and these parameters are dimmed so that you cannot set them in the template. You do not have to individually mark the parameters as device specific. Then, when you attach a device template to a device, Cisco vManage prompts you for the values to use for these parameters. Using optional rows to enter device-specific values is useful when a group of many Cisco vEdge devices provide identical services at their branch or site, but individual routers have their own hostname, IP address, GPS location, and other site or store properties, such as BGP neighbors.

Optional rows are available for some parameters in some feature configuration templates. To treat a parameter or set of parameters as an optional row, click the **Mark as Optional Row** box. For these types of parameters, the feature configuration template has a table listing all the configured parameters. The Optional column indicates which are optional rows,

To manually enter values for device-specific variables or for variables in optional rows when you attach the template to a device:

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Device Templates**, and select the desired device template.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. Click **...**, and click **Attach Devices**. The **Attach Devices** dialog box opens.
4. Choose one or more devices from **Available Devices** and move them to **Selected Devices**.
5. Click **Attach**.
6. In the **Chassis Number** list, select the desired device.
7. Click **...**, and click **Edit Device Template**. The **Update Device Template** dialog box opens.
8. Enter values for the optional parameters. When you are using optional rows, if you do not want to include the parameter for the specific device, do not specify a value.
9. Click **Update**.
10. Click **Next**.

If any devices have the same system IP address, a dialog box appears or an error message is displayed when you click **Next**. Modify the system IP addresses so that there are no duplicates, and click **Save**. Then click **Next** again.



Note You need to shut down the OMP on the device, before changing the system-ip on the device.

11. In the left pane, select the device. The right pane displays the device configuration and the **Config Preview** tab in the upper right corner is selected.
12. Click **Config Diff** to preview the differences between this configuration and the configuration currently running on the device, if applicable. To edit the variable values entered in the previous screen, click **Back**.

- Click **Configure Devices** to push the configuration to the devices.

The Status column displays whether the configuration was successfully pushed. Click the **right angle bracket** to the left of the row to display details of the push operation.

View Device Templates

View a Template

- From the Cisco vManage menu, choose **Configuration > Templates**.
- Click **Device Templates** or **Feature Templates**, and select a template you wish to view.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**, and **Feature Templates** is titled **Feature**.

- Click **...**, and then click **View**.

View Device Templates Attached to a Feature Template

- From the Cisco vManage menu, choose **Configuration > Templates**.
- Click **Feature Templates**, and select a template you wish to view.



Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

- Click **...**, and click **Show Attached Device Templates**.

Device Templates dialog box opens, displaying the names of the device templates to which the feature template is attached.

View Devices Attached to a Device Template

For a device template that you created from feature templates:

- From the Cisco vManage menu, choose **Configuration > Templates**.
- Click **Device Templates**, and select a template you wish to view.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

- Click **...**, and click **Attach Devices**.
- From **Attach Devices**, click **Attached Devices**.

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

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Device Templates**, and select a template you wish to view.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. Click **...**, and then click **Show Attached Devices**.

Attach and Detach a Device Template

To configure a device on the network, you attach a device template to the device. You can attach only one device template to a device, so the template—whether you created it by consolidating individual feature templates or by entering a CLI text-style configuration—must contain the complete configuration for the device. You cannot mix and match feature templates and CLI-style configurations.

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

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

For template operations, the following rules apply:

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



Note You need to recreate the feature templates as the templates created prior to Cisco vManage Release 20.5 fails when attached to the device.

If the device being configured is present and operational on the network, the configuration is sent to the device immediately and takes effect immediately. If the device has not yet joined the network, the pushing of the configuration to the device is scheduled. When the device joins the network, Cisco vManage pushes the configuration immediately after it learns that the device is present in the network.

Attach a Device Template to Devices

You can attach the same templates to multiple devices, and you can do so simultaneously, in a single operation.

To attach a device template to one or more devices:

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Device Templates** and select the desired template.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. Click **...**, and click **Attach Devices**. The **Attach Devices** dialog box opens with the **Select Devices** tab selected
4. In the **Available Devices** column on the left, select a group and search for one or more devices, select a device from the list, or click **Select All**.
5. Click the arrow pointing right to move the device to the **Selected Devices** column on the right.
6. Click **Attach**.
7. If the template contains variables, enter the missing variable values for each device you selected in one of the following ways:
 - Enter the values manually for each device either in the table column or by clicking **...** and **Edit Device Template**. When you are using optional rows, if you do not want to include the parameter for the specific device, do not specify a value.
 - Click **Import File** to upload a CSV file that lists all the variables and defines each variable's value for each device.
8. Click **Update**
9. Click **Next**.
If any devices have the same system IP address, a dialog box appears or an error message is displayed when you click **Next**. Modify the system IP addresses so that there are no duplicates, and click **Save**. Then click **Next** again.
10. In the left pane, select the device, to preview the configuration that is ready to be pushed to the device. The right pane displays the device's configuration and the **Config Preview** tab is selected. Click the **Config Diff** tab to view the differences between this configuration and the configuration currently running on the device, if applicable. Click the **Back** button to edit the variable values entered in the previous screen.
11. If you are attaching a Cisco vEdge device, click **Configure Device Rollback Timer** to configure the time interval at which the device rolls back to its previous configuration if the router loses its control connection to the overlay network. **The Configure Device Rollback Time** dialog box is displayed.
 - a. From the **Devices** drop-down list, select a device.
 - b. To enable the rollback timer, in the **Set Rollback slider**, drag the slider to the left to enable the rollback timer. When you do this, the slider changes in color from gray to green.
 - c. To disable the rollback timer, click the **Enable Rollback** slider. When you disable the timer, the Password field dialog box opens. Enter the password that you used to log in to the vManage NMS.
 - d. In the **Device Rollback Time slider**, drag the slider to the desired value. The default time is 5 minutes. You can configure a time from 6 to 15 minutes.

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

Export a Variables Spreadsheet in CSV Format for a Template

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Device Templates** and select the desired template.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. Click **...**, and click **Export CSV**.

Determine Why a Device Rejects a Template

When you attach a template to a device using the screen, the device might reject the template. One reason that this may occur is because the device template contains incorrect variable values. When a device rejects a template, it reverts to the previous configuration.

To determine why the device rejected the template:

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Device Templates**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. Locate the device. The **Template Status** column indicates why the device rejected the template.

Change the Device Rollback Timer

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

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Device Templates**, and choose a device template.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. Click ..., and click **Change Device Values**.
The right pane displays the device's configuration, and the **Config Preview** tab is selected.
4. In the left pane, click the name of a device.
5. Click **Configure Device Rollback Timer**. The **Configure Device Rollback Time** pop up page is displayed.
6. From the **Devices** drop-down list, select a device.
7. To enable the rollback timer, in the **Set Rollback slider** drag the slider to the left to enable the rollback timer. When you do this, the slider changes in color from gray to green.
8. To disable the rollback timer, click **Enable Rollback slider**. When you disable the timer, the **Password** field dialog box appears. Enter the password that you used to log in to the vManage NMS.
9. In the **Device Rollback Time** slider, drag the slider to the desired value. The default time is 5 minutes. You can configure a time from 6 to 15 minutes.
10. To exclude a device from the rollback timer setting, click **Add Exception** and select the devices to exclude.
11. The table of the **Configure Device Rollback Time** dialog box lists all the devices to which you are attaching the template and their rollback time. To delete a configured rollback time, click the **Trash** icon of the device name.
12. Click **Save**.
13. Click **Configure Devices** to push the configuration to the devices. The **Status** column displays whether the configuration was successfully pushed. Click (+) to display details of the push operation.

Preview Device Configuration and View Configuration Differences

For a configuration that you have created from the CLI:

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Device Templates**, and choose the desired device template.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. Click ..., and click **Change Device Values**.
The right pane displays the device's configuration, and **Config Preview** is selected.
4. Click the name of a device.
5. Click **Config Diff** to view the differences between this configuration and the configuration currently running on the device, if applicable. Click **Back** to edit the variable values entered in the previous screen.

- Click **Configure Devices** to push the configuration to the devices. The Status column displays whether the configuration was successfully pushed. Click the right angle bracket to display details of the push operation.

Change Variable Values for a Device

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

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

- From the Cisco vManage menu, choose **Configuration > Templates**.
- Click **Device Templates**, and choose the desired device template.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

- Click **...**, and click **Change Device Values**.
The screen displays a table of all the devices that are attached to that device template.
- For the desired device, click **...**, and click **Edit Device Template**.
- In the **Update Device Template** dialog box, enter values for the items in the variable list.
- Click **Update**.
- Click **Next**.
- Click **Configure Devices** to push the configuration to the device. The Status column displays if the configuration was successfully pushed or not. Click the right angle bracket to display the details of the push operation.

Default Device Templates

Table 41: Feature History

Feature Name	Release Information	Description
Default Device Templates	Cisco SD-WAN Release 20.1.1	<p>A default device template provides basic information that you can use to bring up devices in a deployment quickly.</p> <p>This feature is supported on the Cisco Cloud Services Router 1000V Series, Cisco C1111-8PLTELA Integrated Services Routers, and Cisco 4331 Integrated Services Routers.</p>

A default device template provides basic information that you can use to bring up devices in a deployment. It provides a way for you to quickly provision devices with the minimum information that they need to operate in your network.

You cannot directly edit or update information in a device default template, but you can copy the template and then edit the copy.

To use a default device template:

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Device Templates**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. From the **Template Type** drop-down list, select **Default**.

A list of default device templates displays.

4. Perform any of these actions:

- To attach a default device template to devices, click **...**, and select **Attach Devices**.

In the **Attach Devices** dialog box, select the devices that you want attach, and then click **Attach**.

- To view the configuration settings for a default device template, click **...**, and choose **View**.

- To copy a default device template, click **...**, and choose **View**.

In the **Template Copy** dialog box, enter a unique name and a description for the copy that you are creating, and then click **Copy**.

The copied version becomes a feature template that you can edit.

- To create an Excel file in CSV format that contains device-specific settings from a device template, click **...**, and choose **Export CSV**. Use the dialog box that displays to open or save the CSV file.

You can use this CSV file as a reference for device-specific settings when you create other device templates.

Configuring Devices using vManage

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

The screenshot shows the Cisco vManage interface for configuring devices. The main content area displays a table titled 'WAN Edge List' under the 'CONFIGURATION | DEVICES' section. The table has the following columns: State, Device Model, Chassis Number, Serial No./Token, Hostname, System IP, and Site ID. There are 8 rows of data. A 'Change Mode' dropdown menu is located above the table. The left sidebar shows the navigation menu with 'Configuration' and 'Devices' highlighted.

State	Device Model	Chassis Number	Serial No./Token	Hostname	System IP	Site ID
✓	vEdge Cloud	f3967a34-d454-49cd-8494-05c73d88c9...	12345703	vm11	172.16.255.21	100
✓	vEdge Cloud	537fcec5-00f8-4062-a894-9db7d11f4fd4	12345715	vm1	172.16.255.11	100
✓	vEdge Cloud	3de2abff-0251-4718-b4d2-023d2d39b9...	12345711	vm4	172.16.255.14	400
✓	ISR4331	ISR4331-FD02106254L	0153B8A8	ISR4331-SDWAN-2	172.16.255.129	1800
✓	CSR1000v	CSR-e109d8c2-7541-40d3-b3f7-362116...	12345607	CSR-cEdge2	172.16.255.134	1900
✗	CSR1000v	CSR-97a0fe05-a03e-4a1c-afb3-4f9e714...	1234560A	CSR-cEdge1	172.16.255.130	1600
✓	ISR4221	ISR4221/K9-FOC22034WR7	0254FAFB	ISR4221	172.16.255.139	2000
✓	C1111-8P	C1111-8P-FOC215124MH	0160C45F	C1111-8P	172.16.255.138	2001

3668731

Change Configuration Modes

A device can be in either of these configuration modes:

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

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

To toggle a router from vManage mode to CLI mode:

1. From the Cisco vManage menu, choose **Configuration > Devices**.
2. Click **WAN Edge List**, and select a device.
3. Click the **Change Mode** drop-down list and select **CLI mode**.



Note Starting from Cisco IOS XE SD-WAN Release 17.11.1a, click the ... icon adjacent to the device that you want to change from the vManage mode to the CLI mode and click **Config Lock (Provision Device)**.

You can use the **Config Lock (Provision Device)** only if a template is attached to a device.

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

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

1. From the Cisco vManage menu, choose **Configuration > Devices**.
2. Click **Controllers**, and select a device.
3. Click the **Change Mode** drop-down list.
4. Select **CLI mode** and then select the device type. The **Change Mode - CLI** window opens.
5. From the **vManage mode** pane, select the device and click the right arrow to move the device to the **CLI mode** pane.
6. Click **Update to CLI Mode**.

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



Note Starting from Cisco IOS XE SD-WAN Release 17.11.1a, click the ... icon adjacent to the device that you want to change from the vManage mode to the CLI mode and click **Config Lock (Provision Device)**.

You can use the **Config Lock (Provision Device)** only if a template is attached to a device.

Upload WAN Edge Router Authorized Serial Number File

The WAN eEdge router authorized serial number file contains the chassis number and the certificate serial numbers of all valid Cisco vEdge devices in the overlay network. You retrieve a serial number file from the Cisco Plug-and-Play (PnP) portal and upload it to Cisco vManage. (For more information about Cisco PnP, see [Cisco Plug and Play Support Guide for Cisco SD-WAN Products](#).) From Cisco vManage, you send the file to the controllers in the network. This file is required to allow the Cisco SD-WAN overlay network components to validate and authenticate each other and to allow the overlay network to become operational.

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

1. From the Cisco vManage menu, choose **Configuration > Devices**.
2. Click **WAN Edge List**, and click **Upload WAN Edge List**.
3. Under **Upload WAN Edge List** screen:
 - a. Click **Choose File** and select the WAN edge router authorized serial number file you received from Cisco PnP.
 - b. To automatically validate the routers and send their chassis and serial numbers to the controllers, ensure that the **Validate the uploaded vEdge List and send to controllers** check box is selected. If you do not select this option, you must individually validate each router in **Configuration > Certificates > WAN Edge List**.
 - c. Click **Upload**.

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

Starting from Cisco vManage Release 20.9.2, you can monitor the newly added WAN Edge devices in the **Monitor > Devices** page.

Upload WAN Edge Router Serial Numbers from Cisco Smart Account

To allow Cisco SD-WAN overlay network components to validate and authenticate each other and to allow the overlay network to become operational, Cisco SD-WAN requires chassis numbers of all valid Cisco vEdge devices in the overlay network.

In addition, certificate serial numbers, are required for all devices.

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

1. From the Cisco vManage menu, choose **Configuration > Devices**.
2. Click **WAN Edge List**, and click **Sync Smart Account**.
3. In the **Sync Smart Account** window:
 - a. Enter the **Username** and **Password** for your Smart account.
 - b. To automatically validate the routers and send their chassis and serial numbers to the controllers, check the **Validate the Uploaded WAN Edge List and Send to Controllers** check box. If you do not select this option, you must individually validate each router in **Configuration > Certificates > WAN Edge List**.
 - c. Click **Sync**.

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

Starting from Cisco vManage Release 20.9.2, you can monitor the newly added WAN Edge devices in the **Monitor > Devices** page.

Generate Bootstrap Configuration for a vEdge Cloud Router

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

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

1. From the Cisco vManage menu, choose **Configuration > Devices**.
2. Click **WAN Edge List**, and click **Export Bootstrap Configuration**.
3. In the **Export Bootstrap Configuration** window, in the **Bootstrap Configuration** field, click **Cloud-Init** or **Encoded String**, depending the Hypervisor you are using to bring up the vEdge Cloud router.
4. Select the devices to configure from the **Available Devices** pane, or click **Select All** to select all devices.
5. Click the right arrow to move the devices to the **Selected Devices** pane.
6. Click **Generate Configuration**. The configurations are downloaded to the vManage NMS.
7. Provision the vEdge Cloud router instance in AWS, KVM, or ESXi with the bootstrap configuration. By default, ge0/0 is the device's tunnel interface and is a DHCP client. To use an interface other than ge0/0

as the tunnel interface or to use a static IP as the IP address, reconfigure the device through the CLI. For more information about configuring interfaces, see **Configure Network Interfaces**.

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

Export Device Data in CSV Format

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

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

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

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

View and Copy Device Configuration

View a Device's Running Configuration

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

To view a device's running configuration:

1. From the Cisco vManage menu, choose **Configuration > Devices**.
2. Click **WAN Edge List** or **Controllers**, and select the device.
3. Click **...**, and click **Running Configuration**.

View a Device's Local Configuration

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

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

1. From the Cisco vManage menu, choose **Configuration > Devices**.
2. Click **WAN Edge List** or **Controllers**, and select the device.
3. Click **...**, and click **Local Configuration**.

Copy Router Configuration

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

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

1. From the Cisco vManage menu, choose **Configuration > Certificates**.
2. Mark the new Cisco vEdge device as invalid.
3. From the Cisco vManage menu, choose **Configuration > Devices**.
4. Under **WAN Edge List**, select the old router.
5. Click **...**, and click **Copy Configuration**.
6. In the **Copy Configuration** window, select the new router.
7. To confirm the copy of the configuration, click **Update**.

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

1. From the Cisco vManage menu, choose **Configuration > Certificates**.
2. Mark the new router as valid.
3. Click **Send to Controller**.

Delete a WAN Edge Router

Delete a router if you need to remove it from your deployment. Doing so removes from the WAN edge router serial number list any of the following items that are stored for the router:

- Chassis number
- Certificate serial number
- Subject SUDI serial number



Note Deleting a router also permanently removes the router configuration from the vManage NMS.

To delete a router:

1. From the Cisco vManage menu, choose **Configuration > Certificates**.
2. Mark the WAN Edge router as invalid.
3. From the Cisco vManage menu, choose **Configuration > Devices**.
4. Click **WAN Edge List**, and select the router.
5. Click **...**, and click **Delete WAN Edge**.
6. To confirm deletion of the device, click **OK**.

7. From the Cisco vManage menu, choose **Configuration > Certificates**.
8. Click **Send to Controller**.

Decommission a Cloud Router

Decommissioning a cloud router (such as a vEdge Cloud router) removes the device's serial number from Cisco vManage and generates a new token for the device. To do so:

1. From the Cisco vManage menu, choose **Configuration > Devices**.
2. Click **WAN Edge List**, and select a cloud router.
3. Click **...**, and click **Decommission WAN Edge**.
4. To confirm the decommissioning of the router, click **OK**.

View Template Log and Device Bringup

View Log of Template Activities

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

To view a log of template activities:

1. From the Cisco vManage menu, choose **Configuration > Devices**.
2. Click **WAN Edge List** or **Controllers**, and select the device.
3. Click **...**, and click **Template Log**.

View Status of Device Bringup

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

To view the status of a device bringup:

1. From the Cisco vManage menu, choose **Configuration > Devices**.
2. Click **WAN Edge List** or **Controllers**, and select the device.
3. Click **...**, and click **Device Bring Up**.

Add a Cisco vBond Orchestrator

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

1. From the Cisco vManage menu, choose **Configuration > Devices**.

2. Click **Controllers**.
3. Click **Add Controller** drop-down list, and select **vBond**.
4. In the Add vBond window:
 - a. Enter **vBond Management IP Address** of the vBond controller.
 - b. Enter the **Username** and **Password** to access the vBond orchestrator.
 - c. To allow the certificate-generation process to occur automatically, check the **Generate CSR** check box.
 - d. Click **Add**.
5. Repeat Steps 2, 3 and 4 to add additional Cisco vBond Orchestrators.

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

Configure Cisco vSmart Controllers

Add a vSmart Controller

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

1. From the Cisco vManage menu, choose **Configuration > Devices**.
2. Click **Controllers**.
3. Click the **Add Controller** drop-down list and select **vSmart**.
4. In the **Add vSmart** window:
 - a. Enter the system IP address of the Cisco vSmart Controller.
 - b. Enter the username and password to access the Cisco vSmart Controller.
 - c. Select the protocol to use for control-plane connections. The default is DTLS. The DTLS (Datagram Transport Layer Security) protocol is designed to provide security for UDP communications.
 - d. If you select TLS, enter the port number to use for TLS connections. The default is 23456.
The TLS (Transport Socket Layer) protocol that provides communications security over a network.
 - e. Check the **Generate CSR** check box to allow the certificate-generation process to occur automatically.
 - f. Click **Add**.
5. Repeat Steps 2, 3 and 4 to add additional Cisco vSmart Controllers. The vManage NMS can support up to 20 Cisco vSmart Controllers in the network.

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

Edit Controller Details

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

1. From the Cisco vManage menu, choose **Configuration > Devices**.
2. Click **Controllers**, and select the controller.
3. Click **...**, and click **Edit**.
4. In the **Edit** window, edit the IP address and the login credentials.
5. Click **Save**.

Delete a Controller

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

To delete a controller:

1. From the Cisco vManage menu, choose **Configuration > Devices**.
2. Click **Controllers**, and select the controller.
3. Click **...**, and click **Invalidate**.
4. To confirm the removal of the device and all its control connections, click **OK**.

Configure Reverse Proxy on Controllers

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

1. From the Cisco vManage menu, choose **Configuration > Devices**.
2. Click **Controllers**, and select the controller.
3. Click **...**, and click **Add Reverse Proxy**.
The **Add Reverse Proxy** dialog box is displayed.
4. Click **Add Reverse Proxy**.
5. Configure the private IP address and port number for the device. The private IP address is the IP address of the transport interface in VPN 0. The default port number is 12346. This is the port used to establish the connections that handle control and traffic in the overlay network.
6. Configure the proxy IP address and port number for the device, to create the mapping between the private and public IP addresses and port numbers.
7. If the Cisco vManage NMS or Cisco vSmart Controller has multiple cores, repeat Steps 5 and 6 for each core.
8. Click **Add**.

To enable reverse proxy in the overlay network, from the Cisco vManage menu, choose **Administration > Settings**. Then from the Reverse Proxy bar, click **Edit**. Click **Enabled**, and click **Save**.



CHAPTER 7

Network Hierarchy and Resource Management

Table 42: Feature History

Feature Name	Release Information	Description
Network Hierarchy and Resource Management	Cisco SD-WAN Release 20.9.1 Cisco vManage Release 20.9.1	This feature enables you to create a network hierarchy in Cisco vManage to represent the geographical locations of your network. The network hierarchy and the associated resource IDs, including region IDs and site IDs, help you apply configuration settings to a device. In addition, the introduction of the resource manager in Cisco vManage automatically manages these resource IDs, thereby simplifying the overall user experience of Cisco SD-WAN. Note that you can create a region only if you enable the Multi-Region Fabric option in Cisco vManage.
Network Hierarchy and Resource Management (Phase II)	Cisco IOS XE Release 17.10.1a Cisco vManage Release 20.10.1	The following enhancements are introduced in the Network Hierarchy and Resource Management feature. <ul style="list-style-type: none"> • Creation of a system IP pool on the Configuration > Network Hierarchy page • Automatic assignment of site ID, system IP, and hostname to a device in the Quick Connect workflow • Display of detailed information on the Configuration > Network Hierarchy page, including site ID pool, region ID pool, and the list of devices associated with a site
Support for Software Defined Remote Access Pools	Cisco IOS XE Release 17.11.1a Cisco vManage Release 20.11.1	Remote access refers to enabling secure access to an organization's network from devices at remote locations. The resource pool manager manages the IPv4 and IPv6 private IP address pools for Cisco SD-WAN remote access devices. You can create a software defined remote access pool using the Configuration > Network Hierarchy page.

- [Information About Network Hierarchy and Resource Management, on page 178](#)

- [Supported Devices for Network Hierarchy and Resource Management, on page 179](#)
- [Restrictions for Network Hierarchy and Resource Management, on page 179](#)
- [Manage a Network Hierarchy, on page 179](#)
- [Assign Resource IDs to Devices, on page 184](#)

Information About Network Hierarchy and Resource Management

Overview of Network Hierarchy

You can create a network hierarchy in Cisco vManage to represent the geographical locations of your network. Your network hierarchy can contain three types of nodes—regions, areas, and sites. The resource IDs assigned to the nodes help you identify where to apply configuration settings later.

By default, there is one node called global in the network hierarchy.

The network hierarchy has a predetermined hierarchy with three types of nodes:

- **Region:** It represents a region in a multiregion fabric-based Cisco SD-WAN deployment. The Multi-Region Fabric feature provides the option to divide the architecture of the Cisco SD-WAN overlay network into multiple regional networks that operate distinctly from one another, and a central core-region network for managing inter-regional traffic.

You can create a region only if you enable the **Multi-Region Fabric** option in Cisco vManage. For complete information about the Multi-Region Fabric feature, see the [Cisco SD-WAN Multi-Region Fabric \(also Hierarchical SD-WAN\) Configuration Guide](#).

- **Area:** An area is a logical grouping of nodes in a network hierarchy. You can group sites, regions, other areas, or any combination of these into an area.
- **Site:** A site is the lowest level of node or the leaf node in a network hierarchy. You cannot create a child node under a site. You can only associate devices to a site.

For complete information about creating and managing different nodes in a network hierarchy, see [Manage a Network Hierarchy](#).

Overview of Resource Management

The resource manager in Cisco vManage manages the resource IDs, that is, region IDs and site IDs. It automatically generates a region ID for a region that you create on the **Configuration > Network Hierarchy** page. Similarly, it generates a site ID for a site if you do not specify it.

You can assign a site ID and a region ID to a device. For complete information about assigning resource IDs to devices, see [Assign Resource IDs to Devices](#).

If you upgrade from an earlier version of Cisco vManage to Cisco vManage Release 20.9.1, the resource manager in Cisco vManage automatically creates sites based on the site IDs of the existing devices in your setup. Sites are named as SITE_<id>. Cisco vManage displays these sites under the global node on the **Network Hierarchy** page. It also associates the existing devices with their sites in the network hierarchy.

Benefits of Network Hierarchy and Resource Management

- Automates the management of regions and sites.
- Saves the manual effort in an upgrade scenario when Cisco vManage discovers all your existing sites and displays them in the network hierarchy.
- Simplifies the onboarding and configuration of devices.

Supported Devices for Network Hierarchy and Resource Management

This feature is supported on Cisco IOS XE SD-WAN devices and Cisco vEdge devices.

Restrictions for Network Hierarchy and Resource Management

- You can delete a node only if it does not have any child node. For example, you can delete a site only if no devices are associated with it.
- A site is the lowest level of a node or the leaf node in a network hierarchy. You cannot create a child node under a site.
- You cannot create more than one region node between the global node and a site node.
- You cannot create a region in a multitenant deployment.

Manage a Network Hierarchy

The Network Hierarchy and Resource Management feature enables you to do the following:

- Create a region
- Create an area
- Create, edit, and delete a site

Create a Region in a Network Hierarchy

Before You Begin

Ensure that the **Multi-Region Fabric** option in Cisco vManage is enabled.

1. From the Cisco vManage menu, choose **Administration > Settings**.
2. Click **Edit** adjacent to the **Multi-Region Fabric** option.
3. Click **Enabled**, and then click **Save**.

Create a Region

1. From the Cisco vManage menu, choose **Configuration > Network Hierarchy**.
2. Click ... adjacent to a node (global or area) in the left pane and choose **Add MRF Region**.



Note In Cisco vManage Release 20.9.x, you can also use the **Add Node** option to add a region.

3. In the **Name** field, enter a name for the region. The name must be unique and can contain only letters, the digits 0 through 9, hyphens (-), underscores (_), and periods (.).
4. In the **Description** field, enter a description of the region.
5. From the **Parent** drop-down list, choose a parent node.
6. Click **Add**.

Create an Area in a Network Hierarchy

1. From the Cisco vManage menu, choose **Configuration > Network Hierarchy**.
2. Click ... adjacent to a node (global, region, or area) in the left pane and choose **Add Area**.



Note In Cisco vManage Release 20.9.x, you can also use the **Add Node** option to add an area.

3. In the **Name** field, enter a name for the area. The name must be unique and can contain only letters, the digits 0 through 9, hyphens (-), underscores (_), and periods (.).
4. In the **Description** field, enter a description of the area.
5. From the **Parent** drop-down list, choose a parent node.
6. Click **Add**.

Create a Site in a Network Hierarchy

1. From the Cisco vManage menu, choose **Configuration > Network Hierarchy**.
2. Click ... adjacent to a node (global, region, or area) in the left pane and choose **Add Site**.



Note In Cisco vManage Release 20.9.x, you can also use the **Add Node** option to add a site.

3. In the **Name** field, enter a name for the site. The name must be unique and can contain only letters, the digits 0 through 9, hyphens (-), underscores (_), and periods (.).
4. In the **Description** field, enter a description of the site.
5. From the **Parent** drop-down list, choose a parent node.

6. In the **Site ID** field, enter a site ID.
If you do not enter the site ID, Cisco vManage generates a site ID for the site.
7. Click **Add**.

Edit a Region

1. From the Cisco vManage menu, choose **Configuration > Network Hierarchy**.
2. Click ... adjacent to the region name and choose **Edit MRF Region**.
3. Edit the options as needed. You can edit the name, description, and parent of the region.
4. Click **Save**.

Delete a Region

1. From the Cisco vManage menu, choose **Configuration > Network Hierarchy**.
2. Click ... adjacent to the region name and choose **Delete MRF Region**.
3. In the confirmation dialog box, click **Yes**.

Edit an Area

1. From the Cisco vManage menu, choose **Configuration > Network Hierarchy**.
2. Click ... adjacent to the area name and choose **Edit Area**.
3. Edit the options as needed. You can edit the name, description, and parent of the area.
4. Click **Save**.

Delete an Area

1. From the Cisco vManage menu, choose **Configuration > Network Hierarchy**.
2. Click ... adjacent to the area name and choose **Delete Area**.
3. In the confirmation dialog box, click **Yes**.

Edit a Site

1. From the Cisco vManage menu, choose **Configuration > Network Hierarchy**.
2. Click ... adjacent to the site name and choose **Edit Site**.
3. Edit the options as needed. You can edit only the name, description, and parent of the site.
4. Click **Save**.

Delete a Site

1. From the Cisco vManage menu, choose **Configuration > Network Hierarchy**.
2. Click ... adjacent to the site name and choose **Delete Site**.
3. In the confirmation dialog box, click **Yes**.

Create a System IP Pool

Minimum releases: Cisco IOS XE Release 17.10.1a, Cisco vManage Release 20.10.1

1. From the Cisco vManage menu, choose **Configuration > Network Hierarchy**.
The page displays the site pool and region pool for the Global node.
2. Click **Add**.
3. In the **Pool Name** field, enter a name for the pool.
4. In the **Pool Description** field, enter a description of the pool.
5. From the **Pool Type** drop-down list, choose **System IP**.
6. In the **IP Subnet*** field, enter an IP address.
7. In the **Prefix Length*** field, enter the prefix length of the system IP pool.
8. Click **Add**.



Note You can create only one system IP pool. If you want to make any changes to the pool, you must edit the existing pool.

Edit a System IP Pool

Minimum releases: Cisco IOS XE Release 17.10.1a, Cisco vManage Release 20.10.1

1. From the Cisco vManage menu, choose **Configuration > Network Hierarchy**.
The page displays the site pool and region pool for the Global node. The system IP pool is also displayed if you have already created it.
2. Click ... adjacent to the system IP name and choose **Edit**.
3. Edit the options as needed.



Note You can only expand the pool range and cannot enter a lower IP address than the already specified IP address.

4. Click **Save**.

Create a Remote Access Pool

Minimum supported release: Cisco vManage Release 20.11.1

The resource pool manager supports creation of IPv4 and IPv6 private IP pools for Cisco SD-WAN remote access devices. In the remote access configuration you can select the remote access private IP Pool by defining the number of IP addresses.

For more information on Software Defined Remote Access, see [Cisco SD-WAN Remote Access](#).

1. From the Cisco vManage menu, choose **Configuration > Network Hierarchy**.
The page displays the site pool and region pool for the Global node.
2. Click **Add Pool**.
3. In the **Pool Name** field, enter a name for the pool.
4. In the **Pool Description** field, enter a description of the pool.
5. From the **Pool Type** drop-down list, choose **Remote Access**.
6. Choose the **IP Type** by clicking the radio button next to **IPv4** or **IPv6**.
7. In the **IP Subnet** field, enter an IP subnet.
8. In the **Prefix Length** field, enter the prefix length of the remote access pool.
9. Click **Add**.

Edit a Remote Access Pool

Minimum supported release: Cisco vManage Release 20.11.1

You can edit a remote access pool only when you want to expand the pool range.

1. From the Cisco vManage menu, choose **Configuration > Network Hierarchy**.
The page displays the site pool and region pool for the Global node. The remote access pool is also displayed if you have already created it.
2. Click ... adjacent to the remote access pool name and choose **Edit**.
3. Edit the options as needed.



Note When you edit a remote access pool, the new pool range cannot be less than the existing pool range

4. Click **Save**.

Delete a Pool

Minimum supported release: Cisco vManage Release 20.11.1

1. From the Cisco vManage menu, choose **Configuration > Network Hierarchy**.

2. In the Global page, click ... adjacent to the pool name and choose **Delete**.
3. In the confirmation dialog box, click **Yes**.



Note You can delete a pool only when the pool resources are not in use.

Assign Resource IDs to Devices

The Network Hierarchy and Resource Management feature enables you to do the following:

- Assign a site ID to a device
- Assign a region ID to a device

Assign a Site ID to a Device

You can assign a site ID to a device using one of the following ways.

Use the Quick Connect Workflow

1. From the Cisco vManage menu, choose **Workflows > Workflow Library**.
2. Start the **Quick Connect** workflow.
3. Follow the instructions provided in the workflow.
4. On the **Add and Review Device Configuration** page, enter the site ID of the device.



-
- Note**
- You can use any of the existing site IDs that are available in the network hierarchy or enter a new site ID. If you enter a new site ID without creating a node in the network hierarchy, the site is automatically created and listed on the **Configuration > Network Hierarchy** page.
 - (Minimum releases: Cisco IOS XE Release 17.10.1a, Cisco vManage Release 20.10.1) If you want Cisco vManage to automatically generate a site ID for the device, do not make any change to the default value, **AUTO**.
-

Use a Template

1. From the Cisco vManage menu, choose **Configuration > Devices > WAN Edge List**.
2. Check if a device is attached to a device template.
3. From the Cisco vManage menu, choose **Configuration > Templates > Feature Templates**.
4. Click ... adjacent to the System feature template and choose **Edit**.
5. Click the **Basic Configuration** tab and set the scope of the **Site ID** field to **Global** and enter the site ID.

6. Click **Update**.
7. Click **Configure Devices** to push the configuration to the device.

In Step 5, if you set the scope of the **Site ID** field to **Device Specific**, do the following:

1. From the Cisco vManage menu, choose **Configuration > Templates > Device Templates**.
2. Click ... adjacent to the device template and choose **Edit Device Template**.
3. In the **Site ID** field, enter the site ID.

You can use any of the existing site IDs that are available in the network hierarchy or enter a new site ID. If you enter a new site ID without creating a node in the network hierarchy, the site is automatically created and listed on the **Configuration > Network Hierarchy** page.

4. Click **Update**.
5. Click **Configure Devices** to push the configuration to the device.

Use a Configuration Group

The configuration group flow is applicable only for the Cisco IOS XE SD-WAN devices.

1. From the Cisco vManage menu, choose **Configuration > Templates > Configuration Groups**.
2. Click ... adjacent to the configuration group name and choose **Edit**.
3. Click **Associated Devices**.
4. Choose a device that is associated with the configuration group and click **Deploy**.

The **Deploy Configuration Group** workflow starts.

5. Follow the instructions provided in the workflow.
6. On the **Add and Review Device Configuration** page, enter the site ID of the device.

You can use any of the existing site IDs that are available in the network hierarchy or enter a new site ID. If you enter a new site ID without creating a node in the network hierarchy, the site is automatically created and listed on the **Configuration > Network Hierarchy** page.

Assign a Region ID to a Device

Before You Begin

- Have access to the **Multi-Region Fabric** feature.
- Ensure that the region is available in the network hierarchy.

Assign a Region ID

1. From the Cisco vManage menu, choose **Configuration > Devices > WAN Edge List**.
2. Check if the corresponding device is attached to a device template.
3. From the Cisco vManage menu, choose **Configuration > Templates > Feature Templates**.

4. Click ... adjacent to the System feature template and choose **Edit**.
5. Click the **Basic Configuration** tab and set the scope of the **Region ID** field to **Global** and enter the region ID.

You can use any of the existing region IDs that are available in the network hierarchy. If the specified region ID is not available in the network hierarchy, the template push operation to the devices fails.

6. Click **Update**.
7. Click **Configure Devices** to push the configuration to the device.

In Step 5, if you set the scope of the **Region ID** field to **Device Specific**, do the following:

1. From the Cisco vManage menu, choose **Configuration > Templates > Device Templates**.
2. Click ... adjacent to the device template and choose **Edit Device Template**.
3. In the **Region ID** field, enter the region ID.
4. Click **Update**.
5. Click **Configure Devices** to push the configuration to the device.

Assign a System IP to a Device

Minimum releases: Cisco IOS XE Release 17.10.1a, Cisco vManage Release 20.10.1

1. From the Cisco vManage menu, choose **Workflows > Workflow Library**.
2. Start the **Quick Connect** workflow.
3. Follow the instructions provided in the workflow.
4. On the **Add and Review Device Configuration** page, enter the system IP of the device. If you want Cisco vManage to automatically generate a system IP for the device, do not make any change to the default value, **AUTO**.

Assign a Hostname to a Device

Minimum releases: Cisco IOS XE Release 17.10.1a, Cisco vManage Release 20.10.1

1. From the Cisco vManage menu, choose **Workflows > Workflow Library**.
2. Start the **Quick Connect** workflow.
3. Follow the instructions provided in the workflow.
4. On the **Add and Review Device Configuration** page, enter the hostname of the device. If you want Cisco vManage to automatically generate a hostname for the device, do not make any change to the default value, **AUTO**.



CHAPTER 8

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 188](#)
- [Configure Interfaces in the WAN Transport VPN \(VPN 0\), on page 191](#)
- [Extend the WAN Transport VPN, on page 195](#)
- [Configure GRE Interfaces and Advertise Services to Them, on page 198](#)
- [Configure the System Interface, on page 202](#)
- [Configure Control Plane High Availability, on page 203](#)
- [Configure Other Interfaces, on page 203](#)
- [Configure Interface Properties, on page 205](#)
- [Enable DHCP Server using Cisco vManage, on page 207](#)
- [Configuring PPPoE, on page 212](#)
- [Configure PPPoE Over ATM, on page 218](#)
- [Configuring VRRP , on page 220](#)
- [Network Interface Configuration Examples for Cisco vEdge Devices, on page 225](#)
- [Configure VPN Ethernet Interface, on page 240](#)

- [VPN Interface Bridge, on page 256](#)
- [VPN Interface Ethernet PPPoE, on page 262](#)
- [VPN Interface GRE, on page 271](#)
- [VPN Interface IPsec \(for Cisco vEdge Devices\), on page 274](#)
- [VPN Interface PPP, on page 279](#)
- [VPN Interface PPP Ethernet, on page 287](#)
- [Cellular Interfaces, on page 292](#)
- [WiFi Radio, on page 304](#)
- [WiFi SSID, on page 306](#)
- [Interface CLI Reference, on page 308](#)

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.

Create a VPN Template



Note You can configure a static route through the VPN template.

Step 1 From the Cisco vManage menu, choose **Configuration > Templates**.

Step 2 Click **Device Templates**, and click **Create Template**.

Note In Cisco vManage Release 20.7.x and earlier releases **Device Templates** is called **Device**.

Step 3 From the **Create Template** drop-down list, choose **From Feature Template**.

Step 4 From the **Device Model** drop-down list, choose the type of device for which you wish to create the template.

Step 5 To create a template for VPN 0 or VPN 512:

- a. Click **Transport & Management VPN**, or scroll to the **Transport & Management VPN** section.
- b. From the VPN 0 or VPN 512 drop-down list, click **Create Template**. The VPN template form appears.
The form contains fields for naming the template, and fields for defining VPN parameters.


Step 6 To create a template for VPNs 1 through 511, and 513 through 65530:



- a. Click **Service VPN**, or scroll to the **Service VPN** section.
- b. Click the **Service VPN** drop-down list.
- c. From the **VPN** drop-down list, click **Create Template**. The VPN template form displays.
The form contains fields for naming the template, and fields for defining VPN parameters.

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

Step 8 In **Template Description**, 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 list and select one of the following:

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

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

Configure Basic VPN Parameters

To configure basic VPN parameters, choose **Basic Configuration** 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	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 **DNS** and configure the following parameters:

Parameter Name	Options	Description
Primary DNS Address		Click 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 the Mark as Optional Row check box to mark this configuration as device-specific. To include this configuration for a device, enter the requested variable values when you attach a device template to a device, or create a template variables spreadsheet to apply the variables.
	Hostname	Enter the hostname of the DNS server. The name can be up to 128 characters.
	List of IP Addresses	Enter up to eight IP addresses to associate with the hostname. Separate the entries with commas.
To save the DNS server configuration, click Add .		

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

CLI Equivalent

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

Configure Interfaces in the WAN Transport VPN (VPN 0)

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

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

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



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

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

Tunnel interfaces on Cisco vEdge devices must have an IP address, a color, and an encapsulation type. The IP address can be either an IPv4 or IPv6 address. To enable dual stack in releases before Cisco SD-WAN Release 20.3.2, configure both address types.

To use dual stack with Cisco vEdge devices from Cisco SD-WAN Release 20.3.2, configure all controllers with both IPv4 and IPv6 addresses. In addition, configure DNS for the Cisco vBond Orchestrator interface to resolve IPv4 and IPv6 address types so that controllers can reach the Cisco vBond Orchestrator through either IP address type.



Note Starting from Cisco vManage Release 20.6.1, in case of a dual-stack configuration, if an IPv4 address or the fully qualified domain name (FQDN) is not available, but an IPv6 address is available, then the IPv6 address is used to connect to the Cisco vBond Orchestrator.

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.

From Cisco SD-WAN Release 20.3.2, Cisco vEdge devices do not support dual stack on the same TLOC or interface. Only one address type can be provisioned for a TLOC or interface. Using a second address type requires a second TLOC or interface on which it can be provisioned.

```
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.

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 of 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	—	—

Service	Cisco vEdge device	Cisco vSmart Controller	Cisco vSmart Controller
dns	X	—	—
https	—	X	—
icmp	X	X	X
netconf	—	X	—
ntp	X	—	—
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 explicit ACLs are handled, see Configuring Localized Data Policy for IPv4 or Configuring Localized Data Policy for IPv6 .

For each transport tunnel on a vEdge router and for each encapsulation type on a single transport tunnel, the Cisco SD-WAN software creates a TLOC, which consists of the router' system IP address, the color, and the encapsulation. The OMP session running on the tunnel sends the TLOC, as a TLOC route, to the Cisco vSmart

Controller, which uses it to determine the overlay network topology and to determine the best paths for data traffic across the overlay network.

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

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

VPN	INTERFACE	IP ADDRESS	IF ADMIN STATUS	IF OPER STATUS	ENCAP TYPE	PORT TYPE	MTU	HWADDR	SPEED MBPS	DUPLEX	TCP MSS ADJUST	UPTIME	RX PACKETS	TX PACKETS
0	ge0/1	10.0.5.21/24	Up	Up	null	transport	1500	00:0c:29:6c:30:c1	10	full	0	0:04:03:41	260025	260145
0	ge0/2	-	Down	Up	null	service	1500	00:0c:29:6c:30:cb	10	full	0	0:04:03:41	3506	1
0	ge0/3	-	Down	Up	null	service	1500	00:0c:29:6c:30:d5	10	full	0	0:04:03:41	260	1
0	ge0/4	-	Down	Up	null	service	1500	00:0c:29:6c:30:df	10	full	0	0:04:03:41	260	1
0	ge0/5	-	Down	Up	null	service	1500	00:0c:29:6c:30:e9	10	full	0	0:04:03:41	260	1
0	ge0/6	10.0.7.21/24	Up	Up	null	service	1500	00:0c:29:6c:30:f3	10	full	0	0:04:03:41	265	2
0	ge0/7	10.0.100.21/24	Up	Up	null	service	1500	00:0c:29:6c:30:fd	10	full	0	0:04:03:41	278	2
0	system	172.16.255.21/32	Up	Up	null	loopback	1500	00:00:00:00:00:00	10	full	0	0:04:03:37	0	0

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

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

VPN	INTERFACE	AF	TYPE	IPV6 ADDRESS	IF ADMIN STATUS	IF OPER STATUS	ENCAP TYPE	PORT TYPE	MTU	HWADDR	SPEED MBPS	DUPLEX	TCP MSS ADJUST	UPTIME	RX PACKETS	TX 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 device-1 so that Cisco vEdge device-1 can participate independently in the overlay network. You can also make a similar configuration for vEdge-1 so that the WAN transport can be extended from Cisco vEdge device-1 to Cisco vEdge device-2.



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 TLOC only on a physical interface.

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

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

```
vEdge-1# show running-config vpn 0
interface ge0/2.101
  ip address 192.168.19.15/24
  mtu 1496
  tunnel-interface
    color lte
  ...
!
no shutdown
!
interface ge0/2.102
  ip address 192.168.20.15/24
  mtu 1496
  tunnel-interface
    color mpls
  ...
!
no shutdown
!
ip route 0.0.0.0/0 192.168.19.16
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 192.168.19.16/24
  mtu 1496
  tloc-extension ge0/0
  no shutdown
!
interface ge0/2.102
  ip address 192.168.20.16/24
```

Configure GRE Interfaces and Advertise Services to Them

```

mtu 1496
tloc-extension ge0/3
no shutdown
!

```

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

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

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

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

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

SYSTEM IP	SITE ID	STATE	SOURCE TLOC COLOR	REMOTE TLOC COLOR	SOURCE IP	DST PUBLIC IP	DST PUBLIC PORT	ENCAP	DETECT MULTIPLIER	TX INTERVAL(msec)	UPTIME	TRANSI-TIONS
172.16.255.11	100	up	lte	lte	192.168.19.15	10.0.101.1	12346	ipsec	20	1000	0:00:20:26	0
172.16.255.11	100	up	lte	3g	192.168.19.15	10.0.101.2	12346	ipsec	20	1000	0:00:20:26	0
172.16.255.11	100	up	lte	gold	192.168.19.15	10.0.101.3	12346	ipsec	20	1000	0:00:20:26	0
172.16.255.11	100	up	lte	red	192.168.19.15	10.0.101.4	12346	ipsec	20	1000	0:00:20:26	0
172.16.255.11	100	up	mpls	lte	192.168.20.15	10.0.101.1	12346	ipsec	20	1000	0:00:20:26	0
172.16.255.11	100	up	mpls	3g	192.168.20.15	10.0.101.2	12346	ipsec	20	1000	0:00:20:26	0
172.16.255.11	100	up	mpls	gold	192.168.20.15	10.0.101.3	12346	ipsec	20	1000	0:00:20:26	0
172.16.255.11	100	up	mpls	red	192.168.20.15	10.0.101.4	12346	ipsec	20	1000	0:00:20:26	0
172.16.255.14	400	up	lte	lte	192.168.19.15	10.1.14.14	12360	ipsec	20	1000	0:00:20:26	0
172.16.255.14	400	up	mpls	lte	192.168.20.15	10.1.14.14	12360	ipsec	20	1000	0:00:20:26	0
172.16.255.21	100	up	lte	lte	192.168.19.15	10.0.111.1	12346	ipsec	20	1000	0:00:20:26	0
172.16.255.21	100	up	lte	3g	192.168.19.15	10.0.111.2	12346	ipsec	20	1000	0:00:20:26	0
172.16.255.21	100	up	mpls	lte	192.168.20.15	10.0.111.1	12346	ipsec	20	1000	0:00:20:26	0
172.16.255.21	100	up	mpls	3g	192.168.20.15	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.

Create a GRE tunnel by configuring a GRE interface. GRE interfaces are logical interfaces, and you configure them just like any other physical interface. A GRE interface is a logical interface, you must bind it to a physical interface or a PPPoE 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 or PPPoE interface (in the **tunnel-source** command) or the name of the physical interface or PPPoE 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, see 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 a reachable GRE interface, you cannot delete the GRE interface from the configuration. To delete the GRE interface, you must first delete the service. You can, however, reconfigure the service itself, by modifying the **service** command.

Then, create a data policy on the Cisco vSmart Controller that applies to the service VPN. In the action portion of the data policy, you must explicitly configure the policy to service the packets destined for the GRE tunnel. To do this, include the **local** option in the **set service** command:

```
vSmart(config-policy-data-policy-vpn-list-vpn-sequence)# action accept
vSmart(config-action-accept)# set service service-name local
```


If the GRE tunnel used to reach the service is down, packet routing falls back to using standard routing. To drop packets when a GRE tunnel to the service is unreachable, add the **restrict** option:

```
vSmart(config-policy-data-policy-vpn-list-vpn-sequence)# action accept
vSmart(config-action-accept)# set service service-name local restrict
```

To monitor GRE tunnels and their traffic, use the following commands:

- **show interface** —List data traffic transmitted and received on GRE tunnels.
- **show tunnel gre-keepalives** —List GRE keepalive traffic transmitted and received on GRE tunnels.
- **show tunnel statistics** —List both data and keepalive traffic transmitted and received on GRE tunnels.

The following figure illustrates an example of configuring a GRE tunnel in VPN 0, to allow traffic to be redirected to a service that is not located at the same site as the vEdge router. In this example, local traffic is directed to the GRE tunnel using a centralized data policy, which is configured on the Cisco vSmart Controller.



The configuration looks like this:

```
vEdge# show running-config vpn 0
vpn 0
  interface gre1
    ip address 172.16.111.11/24
    keepalive 60 10
    tunnel-source 172.16.255.11
    tunnel-destination 10.1.2.27
    no shutdown
  !
!
vEdge# show running-config vpn 1 service
vpn 1
  service FW interface gre1
```

```
vSmart# show running-config policy
policy
  lists
    prefix-list for-firewall
      ip-prefix 172.16.1.0/24
    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 172.16.1.0/24 vpn 0 interface gre1
```

The **show interface** command displays the GRE interface in VPN 0:

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

VPN	INTERFACE	IP ADDRESS	IF ADMIN STATUS	IF OPER STATUS	ENCAP TYPE	PORT	TYPE	MTU	HWADDR	SPEED MBPS	DUPLEX	TCP MSS ADJUST	UPTIME	RX PACKETS	TX PACKETS
0	gre1	172.16.111.11/24	Up	Down	null	service		1500	0a:00:05:0b:00:00	-	-	1420	-	0	0
0	ge0/1	10.0.26.11/24	Up	Up	null	service		1500	00:0c:29:ab:b7:62	10	full	1420	0:03:35:14	89	5
0	ge0/2	10.0.5.11/24	Up	Up	null	transport		1500	00:0c:29:ab:b7:6c	10	full	1420	0:03:35:14	9353	18563
0	ge0/3	-	Down	Up	null	service		1500	00:0c:29:ab:b7:76	10	full	1420	0:03:57:52	99	0
0	ge0/4	10.0.7.11/24	Up	Up	null	service		1500	00:0c:29:ab:b7:80	10	full	1420	0:03:35:14	89	5
0	ge0/5	-	Down	Up	null	service		1500	00:0c:29:ab:b7:8a	10	full	1420	0:03:57:52	97	0
0	ge0/6	-	Down	Up	null	service		1500	00:0c:29:ab:b7:94	10	full	1420	0:03:57:52	85	0
0	ge0/7	10.0.100.11/24	Up	Up	null	service		1500	00:0c:29:ab:b7:9e	10	full	1420	0:03:56:30	3146	2402
0	system	172.16.255.11/32	Up	Up	null	loopback		1500	00:00:00:00:00:00	10	full	1420	0:03:34:15	0	0

You can also view the GRE tunnel information:

```
vEdge# show tunnel gre-keepalives
```

VPN	IF NAME	SOURCE IP	DEST IP	ADMIN STATE	OPER STATE	KA ENABLED	REMOTE TX PACKETS	REMOTE RX PACKETS	TX PACKETS	RX PACKETS	TX ERRORS	RX ERRORS	TRANSITIONS
0	gre1	10.0.5.11	10.1.2.27	up	down	true	0	0	442	0	0	0	0

```
vEdge# show tunnel statistics
```

```
tunnel statistics gre 10.0.5.11 10.1.2.27 0 0
tunnel-mtu 1460
tx_pkts 451
tx_octets 54120
rx_pkts 0
rx_octets 0
tcp-mss-adjust 1380
```

Configure the System Interface

For each Cisco vEdge device, you configure a system interface with the **system system-ip** command. The system interface's IP address is a persistent address that identifies the Cisco vEdge device. It is similar to a router ID on a regular router, which is the address used to identify the router from which packets originated.

```
vEdge(config)# system system-ip ipv4-address
```

Specify the system IP address as an IPv4 address in decimal four-part dotted notation. Specify just the address; the prefix length (/32) is implicit.

The system IP address can be any IPv4 address except for 0.0.0.0/8, 127.0.0.0/8, and 224.0.0.0/4, and 240.0.0.0/4 and later. Each device in the overlay network must have a unique system IP address. You cannot use this same address for another interface in VPN 0.

The system interface is placed in VPN 0, as a loopback interface named **system**. Note that this is not the same as a loopback address that you configure for an interface.

To display information about the system interface, use the **show interface** command. For example:

```
vEdge# show running-config system system-ip
system
 system-ip 172.16.255.11
!
vEdge# show interface vpn 0
```

VPN	INTERFACE	IP ADDRESS	IF ADMIN STATUS	IF OPER STATUS	ENCAP TYPE	PORT TYPE	MTU	HWADDR	SPEED MBPS	DUPLEX	TCP MSS ADJUST	UPTIME	RX PACKETS	TX PACKETS
0	ge0/1	10.0.26.11/24	Up	Up	null	service	1500	00:0c:29:ab:b7:62	1000	full	1420	0:10:32:16	1606	8
0	ge0/2	10.0.5.11/24	Up	Up	null	transport	1500	00:0c:29:ab:b7:6c	1000	full	1420	0:10:32:16	307113	303457
0	ge0/3	-	Down	Up	null	service	1500	00:0c:29:ab:b7:76	1000	full	1420	0:10:47:49	1608	0
0	ge0/4	10.0.7.11/24	Up	Up	null	service	1500	00:0c:29:ab:b7:80	1000	full	1420	0:10:32:16	1612	8
0	ge0/5	-	Down	Up	null	service	1500	00:0c:29:ab:b7:8a	1000	full	1420	0:10:47:49	1621	0
0	ge0/6	-	Down	Up	null	service	1500	00:0c:29:ab:b7:94	1000	full	1420	0:10:47:49	1600	0
0	ge0/7	10.0.100.11/24	Up	Up	null	service	1500	00:0c:29:ab:b7:9e	1000	full	1420	0:10:47:31	3128	1165
0	system	172.16.255.11/32	Up	Up	null	loopback	1500	00:00:00:00:00:00	10	full	1420	0:10:31:58	0	0

The system IP address is used as one of the attributes of the OMP TLOC. Each TLOC is uniquely identified by a 3-tuple comprising the system IP address, a color, and an encapsulation. To display TLOC information, use the **show omp tlocs** command.

For device management purposes, it is recommended as a best practice that you also configure the same system IP address on a loopback interface that is located in a service-side VPN that is an appropriate VPN for management purposes. You use a loopback interface because it is always reachable when the router is operational and when the overlay network is up. If you were to configure the system IP address on a physical interface, both the router and the interface would have to be up for the router to be reachable. You use a service-side VPN because it is reachable from the data center. Service-side VPNs are VPNs other than VPN 0 (the WAN transport VPN) and VPN 512 (the management VPN), and they are used to route data traffic.

Here is an example of configuring the system IP address on a loopback interface in VPN 1:

```
vEdge# config
Entering configuration mode terminal
vEdge(config)# vpn 1
vEdge(config-vpn-1)# interface loopback0 ip address 172.16.255.11/32
vEdge(config-vpn-1)# no shutdown
vEdge(config-interface-loopback0)# commit and-quit
Commit complete.
vEdge# show interface
```

VPN	INTERFACE	IP ADDRESS	IF ADMIN STATUS	IF OPER STATUS	ENCAP TYPE	PORT TYPE	MTU	HWADDR	SPEED MBPS	DUPLEX	TCP MSS ADJUST	UPTIME	RX PACKETS	TX PACKETS
0	ge0/1	10.0.26.11/24	Up	Up	null	service	1500	00:0c:29:ab:b7:62	1000	full	1420	0:10:27:33	1597	8
0	ge0/2	10.0.5.11/24	Up	Up	null	transport	1500	00:0c:29:ab:b7:6c	1000	full	1420	0:10:27:33	304819	301173
0	ge0/3	-	Down	Up	null	service	1500	00:0c:29:ab:b7:76	1000	full	1420	0:10:43:07	1599	0
0	ge0/4	10.0.7.11/24	Up	Up	null	service	1500	00:0c:29:ab:b7:80	1000	full	1420	0:10:27:33	1603	8


```

0   ge0/5   -           Down   Up     null  service  1500  00:0c:29:ab:b7:8a  1000  full  1420  0:10:43:07  1612  0
0   ge0/6   -           Down   Up     null  service  1500  00:0c:29:ab:b7:94  1000  full  1420  0:10:43:07  1591  0
0   ge0/7   10.0.100.11/24  Up     Up     null  service  1500  00:0c:29:ab:b7:9e  1000  full  1420  0:10:42:48  3118  1164
0   system  172.16.255.11/32  Up     Up     null  loopback  1500  00:00:00:00:00:00  10    full  1420  0:10:27:15  0      0
1   ge0/0   10.2.2.11/24    Up     Up     null  service  1500  00:0c:29:ab:b7:58  1000  full  1420  0:10:27:30  5734  4204
1   loopback0  172.16.255.11/32  Up     Up     null  service  1500  00:00:00:00:00:00  10    full  1420  0:00:00:28  0      0
512 eth0     10.0.1.11/24    Up     Up     null  service  1500  00:50:56:00:01:0b  1000  full  0      0:10:43:03  20801  14368

```

Configure Control Plane High Availability

A highly available Cisco SD-WAN network contains two or more Cisco vSmart Controllers in each domain. A Cisco SD-WAN domain can have up to eight Cisco vSmart Controllers, and each Cisco vEdge device, by default, connects to two of them. You change this value on a per-tunnel basis:

```
vEdge(config-tunnel-interface)# max-controllers number
```

When the number of Cisco vSmart Controllers in a domain is greater than the maximum number of controllers that a domain's Cisco vEdge devices are allowed to connect to, the Cisco SD-WAN software load-balances the connections among the available Cisco vSmart Controllers.

Configure Other Interfaces

Configure Interfaces in the Management (VPN 512)

On all Cisco SD-WAN devices, VPN 512 is used for out-of-band management, by default as part of the factory-default configuration. On Cisco vEdge devices the interface type for management interfaces is **mgmt**, and the initial address for the interface is 192.168.1.1.

```

vEdge# show running-config vpn 512
vpn 512
 interface mgmt0
   ip dhcp-client
   no shutdown
 !
 !

```

To display information about the configured management interfaces, use the **show interface** command. For example:

```

vEdge# show interface vpn 512

```

VPN	INTERFACE	IP ADDRESS	IF ADMIN STATUS	IF OPER STATUS	ENCAP TYPE	PORT TYPE	MTU	HWADDR	SPEED MBPS	DUPLEX	TCP MSS ADJUST	UPTIME	RX PACKETS	TX PACKETS
512	mgmt0	192.168.1.1/24	Up	Up	null	service	1500	00:50:56:00:01:1f	1000	full	0	0:04:08:01	1131	608



Note VPN 512 is not advertised in the overlay. It is local to the device. If you need a management VPN that is reachable through the overlay, create a VPN with a number other than 512.

Configure Service-Side Interfaces for Carrying Data Traffic

On Cisco vEdge devices, the VPNs other than 0 and 512 are service-side VPNs, and the interfaces in these VPNs connect the router to service-side LANs and WLANs. These interfaces are the interfaces that carry data traffic between vEdge routers and sites across the overlay network. At a minimum, for these interfaces, you must configure an IPv4 address, and you must enable the interface:

```
vEdge(config)# vpn vpn-id
vEdge(config-vpn)# interface ge slot / port
vEdge(config-interface)# ip address prefix/length
vEdge(config-interface)# no shutdown
```

For service-side interfaces, you can configure up to four secondary IP addresses.

```
vEdge(config)# vpn vpn-id
vEdge(config-vpn)# interface ge slot/port
vEdge(config-interface)# ip secondary-address ipv4-address
```

To display information about the configured data traffic interfaces, use the **show interface** command.

```
vEdge# show interface vpn 1
```

VPN	INTERFACE	IP ADDRESS	IF ADMIN STATUS	IF OPER STATUS	ENCAP TYPE	PORT TYPE	MTU	HWADDR	SPEED MBPS	DUPLEX	TCP MSS ADJUST	UPTIME	RX PACKETS	TX PACKETS
1	ge0/1	10.192.1.1/28	Up	Up	null	service	1500	00:0c:bd:05:f0:84	100	full	0	1:05:44:07	399	331
1	loopback1	10.255.1.1/32	Up	Up	null	service	1500	00:00:00:00:00:00	10	full	0	1:05:44:07	0	0

For some protocols, you specify an interface as part of the protocol's configuration. In these cases, the interface used by the protocol must be the same as one of the interfaces configured in the VPN. As example is OSPF, where you place interfaces in OSPF areas. In this example, the interface **ge0/0** is configured in VPN 1, and this interface is configured to be in the OSPF backbone area:

```
vEdge# show running-config vpn 1
vpn 1
router
ospf
router-id 172.16.255.21
timers spf 200 1000 10000
redistribute static
redistribute omp
area 0
interface ge0/0
exit
exit
!
!
interface ge0/0
ip address 10.2.3.21/24
no shutdown
!
!
```

Configure Loopback Interfaces

Use the interface name format **loopback string**, where *string* can be any alphanumeric value and can include underscores (_) and hyphens (-). The total interface name, including the string "loopback", can be a maximum of 16 characters long. (Note that because of the flexibility of interface naming in the CLI, the interfaces **lo0** and **loopback0** are parsed as different strings and as such are not interchangeable. For the CLI to recognize as interface as a loopback interface, its name must start with the full string **loopback**.)

One special use of loopback interfaces is to configure data traffic exchange across private WANs, such as MPLS or metro Ethernet networks. To allow a router that is behind a private network to communicate directly over the private WAN with other edge routers, you direct data traffic to a loopback interface that is configured as a tunnel interface rather than to an actual physical WAN interface.

Configure Interface Properties

Set the Interface Speed

When a Cisco vEdge device comes up, the Cisco SD-WAN software autodetects the SFPs present in the router and sets the interface speed accordingly. The software then negotiates the interface speed with the device at the remote end of the connection to establish the actual speed of the interface. To display the hardware present in the router, use the **show hardware inventory** command:

```
vEdge# show hardware inventory
```

HW TYPE	HW DEV INDEX	VERSION	PART NUMBER	SERIAL NUMBER	DESCRIPTION
Chassis	0	3.1	vEdge-1000	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
```

VPN	INTERFACE	IP ADDRESS	IF ADMIN STATUS	IF OPER STATUS	ENCAP TYPE	PORT TYPE	MTU	HWADDR	SPEED MBPS	DUPLEX	TCP MSS ADJUST	UPTIME	RX PACKETS	TX PACKETS
0	ge0/0	192.168.1.4/24	Up	Up	null	transport	1500	00:0c:bd:05:f0:83	1000	full	1300	0:06:10:59	2176305	2168760
0	ge0/2	-	Down	Down	null	service	1500	00:0c:bd:05:f0:81	-	-	0	-	0	0
0	ge0/3	-	Down	Down	null	service	1500	00:0c:bd:05:f0:82	-	-	0	-	0	0
0	ge0/4	-	Down	Down	null	service	1500	00:0c:bd:05:f0:87	-	-	0	-	0	0
0	ge0/5	-	Down	Down	null	service	1500	00:0c:bd:05:f0:88	-	-	0	-	0	0
0	ge0/6	-	Down	Down	null	service	1500	00:0c:bd:05:f0:85	-	-	0	-	0	0
0	ge0/7	-	Down	Down	null	service	1500	00:0c:bd:05:f0:86	-	-	0	-	0	0
0	system	10.255.1.1/32	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	10.255.1.1/32	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 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 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.



Note If you set an MTU on Cisco vEdge hardware device, when a packet whose size is larger than the MTU is received, the vEdge interface drops the packet. This is true, if the "Do Not Fragment" bit is set or not. However, this behavior is not true for vEdge Cloud devices.



Note From Cisco SD-WAN release 20.5 and later releases, PMTU discovery on Cisco vEdge devices is enabled for asymmetric networks. PMTU is calculated based on the egress path MTU.

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. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Device Templates**, and then click **Create Template**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. From the **Create Template** drop-down list, choose **From Feature Template**.
4. From the **Device Model** drop-down list, select the type of device for which you are creating the template.

5. Click **Service VPN** or scroll to the **Service VPN** section.
6. Click **Service VPN** drop-down list.
7. From **Additional VPN Templates**, click **VPN Interface**.
8. From the **Sub-Templates** drop-down list, choose **DHCP Server**.
9. From the **DHCP Server** drop-down list, click **Create Template**. The DHCP-Server template form is displayed.

This form contains fields for naming the template, and fields for defining the DHCP Server parameters.

10. In **Template Name**, enter a name for the template.
The name can be up to 128 characters and can contain only alphanumeric characters.
11. In **Template Description**, 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 list.

Minimum DHCP Server Configuration

To configure DHCP server functionality, select **Basic Configuration** and configure the following parameters. Parameters marked with an asterisk as required to configure DHCP servers.

Table 43:

Parameter Name	Description
Address Pool*	Enter the IPv4 prefix range, in the format <i>prefix/length</i> , for the pool of addresses in the service-side network for which the router interface acts as DHCP server.
Exclude Addresses	Enter one or more IP addresses to exclude from the DHCP address pool. To specify multiple individual addresses, list them separated by a comma. To specify a range of addresses, separate them with a hyphen.
Maximum Leases	Specify the number of IP addresses that can be assigned on this interface. <i>Range:</i> 0 through 4294967295
Lease Time	Specify how long a DHCP-assigned IP address is valid. <i>Range:</i> 0 through 4294967295 seconds
Offer Time	Specify how long the IP address offered to a DHCP client is reserved for that client. By default, an offered IP address is reserved indefinitely, until the DHCP server runs out of addresses. At that point, the address is offered to another client. <i>Range:</i> 0 through 4294967295 seconds <i>Default:</i> 600 seconds
Administrative State	Select Up to enable or Down to disable the DHCP functionality on the interface. By default, DHCP server functionality is disabled on an interface.

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

```

vpn vpn-id
interface geslot/port
dhcp-server address-pool prefix/length admin-state (down | up)
    exclude ip-address
    lease-time seconds
    max-leases number
    offer-time minutes

```

Configure Static Leases

To configure a static lease to assign a static IP address to a client device on the service-side network, click **Static Lease**, and click **Add New Static Lease** and configure the following parameters:

Table 44:

Parameter Name	Description
MAC Address	Enter the MAC address of the client to which the static IP address is being assigned.
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 **pencil** icon.

To remove a static lease, click **trash** icon.

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 **Advanced** and then configure the following parameters:

Table 45:

Parameter Name	Description
Interface MTU	Specify the maximum MTU size of packets on the interface. <i>Range:</i> 68 to 65535 bytes
Domain Name	Specify the domain name that the DHCP client uses to resolve hostnames.
Default Gateway	Enter the IP address of a default gateway in the service-side network.
DNS Servers	Enter one or more IP address for a DNS server in the service-side network. Separate multiple entries with a comma. You can specify up to eight addresses.
TFTP Servers	Enter the IP address of a TFTP server in the service-side network. You can specify one or two addresses. If two, separate them with a comma.

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

CLI equivalent:

```
vpn vpn-id
interface geslot/port
dhcp-server options
    default-gateway ip-address
    dns-servers ip-address
    domain-name domain-name
    interface-mtu mtu
    tftp-servers ip-address
```

Release Information

Introduced in Cisco vManage 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.

```
vml# 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 10.255.1.1-10.255.1.10**).

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.

Create a VPN-Interface-PPP feature template to configure PPP parameters for the PPP virtual interface:

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Feature Templates**, and click **Add Template**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

3. Choose Cisco vEdge device Cloud or a router model.
4. Choose the **VPN-Interface-PPP** template.
5. In the template, configure the following parameters:

Table 46:

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.
IP MTU	Click Advanced , 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. Starting from Cisco vManage Release 20.9.1, there is 8 bytes overheads deduced based on the specified IP MTU value when configuration is pushed to the device.
Save	To save the feature template, click Save .

To create a VPN-Interface-PPP-Ethernet feature template to enable the PPPoE client on the physical interfaces:

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Feature Templates**, and click **Add Template**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

3. Choose Cisco vEdge device Cloud or a router model.

4. Choose the **VPN-Interface-PPP-Ethernet** template.
5. In the template, configure the following parameters:

Parameter Field	Procedure
Template Name	Enter a name for the template. It can be up to 128 alphanumeric characters.
Description	Enter a description for the template. It can be up to 2048 alphanumeric characters.
Shutdown	Click No to enable the PPPoE-enabled interface.
Interface Name	Enter the name of the physical interface in VPN 0 to associate with the PPP interface.
Description (optional)	Enter a description for the PPPoE-enabled interface.
IP Configuration	Assign an IP address to the physical interface: <ul style="list-style-type: none"> • To use DHCP, select Dynamic. The default administrative distance of routes learned from DHCP is 1. • To configure the IP address directly, enter of the IPv4 address of the interface.
DHCP Helper (optional)	Enter up to four IP addresses for DHCP servers in the network.
Save	To save the feature template, click Save .

To create a VPN feature template to configure the PPPoE-enabled interface in VPN 0, the transport VPN:

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Feature Templates**, and click **Add Template**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

3. Choose Cisco vEdge device Cloud or a router model.
4. Choose the **VPN** template.
5. In the template, configure the following parameters:

Parameter Field	Procedure
Template Name	Enter a name for the template. It can be up to 128 alphanumeric characters.
Description	Enter a description for the template. It can be up to 2048 alphanumeric characters.
VPN Identifier	Enter VPN identifier 0.
Name	Enter a name for the VPN.
Other interface parameters	Configure the desired interface properties.

Parameter Field	Procedure
Save	To save the feature template, click Save .

To create a device template that incorporates the VPN-Interface-PPP, VPN-Interface-PPP-Ethernet, and VPN feature templates:

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Device Templates**, and then click **Create Template**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. From the **Create Template** drop-down list, choose **From Feature Template**.
4. From the **Device Model** drop-down list, choose 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 **Transport & Management VPN**, 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 **Additional VPN 0 Templates**, click the plus sign (+) next to **VPN Interface PPP**.
8. From **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. To create the device template, click **Create**.

To attach a device template to a device:

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Device Templates**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. Choose a template.
4. Click **...**, and click **Attach Device**.

5. 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

Table 47: Feature History

Feature Name	Release Information	Feature Description
Assign Static IP Address to PPP Interface.	Cisco SD-WAN Release 20.4.1 Cisco vManage Release 20.4.1	This feature enables you to assign a static IP address to a PPP interface and configure PPP interface echo requests.

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. Enable the PPP interface to be operationally up:

```
vEdge(config-interface-ppp) # no shutdown
```
4. 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
```
5. Configure a tunnel interface for the PPP interface:

```
vEdge(config-interface-ppp) # tunnel-interface color color
```
6. Optionally, configure the name of the access concentrator used by PPPoE to route connections to the internet:

```
vEdge(config-interface-ppp) # ppp ac-name name
```
7. Link a physical Gigabit Ethernet interface in VPN 0 to the PPP interface:

```
vEdge(config-vpn) # interface slot|port  

vEdge(config-interface-ge) # pppoe-client ppp-interface ppp number
```
8. Enable the physical Gigabit Ethernet interface to be operationally up:

```
vEdge(config-interface-ge) # no shutdown
```
9. Optionally, configure a static IP address for the PPP interface:

```
vEdge(config-vpn)# interface ppp
vEdge(config-interface-ppp)# ppp local-ip ipv4-address
```

10. Optionally, configure the number of consecutive echo requests after which the PPP interface terminates if no responses are received:

```
vEdge(config-interface-ppp)# ppp lcp-echo-failure number
```

11. Optionally, configure the number of seconds between echo requests that the PPP interface sends:

```
vEdge(config-interface-ppp)# ppp lcp-echo-interval seconds
```

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==
 !
 ppp ac-name ac_name
 ppp local-ip 10.1.5.15
 ppp lcp-echo-failure 5
 ppp lcp-echo-interval 25
 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.

```
vEdge# show ppp interface
```

VPN	IFNAME	PPPOE INTERFACE	INTERFACE IP	GATEWAY IP	PRIMARY DNS	SECONDARY DNS	MTU
0	ppp10	ge0/1	10.0.0.11	10.255.255.254	10.8.8.8	10.8.4.4	1150

To view PPPoE session information, use the **show pppoe session** command.

```
vEdge# show pppoe session
```

VPN	IFNAME	SESSION ID	SERVER MAC	LOCAL MAC	PPP INTERFACE	SERVICE AC NAME	SERVICE 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 -

```

Configure PPPoE Over ATM

Table 48: Feature History

Feature Name	Release Information	Description
Configure PPPoE over ATM	Cisco IOS XE Release 17.4.1a Cisco vManage Release 20.4.1	This feature provides support for configuring PPPoEoA on Cisco IOS XE SD-WAN devices. PPPoEoA uses AAL5MUX encapsulation which delivers better efficiency compared to other encapsulation methods.

You can configure PPPoE over ATM interfaces (PPPoEoA) on Cisco IOS XE SD-WAN devices that support ADSL. PPPoEoA uses ATM Adaptation Layer 5 Multiplexed Encapsulation (AAL5MUX) encapsulation to carry PPPoE over ATM permanent virtual circuits (PVCs), providing efficiency gain over AAL5 LLC/SNAP encapsulation.

PPPoEoA over AAL5MUX reduces Subnetwork Access Protocol (SNAP) encapsulation bandwidth usage, using multiplexed (MUX) encapsulation to reduce the number of cells needed to carry voice packets. Deploying the PPPoEoA over ATM AAL5MUX feature in a VoIP environment results in improved throughput and bandwidth usage.

Supported Platforms for PPPoE Over ATM

The following platforms support PPPoE over ATM:

- Cisco 1100 4G/6G Series Integrated Services routers.
- Cisco 1100 Series Integrated Service routers.
- Cisco 1109 Series Integrated Service routers.
- Cisco 111x Series Integrated Service routers.
- Cisco 1111x Series Integrated Service routers.
- Cisco 1120 Series Integrated Service routers.
- Cisco 1160 Series Integrated Service routers.

Configure PPPoE Over ATM using Cisco vManage

You can configure PPPoE using in Cisco vManage using the device CLI template.

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. From **Device Templates**, click **Create Template**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. From the **Create Template** drop-down list, select **CLI Template**.
4. From the **Device Model** drop-down list, select the type of device for which you are creating the template.
5. In **Template Name**, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
6. In **Template Description**, enter a description of the template. The description can be up to 2048 characters and can contain only alphanumeric characters.
7. Choose **Device configuration**. Using this option, you can provide IOS-XE configuration commands that appear in the output of the `show sdwan running-config` command.
8. (Optional) To load the running config of a connected device, select it from the Load Running config from reachable device list and click **Search**.
9. In **CLI Configuration**, enter the configuration either by typing it, cutting and pasting it, or uploading a file. The configuration for PPPoEoA is available in the [Configure PPPoE Over ATM on the CLI](#) section.
10. To convert an actual configuration value to a variable, select the value and click **Create Variable**. Enter the variable name, and click **Create Variable**. You can also type the variable name directly, in the format `{{variable-name}}`; for example, `{{hostname}}`.
11. Click **Add**. The new device template is displayed in the Device Template table. The **Type** column shows **CLI** to indicate that the device template was created from CLI text.

Configure PPPoE Over ATM on the CLI

This section provides example CLI configurations to configure Ppoe over ATM on the CLI.

```
Device(config)# interface atm number
Device(config)# no ip address
Device(config)# interface atm number point-to-point
Device(config)# no atm enable-ilmi-trap
Device(config)# encapsulation aal5mux pppoe-client
Device(config)# pppoe-client dial-pool-number number
Device(config)# interface Dialer dialer-rotary-group-number
Device(config)# mtu bytes
Device(config)# ip address negotiated
Device(config-if)# encapsulation encapsulation-type
Device(config)# load-interval seconds
Device(config)# dialer pool number
Device(config)# dialer-group group-number
Device(config)# ppp mtu adaptive
Device(config)# ppp chap hostname hostname
Device(config)# ppp chap password secret
Device(config)# ppp ipcp address required
Device(config)# ppp link reorders
```

Configuration Example for Configuring PPPoE Over ATM Interfaces

This example shows configuring PPPoE over ATM interfaces.

```

Device(config)# interface ATM0/1/0
Device(config)# no ip address
Device(config)# no atm enable-ilmi-trap
!
Device(config)# interface ATM0/1/0.10 point-to-point
Device(config)# no atm enable-ilmi-trap
Device(config)# cdp enable
Device(config)# pvc 22/62
Device(config)# ubr 1045
Device(config-if)# encapsulation aal5mux pppoe-client
Device(config)# pppoe-client dial-pool-number 120
!
!
Device(config)# interface Dialer 120
Device(config)# mtu 1492
Device(config)# ip address negotiated
Device(config)# ip nat outside
Device(config-if)# encapsulation ppp
Device(config)# load-interval 30
Device(config)# dialer pool 120
Device(config)# dialer-group 1
Device(config)# ppp mtu adaptive
Device(config)# ppp chap hostname test@cisco.com
Device(config)# ppp chap password 0 cisco
Device(config)# ppp ipcp address required
Device(config)# ppp link reorders
!

```

Configuring VRRP

Table 49: Feature History

Feature Name	Release Information	Description
Support for Multiple VRRP Groups on the Same LAN Interface or Sub-interface	Cisco SD-WAN Release 20.3.1 Cisco vManage Release 20.3.1	This feature increases support from one VRRP group per interface to five VRRP groups per interface. Multiple VRRP groups are useful for providing redundancy and for load balancing.



Note The x710 NIC must have the `t->system-> vrrp-adv-t-with-phymac` command configured, for VRRP to function.

The Virtual Router Redundancy Protocol (VRRP) is a LAN-side protocol that 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.

VRRP is only supported with service-side VPNs (VPN 0 and 512 reserved) and if sub-interfaces are used, then the VRRP 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 512 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
```

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

VPN	IF NAME	ID	GROUP	TRACK VIRTUAL PREFIX LIST	PREFIX LIST	VIRTUAL MAC	PRIORITY	VRRP STATE	OMP STATE	ADVERTISEMENT TIMER	DOWN TIMER	MASTER LAST
1	ge0/6.2	2		10.2.2.1	00:0c:29:ab:b7:94	100	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	100	init	down	1	3	
				0000-00-00T00:00:00+00:00	-	-						

In the following example, Router-1 is the primary VRRP, because it has a higher priority value than Router 2:

```
Router-1# show running-config vpn 1
vpn 1
!
interface ge0/1.15
ip address 10.10.1.2/24
mtu 1496
no shutdown
vrrp 15
priority 110
track-omp
ipv4 10.20.23.1
!
!
```

```
Router-1# show vrrp vpn 1
```

VPN	IF NAME	ID	GROUP	TRACK VIRTUAL PREFIX LIST	PREFIX LIST	VIRTUAL MAC	PRIORITY	VRRP STATE	OMP STATE	ADVERTISEMENT TIMER	DOWN TIMER	MASTER LAST
1	ge0/1.1	1		10.20.22.1	00:0c:bd:08:79:a4	100	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	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	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	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	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	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	100	backup	up	1	3	
				2016-01-13T03:10:56+00:00	-	-						

```
Router-1# show vrrp vpn 1 interfaces ge0/1.15 groups 15
```

GROUP	TRACK PREFIX LIST	PREFIX LIST	VIRTUAL IP	VIRTUAL MAC	PRIORITY	VRRP STATE	OMP STATE	ADVERTISEMENT TIMER	DOWN TIMER	MASTER LAST	STATE CHANGE
1	10.20.33.1	00:0c:bd:08:79:a4	110	00:0c:bd:08:79:a4	110	master	up	1	3		

```
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
GROUP		TRACK	PREFIX			VRRP	OMP	ADVERTISEMENT	DOWN	
IF NAME	ID	VIRTUAL IP	PREFIX LIST	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
GROUP		TRACK	PREFIX			VRRP	OMP	ADVERTISEMENT	DOWN	
IF NAME	ID	VIRTUAL IP	PREFIX LIST	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	-	-	-	-	-	-	-	-	-	-

Multiple VRRP Groups on One Interface

Cisco SD-WAN supports configuring multiple VRRP groups on an interface. A use case for configuring this is where primary and secondary IP addresses have been assigned to a single interface. On one interface, you can configure:

- One primary IP address
- Up to four secondary IP addresses

To support each of these IP addresses, you can configure up to 5 VRRP groups (each with a unique group ID) on an interface, subinterface, or integrated routing and bridging (IRB) interface that supports VRRP groups.

The following is an example of configuring 5 VRRP groups on 1 interface.

```

vpn 2
interface ge0/4.2
 ip address 10.0.1.10/24
 ip secondary-address 10.0.2.10/24
 ip secondary-address 10.0.3.10/24
 ip secondary-address 10.0.4.10/24
 mtu 1496
 no shutdown

```

```

vrrp 1
  priority 101
  ipv4 10.0.1.1
!
vrrp 2
  ipv4 10.0.1.2
!
vrrp 3
  priority 101
  ipv4 10.0.2.1
!
vrrp 4
  ipv4 10.0.3.1
!
vrrp 5
  ipv4 10.0.4.1
!
!
!


```

Network Interface Configuration Examples for Cisco vEdge Devices

This topic provides examples of configuring interfaces on Cisco vEdge devices to allow the flow of data traffic across both public and private WAN transport networks.

Connect to a Public WAN

This example shows a basic configuration for two connected to the same public WAN network (such as the Internet). The Cisco vSmart Controller and Cisco vBond Orchestrator are also connected to the public WAN network, and the Cisco vSmart Controller is able to reach all destinations on the public WAN.



For Cisco vEdge device-1, the interface ge0/1 connects to the public WAN, so it is the interface that is configured as a tunnel interface. The tunnel has a color of biz-internet, and the encapsulation used for data traffic is IPsec. The Cisco SD-WAN software creates a single TLOC for this interface, comprising the interface's IP address, color, and encapsulation, and the TLOC is sent to the Cisco vSmart Controller over the OMP session running on the tunnel. The configuration also includes a default route to ensure that the router can reach the Cisco vBond Orchestrator and Cisco vSmart Controller.

```

vpn 0
  interface ge0/1
  ip address 172.16.13.3/24
  tunnel-interface
  encapsulation ipsec
  color biz-internet
  allow-service dhcp
  allow-service dns
  allow-service icmp

```

```

        no allow-service sshd
        no allow-service ntp
        no allow-service stun
    !
    no shutdown
    !
    ip route 0.0.0.0/0 0.0.0.0
    !

```

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 0.0.0.0
!

```

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 0.0.0.0
!

vpn 0
interface ge0/1
ip address 172.16.16.6/24
tunnel-interface
!
no shutdown
!
ip route 0.0.0.0/0 0.0.0.0
!

```

Use the **show interface** command to check that the interfaces are operational and that the tunnel connections have been established. In the Port Type column, tunnel connections are marked as "transport".

```
vEdge-1# show interface vpn 0
```

RX		TX		IF	IF					TCP					
VPN	INTERFACE	IP ADDRESS	STATUS	ADMIN	OPER	ENCAP	PORT	TYPE	MTU	HWADDR	SPEED	DUPLEX	MSS	ADJUST	UPTIME
	PACKETS	PACKETS				TYPE					MBPS				
0	ge0/0	172.16.13.3/24	Up	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	172.17.1.15/24	Up	Up	null	service	1500	00:0c:29:7d:1e:3a	10	full	0	0:02:26:20
217	1											
0	ge0/7	10.0.100.15/24	Up	Up	null	service	1500	00:0c:29:7d:1e:44	10	full	0	0:02:25:02
850	550											
0	system	172.16.255.3/32	Up	Up	null	loopback	1500	00:00:00:00:00:00	10	full	0	0:02:13:31
0	0											

Use the **show control connections** command to check that the Cisco vEdge device has a DTLS or TLS session established to the Cisco vSmart Controller.

vEdge-1# show control connections

PEER	PEER	PEER	SITE	DOMAIN	PEER	PRIVATE	PEER	PUBLIC	LOCAL	COLOR
TYPE	PROTOCOL	SYSTEM IP	ID	ID	PRIVATE IP	PORT	PUBLIC IP	PORT	LOCAL	COLOR
STATE		UPTIME								
vsmart	dtls	172.16.255.19	100	1	10.0.5.19	12346	10.0.5.19	12346	biz-internet	
up		0:02:13:13								
vsmart	dtls	172.16.255.20	200	1	10.0.12.20	12346	10.0.12.20	12346	biz-internet	
up		0:02:13:13								

Use the **show bfd sessions** command to display information about the BFD sessions that have been established between the local Cisco vEdge device and remote routers:

vEdge-1# show bfd sessions

DST PUBLIC	DETECT	TX	SOURCE TLOC	REMOTE TLOC	DST PUBLIC
SYSTEM IP	SITE ID	STATE	COLOR	COLOR	IP
PORT	ENCAP	MULTIPLIER	INTERVAL (msec)	UPTIME	SOURCE IP
					TRANSITIONS
172.16.255.11	100	up	biz-internet	biz-internet	10.1.15.15
12346	ipsec	20	1000	0:02:24:59	1
172.16.255.14	400	up	biz-internet	biz-internet	10.1.15.15
12360	ipsec	20	1000	0:02:24:59	1
172.16.255.16	600	up	biz-internet	biz-internet	10.1.15.15
12346	ipsec	20	1000	0:02:24:59	1
172.16.255.21	100	up	biz-internet	biz-internet	10.1.15.15
12346	ipsec	20	1000	0:02:24:59	1

Use the **show omp tlocs** command to list the TLOCs that the local router has learned from the Cisco vSmart Controller:

vEdge-1# show omp tlocs

C -> chosen
 I -> installed
 Red -> redistributed
 Rej -> rejected
 L -> looped
 R -> resolved
 S -> stale
 Ext -> extranet
 Inv -> invalid

ADDRESS	PRIVATE	BFD	ENCAP	FROM PEER	STATUS	PUBLIC IP	PORT
FAMILY	TLOC IP	COLOR					
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					
10.1.14.14	172.16.255.14	biz-internet	ipsec	172.16.255.19	C,I,R	10.1.14.14	12360
10.1.14.14	12360	up		172.16.255.20	C,R	10.1.14.14	12360


10.1.14.14	12360	up							
	172.16.255.16	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 0.0.0.0
```

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 0.0.0.0

```

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 0.0.0.0

```

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 0.0.0.0

vpn 0
 interface ge0/1
   ip address 172.16.16.6/24
   tunnel-interface
   !
   no shutdown
 ip route 0.0.0.0/0 0.0.0.0

```

Connect to Public and Private WANs, with Separation of Network Traffic

In this example, two Cisco vEdge devices at two different sites each connect to the same public WAN (here, the Internet) and the same private WAN (here, an MPLS network). We want to separate the MPLS network completely so that it is not reachable by the Internet. The Cisco vSmart Controller and Cisco vBond Orchestrator are hosted in the provider's cloud, which is reachable only via the Internet. A third Cisco vEdge device at a third site connects only to the public WAN (Internet).



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

- Complete traffic separation exists between private-WAN (MPLS) traffic and public-WAN (Internet) traffic.
- Each site (that is, each Cisco vEdge device) must have a connection to the Internet, because this is the only way that the overlay network can come up.

To maintain complete separation between the public and private networks so that all MPLS traffic stays within the MPLS network, and so that only public traffic passes over the Internet, we create two overlays, one for the private MPLS WAN and the second for the public Internet. For the private overlay, we want to create data traffic tunnels (which run IPsec and BFD sessions) between private-WAN TLOCs, and for the public overlay we want to create these tunnel connections between Internet TLOCs. To make sure that no data traffic tunnels are established between private-WAN TLOCs and Internet TLOCs, or vice versa, we associate the **restrict** attribute with the color on the private-WAN TLOCs. When a TLOC is marked as restricted, a TLOC on the local router establishes tunnel connections with a remote TLOC only if the remote TLOC has the same color. Put another way, BFD sessions come up between two private-WAN TLOCs and they come up between two public-WAN TLOCs, but they do not come up between an MPLS TLOC and an Internet TLOC.

Each site must have a connection to the public (Internet) WAN so that the overlay network can come up. In this topology, the Cisco vSmart Controller and Cisco vBond Orchestrator are reachable only via the Internet, but the MPLS network is completely isolated from the Internet. This means that if a Cisco vEdge device were to connect just to the MPLS network, it would never be able to discover the Cisco vSmart Controller and Cisco vBond Orchestrators and so would never be able to establish control connections in the overlay network. In order for a Cisco vEdge device in the MPLS network to participate in overlay routing, it must have at least one tunnel connection, or more specifically, one TLOC, to the Internet WAN. (Up to seven TLOCs can be configured on each Cisco vEdge device). The overlay network routes that the router learns over the public-WAN tunnel connection populate the routing table on the Cisco vEdge device and allow the router and all its interfaces and TLOCs to participate in the overlay network.

By default, all tunnel connections attempt to establish control connections in the overlay network. Because the MPLS tunnel connections are never going to be able to establish these connections to the Cisco vSmart Controller or Cisco vBond Orchestrators, we include the **max-control-connections 0** command in the configuration. While there is no harm in having the MPLS tunnels attempt to establish control connections, these attempts will never succeed, so disabling them saves resources on the Cisco vEdge device. Note that **max-control-connections 0** command works only when there is no NAT device between the Cisco vEdge device and the PE router in the private WAN.

Connectivity to sites in the private MPLS WAN is possible only by enabling service-side routing.

Here is the configuration for the tunnel interfaces on vEdge-1. This snippet does not include the service-side routing configuration.

```
vpn 0
  interface ge0/1
    ip address 172.16.13.3/24
    tunnel-interface
      encapsulation ipsec
```

```

        color biz-internet
    !
    no shutdown
!
interface ge0/2
 ip address 10.10.23.3/24
 tunnel-interface
   encapsulation ipsec
   color mpls restrict
   max-control-connections 0
!
 no shutdown
!
 ip route 0.0.0.0/0 0.0.0.0

```

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 0.0.0.0
!

```

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 0.0.0.0
!

```

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 0.0.0.0
!

```


```

vpn 0
 interface ge0/1
   ip address 172.16.16.6/24
   tunnel-interface
   !
   no shutdown
   !
 ip route 0.0.0.0/0 0.0.0.0
 !

```

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 0.0.0.0
!

```

The **show control connections** command lists two DTLS sessions to the Cisco vSmart Controller, one from the public tunnel (color of **biz-internet**) and one from the private tunnel (color of **mpls**):

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

PEER	PEER	PEER	SITE	DOMAIN	PEER	PRIVATE	PEER
PEER	PEER	PEER	SITE	DOMAIN	PEER	PRIVATE	PEER
TYPE	PROTOCOL	SYSTEM IP	ID	ID	PRIVATE IP	PORT	PUBLIC IP
PORT	LOCAL COLOR	STATE	UPTIME				
vsmart	dtls	10.255.1.9	900	1	172.16.8.2	12346	172.16.8.2
12346	mpls	up		0:01:41:17			
vsmart	dtls	10.255.1.9	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
```


DST PUBLIC	DETECT	TX	SOURCE TLOC	REMOTE TLOC	DST PUBLIC
SYSTEM IP	SITE ID	STATE	COLOR	COLOR	IP
PORT	ENCAP	MULTIPLIER	INTERVAL (msec)	UPTIME	SOURCE IP
					TRANSITIONS
10.255.1.5	500	up	mpls	biz-internet	10.10.23.3
12346	ipsec	3	1000	0:06:07:19	1
10.255.1.5	500	up	biz-internet	biz-internet	172.16.31.3
12360	ipsec	3	1000	0:06:07:19	1
10.255.1.6	600	up	mpls	biz-internet	10.10.23.3
12346	ipsec	3	1000	0:06:07:19	1
10.255.1.6	600	up	biz-internet	biz-internet	172.16.31.3
12346	ipsec	3	1000	0:06:07:19	1

Exchange Data Traffic within a Single Private WAN

When the Cisco vEdge device is connected to a private WAN network, such as an MPLS or a metro Ethernet network, and when the carrier hosting the private network does not advertise the router's IP address, remote Cisco vEdge devices on the same private network but at different sites can never learn how to reach that router and hence are not able to exchange data traffic with it by going only through the private network. Instead, the remote routers must route data traffic through a local NAT and over the Internet to a Cisco vBond Orchestrator, which then provides routing information to direct the traffic to its destination. This process can add significant overhead to data traffic exchange, because the Cisco vBond Orchestrator may physically be located at a different site or a long distance from the two Cisco vEdge devices and because it may be situated behind a DMZ.

To allow Cisco vEdge devices at different overlay network sites on the private network to exchange data traffic directly using their private IP addresses, you configure their WAN interfaces to have one of eight private colors, **metro-ethernet**, **mpls**, and **private1** through **private6**. Of these four colors, the WAN interfaces on the Cisco vEdge devices must be marked with the same color so that they can exchange data traffic.

To illustrate the exchange of data traffic across private WANs, let's look at a simple topology in which two Cisco vEdge devices are both connected to the same private WAN. The following figure shows that these two Cisco vEdge devices are connected to the same private MPLS network. The vEdge-1 router is located at Site 1, and vEdge-2 is at Site 2. Both routers are directly connected to PE routers in the carrier's MPLS cloud, and you want both routers to be able to communicate using their private IP addresses.



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

- The Cisco vEdge devices are in different sites; that is, they are configured with different site IDs.
- The Cisco vEdge devices are directly connected to the PE routers in the carrier's MPLS cloud.
- The MPLS carrier does not advertise the link between the Cisco vEdge device and its PE router.

To be clear, if the situation were one of the following, no special configuration would be required:

- vEdge-1 and vEdge-2 are configured with the same site ID.
- vEdge-1 and vEdge-2 are in different sites, and the Cisco vEdge device connects to a CE router that, in turn, connects to the MPLS cloud.
- vEdge-1 and vEdge-2 are in different sites, the Cisco vEdge device connects to the PE router in the MPLS cloud, and the private network carrier advertises the link between the Cisco vEdge device and the PE router in the MPLS cloud.
- vEdge-1 and vEdge-2 are in different sites, and you want them to communicate using their public IP addresses.

In this topology, because the MPLS carrier does not advertise the link between the Cisco vEdge device and the PE router, you use a loopback interface on 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.

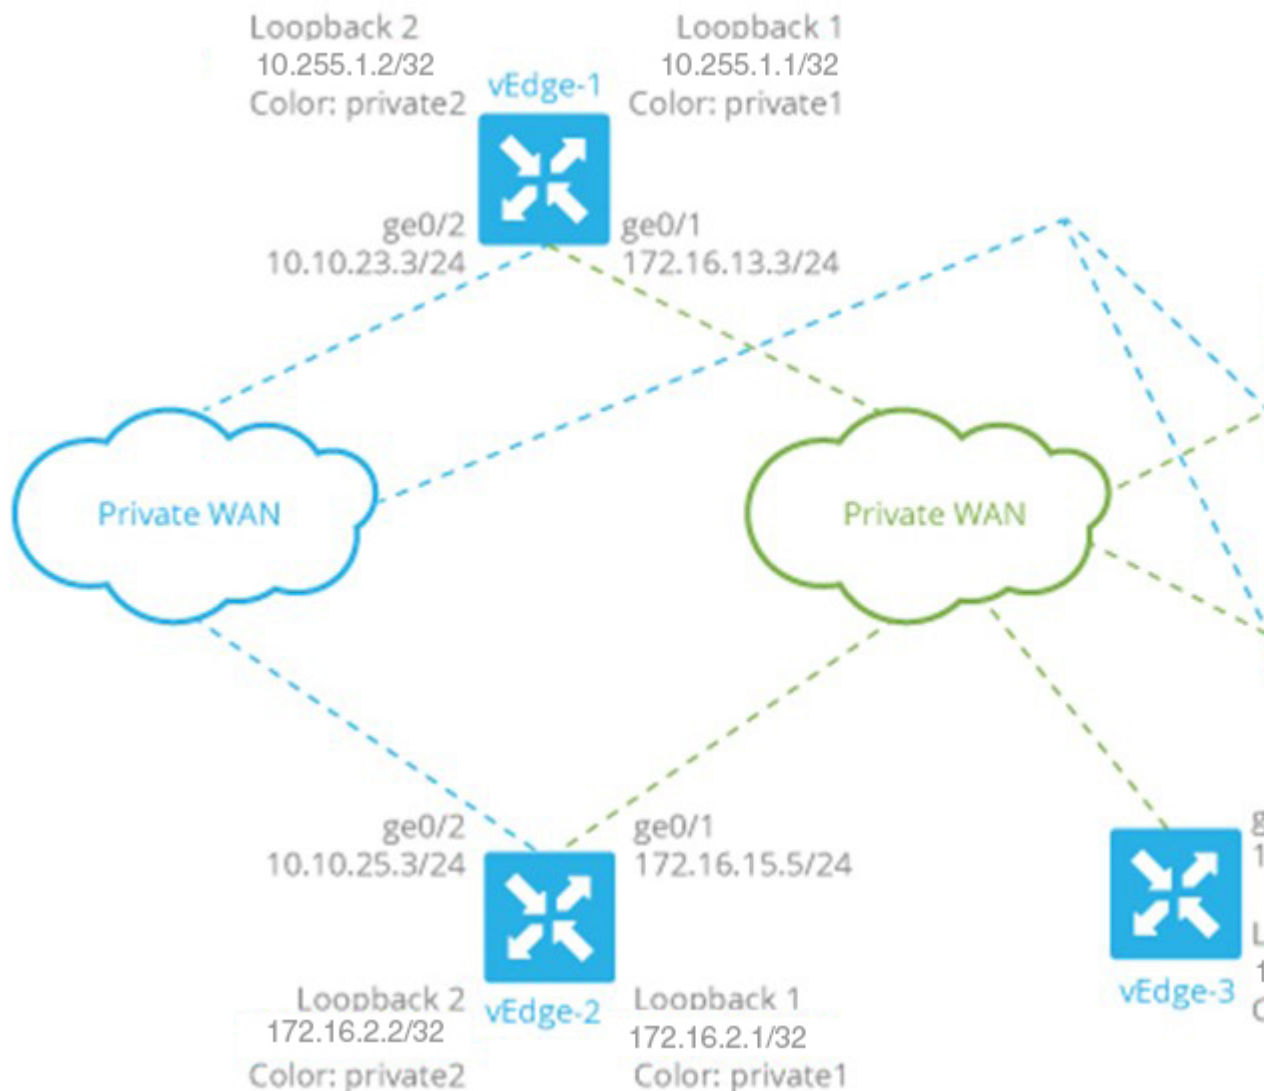
```
vedge-1# show interface
```

VPN	TCP		RX PACKETS	IF STATUS	IF STATUS	ENCAP TYPE	PORT	TYPE	MTU	HWADDR	SPEED MBPS
	MSS	ADJUST									
0	ge0/0		10.1.15.15/24	Up	Up	null	transport	1500	00:0c:29:7d:1e:fe	10	full
0	0	0:07:38:49	213199	243908							
0	ge0/1		10.1.17.15/24	Up	Up	null	service	1500	00:0c:29:7d:1e:08	10	full
0	0	0:07:38:49	197	3							
0	ge0/2		-	Down	Down	null	service	1500	00:0c:29:7d:1e:12	-	-
0	0	-	1	1							
0	ge0/3		10.0.20.15/24	Up	Up	null	service	1500	00:0c:29:7d:1e:1c	10	full
0	0	0:07:38:49	221	27							
0	ge0/6		172.17.1.15/24	Up	Up	null	service	1500	00:0c:29:7d:1e:3a	10	full
0	0	0:07:38:49	196	3							
0	ge0/7		10.0.100.15/24	Up	Up	null	service	1500	00:0c:29:7d:1e:44	10	full
0	0	0:07:44:47	783	497							
0	loopback1		172.16.255.25/32	Up	Up	null	transport	1500	00:00:00:00:00:00	10	full
0	0	0:00:00:20	0	0							
0	system		172.16.255.15/32	Up	Up	null	loopback	1500	00:00:00:00:00:00	10	full
0	0	0:07:38:25	0	0							
1	ge0/4		10.20.24.15/24	Up	Up	null	service	1500	00:0c:29:7d:1e:26	10	full
0	0	0:07:38:46	27594	27405							
1	ge0/5		172.16.1.15/24	Up	Up	null	service	1500	00:0c:29:7d:1e:30	10	full
0	0	0:07:38:46	196	2							
512	eth0		10.0.1.15/24	Up	Up	null	service	1500	00:50:56:00:01:05	1000	full
0	0	0:07:45:55	15053	10333							

To allow Cisco vEdge device at different overlay network sites on the private network to exchange data traffic directly, you use a loopback interface on the each Cisco vEdge device to handle the data traffic instead of using the physical interface that connects to the WAN. You associate the same tag, called a carrier tag, with each loopback interface so that all the routers learn that they are on the same private WAN. Because the loopback interfaces are advertised across the overlay network, the vEdge routers are able to learn reachability information, and they can exchange data traffic over the private network. To allow the data traffic to actually be transmitted out the WAN interface, you bind the loopback interface to a physical WAN interface, specifically to the interface that connects to the private network. Remember that this is the interface that the private network does not advertise. However, it is still capable of transmitting data traffic.

Exchange Data Traffic between Two Private WANs

This example shows a topology with two different private networks, possibly the networks of two different network providers, and all the Cisco SD-WAN devices are located somewhere on one or both of the private networks. Two Cisco vEdge devices are located at two different sites, and they both connect to both private networks. A third Cisco vEdge device connects to only one of the private WANs. The Cisco vBond Orchestrator and Cisco vSmart Controller both sit in one of the private WANs, perhaps in a data center, and they are reachable over both private WANs. For the Cisco vEdge devices to be able to establish control connections, the subnetworks where the Cisco vBond Orchestrator and Cisco vSmart Controller devices reside must be advertised into each private WAN. Each private WAN CPE router then advertises these subnets in its VRF, and each Cisco vEdge device learns those prefixes from each PE router that it is connected to.



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: {10.255.1.1, private2, ipsec} and {10.255.1.2, 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 10.255.1.2/32
    tunnel-interface
      encapsulation ipsec
      color private1
      bind ge0/1
    !
    no shutdown
  !
  interface loopback2
    ip address 10.255.1.1/32
    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 0.0.0.0

```

The configuration for vEdge-2 is similar:

```

vpn 0
  interface loopback1
    ip address 172.16.2.1/32
  tunnel-interface
    encapsulation ipsec
    color private1
    bind ge0/1
  !
  no shutdown
!
  interface loopback2
    ip address 172.16.2.2/32
  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 0.0.0.0
!

```

The vEdge-3 router connects only to the green private WAN:

```

vpn 0
  interface loopback1
    ip address 192.168.3.3/32
  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 0.0.0.0
!

```

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 0.0.0.0
 !

vpn 0
 interface ge0/1
   ip address 172.16.16.6/24
   tunnel-interface
   !
   no shutdown
   !
 ip route 0.0.0.0/0 0.0.0.0
 !

```

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
```

VPN	IFNAME	PPPOE INTERFACE	INTERFACE IP	GATEWAY IP	PRIMARY DNS	SECONDARY DNS	MTU
0	ppp10	ge0/3	10.0.0.11	10.255.255.254	10.8.8.8	10.8.4.4	1150

Use the **show pppoe session** and **show pppoe statistics** commands to view information about PPPoE sessions:

```
vEdge# show pppoe session
```

VPN	IFNAME	SESSION ID	SERVER MAC	LOCAL MAC	PPP INTERFACE	AC NAME	SERVICE 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 Ethernet Interface

Step 1 From the Cisco vManage menu, choose **Configuration > Templates**.

Step 2 Click **Device Templates**, and click **Create Template**.

Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

Step 3 From the **Create Template** drop-down list, choose **From Feature Template**.

Step 4 From the **Device Model** drop-down list, choose 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 **Transport & Management VPN** or scroll to the **Transport & Management VPN** section.
- b. Under **Additional VPN 0 Templates**, click **VPN Interface**.
- c. From the **VPN Interface** drop-down list, click **Create Template**. The **VPN Interface Ethernet** template form displays.

This form contains fields for naming the template, and fields for defining the VPN Interface Ethernet parameters.

Step 6 To create a template for VPNs 1 through 511, and 513 through 65530:

- a. Click **Service VPN**, or scroll to the **Service VPN** section.
- b. Click the **Service VPN** drop-down list.
- c. Under **Additional VPN** templates, click **VPN Interface**.
- d. From the **VPN Interface** drop-down list, click **Create Template**. The VPN Interface Ethernet template form displays.

This form contains fields for naming the template, and fields for defining the VPN Interface Ethernet parameters.

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

Step 8 In **Template Description**, 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 **Basic Configuration** 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.

Parameter Name	IPv4 or IPv6	Options	Description
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.
Secondary IP Address	IPv4		Click Add to enter up to four secondary IPv4 addresses for a service-side interface.
IPv6 Address	IPv6		Click Add to enter up to two secondary IPv6 addresses for a service-side interface.
DHCP Helper	Both		To designate the interface as a DHCP helper on a router, enter up to eight IP addresses, separated by commas, for DHCP servers in the network. A DHCP helper interface forwards BootP (broadcast) DHCP requests that it receives from the specified DHCP servers.
Block Non-Source IP	Yes / No		Click Yes to have the interface forward traffic only if the source IP address of the traffic matches the interface's IP prefix range. Click No to allow other traffic.
Bandwidth Upstream			For Cisco vEdge devices and vManage: For transmitted traffic, set the bandwidth above which to generate notifications. Range: 1 through (232 / 2) – 1 kbps
Bandwidth Downstream			For Cisco vEdge devices and vManage: For received traffic, set the bandwidth above which to generate notifications. Range: 1 through (232 / 2) – 1 kbps

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

CLI Equivalent

```
vpn vpn-id
  interface interface-name
    bandwidth-downstream kbps
```



```

bandwidth-upstream kbps
block-non-source-ip
description text
dhcp-helper ip-address
(ip address ipv4-prefix/length | ip dhcp-client [dhcp-distance number])
(ipv6 address ipv6-prefix/length | ipv6 dhcp-client [dhcp-distance number]
[dhcp-rapid-commit])
secondary-address ipv4-address
[no] shutdown

```

Create a Tunnel Interface

On Cisco vEdge devices, you can configure up to eight tunnel interfaces. This means that each Cisco vEdge device router can have up to eight 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 **Interface Tunnel** and configure the following parameters:

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

Parameter Name	Cisco vEdge devices Only	Description
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)
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

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
```

Create Tunnel Groups

By default, WAN Edge routers try to build tunnels with all other TLOCs in the network, regardless of color. When the restrict option is used with the color designation under the tunnel configuration, the TLOC is restricted to only building tunnels to TLOCs of the same color. For more information on the restrict option see, [Configure Interfaces in the WAN Transport VPN\(VPN0\)](#).

The tunnel group feature is similar to the restrict option but gives more flexibility because once a tunnel group ID is assigned under a tunnel, only TLOCs with the same tunnel group IDs can form tunnels with each other irrespective of color.

If a TLOC is associated with a tunnel group ID, it continues to form tunnels with other TLOCs in the network that are not associated with any tunnel group IDs.



Note The restrict option can still be used in conjunction with this feature. If used, then an interface with a tunnel group ID and restrict option defined on an interface will only form a tunnel with other interfaces with the same tunnel group ID and color.

Configure Tunnel Groups on Cisco vEdge devices Using CLI

To configure tunnel groups on Cisco vEdge devices:

```
vEdge(config)# vpn 0
vEdge(config-vpn-0)# interface interface name
vEdge(config-interface-interface name)# tunnel-interface
vEdge(config-tunnel-interface)# group group-id
```

Limit Keepalive Traffic on a Tunnel Interface

By default, Cisco vEdge devices send a Hello packet once per second to determine whether the tunnel interface between two devices is still operational and to keep the tunnel alive. The combination of a hello interval and a hello tolerance determines how long to wait before declaring a DTLS or TLS tunnel to be down. The default hello interval is 1 second, and the default tolerance is 12 seconds. With these default values, if no Hello packet is received within 11 seconds, the tunnel is declared down at 12 seconds.

If the hello interval or the hello tolerance, or both, are different at the two ends of a DTLS or TLS tunnel, the tunnel chooses the interval and tolerance as follows:

- For a tunnel connection between two controller devices, the tunnel uses the lower hello interval and the higher tolerance interval for the connection between the two devices. (Controller devices are vBond controllers, vManage NMSs, and vSmart controllers.) This choice is made in case one of the controllers

has a slower WAN connection. The hello interval and tolerance times are chosen separately for each pair of controller devices.

- For a tunnel connection between a Cisco vEdge device and any controller device, the tunnel uses the hello interval and tolerance times configured on the router. This choice is made to minimize the amount of traffic sent over the tunnel, to allow for situations where the cost of a link is a function of the amount of traffic traversing the link. The hello interval and tolerance times are chosen separately for each tunnel between a Cisco vEdge device and a controller device.

To minimize the amount of keepalive traffic on a tunnel interface, increase the Hello packet interval and tolerance on the tunnel interface:

```
vEdge(config-tunnel-interface)# hello-interval milliseconds
vEdge(config-tunnel-interface)# hello-tolerance seconds
```

The default hello interval is 1000 milliseconds, and it can be a time in the range 100 through 600000 milliseconds (10 minutes). The default hello tolerance is 12 seconds, and it can be a time in the range 12 through 600 seconds (10 minutes). The hello tolerance interval must be at most one-half the OMP hold time. The default OMP hold time is 60 seconds, and you configure it with the **omp timers holdtime** command.

Configure Multiple Tunnel Interfaces on a vEdge Router

On a Cisco vEdge device, you can configure up to eight tunnel interfaces in the transport interface (VPN 0). This means that each Cisco vEdge device can have up to eight TLOCs.

When a Cisco vEdge device has multiple TLOCs, each TLOC is preferred equally and traffic to each TLOC is weighted equally, resulting in ECMP routing. ECMP routing is performed regardless of the encapsulation used on the transport tunnel, so if, for example, a router has one IPsec and one GRE tunnel, with ECMP traffic is forwarded equally between the two tunnels. You can change the traffic distribution by modifying the preference or the weight, or both, associated with a TLOC. (Note that you can also affect or change the traffic distribution by applying a policy on the interface that affects traffic flow.)

```
vEdge(config)# vpn 0
vEdge(config-vpn-0)# interface interface-name
vEdge(config-tunnel-interface) encapsulation (gre | ipsec)
vEdge(config-encapsulation)# preference number
vEdge(config-encapsulation)# weight number
```

The **preference** command controls the preference for directing inbound and outbound traffic to a tunnel. The preference can be a value from 0 through 4294967295 ($2^{32} - 1$), and the default value is 0. A higher value is preferred over a lower value.

When a Cisco vEdge device has two or more tunnels, if all the TLOCs have the same preference and no policy is applied that affects traffic flow, all the TLOCs are advertised into OMP. When the router transmits or receives traffic, it distributes traffic flows evenly among the tunnels, using ECMP.

When a Cisco vEdge device has two or more tunnels, if the TLOCs have different preferences and a policy that affects traffic flow is not applied, all the TLOCs are advertised to Cisco vSmart Controller via OMP for further processing based on the control policy applied on Cisco vSmart Controller for the corresponding vEdge site-id. When the router transmits or receives traffic, it sends traffic to or receives traffic from only the TLOC with the highest preference. When there are three or more tunnels and two of them have the same preference, traffic flows are distributed evenly between these two tunnels.

A remote Cisco vEdge device trying to reach one of these prefixes selects which TLOC to use from the set of TLOCs that have been advertised. So, for example, if a remote router selects a GRE TLOC on the local router, the remote router must have its own GRE TLOC to be able to reach the prefix. If the remote router

has no GRE TLOC, it is unable to reach the prefix. If the remote router has a single GRE TLOC, it selects that tunnel even if there is an IPsec TLOC with a higher preference. If the remote router has multiple GRE TLOCs, it selects from among them, choosing the one with the highest preference or using ECMP among GRE TLOCs with equal preference, regardless of whether there is an IPsec TLOC with a higher preference.

The **weight** command controls how traffic is balanced across multiple TLOCs that have equal 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

For information on how to configure NAT, see the [Cisco SD-WAN NAT Configuration Guide, Cisco IOS XE Release 17.x](#).

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
```

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

To create Network Address Translation (NAT) pools of IP addresses in VPNs, use the **VPN Interface NAT Pool** template for Cisco vEdge devices. To configure NAT pool interfaces in a VPN using Cisco vManage templates:

1. Create a **VPN Interface NAT Pool** 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 Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Feature Templates**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled **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. This form contains fields for naming the template, fields for defining the 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 50:

Parameter Name	Description
Shutdown*	Yes Click No to enable the interface. No

Parameter Name	Description
Interface Name (1...31)*	Enter a number for the NAT pool interface to use for service-side NAT. For example, <i>natpool22</i> . Range: 1-31
Description	Enter a description for the interface.
IPv4 Address*	Enter the IPv4 address of the interface. The address length determines the number of NAT addresses that the router use at the same time. A Cisco vEdge device router can support a maximum of 250 NAT IP addresses.
Refresh Mode	Select how NAT mappings are refreshed:
bi-directional	Keep active the NAT mappings for inbound and outbound traffic.
outbound	Keep active the NAT mappings for outbound traffic. This is the default.
UDP Timeout	Enter the time when NAT translations over UDP sessions time out. <i>Default: 1 minute</i> Range: 1-65536 minutes
TCP Timeout	Enter the time when NAT translations over TCP sessions time out. <i>Default: 60 minutes (1 hour)</i> Range:1-65536 minutes
Block ICMP	Select whether a Cisco vEdge device that is acting as a NAT device should receive inbound ICMP error messages. By default, the router blocks these error messages. Click Off to receive the ICMP error messages.
Direction	Select the direction in which the NAT interface performs address translation:
inside	Translate the source IP address of packets that are coming from the service side of the Cisco vEdge device and that are destined to transport side of the router. This is the default.
outside	Translate the source IP address of packets that are coming to the Cisco vEdge device from the transport side of the Cisco vEdge device and that are destined to a service-side device.
Overload	Click No to disable dynamic NAT. By default, dynamic NAT is enabled.

Configure a Tracker Interface

1. To create one or more tracker interfaces, click **Tracker**, 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**.
4. 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. Click **Port Forward**.
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 51:

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.
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, click **ACL/QoS**, 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.
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 ARP. 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
```

Configuring VRRP

To have an interface run the Virtual Router Redundancy Protocol (VRRP), which allows multiple routers to share a common virtual IP address for default gateway redundancy, select the VRRP tab. Then click **Add New VRRP** and configure the following parameters:

Parameter Name	Description
Group ID	Enter the virtual router ID, which is a numeric identifier of the virtual router. You can configure a maximum of 24 groups. Range: 1 through 255
Priority	Enter the priority level of the router. The router with the highest priority is elected as primary VRRP router. If two routers have the same priority, the one with the higher IP address is elected as primary VRRP router. Range: 1 through 254 Default: 100

Parameter Name	Description
Timer (milliseconds)	<p>Specify how often the primary VRRP router sends VRRP advertisement messages. If subordinate routers miss three consecutive VRRP advertisements, they elect a new primary VRRP routers.</p> <p>Range: 100 through 40950 milliseconds</p> <p>Default: 100 msec</p> <p>Note When the timer is 100 ms for the VRRP feature template on Cisco IOS XE SD-WAN devices, the VRRP fails if the traffic is high on LAN interface.</p>
Track OMP Track Prefix List	<p>By default, VRRP uses of the state of the service (LAN) interface on which it is running to determine which router is the primary virtual router. 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, configure one of the following:</p> <p>Track OMP—Click On for VRRP to track the Overlay Management Protocol (OMP) session running on the WAN connection. If the primary VRRP router loses all its OMP sessions, VRRP elects a new default gateway from those that have at least one active OMP session.</p> <p>Track Prefix List—Track both the OMP session and a list of remote prefixes, which is defined in a prefix list configured on the local router. If the primary VRRP router loses all its OMP sessions, VRRP failover occurs as described for the Track OMP option. In addition, if reachability to all of 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 routers determine the primary VRRP router.</p>
IP Address	Enter the IP address of the virtual router. This address must be different from the configured interface IP addresses of both the local router and the peer running VRRP.

CLI Equivalent

```

vpn vpn-id
  interface geslot/port[.subinterface]
    vrrp group-number
      ipv4 ip-address
      priority number
      timer seconds
      (track-omp | track-prefix-list list-name)

```

Configure a Prefix List for VRRP

You can configure prefix list tracking for VRRP using device and feature templates. To configure a prefix list, do the following:

1. From the Cisco vManage menu, choose **Configuration > Policy**.
2. Click **Localized Policy**.

3. From the **Custom Options** drop-down list, click **Lists**.
4. Click **Prefix** from the left pane, and click **New Prefix List**.
5. In **Prefix List Name**, enter a name for the prefix list.
6. Choose **IPv4** as the **Internet Protocol**.
7. In **Add Prefix**, enter the prefix entries separated by commas.
8. Click **Add**.
9. Click **Next** and configure **Forwarding Classes/QoS**.
10. Click **Next** and configure **Access Control Lists**.
11. Click **Next** and in **Route Policy** pane, select a relevant route policy and click **...**, and click **Edit** to add the newly added prefix list.
12. From the **Match** pane, click **AS Path List** and in the **Address**, choose the newly added prefix list.
13. Click **Save Match and Actions**.
14. Click **Next** and enter the **Policy Name** and **Policy Description** in the **Policy Overview** screen.
15. Click **Save Policy**.

Configure a Prefix List for VRRP in the Device Template

To configure the Prefix List to the VRRP and the localized policy in the device template, do the following:

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Device Templates**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. Select a relevant device template and click **...** and click **Edit** to edit the template details.
4. From **Policy**, select the policy with the newly added prefix list.
5. Click **Update**.
6. Click **Feature Templates**.
7. Select a relevant device template and click **...** and click **Edit** to edit the template details.
8. Click **VRRP**.
9. Select a relevant group ID and click the pen icon to associate the new prefix-list to the VRRP details.
10. Click the **Track Prefix List** drop-down list and enter the newly added prefix-list name.
11. Click **Save Changes**.
12. Click **Update** to save the changes.
13. Click **Device Templates** and select the policy with the newly added prefix list.

14. Click ... and click **Attach Devices**.
15. From **Available Devices**, double-click the relevant device to move it to **Selected Devices**, and then click **Attach**.

Configure Advanced Properties

To configure other interface properties, select the **Advanced** tab and configure the following parameters:

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 TCP SYN packets passing through the router. By default, the MSS is dynamically adjusted based on the interface or tunnel MTU such that TCP SYN packets are never fragmented. Range: 552 to 1460 bytes Default: None
Speed	Specify the speed of the interface for use when the remote end of the connection does not support autonegotiation. Values: 10, 100, 1000, or 10000 Mbps
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. Note Clear-Dont-Fragment clears the DF bit when there is fragmentation needed and the DF bit is set. For packets not requiring fragmentation, the DF bit is not affected.

Parameter Name	Description
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 (20 minutes)
Autonegotiation	Click Off to turn autonegotiation off. 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 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. Note that TLOC extension over L3 is only supported for Cisco IOS XE routers. If configuring TLOC extension over L3 for a Cisco IOS XE router, enter the IP address of the L3 interface.
Power over Ethernet	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 vpn-id
  interface interface-name
    arp-timeout seconds (on vEdge routers only)
    [no] autonegotiate
    clear-dont-fragment
    duplex (full | half)
    flow-control control
    icmp-redirect-disable (on vEdge routers only)
    mac-address mac-address
    mtu bytes
    pmtu
    pppoe-client (on vEdge 100m and vEdge 100wm routers only)
    ppp-interface pppnumber
    speed speed
    static-ingress-qos number (on vEdge routers only)
    tcp-mss-adjust bytes
    tloc-extension interface-name (on vEdge routers only)
    tracker tracker-name (on vEdge routers only)

```

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. From the Cisco vManage menu, choose **Configuration** > **Templates**.
2. Click **Device Templates**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. From the **Create Template** drop-down list, select **From Feature Template**.
4. From the **Device Model** drop-down list, select the type of device for which you are creating the template.
5. Click **Service VPN** or scroll to the **Service VPN** section.
6. Click the **Service VPN** drop-down list.
7. From **Additional VPN Templates**, click **VPN Interface Bridge**.
8. From the **VPN Interface Bridge** drop-down list, 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 **Template Name**, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
10. In **Template Description**, 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 52:

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

Release Information

Introduced in Cisco vManage NMS in Release 15.3. In Release 18.2, add support for disabling ICMP redirect messages.

Create a Bridging Interface

To configure an interface to use for bridging servers, select **Basic Configuration** and configure the following parameters. Parameters marked with an asterisk are required to configure bridging.

Table 53:

Parameter Name	Description
Shutdown*	Click No to enable the interface.
Interface name*	Enter the name of the interface, in the format irb number . The IRB interface number can be from 1 through 63, and must be the same as the VPN identifier configured in the Bridge feature template for the bridging domain that the IRB is connected to.
Description	Enter a description for the interface.
IPv4 Address*	Enter the IPv4 address of the router.

Parameter Name	Description
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.
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.
Secondary IP Address (on Cisco vEdge devices)	Click Add to configure up to four secondary IPv4 addresses for a service-side interface.

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

CLI equivalent:

```
vpn vpn-id interface irbnumber description "text description" dhcp-helper ip-addresses
ip address prefix/length mac-address mac-address mtu bytes secondary-address ipv4-address
[no] shutdown tcp-mss-adjust bytes
```

Apply Access Lists

Apply Access Lists

To apply access lists to IRB interfaces, select the ACL tab and configure the following parameters. The ACL filter determines what is allowed in or out of a bridging domain:

Table 54:

Parameter Name	Description
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.

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

CLI equivalent:

```
vpn vpn-id interface irbnumber access-list acl-name (in | out)
```

Configure VRRP

To have an interface run the Virtual Router Redundancy Protocol (VRRP), which allows multiple routers to share a common virtual IP address for default gateway redundancy, choose **VRRP**. Then click **Add New VRRP** and configure the following parameters:

Table 55:

Parameter Name	Description
Group ID	Enter the virtual router ID, which is a numeric identifier of the virtual router. You can configure a maximum of 24 groups. <i>Range:</i> 1 through 255
Priority	Enter the priority level of the router. There router with the highest priority is elected as primary VRRP router. If two Cisco vEdge devices have the same priority, the one with the higher IP address is elected as primary VRRP router. <i>Range:</i> 1 through 254 <i>Default:</i> 100
Timer (milliseconds)	Specify how often the primary VRRP router sends VRRP advertisement messages. If subordinate routers miss three consecutive VRRP advertisements, they elect a new primary VRRP router. <i>Range:</i> 100 through 40950 milliseconds <i>Default:</i> 100 msec Note When the timer is 100 ms for the VRRP feature template on s, the VRRP fails if the traffic is high on LAN interface.
Track OMP Track Prefix List	By default, VRRP uses of the state of the service (LAN) interface on which it is running to determine which Cisco vEdge device is the primary virtual router. if a Cisco vEdge device 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, configure one of the following: Track OMP—Click On for VRRP to track the Overlay Management Protocol (OMP) session running on the WAN connection. If the primary VRRP router loses all its OMP sessions, VRRP elects a new default gateway from those that have at least one active OMP session. Track Prefix List—Track both the OMP session and a list of remote prefixes, which is defined in a prefix list configured on the local router. If the primary VRRP router loses all its OMP sessions, VRRP failover occurs as described for the Track OMP option. In addition, if reachability to all of 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 Cisco vEdge devices determine the primary VRRP router.
IP Address	Enter the IP address of the virtual router. This address must be different from the configured interface IP addresses of both the local Cisco vEdge device and the peer running VRRP.

To save the VRRP configuration, click **Add**.

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

CLI equivalent:

```
vpn vpn-id
interface irbnumber[.subinterface]
  vrrp group-number
  ipv4 ip-address
  priority number
```

```
timer seconds
(track-omp | track-prefix-list list-name)
```

Add ARP Table Entries

To configure static Address Resolution Protocol (ARP) table entries on the interface, choose **ARP**. Then click **Add New ARP** and configure the following parameters:

Table 56:

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 Advanced Properties

To configure other interface properties, click **Advanced** and configure the following parameters:

Table 57:

Parameter Name	Description
MAC Address	MAC addresses can be static or dynamic. A static MAC address is manually configured as opposed to a dynamic MAC address that is one learned via an ARP request. You can configure a static MAC on a router's interface or indicate a static MAC that identifies a router's interface. Specify a MAC address to associate with the interface, in colon-separated hexadecimal notation.
IP MTU	Similar to MTU, IP MTU only affects IP packets. If an IP packet exceeds the IP MTU, then the packet will be fragmented. Specify the maximum MTU size of packets on the interface. <i>Range:</i> 576 through 1804 <i>Default:</i> 1500 bytes

Parameter Name	Description
TCP MSS	<p>TCP MSS will affect any packet that contains an initial TCP header that flows through the router. When configured, TCP MSS will be examined against the MSS exchanged in the three-way handshake. The MSS in the header will be lowered if the configured setting is lower than what is in the header. If the header value is already lower, it will flow through unmodified. The end hosts will use the lower setting of the two hosts. If the TCP MSS is to be configured, it should be set it at 40 bytes lower than the minimum path MTU.</p> <p>Specify the maximum segment size (MSS) of TPC SYN packets passing through the Cisco vEdge device. By default, the MSS is dynamically adjusted based on the interface or tunnel MTU such that TCP SYN packets are never fragmented.<i>Range:</i> 552 to 1460 bytes<i>Default:</i> None</p>
Clear-Dont-Fragment	<p>Configure Clear-Dont-Fragment if there are packets arriving on an interface with the DF bit set. If these packets are larger than the MTU will allow, they are dropped. If you clear the df-bit, the packets will be fragmented and sent.</p> <p>Click On to clear the Dont 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.</p> <p>Note Clear-Dont-Fragment clears the DF bit when there is fragmentation needed and the DF bit is set. For packets not requiring fragmentation, the DF bit is not affected.</p>
ARP Timeout	<p>ARP Timeout controls how long we maintain the ARP cache on a router.</p> <p>Specify how long it takes for a dynamically learned ARP entry to time out.</p> <p><i>Range:</i> 0 through 2678400 seconds (744 hours)<i>Default:</i> 1200 seconds (20 minutes)</p>
ICMP Redirect	<p>ICMP Redirects are sent by a router to the sender of an IP packet when a packet is being routed sub-optimally.</p> <p>The ICMP Redirect informs the sending host to forward subsequent packets to that same destination through a different gateway.</p> <p>To disable ICMP redirect messages on the interface, click Disable. By default, an interface allows ICMP redirect messages.</p>

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

CLI equivalent:

```
vpn vpn-id interface irbnumber arp-timeout seconds clear-dont-fragment
icmp-redirect-disable mac-address mac-address mtu bytes tcp-mss-adjust bytes
```

VPN Interface Ethernet PPPoE

Use the PPPoE template for Cisco IOS 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 section.
2. Create a VPN feature template to configure VPN parameters. See VPN help topic.

Navigate to the Template Screen and Name the Template

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Device Templates**, and click **Create Template**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. From the **Create Template** drop-down list, choose **From Feature Template**.
4. From the **Device Model** drop-down list, select the type of device for which you are creating the template.
5. Click **Transport & Management VPN** or scroll to the **Transport & Management VPN** section.
6. Under **Additional VPN 0 Templates**, click **VPN Interface Ethernet PPPoE**.
7. From the **VPN Interface Ethernet PPPoE** drop-down list, click **Create Template**. The VPN Interface Ethernet PPPoE template form is displayed.

This form contains fields for naming the template, and fields for defining the Ethernet PPPoE parameters.



8. In **Template Name**, enter a name for the template.
The name can be up to 128 characters and can contain only alphanumeric characters.
9. In **Template Description**, 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 list and select one of the following:

Table 58:

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

Configure PPPoE Functionality

To configure basic PPPoE functionality, click **Basic Configuration** and configure the following parameters. Required parameters are indicated with an asterisk.

Table 59:

Parameter Name	Description
Shutdown*	Click No to enable the GigabitEthernet interface.
Ethernet Interface Name	Enter the name of a GigabitEthernet interface. For IOS XE routers, you must spell out the interface names completely (for example, GigabitEthernet0/0/0).
VLAN ID	VLAN tag of the sub-interface.
Description	Enter a description of the Ethernet-PPPoE-enabled interface.
Dialer Pool Member	Enter the number of the dialer pool to which the interface belongs. <i>Range:</i> 100 to 255.
PPP Maximum Payload	Enter the maximum receive unit (MRU) value to be negotiated during PPP Link Control Protocol (LCP) negotiation. <i>Range:</i> 64 through 1792 bytes

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

Configure the PPP Authentication Protocol

To configure the PPP Authentication Protocol, click **PPP** and configure the following parameters. Required parameters are indicated with an asterisk.

Table 60:

Parameter Name	Description
PPP Authentication Protocol	<p>Select the authentication protocol used by the MLP:</p> <ul style="list-style-type: none"> • CHAP—Enter the hostname and password provided by your Internet Service Provider (ISP). <i>hostname</i> can be up to 255 characters. • PAP—Enter the username and password provided by your ISP. <i>username</i> can be up to 255 characters. • PAP and CHAP—Configure both authentication protocols. Enter the login credentials for each protocol. To use the same username and password for both, click Same Credentials for PAP and CHAP.

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

Create a Tunnel Interface

On IOS XE routers, you can configure up to eight tunnel interfaces. This means that each router can have up to eight 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 **Tunnel Interface** and configure the following parameters:

Table 61:

Parameter Name	Description
Tunnel Interface	Click On to create a tunnel interface.
Color	Select a color for the TLOC.

Parameter Name	Description
Control Connection	<p>By default, Control Connection is set to On, which establishes a control connection for the TLOC. If the router has multiple TLOCs, click No to have the tunnel not establish control connection for the TLOC.</p> <p>Note We recommend a minimum of 650-700 Kbps bandwidth with default 1 sec hello-interval and 12 sec hello-tolerance parameters configured to avoid any data/packet loss in connection traffic.</p> <p>For each BFD session, an additional average sized BFD packet of 175 Bytes consumes 1.4 Kbps of bandwidth.</p> <p>A sample calculation of the required bandwidth for bidirectional BFD packet flow is given below:</p> <ul style="list-style-type: none"> • 650 – 700 Kbps per device for control connections. • 175 Bytes (or 1.4 Kbps) per BFD session on the device (request) • 175 Bytes (or 1.4 Kbps) per BFD session on the device (response) <p>If the path MTU discovery (PMTUD) is enabled, bandwidth for send/receive BFD packets per tunnel for every 30 secs:</p> <p>A 1500 Bytes BFD request packet is sent per tunnel every 30 secs: $1500 \text{ Bytes} * 8 \text{ bits/1 byte} * 1 \text{ packet} / 30 \text{ secs} = 400 \text{ bps (request)}$</p> <p>A 147 Bytes BFD packet is sent in response: $147 \text{ Bytes} * 8 \text{ bits/1 byte} * 1 \text{ packet} / 30 \text{ secs} = 40 \text{ bps (response)}$</p> <p>Therefore, a device with 775 BFD sessions (for example) requires a bandwidth of:</p> $700k + (1.4k*775) + (400 *775) + (1.4k*775) + (40 *775) = \sim 3,5 \text{ MBps}$
Maximum Control Connections	<p>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.</p> <p><i>Range: 0 through 8 Default: 2</i></p>
Cisco vBond Orchestrator As STUN Server	<p>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.</p>
Exclude Controller Group List	<p>Set the Cisco vSmart Controllers that the tunnel interface is not allowed to connect to. <i>Range: 0 through 100</i></p>
Cisco vManage Connection Preference	<p>Set the preference for using a tunnel interface to exchange control traffic with the Cisco vManage NMS. <i>Range: 0 through 8 Default: 5</i></p>

Parameter Name	Description
Port Hop	Click On to enable port hopping, or click Off to disable it. When a router is behind a NAT, port hopping rotates through a pool of preselected OMP port numbers (called base ports) to establish DTLS connections with other routers when a connection attempt is unsuccessful. The default base ports are 12346, 12366, 12386, 12406, and 12426. To modify the base ports, set a port offset value. <i>Default: Enabled</i>
Low-Bandwidth Link	Select to characterize the tunnel interface as a low-bandwidth link.
Allow Service	Select On or Off for each service to allow or disallow the service on the interface.

To configure additional tunnel interface parameters, click **Advanced Options** and configure the following parameters:

Table 62:

Parameter Name	Description
GRE	Use GRE encapsulation on the tunnel interface. By default, GRE is disabled. If you select both IPsec and GRE encapsulations, two TLOCs are created for the tunnel interface that have the same IP addresses and colors, but that differ by their encapsulation.
IPsec	Use IPsec encapsulation on the tunnel interface. By default, IPsec is enabled. If you select both IPsec and GRE encapsulations, two TLOCs are created for the tunnel interface that have the same IP addresses and colors, but that differ by their encapsulation.
IPsec Preference	Specify a preference value for directing traffic to the tunnel. A higher value is preferred over a lower value. <i>Range: 0 through 4294967295. Default: 0</i>
IPsec Weight	Enter a weight to use to balance traffic across multiple TLOCs. A higher value sends more traffic to the tunnel. <i>Range: 1 through 255. Default: 1</i>
Carrier	Select the carrier name or private network identifier to associate with the tunnel. <i>Values: carrier1, carrier2, carrier3, carrier4, carrier5, carrier6, carrier7, carrier8, default. Default: default</i>
Bind Loopback Tunnel	Enter the name of a physical interface to bind to a loopback interface.

Parameter Name	Description
Last-Resort Circuit	<p>Select to use the tunnel interface as the circuit of last resort.</p> <p>Note An interface configured as a circuit of last resort is expected to be down and is skipped while calculating the number of control connections, the cellular modem becomes dormant, and no traffic is sent over the circuit.</p> <p>When the configurations are activated on the edge device with cellular interfaces, then all the interfaces begin the process of establishing control and BFD connections. When one or more of the primary interfaces establishes a BFD connection, the circuit of last resort shuts itself down.</p> <p>Only when all the primary interfaces lose their connections to remote edges, then the circuit of last resort activates itself triggering a BFD TLOC Down alarm and a Control TLOC Down alarm on the edge device. The last resort interfaces are used as backup circuit on edge device and are activated when all other transport links BFD sessions fail. In this mode the radio interface is turned off, and no control or data connections exist over the cellular interface.</p> <p>Note Configuring administrative distance values on primary interface routes is not supported.</p>
NAT Refresh Interval	Enter the interval between NAT refresh packets sent on a DTLS or TLS WAN transport connection. <i>Range:</i> 1 through 60 seconds. <i>Default:</i> 5 seconds
Hello Interval	Enter the interval between Hello packets sent on a DTLS or TLS WAN transport connection. <i>Range:</i> 100 through 10000 milliseconds. <i>Default:</i> 1000 milliseconds (1 second)
Hello Tolerance	Enter the time to wait for a Hello packet on a DTLS or TLS WAN transport connection before declaring that transport tunnel to be down. <i>Range:</i> 12 through 60 seconds. <i>Default:</i> 12 seconds

Configure the Interface as a NAT Device

To configure an interface to act as a NAT device for applications such as port forwarding, select **NAT**, click **On** and configure the following parameters:

Table 63:

Parameter Name	Description
NAT	Click On to have the interface act as a NAT device.
Refresh Mode	Select how NAT mappings are refreshed, either outbound or bidirectional (outbound and inbound). <i>Default:</i> Outbound
UDP Timeout	Specify when NAT translations over UDP sessions time out. <i>Range:</i> 1 through 65536 minutes. <i>Default:</i> 1 minutes

Parameter Name	Description
TCP Timeout	Specify when NAT translations over TCP sessions time out. <i>Range:</i> 1 through 65536 minutes. <i>Default:</i> 60 minutes (1 hour)
Block ICMP	Select On to block inbound ICMP error messages. By default, a router acting as a NAT device receives these error messages. <i>Default:</i> Off
Respond to Ping	Select On to have the router respond to ping requests to the NAT interface's IP address that are received from the public side of the connection.

To create a port forwarding rule, click **Add New Port Forwarding Rule** and configure the following parameters. You can define up to 128 port-forwarding rules to allow requests from an external network to reach devices on the internal network.

Table 64:

Parameter Name	Description
Port Start Range	Enter a port number to define the port or first port in the range of interest. <i>Range:</i> 0 through 65535
Port End Range	Enter the same port number to apply port forwarding to a single port, or enter a larger number to apply it to a range of ports. <i>Range:</i> 0 through 65535
Protocol	Select the protocol to which to apply the port-forwarding rule, either TCP or UDP. To match the same ports for both TCP and UDP traffic, configure two rules.
VPN	Specify the private VPN in which the internal server resides. This VPN is one of the VPN identifiers in the overlay network. <i>Range:</i> 0 through 65530
Private IP	Specify the IP address of the internal server to which to direct traffic that matches the port-forwarding rule.

To save a port forwarding rule, click **Add**.

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

Apply Access Lists

To apply a rewrite rule, access lists, and policers to a router interface, click **ACL** and configure the following parameters:

Table 65:

Parameter Name	Description
Shaping rate	Configure the aggregate traffic transmission rate on the interface to be less than line rate, in kilobits per second (kbps).
QoS map	Specify the name of the QoS map to apply to packets being transmitted out the interface.

Parameter Name	Description
Rewrite Rule	Click On , and specify the name of the rewrite rule to apply on the interface.
Ingress ACL – IPv4	Click On , and specify the name of the access list to apply to IPv4 packets being received on the interface.
Egress ACL – IPv4	Click On , and specify the name of the access list to apply to IPv4 packets being transmitted on the interface.
Ingress ACL – IPv6	Click On , and specify the name of the access list to apply to IPv6 packets being received on the interface.
Egress ACL – IPv6	Click On , and specify the name of the access list to apply to IPv6 packets being transmitted on the interface.
Ingress Policer	Click On , and specify the name of the policer to apply to packets being received on the interface.
Egress Policer	Click On , and specify the name of the policer to apply to packets being transmitted on the interface.

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

Configure Other Interface Properties

To configure other interface properties, click **Advanced** and configure the following properties:

Table 66:

Parameter Name	Description
Bandwidth Upstream	For transmitted traffic, set the bandwidth above which to generate notifications. <i>Range:</i> 1 through $(2^{32} / 2) - 1$ kbps
Bandwidth Downstream	For received traffic, set the bandwidth above which to generate notifications. <i>Range:</i> 1 through $(2^{32} / 2) - 1$ kbps
IP MTU	Specify the maximum MTU size of packets on the interface. <i>Range:</i> 576 through 1804. <i>Default:</i> 1500 bytes
TCP MSS	Specify the maximum segment size (MSS) of TCP SYN packets passing through the router. By default, the MSS is dynamically adjusted based on the interface or tunnel MTU such that TCP SYN packets are never fragmented. <i>Range:</i> 552 to 1460 bytes. <i>Default:</i> None
TLOC Extension	Enter the name of the physical interface on the same router that connects to the WAN transport circuit. This configuration then binds this service-side interface to the WAN transport. A second router at the same site that itself has no direct connection to the WAN (generally because the site has only a single WAN connection) and that connects to this service-side interface is then provided with a connection to the WAN.

Parameter Name	Description
Tracker	Enter the name of a tracker to track the status of transport interfaces that connect to the internet.
IP Directed-Broadcast	Enables translation of a directed broadcast to physical broadcasts. An IP directed broadcast is an IP packet whose destination address is a valid broadcast address for some IP subnet but which originates from a node that is not itself part of that destination subnet.

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

Release Information

Introduced in Cisco vManage NMS in Release 18.4.1.

VPN Interface GRE

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 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 can 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.
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.
3. Create a data policy on the Cisco vSmart Controller that applies to the service VPN, including a **set-service service-name local** command.

Navigate to the Template Screen and Name the Template

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Device Templates**, and click **Create Template**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. From the **Create Template** drop-down list, select **From Feature Template**.
4. From the **Device Model** drop-down list, 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 **Transport & Management VPN** or scroll to the **Transport & Management VPN** section.
 - b. Under **Additional VPN 0 Templates**, click **VPN Interface GRE**.

- c. From the **VPN Interface GRE** drop-down list, click **Create Template**. The VPN Interface GRE template form is displayed.

This form contains fields for naming the template, and fields for defining the VPN Interface GRE parameters.

6. To create a template for VPNs 1 through 511, and 513 through 65530:
 - a. Click **Service VPN** or scroll to the **Service VPN** section.
 - b. Click the **Service VPN** drop-down list.
 - c. Under **Additional VPN templates**, click **VPN Interface GRE**.
 - d. From the **VPN Interface GRE** drop-down list, click **Create Template**. The VPN Interface GRE template form is displayed.

This form contains fields for naming the template, and fields for defining the VPN Interface GRE parameters.
7. In **Template Name**, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
8. In **Template Description**, 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 the parameter scope.

Configuring a Basic GRE Interface

To configure a basic GRE interface, click **Basic Configuration** and then configure the following parameters. Parameters marked with an asterisk are required to configure a GRE interface.

Table 67:

Parameter Name	Description
Shutdown*	Click Off to enable the interface.
Interface Name*	Enter the name of the GRE interface, in the format gre number . <i>number</i> can be from 1 through 255.
Description	Enter a description of the GRE interface.
Source*	Enter the source of the GRE interface: <ul style="list-style-type: none"> • GRE Source IP Address—Enter the source IP address of the GRE tunnel interface. This address is on the local router. • Tunnel Source Interface—Enter the physical interface that is the source of the GRE tunnel.
Destination*	Enter the destination IP address of the GRE tunnel interface. This address is on a remote device.

Parameter Name	Description
GRE Destination IP Address*	Enter the destination IP address of the GRE tunnel interface. This address is on a remote device
IPv4 Address	Enter an IPv4 address for the GRE tunnel.
IP MTU	Specify the maximum MTU size of packets on the interface. <i>Range: 576 through 1804 Default: 1500 bytes</i>
Clear-Dont-Fragment	Click On to clear the Don't Fragment bit in the IPv4 packet header for packets being transmitted out the interface.
TCP MSS	Specify the maximum segment size (MSS) of TCP SYN packets passing through the Cisco vEdge device. By default, the MSS is dynamically adjusted based on the interface or tunnel MTU such that TCP SYN packets are never fragmented. <i>Range: 552 to 1460 bytes Default: None</i>

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

CLI equivalent:

```
vpn vpn-id interface grenumber clear-dont-fragment description text
ip address ipv4-prefix/length keepalive seconds retries mtu bytes
policer policer-name (in |out)
  qos-map name rewrite-rule name shaping-rate name
  [no] shutdown tcp-mss-adjust bytes tunnel-destination ip-address
  ( tunnel-source ip-address | tunnel-source-interface interface-name)
```

Configure Interface Access Lists

To configure access lists on a GRE interface, click **ACL** and configure the following parameters:

Table 68:

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 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
```

Configure Tracker Interface

To configure a tracker interface to track the status of a GRE interface, select **Advanced** and configure the following parameter:

Table 69:

Parameter Name	Description
Tracker	Enter the name of a tracker to track the status of GRE interfaces that connect to the Internet.

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. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Device Templates**, and from the **Create Template** drop-down list, select **From Feature Template**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. From the **Device Model** drop-down list, select the vEdge device for which you are creating the template.
4. Click **Transport and Management VPN** and the page scrolls to the Transport and Management VPN section.
5. Under **Additional VPN 0 Templates**, click **VPN Interface IPsec**.
6. From the **VPN Interface IPsec** drop-down list, click **Create Template**. The VPN Interface IPsec template form is displayed.
This form contains fields for naming the template, and fields for defining the VPN Interface IPsec parameters.
7. In **Template Name**, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
8. In **Template Description**, 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 list and select one of the following:

Table 70:

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

Configure a Basic IPsec Tunnel Interface

To configure an IPsec tunnel to use for IKE sessions, select the Basic Configuration tab and configure the following parameters. Parameters marked with an asterisk are required to configure an IPsec tunnel.

Table 71:

Parameter Name	Description
Shutdown*	Click No to enable the interface.
Interface Name*	Enter the name of the IPsec interface, in the format ipsec number . <i>number</i> can be from 1 through 256.
Description	Enter a description of the IPsec interface.
IPv4 Address*	Enter the IPv4 address of the IPsec interface, in the format <i>ipv4-prefix/length</i> . The address must be a /30.
Source*	<p>Set the source of the IPsec tunnel that is being used for IKE key exchange:</p> <ul style="list-style-type: none"> • Click IP Address—Enter the IPv4 address that is the source tunnel interface. This address must be configured in VPN 0. • Click Interface—Enter the name of the physical interface that is the source of the IPsec tunnel. This interface must be configured in VPN 0.

Parameter Name	Description
Destination: IPsec Destination IP Address/FQDN*	Set the destination of the IPsec tunnel that is being used for IKE key exchange. Enter either an IPv4 address or the fully qualified DNS name that points to the destination.
TCP MSS	Specify the maximum segment size (MSS) of TCP SYN packets passing through the Cisco vEdge device. By default, the MSS is dynamically adjusted based on the interface or tunnel MTU such that TCP SYN packets are never fragmented. <i>Range: 552 to 1460 bytes</i> <i>Default: None</i>
IP MTU	Specify the maximum MTU size of packets on the interface. <i>Range: 576 through 1804</i> <i>Default: 1500 bytes</i>

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

CLI equivalent:

```
vpn vpn-id
 interface ipsec number ip address ipv4-prefix/length mtu bytes
   no shutdown
   tcp-mss-adjust bytes tunnel-destination ipv4-address
   ( tunnel-source ip-address | tunnel-source-interface interface-name)
```

Configure Dead-Peer Detection

To configure IKE dead-peer detection to determine whether the connection to an IKE peer is functional and reachable, click **DPD** and the page scrolls to the section. Configure the following parameters:

Table 72:

Parameter Name	Description
DPD Interval	Specify the interval for IKE to send Hello packets on the connection. <i>Range: 0 to 30 seconds. Default: 10 seconds</i>
DPD Retries	Specify how many unacknowledged packets to accept before declaring an IKE peer to be dead and then tearing down the tunnel to the peer. <i>Range: 0 to 255. Default: 3</i>

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

CLI equivalent:

```
vpn vpn-id interface ipsec number dead-peer-detection seconds retries number
```

Configure IKE

To configure IKE, click **IKE** 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 73:

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

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

CLI equivalent:

```
vpn vpn-id interface ipsec number ike authentication-type type
  local-id id
  pre-shared-secret password
  remote-id id cipher-suite suite group number mode mode rekey-interval seconds
  version 1
```

To configure IKEv2, configure the following parameters:

Table 74:

Parameter Name	Description
IKE Version	Enter 2 to select IKEv2.

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

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

CLI equivalent:

```
vpn vpn-id interface ipsec number ike authentication-type type
    local-id id
    pre-shared-secret password
    remote-id id cipher-suite suite group number rekey-interval seconds
version 2
```

Configure IPsec Tunnel Parameters

To configure the IPsec tunnel that carries IKE traffic, click **IPsec** and configure the following parameters:

Table 75:

Parameter Name	Description
IPsec Rekey Interval	Specify the interval for refreshing IKE keys. <i>Range:</i> 3600 through 1209600 seconds (1 hour through 14 days). <i>Default:</i> 14400 seconds (4 hours)
IKE Replay Window	Specify the replay window size for the IPsec tunnel. <i>Values:</i> 64, 128, 256, 512, 1024, 2048, 4096, 8192 bytes. <i>Default:</i> 32 bytes
IPsec Cipher Suite	Specify the authentication and encryption to use on the IPsec tunnel. <i>Values:</i> aes256-cbc-sha1 , aes256-gcm , null-sha1 . <i>Default:</i> aes256-gcm
Perfect Forward Secrecy	Specify the PFS settings to use on the IPsec tunnel. <i>Values:</i> • group-2: Use the 1024-bit Diffie-Hellman prime modulus group. • group-14: Use the 2048-bit Diffie-Hellman prime modulus group. • group-15: Use the 3072-bit Diffie-Hellman prime modulus group. • group-16: Use the 4096-bit Diffie-Hellman prime modulus group. • none: Disable PFS. <i>Default:</i> group-16

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

CLI equivalent:

```
vpn vpn-id interface ipsec number ipsec cipher-suite suite perfect-forward-secrecy
pfs-setting rekey-interval seconds replay-window number
```

Release Information

Introduced in Cisco vManage Release 17.2. In Release 17.2.3, add support for PFS. In Release 18.2, 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 section.
2. Create a VPN Interface PPP Ethernet feature template to configure a PPPoE-enabled interface. See VPN Interface PPP Ethernet.
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. From the Cisco vManage menu, choose **Configuration** > **Templates**.
2. Click **Device Templates**, and click **Create Template**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. From the **Create Template** drop-down list, choose **From Feature Template**.
4. From the **Device Model** drop-down list, select the type of device for which you are creating the template.
5. Click **Transport & Management VPN** or scroll to the **Transport & Management VPN** section.
6. Under **Additional VPN 0 Templates**, click **VPN Interface PPP**.
7. From the **VPN Interface PPP** drop-down list, click **Create Template**. The VPN Interface PPP template form is displayed.

This form contains fields for naming the template, and fields for defining the 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 list and select one of the following:

Table 76:

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

Configure a PPP Virtual Interface

To configure a PPP virtual interface, click **Basic Configuration** 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 77:

Parameter Name	Description
Shutdown*	Click No to enable the PPP virtual interface.
PPP Interface Name*	Enter the number of the PPP interface. It can be a number from 1 through 31.
Description	Enter a description for the PPP virtual interface.
Bandwidth Upstream	For transmitted traffic, set the bandwidth above which to generate notifications. <i>Range:</i> 1 through $(2^{32} / 2) - 1$ kbps
Bandwidth Downstream	For received traffic, set the bandwidth above which to generate notifications. <i>Range:</i> 1 through $(2^{32} / 2) - 1$ kbps

Parameter Name	Description
Block Non-Source IP	Click Yes to have the interface forward traffic only if the source IP address of the traffic matches the interface's IP prefix range.

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

CLI equivalent:

```
vpn 0
  interface pppnumber bandwidth-downstream kbps bandwidth-upstream kbps block-non-source-ip
  ppp
    no shutdown
```

Configure the Access Concentrator Name and Authentication Protocol

To configure the access concentrator name, click **PPP** and configure the following parameters:

Table 78:

Parameter Name	Description
AC Name	Name of the access concentrator used by PPPoE to route connections to the Internet.
Authentication Protocol	Select the authentication protocol used by PPPoE: <ul style="list-style-type: none"> • CHAP—Enter the hostname and password provided by your Internet Service Provider (ISP). <i>hostname</i> can be up to 255 characters. • PAP—Enter the username and password provided by your ISP. <i>username</i> can be up to 255 characters. • PAP and CHAP—Configure both authentication protocols. Enter the login credentials for each protocol. To use the same username and password for both, click Same Credentials for PAP and CHAP.

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

CLI equivalent:

```
vpn 0
  interface pppnumber ppp
    ac-name name
    authentication
      chap hostname name password password
      pap password password sent-username name
```

Create a Tunnel Interface

On Cisco vEdge devices, you can configure up to four tunnel interfaces. This means that each Cisco vEdge device can have up to four TLOCs.

For the control plane to establish itself so that the overlay network can function, you must configure WAN transport interfaces in VPN 0.

To configure a tunnel interface for the PPP interface, select the **Tunnel Interface** tab and configure the following parameters:

Table 79:

Parameter Name	Description
Tunnel Interface	Click On to create a tunnel interface.
Color	Select a color for the TLOC.
Control Connection	<p>By default, Control Connection is set to On, which establishes a control connection for the TLOC. If the router has multiple TLOCs, click No to have the tunnel not establish control connection for the TLOC.</p> <p>Note We recommend a minimum of 650-700 Kbps bandwidth with default 1 sec hello-interval and 12 sec hello-tolerance parameters configured to avoid any data/packet loss in connection traffic.</p> <p>For each BFD session, an additional average sized BFD packet of 175 Bytes consumes 1.4 Kbps of bandwidth.</p> <p>A sample calculation of the required bandwidth for bidirectional BFD packet flow is given below:</p> <ul style="list-style-type: none"> • 650 – 700 Kbps per device for control connections. • 175 Bytes (or 1.4 Kbps) per BFD session on the device (request) • 175 Bytes (or 1.4 Kbps) per BFD session on the device (response) <p>If the path MTU discovery (PMTUD) is enabled, bandwidth for send/receive BFD packets per tunnel for every 30 secs:</p> <p>A 1500 Bytes BFD request packet is sent per tunnel every 30 secs: $1500 \text{ Bytes} * 8 \text{ bits/1 byte} * 1 \text{ packet} / 30 \text{ secs} = 400 \text{ bps (request)}$</p> <p>A 147 Bytes BFD packet is sent in response: $147 \text{ Bytes} * 8 \text{ bits/1 byte} * 1 \text{ packet} / 30 \text{ secs} = 40 \text{ bps (response)}$</p> <p>Therefore, a device with 775 BFD sessions (for example) requires a bandwidth of: $700\text{k} + (1.4\text{k}*775) + (400 *775) + (1.4\text{k}*775) + (40 *775) = \sim 3,5 \text{ MBps}$</p>
Maximum Control Connections	<p>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.</p> <p><i>Range: 0 through 8 Default: 2</i></p>
vBond As STUN Server	Click On to enable Session Traversal Utilities for NAT (STUN) to allow the tunnel interface to discover its public IP address and port number when the Cisco vEdge device is located behind a NAT.
Exclude Controller Group List	Set the Cisco vSmart Controller that the tunnel interface is not allowed to connect to. <i>Range: 0 through 100</i>

Parameter Name	Description
vManage Connection Preference	Set the preference for using a tunnel interface to exchange control traffic with the vManage NMS. <i>Range: 0 through 8 Default: 5</i>
Low-Bandwidth Link	Select to characterize the tunnel interface as a low-bandwidth link.
Allow Service	Select On or Off for each service to allow or disallow the service on the interface.

To configure additional tunnel interface parameters, click **Advanced Options** and configure the following parameters:

Table 80:

Parameter Name	Description
GRE	Use GRE encapsulation on the tunnel interface. By default, GRE is disabled. If you select both IPsec and GRE encapsulations, two TLOCs are created for the tunnel interface that have the same IP addresses and colors, but that differ by their encapsulation.
IPsec	Use IPsec encapsulation on the tunnel interface. By default, IPsec is enabled. If you select both IPsec and GRE encapsulations, two TLOCs are created for the tunnel interface that have the same IP addresses and colors, but that differ by their encapsulation.
IPsec Preference	Specify a preference value for directing traffic to the tunnel. A higher value is preferred over a lower value. <i>Range: 0 through 4294967295 Default: 0</i>
IPsec Weight	Enter a weight to use to balance traffic across multiple TLOCs. A higher value sends more traffic to the tunnel. <i>Range: 1 through 255 Default: 1</i>
Carrier	Select the carrier name or private network identifier to associate with the tunnel. <i>Values: carrier1, carrier2, carrier3, carrier4, carrier5, carrier6, carrier7, carrier8, default Default: default</i>
Bind Loopback Tunnel	Enter the name of a physical interface to bind to a loopback interface.
Last-Resort Circuit	Select to use the tunnel interface as the circuit of last resort.
NAT Refresh Interval	Enter the interval between NAT refresh packets sent on a DTLS or TLS WAN transport connection. <i>Range: 1 through 60 seconds Default: 5 seconds</i>
Hello Interval	Enter the interval between Hello packets sent on a DTLS or TLS WAN transport connection. <i>Range: 100 through 10000 milliseconds Default: 1000 milliseconds (1 second)</i>

Parameter Name	Description
Hello Tolerance	Enter the time to wait for a Hello packet on a DTLS or TLS WAN transport connection before declaring that transport tunnel to be down. <i>Range: 12 through 60 seconds Default: 12 seconds</i>

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, click **NAT** and configure the following parameters:

Table 81:

Parameter Name	Description
NAT	Click On to have the interface act as a NAT device.
Refresh Mode	Select how NAT mappings are refreshed, either outbound or bidirectional (outbound and inbound). <i>Default: Outbound</i>
UDP Timeout	Specify when NAT translations over UDP sessions time out. <i>Range: 1 through 65536 minutes</i> <i>Default: 1 minutes</i>
TCP Timeout	Specify when NAT translations over TCP sessions time out. <i>Range: 1 through 65536 minutes</i> <i>Default: 60 minutes (1 hour)</i>
Block ICMP	Select On to block inbound ICMP error messages. By default, a Cisco vEdge device acting as a NAT device receives these error messages. <i>Default: Off</i>
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 82:

Parameter Name	Description
Port Start Range	Enter a port number to define the port or first port in the range of interest. <i>Range:</i> 0 through 65535
Port End Range	Enter the same port number to apply port forwarding to a single port, or enter the larger number to apply it to a range or ports. <i>Range:</i> 0 through 65535
Protocol	Select the protocol to which to apply the port-forwarding rule, either TCP or UDP. To match the same ports for both TCP and UDP traffic, configure two rules.
VPN	Specify the private VPN in which the internal server resides. This VPN is one of the VPN identifiers in the overlay network. <i>Range:</i> 0 through 65535
Private IP	Specify the IP address of the internal server to which to direct traffic that matches the port-forwarding rule.

To save a port forwarding rule, click **Add**.

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

CLI equivalent:

```
vpn vpn-id
interface interface-name nat block-icmp-error port-forward port-start port-number1
port-end port-number2 proto (tcp | udp)
private-ip-address ip-address private-vpn vpn-id refresh (bi-directional | outbound)

respond-to-ping tcp-timeout minutes
udp-timeout minutes
```

Apply Access Lists

To apply a rewrite rule, access lists, and policers to a router interface, select the **ACL** tab and configure the following parameters:

Table 83:

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.

Parameter Name	Description
Egress ACL – IPv6	Click On , and specify the name of the access list to apply to IPv6 packets being transmitted on the interface.
Ingress Policer	Click On , and specify the name of the policer to apply to packets being received on the interface.
Egress Policer	Click On , and specify the name of the policer to apply to packets being transmitted on the interface.

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

CLI equivalent:

```
vpn 0
interface pppnumber access-list acl-name (in | out)
    ipv6 access-list acl-name (in | out)
    policer policer-name (in | out)
    rewrite-rule name
```

Configure Other Interface Properties

To configure other interface properties, click **Advanced** and configure the following properties:

Table 84:

Parameter Name	Description
MAC Address	Specify a MAC address to associate with the interface, in colon-separated hexadecimal notation.
IP MTU	Specify the maximum MTU size of packets on the interface. <i>Range:</i> 576 through 1804 <i>Default:</i> 1500 bytes
TCP MSS	Specify the maximum segment size (MSS) of TCP SYN packets passing through the Cisco vEdge device. By default, the MSS is dynamically adjusted based on the interface or tunnel MTU such that TCP SYN packets are never fragmented. <i>Range:</i> 552 to 1460 bytes <i>Default:</i> None
Clear Dont Fragment	Click On to clear the Don't Fragment bit in the IPv4 packet header for packets being transmitted out the interface. When the DF bit is cleared, packets larger than that interface's MTU are fragmented before being sent.
TLOC Extension	Enter the name of the physical interface on the same router that connects to the WAN transport circuit. This configuration then binds this service-side interface to the WAN transport. A second Cisco vEdge device at the same site that itself has no direct connection to the WAN (generally because the site has only a single WAN connection) and that connects to this service-side interface is then provided with a connection to the WAN.
Tracker	Enter the name of a tracker to track the status of transport interfaces that connect to the internet.

Parameter Name	Description
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. From the Cisco vManage menu, choose **Configuration** > **Templates**.
2. Click **Device Templates**, and click **Create Template**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. From the **Create Template** drop-down list, select **From Feature Template**.
4. From the **Device Model** drop-down list, select the type of device for which you are creating the template.
5. Click **Transport & Management VPN** or scroll to the **Transport & Management VPN** section.
6. Under **Additional VPN 0 Templates**, click **VPN Interface PPP**.

- From the **VPN Interface PPP Ethernet** drop-down list, click **Create Template**. The **VPN Interface PPP Ethernet** template form is displayed.

This form contains fields for naming the template, and fields for defining the VPN Interface PPP parameters.

- 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.
- 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 list and select one of the following:

Table 85:

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

Configure a Basic PPPoE-Enabled Interface

To create a PPPoE-enabled interface on a Cisco vEdge device, select the **Basic Configuration** tab and configure the following parameters. Parameters marked with an asterisk are required to configure the interface.

Table 86:

Parameter Name	Description
Shutdown*	Click No to enable the PPPoE-enabled interface.

Parameter Name	Description
Interface Name*	Enter the name of the physical interface in VPN 0 to associate with the PPP interface. For Cisco IOS 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.
IPv6 Configuration*	To configure a static address for an interface in VPN 0, click Static and enter an IPv6 address. To set the interface as a DHCP client so that the interface to receive its IP address from a DHCP server, click Dynamic. You can optionally set the DHCP distance to specify the administrative distance of routes learned from a DHCP server. The default DHCP distance is 1. You can optionally enable DHCP rapid commit, to speed up the assignment of IP addresses.
DHCP Helper	Enter up to eight IP addresses for DHCP servers in the network, separated by commas, to have the interface be a DHCP helper. A DHCP helper interface forwards BOOTP (Broadcast) DHCP requests that it receives from the specified DHCP servers.
Bandwidth Upstream	For transmitted traffic, set the bandwidth above which to generate notifications. <i>Range:</i> 1 through $(2^{32} / 2) - 1$ kbps
Bandwidth Downstream	For received traffic, set the bandwidth above which to generate notifications. <i>Range:</i> 1 through $(2^{32} / 2) - 1$ kbps

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

CLI equivalent:

```

vpn 0
  interface pppnumber bandwidth-downstream kbps bandwidth-upstream kbps description text
  dhcp-helper ip-address
    ( ip address ipv4-prefix/length | ip-dhcp-client [dhcp-distance number])
    ( ipv6 address ipv6-prefix/length | ipv6 dhcp-client [dhcp-distance number] [
dhcp-rapid-commit]
  pppoe-client ppp-interface pppnumber
  [no] shutdown

```

Apply Access Lists

To configure a shaping rate to a PPPoE-enabled interface and to apply a QoS map, a rewrite rule, access lists, and policers to the interface, click **ACL/QOS** and configure the following parameters:

Table 87:

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.
Ingress ACL – IPv6	Egress ACL – IPv6
Egress ACL – IPv6	Egress ACL – IPv6
Ingress Policer	Click On and specify the name of the policer to apply to packets being received on the interface.
Egress Policer	Click On , and specify the name of the policer to apply to packets being transmitted on the interface.

To save the feature temp

CLI equivalent:

```
vpn 0
interface pppnumber access-list acl-list (in | out)
  policer policer-name (in |out)
  qos-map name rewrite-rule name shaping-rate name
```

Configure Other Interface Properties

To configure other interface properties, click **Advanced** and configure the following properties:

Table 88:

Parameter Name	Description
Duplex	Choose full or half to specify whether the interface runs in full-duplex or half-duplex mode. <i>Default:</i> Full
MAC Address	Specify a MAC address to associate with the interface, in colon-separated hexadecimal notation.
IP MTU	Specify the maximum MTU size of packets on the interface. <i>Range:</i> 576 through 1804 <i>Default:</i> 1500 bytes

Parameter Name	Description
PMTU Discovery	Click On to enable path MTU discovery on the interface. PMTU determines the largest MTU size that the interface supports so that packet fragmentation does not occur.
Flow Control	Select a setting for bidirectional flow control, which is a mechanism for temporarily stopping the transmission of data on the interface. <i>Values:</i> autonet, both, egress, ingress, none <i>Default:</i> autoneg
TCP MSS	Specify the maximum segment size (MSS) of TPC SYN packets passing through the Cisco vEdge device. By default, the MSS is dynamically adjusted based on the interface or tunnel MTU such that TCP SYN packets are never fragmented. <i>Range:</i> 552 to 1460 bytes <i>Default:</i> None
Speed	Specify the speed of the interface, for use when the remote end of the connection does not support autonegotiation. <i>Values:</i> 10, 100, or 1000 Mbps <i>Default:</i> Autonegotiate (10/100/1000 Mbps)
Static Ingress QoS	Specify a queue number to use for incoming traffic. <i>Range:</i> 0 through 7
ARP Timeout	Specify how long it takes for a dynamically learned ARP entry to time out. <i>Range:</i> 0 through 2678400 seconds (744 hours) <i>Default:</i> 1200 seconds (20 minutes)
Autonegotiation	Click Off to turn off autonegotiation. By default, an interface runs in autonegotiation mode.
TLOC Extension	Enter the name of a physical interface on the same router that connects to the WAN transport. This configuration then binds this service-side interface to the WAN transport. A second Cisco vEdge device at the same site that itself has no direct connection to the WAN (generally because the site has only a single WAN connection) and that connects to this service-side interface is then provided with a connection to the WAN.
Power over Ethernet (on Cisco vEdge 100m and Cisco vEdge 100wm routers)	Click On to enable PoE on the interface.
ICMP Redirect	Click Disable to disable ICMP redirect messages on the interface. By default, an interface allows ICMP redirect messages.

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

CLI equivalent:

```
vpn 0
  interface pppnumber arp-timeout seconds
    [no] autonegotiate duplex (full | half)
    flow-control control icmp-redirect-disable mac-address mac-address mtu bytes pmtu
pppoe-client
  ppp-interface pppnumber speed speed
  static-ingress-qos number tcp-mss-adjust bytes tloc-extension interface-name
```

Release Information

Introduced in vManage NMS Release 15.3. In Release 16.3, add support for IPv6. In Release 18.2, add support for disabling ICMP redirect messages.

Cellular Interfaces

To enable LTE connectivity, configure cellular interfaces on a router that has a cellular module. The cellular module provides wireless connectivity over a service provider's cellular network. One use case is to provide wireless connectivity for branch offices.

A cellular network is commonly used as a backup WAN link, to provide network connectivity if all the wired WAN tunnel interfaces on the router become unavailable. You can also use a cellular network as the primary WAN link for a branch office, depending on usage patterns within the branch office and the data rates supported by the core of the service provider's cellular network.

When you configure a cellular interface on a device, you can connect the device to the Internet or another WAN by plugging in the power cable of the device. The device then automatically begins the process of joining the overlay network, by contacting and authenticating with Cisco vBond Orchestrators, Cisco vSmart Controllers, and Cisco vManage systems.

vEdge routers support LTE and CDMA radio access technology (RAT) types.

Configure Cellular Interfaces Using Cisco vManage

To configure cellular interfaces using Cisco vManage templates:

1. Create a VPN Interface Cellular feature template to configure cellular module parameters, as described in this section.
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.

Create VPN Interface Cellular

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Device Templates**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. From the **Create Template** drop-down list, choose **From Feature Template**.
4. From the **Device Model** drop-down list, select the type of device for which you are creating the template.
5. Click **Transport & Management VPN** or scroll to the **Transport & Management VPN** section.
6. Under **Additional Cisco VPN 0 Templates**, click **VPN Interface Cellular**.
7. From the **VPN Interface Cellular** drop-down list, click **Create Template**. The VPN Interface Cellular template form is displayed.

This form contains fields for naming the template, and fields for defining the VPN Interface Cellular parameters.

8. In **Template Name**, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
9. In **Template Description**, 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 list.

Configure Basic Cellular Interface Functionality

To configure basic cellular interface functionality, click **Basic Configuration** 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 89:

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 . For Cisco ISR 1100 and ISR 1100X Series Routers operating with an LTE cellular module (LTE dongle), configure the value as lte .
Interface Name*	Enter the name of the interface. It must be cellular0 .
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.

Parameter Name	Description
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, click **Tunnel**, and configure the following parameters. Parameters marked with an asterisk (*) are required to configure a cellular interface.

Parameter Name	Description
Tunnel Interface*	From the drop-down, select Global . Click On to create a tunnel interface.
Per-tunnel QoS	From the drop-down, select Global . Click On to create per-tunnel QoS. You can apply a Quality of Service (QoS) policy on individual tunnels, and is only supported for hub-to-spoke network topologies.
Per-tunnel QoS Aggregator	From the drop-down, select Global . Click On to create per-tunnel QoS. Note 'bandwidth downstream' is required for per-Tunnel QoS feature to take effect as spoke role.

Parameter Name	Description
Color*	From the drop-down, select Global . Select a color for the TLOC. The color typically used for cellular interface tunnels is lte .
Groups	From the drop-down, select Global . Enter the list of groups in the field.
Border	From the drop-down, select Global . Click On to set TLOC as border 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 allowed 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 Cisco vManage. Range: 0 through 9 Default: 5
Port Hop	From the drop-down, select Global . Click Off to allow port hopping on tunnel interface. Default: On , which disallows port hopping on tunnel interface.
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 90:

Parameter Name	Description
GRE	From the drop-down, select Global . Click On to 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
GRE Preference	From the drop-down, select Global . Enter a value to set GRE preference for TLOC. Range: 0 to 4294967295
GRE Weight	From the drop-down, select Global . Enter a value to set GRE weight for TLOC. Default: 1
IPsec	From the drop-down, select Global . Click On to 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	From the drop-down, select Global . 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	From the drop-down, select Global . Enter a value to set weight for balancing traffic across multiple TLOCs. A higher value sends more traffic to the tunnel. Range: 1 through 255. Default: 1
Carrier	From the drop-down, select Global . From the Carrier drop-down, 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	From the drop-down, select Global . Click On to use the tunnel interface as the circuit of last resort. By default, it is disabled. Note An interface configured as a circuit of last resort is expected to be down and is skipped while calculating the number of control connections, the cellular modem becomes dormant, and no traffic is sent over the circuit. When the configurations are activated on the edge device with cellular interfaces, then all the interfaces begin the process of establishing control and BFD connections. When one or more of the primary interfaces establishes a BFD connection, the circuit of last resort shuts itself down. Only when all the primary interfaces lose their connections to remote edges, then the circuit of last resort activates itself triggering a BFD TLOC Down alarm and a Control TLOC Down alarm on the edge device. The last resort interfaces are used as backup circuit on edge device and are activated when all other transport links BFD sessions fail. In this mode the radio interface is turned off, and no control or data connections exist over the cellular interface.

Parameter Name	Description
Hold Time	From the drop-down, select Global . Enter a value to set last resort hold down time for TLOC. Range: 100 to 10000 msec. Default: 7000 ms.
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, click **NAT**, and configure the following parameters:

Table 91:

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)

Parameter Name	Description
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 92:

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, click **ACL/QoS** and configure the following parameters:

Table 93: Access Lists Parameters

Parameter Name	Description
Shaping rate	Configure the aggregate traffic transmission rate on the interface to be less than line rate, in kilobits per second (kbps).
QoS map	Specify the name of the QoS map to apply to packets being transmitted out the interface.
Rewrite rule	Click On , and specify the name of the rewrite rule to apply on the interface.
Ingress ACL – IPv4	Click On , and specify the name of an IPv4 access list to packets being received on the interface.
Egress ACL– IPv4	Click On , and specify the name of an IPv4 access list to packets being transmitted on the interface.
Ingress ACL – IPv6	Click On , and specify the name of an IPv6 access list to packets being received on the interface.
Egress ACL– IPv6	Click On , and specify the name of an IPv6 access list to packets being transmitted on the interface.
Ingress policer	Click On , and specify the name of the policer to apply to packets being received on the interface.
Egress policer	Click On , and specify the name of the policer to apply to packets being transmitted on the interface.

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

CLI equivalent:

```

vpn 0
 interface cellular0
   access-list acl-name (in | out)
   ipv6 access-list acl-name (in | out)
   policer policer-name (in |out)
   qos-map name rewrite-rule name shaping-rate name

```

Add ARP Table Entries

To configure static Address Resolution Protocol (ARP) table entries on the interface, click **ARP**. Then click **Add New ARP** and configure the following parameters:

Table 94:

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, click **Advanced** and configure the following parameters.

Table 95: Cellular Interfaces Advanced Parameters

Parameter Name	Description
PMTU Discovery	Click On to enable path MTU discovery on the interface, to allow the router to determine the largest MTU size supported without requiring packet fragmentation.
TCP MSS	Specify the maximum segment size (MSS) of TCP SYN packets passing through the router. By default, the MSS is dynamically adjusted based on the interface or tunnel MTU such that TCP SYN packets are never fragmented. 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.
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 Cisco vManage 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 the CLI

To configure a cellular interface on a Cisco vEdge device that has a cellular module:

1. Create a cellular profile:

```
vEdge(config)# cellular cellular number
vEdge(config-cellular)# profile profile-id
```

Each Cisco vEdge device has only one LTE module, so *number* must be 0. The profile identifier can be a value from 1 through 15.

2. If your ISP requires that you configure profile properties, configure one or more of the following:

```
vEdge(config-profile)# apn
    name
vEdge(config-profile)# auth auth-method
vEdge(config-profile)# ip-addr ip-address
vEdge(config-profile)# name name
vEdge(config-profile)# pdn-type type
vEdge(config-profile)# primary-dns ip-address
vEdge(config-profile)# secondary-dns ip-address
vEdge(config-profile)# user-name username
vEdge(config-profile)# user-pass password
```

1. Create the cellular interface:

```
vEdge(config)# vpn 0 interface cellular0
```

2. Enable the cellular interface:

```
vEdge(config-interface)# no shutdown
```

3. For cellular interfaces, you must use a DHCP client to dynamically configure the IP address. This is the default option. To explicitly configure this:

```
vEdge(config-interface)# ip dhcp-client [dhcp-distance number]
```

number is the administrative distance of routes learned from a DHCP server. You can configure it to a value from 1 through 255.

4. Associate the cellular profile with the cellular interface:

```
vEdge(config-interface)# profile profile-id
```

The profile identifier is the number you configured in Step 1.

5. Set the interface MTU:

```
vEdge(config-interface)# mtu bytes
```

The MTU can be 1428 bytes or smaller.

6. By default, the radio access technology (RAT) type is LTE. For 2G/3G networks, change it to CDMA:

```
vEdge(config-interface)# technology cdma
```

If you are using the interface for ZTP, change the technology to **auto**:

```
vEdge(config-interface)# technology auto
```

7. Configure any other desired interface properties.

8. Create a tunnel interface on the cellular interface:

```
vEdge(config-interface)# tunnel-interface
vEdge(config-tunnel-interface)# color color
vEdge(config-tunnel-interface)# encapsulation (gre | ipsec)
```

9. By default, the tunnel interface associated with a cellular interface is not considered to be the circuit of last resort. To allow the tunnel to be the circuit of last resort:

```
vEdge(config-tunnel-interface)# last-resort-circuit
```

An interface configured as a circuit of last resort is expected to be down and is skipped while calculating the number of control connections, the cellular modem becomes dormant, and no traffic is sent over the circuit.

When the configurations are activated on the edge device with cellular interfaces, then all the interfaces begin the process of establishing control and BFD connections. When one or more of the primary interfaces establishes a BFD connection, the circuit of last resort shuts itself down.

Only when all the primary interfaces lose their connections to remote edges, then the circuit of last resort activates itself triggering a BFD TLOC Down alarm and a Control TLOC Down alarm on the edge device. The last resort interfaces are used as backup circuit on edge device and are activated when all other transport links BFD sessions fail. In this mode the radio interface is turned off, and no control or data connections exist over the cellular interface.

10. To minimize the amount of control plane keepalive traffic on the cellular interface, increase the Hello packet interval and tolerance on the tunnel interface:

```
vEdge(config-tunnel-interface)# hello-interval milliseconds
vEdge(config-tunnel-interface)# hello-tolerance seconds
```

The default hello interval is 1000 milliseconds, and it can be a time in the range 100 through 600000 milliseconds (10 minutes). The default hello tolerance is 12 seconds, and it can be a time in the range 12 through 600 seconds (10 minutes). To reduce outgoing control packets on a TLOC, it is recommended that on the tunnel interface you set the hello interval to 60000 milliseconds (10 minutes) and the hello tolerance to 600 seconds (10 minutes) and include the **no track-transport** to disable regular checking of the DTLS connection between the Cisco vEdge device and the vBond orchestrator. For a tunnel connection between a Cisco vEdge device and any controller device, the tunnel uses the hello interval and tolerance times configured on the Cisco vEdge device. This choice is made to minimize the amount of traffic sent over the tunnel, to allow for situations where the cost of a link is a function of the amount of traffic traversing the link. The hello interval and tolerance times are chosen separately for each tunnel between a Cisco vEdge device and a controller device. Another step taken to minimize the amount of control plane traffic is to not send or receive OMP control traffic over a cellular interface when other interfaces are available. This behavior is inherent in the software and is not configurable.

11. If the Cisco vEdge device has two or more cellular interfaces, you can minimize the amount of traffic between the vManage NMS and the cellular interfaces by setting one of the interfaces to be the preferred one to use when sending updates to the vManage NMS and receiving configurations from the vManage NMS:

```
vEdge(config-tunnel-interface)# vmanage-connection-preference number
```

The preference can be a value from 0 through 8. The default preference is 5. To have a tunnel interface never connect to the vManage NMS, set the number to 0. At least one tunnel interface on the Cisco vEdge device must have a nonzero vManage connection preference.

12. Configure any other desired tunnel interface properties.
13. To minimize the amount of data plane keepalive traffic on the cellular interface, increase the BFD Hello packet interval:

```
vEdge(bfd-color-lte)# hello-interval milliseconds
```

The default hello interval is 1000 milliseconds (1 second), and it can be a time in the range 100 through 300000 milliseconds (5 minutes).

To determine the status of the cellular hardware, use the **show cellular status** command.

To determine whether a Cisco vEdge device has a cellular module, use the **show hardware inventory** command.

To determine whether a cellular interface is configured as a last-resort circuit, use the **show control affinity config** and **show control local-properties** commands.



Note If you want to remove a property from the cellular profile, delete the profile entirely from the configuration, and create it again with only the required parameters.



Note An interface configured as a circuit of last resort is expected to be down and is skipped while calculating the number of control connections, the cellular modem becomes dormant, and no traffic is sent over the circuit.

When the configurations are activated on the edge device with cellular interfaces, then all the interfaces begin the process of establishing control and BFD connections. When one or more of the primary interfaces establishes a BFD connection, the circuit of last resort shuts itself down.

Only when all the primary interfaces lose their connections to remote edges, then the circuit of last resort activates itself triggering a BFD TLOC Down alarm and a Control TLOC Down alarm on the edge device. The last resort interfaces are used as backup circuit on edge device and are activated when all other transport links BFD sessions fail. In this mode the radio interface is turned off, and no control or data connections exist over the cellular interface.

Best Practices for Configuring Cellular Interfaces

Cellular technology on edge devices can be used in a number of ways:

- Circuit of last resort: An interface configured as a circuit of last resort is expected to be down and is skipped while calculating the number of control connections, the cellular modem becomes dormant, and no traffic is sent over the circuit.

When the configurations are activated on the edge device with cellular interfaces, then all the interfaces begin the process of establishing control and BFD connections. When one or more of the primary interfaces establishes a BFD connection, the circuit of last resort shuts itself down.

Only when all the primary interfaces lose their connections to remote edges, then the circuit of last resort activates itself triggering a BFD TLOC Down alarm and a Control TLOC Down alarm on the edge device. The last resort interfaces are used as backup circuit on edge device and are activated when all other transport links BFD sessions fail. In this mode the radio interface is turned off, and no control or data connections exist over the cellular interface.

Use the **last-resort-circuit** command to configure a cellular interface to be a circuit of last resort.

- **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 Cisco vManage control traffic over a non-cellular interface: When an edge device has both cellular and non-cellular transport interfaces, by default, the edge device chooses one of the interfaces to use to exchange control traffic with the Cisco vManage. You can configure the edge device to never use the cellular interface to exchange traffic with the Cisco vManage, 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 interfaces have a Cisco 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 Cisco vManage.



Note At least one tunnel interface on the edge device must have a non-0 Cisco vManage connection preference value. Otherwise, the device has no control connections.

WiFi Radio

Use the WiFi Radio template for all devices that support wireless LANs (WLANs).

To configure WLAN radio parameters using Cisco vManage templates:

1. Create a WiFi Radio template to configure WLAN radio parameters, as described in this article.
2. Create a Wifi SSID template to configure an SSID and related parameters.

Create WLAN Feature Template

1. From the Cisco vManage menu, choose **Configuration** > **Templates**.
2. Click **Device Templates**, and click **Create Template**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. From the **Create Template** drop-down list, choose **From Feature Template**.
4. From the **Device Model** drop-down list, select the device model that supports wireless LANs (WLANs).
5. Click **WLAN**, or scroll to the **WLAN** section.
6. From the **WiFi Radio** drop-down list, click **Create Template**. The **WiFi Radio** template form is displayed. This form contains fields for naming the template, and fields for defining the WiFi Radio parameters.
7. In **Template Name**, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
8. In **Template Description**, 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 the WLAN Radio Frequency

To configure the WLAN radio frequency, click **Basic Configuration**, and configure the following parameters. Parameters marked with an asterisk are required to configure the radio.

Table 96:

Parameter Name	Description
Select Radio*	Select the radio band. It can be 2.4 GHz or 5 GHz.
Country*	Select the country where the router is installed.
Channel Bandwidth	Select the IEEE 802.11n and 802.11ac channel bandwidth. For a 5-GHz radio band, the default value is 80 MHz, and for 2.4 GHz, the default is 20 MHz.
Channel	Select the radio channel. The default is "auto", which automatically selects the best channel. For 5-GHz radio bands, you can configure dynamic frequency selection (DFS) channels.
Guard Interval	Select the guard interval. For a 5-GHz radio band, the default value is the short guard interval (SGI) of 400 ns, and for 2.4 GHz, the default is 800 ns.

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

CLI equivalent:

```
wlan frequency channel channel channel-bandwidth megahertz country country guard-interval
nanoseconds
```

Release Information

Introduced in vManage NMS Release 16.3.

WiFi SSID

You can use the WiFi SSID template for all devices that support wireless LANs (WLANs)

To configure SSIDs on the WLAN radio using vManage templates:

1. Create a WiFi SSID template to configure the VAP interfaces to use as SSIDs, as described in this article.
2. Create a WiFi Radio template to configure WLAN radio parameters.
3. Create a Bridge template to assign the VAP interface to a bridging domain.
4. Create a device template that incorporates the WiFi Radio feature template and the Wifi SSID feature template.

Navigate to the Template Screen and Name the Template

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Device Templates**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. From the **Create Template** drop-down list, choose **From Feature Template**.
4. From the **Device Model** drop-down list, select a device that supports wireless LANs (WLANs).
5. Click **WLAN**, or scroll to the **WLAN** section.
6. Under **Additional WiFi Radio Templates**, click **WiFi SSID**.
7. From the **WiFi SSID** drop-down list, click **Create Template**. The **WiFi SSID** template form is displayed. This form contains fields for naming the template, and fields for defining the WiFi SSID parameters.
8. In **Template Name**, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
9. In **Template Description**, 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 list.

WLAN SSID Configuration

To configure SSIDs on a device, configure the following parameters in **Basic Configuration**. Parameters marked with an asterisk are required to configure the SSIDs.

Table 97:

Parameter Name	Description
Interface Name*	Select the VAP interface name.
Shutdown*	Click No to enable the interface.
Description (optional)	Enter a description for the interface.
SSID*	Enter the name of the SSID. It can be a string from 4 through 32 characters. The SSID must be unique. You can configure up to four SSIDs. Each SSID is called a virtual access point (VAP) interface. To a client, each VAP interfaces appears as a different access point (AP) with its own SSID. To provide access to different networks, assign each VAP to a different VLAN.
Maximum Clients	Enter the maximum number of clients allowed to connect to the WLAN. <i>Range:</i> 1 through 50 <i>Default:</i> 25
Data Security	Select the security type to enable user authentication or enterprise WPA security. For user authentication, select from WPA Personal, WPA/WPA2 Personal, or WPA2 Personal, and then enter a clear text or an AES-encrypted key. For enterprise security, select from WPA Enterprise, WPA/WPA2 Enterprise, or WPA2 Enterprise, and then enter a RADIUS server tag.
RADIUS Server	If you select one of the enterprise security methods based on using a RADIUS authentication server, enter the RADIUS server tag.
WPA Personal Key	If you select one of the personal security methods based on preshared keys, enter either a clear text or an AES-encrypted password.
Management Security	If you select one of the WPA2 security methods, select the encryption of management frames to be none, optional, or required.

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

CLI equivalent:

```
wlan frequency interface vapnumber data-security security
description text mgmt-security security radius-servers tag
no shutdown
ssid ssid wpa-personal-key password
```

Release Information

Introduced in Cisco vManage Release 16.3.

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



CHAPTER 9

IPv6 Functionality

This chapter describes the options for enabling IPv6 functionality for Cisco SD-WAN templates and policies. Use the information in this chapter if your deployment uses IPv6.

Configure IPv6 Functionality for an Interface or Subinterface Template

To configure IPv6 functionality for an interface or subinterface template, perform the following steps.

Cisco SD-WAN supports dual stack: you can configure IPv4 and IPv6 in the same deployment. You can configure up to three global IPv6 addresses per interface.

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Feature Templates**, and click **Add Template** to select an appropriate device model.



Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled as **Feature**.

3. Select **Cisco VPN Interface Ethernet** from the list of templates.
4. From **Basic Configuration**, click **IPv6** and configure the parameters that the following table describes:

Parameter Name	Description
Static	This radio button is selected by default because IPv6 addresses are static.
IPv6 Address	Enter the IPv6 address of the interface or subinterface.

Configure IPv6 Functionality for an OMP Template

To configure IPv6 functionality for an Overlay Management Protocol (OMP) template, follow these steps:

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Feature Templates**, and click **Add Template** to select an appropriate device model.



Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled as **Feature**.

3. Select **Cisco OMP** from the list of templates.
4. Click **Advertise** and choose **IPv6** to configure the parameters that the following table describes:

Parameter Name	Description
Connected	Click Off to disable advertising connected routes to OMP. By default, connected routes are advertised to OMP.
Static	Click Off to disable advertising static routes to OMP. By default static routes are advertised to OMP.
BGP	Click On to advertise BGP routes to OMP. By default, BGP routes are not advertised to OMP.

Configure IPv6 Functionality for a BGP Template

To configure IPv6 functionality for a Border Gateway Protocol (BGP) template, follow these steps:

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Feature Templates**, and click **Add Template** to select an appropriate device model.



Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled as **Feature**.

3. Select **Cisco BGP** from the list of templates.
4. Click **Unicast Address Family** and choose **IPv6** to configure the parameters that the following table describes:

Tab	Parameter Name	Description
	Maximum Paths	Specify the maximum number of parallel IBGP paths that can be installed into a route table to enable IBGP multipath load sharing. <i>Range:</i> 0 to 32
	Address Family	Enter the BGP IPv6 unicast address family.
RE-DISTRIBUTE		Click the Redistribute tab, and then click Add New Redistribute .
	Protocol	Select the protocols from which to redistribute routes into BGP, for all BGP sessions. Options are Connected, NAT, OMP, OSPF, and Static. At a minimum, select the following: <ul style="list-style-type: none"> • For service-side BGP routing, select OMP. By default, OMP routes are not redistributed into BGP. • For transport-side BGP routing, select Connected, and then under Route Policy, specify a route policy that has BGP advertise the loopback interface address to its neighbors.

Tab	Parameter Name	Description
	Route Policy	Enter the name of the route policy to apply to redistributed routes.
		Click Add to save the redistribution information.
NETWORK		Click the Network tab, and then click Add New Network .
	Network Prefix	Enter a network prefix, in the format of <i>prefix/length</i> , to be advertised by BGP.
		Click Add to save the network prefix.
AGGREGATE ADDRESS		Click the Aggregate Address tab, and then click Add New Aggregate Address .
	Aggregate Prefix	Enter the prefix of the addresses to aggregate for all BGP sessions, in the format <i>prefix/length</i> .
	AS Set Path	Click On to generate set path information for the aggregated prefixes.
	Summary Only	Click On to filter out more specific routes from BGP updates.
		Click Add to save the aggregate address.

1. In the Neighbor area, click **IPv6**, create a new neighbor or edit an existing one, and then configure the parameters that the following table describes.

Parameters marked with an asterisk are required.

Parameter Name	Description
IPv6 Address*	Specify the IPv6 address of the BGP neighbor.
Description	Enter a description of the BGP neighbor.
Remote AS*	Enter the AS number of the remote BGP peer.
Address Family	Select Global from the drop-down list, click On and select the address family. Enter the address family information.
Shutdown	To shut down a BGP neighbor when you push the template, select Global from the drop-down list and then click Yes . <i>Default: Off</i>

Configure IPv6 Functionality for a VRRP Template

To configure IPv6 functionality for a Virtual Router Redundancy Protocol (VRRP) template, follow these steps:

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Feature Templates**, and click **Add Template** to select an appropriate device model.



Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled as **Feature**.

3. Select **Cisco VPN Interface Ethernet** from the list of templates.
4. Click **VRRP** and choose **IPv6**.
5. Click **New VRRP**.
6. Configure the parameters that the following table describes:

Parameter Name	Description
Group ID	Enter a virtual router ID, which represents a group of routers. Range: 1 through 255
Priority	Enter the priority level of the router within a VRRP group. <ul style="list-style-type: none"> • <i>Range:</i> 1 through 254 • <i>Default:</i> 100
Timer	Not used.
Track OMP	Select On to track the Overlay Management Protocol (OMP) session running on the WAN connection when determining the primary VRRP virtual router. <i>Default:</i> Off
Track Prefix List	Enter a value to track a list of IPv6 remote prefixes. This value is an alphanumeric string that is configured under Policy.
Link Local IPv6 Address	Enter a virtual link local IPv6 address, which represents the link local address of the group. The address should be in standard link local address format. For example, FE80::AB8.
Global IPv6 Address	Enter a virtual global unicast IPv6 address, which represents the global address of the group. The address should be an IPv6 global prefix address that has the same mask as the interface forwarding address on which the VRRP group is configured. For example, 2001::2/124. You can configure up to 3 global IPv6 addresses.

Configure IPv6 Functionality for an ACL Template or a QoS Template

To configure IPv6 functionality for an ACL and QoS template, follow these steps:

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Feature Templates**, and click **Add Template** to select an appropriate device model.



Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled as **Feature**.

3. Select **Cisco VPN Interface Ethernet** from the list of templates.
4. From **ACL/QoS**, configure the parameters that the following table describes:

Parameter Name	Description
Ingress ACL – IPv6	Click On to enable the IPv6 ingress access list.
IPv6 Ingress Access List	Enter the name of the IPv6 ingress access list.
Egress ACL – IPv6	Click On to enable the IPv6 egress access list.
IPv6 Egress Access List	Enter the name of the IPv6 egress access list.

Configure IPv6 Functionality for a Logging Template

To configure IPv6 functionality for a Logging template, follow these steps:

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Feature Templates**, and click **Add Template** and then select an appropriate device model.



Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled as **Feature**.

3. Select **Cisco Logging** from the list of templates.
4. From **Server**, click **IPv6**.
5. Configure the parameters that the following table describes.

Parameter Name	Description
IPv6 Hostname/IPv6 Address	Host name or IP address of the server to direct the logging information.
VPN ID	VPN ID of the VPN source interface.
Source Interface	Name of the source interface.
Priority	Choose the maximum severity of messages that are logged.

Configure IPv6 Functionality for a New Prefix List

To configure an IPv6 address for a new prefix list, follow these steps:

1. From the Cisco vManage menu, choose **Configuration > Policies**.

2. From the **Custom Options** drop-down list, select **Lists**. You can make this selection for a Centralized Policy or a Localized Policy
3. Select **Prefix** from the list on the left and then select **New Prefix List**.
4. Click **IPv6** and enter the IPv6 address in **Add Prefix**.

Configure IPv6 Functionality for a Data Prefix

To configure an IPv6 address for a new prefix list, follow these steps:

1. From the Cisco vManage menu, choose **Configuration > Policies**.
2. From the Custom Options drop-down list, select **Lists**. You can make this selection for a Centralized Policy or a Localized Policy
3. Select **Data Prefix** from the list on the left and then select **New Data Prefix List**.
4. From **Internet Protocol**, click **IPv6** and enter the IPv6 address in **Add Prefix**.

Configure IPv6 Functionality for a Centralized Policy

To configure a centralized policy to apply to IPv6 address families, follow these steps:

1. From the Cisco vManage menu, choose **Configuration > Policies**.
2. From the Custom Options drop-down menu, select **Traffic Policy** under Centralized Policy.
3. Select **Traffic Data**.
4. Select **Add Policy** and click **Create New**.
5. Click **Sequence Type** and then select **Traffic Engineering**.
6. Click **Sequence Rule**.
7. From the **Protocol** drop-down list, select **IPv6** to apply the policy only to IPv6 address families, or select **Both** to apply the policy IPv4 and IPv6 address families.
8. Click **Sequence Type** and then select **QoS**.
9. Click **Sequence Rule**.
10. From the Protocol drop-down list, click **IPv6** to apply the policy only to IPv6 address families, or select **Both** to apply the policy IPv4 and IPv6 address families.

Configure IPv6 Functionality for a Localized Policy

To configure a localized policy to apply to IPv6 address families, follow these steps:

1. From the Cisco vManage menu, choose **Configuration > Policies**.
2. From the **Custom Options** drop-down list, select **Access Control Lists** under Localized Policy.
3. Click **Add Access Control List Policy** and choose **Add IPv6 ACL Policy**. The policy you create will apply only to IPv6 address families.



CHAPTER 10

Configure a Cellular Gateway

Table 98: Feature History

Feature Name	Release Information	Feature Description
Cellular Gateway Configuration	Cisco vManage Release 20.4.1 Cisco IOS XE Release 17.4.1a (on devices)	This feature provides templates for configuring a supported cellular gateway as an IP pass-through device. This release supports the Cisco Cellular Gateway CG418-E and CG522-E.

You can configure a supported cellular gateway as an IP pass-through device. By positioning the configured device in an area in your facility that has a strong LTE signal, the signal can be extended over an Ethernet connection to a routing infrastructure in a location with a weaker LTE signal.

To configure a cellular gateway in Cisco vManage:

1. Create a device template for the **Cisco Cellular Gateway CG418-E** device.

See "Create a Device Template from Feature Templates" in *Systems and Interfaces Configuration Guide*.

After you enter a description for the feature template:

- a. From the Cisco vManage menu, choose **Configuration > Templates**.
- b. Click **Device Templates**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

- c. From the **Create Template** drop-down list choose **From Feature Template**.
- d. From the **Device Model** drop-down list select the type of device for which you are creating the template.
- e. Choose **Cellular Gateway > Cellular Gateway Platform > Create Template**. Then configure the Cellular Gateway Platform feature template as shown in the following table.

Table 99: Cellular Gateway Platform Template Parameters

Parameter Name	Description
Basic Configuration Tab	
Time Zone	Choose the time zone to use for the device. The device uses this time zone for clock synchronization when NTP is configured.
Management Interface	Enter the IPv4 address of the management interface for accessing the device.
Admin-Password	Enter the admin user password for logging in to the device by using an SSH client or a console port.
NTP-Servers	Configure one or more NTP servers to which the device synchronizes its clock.
Cellular Configuration Tab	
IP-Src-Violation	Choose v4 only , v6 only , or v4 and v6 to enable the IP source violation feature for the corresponding IP address types. Choose None if you do not want to enable this feature.
Auto-SIM	Choose On to enable the auto-SIM feature. When this feature is enabled, the device automatically detects the service provider to which SIMs in the device belong and automatically loads the appropriate firmware for that provider.
Primary SIM Slot	Choose the slot that contains the primary SIM card for the device. If the device loses service to this slot, it fails over to the secondary slot.
Failover-Timer (minutes)	Enter the number of minutes that the device waits before trying to communicate with the primary SIM slot after the device detects loss of service to this slot.
Max-Retry	Enter the number of consecutive unsuccessful attempts by the device to communicate with the primary SIM before failing over to the secondary slot

- f. Choose **Cellular Gateway > Cellular Gateway Profile** and choose **Create Template** from the Cellular Gateway Profile drop-down list. Then configure the Cellular Gateway Profile feature template as shown in the following table.

Table 100: Cellular Gateway Profile Template Parameters

Parameter Name	Description
Basic Configuration Tab	
SIM	<p>Choose a SIM slot and configure the following options to create a profile for the SIM in that slot. This profile indicates to the service provider which of its cellular networks the SIM should attach to.</p> <ul style="list-style-type: none"> • Profile ID: Enter a unique ID for the profile • Access Point Name: Enter the name of the access point for this profile • Packet Data Network Type: Choose the type of network for data services for this profile (IPv4, IPv6, or IPv4v6) • Authentication: Choose the authentication method that this profile uses for data, and enter the user name and password for this method in the Profile Username and Profile Password fields that display <p>You can configure one profile for each SIM slot in the device.</p>
Add Profile	<p>Click to add an access point name (APN) profile that the cellular device uses to attach to a cellular network.</p> <p>You can add up to 16 profiles.</p>
Profile ID	<p>Enter a unique identifier for the profile.</p> <p>Valid values: Integers 1 through 16.</p>
Access Point Name	Enter a name to identify the cellular access point.
Packet Data Network Type	Choose the packet data network (PDN) type of the cellular network (IPv4 , IPv6 , or IPv4v6).
Authentication	Choose the authentication method that is used to attach to the cellular access point (none , pap , chap , pap_chap).
Profile Username	If you choose an authentication method other than none , enter the user name to use for authentication when attaching to the cellular access point.
Password	If you choose an authentication method other than none , enter the password to use for authentication when attaching to the cellular access point.

Parameter Name	Description
Add	Click to add the profile your are configuring.
Advanced Configuration Tab	
Attach Profile	Choose the profile that the device uses to connect to the cellular network.
Cellular 1/1 Profile	Choose the profile that the device uses for data connectivity over the cellular network.

2. Attach the device template to the device.

See "Attach and Detach a Device Template" in *Systems and Interfaces Configuration Guide*.



CHAPTER 11

Track Static Routes for Service VPNs

Table 101: Feature History

Feature Name	Release Information	Description
Static Route Tracker for Service VPNs for Cisco vEdge Devices	Cisco SD-WAN Release 20.4.1 Cisco vManage Release 20.4.1	This feature enables you to configure IPv4 static route endpoint tracking for service VPNs. For static routes, endpoint tracking determines whether the configured endpoint is reachable before adding that route to the route table of the device.
TCP/UDP Endpoint Tracker and Dual Endpoint Static Route Tracker for Cisco vEdge devices	Cisco SD-WAN Release 20.7.1 Cisco vManage Release 20.7.1	This feature enables you to configure the TCP/UDP static route endpoint trackers. Using this feature you can also configure IPv4, TCP/UDP dual endpoint static-route tracker groups for service VPNs to enhance the reliability of probes.

- [Information About Static Route Tracking, on page 323](#)
- [Restrictions for IPv4 Static Route Tracking, on page 324](#)
- [Workflow to Configure IPv4 Static Route Tracking, on page 324](#)
- [Configure Static Routes Using CLI, on page 328](#)
- [Configuration Examples for Static Route Tracking Using the CLI, on page 330](#)
- [Verify Static Route Tracking Configuration Using CLI, on page 331](#)

Information About Static Route Tracking

Static-route tracking for service VPNs enables you to track the availability of the configured endpoint address to determine if the static route can be included in the routing table of a device. This is applicable when a site uses a static route in a service VPN to advertise its route over Overlay Management Protocol (OMP). The static route tracker periodically sends ICMP ping probes to the configured endpoint. If the tracker does not receive a response, the static route is not included in the routing table and is not advertised to OMP. You can

configure an alternative next-hop address or a static route with a higher administrative distance to provide a backup path. This path is advertised over OMP.



Note From Cisco SD-WAN Release 20.7.1, you can configure TCP/UDP individual endpoint trackers and configure a tracker group with dual endpoints (using two trackers), and associate the trackers and tracker group to a static route. Dual endpoints help in avoiding false negatives that might be introduced because of the unavailability of the routes.

Restrictions for IPv4 Static Route Tracking

- Only one endpoint tracker is supported per static route per next-hop address.
- IPv6 static routes are not supported.
- You cannot link the same endpoint-tracker to static routes in different VPNs. Endpoint-tracker is identified by a name and can be used for multiple static routes in a single VPN.

Workflow to Configure IPv4 Static Route Tracking

1. Configure an endpoint tracker using the System template.
2. Configure a static route using the VPN template.
3. Apply the tracker to the next-hop address.

Create a Static Route Tracker

Use the **System Template** to create a tracker for static routes.

1. From Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Feature Templates**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

3. Navigate to the **Cisco System** template for the device.



Note For information about creating a System template, see [Create System Template](#).

4. Click **Tracker**. Click **New Endpoint Tracker** to configure the tracker parameters.

Table 102: Tracker Parameters

Field	Description
Name	Name of the tracker. The name can be up to 128 alphanumeric characters.
Threshold	Wait time for the probe to return a response before declaring that the configured endpoint is down. Range is from 100 to 1000 milliseconds. Default is 300 milliseconds.
Interval	Time interval between probes to determine the status of the configured endpoint. Default is 60 seconds (1 minute). Range is from 10 to 600 seconds.
Multiplier	Number of times probes are sent before declaring that the endpoint is down. Range is from 1 to 10. Default is 3.
Tracker Type	From the drop-down, choose Global . From the Tracker Type field drop-down, choose Static Route . From Cisco SD-WAN Release 20.7.1, you can configure a tracker group with dual endpoints on Cisco vEdge devices and associate this tracker group to a static route.
Endpoint Type	Choose endpoint type IP Address.
End-Point Type: IP Address	IP address of the static route end point. This is the destination on the internet to which the router sends probes to determine the status of the route.

5. Click **Add**.
6. Click **Save**.
7. To create a tracker group, click **New Endpoint Tracker**.

From the **Tracker Type** drop-down list, choose **tracker-group** and configure the tracker group parameters.



Note Ensure that you have created two trackers to form a tracker group.

Table 103: Tracker Group Parameters

Fields	Description
Name	Name of the tracker group.
Tracker Type	From the drop-down, choose Global . From the Tracker Type field drop-down, choose Static Route . From Cisco SD-WAN Release 20.7.1, you can configure a tracker group with dual endpoints on Cisco vEdge devices and associate this tracker group to a static route.

Fields	Description
Tracker Elements	This field is displayed only if you chose Tracker-group as the tracker type. Add the existing interface tracker names (separated by a space). When you add this tracker to the template, the tracker group is associated with these individual trackers, and you can then associate the tracker group to a static route.
Tracker Boolean	From the drop-down list, choose Global . This field is displayed only if you chose tracker-group as the Tracker Type . By default, the OR option is selected. Choose AND or OR . OR ensures that the static route status is reported as active if either one of the associated trackers of the tracker group report that the route is active. If you select AND , the static route status is reported as active if both the associated trackers of the tracker group report that the route is active.

8. Click **Add**.

9. Click **Save**.



Note Complete all the mandatory actions before you save the template.

Configure a Next Hop Static Route with Tracker

Use the **VPN** template to associate a tracker to a static route next hop.



Note You can apply only one tracker per static route next hop.

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Feature Templates**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

3. Navigate to the **Cisco VPN Template** for the device.



Note For information about creating a VPN template, see [Create VPN Template](#).

4. Enter **Template Name** and **Description** as required.

- In Basic Configuration, by default, VPN is set to 0. Set a VPN value within (1–511, 513–65530) range for service VPNs, for service-side data traffic on Cisco IOS XE SD-WAN devices.



Note You can configure static route tracker only on service VPNs.

- Click **IPv4 Route**.
- Click **New IPv4 Route**.
- In the **IPv4 Prefix** field, enter a value.
- Click **Next Hop**.
- Click **Add Next Hop** and enter values for the fields listed in the table.

Parameter Name	Description
Address	Specify the next-hop IPv4 address.
Distance	Specify the administrative distance for the route.
Tracker	Enter the name of the gateway tracker to determine whether the next hop is reachable before adding that route to the route table of the device.
Add Next Hop	Enter the name of the gateway tracker with the next hop address to determine whether the next hop is reachable before adding that route to the route table of the device.

- Click **Add** to create the static route with the next-hop tracker.



Note Configuring a static route with a next-hop 'X.X.X.255' is not supported. Cisco vEdge device does not implement RFC 3021.

- Click **Save**.



Note You need to fill all the mandatory fields in the form to save the VPN template.

Monitor Static Route Tracker Configuration

View Static Route Tracker

To view information about a static tracker on a transport interface:

- From the Cisco vManage menu, choose **Monitor > Devices**.

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

2. Choose a device from the list of devices.
3. Click **Real Time**.
4. From the **Device Options** drop-down list, choose **Static Route Tracker Info**.

Configure Static Routes Using CLI

The following sections provide information about how to configure static routes using the CLI.

Configure a Static Route Tracker



Note You can configure static route tracking using the Cisco vManage CLI Add-on feature templates and CLI device templates. For more information on configuring using CLI templates, see [CLI Templates](#).

```
Device# config terminal
Device(config)# system tracker <tracker-name>
Device(config-tracker-trackername)# tracker-type <tracker-type>
Device(config-tracker-trackername)# endpoint-ip <ip-address>
Device(config-tracker-trackername)# threshold <value>
Device(config-tracker-trackername)# multiplier <value>
Device(config-tracker-trackername)# interval <value>
Device(config-tracker-trackername)# exit
```

Configure a Static Route Tracker with TCP Port as the Endpoint

```
Device# config terminal
Device(config)# system tracker <tracker-name>
Device(config-tracker-trackername)# tracker-type <tracker-type>
Device(config-tracker-trackername)# endpoint-ip <ip-address> tcp <port-number>
Device(config-tracker-trackername)# threshold <value>
Device(config-tracker-trackername)# multiplier <value>
Device(config-tracker-trackername)# interval <value>
Device(config-tracker-trackername)# exit
```

Configure a Static Route Tracker with UDP Port as the Endpoint

```
Device# config terminal
Device(config)# system tracker <tracker-name>
Device(config-tracker-trackername)# tracker-type <tracker-type>
Device(config-tracker-trackername)# endpoint-ip <ip-address> udp <port-number>
Device(config-tracker-trackername)# threshold <value>
Device(config-tracker-trackername)# multiplier <value>
Device(config-tracker-trackername)# interval <value>
Device(config-tracker-trackername)# exit
```

Configure Tracker Groups



Note You can create tracker groups to probe static routes from Cisco SD-WAN Release 20.7.1 and Cisco vManage Release 20.7.1.

```
Device# config terminal
Device(config)# system tracker <tracker-name1>
Device(config-tracker-trackername1)# tracker-type <tracker-type>
Device(config-tracker-trackername1)# endpoint-ip <ip-address> tcp <port-number>
Device(config-tracker-trackername1)# threshold <value>
Device(config-tracker-trackername1)# multiplier <value>
Device(config-tracker-trackername1)# interval <value>
Device(config-tracker-trackername1)# exit

Device(config)# system tracker <tracker-name2>
Device(config-tracker-trackername2)# tracker-type <tracker-type>
Device(config-tracker-trackername2)# endpoint-ip <ip-address> udp <port-number>
Device(config-tracker-trackername2)# threshold <value>
Device(config-tracker-trackername2)# multiplier <value>
Device(config-tracker-trackername2)# interval <value>
Device(config-tracker-trackername2)# exit

Device(config)# system tracker <tracker-group-name>
Device(config-tracker-tracker-group-name)# tracker-type <tracker-group>
Device(config-tracker-tracker-group-name)# tracker-elements <tracker-name1> <tracker-name2>
Device(config-tracker-tracker-group-name)# boolean {and | or}
Device(config-tracker-tracker-group-name)# exit
```

Configure a Next Hop Static Route with Tracker

```
Device(config)# system
Device(config)# vpn <vpn-number>
Device(config-vpn-vpn-number)# ip route <ipv4address/prefix> <ip-address>
<administrative-distance> tracker <tracker-name>
```



Note Configuring a static route with a next-hop 'X.X.X.255' is not supported. Cisco vEdge device does not implement RFC 3021.



Note

- Use the **ip route** command to bind a tracker or tracker group with a static route and to configure a backup route for administrative distance that is higher than the default value of 1.
- You can apply only one tracker to an endpoint.
- A tracker group can have a mix of endpoint trackers. For example, you can create a tracker group with an IP address tracker and UDP tracker.

Configuration Examples for Static Route Tracking Using the CLI

Configure Tracker

This example shows how to configure a single static route tracker:

```
config terminal
!
system tracker tracker1
!
tracker-type static-route
endpoint-ip 10.1.1.1
threshold 100
multiplier 5
interval 60
exit
!
vpn 1
ip route 192.0.2.0/24 10.20.24.17 tracker tracker1
ip route 172.16.0.0/12 10.20.24.16 100
```

This example shows how to configure a tracker with TCP port as endpoint:

```
config terminal
!
system tracker tcp-10001
!
tracker-type static-route
endpoint-ip 10.0.0.1 tcp 10001
threshold 100
interval 10
multiplier 1
exit
!
vpn 1
ip route 192.0.0.4/24 10.20.25.18 tracker tcp-10001
```

This example shows how to configure a tracker with UDP port as endpoint:

```
config terminal
!
system tracker udp-10001
!
tracker-type static-route
endpoint-ip 10.0.0.1 udp 10001
threshold 100
interval 10
multiplier 1
exit
!
vpn 1
ip route 192.0.0.4/24 10.20.30.19 tracker udp-10001
```

Configure Tracker Groups

This example shows how to configure a tracker group with two trackers (two endpoints). You can create tracker groups to probes static routes from Cisco SD-WAN Release 20.7.1.


```

config terminal
!
 system tracker tcp-10001
!
   tracker-type static-route
   endpoint-ip 10.1.1.1 tcp 10001
   threshold 100
   multiplier 5
   interval 20
!
 system tracker udp-10002
!
   tracker-type static-route
   endpoint-ip 10.2.2.2 udp 10002
   threshold 100
   multiplier 5
   interval 20
!
system tracker group-tcp-10001-udp-10002
!
   tracker-type tracker-group
   boolean and
   tracker-elements tcp-10001 udp-10002
   exit
!
vpn 1
 ip route 192.168.2.0/16 10.20.24.17 tracker group-tcp-10001-udp-10002
 ip route 192.168.2.0/16 10.20.24.16 100

```

**Note**

- You must configure an administrative distance when you are configuring through CLI templates.
- Use the **ip route** command to bind the tracker or tracker group with a static route and to configure a backup route for administrative distance when it is higher than the default value of 1.
- You can apply only one tracker to an endpoint.
- Configuring a static route with a next-hop 'X.X.X.255' is not supported.
Cisco vEdge device does not implement RFC 3021.

Verify Static Route Tracking Configuration Using CLI

Command Verification

Use the following command to verify if the configuration is committed. The following sample configuration shows tracker definition for a static route tracker and it's application to an IPv4 static route:

```

Device# show running-config system tracker
system
 tracker tracker1
 endpoint-ip 10.1.1.1
 interval 60
 multiplier 5
 tracker-type static-route

```

```

tracker tracker2
endpoint-ip 10.1.1.12
interval 40
multiplier 2
tracker-type static-route

```

Use the following command to verify the IPv4 route:

```
Device# show running-config vpn 1 ip route
```

```

vpn 1
ip route 10.20.30.0/24 10.20.30.1
ip route 192.168.2.0/16 10.20.24.16 100
ip route 192.168.2.0/16 10.20.24.17 tracker tracker1
!

```

The following is a sample output from the **show tracker static-route** command displaying individual static route tracker status:

```

Device# show tracker static-route
TRACKER          RTT IN
NAME            VPN  STATUS  MSEC
-----
tcp-10001      1    UP      0
udp-10002      1    UP      0

```

The following is a sample output from the **show tracker static-route-group** command displaying tracker group status:

```

Device# show tracker static-route-group
TRACKER NAME          VPN  BOOLEAN  STATUS  TRACKER ELEMENT NAME  TRACKER ELEMENT STATUS  TRACKER ELEMENT RTT
-----
group-tcp-10001-udp-10002  1    and      UP      tcp-10001            UP      0
                           udp-10002            UP      0

```

The following is a sample output from the **show ip route static** command:

```

Device# show ip route static
Codes Proto-sub-type:
  IA -> ospf-intra-area, IE -> ospf-inter-area,
  E1 -> ospf-external1, E2 -> ospf-external2,
  N1 -> ospf-nssa-external1, N2 -> ospf-nssa-external2,
  e -> bgp-external, i -> bgp-internal
Codes Status flags:
  F -> fib, S -> selected, I -> inactive,
  B -> blackhole, R -> recursive, L -> import

```

VPN	PREFIX	PROTOCOL	SUB TYPE	IF NAME	NEXTHOP ADDR	NEXTHOP VPN	NEXTHOP TLOC IP	COLOR
1	192.168.2.0/16	STATIC	-	ge0/4	10.20.24.17	-	-	-
-	F,S							
1	192.168.2.0/16	STATIC	-	ge0/4	10.20.24.16	-	-	-
-	F,S							



CHAPTER 12

VRRP Interface Tracking

Table 104: Feature History

Feature Name	Release Information	Description
VRRP Interface Tracking for Cisco vEdge Devices	Cisco SD-WAN Release 20.4.1 Cisco vManage Release 20.4.1	This feature enables VRRP to set the edge as active or standby based on the WAN Interface or SIG tracker events and increase the TLOC preference value on a new VRRP active to ensure traffic symmetry, for Cisco vEdge Devices. In this release, you can configure VRRP interface tracking using only the CLI template.
VRRP Interface Tracking for Cisco vEdge Devices.	Cisco SD-WAN Release 20.7.1 Cisco vManage Release 20.7.1	Starting this release, you can configure VRRP interface tracking through Cisco vManage feature template on Cisco vEdge Devices.

- [Information About VRRP Interface Tracking, on page 333](#)
- [Restrictions and Limitations, on page 334](#)
- [VRRP Tracking Use Cases, on page 334](#)
- [Workflow to Configure VRRP Tracking, on page 335](#)
- [Configure an Object Tracker, on page 335](#)
- [Configure VRRP for a VPN Interface Template and Associate Interface Object Tracker, on page 336](#)
- [Configure VRRP Tracking Using CLI Templates, on page 337](#)
- [Configuration Example for VRRP Object Tracking Using CLI, on page 338](#)
- [Configuration Examples for SIG Object Tracking, on page 339](#)
- [Verify VRRP Tracking, on page 339](#)

Information About VRRP Interface Tracking

The Virtual Router Redundancy Protocol (VRRP) is a LAN-side protocol that provides redundant gateway service for switches and other IP end stations. In Cisco SD-WAN, you can configure VRRP on interfaces and subinterfaces, within a VPN.

For more information, see [Configuring VRRP](#).

The VRRP Tracking feature enables switching to a backup or a secondary VRRP router in the following scenarios:

- If a single tunnel (or two tunnels - when you configure redundancy using Transport Locators (TLOC)) on a vEdge device goes down. In this case, the VRRP priority decrements and the secondary router becomes the primary router. VRRP notifies this change to the overlay through Overlay Management Protocol (OMP).
- VRRP can track up to one interface object or Secure Internet Gateway (SIG) object for a group. The interface object can have up to four interfaces. Hence, a group can track up to four tunnel interfaces. The VRRP priority decrements only if all the interfaces of an interface object go down.

Restrictions and Limitations

- VRRP is only supported with service-side VPNs. If you are using subinterfaces, configure VRRP physical interfaces in VPN 0.
- VRRP tracking is enabled on either a physical uplink interface or a logical tunnel interface (IPSEC or GRE or both).
- The VRRP Tracking feature does not support IP prefix as an object.
- You can track a maximum of four interfaces simultaneously using a single tracker. VRRP state transition gets triggered only if all four interfaces go down.
- You can use the same tracker under multiple VRRP groups or VPNs.
- You cannot configure **tloc-change** and **increase-preference** on more than one VRRP group.
- In Cisco SD-WAN release 20.6.1 and earlier releases, you can configure VRRP tracking only through Cisco vManage CLI template.



Note Starting from Cisco SD-WAN release 20.7.1, you can configure VRRP tracking using Cisco vManage feature template as well.



Note In Cisco SD-WAN release 20.6.1 and earlier releases, to update any existing VRRP configuration and add VRRP tracking, convert the configuration and the VRRP tracking commands to the CLI template.

VRRP Tracking Use Cases

The VRRP state is determined based on the tunnel link status. If the tunnel or interface is down on the primary VRRP, then the traffic is directed to the secondary VRRP. The secondary VRRP router in the LAN segment becomes primary VRRP to provide gateway for the service-side traffic.

Zscaler Tunnel Use Case 1—Primary VRRP, Single Internet Provider

The primary and secondary Zscaler tunnels are connected through a single internet provider to the primary VRRP. The primary and secondary VRRP routers are connected through using TLOC extension. In this scenario, the VRRP state transition occurs if the primary and secondary tunnels go down on primary VRRP. The predetermined priority value decrements when the tracking object is down, which triggers the VRRP state transition. To avoid asymmetric routing, VRRP notifies this change to the Overlay through OMP.

Zscaler Tunnel Use Case 2—VRRP Routers in TLOC Extension, Dual Internet Providers

The primary and secondary VRRP routers are configured in TLOC extension high availability mode. The primary and secondary Zscaler tunnels are directly connected with primary and secondary VRRP routers, respectively, using dual internet providers. In this scenario too, the VRRP state transition occurs if the primary and secondary tunnels go down on primary VRRP. The predetermined priority value decrements when the tracking object is down, which triggers the VRRP state transition. VRRP notifies this change to the Overlay through OMP.

TLOC Preference

Transport Locators (TLOCs) connect an OMP route to a physical location. A TLOC is directly reachable using an entry in the routing table of the physical network, or represented by a prefix beyond a NAT device.

The TLOC change preference is an optional configuration under VRRP group. If you configure TLOC change preference value using the `tloc-change-pref` command, the value increases by 1 when a node becomes the primary node. The configured or default TLOC preference is applied back on standby state.



Note We recommend that you use the same TLOC preference value for all TLOCs in a site. For a Cisco vEdge device, the default TLOC preference for the tunnel interface can be modified irrespective of whether VRRP is configured or not. However, if you want to use the VRRP tracking feature and utilize the advantage of TLOC preference values for VRRP tracking, ensure that the default tunnel preference is same on both the VRRP routers.

Workflow to Configure VRRP Tracking

1. Configure an object tracker. For more information, see [Configure an Object Tracker, on page 335](#).
2. Configure VRRP for a VPN Interface template and associate the object tracker with the template. For more information, see [Configure VRRP for a VPN Interface Template and Associate Interface Object Tracker, on page 336](#).

Configure an Object Tracker

Use the **System** template to configure an object tracker.

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Feature**.
3. Navigate to the **System** template for the device.



Note To create a **System** template, see [Create System Template](#)

- Click **Tracker**, and click **New Object Tracker** to configure the tracker parameters.

Table 105: Tracker Parameters

Field	Description
Tracker Type	Choose Interface or SIG to configure the Object tracker.
Tracker List	Enter the name of the tracker list.
Interface	Choose global or device-specific tracker interface name.

- Click **Add**.
- Click **Save**.

Configure VRRP for a VPN Interface Template and Associate Interface Object Tracker

To configure VRRP for a **VPN** template, do the following:

- From the Cisco vManage menu, choose **Configuration > Templates**.
- Click **Feature Templates**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

- Navigate to the **VPN Interface Ethernet** template for the device.



Note For information about creating a new **VPN Interface Ethernet** template, see [Configure VPN Ethernet Interface](#).

- Click **VRRP** and choose **IPv4**.
- Click **New VRRP** to create a new VRRP or edit the existing VRRP and configure the following parameters:

Parameter Name	Description
TLOC Preference Change	(Optional) Choose On or Off to set whether the TLOC preference can be changed or not.

6. Click the **Add Tracking Object** link, and in the **Tracking Object** dialog box that is displayed, click **Add Tracking Object**.
7. In the **Tracker Name** field, enter the name of the tracker.
8. From the **Action** drop-down list, choose **Decrement** and enter the **Decrement Value**.
9. Click **Add**.
10. Click **Add** to save the VRRP details.
11. Click **Save**.

Configure VRRP Tracking Using CLI Templates

You can configure VRRP tracking using the CLI add-on feature templates and CLI device templates. For more information, see [CLI Templates](#).

VRRP Object Tracking Using CLI

Configure Track List Interface

Use the following configuration to add an interface to a track list using Cisco vManage device CLI template:

```
Device# config terminal
Device(config)# system
Device(config-system)# track-list zsl interface ge0/1 gre1 ipsec1
Device(config-track-list-zsl)# commit
Device(config-system-tracker-list-zsl)# exit
Device(config-system)# exit
```

Configure Interface Tracking and Priority Decrement

```
Device(config)# vpn 1
Device(config-vpn-1)# name vpn-name
Device(config-vpn-1)# interface ge0/2
Device(config-interface-ge0/2)# ip address 172.16.10.1/24
Device(config-interface-ge0/2)# no shutdown
Device(config-interface-ge0/2)# vrrp 100
Device(config-vrrp-100)# track zsl decrement 10
Device(config-vrrp-track-zsl)# exit
Device(config-vrrp-100)# ipv4 172.16.10.100
Device(config-vrrp-100)# tloc-change-pref
```

SIG Container Tracking

The following example shows how to configure a track list and tracking for SIG containers using the Cisco vManage device CLI template.



Note In SIG Object Tracking, you can only set *global* as the variable for Service Name.

Configure Track List for SIG Container

```
Device# config terminal
Device(config)# system
Device(config-system)# track-list SIG sig-container global
Device(config-system-tracker-list-SIG)# exit
Device(config-system)# exit
```

Configure SIG Container Tracking and Priority Decrement

```
Device(config)# vpn 1
Device(config-vpn-1)# name vpn-name
Device(config-vpn-1)# interface ge0/2
Device(config-interface-ge0/2)# ip address 172.16.10.1/24
Device(config-interface-ge0/2)# no shutdown
Device(config-interface-ge0/2)# vrrp 100
Device(config-vrrp-100)# track SIG decrement 10
Device(config-vrrp-track-zs1)# exit
Device(config-vrrp-100)# ipv4 172.16.10.100
Device(config-vrrp-100)# tloc-change-pref
```

Configure SIG Container Tracking for VRRP Group

```
Device(config-vpn-1)# int ge0/4
Device(config-interface-ge0/4)# vrrp 10
Device(config-vrrp-10)# track SIG decrement 10
Device(config-track-SIG)# commit
Commit complete.
Device(config-track-SIG)#
```

Configuration Example for VRRP Object Tracking Using CLI

Interface Object Tracking Using CLI

This example shows how to addan interface to a track list using Cisco vManage device CLI template:

```
Configure terminal
 system
 track-list zs1 interface ge0/1 gre1 ipsec1
 commit
 exit
```

Configure Interface Tracking and Priority Decrement

```
vpn 1
name vpn-name
interface ge0/2
ip address 172.16.10.1/24
no shutdown
vrrp 100
track zs1 decrement 10
exit
ipv4 172.16.10.100
tloc-change-pref
```


Configuration Examples for SIG Object Tracking

Configure Track List for SIG Container

```
config terminal
system
track-list SIG sig-container global
exit
exit
```

Configure SIG Container Tracking and Priority Decrement

```
vpn 1
name vpn-name
interface ge0/2
ip address 172.16.10.1/24
no shutdown
vrrp 100
track SIG decrement 10
exit
ipv4 172.16.10.100
tloc-change-pref
```

Verify VRRP Tracking

Device# show vrrp

The following is a sample output for the **show vrrp** command:

```
vrrp vpn 1
interfaces ge0/4
groups 10
virtual-ip          10.1.1.2
virtual-mac        00:00:5e:00:01:0a
priority           100
real-priority      100
vrrp-state         init
omp-state          up
advertisement-timer 1
primary-down-timer 3
last-state-change-time 0000-00-00T00:00:00+00:00
```

Device# show vrrp detail

The following is a sample output for the **show vrrp detail** command:

```
OMP status: up

group-id: 10, track-omp: no, initialized: yes
address: 10.20.24.1
track-prefix-list: -, resolved: -
state: Primary, down-reason: none, cfg-priority: 100, priority: 100
adv-timer: 1, primary-down-timer: 3, sock-fd: 23, addr-count: 1
adv-timer: Enabled (e: 4 v: 10 c: 1)
primary-down-timer: Disabled (e: -1 v: 30 c: 3)
virtual-mac: 0x0 0x0 0x5e 0x0 0x1 0xa
TLOC Change Preference: Configured
TLOC Change Preference value: 1
TLOC Real Preference value: 1
```

```

Group current adaptive priority: 0
Total Tracking object : 1 (head: 0x7f0f6d6771c0)
Group Address: 0x7f0f6d624100
  Name: zs1
  Decrement: 18
  Adaptive direction: 0
  List Entry :0x7f0f6d687230

```

Track List:

```

  Name: zs1
  Total Tracking Objects: 0
  VRRP Daemon: 0x7f0f6d68e140
  Tracking Object: 0x7f0f6d677270
  Type: 1
  VRRP Daemon: 0x7f0f6d68e140
  Total Interface: 1
    Interface: ge0_1(0x7f0f6d66a700)
  Interface Created: Yes
  Operational State: UP

```

Device# show run system

The following is a sample output for the **show run system** command:

```

system
host-name                vm6
system-ip                172.16.255.16
site-id                 600
no admin-tech-on-failure
route-consistency-check
organization-name       "vIPtela Inc Regression"
track-list SIG
  container global
!
track-list zs1
  track-interface ge0/1 ge0/7
!

```



CHAPTER 13

Configure a Cisco vEdge Device as an NTP Parent

Table 106: Feature History

Feature Name	Release Information	Feature Description
Configure a Cisco vEdge Device as an NTP Parent and Optionally to Support NTP in Symmetric Active Mode.	Cisco SD-WAN Release 20.4.1 Cisco vManage Release 20.4.1	This feature enables configuring a Cisco vEdge device as an NTP parent and configuring the device to support NTP in symmetric active mode.

You can configure a Cisco vEdge device as an NTP parent. You also can configure an NTP parent device to support NTP in symmetric active mode.

- [Configure an NTP Parent, on page 341](#)
- [Configure Support for NTP in Symmetric Active Mode, on page 342](#)

Configure an NTP Parent

Starting with Cisco SD-WAN Release 20.4.1, you can configure a supported Cisco vEdge device as an NTP parent device by using the device CLI template. A device that is configured in this way acts as the NTP server to which other nodes in the deployment synchronize their clocks. You can configure multiple devices as NTP parents. The NTP server functionality is supported for IPv4, but not for IPv6.

You also can configure a device that is configured as an NTP parent device to support NTP in symmetric active mode. See "Configure Support for NTP in Symmetric Active Mode."

Use the following commands to configure device as an NTP parent device using a Cisco vEdge device device CLI template. For more information about configuring device CLI template, see "Create a Device CLI Template" in *Systems and Interfaces Configuration Guide*.

```
Device# config terminal
Device(config)# system
Device(config-system) ntp
Device(config-ntp) # parent
Device(config-parent) # enable
Device(config-parent) # source-interface loopback511
Device(config-parent) # stratum 6
Device(config-parent) # vpn 511
Device(config-parent) # exit
```

Restrictions and Limitations

- You can configure a device as an NTP parent only through a Cisco vManage CLI template. Cisco vManage feature templates do not support this configuration.
- The source interface must be in the same VPN that the vpn keyword defines.

Verify Configuration

Use the following show command to verify NTP parent configuration. The sample output shows that the server also is configured to support NTP in symmetric active mode.

```
Device# show running-config system ntp

system
 ntp
  keys
  authentication 101 md5 $8$vV6PtHeLdiEcLqDNLqV/mCWN5X92yT8PUPowDCQgS4c=
  authentication 108 md5 $8$NTzFC6sRZiFUyeHw/pOY2dEoiO6dxphecDs7YnRKeuY=
  trusted 101 108
!
parent
 enable
 stratum 6
 source-interface loopback511
 vpn 511
 exit
server 10.20.25.1
 source-interface ge0/1
 vpn 511
 version 4
 exit
peer 172.16.10.100
 key 101
 vpn 511
 version 4
 source-interface ge0/1
 exit
```

Configure Support for NTP in Symmetric Active Mode

Starting with Cisco SD-WAN Release 20.4.1, you can configure a Cisco vEdge device that is configured as an NTP parent to support NTP in symmetric active mode by using the device CLI template. When a device is configured in this way, it synchronizes its time with another device that is defined with this mode if it cannot reach its original NTP parent.

Use the following commands to configure a device to support NTP in symmetric active mode by using a Cisco vManage device CLI template. For more information about configuring device CLI template, see "Create a Device CLI Template" in *Systems and Interfaces Configuration Guide*.

```
Device# config terminal
Device# system
Device(config-system) ntp
Device(config-ntp)# peer 172.16.10.1
Device(config-peer)# key 101
Device(config-peer)# vpn 511
Device(config-peer)# version 4
Device(config-parent)# source-interface ge0/1
Device(config-parent)# exit
```

Restrictions and Limitations

- You can configure a device support NTP in symmetric active mode only through a Cisco vManage CLI template. Cisco vManage feature templates do not support this configuration.
- You can configure up to two devices to support NTP in symmetric active mode.
- A device that is configured as an NTP peer should also be configured as an NTP parent.
- The source interface must be in the same VPN that the vpn keyword defines.
- Each peer must use the same source interface.

Use the following show command to verify NTP parent functional configuration. In the **show ntp peer** command output, the server with the .LOCL. REFID is the NTP parent.

```
Device# show ntp peer
```

INDEX	REMOTE	REFID	ST	TYPE	WHEN	POLL	REACH	DELAY	OFFSET	JITTER
1	*10.20.25.1	.GNSS.	1	u	186	1024	377	226.712	0.793	2.381
2	172.16.10.1	(loop)	3	s	760	1024	376	0.126	-1.307	1.397
3	172.16.10.10	.LOCL.	6	l	52h	64	0	0.000	0.000	0.000



CHAPTER 14

Dynamic On-Demand Tunnels

Table 107: Feature History

Feature Name	Release Information	Description
Dynamic On-Demand Tunnels	Cisco SD-WAN Release 20.3.1 Cisco vManage Release 20.3.1	This feature enables you to configure an Inactive state for tunnels between edge devices, reducing performance demands on devices and reducing network traffic.

Cisco SD-WAN supports dynamic on-demand tunnels between any two Cisco SD-WAN spoke devices. These tunnels are triggered to be set up only when there is traffic between the two devices. After the flow of traffic between the devices stops, a user-configurable inactivity timer starts, and after the configured time, the tunnel between the devices is removed. The on-demand link between the two devices is then considered to be inactive. In this inactive state, it does not use network bandwidth and does not affect device performance.

Backup Route and Reactivating the Tunnel

To enable two spoke device peers to use on-demand tunnels, they must have an alternate route, a backup route, through a hub. Using the backup route, either spoke device can resume traffic flow between the two spokes, which reactivates the tunnel to handle the traffic directly from peer to peer.

Advantages

On-demand tunnels offer the following advantages:

- Improved performance, especially for less-powerful platforms operating in a full-mesh network.
 - Improved latency in hub-and-spoke deployments when on-demand tunnels are used between spokes.
 - Reduced bandwidth use in the network because tunnels in Inactive state do not require Bidirectional Forwarding Detection (BFD) probes, so there is less BFD traffic produced in the network.
 - Direct tunnels between spokes, while also optimizing CPU and memory usage.
- [On-Demand Tunnel Mechanism in Detail, on page 346](#)
 - [Notes and Limitations, on page 347](#)
 - [Configure On-Demand Tunnels, on page 348](#)

On-Demand Tunnel Mechanism in Detail

When you configure a site to use dynamic tunnels, the on-demand functionality is enabled. In this mode of operation, Cisco SD-WAN edge routers do not bring up direct tunnels to other sites that are also enabled with on-demand functionality.

Cisco SD-WAN selects one or more edge routers (typically centrally located routers) to act as backup forwarding node(s), providing a secondary path for traffic between two nodes. The backup node(s) are not enabled for on-demand. All on-demand sites form static tunnels with the backup node(s). The backup node(s) provide a static backup route for traffic between two nodes that have on-demand enabled.

The first packet of traffic between two nodes is routed through the static backup path, and triggers the on-demand tunnel to become active between the sites. The backup path continues to forward traffic until the direct path becomes active.

All on-demand sites learn the TLOCs and prefixes of all other on-demand remote sites. The prefixes also have a backup path set up through Cisco vSmart Controller control policy. So in the control plane, the on-demand tunnel network has the same state as a full-mesh tunnel network, including a backup path. The control plane downloads to the data plane, routes, with the backup path and remote TLOCs that represent a potential direct path between any two sites, but it does not set up a direct path tunnel to remote TLOCs.

Traffic from either end of the on-demand tunnel triggers setting up the tunnel. This enables on-demand tunnels to accommodate network address translation (NAT) traversal.

The on-demand tunnel feature introduces two states for the on-demand branch site:


- **Inactive:** The on-demand tunnel is not set up with the remote site. There is no active traffic to or from the remote site. Remote site TLOCs are inactive - no bidirectional forwarding detection (BFD) is set up, the prefix is installed with the inactive paths, and the backup path is set as the path to forward any traffic. The inactive path detects flows and triggers a direct site-to-site tunnel to be set up.
- **Active:** The on-demand direct site-to-site tunnel is set up to the remote site. There is active traffic to or from the remote site. This state is identical to the case of a typical tunnel, where the remote TLOCs have BFD set up, and the prefix is installed with the direct path tunnel. In this state, tunnel activity is tracked. If there is no traffic for the “idle time” duration (default 10 minutes), the direct site-to-site tunnel is removed and the state changes to Inactive.

Steps in Illustrations

The figures below show the following steps that occur between two edge routers with an on-demand tunnel configured.

1. There is no active tunnel between the two edge routers. edge-1 and edge-4 are in Inactive state.
2. The host behind edge-1 initiates traffic toward the host behind edge-4.
3. edge-1 forwards the traffic through the backup path using the hub or backup node to edge-4.
4. edge-1 provisions the on-demand tunnel and begins bidirectional forwarding detection (BFD). edge-4 is now in Active state on edge-1.
5. When edge-4 receives the return traffic for the host behind edge-1, it forwards the traffic through the backup path using the hub or backup node to edge-1.
6. edge-4 provisions the on-demand tunnel and begins BFD. edge-1 is now in Active state on edge-4.

7. At this point, the on-demand tunnel between edge-1 and edge-4 is up, and BFD is up.
8. Traffic between the two edge devices takes the direct route through the on-demand tunnel.
9. Both edge-1 and edge-4 track the traffic activity on the on-demand tunnel in both directions.
If there is no traffic for the idle timeout duration, the on-demand tunnel is deleted, and the edge-1 and edge-4 devices go back to the Inactive state.



520715

520716

Notes and Limitations

- On-demand tunnel Performance Routing (PfR) statistics collection starts fresh every time an on-demand tunnel is setup. PFR statistics are not cached for deleted on-demand tunnels after idle timeout.
- Out of order (OOO) packets may occur when traffic moves from the backup path to the direct on-demand tunnel. Packets are forwarded by the Cisco SD-WAN router as they are received.

- Unidirectional flows do not trigger on-demand tunnel setup. They continue to use the backup path.
- Multicast traffic does not trigger on-demand tunnel setup. It continues to use the backup path.
- Do not configure a data policy that applies a **set floc-list** action to an on-demand site TLOC. If configured, traffic will be dropped.
- On-demand tunnels are not supported when the Pair Wise Key (PWK) IPSEC feature is enabled.
- All TLOCs in the system will be reset (disabled and enabled) when **on-demand enable** or **no on-demand enable** is executed.
- When an edge device provisions on-demand tunnels, it provisions to all the TLOCs on the remote site.
- For a multi-home site to be in on-demand mode, you must configure on-demand enable on all of the systems at the site.
- All edge devices using on-demand tunnels are kept active if there is a service or user traffic on any on-demand tunnel in either direction.
- On-demand tunnels can be enabled between two sites only if both sites are enabled with on-demand mode.
- The first packet to any host behind a remote site triggers on-demand tunnel setup to that remote site. Return traffic from that host triggers tunnel setup in the opposite direction.
- All prefixes from on-demand remote sites must also have a backup path configured. If not, sites will not be able set up on-demand tunnels. The backup path is a static tunnel and must be always UP.
- The setup or removal of on-demand tunnels does not affect overlay route (OMP) updates by Cisco vSmart Controller, or service/LAN-side route updates (examples: OSPF or BGP).
- If either the local site or the remote site is not in on-demand mode, static tunnels are set up between the sites.

Configure On-Demand Tunnels

Prerequisites for On-Demand Tunnels

There are several prerequisites for using on-demand tunnels:

- [Prerequisites: Cisco vSmart Controller Centralized Control Policy, on page 348](#)
- [Prerequisites: OMP Settings, on page 350](#)
- [Prerequisites: Hub Device, on page 350](#)
- [Prerequisites: Spoke Devices, on page 351](#)

Prerequisites: Cisco vSmart Controller Centralized Control Policy

1. The Cisco vSmart Controller centralized control policy must include the **floc-action backup** action.

Explanation: This ensures that the backup path through the hub for communication between all of the spoke devices.

2. The Cisco vSmart Controller centralized control policy must accept all spoke prefix routes.
3. The Cisco vSmart Controller centralized control policy must accept TLOCs of all spokes.

For information about configuring a Cisco vSmart Controller **centralized control policy**, see the Policies configuration guides on the [Cisco SD-WAN Configuration Guides page](#).

CLI Example, Centralized Control Policy Addressing Prerequisites

```

viptela-policy:policy
control-policy Dynamic-Tunnel-Control-Policy
  sequence 100
  match route
    site-list Branches
  !
  action accept
  set
    tloc-action backup
    tloc-list Hub-TLOCs
  !
  !
  sequence 200
  match tloc
  !
  action accept
  !
default-action accept
!
lists
site-list Branches
  site-id 200
  site-id 300
!
tloc-list Hub-TLOCs
  tloc 10.0.0.1 color mpls encap ipsec
  tloc 10.0.0.1 color public-internet encap ipsec
!
!
apply-policy
  site-list Branches
  control-policy Dynamic-Tunnel-Control-Policy out
!
!

```

Cisco vManage Procedure

1. From the Cisco vManage menu, choose **Configuration > Policies**.
2. Select **Centralized Policy**.
3. Click **Add Topology** and select **Custom Control (Route & TLOC)**.
4. From **Match Conditions**, in **Site**, select one or more site lists, and click **Accept**.
5. From **Actions**, in **TLOC Action**, select the **Backup** action.
6. From **TLOC List**, select an existing TLOC list or create a new one.

Prerequisites: OMP Settings

1. The Cisco vSmart Controller `send-path-limit` must be more than the default 4.

Explanation: When on-demand tunnels are enabled, spokes use backup paths through the hub, so a higher path limit is necessary. The direct paths as well as the backup paths need to be advertised. To accommodate this, increase the Cisco vSmart Controller `send-path-limit` to advertise all available paths. We recommend to use the maximum possible value.



Note If there are too many Hub TLOCs configured in the On-Demand Tunnel control policy, the recommended value for **send-path-limit** is not enough always. In such cases, the On-Demand Tunnel feature will not work at all.

Starting from Cisco vManage Release 20.8.1 and Cisco IOS XE Release 17.8.1a, the maximum **send-path-limit** is 32. In Cisco vManage Release 20.7.x and earlier releases, the maximum **send-path-limit** is 16.

For information about configuring the vSmart **send-path-limit**, see the Routing Configuration guides on the [Cisco SD-WAN Configuration Guides page](#).

CLI Example

```
omp
no shutdown
send-path-limit 16
graceful-restart
```

Cisco vManage Procedure

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Feature Templates**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

3. Click **Add Template**.
4. Select a device and click **Cisco OMP**.
5. In **Basic Configuration**, set the **Number of Paths Advertised per Prefix** to 16 (recommended).

Prerequisites: Hub Device

1. On the hub device, the Traffic Engineering service (service TE) must be enabled.

Explanation: This ensures that the Cisco SD-WAN Overlay Management Protocol (OMP) on the spoke devices accepts the backup path through the hub, which is being added as an intermediate path between the two spoke devices. Without this, the backup path through the hub would be considered invalid and unresolved by the spoke devices.

CLI Example (Cisco vEdge Devices)

```
vpn 0
  service TE
exit
```

CLI Example (Cisco IOS XE SD-WAN Devices)

```
sdwan
  service TE vrf global
exit
```

Cisco vManage Procedure

1. In Cisco vManage, open **Configuration > Templates**.
2. Click **Feature Templates**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

3. Click **Add Template**.
4. Select a platform.
5. From **VPN**, select **VPN**.
6. Ensure that in **Basic Configuration**, the **VPN** field is set to 0.
7. From **Service**, click **New Service** and select **TE**.
8. Click **Add**, and then click **Update**. The TE service appears in the table of services.
9. Apply the VPN-0 template to the hub.

Prerequisites: Spoke Devices

1. On spoke devices, the `ecmp-limit` must be more than the default 4. Recommended: 16

Explanation: When on-demand tunnels are enabled, spoke devices create both direct and backup paths. To accommodate the need for more paths, increase the `ecmp-limit`.

CLI Example

```
omp
  no shutdown
  ecmp-limit 16
```



Note You can view the current `ecmp-limit` using the **show running-config omp** command.

Cisco vManage Procedure

1. From the Cisco vManage menu, choose **Configuration > Templates**.

2. Click **Feature Templates**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

3. Click **Add Template**.
4. Select a device and click **Cisco OMP**.
5. In **Basic Configuration**, set the **ECMP Limit** field to 16 (recommended).

Configure On-Demand Tunnels Using Cisco vManage



-
- Note**
- See the [Prerequisites for On-Demand Tunnels](#).
 - Do not enable on-demand on the hub device.
-

On the spoke devices, enable on-demand at the system level on all VPN-0 transport interfaces. In the case of multi-homed sites, enable on-demand on all systems in the site.

1. From the Cisco vManage menu, choose **Configuration > Templates**.
2. Click **Feature Templates**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

3. Click **Add Template**.
4. Select a device.
5. From **Basic Information**, select **Cisco System**.
6. Click **Advanced**.
7. Enable **On-demand Tunnel**.
8. (optional) Configure the **On-demand Tunnel Idle Timeout** time. The default idle timeout value is 10 minutes. Range: 1 to 65535 minutes
9. Attach the System feature template to the device template for the spoke device.

Configure On-Demand Tunnels Using the CLI

**Note**

- See [Prerequisites for On-Demand Tunnels, on page 348](#).
- Do not enable on-demand on the hub device

1. On the spoke devices, enable on-demand tunnels at the system level. In the case of multi-homed sites, enable on-demand on all systems in the site.

The default idle timeout value is 10 minutes. Range: 1 to 65535 minutes

Example

```
system
  on-demand enable
  on-demand idle-timeout 10
```

View Current Status of On-Demand Tunnels in Cisco vManage

1. From the Cisco vManage menu, choose **Monitor > Devices**.
Cisco vManage Release 20.6.x and earlier: From the Cisco vManage menu, choose **Monitor > Network**.
2. Select a device.
3. Select **Real Time**.
4. For **Device Options**, select one of the following:
 - **On Demand Local**: Displays the status of on-demand tunnels on the specified device.
 - **On Demand Remote**: Displays the status of on-demand tunnels on the specified device, and on all connected devices.

The output is equivalent to executing the `show [sdwan] system on-demand [remote-system] [system-ip ip-address]` CLI command. It displays the status of on-demand tunnels.

View Chart of On-Demand Tunnel Status Over Time in Cisco vManage

1. From the Cisco vManage menu, choose **Monitor > Devices**.
Cisco vManage Release 20.6.x and earlier: From the Cisco vManage menu, choose **Monitor > Network**.
2. Select a device.
3. From **WAN**, choose **Tunnel**.
4. From the **Chart Options** drop-down list, select **On-Demand Tunnel Status**. The chart shows the status of tunnels as ACTIVE or INACTIVE. INACTIVE indicates that an on-demand tunnel is in Inactive mode.

For more information, see the Monitor and Maintain guide on the [Cisco SD-WAN Configuration Guides page](#).



CHAPTER 15

Cisco SD-WAN Multitenancy

- [Overview of Cisco SD-WAN Multitenancy, on page 355](#)
- [Supported Devices and Controller Specifications, on page 359](#)
- [Restrictions, on page 360](#)
- [Initial Setup for Multitenancy, on page 361](#)
- [Expand a Multitenant Deployment to Support More Tenants and Tenant Devices, on page 369](#)
- [Manage Tenants, on page 371](#)
- [Cisco vManage Dashboard for Multitenancy, on page 376](#)
- [Manage Tenant WAN Edge Devices, on page 380](#)
- [Tenant-Specific Policies on Cisco vSmart Controllers, on page 381](#)
- [Manage Tenant Data, on page 382](#)
- [View OMP Statistics per Tenant on a Cisco vSmart Controller, on page 385](#)
- [View Tenants Associated with a Cisco vSmart Controller, on page 386](#)
- [Migrate Single-Tenant Cisco SD-WAN Overlay to Multitenant Cisco SD-WAN Deployment, on page 386](#)
- [Migrate Multitenant Cisco SD-WAN Overlay, on page 390](#)
- [Upgrade Cisco SD-WAN Controller and Edge Device Software, on page 392](#)
- [Multitenant Cisco vManage: Disaster Recovery, on page 393](#)
- [Multitenant Cisco vManage: Disaster Recovery in an Overlay Network with Virtual Routers, on page 398](#)
- [Multitenant Cisco vManage: Disaster Recovery After a Failed Data Center Becomes Operational, on page 404](#)
- [Replace Faulty Cisco vSmart Controller, on page 408](#)

Overview of Cisco SD-WAN Multitenancy

With Cisco SD-WAN multitenancy, a service provider can manage multiple customers, called tenants, from Cisco vManage. The tenants share the same set of underlying Cisco SD-WAN controllers: Cisco vManage, Cisco vBond Orchestrator, and Cisco vSmart Controller. The tenant data is logically isolated on these shared controllers.

The service provider accesses Cisco vManage using a domain name mapped to the IP address of a Cisco vManage cluster and manages the multitenant deployment. Each tenant is provided a subdomain to access a tenant-specific Cisco vManage view and manage the tenant deployment. For example, a service provider using the domain name `managed-sp.com`, can assign tenants Customer1 and Customer2 the subdomains

customer1.managed-sp.com and customer2.managed-sp.com and manage them on the same set of Cisco SD-WAN controllers, instead of providing each customer a single-tenant setup with a dedicated set of Cisco SD-WAN controllers.

Following are the key features of Cisco SD-WAN multitenancy:

- Full enterprise multitenancy: Cisco SD-WAN supports multitenancy and offers enterprises the flexibility of segregated roles such as service provider and tenants. Service providers can use multitenancy to provide Cisco SD-WAN service offerings to their customers.
- Multi-tenant Cisco vManage
- Multi-tenant Cisco vBond Orchestrators
- Multi-tenant Cisco vSmart Controllers
- Tenant-specific WAN Edge Devices
- Overlapping VPN numbers: A particular VPN or a set of common VPNs is assigned to a specific tenant, with their own configurations and monitoring dashboard environment. These VPN numbers can overlap where they are used by other tenants.
- On-prem and cloud deployment models: Cisco SD-WAN controllers can be deployed in an organization data center on servers running the VMware ESXi 6.7 or later, or the Kernel-based Virtual Machine (KVM) hypervisor. Cisco SD-WAN controllers can also be hosted on Amazon Web Services (AWS) servers by Cisco CloudOps.
- Tenant-specific Cisco vAnalytics: Cisco vAnalytics is a cloud-based service that offers insights into the performance of applications and the underlying SD-WAN network infrastructure. Each tenant can obtain Cisco vAnalytics insights for their overlay network by requesting a tenant-specific Cisco vAnalytics instance and enabling data collection on Cisco vManage. The service provider must enable cloud services on Cisco vManage in the provider view to facilitate the onboarding of the Cisco vAnalytics instance for the tenant overlay network.

Multi-tenant Cisco vManage

Cisco vManage is deployed and configured by the service provider. The provider enables multitenancy and creates a Cisco vManage cluster to serve tenants. Only the provider can access a Cisco vManage instance through the SSH terminal.

Cisco vManage offers service providers an overall view of the SD-WAN multi-tenant deployment and allows a provider to manage the shared Cisco vBond Orchestrator and Cisco vSmart Controller devices. Cisco vManage also allows service providers to monitor and manage the deployments of each tenant.

Cisco vManage allows tenants to monitor and manage their deployment. Through Cisco vManage, tenants can deploy and configure WAN edge devices. Tenants can also configure custom policies on assigned Cisco vSmart Controllers.

Multi-tenant Cisco vBond Orchestrators

Cisco vBond Orchestrators are deployed and configured by the service provider. Only the provider can access a Cisco vBond Orchestrator through the SSH terminal.

Cisco vBond Orchestrators serve WAN edge devices of multiple tenants as the devices are added to the overlay network.

Multi-tenant Cisco vSmart Controllers

Cisco vSmart Controllers are deployed by the service provider. Only the provider can create and attach device and feature templates to Cisco vSmart Controllers, and can access a Cisco vSmart Controller through the SSH terminal.

- When a tenant is created, Cisco vManage assigns two Cisco vSmart Controllers for the tenant. The Cisco vSmart Controllers form an active-active cluster.

Each tenant is assigned only two Cisco vSmart Controllers. Before a tenant is created, two Cisco vSmart Controllers must be available to serve the tenant.

- When more than one pair of Cisco vSmart Controllers are available to serve a tenant, Cisco vManage assigns to the tenant the pair of Cisco vSmart Controllers connected to the lowest number of forecast devices. If two pairs of Cisco vSmart Controllers are connected to the same number of devices, Cisco vManage assigns to the tenant the pair of Cisco vSmart Controllers serving the lowest number of tenants.
- From Cisco vManage Release 20.9.1, while onboarding a tenant to a multitenant deployment, you can choose the pair of multitenant Cisco vSmart Controllers that serve the tenant. After onboarding a tenant, you can migrate the tenant to a different pair of multitenant Cisco vSmart Controllers, if necessary. For more information, see [Flexible Tenant Placement on Multitenant Cisco vSmart Controllers](#).
- Each pair of Cisco vSmart Controllers can serve a maximum of 24 tenants.
- Tenants can configure custom policies on the Cisco vSmart Controllers assigned to them. Cisco vManage notifies the Cisco vSmart Controllers to pull the policy templates. Cisco vSmart Controllers pull the templates and deploy the policy configuration for the specific tenant.
- Only the provider can view events, audit logs, and OMP alarms for a Cisco vSmart Controller on Cisco vManage.

Tenant-Specific WAN Edge Devices

A tenant or the provider acting on behalf of a tenant can add WAN edge devices to the tenant network, configure the devices, and remove the devices from the tenant network, or access the device through the SSH terminal.

A provider can manage the WAN edge devices only from [provider-as-tenant](#) view. In the [provider](#) view, Cisco vManage does not show any WAN edge device information.

Cisco vManage reports WAN edge device events, logs, and alarms only in the [Tenant Role](#) and the [provider-as-tenant](#) views.

User Roles in Multitenant Environment

A multi-tenant environment includes the service provider and tenant roles. Each role has distinct privileges, views, and functions.

Provider Role

The provider role entitles system-wide administrative privileges. A user with the provider role has the default username **admin**. The provider user can access Cisco vManage using the domain name of the service provider or by using the Cisco vManage IP address. When using a domain name, the domain name has the format `https://managed-sp.com`.

The **admin** user is part of the user group **netadmin**. Users in this group are permitted to perform all operations on the controllers and the WAN edge devices of the tenants. You can add additional users to the **netadmin** group.

You cannot modify the privileges of the **netadmin** group. On Cisco vManage, you can view the privileges of the user group from the **Administration > Manage Users > User Groups** page.



Note When you create a new provider user in Cisco vManage, including a **netadmin** user, by default, the user is not allowed SSH access to the Cisco vManage VM. To enable SSH access, configure SSH authentication using a AAA template and push the template to Cisco vManage. For more information on enabling SSH authentication, see [SSH Authentication using vManage on Cisco vEdge Devices](#).

For more information about configuring users and user groups, see [Configure User Access and Authentication](#).

Cisco vManage offers two views to a provider:

- **Provider View**

When a provider user logs in to multi-tenant Cisco vManage as **admin** or another **netadmin** user, Cisco vManage presents the provider view and displays the provider dashboard.

You can perform the following functions from the provider view:

- Provision and manage Cisco vManage, Cisco vBond Orchestrators and Cisco vSmart Controllers.
- Add, modify, or delete tenants.
- Monitor the overlay network.

- **Provider-as-Tenant View**

When a provider user selects a specific tenant from the **Select Tenant** drop-down list at the top of the provider dashboard, Cisco vManage presents the provider-as-tenant view and displays the tenant dashboard for the selected tenant. The provider user has the same view of Cisco vManage as a tenant user would when logged in as **tenantadmin**. From this view, the provider can manage the tenant deployment on behalf of the tenant.

In the provider dashboard, a table of tenants presents a status summary for each tenant. A provider user can also launch the provider-as-tenant view by clicking on a tenant name in this table.

Tenant Role

The tenant role entitles tenant administrative privileges. A user with the tenant role has the default username **tenantadmin**. The default password is **Cisco#123@Viptela**. We recommend that you change the default password on first login. For information on changing the default password, see [Hardware and Software Installation](#).

The **tenantadmin** user is part of the user group **tenantadmin**. Users in this group are permitted to perform all operations on the WAN edge devices of the tenants. You can add additional users to the **tenantadmin** group.

You cannot modify the privileges of the **tenantadmin** group. On Cisco vManage, you can view the privileges of the user group from the **Administration > Manage Users > User Groups** page.

For more information about configuring users and user groups, see [Configure User Access and Authentication](#).

A tenant user can log in to Cisco vManage using a dedicated URL and the default username **tenantadmin**. For example, the dedicated URL of a tenant could be `https://customer1.managed-sp.com` for a provider using the domain name `https://managed-sp.com`. When the user logs in, Cisco vManage presents the tenant view and displays the tenant dashboard.



Tip If you cannot access the dedicated tenant URL, update the subdomain details in the `/etc/hosts` file on the local machine. Alternatively, if you use an external DNS server, add a DNS entry for the tenant subdomain.

A tenant user with administrative privileges can perform the following functions:

- Provision and manage tenant routers
- Monitor overlay network of the tenant
- Create custom policies on the assigned Cisco vSmart Controllers
- Upgrade the software on the tenant routers.

Supported Devices and Controller Specifications

The following Cisco SD-WAN edge devices support multitenancy.

Table 108: Supported Devices

Platform	Device Models
Cisco vEdge device	<ul style="list-style-type: none"> • vEdge 100, vEdge 100b, vEdge 100m, vEdge 100wm, vEdge 1000, vEdge 2000, vEdge 5000, vEdge Cloud • ISR1100-6G/ISR1100-4G, ISR1100-4GLTENA, ISR1100-4GLTEGB

The following hypervisors are supported for multitenancy:

- VMware ESXi 6.7 or later
- KVM
- AWS (cloud-hosted and managed by Cisco CloudOps)
- Microsoft Azure (cloud-hosted and managed by Cisco CloudOps)

From Cisco vManage Release 20.6.1, a multitenant Cisco vManage instance can have one of the following three personas. The personas enable a predefined set of services on the Cisco vManage instance.

Table 109: Cisco vManage Personas

Persona	Services
Compute+Data	Cluster Oracle, Service Proxy, Messaging Service, Coordination Service, Configuration Database, Data Collection Agent, Statistics Database, and Application Server
Data	Cluster Oracle, Service Proxy, Application Server, Data Collection Agent, and Statistics Database
Compute	Cluster Oracle, Service Proxy, Messaging Service, Coordination Service, Configuration Database, and Application Server

The supported hardware specifications for the Cisco vBond Orchestrator, Cisco vManage, and the Cisco vSmart Controllers are as follows:

Hardware Specifications to Support 50 Tenants and 1000 Devices

For more information on supported hardware specifications for the Cisco vBond Orchestrator, Cisco vManage, and the Cisco vSmart Controllers see, [Cisco SD-WAN Controller Compatibility Matrix and Recommended Computing Resources](#).

Hardware Specifications to Support 75 Tenants and 2500 Devices

For more information on supported hardware specifications for the Cisco vBond Orchestrator, Cisco vManage, and the Cisco vSmart Controllers see, [Cisco SD-WAN Controller Compatibility Matrix and Recommended Computing Resources](#).

Hardware Specifications to Support 100 Tenants and 5000 Devices

For more information on supported hardware specifications for the Cisco vBond Orchestrator, Cisco vManage, and the Cisco vSmart Controllers see, [Cisco SD-WAN Controller Compatibility Matrix and Recommended Computing Resources](#).

Hardware Specifications to Support 150 Tenants and 7500 Devices

For more information on supported hardware specifications for the Cisco vBond Orchestrator, Cisco vManage, and the Cisco vSmart Controllers see, [Cisco SD-WAN Controller Compatibility Matrix and Recommended Computing Resources](#).

Restrictions

- Do not use a user-configured system IP address to connect to a device through SSH. Instead, use the IP address of the `vmanage_system` interface; this IP address is assigned by Cisco vManage.

To find the IP address of the `vmanage_system` interface, use one of the following methods:

- Launch the device SSH terminal from Cisco vManage and find the `vmanage_system` IP address from the first line of the log-in prompt.

- Run the **show interface description** command and find the `vmanage_system` IP address from the command output.
- If you add a second tenant immediately after adding a tenant, Cisco vManage adds them sequentially, and not in parallel.
- If you are adding a WAN edge device that you had previously invalidated and deleted from an overlay network, you must reset the device software after adding the device. To reset the software on a Cisco vEdge device, use the command **request platform software reset**.

Initial Setup for Multitenancy

Prerequisites

- Download and install software versions as recommended in the following table:

Table 110: Minimum Software Prerequisites for Cisco SD-WAN Multitenancy

Device	Software Version
Cisco vManage	Cisco vManage Release 20.6.1
Cisco vBond Orchestrator	Cisco SD-WAN Release 20.6.1
Cisco vSmart Controller	Cisco SD-WAN Release 20.6.1
Cisco vEdge Device	Cisco SD-WAN Release 20.6.1

A configuration in which one or more controllers, or WAN edge devices, are running software versions earlier than those mentioned in the table above is not supported.

- Do not migrate an existing single-tenant Cisco vManage instance into multitenant mode, even if you invalidate or delete all devices from the existing Cisco vManage instance. Instead, download and install a new Cisco vManage software image.



Note After you enable Cisco vManage for multitenancy, you cannot migrate it back to single tenant mode.

- Follow the recommended hardware specifications in the *Supported Devices and Controller Specifications* section of this document.
 - Log in to Cisco vManage as the provider **admin** user.
1. Create Cisco vManage cluster.
 - a. To support 50 tenants and 1000 devices across all tenants, [Create a 3-Node Cisco vManage Cluster](#).
 - b. To support 100 tenants and 5000 devices across all tenants, [Create a 6-Node Cisco vManage Cluster](#).

- c. From Cisco SD-WAN Release 20.6.3, Cisco vManage Release 20.6.3, to support 150 tenants and 7500 devices across all tenants, [Create a 6-Node Cisco vManage Cluster](#).
2. Create and configure Cisco vBond Orchestrator instances. See [Deploy Cisco vBond Orchestrator](#).
While configuring Cisco vBond Orchestrator instances, configure the service provider organization name (`sp-organization-name`) and the organization name (`organization-name`). See [Configure Organization Name in Cisco vBond Orchestrator](#).

```
sp-organization-name multitenancy
organization-name multitenancy
```
 3. Create Cisco vSmart Controller instances. See [Deploy the Cisco vSmart Controller](#).
 - To support 50 tenants and 1000 devices across all tenants, deploy 6 Cisco vSmart Controller instances.
 - To support 100 tenants and 5000 devices across all tenants, deploy 10 Cisco vSmart Controllers.
 - From Cisco SD-WAN Release 20.6.3, Cisco vManage Release 20.6.3, to support 150 tenants and 7500 devices across all tenants, deploy 16 Cisco vSmart Controllers.
 - a. [Add Cisco vSmart Controller](#) to the overlay network.
 4. Onboard new tenants. See [Add a New Tenant, on page 372](#).

Create a 3-Node Cisco vManage Cluster

1. Download the Cisco vManage Release 20.6.1 or later software image from [Cisco Software Download](#).
2. Create three Cisco vManage instances (say vManage1, vManage2, and vManage3) by installing the downloaded software image file. See [Deploy Cisco vManage](#).



Important

- Deploy Cisco vManage servers having the hardware specifications in the table *Hardware Specifications to Support 50 Tenants and 1000 Devices* from the *Supported Devices and Controller Specifications* section of this document.
- Choose the **Compute+Data** persona for each Cisco vManage instance.

3. Complete the following operations on vManage1:
 - a. Configure the following using the CLI:
 - System IP address
 - Site ID
 - Service Provider organization name (`sp-organization-name`)
 - Organization-name
 - vBond IP address
 - VPN 0 Transport/Tunnel interface

- VPN 0 Out-of-band (OOB) interface: Ensure that you assign a static IP address to this interface. Do not enable DHCP.
- VPN 512 Management interface



Note Configure only one default route in VPN 0.

- [Enable Multitenancy on Cisco vManage, on page 367.](#)
- (Optional) Using the CLI, install the Root CA certificate for vManage1.



Note Skip this step if you are using a Symantec or Cisco PKI certificate.

- Complete the following through the Cisco vManage GUI:
 - [Generate a Certificate Signing Request](#)
 - After Symantec or your enterprise root CA has signed the certificate, [install the signed certificate](#).
- [Configure the Cluster IP Address of the Cisco vManage Server.](#)

Before proceeding to the next step, ensure that the **vManage IP Address** field on the **Administration > Cluster Management** page shows the OOB interface address.

- Complete the following operations on vManage2 and vManage 3:



Important Do not enable multitenancy on vManage2 and vManage3.

- Configure the following using the CLI:
 - System IP address
 - Site ID
 - Service Provider organization name (`sp-organization-name`)
 - Organization-name
 - vBond IP address
 - VPN 0 Transport/Tunnel interface
 - VPN 0 Out-of-band (OOB) interface: Ensure that you assign a static IP address to this interface. Do not enable DHCP.
 - VPN 512 Management interface
- (Optional) Using the CLI, install the Root CA certificate for vManage1.



Note Skip this step if you are using a Symantec or Cisco PKI certificate.

- c. Complete the following through the Cisco vManage GUI:
 1. [Generate a Certificate Signing Request](#)
 2. After Symantec or your enterprise root CA has signed the certificates, [install signed certificate](#).
- d. [Log in to the Cisco vManage Web Application Server](#).
- e. Ping the OOB interfaces on the other two Cisco vManage instances and ensure they are reachable.
- f. [Configure the Cluster IP Address of the Cisco vManage Server](#).

Before proceeding to the next step, ensure that the **vManage IP Address** field on the **Administration > Cluster Management** page shows the OOB interface address.

5. Log in to the vManage1 GUI and [add vManage2 to the cluster](#).

vManage2 reboots before being added to the cluster.

While vManage2 is being added to the cluster, on the **Administration > Cluster Management** page, the **Configure Status** for vManage2 shows **Pending**. You can monitor the System Generated Cluster Sync transaction to check the progress of the adding vManage2 to the cluster.

When the operation is completed, on the **Administration > Cluster Management** page, you can view both vManage1 and vManage2, and their node personas.

6. Repeat **Step 5** and add vManage3 to the cluster.



Note After rebooting, you have to select persona (non-cloud setup) from CLI and services starts running on the node according to the selected persona.

Create a 6-Node Cisco vManage Cluster

1. Download the Cisco vManage Release 20.6.1 or later software image from [Cisco Software Download](#).
2. Create six Cisco vManage instances by installing the downloaded software image file. See [Deploy Cisco vManage](#).

**Important**

- To support 100 tenants and 5000 devices across all tenants, deploy Cisco vManage servers having the hardware specifications in the table *Hardware Specifications to Support 100 Tenants and 5000 Devices* from the *Supported Devices and Controller Specifications* section of this document.

From Cisco SD-WAN Release 20.6.3, Cisco vManage Release 20.6.3, to support 150 tenants and 7500 devices across all tenants, deploy Cisco vManage servers having the hardware specifications in the table *Hardware Specifications to Support 150 Tenants and 7500 Devices* from the *Supported Devices and Controller Specifications* section of this document.

- Choose the **Compute+Data** persona for three Cisco vManage instances (say vManage1, vManage2, and vManage3). Choose the **Data** persona for the other three Cisco vManage instances (say vManage4, vManage5, and vManage6).

3. Complete the following operations on vManage1:**a. Configure the following using the CLI:**

- System IP address
- Site ID
- Service Provider organization name (sp-organization-name)
- Organization-name
- vBond IP address
- VPN 0 Transport/Tunnel interface
- VPN 0 Out-of-band (OOB) interface: Ensure that you assign a static IP address to this interface. Do not enable DHCP.
- VPN 512 Management interface



Note Configure only one default route in VPN 0.

- b.** [Enable Multitenancy on Cisco vManage, on page 367.](#)
- c.** (Optional) Using the CLI, install the Root CA certificate for vManage1.



Note Skip this step if you are using a Symantec or Cisco PKI certificate.

- d.** Complete the following through the Cisco vManage GUI:
 1. [Generate a Certificate Signing Request](#)
 2. After Symantec or your enterprise root CA has signed the certificate, [install the signed certificate.](#)
- e.** [Configure the Cluster IP Address of the Cisco vManage Server.](#)

Before proceeding to the next step, ensure that the **vManage IP Address** field on the **Administration > Cluster Management** page shows the OOB interface address.

4. Complete the following operations on vManage2 through vManage6:



Important Do not enable multitenancy on vManage2 through vManage6.

- a. Configure the following using the CLI:
- System IP address
 - Site ID
 - Service Provider organization name (`sp-organization-name`)
 - Organization-name
 - vBond IP address
 - VPN 0 Transport/Tunnel interface
 - VPN 0 Out-of-band (OOB) interface: Ensure that you assign a static IP address to this interface. Do not enable DHCP.
 - VPN 512 Management interface
- b. (Optional) Using the CLI, install the Root CA certificate for vManage1.



Note Skip this step if you are using a Symantec or Cisco PKI certificate.

- c. Complete the following through the Cisco vManage GUI:
1. [Generate a Certificate Signing Request](#)
 2. After Symantec or your enterprise root CA has signed the certificate, [install the signed certificate](#).
- d. [Log in to the Cisco vManage Web Application Server](#).
- e. Ping the OOB interfaces on the other Cisco vManage instances and ensure they are reachable.
- f. [Configure the Cluster IP Address of the Cisco vManage Server](#).

Before proceeding to the next step, ensure that the **vManage IP Address** field on the **Administration > Cluster Management** page shows the OOB interface address.

5. Log in to the vManage1 GUI and [add vManage2 to the cluster](#).

vManage2 reboots before being added to the cluster.

While vManage2 is being added to the cluster, on the **Administration > Cluster Management** page, the **Configure Status** for vManage2 shows **Pending**. You can monitor the System Generated Cluster Sync transaction to check the progress of the adding vManage2 to the cluster.

When the operation is completed, on the **Administration > Cluster Management** page, you can view both vManage1 and vManage2, and their node personas.

6. Repeat **Step 5** and add vManage3 through vManage6 to the cluster.

Enable Multitenancy on Cisco vManage

Prerequisites

Do not migrate an existing single-tenant Cisco vManage into multitenant mode, even if you invalidate or delete all devices from the existing Cisco vManage. Instead, download and install a new software image of Cisco vManage Release 20.6.1 or a later release.



Note After you enable multitenancy on Cisco vManage, you cannot migrate it back to single tenant mode.

1. Launch Cisco vManage using the URL `https://vmanage-ip-address:port`. Log in as the provider **admin** user.
2. From the Cisco vManage menu, choose **Administration > Settings**.
3. In the **Tenancy Mode** bar, click the **Edit**.
4. In the **Tenancy** field, click **Multitenant**.
5. In the **Domain** field, enter the domain name of the service provider (for example, managed-sp.com).
6. Enter a **Cluster Id** (for example, cluster-1 or 123456).
7. Click **Save**.
8. Click **Proceed** to confirm that you want to change the tenancy mode.

Cisco vManage reboots in multitenant mode and when a provider user logs in to Cisco vManage, the provider dashboard appears.



Note The **Domain** and **Cluster Id** values created in steps 5 and 6 serve as the Provider FQDN. Ensure these values conform to current DNS naming conventions. You can not modify these values after the configuration is saved. To change these values, a new Cisco vManage cluster need to be deployed. For more details on Provider and Tenant DNS requirements, see step 3.d in [Add a New Tenant](#).

Add Cisco vSmart Controller

1. Log in to Cisco vManage as the provider **admin** user.
2. From the Cisco vManage menu, choose **Configuration > Devices**.
3. Click **Controllers**.
4. Click **Add Controller** and click **vSmart**.

5. In the **Add vSmart** dialog box, do the following:
 - a. In the **vSmart Management IP Address** field, enter the system IP address of the Cisco vSmart Controller.
 - b. Enter the **Username** and **Password** required to access the Cisco vSmart Controller.
 - c. Select the protocol to use for control-plane connections. The default is **DTLS**.
If you select **TLS**, enter the port number to use for TLS connections. The default is 23456.
 - d. Check the **Generate CSR** check box for Cisco vManage to create a Certificate Signing Request.
 - e. Click **Add**.
6. From the Cisco vManage menu, choose **Configuration > Certificates**.
For the newly added Cisco vSmart Controller, the **Operation Status** reads **CSR Generated**.
 - a. For the newly added Cisco vSmart Controller, click **More Options** icon and click **View CSR**.
 - b. Submit the CSR to the Certificate Authority (CA) and obtain a signed certificate.
7. From the Cisco vManage menu, choose **Configuration > Certificates**.
8. Click **Install Certificate**.
9. In the **Install Certificate** dialog box, paste the **Certificate Text** or click **Select a file** upload the certificate file. Click **Install**.

Cisco vManage installs the certificate on the Cisco vSmart Controller. Cisco vManage also sends the serial number of the certificate to other controllers.

On the **Configuration > Certificates** page, the **Operation Status** for the newly added Cisco vSmart Controller reads as **vBond Updated**.

On the **Configuration > Devices** page, the new controller is listed in the Controller table with the controller type, hostname of the controller, IP address, site ID, and other details. The **Mode** is set to **CLI**.
10. Change the mode of the newly added Cisco vSmart Controller to **vManage** by attaching a template to the device.
 - a. From the Cisco vManage menu, choose **Configuration > Templates**.
 - b. Click **Device Templates**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled as **Device**

- c. Find the template to be attached to the Cisco vSmart Controller.
- d. Click **...**, and click **Attach Devices**.
- e. In the **Attach Devices** dialog box, move the new controller to the **Selected Device** list and click **Attach**.
- f. Verify the **Config Preview** and click **Configure Devices**.

Cisco vManage pushes the configuration from the template to the new controller.

In the **Configuration > Devices** page, the **Mode** for the Cisco vSmart Controller shows **vManage**. The new Cisco vSmart Controller is ready to be used in your multitenant deployment.

Expand a Multitenant Deployment to Support More Tenants and Tenant Devices

As a service provider, suppose you have deployed a Cisco SD-WAN multitenant overlay to support 50 tenants and 1000 devices. If you need to support more tenants or more devices, you can expand the Cisco vManage cluster and add additional Cisco vSmart Controllers to the overlay to support up to 100 tenants and 5000 devices. From Cisco SD-WAN Release 20.6.3, Cisco vManage Release 20.6.3, you can expand the Cisco vManage cluster and add additional Cisco vSmart Controllers to the overlay to support up to 150 tenants and 7500 devices.

Prerequisites

A multitenant Cisco SD-WAN overlay that supports up to 50 tenants and 1000 devices, deployed according to the steps outlined in the *Initial Setup for Multitenancy* section of this document.

1. [Expand a 3-Node Cluster to a 6-node Cluster](#).
2. To support up to 100 tenants and 5000 devices, you must have 10 Cisco vSmart Controllers in the overlay. So, deploy 4 Cisco vSmart Controllers in addition to the 6 existing Cisco vSmart Controllers in the overlay.
To support up to 150 tenants and 7500 devices, you must have 16 Cisco vSmart Controllers in the overlay. So, deploy 10 Cisco vSmart Controllers in addition to the 6 existing Cisco vSmart Controllers in the overlay.
 - a. Create Cisco vSmart Controller instances. See [Deploy the Cisco vSmart Controller](#).
 - b. [Add Cisco vSmart Controller](#) to the overlay network.

You can now add more tenants or allow your existing tenants to add more devices subject to the relevant limits.

Expand a 3-Node Cluster to a 6-node Cluster



Note You can only expand a 3-node Cisco vManage cluster to a 6-node Cisco vManage cluster. Expansion of the 3-node cluster to other cluster sizes is not supported.

1. To support 100 tenants and 5000 devices: Upgrade the three Cisco vManage servers in the existing 3-node cluster to the hardware specifications in the table *Hardware Specifications to Support 100 Tenants and 5000 Devices* from the *Supported Devices and Controller Specifications* section of this document.

From Cisco SD-WAN Release 20.6.3, Cisco vManage Release 20.6.3, to support 150 tenants and 7500 devices: Upgrade the three Cisco vManage servers in the existing 3-node cluster to the hardware specifications in the table *Hardware Specifications to Support 150 Tenants and 7500 Devices* from the *Supported Devices and Controller Specifications* section of this document.

2. Download the Cisco vManage Release 20.6.1 or a later release software image from [Cisco Software Download](#).
3. Create three Cisco vManage instances (say vManage1, vManage2, and vManage3) by installing the downloaded software image file. See [Deploy Cisco vManage](#).

**Important**

- Deploy Cisco vManage servers having the hardware specifications in the table *Hardware Specifications to Support 100 Tenants and 5000 Devices* from the *Supported Devices and Controller Specifications* section of this document.

From Cisco SD-WAN Release 20.6.3, Cisco vManage Release 20.6.3, to support 150 tenants and 7500 devices, deploy Cisco vManage servers having the hardware specifications in the table *Hardware Specifications to Support 150 Tenants and 7500 Devices* from the *Supported Devices and Controller Specifications* section of this document.

- Choose the **Data** persona for each Cisco vManage instance.

4. Complete the following operations on vManage1 through vManage3:

**Important**

Do not enable multitenancy on vManage1 through vManage3.

- a. Configure the following using the CLI:

- System IP address
- Site ID
- Service Provider organization name (`sp-organization-name`)
- Organization-name
- vBond IP address
- VPN 0 Transport/Tunnel interface
- VPN 0 Out-of-band (OOB) interface: Ensure that you assign a static IP address to this interface. Do not enable DHCP.
- VPN 512 Management interface

**Note**

Configure only one default route in VPN 0.

- b. (Optional) Using the CLI, install the Root CA certificate for vManage1.

**Note**

Skip this step if you are using a Symantec or Cisco PKI certificate.

- c. Complete the following through the Cisco vManage GUI:

1. [Generate a Certificate Signing Request](#)
2. After Symantec or your enterprise root CA has signed the certificate, [install the signed certificate](#).
- d. [Log in to the Cisco vManage Web Application Server](#).
- e. Ping the OOB interfaces on the other Cisco vManage instances and ensure they are reachable.
- f. [Configure the Cluster IP Address of the Cisco vManage Server](#).

Before proceeding to the next step, ensure that the **vManage IP Address** field on the **Administration > Cluster Management** page shows the OOB interface address.

5. Log in to the GUI of the existing 3-node Cisco vManage cluster and [add vManage1 to the cluster](#).
vManage1 reboots before being added to the cluster.

While vManage1 is being added to the cluster, on the **Administration > Cluster Management** page, the **Configure Status** for vManage1 shows **Pending**. You can monitor the System Generated Cluster Sync transaction to check the progress of the adding vManage1 to the cluster.

When the operation is completed, on the **Administration > Cluster Management** page, you can view vManage1 and its node persona listed along with the three Cisco vManage instances that were part of the original 3-node cluster.

6. Repeat **Step 4** and add vManage2 and vManage3 to the cluster.

Manage Tenants

Table 111: Feature History

Feature Name	Release Information	Description
Tenant Device Forecasting	Cisco SD-WAN Release 20.6.1 Cisco vManage Release 20.6.1	With this feature, a service provider can control the number of WAN edge devices a tenant can add to their overlay network. By doing so, the provider can utilize Cisco SD-WAN controller resources efficiently.

Tenant Device Forecasting

While adding a new tenant to the multitenant Cisco SD-WAN deployment, a service provider can forecast the number of WAN edge devices that the tenant may deploy in their overlay network. Cisco vManage enforces this forecast limit. If the tenant tries to add devices beyond this limit, Cisco vManage responds with an appropriate error message and the device addition fails.

In a multitenant deployment, a tenant can add a maximum of 1000 devices to their overlay network.



Note From Cisco SD-WAN Release 20.6.2, Cisco vManage Release 20.6.2, you can modify the device forecast for a tenant after the tenant is added. This modification is not supported in Cisco SD-WAN Release 20.6.1, Cisco vManage Release 20.6.1.

Benefits:

- The service provider can ensure that the Cisco SD-WAN controller resources are used more efficiently.
- Depending on the configuration, a multitenant deployment can support a fixed number of WAN edge devices across all tenants. By forecasting the number of devices a tenant may add, the service provider can assign a quota for each tenant from the overall pool of edge devices that the deployment can support.

Add a New Tenant

Prerequisites

- At least two Cisco vSmart Controllers must be operational and in the `vManage` mode before you can add new tenants.

A Cisco vSmart Controller enters the `vManage` mode when you push a template onto the controller from Cisco vManage. A Cisco vSmart Controller in the `CLI` mode cannot serve multiple tenants.

- Each pair of Cisco vSmart Controllers can serve a maximum of 24 tenants and a maximum of 1000 tenant devices. Ensure that there at least two Cisco vSmart Controllers that can serve a new tenant. If no pair of Cisco vSmart Controllers in the deployment can serve a new tenant, add two Cisco vSmart Controllers and change their mode to `vManage`.
- If you add a second tenant immediately after adding a tenant, Cisco vManage adds them sequentially, and not in parallel.
- Each tenant must have a unique Virtual Account (VA) on **Plug and Play Connect** on [Cisco Software Central](#). The tenant VA should belong to the same Smart Account (SA) as the provider VA.
- For an on-premises deployment, create a Cisco vBond Orchestrator controller profile for the tenant on **Plug and Play Connect**. The fields in the following table are mandatory.

Table 112: Controller Profile Fields

Field	Description/Value
Profile Name	Enter a name for the controller profile.
Multi-Tenancy	From the drop-down list, select Yes .
SP Organization Name	Enter the provider organization name.
Organization Name	Enter the tenant organization name in the format <code><SP Org Name>-<Tenant Org Name></code> . Note The organization name can be up to 64 characters.

Field	Description/Value
Primary Controller	Enter the host details for the primary Cisco vBond Orchestrator.

For a cloud deployment, the Cisco vBond Orchestrator controller profile is created automatically as part of the tenant creation process.

1. Log in to Cisco vManage as the provider **admin** user.
2. From the Cisco vManage menu, choose **Administration > Tenant Management**.
3. Click **Add Tenant**. In the **Add Tenant** dialog box:
 - a. Enter a name for the tenant.
For a cloud deployment, the tenant name should be same as the tenant VA name on **Plug and Play Connect**.
 - b. Enter a description of the tenant.
The description can be up to 256 characters and can contain only alphanumeric characters.
 - c. Enter the name of the organization.
The organization name is case-sensitive. Each tenant or customer must have a unique organization name.
Enter the organization name in the following format:
<SP Org Name>--<Tenant Org Name>
For example, if the provider organization name is 'multitenancy' and the tenant organization name is 'Customer1', while adding the tenant, enter the organization name as **multi tenancy-Customer1**.



Note The organization name can be up to 64 characters.

- d. In the **URL Subdomain Name** field, enter the fully qualified sub-domain name of the tenant.
 - The sub-domain name must include the domain name of the service provider. For example, for the `managed-sp.com` service provider, a valid domain name can be `customer1.managed-sp.com`.



Note The service provider name is shared amongst all tenants. Hence, ensure that the URL naming convention follows the same domain name convention that was provided while enabling multitenancy from **Administration > Settings > Tenancy Mode**.

- For an on-premises deployment, add the fully qualified sub-domain name of the tenant to the DNS. Map the fully qualified sub-domain name to the IP addresses of the three Cisco vManage instances in the Cisco vManage cluster.

- **Provider Level:** Create DNS A record and map it to the IP addresses of the Cisco vManage instances running in the Cisco vManage cluster. The A record is derived from the domain and Cluster ID that was created in steps 5 and 6 in [Enable Multitenancy on Cisco vManage](#). For example, if domain is **sdwan.cisco.com** and Cluster ID is **vmanage123**, then A record will need to be configured as **vmanage123.sdwan.cisco.com**.



Note If you fail to update DNS entries, it will result in authentication errors when logging in to Cisco vManage. Validate DNS is configured correctly by executing **nslookup vmanage123.sdwan.cisco.com**.

- **Tenant Level:** Create DNS CNAME records for each tenant created and map them to the FQDN created at the Provider Level. For example, if domain is **sdwan.cisco.com** and tenant name is **customer1** the CNAME record will need to be configured as **customer1.sdwan.cisco.com**.



Note Cluster ID is not required for CNAME record. Validate DNS is configured correctly by executing **nslookup customer1.sdwan.cisco.com**.

For a cloud deployment, the fully qualified sub-domain name of the tenant is automatically added to the DNS as part of the tenant creation process. After you add a tenant, it could take up to an hour before the fully qualified sub-domain name of the tenant can be resolved by the DNS.

- In the **Number of Devices** field, enter the number of WAN edge devices that the tenant can deploy. If the tenant tries to add WAN edge devices beyond this number, Cisco vManage reports an error and the device addition fails.
- Click **Save**.

The **Create Tenant** screen appears, and the **Status** of the tenant creation reads **In progress**. To view status messages related to the creation of a tenant, click the > button to the left of the status.

Cisco vManage does the following:

- creates the tenant
- assigns two Cisco vSmart Controllers to serve the tenant and pushes a CLI template to these controllers to configure tenant information
- sends the tenant and Cisco vSmart Controller information to Cisco vBond Orchestrators.

What to do next:

After the **Status** column changes to **Success**, you can view the tenant information on the **Administration > Tenant Management** page.

Modify Tenant Information

1. Log in to Cisco vManage as the provider **admin** user.

2. From the Cisco vManage menu, choose **Administration > Tenant Management**.
3. In the left pane, click the name of the tenant.
The tenant information is displayed in a pane on the right.
4. To modify tenant data, do as follows:
 - a. In the right pane, click the pencil icon.
 - b. In the **Edit Tenant** dialog box, you can modify the following:
 - **Description**: The description can be up to 256 characters and can contain only alphanumeric characters.
 - **Forecasted Device**: The number of WAN edge devices that the tenant can deploy.
A tenant can add a maximum of 1000 devices.



Note This option is available from Cisco SD-WAN Release 20.6.2, Cisco vManage Release 20.6.2.

If you increase the number of devices that a tenant can deploy, you must add the required number of device licenses to the tenant virtual account on **Plug and Play Connect** on [Cisco Software Central](#).

Before you increase the number of devices that a tenant can deploy, ensure that the Cisco vSmart Controller pair assigned to the tenant can support this increased number. A pair of Cisco vSmart Controllers can support a maximum of 24 tenants and 1000 devices across all these tenants.

-
- **URL Subdomain Name**: Modify the fully qualified sub-domain name of the tenant.
- c. Click **Save**

Delete a Tenant

Before you delete a tenant, delete all tenant WAN edge devices. See [Delete a WAN Edge Device from a Tenant Network, on page 381](#).

1. Log in to Cisco vManage as the provider **admin** user.
2. From the Cisco vManage menu, choose **Administration > Tenant Management**.
3. In the left pane, click the name of the tenant.
The tenant information is displayed in a pane on the right.
4. To delete the tenant, do as follows:
 - a. In the right pane, click the trash icon.
 - b. In the **Delete Tenant** dialog box, enter the provider **admin** password and click **Save**.

Cisco vManage Dashboard for Multitenancy

After enabling Cisco vManage for multitenancy, you can view the multitenant dashboard when you log in to Cisco vManage. Cisco vManage multitenant dashboard is a portal where the provider or tenant can view and provision the underlying system.

The bar at the top of every Cisco vManage multitenant screen includes icons that allow smooth navigation.

View Tenant Activity, Device, and Network Information

When you log in to a multitenant Cisco vManage as an administrator, the provider dashboard displays the following components. To return to the provider dashboard from other Cisco vManage screens, click **Dashboard**.

- **Device pane** — runs across the top of the multitenant dashboard screen. The device pane displays the number of active Cisco vSmart Controllers, Cisco vBond Orchestrators, and Cisco vManage instances, the connectivity status of devices, and information on certificates that have expired or about to expire.
- **Tenants pane** — displays the total number of tenants and a summary of the control status, site health, router health, and Cisco vSmart Controller status of all tenants.
- **Table of tenants in the overlay network** — List of individual tenants, with separate information about the control status, site health, WAN edge device health, and Cisco vSmart Controller status for each tenant.

To display tenant-specific status summary information,

1. Click a tenant name from the tenant list.

A dialog box opens on the right side of the screen that provides additional information about the status of the tenant.

2. To access the tenant dashboard for the selected tenant, click **<Tenant name> Dashboard**.

Cisco vManage presents the provider-as-tenant view and displays the tenant dashboard. To return to the provider view, click **Provider** at the top of page.

3. To close the dialog box, click the tenant name from the tenant list.

View Detailed Information of a Tenant Setup

Cisco vManage displays the tenant dashboard, which provides information about a tenant deployment when

- a provider **admin** user selects a specific tenant from the **Select Tenant** drop-down list in the provider dashboard. This view is called the provider-as-tenant view.
- a **tenantadmin** user logs in to Cisco vManage. This view is called the tenant view.

View All Network Connections in the Tenant Overlay Network

The **Device** pane runs across the top of the tenant dashboard and displays the number of control connections from Cisco vManage to the Cisco vSmart Controllers and routers in the overlay network of a tenant. For each WAN edge device, the Device pane shows

- Total number of control connections between Cisco vSmart Controllers and WAN edge devices
- Number of valid control connections between Cisco vSmart Controllers and WAN edge devices
- Number of invalid control connections between Cisco vSmart Controllers and WAN edge devices

Click a connection number, or the **Up** or **Down** arrow, to display a table with detailed information about each connection. Click the **More Actions** icon at the right of each table row to access the **Device Dashboard** or **Real Time** view from the **Monitor > Network** screen, or access the **Tools > SSH Terminal** Screen.

Click a connection number, or the **Up** or **Down** arrow, to display a table with detailed information about each connection. Click the **More Actions** icon at the right of each table row to access the **Device Dashboard** or **Real Time** view from the **Monitor > Devices** screen, or access the **Tools > SSH Terminal** Screen.



Note In Cisco vManage Release 20.6.1 and earlier releases, you can view information related to the **Monitor > Devices** page under the **Monitor > Network** page.

View Information About Device Reboots

The **Reboot** pane displays the total number of reboots in the last 24 hours for all devices in the network. It includes soft and cold reboots and reboots that occurred as a result of power-cycling a device. For each reboot, the following information is listed:

- System IP and hostname of the device that rebooted.
- Time when the device was rebooted.
- Reason for the device reboot

If the same device reboots more than once, each reboot option is reported separately.

Click the **Reboot** pane to open the **Reboot** dialog box. In the **Reboot** dialog box, click the **Crashes** tab. For all device crashes, the following information is listed:

- System IP and hostname of the device on which the crash occurred.
- Crash index of the device
- Core time when the device crashed.
- File name of the device crash log

View Network Connections

The **Control Status** pane displays whether Cisco vSmart Controller and WAN edge devices are connected. Each Cisco vSmart Controller must connect to all other Cisco vSmart Controllers in the network. Each WAN edge device must connect to the maximum number of configured Cisco vSmart Controllers. The **Control Status** pane displays three network connection counts:

- Control Up — total number of devices with the required number of operational control plane connections to a Cisco vSmart Controller
- Partial — total number of devices with some, but not all, operational control plane connection to Cisco vSmart Controllers.

- Control Down — total number of devices with no control plane connection to a Cisco vSmart Controller

To display a table with device details, click a row from the **Control Status** dialog box. Click the **More Actions** icon at the right of each table row to access the **Device Dashboard** or **Real Time** view from the **Monitor > Devices** screen.



Note In Cisco vManage Release 20.6.1 and earlier releases, you can view information related to the **Monitor > Devices** page under the **Monitor > Network** page.

View State of Data Connections for a Site

The **Site Health** pane displays the state of data connections for a site. When a site has multiple WAN edge devices, this pane displays the state for the entire site and not for individual devices. The Site Health pane displays three connectivity states:

- Full WAN Connectivity — total number of sites where all BFD sessions on all routers are in the up state.
- Partial WAN Connectivity — total number of sites where tunnel and all BFD sessions on all routers are in the down state. These sites still have limited data plane connectivity.
- No WAN Connectivity — total number of sites where all BFD sessions on all routers are in the down state. These sites have no data plane connectivity.

To display a table with detailed information about each site, node, or tunnel, click a row from the **Site Health** dialog box. Click the **More Actions** icon at the right of each row in the table to access the **Device Dashboard** or **Real Time** view from the **Monitor > Devices** screen, or access the **Tools > SSH Terminal** screen.



Note In Cisco vManage Release 20.6.1 and earlier releases, you can view information related to the **Monitor > Devices** page under the **Monitor > Network** page.

View Interface Usage for WAN Edge Interfaces

The **Transport Interface Distribution** pane displays interface usage in the last 24 hours for all WAN edge interfaces in VPN 0. It includes all TLOC interfaces. Click the pane to view details of interface usage in the **Transport Interface Distribution** dialog box.

View WAN Edge Device Counts

The **WAN Edge Inventory** pane provides four WAN edge device counts:

- Total — total number of authorized serial numbers for WAN edge devices that have been uploaded on Cisco vManage. The serial number is uploaded on the **Configuration > Devices** screen.
- Authorized — total number of authorized WAN edge devices in the overlay network. These WAN edge devices are marked as **Valid** in the **Configuration > Certificates > WAN Edge List** screen.
- Deployed — total number of deployed WAN edge devices. These are WAN edge devices that are marked as **Valid** and are now operational in the network.

- Staging — total number of WAN edge devices you configure at a staging site before they are made a part of the overlay network. These routers do not take part in any routing decisions and do not affect network monitoring through Cisco vManage.

Click the pane to view hostname, system IP, site ID, and other details of each router from the **WAN Edge Inventory** dialog box.

View Aggregated State of WAN Edge Devices

The **WAN Edge Health** pane offers an aggregated view of the state of WAN edge devices by providing a count of the number of devices in each state, therefore describing the health of the hardware nodes. The three WAN edge device states are:

- Normal — number of WAN edge devices with memory, hardware, and CPU in normal state. Using less than 70% of total memory or total CPU is classified as, normal.
- Warning — number of WAN edge devices with memory, hardware, or CPU in warning state. Using between 70% and 90% of total memory or total CPU is classified as, warning
- Error — number of WAN edge devices with memory, hardware, or CPU in error state. Using more than 90% of total memory or total CPU is classified as, error.

Click a number or the WAN edge device state to display a table with the last 12 or 24 hours of memory usage, CPU utilization, and hardware-related alarms, including temperature, power supply, and PIM modules. Click the **More Actions** icon at the right of each row in the table to access the following:

- **Hardware Environment**
- **Real Time** view from the **Monitor > Network** screen




Note In Cisco vManage Release 20.6.1 and earlier releases, you can view information related to the **Monitor > Devices** page under the **Monitor > Network** page.


- **Tools > SSH Terminal** screen.

View WAN Edge Device Loss, Latency, Jitter

The **Transport Health** pane displays the aggregated average loss, latency, and jitters for all links and all combinations of colors (for example, all LTE-to-LTE links, all LTE-to-3G links).

From the **Type** drop-down arrow, choose loss, latency, or jitter.

Click the  icon to select a time period for which to display the transport health.


Click the  icon to open the **Transport Health** dialog box. This dialog box displays a more detailed view. To display information in a tabular format, click the **Details** tab. You can choose to change the displayed health type and time period.


View SAIE Flow Information of WAN Edge Devices

The **Top Applications** pane displays SD-WAN Application Intelligence Engine (SAIE) flow information for traffic transiting routers in the overlay network.



-
- Note**
- In Cisco vManage Release 20.7.x and earlier releases, the SAIE flow is known as deep packet inspection (DPI).
 - The SAIE flow information is shown only for the last 24 hours. To view SAIE flow information for a time before the last 24 hours, you must check the information for the specific device.
-

Click the  icon to select a time period for which to display data. From the **VPN** drop-down list, select a VPN to display SAIE information for all flows in that VPN.


Click the  icon to open the **Top Applications** dialog box. This dialog box displays a more detailed view of the same information. You can change the VPN and time period.


View Tunnels Data

The **Application-Aware Routing** pane allows you to choose the following tunnel criteria from the **Type** drop-down arrow:

- Loss
- Latency
- Jitter

Based on the tunnel criteria, the pane displays the 10 worst tunnels. For example, if you choose loss, the pane shows 10 tunnels with the greatest average loss over the last 24 hours.

Click the  icon against a row to display a graphical representation of the data. Select a time period for which to display data or click **Custom** to display a drop-down arrow for specifying a custom time period.

Click the  icon to open the **Application-Aware Routing** dialog box. This dialog box displays the 25 worst tunnels based on criteria you choose from the **Type** drop-down arrow, the criteria being loss, latency, and jitter.

Manage Tenant WAN Edge Devices

Add a WAN Edge Device to a Tenant Network



-
- Note** If you are adding a WAN edge device that you had previously invalidated and deleted from an overlay network, you must reset the device software after adding the device. To reset the software on a Cisco vEdge device, use the command **request platform software reset**.
-

1. Log in to Cisco vManage.

If you're a provider user, log in as the **admin**. In the provider dashboard, choose a tenant from the drop-down list to enter the provider-as-tenant view.

If you're a tenant user, log in as the **tenantadmin**.

2. Upload the device serial number file to Cisco vManage.
3. Validate the device and send details to controllers.
4. Create a configuration template for the device and attach the device to the template.

While configuring the device, configure the service provider organization name and the tenant organization name as in the following example:

```
sp-organization-name multitenancy
organization-name multitenancy-Customer1
```



Note Enter the `organization-name` in the format `<SP Org Name>-<Tenant Org Name>`.

5. Bootstrap the device using bootstrap configuration generated through Cisco vManage or manually create the initial configuration on the device.
6. If you are using Enterprise Certificates to authenticate the device, download the CSR from Cisco vManage and get the CSR signed by the Enterprise CA. Install the certificate on Cisco vManage.

Delete a WAN Edge Device from a Tenant Network

1. Log in to Cisco vManage.
If you're a provider user, log in as the **admin**. In the provider dashboard, choose a tenant from the drop-down list to enter the provider-as-tenant view.
If you're a tenant user, log in as the **tenantadmin**.
2. Detach the device from any configuration templates.
3. [Delete a WAN Edge Router](#).

Tenant-Specific Policies on Cisco vSmart Controllers

A provider **admin** user (from the Cisco vManage provider-as-tenant view) or a **tenantadmin** user (from the Cisco vManage tenant view) can create and deploy tenant-specific policies on the Cisco vSmart Controllers serving the tenant. The user can configure a CLI policy or create the policy using the UI policy configuration wizard.

When you activate or deactivate a policy,

1. Cisco vManage identifies the Cisco vSmart Controllers serving the tenant.
2. Cisco vManage notifies the Cisco vSmart Controllers to pull the policy configuration.
3. Cisco vSmart Controllers pull and deploy the policy configuration.
4. Cisco vManage reports the status of the policy pull by the Cisco vSmart Controllers.

Manage Tenant Data

Back Up Tenant Data

The tenant data backup solution of Cisco vManage multitenancy provides the following functionalities:

- [Create, Extract, and List Configuration Data Backup File](#).
- Back up configuration database of a specific tenant with an option to restore it later. See [Restore and Delete Tenant Data Backup File](#).
- Delete back up files of a tenant stored in Cisco vManage. For deleting tenant data backup files, see [Restore and Delete Tenant Data Backup File](#).

The following factors are applicable when using data backup solution:

- The tenant data backup solution operations can be performed by a tenant administrator in the tenant view and or by a provider administrator in the provider-as-tenant view. To know how to access tenant dashboard through different views, see [User Roles in Multitenant Environment, on page 357](#).
- A tenant is allowed to perform the following backup operations at a particular time and must complete an operation before starting a new operation:
 - Back up a single configuration database
 - Download the backup file.
 - Restore or import backup files
 - Delete backup files.
 - List backup files

- A tenant backup file format is as follows:

```
Bkup_tenantId_MMDDYY-HHMMSS_taskIdWithoutDash.tar.gz
```

- The tenant data backup operation is a readonly operation on the configuration database. However, to ensure data consistency and prevent data loss, do not perform any major changes on the network while the operation is in progress.
- Multiple tenants can perform back-up and restore operations in parallel.
- A tenant is not allowed to perform other backup operations when the restore operation of the tenant database is in-progress. So, a tenant can perform a single backup operation and when this operation is in-progress, all new backup operation requests are rejected.

The remaining tenants can continue with their backup operations.

- A tenant must perform backup and restore operations on Cisco vManage instances running identical Cisco vManage software versions.
- A tenant can store a maximum of three backup files in Cisco vManage and can download to store them outside Cisco vManage repository. If the tenant already has three backup files, a subsequent backup operation results in the earliest backup file being deleted and a new backup file being generated.

- Ensure that the following parameter values match in both the backup file and the setup where tenant has requested for a restore operation:
 - Tenant Id
 - Organization Name
 - SP Organization Name
- The tenant data backup solution creates a task in the tenant view of Cisco vManage. Therefore, the tenant can monitor the progress of the operation from the task view of the tenant dashboard.
- A provider cannot back up provider data using this solution. Therefore, the provider can back up all tenants information at once by backing up all tenants configuration database using CLI.

Create, Extract, and List Configuration Data Backup File

1. Log in to Cisco vManage.

If you're a provider user, log in as the **admin**. In the provider dashboard, choose a tenant from the drop-down list to enter the provider-as-tenant view.

If you're a tenant user, log in as the **tenantadmin**.

2. In the address bar, modify the URL path with `dataservice` for the REST API connection.

Example: `https://<tenant_URL>/dataservice`

3. Create a configuration backup file by using the following API:

`https://<tenant_URL>/dataservice/tenantbackup/export.`

4. If the configuration backup file has been created successfully, Cisco vManage task view indicates that the backup file has been generated. You can view the process identifier of the created process or task.

Example:

```
{
  "processId": "72d69805-b987-436f-9b7a-afef2f3f9061",
  "status": "in-progress"
}
```

5. Verify the task status using the obtained process identifier.

Example:

`https://<tenant_URL>/dataservice/device/action/status/72d69805-b987-436f-9b7a-afef2f3f9061`

The verification generates the details of the task in the JSON file format.

6. After the task is completed, extract or download the backed-up file available under the **data** section of the JSON task file.

Example: To extract or download the backup file, use the following API:

`https://<tenant_URL>/dataservice/tenantbackup/download/1570057020772/backup_1570057020772_100919-181838.tar.gz`

7. List backup files stored in Cisco vManage using the following API.

Example: `https://<tenant_URL>/dataservice/tenantbackup/list`

Restore and Delete Tenant Data Backup File

Before you begin:

To run the restore and delete tenant data backup files API, you can download and install Postman tool or any other alternative tool for testing http applications and services. In this document, the procedure to restore and delete tenant data backup files has been explained using the Postman tool. Postman is a software tool used as an API development environment. You can download the tool from the Postman website.

1. Open Google Chrome, or another browser, and enable developer mode on it.
2. Log in to Cisco vManage.

If you're a provider user, log in as the **admin**. In the provider dashboard, choose a tenant from the drop-down list to enter the provider-as-tenant view.

If you're a tenant user, log in as the **tenantadmin**.

3. To get header information of the restore API, do as follows:
 - a. On the right side of the screen, click the **Network** tab to get the network capture view.
 - b. In the network capture view, click the **Name** column to sort the listed items.
 - c. Search and click **index.html**.
 - d. Click the **Headers** tab and expand **Request Headers**.
 - e. Choose all text under **Request Headers** and copy it to the clipboard.
4. Import backup files through the Postman UI:
 - a. Open the Postman UI.
 - b. To disable SSL certificate verification, click **Postman > Preferences > General > Request**. Turn off **SSL Certificate Verification**.
 - c. In the Postman UI, create a new tab.
 - d. Click **Headers** and then click **Bulk Edit**.
 - e. Paste the text that was copied in step 3 from the **Request Headers** block into an editable form.
 - f. From the **GET** method drop-down list, choose **POST**.
 - g. In the **Paste request URL** field, paste the dedicated URL of the tenant and include `dataservice/tenantbackup/import`.
Example: `https://customer1.managed-sp.com/dataservice/tenantbackup/import`
 - h. Click the **Body** tab and select **form-data**.
 - i. Under **KEY** column, enter `bakup.tar.gz`
 - j. Under **VALUE** column, click **Select Files** and select a backup file to be imported.
 - k. To run the API, click **Send**.

In the **Response** section of the Postman UI, you can view the JSON information that indicates the file that was restored.

5. Monitor the restoration of backup files in either of the following ways:
 - a. Use Cisco vManage task view that indicates if backup file has been imported successfully. You can view the process identifier of the created process or task.

Example:

```
{
  "processId": "40adb6c0-eacc-4ad4-ba6c-2c2da2e96d1d",
  "status": "Import Successfully Submitted for tenant 1579026919487"
}
```
 - b. Use the following URL to get the status,


```
https://<tenant_URL>/dataservice/device/action/status/<processId>
```

Example:

```
https://customer1.managed-sp.com/dataservice/device/action/status/40adb6c0-eacc-4ad4-ba6c-2c2da2e96d1d
```
6. Delete tenant data backup file through Postman UI.
 - a. In the Postman UI, create a new tab.
 - b. Click **Headers** and then click **Bulk Edit**.
 - c. Paste the text that was copied in step 3 from the **Request Headers** block into an editable form.
 - d. From the **GET** method drop-down list, choose **DELETE**.
 - e. In the **Paste request URL** field, paste the dedicated URL of the tenant and include `dataservice/tenantbackup/delete?fileName=' filename '`. The filename can either be name of the backup file or all.

Example:

```
https://customer1.managed-sp.com/dataservice/tenantbackup/delete?fileName=bkup_1579026919487_012820-180712_c09230904dfc40edb0d1e50b68b03002.tar.gz
```

Example: `https://customer1.managed-sp.com/dataservice/tenantbackup/delete?fileName=all`
 - f. To run the API, click **Send**.

In the **Response** section of the Postman UI, you can view the JSON information that indicates the files that were deleted.

Example:

```
{
  "Deleted": [
    "bkup_1579026919487_012820-180712_c09230904dfc40edb0d1e50b68b03002.tar.gz"
  ]
}
```

View OMP Statistics per Tenant on a Cisco vSmart Controller

1. Log in to Cisco vManage as the provider **admin** user.
2. From the Cisco vManage menu, choose **Monitor > Devices**.
Cisco vManage Release 20.6.x and earlier: From the Cisco vManage menu, choose **Monitor > Network**.
3. In the table of devices, click on the hostname of a Cisco vSmart Controller.

4. In the left pane, click **Real Time**.
5. In the **Device Options** field, enter **OMP** and select the OMP statistics you wish to view.
6. In the **Select Filters** dialog box, click **Show Filters**.
7. Enter the **Tenant Name** and click **Search**.

Cisco vManage displays the selected OMP statistics for the particular tenant.

View Tenants Associated with a Cisco vSmart Controller

1. Log in to Cisco vManage as the provider **admin** user.
2. Click a **vSmart** connection number to display a table with detailed information about each connection.

Cisco vManage displays a table that provides a summary of the Cisco vSmart Controllers and their connections.

3. For a Cisco vSmart Controller, click **...** and click **Tenant List**.

Cisco vManage displays a summary of tenants associated with the Cisco vSmart Controller.

Migrate Single-Tenant Cisco SD-WAN Overlay to Multitenant Cisco SD-WAN Deployment

Before You Begin

- Before you begin the migration,
 - Migration of a single-tenant overlay to a multitenant deployment is only supported with the Cisco SD-WAN controllers deployed on-premises. Migration is yet to be supported with cloud-hosted Cisco SD-WAN controllers.
 - Ensure that the edge devices in the single-tenant deployment can reach the Cisco vBond Orchestrator in the multitenant deployment
 - Ensure that the template, routing, and policy configuration on the edge devices is synchronized with the current configuration on Cisco vManage
 - Configure a maintenance window for the single-tenant overlay before performing this procedure. See [Configure or Cancel vManage Server Maintenance Window](#).
- Minimum software requirements for the single-tenant overlay to be migrated:

Device	Software Version
Cisco vManage	Cisco vManage Release 20.6.1
Cisco vBond Orchestrator	Cisco SD-WAN Release 20.6.1
Cisco vSmart Controller	Cisco SD-WAN Release 20.6.1

Device	Software Version
Cisco vEdge Device	Cisco SD-WAN Release 20.6.1

- Minimum software requirements for the multitenant deployment to which the single-tenant overlay must be migrated:

Device	Software Version
Cisco vManage	Cisco vManage Release 20.6.1
Cisco vBond Orchestrator	Cisco SD-WAN Release 20.6.1
Cisco vSmart Controller	Cisco SD-WAN Release 20.6.1
Cisco vEdge Device	Cisco SD-WAN Release 20.6.1

- The software versions of the Cisco SD-WAN controllers and WAN edge devices must be identical in both the single-tenant and multitenant deployments.
- We recommend that you use a custom script or a third-party application like Postman to execute the API calls.

Migration Procedure

1. Export the single-tenant deployment and configuration data from a Cisco vManage instance controlling the overlay.

Method	POST
URL	<code>https://ST-vManage-IP-address</code>
Endpoint	<code>/dataservice/tenantmigration/export</code>
Authorization	Admin user credentials.

Body	<p>Required</p> <p>Format: Raw JSON</p> <pre>{ "desc": <tenant_description>, "name": <tenant_name>, "subdomain": <tenant_name>.<domain>, "orgName": <tenant_orgname > }</pre> <p>Field Description:</p> <ul style="list-style-type: none"> • desc: A description of the tenant. The description can be up to 256 characters and can contain only alphanumeric characters. • name: Unique name for the tenant in the multitenant deployment. • subdomain: Fully qualified sub-domain name of the tenant. The sub-domain name must include the domain name of the service provider. For example, if <code>managed-sp.com</code> is the domain name of service provider, and the tenant name is <code>Customer1</code>, the tenant sub-domain name would be <code>customer1.managed-sp.com</code>. • orgName: Name of the tenant organization. The organization name is case-sensitive.
Response	<p>Format: JSON</p> <pre>{ "processId": <vManage_process_ID>, }</pre>

While exporting the data, Cisco vManage attempts to detach any CLI templates from the edge devices in preparation for the migration to the multitenant deployment. If prompted by Cisco vManage, detach CLI templates from the edge devices and execute the export API call again.

2. Check the status of the data export task in Cisco vManage. When the task succeeds, download the data using the URL

`https://ST-vManage-IP-address/dataservice/tenantmigration/download/default.tar.gz`

3. On a multitenant Cisco vManage instance, import the data exported from the single-tenant overlay.

Method	POST
URL	<code>https://MT-vManage-IP-address</code>
Endpoint	<code>/dataservice/tenantmigration/import</code>
Authorization	Provider Admin user credentials.
Body	<p>Required</p> <p>Format: form-data</p> <p>Key Type: File</p> <p>Value: <code>default.tar.gz</code></p>

Response	Format: JSON <pre>{ "processId": <vManage_process_ID>, "migrationTokenURL": <token_URL>, }</pre>
----------	---------------------------------------------------------------------------------------------------------------------

When the task succeeds, on the multitenant Cisco vManage, you can view the devices, templates, and policies imported from the single-tenant overlay.

4. Obtain the migration token using the token URL obtained in response to the API call in **Step 3**.

Method	GET
URL	<code>https://MT-vManage-IP-address</code>
Endpoint	migrationTokenURL obtained in Step 3 .
Authorization	Provider Admin user credentials.
Response	The migration token as a large blob of encoded text.

5. On the single-tenant Cisco vManage instance, initiate the migration of the overlay to the multitenant deployment.

Method	POST
URL	<code>https://ST-vManage-IP-address</code>
Endpoint	<code>dataservice/tenantmigration/networkMigration</code>
Authorization	Admin user credentials.
Body	Required Format: Raw text Content: Migration token obtained in Step 4 .
Response	Format: JSON <pre>{ "processId": <vManage_process_ID>, }</pre>

In Cisco vManage, check the status of the migration task. As part of the migration task, the address of the multitenant vBond Orchestrator, and the service provider and tenant organization names are pushed to the WAN edge devices of the single-tenant overlay. If the task succeeds, WAN edge devices form control connections to controllers in the multitenant deployment; the WAN edge devices are no longer connected to the controllers of the single-tenant overlay.

Attach any CLI templates detached from the edge devices (in Step 1) after migration to the multitenant deployment. Before you attach the templates, update the Cisco vBond Orchestrator IP address and the Organization name to match the configuration of the multitenant deployment.



Note In the single-tenant deployment, if Cisco vManage-signed certificates are installed on cloud-based WAN edge devices, the certificates are cleared when the devices are migrated to the multitenant deployment. You must re-certify the devices on the multitenant Cisco vManage. If enterprise certificates are installed on the cloud-based WAN edge devices, the certificates are not affected by the migration. For more information, see [Enterprise Certificates](#).

Migrate Multitenant Cisco SD-WAN Overlay

Table 113: Feature History

Feature Name	Release Information	Description
Migrate Multitenant Cisco SD-WAN Overlay	Cisco SD-WAN Release 20.6.1 Cisco vManage Release 20.6.1	This feature enables you to migrate a multitenant Cisco SD-WAN overlay comprising shared Cisco vManage instances and Cisco vBond Orchestrators, and tenant-specific Cisco vSmart Controllers to a multitenant overlay comprising shared Cisco vManage instances, Cisco vBond Orchestrators, and Cisco vSmart Controllers.

Prerequisites

Minimum software requirements for Cisco SD-WAN controllers and WAN edge devices in the multitenant overlay to be migrated:

Device	Software Version
Cisco vManage	Cisco vManage Release 20.3.3
Cisco vBond Orchestrator	Cisco SD-WAN Release 20.3.3
Cisco vSmart Controller	Cisco SD-WAN Release 20.3.3
Cisco vEdge Device	Cisco SD-WAN Release 20.3.3

Restrictions

- This migration procedure applies only to Cisco SD-WAN controllers deployed on premises.
- The multitenant overlay can only be migrated to a setup in which Cisco vManage instances run Cisco vManage Release 20.6.1 software and Cisco SD-WAN controllers run Cisco SD-WAN Release 20.6.1 software.
- This migration procedure cannot be used to merge two or more multitenant overlays. Only one multitenant overlay can be migrated to the new setup at a time.

Migration Procedure

1. Upgrade the software on the three Cisco vManage instances in the cluster to Cisco vManage Release 20.6.1. For more information, see [Upgrade Cisco vManage Cluster](#).



Note Run the command **request nms configuration-db upgrade** on only one of the Cisco vManage instances.

2. After the Cisco vManage software is upgraded to Cisco vManage Release 20.6.1, log in to the Cisco vManage GUI.
You're prompted to set a new password.
 - a. Enter a new password that adheres to the password guidelines.
3. Upload the Cisco SD-WAN Release 20.6.1 software to Cisco vManage. For more information, see [Add an Image to the Software Repository](#).
4. Upgrade the Cisco vBond Orchestrator software to Cisco SD-WAN Release 20.6.1. For more information, see [Upgrade the Software Image on a Device](#).
5. Create two Cisco vSmart Controller instances running Cisco SD-WAN Release 20.6.1 software. See [Deploy the Cisco vSmart Controller](#).



Note With two Cisco vSmart Controller instances, you can support up to 24 tenants. To support up to 50 tenants, create six Cisco vSmart Controller instances.

- a. [Add Cisco vSmart Controller](#) to the overlay network.

The **Provider Dashboard** shows the new Cisco vSmart Controllers running Cisco SD-WAN Release 20.6.1 software. The **Tenant Dashboard** shows the older Cisco vSmart Controllers running Cisco SD-WAN Release 20.3.3 software.

6. Enable maintenance window on Cisco vManage. For more information, see [Configure or Cancel vManage Server Maintenance Window](#).
A maintenance window of 3 to 4 hours is recommended.
7. Migrate the tenant configuration from the older tenant-specific Cisco vSmart Controllers running Cisco SD-WAN Release 20.3.3 software to the new shared Cisco vSmart Controllers running Cisco SD-WAN Release 20.6.1 software.

Method	POST
URL	<code>https://<vmanageip>:<port></code>
Endpoint	<code>dataservice/tenant/vsmart-mt/migrate</code>
Authorization	Provider admin user credentials.
Body	Required Format: Raw JSON { }

Response	Format: JSON <pre>{ "processId": <vManage_process_ID>, }</pre>
----------	---------------------------------------------------------------------------

In Cisco vManage, check the status of the migration task using the `processId` from the API response. During the migration task, the following changes are effected:

- a. The older Cisco vSmart Controllers are invalidated and deleted from the overlay network.
 - b. In the tenant view, the older Cisco vSmart Controllers are removed from the **Tenant Dashboard**, and the **Devices** and the **Certificates** page.
 - c. The tenant WAN edge devices are connected to the new Cisco vSmart Controllers.
8. (Optional) Upgrade the Cisco vEdge device software to Cisco SD-WAN Release 20.6.1. For more information, see [Upgrade the Software Image on a Device](#) and [Activate a New Software Image](#).



Tip It is not necessary to upgrade the tenant WAN edge device software in the same maintenance window in which you migrate the multitenant overlay. However, we recommend that you upgrade the tenant WAN edge device software within a few weeks of the migration.

Verify the Migration

1. In the provider view, perform the following checks:
 - a. From the **Main Dashboard** page, verify whether the tenant WAN edge devices are connected to the new multitenant Cisco vSmart Controllers.
 - b. [View Tenants Associated with a Cisco vSmart Controller, on page 386](#).
 - c. On the Cisco vSmart Controller CLI, run the command **show control connections**. In the command output, verify that control connections are established between the Cisco vSmart Controller and the tenant WAN edge devices.
2. In the provider-as-tenant view, verify whether the multitenant Cisco vSmart Controllers appear on the **Tenant Dashboard**.

Upgrade Cisco SD-WAN Controller and Edge Device Software

Prerequisites

Minimum software requirements for Cisco SD-WAN controllers and WAN edge devices:

Device	Software Version
Cisco vManage	Cisco vManage Release 20.4.1 or later
Cisco vBond Orchestrator	Cisco SD-WAN Release 20.4.1 or later

Device	Software Version
Cisco vSmart Controller	Cisco SD-WAN Release 20.4.1 or later
Cisco vEdge Device	Cisco SD-WAN Release 20.4.1 or later

Upgrade Procedure

1. Upgrade the software on the three Cisco vManage instances in the cluster to Cisco vManage Release 20.6.1 or a later release. For more information, see [Upgrade Cisco vManage Cluster](#).



Note Skip the step to upgrade the configuration-db service using the command `request nms configuration-db upgrade`.

2. After the Cisco vManage software is upgraded to Cisco vManage Release 20.6.1 or a later release, log in to the Cisco vManage GUI.
3. Upload the Cisco SD-WAN Release 20.6.1 or a later release software to Cisco vManage. For more information, see [Add an Image to the Software Repository](#).
4. Upgrade the Cisco vBond Orchestrator software to Cisco SD-WAN Release 20.6.1 or a later release. For more information, see [Upgrade the Software Image on a Device](#) and [Activate a New Software Image](#).
5. Enable maintenance window on Cisco vManage. For more information, see [Configure or Cancel vManage Server Maintenance Window](#).
6. Upgrade the Cisco vSmart Controller software to Cisco SD-WAN Release 20.6.1 or a later release. For more information, see [Upgrade the Software Image on a Device](#) and [Activate a New Software Image](#).
7. Upgrade the Cisco vEdge device software to Cisco SD-WAN Release 20.6.1 or a later release. For more information, see [Upgrade the Software Image on a Device](#) and [Activate a New Software Image](#).



Tip We recommend that you upgrade the WAN edge device software in the same maintenance window. If the WAN edge device software is not upgraded within the OMP graceful restart window, traffic may be lost.

Multitenant Cisco vManage: Disaster Recovery

If a Multitenant Cisco vManage cluster or the data center hosting the Cisco vManage nodes in the cluster fail, you can recover from the failure by activating a standby Cisco vManage cluster. You can perform disaster recovery as follows:

1. Deploy and configure a standby Cisco vManage cluster.

The standby Cisco vManage cluster is not part of the overlay network and is not active.

2. Back up the configuration database of the active Cisco vManage cluster periodically.

Choose a Cisco vManage node in the cluster that hosts the configuration database service and back up the configuration database.

3. If the active Cisco vManage cluster fails, restore the most recent configuration database on the standby Cisco vManage cluster, activate the standby Cisco vManage cluster, and remove the previously active Cisco vManage cluster from the overlay network.

Choose a Cisco vManage node in the cluster that hosts the configuration database service and restore the configuration database backed up from the previously active Cisco vManage cluster.

To test disaster recovery, you can simulate a scenario in which the active Cisco vManage cluster fails. One way to simulate such a failure would be by disabling the tunnel interface as described in this document.

Prerequisites

- The number of Cisco vManage nodes in the active and standby clusters must be identical.
- Each Cisco vManage node in the active and standby clusters must run the same Cisco vManage software release.
- Each Cisco vManage node in the active and standby clusters must be able to connect to the WAN transport IP address of the Cisco vBond Orchestrators in the overlay network.
- Initially, the tunnel interfaces of the Cisco vManage nodes in the standby cluster must be disabled.
- The Cisco vManage nodes in the standby cluster must be certified.
- The clock of every Cisco vManage node in the standby cluster must be synchronized with the clocks of the Cisco SD-WAN controllers and WAN edge devices in the overlay network. If NTP is configured on the overlay, configure the same on the standby Cisco vManage nodes.
- The Cisco vManage nodes in the active and standby clusters should use identical neo4j credentials.

Restrictions

- Do not interrupt any active processes while backing up the configuration database.
- If you wish to enable SD-AVC, you must do so before the restoring the configuration database on standby Cisco vManage node.

Configure a Standby Cisco vManage Cluster

1. Configure the standby Cisco vManage nodes with a similar running configuration as the active Cisco vManage nodes. Install local certificates on the standby Cisco vManage nodes.



Note The running configuration on a standby Cisco vManage is usually identical to that of an active Cisco vManage node. However, you must ensure that settings such as the system IP address and the tunnel interface IP address are unique.

2. On the standby Cisco vManage nodes, shut down the transport interface in VPN 0: On the CLI, include the **shutdown** command in the transport interface configuration.
3. Create a standby cluster using the standby Cisco vManage nodes.

With the standby Cisco vManage nodes configured in this manner, the overlay network is not aware of the standby Cisco vManage cluster.

Back Up the Active Cisco vManage Cluster Configuration

Back up the full configuration database of the active Cisco vManage cluster periodically. Additionally, take snapshots of the active Cisco vManage virtual machines.

1. Choose an active Cisco vManage node that hosts the configuration database service and export a backup of the configuration database. On the CLI of the Cisco vManage node, run the following command: **request nms configuration-db backup path *file-path***

The command backs up the configuration database in a `.tar.gz` file and saves the file in the specified *file-path*.

In the following example, the database is backed up to a file named `db_backup.tar.gz` in the `/home/admin/` directory.

```
Active-vManage# request nms configuration-db backup path /home/admin/db_backup
Successfully saved database to /home/admin/db_backup.tar.gz
```

2. Choose a standby Cisco vManage node that hosts the configuration database service and copy the configuration database backup to this node.

In the following example, `db_backup.tar.gz` is copied from the active Cisco vManage node to the `/home/admin/` directory of a standby Cisco vManage node.

```
Active-vManage# request execute vpn 512 scp /home/admin/db_backup.tar.gz
admin@10.126.93.92:/home/admin
The authenticity of host '10.126.93.92 (10.126.93.92)' can't be established.
ECDSA key fingerprint is SHA256:jTjJWQ0UNHvLrBUxWzNjd8mUz819gPf51MeopsgDlAc.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '10.126.93.92' (ECDSA) to the list of known hosts.
viptela 18.4.5

admin@10.126.93.92's password:
db_backup.tar.gz                               100% 399KB 4.4MB/s 00:00
```

Restore Cisco vManage Cluster Using the Configuration Database Backup

Restore the most recent backup of the configuration database from the active Cisco vManage cluster on the standby Cisco vManage node to which you copied this backup.



Note

- The restore operation does not restore all the information included in the configuration database. Cisco vManage configurations such as users and repositories must be configured on the standby Cisco vManage node after the configuration database is restored using the backup.
- When you complete the steps that follow, the previously active Cisco vManage nodes cannot be reused. To reuse the nodes, you must perform additional steps that are beyond the scope of this document.

1. On the CLI of the standby Cisco vManage node, run the following command: **request nms configuration-db restore path *file-path***

In the following example, the configuration database is restored using the backup file `db_backup.tar.gz`.

```
Standby-vManage# request nms configuration-db restore path /home/admin/db_backup.tar.gz
Configuration database is running in a standalone mode
Importing database...Successfully restored database
```

2. Verify that the appropriate services are running on the standby Cisco vManage nodes: On the CLI of each standby Cisco vManage node, run the **request nms all status** command. From the command output, verify the services running on the node.
3. Verify that every standby Cisco vManage node has a list of all the active and standby Cisco vManage nodes.
 - a. From the Cisco vManage menu, choose **Configuration > Devices > Controllers**.
 - b. Verify that the page displays all active and standby Cisco vManage nodes.
4. On the standby Cisco vManage nodes, enable the transport interface on VPN 0.
Use one of the following two methods:
 - a. Enable the transport interface in VPN 0: On the CLI of each standby Cisco vManage node, run the **no shutdown** command.


```
Active-vManage# config
Active-vManage(config)# vpn 0 interface interface-name
Active-vManage(config-interface)# no shutdown
Active-vManage(config-interface)# commit and-quit
```
 - b. Activate the tunnel interface in VPN 0: On the CLI of each standby Cisco vManage node, run the **tunnel-interface** command.


```
Active-vManage# config
Active-vManage(config)# vpn 0 interface interface-name
Active-vManage(config-interface)# tunnel-interface
Active-vManage(config-interface)# commit and-quit
```
5. Add each standby Cisco vManage node to the overlay network.
 - a. From the Cisco vManage menu, choose **Configuration > Devices**.
 - b. Click **Controllers**.
 - c. For a Cisco vBond Orchestrator, click ... and click **Edit**.
 - d. In the **Edit** dialog box, enter the following details of the Cisco vBond Orchestrator: WAN transport IP address, username, and password.
 - e. Repeat **Step 5c** and **Step 5d** for every Cisco vBond Orchestrator.
6. Disconnect the active Cisco vManage nodes from the overlay network.



Note In a lab environment, where you are simulating a disaster scenario, you can perform this step. However, if you cannot reach Cisco vManage instances in an actual disaster scenario, you may not be able to perform this step and can omit the step.

Use one of the following two methods:

- a. Shut down the transport interface in VPN 0: On the CLI of each active Cisco vManage node, run the **shutdown** command.

```
Active-vManage# config
Active-vManage(config)# vpn 0 interface interface-name
```

```
Active-vManage(config-interface)# shutdown
Active-vManage(config-interface)# commit and-quit
```

- b. Deactivate the tunnel interface in VPN 0: On the CLI of each active Cisco vManage node, run the **no tunnel-interface** command.

```
Active-vManage# config
Active-vManage(config)# vpn 0 interface interface-name
Active-vManage(config-interface)# no tunnel-interface
Active-vManage(config-interface)# commit and-quit
```

7. From the standby Cisco vManage, send the updated controller and device list to the Cisco vBond Orchestrators.

Send the list of controllers:

- a. From the Cisco vManage menu, choose **Configuration > Certificates**.
- b. Click **Controllers**.
- c. Click **Send to vBond**.

Wait for the configuration task to complete. When the task is complete,

- The standby Cisco vManage nodes become the active Cisco vManage nodes.
- The previously active Cisco vManage nodes are no longer part of the overlay network.
- The active Cisco vManage nodes have the configuration from the most recent configuration database backup.
- Every controller establishes connection with the other controllers in the network.

- d. Click **WAN Edge List**.
- e. Click **Send to Controllers**.

8. Verify that the following are intact:

- Policies
- Templates
- Controller and WAN edge device lists

9. Verify the valid Cisco vManage nodes.

- a. Log in to the CLI of each Cisco vBond Orchestrator and run the **show orchestrator valid-vmanage-id** command.

In the command output, verify that the chassis numbers of the active and previously active Cisco vManage nodes are listed.

- b. Log in to the CLI of a WAN edge device and run the **show control valid-vmanage-id** command.
In the command output, verify that the chassis numbers of the active and previously active Cisco vManage nodes are listed. Also, check whether the device is connected to the active Cisco vManage nodes and the Cisco vSmart Controllers.

10. Invalidate the previously active Cisco vManage nodes.



Note After you invalidate the Cisco vManage nodes, the nodes cannot be reused without performing additional steps that are beyond the scope of this document.

- a. From the Cisco vManage menu, choose **Configuration > Certificates**.
 - b. Click **Controllers**.
 - c. For each previously active Cisco vManage node, click ... and click **Invalidate**.
11. Verify the valid Cisco vManage nodes.
 - a. Log in to the CLI of each Cisco vBond Orchestrator and run the **show orchestrator valid-vmanage-id** command.
 In the command output, verify that the chassis numbers of only the active Cisco vManage nodes are listed.
 - b. Log in to the CLI of WAN edge device and run the **show control valid-vmanage-id** command.
 In the command output, verify that the chassis numbers of only the active Cisco vManage nodes are listed. Also, check whether the device is connected to the active Cisco vManage nodes and the Cisco vSmart Controllers.

The Cisco vManage cluster that was initially the standby cluster is now the active Cisco vManage cluster.

Multitenant Cisco vManage: Disaster Recovery in an Overlay Network with Virtual Routers

If a Multitenant Cisco vManage cluster or the data center hosting the Cisco vManage nodes in the cluster fail, you can recover from the failure by activating a standby Cisco vManage cluster. You can perform disaster recovery as follows:

1. Deploy and configure a standby Cisco vManage cluster.
 The standby Cisco vManage cluster is not part of the overlay network and is not active.
2. Back up the configuration database of the active Cisco vManage cluster periodically.
 Choose a Cisco vManage node in the cluster that hosts the configuration database service and back up the configuration database.
3. If the active Cisco vManage cluster fails, restore the most recent configuration database on the standby Cisco vManage cluster, activate the standby Cisco vManage cluster, and remove the previously active Cisco vManage cluster from the overlay network.
 Choose a Cisco vManage node in the cluster that hosts the configuration database service and restore the configuration database backed up from the previously active Cisco vManage cluster.

To test disaster recovery, you can simulate a scenario in which the active Cisco vManage cluster fails. One way to simulate such a failure would be by disabling the tunnel interface as described in this document.

The following disaster recovery procedure applies to an overlay network in which Cisco vEdge Cloud routers are deployed at branch locations.

Prerequisites

- The number of Cisco vManage nodes in the active and standby clusters must be identical.
- Each Cisco vManage node in the active and standby clusters must run the same Cisco vManage software release.
- Each Cisco vManage node in the active and standby clusters must be able to connect to the WAN transport IP address of the Cisco vBond Orchestrators in the overlay network.
- Initially, the tunnel interfaces of the Cisco vManage nodes in the standby cluster must be disabled.
- The Cisco vManage nodes in the standby cluster must be certified.
- The clock of every Cisco vManage node in the standby cluster must be synchronized with the clocks of the Cisco SD-WAN controllers and WAN edge devices in the overlay network. If NTP is configured on the overlay, configure the same on the standby Cisco vManage nodes.
- The Cisco vManage nodes in the active and standby clusters should use identical neo4j credentials.

Restrictions

- Do not interrupt any active processes while backing up the configuration database.
- If you wish to enable SD-AVC, you must do so before restoring the configuration database on standby Cisco vManage node.

Configure a Standby Cisco vManage Cluster

1. Configure the standby Cisco vManage nodes with a similar running configuration as the active Cisco vManage nodes. Install local certificates on the standby Cisco vManage nodes.



Note The running configuration on a standby Cisco vManage is usually identical to that of an active Cisco vManage node. However, you must ensure that settings such as the system IP address and the tunnel interface IP address are unique.

2. On the standby Cisco vManage nodes, shut down the transport interface in VPN 0: On the CLI, include the **shutdown** command in the transport interface configuration.
3. Create a standby cluster using the standby Cisco vManage nodes.

With the standby Cisco vManage nodes configured in this manner, the overlay network is not aware of the standby Cisco vManage cluster.

Back Up the Active Cisco vManage Cluster Configuration

Back up the full configuration database of the active Cisco vManage cluster periodically. Additionally, take snapshots of the active Cisco vManage virtual machines.

1. Choose an active Cisco vManage node that hosts the configuration database service and export a backup of the configuration database. On the CLI of the Cisco vManage node, run the following command: **request nms configuration-db backup path *file-path***

The command backs up the configuration database in a `.tar.gz` file and saves the file in the specified *file-path*.

In the following example, the database is backed up to a file named `db_backup.tar.gz` in the `/home/admin/` directory.

```
Active-vManage# request nms configuration-db backup path /home/admin/db_backup
Successfully saved database to /home/admin/db_backup.tar.gz
```

2. Choose a standby Cisco vManage node that hosts the configuration database service and copy the configuration database backup to this node.

In the following example, `db_backup.tar.gz` is copied from the active Cisco vManage node to the `/home/admin/` directory of a standby Cisco vManage node.

```
Active-vManage# request execute vpn 512 scp /home/admin/db_backup.tar.gz
admin@10.126.93.92:/home/admin
The authenticity of host '10.126.93.92 (10.126.93.92)' can't be established.
ECDSA key fingerprint is SHA256:jTjJWQ0UNHv1rBUxWzNjd8mUz819gPf51MeopsgDlAc.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '10.126.93.92' (ECDSA) to the list of known hosts.
viptela 18.4.5

admin@10.126.93.92's password:
db_backup.tar.gz                               100% 399KB 4.4MB/s 00:00
```

Restore Cisco vManage Cluster Using the Configuration Database Backup

Restore the most recent backup of the configuration database from the active Cisco vManage cluster on the standby Cisco vManage node to which you copied this backup.



Note

- The restore operation does not restore all the information included in the configuration database. Cisco vManage configurations such as users and repositories must be configured on the standby Cisco vManage node after the configuration database is restored using the backup.
- When you complete the steps that follow, the previously active Cisco vManage nodes cannot be reused. To reuse the nodes, you must perform additional steps that are beyond the scope of this document.

1. On the CLI of the standby Cisco vManage node, run the following command: **request nms configuration-db restore path *file-path***

In the following example, the configuration database is restored using the backup file `db_backup.tar.gz`.

```
Standby-vManage# request nms configuration-db restore path /home/admin/db_backup.tar.gz
Configuration database is running in a standalone mode
Importing database...Successfully restored database
```

2. Verify that the appropriate services are running on the standby Cisco vManage nodes: On the CLI of each standby Cisco vManage node, run the **request nms all status** command. From the command output, verify the services running on the node.

3. Verify that every standby Cisco vManage node has a list of all the active and standby Cisco vManage nodes.
 - a. From the Cisco vManage menu, choose **Configuration > Devices > Controllers**.
 - b. Verify that the page displays all active and standby Cisco vManage nodes.
4. Log in to the CLI of each Cisco vBond Orchestrator and run the **show orchestrator valid-vmanage-id** command.
In the command output, verify that the chassis numbers of the active and previously active Cisco vManage nodes are listed.
5. Log in to the CLI of Cisco vEdge Cloud router and run the **show control valid-vmanage-id** command.
In the command output, verify that the chassis numbers of the active and previously active Cisco vManage nodes are listed.
6. On the standby Cisco vManage nodes, enable the transport interface on VPN 0.
Use one of the following two methods:
 - a. Enable the transport interface in VPN 0: On the CLI of each standby Cisco vManage node, run the **no shutdown** command.

```
Active-vManage# config
Active-vManage(config)# vpn 0 interface interface-name
Active-vManage(config-interface)# no shutdown
Active-vManage(config-interface)# commit and-quit
```
 - b. Activate the tunnel interface in VPN 0: On the CLI of each standby Cisco vManage node, run the **tunnel-interface** command.

```
Active-vManage# config
Active-vManage(config)# vpn 0 interface interface-name
Active-vManage(config-interface)# tunnel-interface
Active-vManage(config-interface)# commit and-quit
```
7. Add each standby Cisco vManage node to the overlay network.
 - a. From the Cisco vManage menu, choose **Configuration > Devices**.
 - b. Click **Controllers**.
 - c. For a Cisco vBond Orchestrator, click **...** and click **Edit**.
 - d. In the **Edit** dialog box, enter the following details of the Cisco vBond Orchestrator: WAN transport IP address, username, and password.
 - e. Repeat **Step 7c** and **Step 7d** for every Cisco vBond Orchestrator.
8. Disconnect the active Cisco vManage nodes from the overlay network.



Note In a lab environment, where you are simulating a disaster scenario, you can perform this step. However, if you cannot reach Cisco vManage instances in an actual disaster scenario, you may not be able to perform this step and can omit this step.

Use one of the following two methods:

- a. Shut down the transport interface in VPN 0: On the CLI of each active Cisco vManage node, run the **shutdown** command.

```
Active-vManage# config
Active-vManage(config)# vpn 0 interface interface-name
Active-vManage(config-interface)# shutdown
Active-vManage(config-interface)# commit and-quit
```

- b. Deactivate the tunnel interface in VPN 0: On the CLI of each active Cisco vManage node, run the **no tunnel-interface** command.

```
Active-vManage# config
Active-vManage(config)# vpn 0 interface interface-name
Active-vManage(config-interface)# no tunnel-interface
Active-vManage(config-interface)# commit and-quit
```

9. From the standby Cisco vManage, send the updated controller and device list to the Cisco vBond Orchestrators.

Send the list of controllers:

- a. From the Cisco vManage menu, choose **Configuration > Certificates**.
- b. Click **Controllers**.
- c. Click **Send to vBond**.

Wait for the configuration task to complete. When the task is complete,

- The standby Cisco vManage nodes become the active Cisco vManage nodes.
- The previously active Cisco vManage nodes are no longer part of the overlay network.
- The active Cisco vManage nodes have the configuration from the most recent configuration database backup.
- Every controller establishes connection with the other controllers in the network.

- d. Click **WAN Edge List**.
- e. Click **Send to Controllers**.

10. Verify that the following are intact:

- Policies
- Templates
- Controller and WAN edge device lists

11. Verify the valid Cisco vManage nodes.

- a. Log in to the CLI of each Cisco vBond Orchestrator and run the **show orchestrator valid-vmanage-id** command.

In the command output, verify that the chassis numbers of the active and previously active Cisco vManage nodes are listed.

- b. Log in to the CLI of a Cisco vEdge Cloud router and run the **show control valid-vmanage-id** command.

In the command output, verify that the chassis numbers of the active and previously active Cisco vManage nodes are listed. Also, check whether the device is connected to the active Cisco vManage nodes and the Cisco vSmart Controllers.

12. Invalidate the previously active Cisco vManage nodes.

The previously active Cisco vManage is the certificate issuer for the cloud WAN edge devices. The active Cisco vManage issues certificates to the cloud WAN edge devices only after the previously active Cisco vManage nodes are invalidated.



Note

- After you invalidate the Cisco vManage nodes, the nodes cannot be reused without performing additional steps that are beyond the scope of this document.
- When you invalidate the previously active Cisco vManage nodes, Cisco vManage marks the nodes as invalid and sends an update to all controllers. However, Cisco vManage does not send an updated list of valid Cisco vManage UUIDs to Cisco vBond Orchestrators immediately because the previously active Cisco vManage is the CA for the cloud WAN edge devices. So, the output of the **show orchestrator valid-vmanage-id** command on a Cisco vBond Orchestrator includes the UUIDs of the invalidated Cisco vManage nodes.

Cisco vManage has a scheduled task that runs every 24 hours and checks to see if all the cloud WAN edges have been moved to the active Cisco vManage. Cisco vManage sends the updated list of valid Cisco vManage UUIDs to Cisco vBond Orchestrator only after the cloud WAN edge devices have been moved to the active Cisco vManage. After this list is received, the output of the **show orchestrator valid-vmanage-id** command on a Cisco vBond Orchestrator does not include the UUIDs of the invalidated Cisco vManage nodes.

- a. From the Cisco vManage menu, choose **Configuration > Certificates**.
- b. Click **Controllers**.
- c. For each previously active Cisco vManage node, click ... and click **Invalidate**.

13. Verify the valid Cisco vManage nodes after 24 hours.

- a. Log in to the CLI of each Cisco vBond Orchestrator and run the **show orchestrator valid-vmanage-id** command.

In the command output, verify that the chassis numbers of only the active Cisco vManage nodes are listed.

- b. Log in to the CLI of WAN edge device and run the **show control valid-vmanage-id** command.

In the command output, verify that the chassis numbers of only the active Cisco vManage nodes are listed. Also, check whether the device is connected to the active Cisco vManage nodes and the Cisco vSmart Controllers.

The Cisco vManage cluster that was initially the standby cluster is now the active Cisco vManage cluster.

Multitenant Cisco vManage: Disaster Recovery After a Failed Data Center Becomes Operational

If a Multitenant Cisco vManage cluster or the data center hosting the Cisco vManage nodes in the cluster fail, you can recover from the failure by activating a standby Cisco vManage cluster. You can perform disaster recovery as follows:

1. Deploy and configure a standby Cisco vManage cluster.

The standby Cisco vManage cluster is not part of the overlay network and is not active.

2. Back up the configuration database of the active Cisco vManage cluster periodically.

Choose a Cisco vManage node in the cluster that hosts the configuration database service and back up the configuration database.

3. If the active Cisco vManage cluster fails, restore the most recent configuration database on the standby Cisco vManage cluster, activate the standby Cisco vManage cluster, and remove the previously active Cisco vManage cluster from the overlay network.

Choose a Cisco vManage node in the cluster that hosts the configuration database service and restore the configuration database backed up from the previously active Cisco vManage cluster.

To test disaster recovery, you can simulate a scenario in which the active Cisco vManage cluster fails. One way to simulate such a failure would be by disabling the tunnel interface as described in this document.

The following procedure applies to a scenario in which an initially active Cisco vManage cluster or the data center hosting the cluster failed and the standby Cisco vManage cluster was configured to be the active Cisco vManage cluster. If the cluster that was initially active becomes operational again, it serves as a standby cluster. By completing the following procedure, you can turn this standby cluster into the active cluster.

Check the Configuration of the Standby vManage NMS

1. Check whether the running configuration of the standby Cisco vManage nodes is similar to the running configuration of the active Cisco vManage nodes. Local certificates must be installed on the standby Cisco vManage nodes.



Note The running configuration on a standby Cisco vManage is usually identical to that of an active Cisco vManage node. However, you must ensure that settings such as the system IP address and the tunnel interface IP address are unique.

2. On the standby Cisco vManage nodes, shut down the transport interface in VPN 0: On the CLI, include the **shutdown** command in the transport interface configuration.

With the standby Cisco vManage nodes configured in this manner, the overlay network is not aware of the standby Cisco vManage cluster.

Back Up the Active Cisco vManage Cluster Configuration

Back up the full configuration database of the active Cisco vManage cluster periodically. Additionally, take snapshots of the active Cisco vManage virtual machines.

1. Choose an active Cisco vManage node that hosts the configuration database service and export a backup of the configuration database. On the CLI of the Cisco vManage node, run the following command: **request nms configuration-db backup path *file-path***

The command backs up the configuration database in a `.tar.gz` file and saves the file in the specified *file-path*.

In the following example, the database is backed up to a file named `db_backup.tar.gz` in the `/home/admin/` directory.

```
Active-vManage# request nms configuration-db backup path /home/admin/db_backup
Successfully saved database to /home/admin/db_backup.tar.gz
```

2. Choose a standby Cisco vManage node that hosts the configuration database service and copy the configuration database backup to this node.

In the following example, `db_backup.tar.gz` is copied from the active Cisco vManage node to the `/home/admin/` directory of a standby Cisco vManage node.

```
Active-vManage# request execute vpn 512 scp /home/admin/db_backup.tar.gz
admin@10.126.93.92:/home/admin
The authenticity of host '10.126.93.92 (10.126.93.92)' can't be established.
ECDSA key fingerprint is SHA256:jTjJWQ0UNHvLrBUxWzNjd8mUz819gPf51MeopsgDlAc.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '10.126.93.92' (ECDSA) to the list of known hosts.
viptela 18.4.5

admin@10.126.93.92's password:
db_backup.tar.gz                               100% 399KB 4.4MB/s 00:00
```

Restore Cisco vManage Cluster Using the Configuration Database Backup

Restore the most recent backup of the configuration database from the active Cisco vManage cluster on the standby Cisco vManage node to which you copied this backup.



Note

- The restore operation does not restore all the information included in the configuration database. Cisco vManage configurations such as users and repositories must be configured on the standby Cisco vManage node after the configuration database is restored using the backup.
- When you complete the steps that follow, the previously active Cisco vManage nodes cannot be reused. To reuse the nodes, you must perform additional steps that are beyond the scope of this document.

1. On the CLI of the standby Cisco vManage node, run the following command: **request nms configuration-db restore path *file-path***

In the following example, the configuration database is restored using the backup file `db_backup.tar.gz`.

```
Standby-vManage# request nms configuration-db restore path /home/admin/db_backup.tar.gz
Configuration database is running in a standalone mode
Importing database...Successfully restored database
```

2. Verify that the appropriate services are running on the standby Cisco vManage nodes: On the CLI of each standby Cisco vManage node, run the **request nms all status** command. From the command output, verify the services running on the node.
3. Verify that every standby Cisco vManage node has a list of all the active and standby Cisco vManage nodes.
 - a. From the Cisco vManage menu, choose **Configuration > Devices > Controllers**.
 - b. Verify that the page displays all active and standby Cisco vManage nodes.
4. On the standby Cisco vManage nodes, enable the transport interface on VPN 0.
Use one of the following two methods:
 - a. Enable the transport interface in VPN 0: On the CLI of each standby Cisco vManage node, run the **no shutdown** command.


```
Active-vManage# config
Active-vManage (config)# vpn 0 interface interface-name
Active-vManage (config-interface)# no shutdown
Active-vManage (config-interface)# commit and-quit
```
 - b. Activate the tunnel interface in VPN 0: On the CLI of each standby Cisco vManage node, run the **tunnel-interface** command.


```
Active-vManage# config
Active-vManage (config)# vpn 0 interface interface-name
Active-vManage (config-interface)# tunnel-interface
Active-vManage (config-interface)# commit and-quit
```
5. Add each standby Cisco vManage node to the overlay network.
 - a. From the Cisco vManage menu, choose **Configuration > Devices**.
 - b. Click **Controllers**.
 - c. For a Cisco vBond Orchestrator, click ... and click **Edit**.
 - d. In the **Edit** dialog box, enter the following details of the Cisco vBond Orchestrator: WAN transport IP address, username, and password.
 - e. Repeat **Step 5c** and **Step 5d** for every Cisco vBond Orchestrator.
6. Disconnect the active Cisco vManage nodes from the overlay network.



Note In a lab environment, where you are simulating a disaster scenario, you can perform this step. However, if you cannot reach Cisco vManage instances in an actual disaster scenario, you may not be able to perform this step and can omit the step.

Use one of the following two methods:

- a. Shut down the transport interface in VPN 0: On the CLI of each active Cisco vManage node, run the **shutdown** command.

```
Active-vManage# config
Active-vManage (config)# vpn 0 interface interface-name
```

```
Active-vManage(config-interface)# shutdown
Active-vManage(config-interface)# commit and-quit
```

- b. Deactivate the tunnel interface in VPN 0: On the CLI of each active Cisco vManage node, run the **no tunnel-interface** command.

```
Active-vManage# config
Active-vManage(config)# vpn 0 interface interface-name
Active-vManage(config-interface)# no tunnel-interface
Active-vManage(config-interface)# commit and-quit
```

7. From the standby Cisco vManage, send the updated controller and device list to the Cisco vBond Orchestrators.

Send the list of controllers:

- a. From the Cisco vManage menu, choose **Configuration > Certificates**.
- b. Click **Controllers**.
- c. Click **Send to vBond**.

Wait for the configuration task to complete. When the task is complete,

- The standby Cisco vManage nodes become the active Cisco vManage nodes.
- The previously active Cisco vManage nodes are no longer part of the overlay network.
- The active Cisco vManage nodes have the configuration from the most recent configuration database backup.
- Every controller establishes connection with the other controllers in the network.

- d. Click **WAN Edge List**.
- e. Click **Send to Controllers**.

8. Verify that the following are intact:

- Policies
- Templates
- Controller and WAN edge device lists

9. Verify the valid Cisco vManage nodes.

- a. Log in to the CLI of each Cisco vBond Orchestrator and run the **show orchestrator valid-vmanage-id** command.

In the command output, verify that the chassis numbers of the active and previously active Cisco vManage nodes are listed.

- b. Log in to the CLI of a WAN edge device and run the **show control valid-vmanage-id** command.

In the command output, verify that the chassis numbers of the active and previously active Cisco vManage nodes are listed. Also, check whether the device is connected to the active Cisco vManage nodes and the Cisco vSmart Controllers.

10. Invalidate the previously active Cisco vManage nodes.



Note After you invalidate the Cisco vManage nodes, the nodes cannot be reused without performing additional steps that are beyond the scope of this document.

- a. From the Cisco vManage menu, choose **Configuration > Certificates**.
 - b. Click **Controllers**.
 - c. For each previously active Cisco vManage node, click ... and click **Invalidate**.
11. Verify the valid Cisco vManage nodes.
- a. Log in to the CLI of each Cisco vBond Orchestrator and run the **show orchestrator valid-vmanage-id** command.
 In the command output, verify that the chassis numbers of only the active Cisco vManage nodes are listed.
 - b. Log in to the CLI of WAN edge device and run the **show control valid-vmanage-id** command.
 In the command output, verify that the chassis numbers of only the active Cisco vManage nodes are listed. Also, check whether the device is connected to the active Cisco vManage nodes and the Cisco vSmart Controllers.

The Cisco vManage cluster that was initially the standby cluster is now the active Cisco vManage cluster.

Replace Faulty Cisco vSmart Controller

To replace a faulty Cisco vSmart Controller with a new instance, follow these steps:

1. Create a Cisco vSmart Controller instance. See [Deploy the Cisco vSmart Controller](#).
2. [Add Cisco vSmart Controller](#) to the overlay network.
3. From the Cisco vManage menu, choose **Configuration > Devices**.
4. Click **Controllers**.
5. For the faulty Cisco vSmart Controllers, click ... and click **Invalidate**.

The **Invalidate** dialog box appears.



Note If you have not added a new Cisco vSmart Controller that can replace the faulty Cisco vSmart Controller, Cisco vManage indicates this through an error message. Click **Cancel** in the **Invalidate** dialog box and add a new Cisco vSmart Controller before invalidating the faulty instance.

6. In the **Invalidate** dialog box, do the following:
 - a. Check the **Replace vSmart** check box.
 - b. From the **Select vSmart** drop-down list, choose the new Cisco vSmart Controller that should replace the faulty instance.

- c. Click **Invalidate**.

Cisco vManage launches the **Invalidate Device** and **Push CLI Template Configuration** task. When these tasks are completed, the faulty Cisco vSmart Controller is invalidated and removed from the overlay network. The tenants that were served by the faulty Cisco vSmart Controller are now served by the new Cisco vSmart Controller that you chose as the replacement.



CHAPTER 16

Flexible Tenant Placement on Multitenant Cisco vSmart Controllers

Table 114: Feature History

Feature Name	Release Information	Description
Flexible Tenant Placement on Multitenant Cisco vSmart Controllers	Cisco vManage Release 20.9.1	With this feature, while onboarding a tenant to a multitenant deployment, you can choose the pair of multitenant Cisco vSmart Controllers that serve the tenant. After onboarding a tenant, you can migrate the tenant to a different pair of multitenant Cisco vSmart Controllers, if necessary.

- [Information About Flexible Tenant Placement on Multitenant Cisco vSmart Controllers, on page 411](#)
- [Restrictions for Flexible Tenant Placement on Multitenant Cisco vSmart Controllers, on page 413](#)
- [Assign Cisco vSmart Controllers to Tenants During Onboarding, on page 413](#)
- [Update Cisco vSmart Controllers Placement For a Tenant, on page 418](#)

Information About Flexible Tenant Placement on Multitenant Cisco vSmart Controllers

Automatic Tenant Placement by Cisco vManage

In Cisco vManage Release 20.8.x and earlier releases, when you onboard a tenant, Cisco vManage assigns a pair of multitenant Cisco vSmart Controllers to the tenant based on an internal algorithm that considers factors such as the following:

- number of tenant WAN edge devices that you forecast for the tenant
- number of tenants served by a pair of multitenant Cisco vSmart Controllers
- number of WAN edge devices connected to a pair of multitenant Cisco vSmart Controllers

After the tenant is onboarded, if the tenant needs to add more devices than you originally forecast, you can modify the forecast if the pair of multitenant Cisco vSmart Controllers serving the tenant can accommodate these additional WAN edge devices. If the Cisco vSmart Controllers cannot accommodate the additional WAN edge devices, you must delete the tenant and onboard the tenant again with the revised device forecast so that Cisco vManage assigns a suitable pair of Cisco vSmart Controllers. If none of the pairs of multitenant Cisco vSmart Controllers can accommodate the revised device forecast, add a new pair of Cisco vSmart Controllers and then onboard the tenant.

Flexible Tenant Placement by Provide Admin User

From Cisco vManage Release 20.9.1, while onboarding a tenant, you have the flexibility to choose the pair of multitenant Cisco vSmart Controllers that are assigned to the tenant. Automatic tenant placement by Cisco vManage continues to be the default behavior with flexible tenant placement as an optional configuration.

To help you with flexible tenant placement, Cisco vManage lists available multitenant Cisco vSmart Controllers and provides the following details, as a percentage, for each controller:

- number of tenants assigned
- number of tenant WAN edge devices connected
- memory utilized
- CPU utilized

A multitenant Cisco vSmart Controller can serve a maximum of 24 tenants and 1000 tenant WAN edge devices across all the tenants. While onboarding a tenant, choose a pair of controllers that can be assigned one more tenant and can also connect to the number of WAN edge devices forecast for the tenant.

After the tenant is onboarded, if the tenant needs to add more devices than you originally forecast and the assigned pair of multitenant Cisco vSmart Controllers cannot connect to these additional WAN edge devices, you can migrate the tenant to another pair of Cisco vSmart Controllers that can serve one more tenant and accommodate the revised WAN edge device forecast for the tenant. If none of the multitenant Cisco vSmart Controllers pairs can accommodate the revised device forecast, you can migrate other tenants to alternative Cisco vSmart Controllers so that the controller utilization is efficient and the tenant assignment is optimal. If the optimization doesn't create the capacity required to accommodate the revised device forecast for the tenant, add a new pair of Cisco vSmart Controllers and then migrate the tenant.

Benefits of Flexible Tenant Placement on Multitenant Cisco vSmart Controllers

- Choose Cisco vSmart Controllers deployed in different failure zones to reduce the probability of both the controllers failing simultaneously. In a cloud environment, choose controllers deployed in different regions.
- Choose Cisco vSmart Controllers deployed in the same geographical region as the tenant WAN edge devices to reduce latency.
- Choose Cisco vSmart Controllers based on the CPU, DRAM, and hard disk resources allocated, and the utilization of these resources.
- Migrate a tenant to a different Cisco vSmart Controller to accommodate changes in the tenant device forecast.

Restrictions for Flexible Tenant Placement on Multitenant Cisco vSmart Controllers

If you wish to migrate a tenant to different pair of Cisco vSmart Controllers, you must change the Cisco vSmart Controllers assigned to the tenant one at a time. Doing so ensures that one of the Cisco vSmart Controllers is available to the tenant WAN edge devices during the migration and prevents disruptions in traffic.

Assign Cisco vSmart Controllers to Tenants During Onboarding

Prerequisites

- At least two Cisco vSmart Controllers must be operational and in the vManage mode before you can add new tenants.

A Cisco vSmart Controller enters the **vManage** mode when you push a template to the controller from Cisco vManage. A Cisco vSmart Controller in the **CLI** mode cannot serve multiple tenants.

- Each pair of Cisco vSmart Controllers can serve a maximum of 24 tenants and a maximum of 1000 tenant devices. Ensure that there are at least two Cisco vSmart Controllers that can serve a new tenant. If no pair of Cisco vSmart Controllers in the deployment can serve a new tenant, add two Cisco vSmart Controllers and change their mode to **vManage**.
- Add up to 16 tenants in a single operation. If you add more than one tenant, during the **Add Tenant** task, Cisco vManage adds the tenants one after another and not in parallel.

While an **Add Tenant** task is in progress, do not perform a second tenant addition operation. If you do so, the second Add Tenant task fails.

- Each tenant must have a unique Virtual Account (VA) on Plug and Play Connect on [Cisco Software Central](#). The tenant VA should belong to the same Smart Account (SA) as the provider VA.
- For an on-premises deployment, create a Cisco vBond Orchestrator controller profile for the tenant on Plug and Play Connect. The fields in the following table are mandatory.

Field	Description
Profile Name	Enter a name for the controller profile.
Multi-Tenancy	From the drop-down list, select Yes .
SP Organization Name	Enter the provider organization name.
Organization Name	Enter the tenant organization name in the format <SP Org Name>-<Tenant Org Name>. The organization name can be up to 64 characters.
Primary Controller	Enter the host details for the primary Cisco vBond Orchestrator.

For a cloud deployment, the Cisco vBond Orchestrator controller profile is created automatically as part of the tenant creation process.

1. Log in to Cisco vManage as the provider **admin** user.
2. From the Cisco vManage menu, choose **Administration > Tenant Management**.
3. Click **Add Tenant**.
4. In the **Add Tenant** slide-in pane, click **New Tenant**.
5. Configure the following tenant details:

Field	Description
Name	Enter a name for the tenant. For a cloud deployment, the tenant name should be same as the tenant VA name on Plug and Play Connect.
Description	Enter a description for the tenant. The description can have up to 256 characters and can contain only alphanumeric characters.
Organization Name	Enter the name of the tenant organization. The organization name can have up to 64 characters. The organization name is case-sensitive. Each tenant or customer must have a unique organization name. Enter the organization name in the following format: <SP Org Name>-<Tenant Org Name> For example, if the provider organization name is 'managed-sp' and the tenant organization name is 'customer1', while adding the tenant, enter the organization name as 'managed-sp-customer1'.

Field	Description
URL Subdomain	

Field	Description
	<p>Enter the fully qualified subdomain name of the tenant.</p> <ul style="list-style-type: none"> The subdomain name must include the domain name of the service provider. For example, for the managed-sp.com service provider, a valid domain name for customer1 is customer1.managed-sp.com. <p>Note The service provider name is shared amongst all tenants. Ensure that the URL naming convention follows the same domain name convention that was followed while enabling multitenancy using Administration > Settings > Tenancy Mode.</p> <ul style="list-style-type: none"> For an on-premises deployment, add the fully qualified subdomain name of the tenant to the DNS. Map the fully qualified subdomain name to the IP addresses of the three Cisco vManage instances in the Cisco vManage cluster. <ul style="list-style-type: none"> Provider DNS: Create a DNS A record and map it to the IP addresses of the Cisco vManage instances running in the Cisco vManage cluster. The A record is derived from the provider's domain name and the cluster ID that was created while enabling multitenancy on Cisco vManage. For example, if the provider's domain name is <code>sdwan.cisco.com</code> and the cluster ID is <code>vmanage123</code>, configure the A record as <code>vmanage123.sdwan.cisco.com</code>. <p>Note If you fail to add the DNS A record, you will experience authentication errors when logging in to Cisco vManage.</p> <p>Validate that the DNS is configured correctly by using the nslookup command. Example: <code>nslookup vmanage123.sdwan.cisco.com</code>.</p> Tenant DNS: Create DNS CNAME records for each tenant that you created and map them to the provider FQDN. For example, if the provider's domain name is <code>sdwan.cisco.com</code> and tenant name is <code>customer1</code>, configure the CNAME record as <code>customer1.sdwan.cisco.com</code>. <p>Cluster ID is not required in the CNAME record.</p> <p>Validate that the DNS is configured correctly by using the nslookup command. Example: <code>nslookup customer1.sdwan.cisco.com</code>. For a cloud deployment, the fully qualified subdomain name of the tenant is automatically added to the DNS as part of the tenant creation process. After you add a tenant, it could take </p>

Field	Description								
	up to an hour before the fully qualified subdomain name of the tenant can be resolved by the DNS.								
Forecasted Devices	<p>Enter the number of WAN edge devices that the tenant can add to the overlay.</p> <p>If the tenant tries to add WAN edge devices beyond this number, Cisco vManage reports an error and the device addition fails.</p>								
Select two vSmarts	<ul style="list-style-type: none"> Automatic tenant placement: Ensure that the Select two vSmarts field has the value Autoplacement. This is the default configuration. Flexible tenant placement: <ul style="list-style-type: none"> a. Click the Select two vSmarts drop-down list. <p>Cisco vManage lists the hostnames of the available Cisco vSmart Controllers. For each Cisco vSmart Controller, Cisco vManage shows whether the controller is reachable and reports the following utilization details:</p> <table border="1"> <tbody> <tr> <td>Tenant hosting capacity</td> <td>Each Cisco vSmart Controller can serve a maximum of 24 tenants. Tenant hosting capacity represents the number of tenants to which the Cisco vSmart Controller is assigned in the form of a percentage. This value indicates whether you can assign another tenant to this controller.</td> </tr> <tr> <td>Used device capacity</td> <td>Each Cisco vSmart Controller can support a maximum of 1000 tenant WAN edge devices. Used device capacity represents the number of tenant WAN edge devices connected to the Cisco vSmart Controller in the form of a percentage of the maximum capacity (1000 WAN edge devices). This value indicates whether the Cisco vSmart Controller can support the number of devices forecast for the tenant that you are onboarding.</td> </tr> <tr> <td>Memory utilized</td> <td>This value represents memory consumption as a percentage.</td> </tr> <tr> <td>CPU utilized</td> <td>This value represents CPU usage as a percentage.</td> </tr> </tbody> </table> b. Select two Cisco vSmart Controllers to assign to the tenant based on the utilization details. <p>To select a Cisco vSmart Controller, check the check box adjacent to its hostname.</p> 	Tenant hosting capacity	Each Cisco vSmart Controller can serve a maximum of 24 tenants. Tenant hosting capacity represents the number of tenants to which the Cisco vSmart Controller is assigned in the form of a percentage. This value indicates whether you can assign another tenant to this controller.	Used device capacity	Each Cisco vSmart Controller can support a maximum of 1000 tenant WAN edge devices. Used device capacity represents the number of tenant WAN edge devices connected to the Cisco vSmart Controller in the form of a percentage of the maximum capacity (1000 WAN edge devices). This value indicates whether the Cisco vSmart Controller can support the number of devices forecast for the tenant that you are onboarding.	Memory utilized	This value represents memory consumption as a percentage.	CPU utilized	This value represents CPU usage as a percentage.
Tenant hosting capacity	Each Cisco vSmart Controller can serve a maximum of 24 tenants. Tenant hosting capacity represents the number of tenants to which the Cisco vSmart Controller is assigned in the form of a percentage. This value indicates whether you can assign another tenant to this controller.								
Used device capacity	Each Cisco vSmart Controller can support a maximum of 1000 tenant WAN edge devices. Used device capacity represents the number of tenant WAN edge devices connected to the Cisco vSmart Controller in the form of a percentage of the maximum capacity (1000 WAN edge devices). This value indicates whether the Cisco vSmart Controller can support the number of devices forecast for the tenant that you are onboarding.								
Memory utilized	This value represents memory consumption as a percentage.								
CPU utilized	This value represents CPU usage as a percentage.								

- To save the tenant configuration, click **Save**.
- To add another tenant, repeat Step 4 to Step 6.

8. To onboard tenants to the deployment, click **Add**.

Cisco vManage initiates the Create Tenant Bulk task to onboard the tenants.

As part of this task, Cisco vManage performs the following activities:

- creates the tenant
- assigns two Cisco vSmart Controllers to serve the tenant and pushes a CLI template to these controllers to configure tenant information
- sends the tenant and Cisco vSmart Controller information to Cisco vBond Orchestrators

When the task is successfully completed, you can view the tenant information, including the Cisco vSmart Controllers assigned to the tenant, on the **Administration > Tenant Management** page.

Update Cisco vSmart Controllers Placement For a Tenant

You can migrate a tenant to a different pair of Cisco vSmart Controllers from the controllers that are currently assigned to the tenant. For instance, if you need to increase the tenant WAN edge device forecast and the controllers assigned to the tenant cannot connect to these revised number of tenant WAN edge devices, you can migrate the tenant to a pair of controllers that can accommodate the revised forecast.

If you wish to migrate a tenant to different pair of Cisco vSmart Controllers, you must change the Cisco vSmart Controllers that are assigned to the tenant one at a time. Doing so ensures that one of the Cisco vSmart Controllers is available to the tenant WAN edge devices during the migration and prevents disruptions in traffic.

1. Log in to Cisco vManage as the provider **admin** user.
2. From the Cisco vManage menu, choose **Administration > Tenant Management**.
3. For the tenant you wish to migrate to a different controller, click **...** adjacent to the tenant organization name.
4. Click **Update vSmart Placement**.
5. In the **Update vSmart Placement** slide-in pane, configure the following:

Field	Description								
Source vSmart (currently applied)	<p data-bbox="833 285 1520 317">a. Click the Source vSmart (currently applied) drop-down list.</p> <p data-bbox="873 331 1503 457">Cisco vManage lists the hostnames of the Cisco vSmart Controllers assigned to the tenant. For each Cisco vSmart Controller, Cisco vManage shows whether the controller is reachable and reports the following utilization details:</p> <table border="1" data-bbox="873 478 1624 1115"> <tbody> <tr> <td data-bbox="873 478 1068 684">Tenant hosting capacity</td> <td data-bbox="1068 478 1624 684">Each Cisco vSmart Controller can serve a maximum of 24 tenants. Tenant hosting capacity represents the number of tenants to which the Cisco vSmart Controller is assigned in the form of a percentage. This value indicates whether you can assign another tenant to this controller.</td> </tr> <tr> <td data-bbox="873 684 1068 982">Used device capacity</td> <td data-bbox="1068 684 1624 982">Each Cisco vSmart Controller can support a maximum of 1000 tenant WAN edge devices. Used device capacity represents the number of tenant WAN edge devices connected to the Cisco vSmart Controller in the form of a percentage of the maximum capacity (1000 devices). This value indicates whether the Cisco vSmart Controller can support the number of devices forecast for the tenant that you are onboarding.</td> </tr> <tr> <td data-bbox="873 982 1068 1066">Memory utilized</td> <td data-bbox="1068 982 1624 1066">This value represents memory consumption as a percentage.</td> </tr> <tr> <td data-bbox="873 1066 1068 1115">CPU utilized</td> <td data-bbox="1068 1066 1624 1115">This value represents CPU usage as a percentage.</td> </tr> </tbody> </table> <p data-bbox="833 1136 1520 1199">b. Check the check box adjacent to the hostname of one of the Cisco vSmart Controllers assigned to the tenant.</p>	Tenant hosting capacity	Each Cisco vSmart Controller can serve a maximum of 24 tenants. Tenant hosting capacity represents the number of tenants to which the Cisco vSmart Controller is assigned in the form of a percentage. This value indicates whether you can assign another tenant to this controller.	Used device capacity	Each Cisco vSmart Controller can support a maximum of 1000 tenant WAN edge devices. Used device capacity represents the number of tenant WAN edge devices connected to the Cisco vSmart Controller in the form of a percentage of the maximum capacity (1000 devices). This value indicates whether the Cisco vSmart Controller can support the number of devices forecast for the tenant that you are onboarding.	Memory utilized	This value represents memory consumption as a percentage.	CPU utilized	This value represents CPU usage as a percentage.
Tenant hosting capacity	Each Cisco vSmart Controller can serve a maximum of 24 tenants. Tenant hosting capacity represents the number of tenants to which the Cisco vSmart Controller is assigned in the form of a percentage. This value indicates whether you can assign another tenant to this controller.								
Used device capacity	Each Cisco vSmart Controller can support a maximum of 1000 tenant WAN edge devices. Used device capacity represents the number of tenant WAN edge devices connected to the Cisco vSmart Controller in the form of a percentage of the maximum capacity (1000 devices). This value indicates whether the Cisco vSmart Controller can support the number of devices forecast for the tenant that you are onboarding.								
Memory utilized	This value represents memory consumption as a percentage.								
CPU utilized	This value represents CPU usage as a percentage.								

Field	Description								
Destination vSmart	<p>a. Click the Destination vSmart drop-down list.</p> <p>Cisco vManage lists the hostnames of the available Cisco vSmart Controllers that are not assigned to the tenant. For each Cisco vSmart Controller, Cisco vManage shows whether the controller is reachable and reports the following utilization details:</p> <table border="1"> <tbody> <tr> <td>Tenant hosting capacity</td> <td>Each Cisco vSmart Controller can serve a maximum of 24 tenants. Tenant hosting capacity represents the number of tenants to which the Cisco vSmart Controller is assigned in the form of a percentage. This value indicates whether you can assign another tenant to this controller.</td> </tr> <tr> <td>Used device capacity</td> <td>Each Cisco vSmart Controller can support a maximum of 1000 tenant WAN edge devices. Used device capacity represents the number of tenant WAN edge devices connected to the Cisco vSmart Controller in the form of a percentage of the maximum capacity (1000 devices). This value indicates whether the Cisco vSmart Controller can support the number of devices forecast for the tenant that you are onboarding.</td> </tr> <tr> <td>Memory utilized</td> <td>This value represents memory consumption as a percentage.</td> </tr> <tr> <td>CPU utilized</td> <td>This value represents CPU usage as a percentage.</td> </tr> </tbody> </table> <p>b. Check the check box adjacent to the hostname of the Cisco vSmart Controller you want to assign to the tenant.</p> <p>If you select a Cisco vSmart Controller that does not have the required capacity to serve the tenant devices, the update operation fails.</p>	Tenant hosting capacity	Each Cisco vSmart Controller can serve a maximum of 24 tenants. Tenant hosting capacity represents the number of tenants to which the Cisco vSmart Controller is assigned in the form of a percentage. This value indicates whether you can assign another tenant to this controller.	Used device capacity	Each Cisco vSmart Controller can support a maximum of 1000 tenant WAN edge devices. Used device capacity represents the number of tenant WAN edge devices connected to the Cisco vSmart Controller in the form of a percentage of the maximum capacity (1000 devices). This value indicates whether the Cisco vSmart Controller can support the number of devices forecast for the tenant that you are onboarding.	Memory utilized	This value represents memory consumption as a percentage.	CPU utilized	This value represents CPU usage as a percentage.
Tenant hosting capacity	Each Cisco vSmart Controller can serve a maximum of 24 tenants. Tenant hosting capacity represents the number of tenants to which the Cisco vSmart Controller is assigned in the form of a percentage. This value indicates whether you can assign another tenant to this controller.								
Used device capacity	Each Cisco vSmart Controller can support a maximum of 1000 tenant WAN edge devices. Used device capacity represents the number of tenant WAN edge devices connected to the Cisco vSmart Controller in the form of a percentage of the maximum capacity (1000 devices). This value indicates whether the Cisco vSmart Controller can support the number of devices forecast for the tenant that you are onboarding.								
Memory utilized	This value represents memory consumption as a percentage.								
CPU utilized	This value represents CPU usage as a percentage.								

6. Click **Update**.

7. To change the other Cisco vSmart Controller that is assigned to the tenant, repeat Step 3 to Step 6.

Cisco vManage initiates the Tenant vSmart Update task to assign the selected Cisco vSmart Controller to the tenant, migrating the tenant details from the Cisco vSmart Controller that was previously assigned. When the task is successfully completed, you can view the tenant information, including the Cisco vSmart Controllers assigned to the tenant, on the **Administration > Tenant Management** page.



CHAPTER 17

Cisco SD-WAN Multitenancy (Cisco SD-WAN Releases 20.4.x and 20.5.x)

Table 115: Feature History

Feature Name	Release Information	Feature Description
Cisco SD-WAN Multitenancy	Cisco SD-WAN Release 20.4.1 Cisco vManage Release 20.4.1	With Cisco SD-WAN multitenancy, a service provider can manage multiple customers, called tenants, from Cisco vManage. In a multitenant Cisco SD-WAN deployment, tenants share Cisco vManage instances, Cisco vBond Orchestrators and Cisco vSmart Controllers. Tenant data is logically isolated on these shared resources.

- [Overview of Cisco SD-WAN Multitenancy, on page 421](#)
- [User Roles in Multitenant Environment, on page 424](#)
- [Hardware Supported and Specifications, on page 425](#)
- [Initial Setup for Multitenancy, on page 426](#)
- [Manage Tenants, on page 430](#)
- [Cisco vManage Dashboard for Multitenancy, on page 433](#)
- [Manage Tenant WAN Edge Devices, on page 437](#)
- [Tenant-Specific Policies on Cisco vSmart Controllers, on page 438](#)
- [Manage Tenant Data, on page 439](#)
- [View OMP Statistics per Tenant on a Cisco vSmart Controller, on page 442](#)
- [View Tenants Associated with a Cisco vSmart Controller, on page 443](#)
- [Migrate Single-Tenant Cisco SD-WAN Overlay to Multitenant Cisco SD-WAN Deployment, on page 443](#)

Overview of Cisco SD-WAN Multitenancy

With Cisco SD-WAN multitenancy, a service provider can manage multiple customers, called tenants, from Cisco vManage. The tenants share Cisco vManage instances, Cisco vBond Orchestrators, and Cisco vSmart Controllers. The domain name of the service provider has subdomains for each tenant. For example, the `multitenancy.com` service provider can manage the tenants `Customer1 (Customer1.multitenancy.com)` and `Customer2 (Customer2.multitenancy.com)`.

Following are the key features of Cisco SD-WAN multitenancy:

- Full enterprise multitenancy: Cisco SD-WAN supports multitenancy and offers enterprises the flexibility of segregated roles such as service provider and tenants. Service providers can use multitenancy to provide Cisco SD-WAN service offerings to their customers.
- Multi-tenant Cisco vManage:
 - Cisco vManage is deployed and configured by the service provider. The provider enables multitenancy and creates a Cisco vManage cluster to serve tenants. Only the provider can access a Cisco vManage instance through the SSH terminal.



Note To connect to a device through SSH, use the IP address of the `vmanage_system` interface; this IP address is assigned by Cisco vManage. Do not use a user-configured system IP address to connect to a device through SSH.

You can find the IP address of the `vmanage_system` interface from the output of the **show interface description** command. Alternatively, you can launch the device SSH terminal from Cisco vManage and find the `vmanage_system` IP address from the first line of the log-in prompt.

- Cisco vManage offers service providers an overall view of the SD-WAN multi-tenant deployment and allows a provider to manage the shared Cisco vBond Orchestrator and Cisco vSmart Controller devices. Cisco vManage also allows service providers to monitor and manage the deployments of each tenant.
- Cisco vManage allows tenants to monitor and manage their deployment. Through Cisco vManage, tenants can deploy and configure WAN edge devices. Tenants can also configure custom policies on assigned Cisco vSmart Controllers.
- Multi-tenant Cisco vBond Orchestrators:
 - Cisco vBond Orchestrators are deployed and configured by the service provider. Only the provider can access a Cisco vBond Orchestrator through the SSH terminal.



Note To connect to a device through SSH, use the IP address of the `vmanage_system` interface; this IP address is assigned by Cisco vManage. Do not use a user-configured system IP address to connect to a device through SSH.

You can find the IP address of the `vmanage_system` interface from the output of the **show interface description** command. Alternatively, you can launch the device SSH terminal from Cisco vManage and find the `vmanage_system` IP address from the first line of the log-in prompt.

- Cisco vBond Orchestrators serve WAN edge devices of multiple tenants as the devices are added to the overlay network.
- Multi-tenant Cisco vSmart Controllers:

- Cisco vSmart Controllers are deployed by the service provider. Only the provider can create and attach device and feature templates to Cisco vSmart Controllers, and can access a Cisco vSmart Controller through the SSH terminal.



Note To connect to a device through SSH, use the IP address of the `vmanage_system` interface; this IP address is assigned by Cisco vManage. Do not use a user-configured system IP address to connect to a device through SSH.

You can find the IP address of the `vmanage_system` interface from the output of the **show interface description** command. Alternatively, you can launch the device SSH terminal from Cisco vManage and find the `vmanage_system` IP address from the first line of the log-in prompt.

- When a tenant is created, Cisco vManage assigns two Cisco vSmart Controllers for the tenant. The Cisco vSmart Controllers form an active-active cluster.

Each tenant is assigned only two Cisco vSmart Controllers. Before a tenant is created, two Cisco vSmart Controllers must be available to serve the tenant.

- Each pair of Cisco vSmart Controllers can serve a maximum of 24 tenants.
- Tenants can configure custom policies on the Cisco vSmart Controllers assigned to them. Cisco vManage notifies the Cisco vSmart Controllers to pull the policy templates. Cisco vSmart Controllers pull the templates and deploy the policy configuration for the specific tenant.
- Only the provider can view events, audit logs, and OMP alarms for a Cisco vSmart Controller on Cisco vManage.

- WAN Edge Devices:

- A tenant or the provider acting on behalf of a tenant can add WAN edge devices to the tenant network, configure the devices, and remove the devices from the tenant network, or access the device through the SSH terminal.



Note To connect to a device through SSH, use the IP address of the `vmanage_system` interface; this IP address is assigned by Cisco vManage. Do not use a user-configured system IP address to connect to a device through SSH.

You can find the IP address of the `vmanage_system` interface from the output of the **show interface description** command. Alternatively, you can launch the device SSH terminal from Cisco vManage and find the `vmanage_system` IP address from the first line of the log-in prompt.

- A provider can manage the WAN edge devices only from [provider-as-tenant](#) view. In the [provider](#) view, Cisco vManage does not present any WAN edge device information.
- Cisco vManage reports WAN edge device events, logs, and alarms only in the [Tenant Role](#) and the [provider-as-tenant](#) views.

- Overlapping VPN numbers: A particular VPN or a set of common VPNs is assigned to a specific tenant, with their own configurations and monitoring dashboard environment. These VPN numbers can overlap where they are used by other tenants.
- On-prem and cloud deployment models: Cisco SD-WAN controllers can be deployed in an organization data center on servers running the VMware vSphere ESXi or the Kernel-based Virtual Machine (KVM) hypervisor. Cisco SD-WAN controllers can also be deployed in the cloud on Amazon Web Services (AWS) servers.

User Roles in Multitenant Environment

A multi-tenant environment includes the service provider and tenant roles. Each role has distinct privileges, views, and functions.

Provider Role

The provider role entitles system-wide administrative privileges. A user with the provider role has the default username **admin**. The provider user can access Cisco vManage using the domain name of the service provider or by using the Cisco vManage IP address. When using a domain name, the domain name has the format `https://multitenancy.com`.

The **admin** user is part of the user group **netadmin**. Users in this group are permitted to perform all operations on the controllers and the Cisco SD-WAN devices of the tenants. You can add additional users to the **netadmin** group.

You cannot modify the privileges of the **netadmin** group. On Cisco vManage, you can view the privileges of the user group from the **Administration > Manage Users > User Groups** page.



Note When you create a new provider user in Cisco vManage, including a **netadmin** user, by default, the user is not allowed SSH access to the Cisco vManage VM. To enable SSH access, configure SSH authentication using a AAA template and push the template to Cisco vManage. For more information on enabling SSH authentication, see [SSH Authentication using vManage on Cisco vEdge Devices](#).

For more information about configuring users and user groups, see [Configure User Access and Authentication](#).

Cisco vManage offers two views to a provider:

• Provider View

When a provider user logs in to multi-tenant Cisco vManage as **admin** or another **netadmin** user, Cisco vManage presents the provider view and displays the provider dashboard.

You can perform the following functions from the provider view:

- Provision and manage Cisco vManage, Cisco vBond Orchestrators and Cisco vSmart Controllers.
- Add, modify, or delete tenants.
- Monitor the overlay network.

• Provider-as-Tenant View

When a provider user selects a specific tenant from the **Select Tenant** drop-down list at the top of the provider dashboard, Cisco vManage presents the provider-as-tenant view and displays the tenant dashboard for the selected tenant. The provider user has the same view of Cisco vManage as a tenant user would when logged in as **tenantadmin**. From this view, the provider can manage the tenant deployment on behalf of the tenant.

In the provider dashboard, a table of tenants presents a status summary for each tenant. A provider user can also launch the provider-as-tenant view by clicking on a tenant name in this table.

Tenant Role

The tenant role entitles tenant administrative privileges. A user with the tenant role has the default username **tenantadmin**. The default password is **Cisco#123@Viptela**. We recommend that you change the default password on first login. For information on changing the default password, see [Hardware and Software Installation](#).

The **tenantadmin** user is part of the user group **tenantadmin**. Users in this group are permitted to perform all operations on the WAN edge devices of the tenants. You can add additional users to the **tenantadmin** group.

You cannot modify the privileges of the **tenantadmin** group. On Cisco vManage, you can view the privileges of the user group from the **Administration > Manage Users > User Groups** page.

For more information about configuring users and user groups, see [Configure User Access and Authentication](#).

A tenant user can log in to Cisco vManage using a dedicated URL and the default username **tenantadmin**. For example, the dedicated URL of a tenant could be `https://Customer1.multitenancy.com` for a provider using the domain name `https://multitenancy.com`. When the user logs in, Cisco vManage presents the tenant view and displays the tenant dashboard.

A tenant user with administrative privileges can perform the following functions:

- Provision and manage tenant routers
- Monitor overlay network of the tenant
- Create custom policies on the assigned Cisco vSmart Controllers
- Upgrade the software on the tenant routers.

Hardware Supported and Specifications

The following platforms support multitenancy.

Table 116: Router Models

Platform	Router Models
Cisco vEdge device	<ul style="list-style-type: none"> • vEdge 100, vEdge 100b, vEdge 100m, vEdge 100wm, vEdge 1000, vEdge 2000, vEdge 5000, vEdge Cloud • ISR1100-6G/ISR1100-4G, ISR1100-4GLTENA, ISR1100-4GLTEGB

The following hypervisors and deployment model are supported for multitenancy.

Table 117: Deployment Model

Specification	Description
Supported hypervisors	VMware, KVM, AWS (cloud-hosted by Cisco)
Cisco vManage Deployment Model	Cluster, 3 vManage instances with each instance running all NMS services.

The supported hardware specifications for the Cisco vBond Orchestrator, Cisco vManage, and the Cisco vSmart Controller are as follows:

Table 118: On-prem Deployment

Server	Cisco vManage	Cisco vBond Orchestrator	Cisco vSmart Controller
Deployment Model	Cluster	N/A	Non-containerized
Number of Instances	3	2	2 per 24 tenants
CPU	32 vCPU	4 vCPU	8 vCPU
DRAM	72 GB	4 GB	16 GB
Hard Disk	1 TB	10 GB	16 GB
NMS Service Distribution	Some services run on all three Cisco vManage instances in the cluster, while some services run on only one of the three instances in the cluster. Therefore, the CPU load may vary among the instances.	N/A	N/A



Note If DPI is enabled, we recommend that the aggregated DPI data across all Cisco vManage instances not exceed 350 GB per day.

Initial Setup for Multitenancy

Prerequisites

- Download and install software versions as recommended in the following table:

Table 119: Software Prerequisites for Cisco SD-WAN Multitenancy

Device	Software Version
Cisco vManage	Cisco vManage Release 20.4.1
Cisco vBond Orchestrator	Cisco SD-WAN Release 20.4.1
Cisco vSmart Controller	Cisco SD-WAN Release 20.4.1
Cisco vEdge Device	Cisco SD-WAN Release 20.4.1

A configuration in which one or more controllers, or WAN edge devices, are running software versions earlier than those mentioned in the table above is not supported.

- Do not migrate an existing single-tenant Cisco vManage instance into multitenant mode, even if you invalidate or delete all devices from the existing Cisco vManage instance. Instead, download and install a new Cisco vManage software image.



Note After you enable Cisco vManage for multitenancy, you cannot migrate it back to single tenant mode.

- Log in to Cisco vManage as the provider **admin** user.
1. Create three Cisco vManage instances and associated configuration templates. See [Deploy Cisco vManage](#).
 - a. While configuring Cisco vManage instances, configure the service provider organization name (`sp-organization-name`) and the organization name (`organization-name`).

Example:

```
sp-organization-name multitenancy
organization-name multitenancy
```
 2. Configure one of the Cisco vManage instances to support multitenancy. See [Enable Multitenancy on Cisco vManage, on page 428](#)
 3. Create a Cisco vManage cluster consisting of three Cisco vManage instances. See [Cluster Management](#).
 - The Cisco vManage cluster must have three Cisco vManage instances. A cluster with more than three instances or fewer than three instances is not a supported configuration for Cisco SD-WAN multitenancy.
 - While creating the Cisco vManage cluster, add the Cisco vManage instance configured to support multitenancy before adding the other two Cisco vManage instances.
 4. Certify all instances of Cisco vManage. See [Generate vManage NMS Certificate](#).
 5. Create and configure Cisco vBond Orchestrator instances. See [Deploy Cisco vBond Orchestrator](#).

While configuring Cisco vBond Orchestrator instances, configure the service provider organization name (`sp-organization-name`) and the organization name (`organization-name`). See [Configure Organization Name in Cisco vBond Orchestrator](#).

```
sp-organization-name multitenancy
organization-name multitenancy
```

6. Create Cisco vSmart Controller instances. See [Deploy the Cisco vSmart Controller](#).
To support 50 tenants and 1000 devices across all tenants, deploy 6 Cisco vSmart Controller instances.
To support 100 tenants and 5000 devices across all tenants, deploy 12 Cisco vSmart Controllers.
 - a. [Add Cisco vSmart Controller](#) to the overlay network.
7. Onboard new tenants. See [Add a New Tenant](#).

Enable Multitenancy on Cisco vManage

1. Launch Cisco vManage using the URL `https://vmanage-ip-address:port`. Log in as the provider **admin** user.
2. From the Cisco vManage menu, choose **Administration > Settings**.
3. In the **Tenancy Mode** bar, click the **Edit**.
4. In the **Tenancy** field, click **Multitenant**.
5. In the **Domain** field, enter the domain name of the service provider (for example, multitenancy.com).
6. Enter a **Cluster Id** (for example, cluster-1 or 123456).
7. Click **Save**.
8. Click **Proceed** to confirm that you want to change the tenancy mode.

Cisco vManage reboots in multitenant mode and when a provider user logs in to Cisco vManage, the provider dashboard appears.

Add Cisco vSmart Controller

1. Log in to Cisco vManage as the provider **admin** user.
2. From the Cisco vManage menu, choose **Configuration > Devices**.
3. Click **Controllers**.
4. Click **Add Controller** and click **vSmart**.
5. In the **Add vSmart** dialog box, do the following:
 - a. In the **vSmart Management IP Address** field, enter the system IP address of the Cisco vSmart Controller.
 - b. Enter the **Username** and **Password** required to access the Cisco vSmart Controller.
 - c. Select the protocol to use for control-plane connections. The default is **DTLS**.
If you select **TLS**, enter the port number to use for TLS connections. The default is 23456.
 - d. Check the **Generate CSR** check box for Cisco vManage to create a Certificate Signing Request.
 - e. Click **Add**.

6. From the Cisco vManage menu, choose **Configuration > Certificates**.
For the newly added Cisco vSmart Controller, the **Operation Status** reads **CSR Generated**.
 - a. For the newly added Cisco vSmart Controller, click **More Options** icon and click **View CSR**.
 - b. Submit the CSR to the Certificate Authority (CA) and obtain a signed certificate.
7. From the Cisco vManage menu, choose **Configuration > Certificates**.
8. Click **Install Certificate**.
9. In the **Install Certificate** dialog box, paste the **Certificate Text** or click **Select a file** upload the certificate file. Click **Install**.

Cisco vManage installs the certificate on the Cisco vSmart Controller. Cisco vManage also sends the serial number of the certificate to other controllers.

On the **Configuration > Certificates** page, the **Operation Status** for the newly added Cisco vSmart Controller reads as **vBond Updated**.

On the **Configuration > Devices** page, the new controller is listed in the Controller table with the controller type, hostname of the controller, IP address, site ID, and other details. The **Mode** is set to **CLI**.
10. Change the mode of the newly added Cisco vSmart Controller to **vManage** by attaching a template to the device.
 - a. From the Cisco vManage menu, choose **Configuration > Templates**.
 - b. Click **Device Templates**.



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled as **Device**

- c. Find the template to be attached to the Cisco vSmart Controller.
- d. Click **...**, and click **Attach Devices**.
- e. In the **Attach Devices** dialog box, move the new controller to the **Selected Device** list and click **Attach**.
- f. Verify the **Config Preview** and click **Configure Devices**.

Cisco vManage pushes the configuration from the template to the new controller.

In the **Configuration > Devices** page, the **Mode** for the Cisco vSmart Controller shows **vManage**. The new Cisco vSmart Controller is ready to be used in your multitenant deployment.

Manage Tenants

Add a New Tenant

Prerequisites

- At least two Cisco vSmart Controllers must be operational and in the `vManage` mode before you can add new tenants.

A Cisco vSmart Controller enters the `vManage` mode when you push a template onto the controller from Cisco vManage. A Cisco vSmart Controller in the `CLI` mode cannot serve multiple tenants.

- Each pair of Cisco vSmart Controllers can serve a maximum of 24 tenants. Ensure that there at least two Cisco vSmart Controllers that can serve a new tenant. If no pair of Cisco vSmart Controllers in the deployment can serve a new tenant, add two Cisco vSmart Controllers and change their mode to `vManage`.
- If you add a second tenant immediately after adding a tenant, Cisco vManage adds them sequentially, and not in parallel.
- Each tenant must have a unique Virtual Account (VA) on **Plug and Play Connect** on [Cisco Software Central](#). The tenant VA should belong to the same Smart Account (SA) as the provider VA.
- For an on-premises deployment, create a Cisco vBond Orchestrator controller profile for the tenant on **Plug and Play Connect**. The fields in the following table are mandatory.

Table 120: Controller Profile Fields

Field	Description/Value
Profile Name	Enter a name for the controller profile.
Multi-Tenancy	From the drop-down list, select Yes .
SP Organization Name	Enter the provider organization name.
Organization Name	Enter the tenant organization name in the format <code><SP Org Name>-<Tenant Org Name></code> . Note The organization name can be up to 64 characters.
Primary Controller	Enter the host details for the primary Cisco vBond Orchestrator.

For a cloud deployment, the Cisco vBond Orchestrator controller profile is created automatically as part of the tenant creation process.

- Log in to Cisco vManage as the provider **admin** user.
- From the Cisco vManage menu, choose **Administration > Tenant Management**.
- Click **Add Tenant**. In the **Add Tenant** dialog box:

- a. Enter a name for the tenant.

For a cloud deployment, the tenant name should be same as the tenant VA name on **Plug and Play Connect**.

- b. Enter a description of the tenant.

The description can be up to 256 characters and can contain only alphanumeric characters.

- c. Enter the name of the organization.

The organization name is case-sensitive. Each tenant or customer must have a unique organization name.

Enter the organization name in the following format:

```
<SP Org Name>-<Tenant Org Name>
```

For example, if the provider organization name is 'multitenancy' and the tenant organization name is 'Customer1', while adding the tenant, enter the organization name as **multitenancy-Customer1**.



Note The organization name can be up to 64 characters.

- d. In the **URL Subdomain Name** field, enter the fully qualified sub-domain name of the tenant.

- The sub-domain name must include the domain name of the service provider. For example, for the `multitenancy.com` service provider, a valid domain name can be `Customer1.multitenancy.com`.



Note The service provider name is shared amongst all tenants. Hence, ensure that the URL naming convention follows the same domain name convention that was provided while enabling multitenancy from the Cisco vManage **Administration > Settings > Tenancy Mode** GUI navigation path.

- For an on-premises deployment, add the fully qualified sub-domain name of the tenant to the DNS. Map the fully qualified sub-domain name to the IP addresses of the three Cisco vManage instances in the Cisco vManage cluster.

When creating fully qualified domain names (FQDN) the following DNS entries are required:

- **Provider Level:** Create DNS A record and map it to the IP addresses of the Cisco vManage instances running in the Cisco vManage cluster. The A record is derived from the domain and Cluster ID that was created in steps 5 and 6 in [Enable Multitenancy on Cisco vManage](#). For example, if domain is `sdwan.cisco.com` and Cluster ID is `vmanage123`, then A record will need to be configured as `vmanage123.sdwan.cisco.com`.



Note If you fail to update DNS entries, it will result in authentication errors when logging in to vManage. Validate DNS is configured correctly by executing `nslookup vmanage123.sdwan.cisco.com`.

- **Tenant Level:** Create DNS CNAME records for each tenant created and map them to the FQDN created at the Provider Level. For example, if domain is **sdwan.cisco.com** and tenant name is **customer1** the CNAME record will need to be configured as **customer1.sdwan.cisco.com**.



Note Cluster ID is not required for CNAME record. Validate DNS is configured correctly by executing **nslookup customer1.sdwan.cisco.com**.

For a cloud deployment, the fully qualified sub-domain name of the tenant is automatically added to the DNS as part of the tenant creation process. After you add a tenant, it could take up to an hour before the fully qualified sub-domain name of the tenant can be resolved by the DNS.

- e. Click **Save**.

The **Create Tenant** screen appears, and the **Status** of the tenant creation reads **In progress**. To view status messages related to the creation of a tenant, click the > button to the left of the status.

Cisco vManage does the following:

- creates the tenant
- assigns two Cisco vSmart Controllers to serve the tenant and pushes a CLI template to these controllers to configure tenant information
- sends the tenant and Cisco vSmart Controller information to Cisco vBond Orchestrators.

What to do next:

After the **Status** column changes to **Success**, you can view the tenant information on the **Administration > Tenant Management** page.

Modify Tenant Information

1. Log in to Cisco vManage as the provider **admin** user.
2. From the Cisco vManage menu, choose **Administration > Tenant Management**.
3. In the left pane, click the name of the tenant.

The tenant information is displayed in a pane on the right.

4. To modify tenant data, do as follows:
 - a. In the right pane, click the pencil icon.
 - b. In the **Edit Tenant** dialog box, modify the tenant name, description, or domain name.
 - c. Click **Save**

Delete a Tenant

Before you delete a tenant, delete all tenant WAN edge devices. See [Delete a WAN Edge Device from a Tenant Network, on page 438](#).

1. Log in to Cisco vManage as the provider **admin** user.
2. From the Cisco vManage menu, choose **Administration > Tenant Management**.
3. In the left pane, click the name of the tenant.
The tenant information is displayed in a pane on the right.
4. To delete the tenant, do as follows:
 - a. In the right pane, click the trash icon.
 - b. In the **Delete Tenant** dialog box, enter the provider **admin** password and click **Save**.

Cisco vManage Dashboard for Multitenancy

After enabling Cisco vManage for multitenancy, you can view the multitenant dashboard when you log in to Cisco vManage. Cisco vManage multitenant dashboard is a portal where the provider or tenant can view and provision the underlying system.

The bar at the top of every Cisco vManage multitenant screen includes icons that allow smooth navigation.

View Tenant Activity, Device, and Network Information

When you log in to a multitenant Cisco vManage as an administrator, the provider dashboard displays the following components. To return to the provider dashboard from other Cisco vManage screens, click **Dashboard** at the left bar.

- Device pane — runs across the top of the multitenant dashboard screen. The device pane displays the number of active Cisco vSmart Controllers, Cisco vBond Orchestrators, and Cisco vManage instances, the connectivity status of devices, and information on certificates that have expired or about to expire.
- Tenants pane — displays the total number of tenants and a summary of the control status, site health, router health, and Cisco vSmart Controller status of all tenants.
- Table of tenants in the overlay network — List of individual tenants, with separate information about the control status, site health, WAN edge device health, and Cisco vSmart Controller status for each tenant.

To display tenant-specific status summary information,

1. Click a tenant name from the tenant list.

A dialog box opens on the right side of the screen that provides additional information about the status of the tenant.

2. To access the tenant dashboard for the selected tenant, click **<Tenant name> Dashboard**.

Cisco vManage presents the provider-as-tenant view and displays the tenant dashboard. To return to the provider view, click **Provider** at the top of page.

3. To close the dialog box, click the tenant name from the tenant list.

View Detailed Information of a Tenant Setup

Cisco vManage displays the tenant dashboard, which provides information about a tenant deployment when

- a provider **admin** user selects a specific tenant from the **Select Tenant** drop-down list in the provider dashboard. This view is called the provider-as-tenant view.
- a **tenantadmin** user logs in to Cisco vManage. This view is called the tenant view.

View All Network Connections in the Tenant Overlay Network

The **Device** pane runs across the top of the tenant dashboard and displays the number of control connections from Cisco vManage to the Cisco vSmart Controllers and routers in the overlay network of a tenant. For each WAN edge device, the Device pane shows

- Total number of control connections between Cisco vSmart Controllers and WAN edge devices
- Number of valid control connections between Cisco vSmart Controllers and WAN edge devices
- Number of invalid control connections between Cisco vSmart Controllers and WAN edge devices

Click a connection number, or the **Up** or **Down** arrow, to display a table with detailed information about each connection. Click the **More Actions** icon at the right of each table row to access the **Device Dashboard** or **Real Time** view from the **Monitor > Devices** screen, or access the **Tools > SSH Terminal** Screen.



Note In Cisco vManage Release 20.6.x and earlier releases, **Real Time** view is part of the **Monitor > Network** screen.

View Information About Device Reboots

The **Reboot** pane displays the total number of reboots in the last 24 hours for all devices in the network. It includes soft and cold reboots and reboots that occurred as a result of power-cycling a device. For each reboot, the following information is listed:

- System IP and hostname of the device that rebooted.
- Time when the device was rebooted.
- Reason for the device reboot

If the same device reboots more than once, each reboot option is reported separately.

Click the **Reboot** pane to open the **Reboot** dialog box. In the **Reboot** dialog box, click **Crashes**. For all device crashes, the following information is listed:

- System IP and hostname of the device on which the crash occurred.
- Crash index of the device
- Core time when the device crashed.

- File name of the device crash log

View Network Connections

The **Control Status** pane displays whether Cisco vSmart Controller and WAN edge devices are connected. Each Cisco vSmart Controller must connect to all other Cisco vSmart Controllers in the network. Each WAN edge device must connect to the maximum number of configured Cisco vSmart Controllers. The **Control Status** pane displays three network connection counts:

- Control Up — total number of devices with the required number of operational control plane connections to a Cisco vSmart Controller
- Partial — total number of devices with some, but not all, operational control plane connection to Cisco vSmart Controllers.
- Control Down — total number of devices with no control plane connection to a Cisco vSmart Controller

To display a table with device details, click a row from the **Control Status** dialog box. Click the **More Actions** icon at the right of each table row to access the **Device Dashboard** or **Real Time** view from the **Monitor > Devices** screen.



Note In Cisco vManage Release 20.6.x and earlier releases, **Real Time** view is part of the **Monitor > Network** screen.

View State of Data Connections for a Site

The **Site Health** pane displays the state of data connections for a site. When a site has multiple WAN edge devices, this pane displays the state for the entire site and not for individual devices. The Site Health pane displays three connectivity states:

- Full WAN Connectivity — total number of sites where all BFD sessions on all routers are in the up state.
- Partial WAN Connectivity — total number of sites where tunnel and all BFD sessions on all routers are in the down state. These sites still have limited data plane connectivity.
- No WAN Connectivity — total number of sites where all BFD sessions on all routers are in the down state. These sites have no data plane connectivity.

To display a table with detailed information about each site, node, or tunnel, click a row from the **Site Health** dialog box. Click the **More Actions** icon at the right of each row in the table to access the **Device Dashboard** or **Real Time** view from the **Monitor > Devices** screen, or access the **Tools > SSH Terminal** screen.



Note In Cisco vManage Release 20.6.x and earlier releases, **Real Time** view is part of the **Monitor > Network** screen.

View Interface Usage for WAN Edge Interfaces

The **Transport Interface Distribution** pane displays interface usage in the last 24 hours for all WAN edge interfaces in VPN 0. It includes all TLOC interfaces. Click the pane to view details of interface usage in the **Transport Interface Distribution** dialog box.

View WAN Edge Device Counts

The **WAN Edge Inventory** pane provides four WAN edge device counts:

- **Total** — total number of authorized serial numbers for WAN edge devices that have been uploaded on Cisco vManage. The serial number is uploaded on the **Configuration > Devices** screen.
- **Authorized** — total number of authorized WAN edge devices in the overlay network. These WAN edge devices are marked as **Valid** in the **Configuration > Certificates > WAN Edge List** screen.
- **Deployed** — total number of deployed WAN edge devices. These are WAN edge devices that are marked as **Valid** and are now operational in the network.
- **Staging** — total number of WAN edge devices you configure at a staging site before they are made a part of the overlay network. These routers do not take part in any routing decisions and do not affect network monitoring through Cisco vManage.

Click the pane to view hostname, system IP, site ID, and other details of each router from the **WAN Edge Inventory** dialog box.

View Aggregated State of WAN Edge Devices

The **WAN Edge Health** pane offers an aggregated view of the state of WAN edge devices by providing a count of the number of devices in each state, therefore describing the health of the hardware nodes. The three WAN edge device states are:

- **Normal** — number of WAN edge devices with memory, hardware, and CPU in normal state. Using less than 70% of total memory or total CPU is classified as, normal.
- **Warning** — number of WAN edge devices with memory, hardware, or CPU in warning state. Using between 70% and 90% of total memory or total CPU is classified as, warning
- **Error** — number of WAN edge devices with memory, hardware, or CPU in error state. Using more than 90% of total memory or total CPU is classified as, error.


Click a number or the WAN edge device state to display a table with the last 12 or 24 hours of memory usage, CPU utilization, and hardware-related alarms, including temperature, power supply, and PIM modules. Click the **More Actions** icon at the right of each row in the table to access the following:


- **Hardware Environment**
- **Real Time** view from the **Monitor > Devices** screen
Cisco vManage Release 20.6.x and earlier: **Real Time** view from the **Monitor > Network** screen
- **Tools > SSH Terminal** screen.

View WAN Edge Device Loss, Latency, Jitter

The **Transport Health** pane displays the aggregated average loss, latency, and jitters for all links and all combinations of colors (for example, all LTE-to-LTE links, all LTE-to-3G links).

From the **Type** drop-down arrow, choose loss, latency, or jitter.

Click the  icon to select a time period for which to display the transport health.


Click the  icon to open the **Transport Health** dialog box. This dialog box displays a more detailed view. To display information in a tabular format, click **Details**. You can choose to change the displayed health type and time period.


View DPI Flow Information of WAN Edge Devices

The **Top Applications** pane displays DPI flow information for traffic transiting routers in the overlay network.



Note DPI flow information is shown only for the last 24 hours. To view DPI flow information for a time before the last 24 hours, you must check the information for the specific device.

Click the  icon to select a time period for which to display data. From the **VPN** drop-down list, select a VPN to display DPI information for all flows in that VPN.


Click the  icon to open the **Top Applications** dialog box. This dialog box displays a more detailed view of the same information. You can change the VPN and time period.


View Tunnels Data

The **Application-Aware Routing** pane allows you to choose the following tunnel criteria from the **Type** drop-down arrow:

- Loss
- Latency
- Jitter

Based on the tunnel criteria, the pane displays the 10 worst tunnels. For example, if you choose loss, the pane shows 10 tunnels with the greatest average loss over the last 24 hours.

Click the  icon against a row to display a graphical representation of the data. Select a time period for which to display data or click **Custom** to display a drop-down arrow for specifying a custom time period.

Click the  icon to open the **Application-Aware Routing** dialog box. This dialog box displays the 25 worst tunnels based on criteria you choose from the **Type** drop-down arrow, the criteria being loss, latency, and jitter.

Manage Tenant WAN Edge Devices

Add a WAN Edge Device to a Tenant Network

1. Log in to Cisco vManage.

If you're a provider user, log in as the **admin**. In the provider dashboard, choose a tenant from the drop-down list to enter the provider-as-tenant view.

If you're a tenant user, log in as the **tenantadmin**.

2. Upload the device serial number file to Cisco vManage.
3. Validate the device and send details to controllers.
4. Create a configuration template for the device and attach the device to the template.

While configuring the device, configure the service provider organization name and the tenant organization name as in the following example:

```
sp-organization-name multitenancy
organization-name multitenancy-Customer1
```



Note Enter the `organization-name` in the format `<SP Org Name>-<Tenant Org Name>`.

5. Bootstrap the device using bootstrap configuration generated through Cisco vManage or manually create the initial configuration on the device.
6. If you are using Enterprise Certificates to authenticate the device, download the CSR from Cisco vManage and get the CSR signed by the Enterprise CA. Install the certificate on Cisco vManage.

Delete a WAN Edge Device from a Tenant Network

1. Log in to Cisco vManage.

If you're a provider user, log in as the **admin**. In the provider dashboard, choose a tenant from the drop-down list to enter the provider-as-tenant view.

If you're a tenant user, log in as the **tenantadmin**.

2. Detach the device from any configuration templates.
3. [Delete a WAN Edge Router](#).

Tenant-Specific Policies on Cisco vSmart Controllers

A provider **admin** user (from the Cisco vManage provider-as-tenant view) or a **tenantadmin** user (from the Cisco vManage tenant view) can create and deploy tenant-specific policies on the Cisco vSmart Controllers serving the tenant. The user can configure a CLI policy or create the policy using the UI policy configuration wizard.

When you activate or deactivate a policy,

1. Cisco vManage identifies the Cisco vSmart Controllers serving the tenant.
2. Cisco vManage notifies the Cisco vSmart Controllers to pull the policy configuration.
3. Cisco vSmart Controllers pull and deploy the policy configuration.
4. Cisco vManage reports the status of the policy pull by the Cisco vSmart Controllers.

Manage Tenant Data

Back Up Tenant Data

The tenant data backup solution of Cisco vManage multitenancy provides the following functionalities:

- [Create, Extract, and List Configuration Data Backup File](#).
- Back up configuration database of a specific tenant with an option to restore it later. See [Restore and Delete Tenant Data Backup File](#).
- Delete back up files of a tenant stored in Cisco vManage. For deleting tenant data backup files, see [Restore and Delete Tenant Data Backup File](#).

The following factors are applicable when using data backup solution:

- The tenant data backup solution operations can be performed by a tenant administrator over tenant view and as a provider. To know how to access tenant dashboard through different views, see [User Roles in Multitenant Environment, on page 424](#).
- A tenant is allowed to perform the following backup operations at a particular time and must complete an operation before starting a new operation:
 - Back up a single configuration database
 - Download the backup file.
 - Restore or import backup files
 - Delete backup files.
 - List backup files

- A tenant backup file format is as follows:

```
Bkup_tenantId_MMDDYY-HHMMSS_taskIdWithoutDash.tar.gz
```

- The tenant data backup operation is a readonly operation on the configuration database. However, to ensure data consistency and prevent data loss, do not perform any major changes on the network.
- When a backup or restore operation for a specific tenant is in-progress, other tenants are allowed to perform the backup and restore operations smoothly.
- A tenant is not allowed to perform other backup operations when the restore operation of the tenant database is in-progress. So, a tenant can perform a single backup operation and when this operation is in-progress, all new backup operation requests are rejected.

The remaining tenants can continue with their backup operations.

- A tenant must use the same Cisco vManage version for backup generation and restore operation.
- A tenant can store a maximum of three backup files in Cisco vManage and can download to store them outside Cisco vManage repository. If the tenant already has three backup files, a subsequent backup operation results in the earliest backup file being deleted and a new backup file being generated.

- Ensure that the following parameter values match in both the backup file and the setup where tenant has requested for a restore operation:
 - Tenant Id
 - Organization Name
 - SP Organization Name
- The tenant data backup solution creates a task in the tenant view of Cisco vManage. Therefore, the tenant can monitor the progress of the operation from the task view of the tenant dashboard.
- A provider cannot back up provider data using this solution. Therefore, the provider can back up all tenants information at once by backing up all tenants configuration database using CLI.

Create, Extract, and List Configuration Data Backup File

1. Log in to Cisco vManage.

If you're a provider user, log in as the **admin**. In the provider dashboard, choose a tenant from the drop-down list to enter the provider-as-tenant view.

If you're a tenant user, log in as the **tenantadmin**.

2. In the address bar, modify the URL path with `dataservice` for the REST API connection.

Example: `https://<tenant_URL>/dataservice`

3. Create a configuration backup file by using the following API:

`https://<tenant_URL>/dataservice/tenantbackup/export.`

4. If the configuration backup file has been created successfully, Cisco vManage task view indicates that the backup file has been generated. You can view the process identifier of the created process or task.

Example:

```
{
  "processId": "72d69805-b987-436f-9b7a-afef2f3f9061",
  "status": "in-progress"
}
```

5. Verify the task status using the obtained process identifier.

Example:

`https://<tenant_URL>/dataservice/device/action/status/72d69805-b987-436f-9b7a-afef2f3f9061`

The verification generates the details of the task in the JSON file format.

6. After the task is completed, extract or download the backed-up file available under the **data** section of the JSON task file.

Example: To extract or download the backup file, use the following API:

`https://<tenant_URL>/dataservice/tenantbackup/download/1570057020772/backup_1570057020772_100919-181838.tar.gz`

7. List backup files stored in Cisco vManage using the following API.

Example: `https://<tenant_URL>/dataservice/tenantbackup/list`

Restore and Delete Tenant Data Backup File

Before you begin:

To run the restore and delete tenant data backup files API, you can download and install Postman tool or any other alternative tool for testing http applications and services. In this document, the procedure to restore and delete tenant data backup files has been explained using the Postman tool. Postman is a software tool used as an API development environment. You can download the tool from the Postman website.

1. Open Google Chrome, or another browser, and enable developer mode on it.
2. Log in to Cisco vManage.

If you're a provider user, log in as the **admin**. In the provider dashboard, choose a tenant from the drop-down list to enter the provider-as-tenant view.

If you're a tenant user, log in as the **tenantadmin**.

3. To get header information of the restore API, do as follows:
 - a. On the right side of the screen, click the **Network** tab to get the network capture view.
 - b. In the network capture view, click the **Name** column to sort the listed items.
 - c. Search and click **index.html**.
 - d. Click the **Headers** tab and expand **Request Headers**.
 - e. Choose all text under **Request Headers** and copy it to the clipboard.
4. Import backup files through the Postman UI:
 - a. Open the Postman UI.
 - b. To disable SSL certificate verification, click **Postman > Preferences > General > Request**. Turn off **SSL Certificate Verification**.
 - c. In the Postman UI, create a new tab.
 - d. Click **Headers** and then click **Bulk Edit**.
 - e. Paste the text that was copied in step 3 from the **Request Headers** block into an editable form.
 - f. From the **GET** method drop-down list, choose **POST**.
 - g. In the **Paste request URL** field, paste the dedicated URL of the tenant and include `dataservice/tenantbackup/import`.
Example: `https://Customer1.multitenancy.com/dataservice/tenantbackup/import`
 - h. Click the **Body** tab and select **form-data**.
 - i. Under **KEY** column, enter `bakup.tar.gz`
 - j. Under **VALUE** column, click **Select Files** and select a backup file to be imported.
 - k. To run the API, click **Send**.

In the **Response** section of the Postman UI, you can view the JSON information that indicates the file that was restored.

5. Monitor the restoration of backup files in either of the following ways:
 - a. Use Cisco vManage task view that indicates if backup file has been imported successfully. You can view the process identifier of the created process or task.

Example:

```
{ "processId": "40adb6c0-eacc-4ad4-ba6c-2c2da2e96d1d",
  "status": "Import Successfully Submitted for tenant 1579026919487"
}
```
 - b. Use the following URL to get the status,


```
https://<tenant_URL>/dataservice/device/action/status/<processId>
```

Example:

```
https://Customer1.multitenancy.com/dataservice/device/action/status/40adb6c0-eacc-4ad4-ba6c-2c2da2e96d1d
```

6. Delete tenant data backup file through Postman UI.
 - a. In the Postman UI, create a new tab.
 - b. Click **Headers** and then click **Bulk Edit**.
 - c. Paste the text that was copied in step 3 from the **Request Headers** block into an editable form.
 - d. From the **GET** method drop-down list, choose **DELETE**.
 - e. In the **Paste request URL** field, paste the dedicated URL of the tenant and include `dataservice/tenantbackup/delete?fileName='filename'`. The filename can either be name of the backup file or all.

Example:

```
https://Customer1.multitenancy.com/dataservice/tenantbackup/delete?fileName=bkup_1579026919487_012820-180712_c09230904dfc40edb0d1e50b68b03002.tar.gz
```

Example:

```
https://Customer1.multitenancy.com/dataservice/tenantbackup/delete?fileName=all
```
 - f. To run the API, click **Send**.

In the **Response** section of the Postman UI, you can view the JSON information that indicates the files that were deleted.

Example:

```
{
  "Deleted": [
    "bkup_1579026919487_012820-180712_c09230904dfc40edb0d1e50b68b03002.tar.gz"
  ]
}
```

View OMP Statistics per Tenant on a Cisco vSmart Controller

1. Log in to Cisco vManage as the provider **admin** user.
2. From the Cisco vManage menu, choose **Monitor > Devices**.
Cisco vManage Release 20.6.x and earlier: From the Cisco vManage menu, choose **Monitor > Network**.
3. In the table of devices, click on the hostname of a Cisco vSmart Controller.

4. In the left pane, click **Real Time**.
5. In the **Device Options** field, enter **OMP** and select the OMP statistics you wish to view.
6. In the **Select Filters** dialog box, click **Show Filters**.
7. Enter the **Tenant Name** and click **Search**.

Cisco vManage displays the selected OMP statistics for the particular tenant.

View Tenants Associated with a Cisco vSmart Controller

1. Log in to Cisco vManage as the provider **admin** user.
2. Click a **vSmart** connection number to display a table with detailed information about each connection.
Cisco vManage displays a table that provides a summary of the Cisco vSmart Controllers and their connections.
3. For a Cisco vSmart Controller, click **...** and click **Tenant List**.
Cisco vManage displays a summary of tenants associated with the Cisco vSmart Controller.

Migrate Single-Tenant Cisco SD-WAN Overlay to Multitenant Cisco SD-WAN Deployment

Table 121: Feature History

Feature Name	Release Information	Description
Migrate Single-Tenant Cisco SD-WAN Overlay to Multitenant Cisco SD-WAN Deployment	Cisco SD-WAN Release 20.5.1 Cisco vManage Release 20.5.1	This feature enables you to migrate a single-tenant Cisco SD-WAN overlay to a multitenant deployment using a sequence of Cisco vManage API calls.

Before You Begin

- Before you begin the migration,
 - Ensure that the edge devices in the single-tenant deployment can reach the Cisco vBond Orchestrator in the multitenant deployment
 - Ensure that the template, routing, and policy configuration on the edge devices is synchronized with the current configuration on Cisco vManage
 - Ensure that the Certificate Authority (CA) on both single-tenant and multitenant Cisco vManages are same.
 - Configure a maintenance window for the single-tenant overlay before performing this procedure. See [Configure or Cancel vManage Server Maintenance Window](#).

- Minimum software requirements for the single-tenant overlay to be migrated:

Device	Software Version
Cisco vManage	Cisco vManage Release 20.5.1
Cisco vBond Orchestrator	Cisco SD-WAN Release 20.5.1
Cisco vSmart Controller	Cisco SD-WAN Release 20.5.1
Cisco vEdge Device	Cisco SD-WAN Release 20.4.1

- Minimum software requirements for the multitenant deployment to which the single-tenant overlay must be migrated:

Device	Software Version
Cisco vManage	Cisco vManage Release 20.5.1
Cisco vBond Orchestrator	Cisco SD-WAN Release 20.5.1
Cisco vSmart Controller	Cisco SD-WAN Release 20.5.1
Cisco vEdge Device	Cisco SD-WAN Release 20.5.1

- We recommend that you use a custom script or a third-party application like Postman to execute the API calls.

Migration Procedure

1. Export the single-tenant deployment and configuration data from a Cisco vManage instance controlling the overlay.

Method	POST
URL	<code>https://ST-vManage-IP-address</code>
Endpoint	<code>/dataservice/tenantmigration/export</code>
Authorization	Admin user credentials.

Body	<p>Required</p> <p>Format: Raw JSON</p> <pre>{ "desc": <tenant_description>, "name": <tenant_name>, "subdomain": <tenant_name>.<domain>, "orgName": <tenant_orname > }</pre> <p>Field Description:</p> <ul style="list-style-type: none"> • desc: A description of the tenant. The description can be up to 256 characters and can contain only alphanumeric characters. • name: Unique name for the tenant in the multitenant deployment. • subdomain: Fully qualified sub-domain name of the tenant. The sub-domain name must include the domain name of the service provider. For example, if <code>multitenancy.com</code> is the domain name of service provider, and the tenant name is <code>Customer1</code>, the tenant sub-domain name would be <code>Customer1.multitenancy.com</code>. • orgName: Name of the tenant organization. The organization name is case-sensitive.
Response	<p>Format: JSON</p> <pre>{ "processId": <vManage_process_ID>, }</pre>

While exporting the data, Cisco vManage attempts to detach any CLI templates from the edge devices in preparation for the migration to the multitenant deployment. If prompted by Cisco vManage, detach CLI templates from the edge devices and execute the export API call again.

2. Check the status of the data export task in Cisco vManage. When the task succeeds, download the data using the URL
<https://ST-vManage-IP-address/dataservice/tenantmigration/download/default.tar.gz>
3. On a multitenant Cisco vManage instance, import the data exported from the single-tenant overlay.

Method	POST
URL	https://MT-vManage-IP-address
Endpoint	/dataservice/tenantmigration/import
Authorization	Provider Admin user credentials.
Body	<p>Required</p> <p>Format: form-data</p> <p>Key Type: File</p> <p>Value: <code>default.tar.gz</code></p>

Response	<pre>Format: JSON { "processId": <vManage_process_ID>, "migrationTokenURL": <token_URL>, }</pre>
----------	------------------------------------------------------------------------------------------------------------------

When the task succeeds, on the multitenant Cisco vManage, you can view the devices, templates, and policies imported from the single-tenant overlay.

- Obtain the migration token using the token URL obtained in response to the API call in **Step 3**.

Method	GET
URL	<code>https://MT-vManage-IP-address</code>
Endpoint	migrationTokenURL obtained in Step 3 .
Authorization	Provider Admin user credentials.
Response	The migration token as a large blob of encoded text.

- On the single-tenant Cisco vManage instance, initiate the migration of the overlay to the multitenant deployment.

Method	POST
URL	<code>https://ST-vManage-IP-address</code>
Endpoint	<code>dataservice/tenantmigration/networkMigration</code>
Authorization	Admin user credentials.
Body	<p>Required</p> <p>Format: Raw text</p> <p>Content: Migration token obtained in Step 4.</p>
Response	<pre>Format: JSON { "processId": <vManage_process_ID>, }</pre>

In Cisco vManage, check the status of the migration task. As part of the migration task, the address of the multitenant vBond Orchestrator, and the service provider and tenant organization names are pushed to the WAN edge devices of the single-tenant overlay. If the task succeeds, WAN edge devices form control connections to controllers in the multitenant deployment; the WAN edge devices are no longer connected to the controllers of the single-tenant overlay.

Attach any CLI templates detached from the edge devices (in Step 1) after migration to the multitenant deployment. Before you attach the templates, update the Cisco vBond Orchestrator IP address and the Organization name to match the configuration of the multitenant deployment.

**Note**

In the single-tenant deployment, if Cisco vManage-signed certificates are installed on cloud-based WAN edge devices, the certificates are cleared when the devices are migrated to the multitenant deployment. You must re-certify the devices on the multitenant Cisco vManage. If enterprise certificates are installed on the cloud-based WAN edge devices, the certificates are not affected by the migration. For more information, see [Enterprise Certificates](#).



CHAPTER 18

Appendix: vManage How-Tos

- [How to Load a Custom vManage Application Server Logo, on page 449](#)

How to Load a Custom vManage Application Server Logo

To change the Cisco vManage web application server logo and load a new custom logo, use the **request nms application-server update-logo** command.

The logo image is located in the upper left corner of all Cisco vManage web application server screens. You can load two files, a larger version, which is displayed on wider browser screens, and a smaller version, which is displayed when the screen size narrows. Both files must be PNG files located on the local device, and both must be 1 MB or smaller in size. For best resolution, it is recommended that the image for the large logo be 180 x 33 pixels, and for the small logo 30 x 33 pixels.

