CISCO VALIDATED DESIGN

Cisco SD-WAN Design Guide

October 2018

Table of Contents

SD-WAN definition

The enterprise landscape is continuously evolving. There is a greater demand for mobile and Internet-of-Things (IoT) device traffic, SaaS applications, and cloud adoption. In addition, security needs are increasing and applications are requiring prioritization and optimization, and as this complexity grows, there is a push to reduce costs and operating expenses. High availability and scale continue to be important.

Legacy WAN architectures are facing major challenges under this evolving landscape. Legacy WAN architectures typically consist of multiple MPLS transports, or an MPLS paired with an Internet or LTE used in an active/backup fashion, most often with Internet or software-as-a-service (SaaS) traffic being backhauled to a central data center or regional hub for Internet access. Issues with these architectures include insufficient bandwidth along with high bandwidth costs, application downtime, poor SaaS performance, complex operations, complex workflows for cloud connectivity, long deployment times and policy changes, limited application visibility, and difficulty in securing the network.

In recent years, software-defined wide-area networking (SD-WAN) solutions have evolved to address these challenges. SD-WAN is part of a broader technology of software-defined networking (SDN). SDN is a centralized approach to network management which abstracts away the underlying network infrastructure from its applications. This de-coupling of data plane forwarding and control plane allows you to centralize the intelligence of the network and allows for more network automation, operations simplification, and centralized provisioning, monitoring, and troubleshooting. SD-WAN applies these principles of SDN to the WAN.

Why deploy SD-WAN

The Cisco® SD-WAN solution is an enterprise-grade WAN architecture overlay that enables digital and cloud transformation for enterprises. It fully integrates routing, security, centralized policy, and orchestration into large-scale networks. It is multi-tenant, cloud-delivered, highly-automated, secure, scalable, and application-aware with rich analytics. The Cisco SD-WAN technology addresses the problems and challenges of common WAN deployments. Some of the benefits include:

- Centralized management and policy management, as well as operational simplicity, resulting in reduced change control and deployment times.
- A mix of MPLS and low-cost broadband or any combination of transports in an active/active fashion, optimizing capacity and reducing bandwidth costs.
- A transport-independent overlay that extends to the data center, branch, or cloud.
- Deployment flexibility. Due to the separation of the control plane and data plane, controllers can be deployed on premises or in the cloud, or a combination of either. Cisco vEdge router deployment can be physical or virtual and can be deployed anywhere in the network.
- Robust and comprehensive security, which includes strong encryption of data, end-to-end network segmentation, router and controller certificate identity with a zero-trust security model, control plane protection, application firewall, and insertion of Cisco Umbrella™, firewalls, and other network services.
- Seamless connectivity to the public cloud and movement of the WAN edge to the branch.
- Application visibility and recognition and application-aware policies with real-time service-level agreement (SLA) enforcement.
- Dynamic optimization of SaaS applications, resulting in improved application performance for users.
- Rich analytics with visibility into applications and infrastructure, which enables rapid troubleshooting and assists in forecasting and analysis for effective resource planning.

Solution architecture overview

The Cisco SD-WAN solution is comprised of separate orchestration, management, control, and data planes.

- The orchestration plane assists in the automatic onboarding of the SD-WAN routers into the SD-WAN overlay.
- The management plane is responsible for central configuration and monitoring.
- The control plane builds and maintains the network topology and makes decisions on where traffic flows.
- The data plane is responsible for forwarding packets based on decisions from the control plane.

Figure 1. Overview of Cisco SD-WAN solution planes

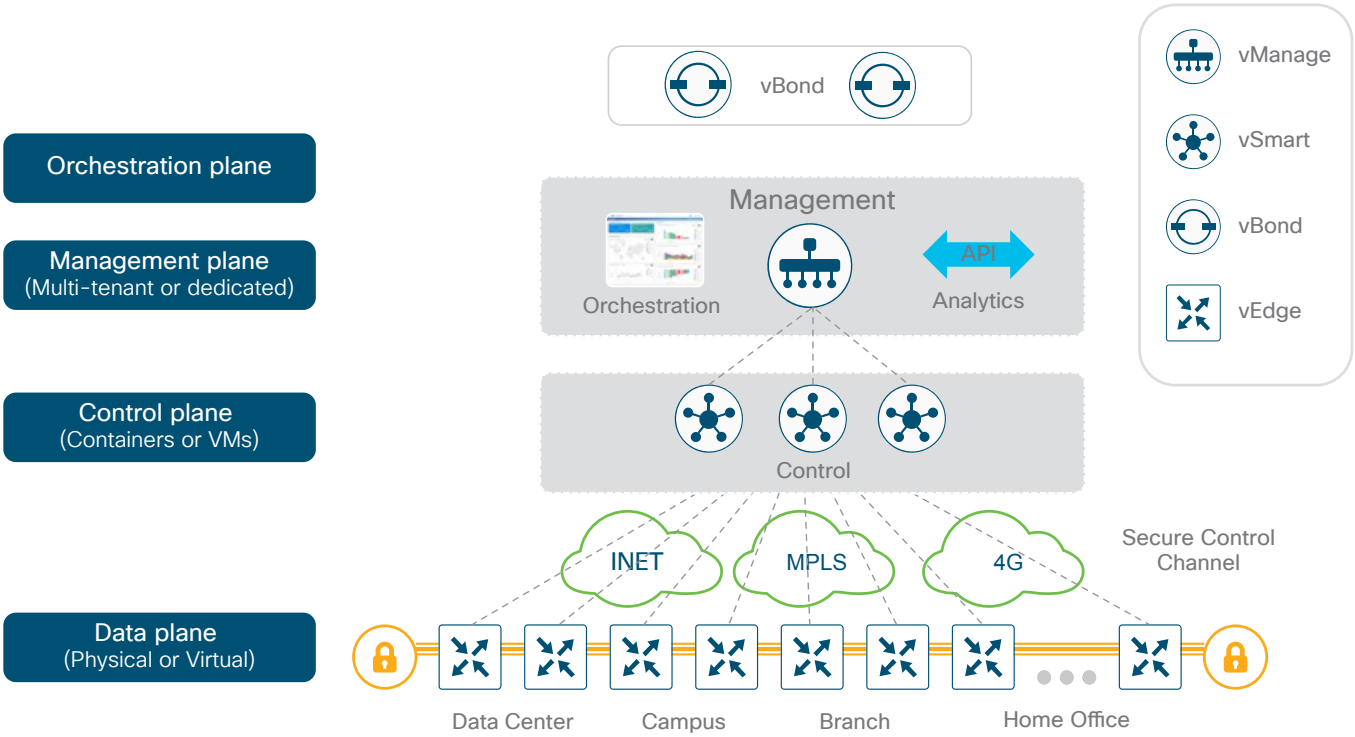

Components

The primary components for the Cisco SD-WAN solution consist of the vManage network management system (management plane), the vSmart controller (control plane), the vBond orchestrator (orchestration plane), and the vEdge router (data plane).

- vManage – This centralized network management system provides a GUI interface to easily monitor, configure, and maintain all Cisco SD-WAN devices and links in the underlay and overlay network.
- vSmart controller – This software-based component is responsible for the centralized control plane of the SD-WAN network. It establishes a secure connection to each vEdge router and distributes routes and policy information via the Overlay Management Protocol (OMP), acting as a route reflector. It also orchestrates the secure data plane connectivity between the vEdge routers by distributing crypto key information, allowing for a very scalable, IKE-less architecture.

- vBond orchestrator – This software-based component performs the initial authentication of vEdge devices and orchestrates vSmart and vEdge connectivity. It also has an important role in enabling the communication of devices that sit behind Network Address Translation (NAT).
- vEdge router – This device, available as either a hardware appliance or software-based router, sits at a physical site or in the cloud and provides secure data plane connectivity among the sites over one or more WAN transports. It is responsible for traffic forwarding, security, encryption, Quality of Service (QoS), routing protocols such as Border Gateway Protocol (BGP) and Open Shortest Path First (OSPF), and more.

The following diagram demonstrates several aspects of the Cisco SD-WAN solution

Figure 2. SD-WAN topology

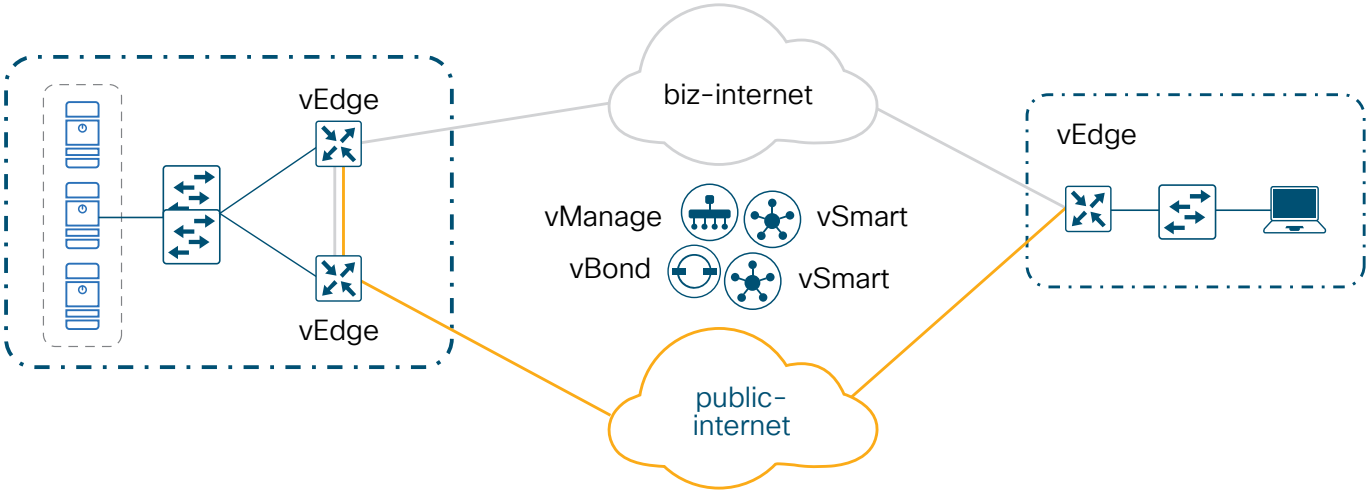

This sample topology depicts two sites and two public Internet transports. The SD-WAN controllers, the two vSmart controllers, and the vBond orchestrator, along with the vManage management GUI that reside on the Internet, are reachable through either transport.

At each site, vEdge routers are used to directly connect to the available transports. Color is used to identify an individual WAN transport; different WAN transports are assigned different colors, such as mpls, private1, biz-internet, metro-ethernet, lte, etc. The topology uses a color called biz-internet for one of the Internet transports and a color called public-internet for the other.

The vEdge routers form a Datagram Transport Layer Security (DTLS) or Transport Layer Security (TLS) control connection to the vSmart controllers and connect to both of the vSmart controllers over each transport. The vEdge routers securely connect to vEdge routers with IPSec tunnels at other sites over each transport. The Bidirectional Forwarding Detection (BFD) protocol is enabled by default and will run over each of these tunnels, detecting loss, latency, jitter, and path failures.

Color

On vEdge routers, the color attribute helps to identify an individual WAN transport tunnel. You cannot use the same color twice on a single vEdge router.

Colors by themselves have significance. The colors metro-ethernet, mpls, and private1, private2, private3, private4, private5, and private6 are considered private colors. They are intended to be used for private networks or in places where you will have no NAT addressing of the transport IP endpoints, as the expectation is that there is no NAT between two endpoints of the same color. When a vEdge router uses a private color, it will attempt to build IPSec tunnels to other vEdge routers using the native, private, underlay IP. The public colors are 3g, biz, internet, blue, bronze, custom1, custom2, custom3, default, gold, green, lte, public-internet, red, and silver. With public colors, vEdge routers will try to build tunnels to the post-NAT IP address (if there is NAT involved).

If you are using a private color and need NAT to communicate to another private color, the carrier setting in the configuration dictates whether you use the private or public IP address. Using this setting, two private colors will establish a session when one or both are using NAT.

Overlay Management Protocol (OMP)

The OMP routing protocol, which is similar to BGP, manages the SD-WAN overlay network. The protocol runs between the vSmart controllers and vEdge routers where control plane information, such as route prefixes, next-hop routes, crypto keys, and policy information, is exchanged over a secure DTLS or TLS connection. The vSmart controller acts a lot like a route reflector; it receives routes from vEdge routers, processes and applies any policy to them, and then advertises the routes to other vEdge routers in the overlay network. If there is no policy defined, the default behavior is a full mesh topology, where each vEdge can connect directly to a vEdge at another site and receive full routing information from each site.

OMP advertises three types of routes:

- OMP routes are prefixes that are learned from the local site, or service side, of a vEdge router. The prefixes are originated as static or connected routes, or from within the OSPF or BGP protocol, and redistributed into OMP so they can be carried across the overlay. OMP routes advertise attributes such as transport location (TLOC) information, which is similar to a BGP next-hop IP address for the route, and other attributes such as origin, originator, preference, site ID, tag, and VPN. An OMP route is only installed in the forwarding table if the TLOC to which it points is active.
- TLOC routes are the logical tunnel termination points on the vEdge routers that connect into a transport network. A TLOC route is uniquely identified and represented by a three-tuple, consisting of system IP address, link color, and encapsulation (Generic Routing Encapsulation [GRE] or IPSec). In addition to system IP address, color, and encapsulation, TLOC routes also carry attributes such as TLOC private and public IP addresses, carrier, preference, site ID, tag, and weight. For a TLOC to be considered in an active state on a particular vEdge, an active BFD session must be associated with that vEdge TLOC.
- Service routes represent services (firewall, IPS, application optimization, etc.) that are connected to the vEdge local-site network and are available for other sites for use with service insertion. In addition, these routes also include VPNs; the VPN labels are sent in this update type to tell the vSmart controllers what VPNs are serviced at a remote site.

See Unicast Overlay Routing Overview for additional information on OMP routing and path selection.

The following diagram shows DTLS/TLS tunnels established between the vEdge routers and the vSmart controller, over which OMP runs. The TLOCs are indicated by the colored circles, T1–T4. IPSec tunnels are established between the TLOCs over each transport. Once the IPSec tunnels are established, BFD is enabled across each of them.

Figure 3. TLOC routes and OMP

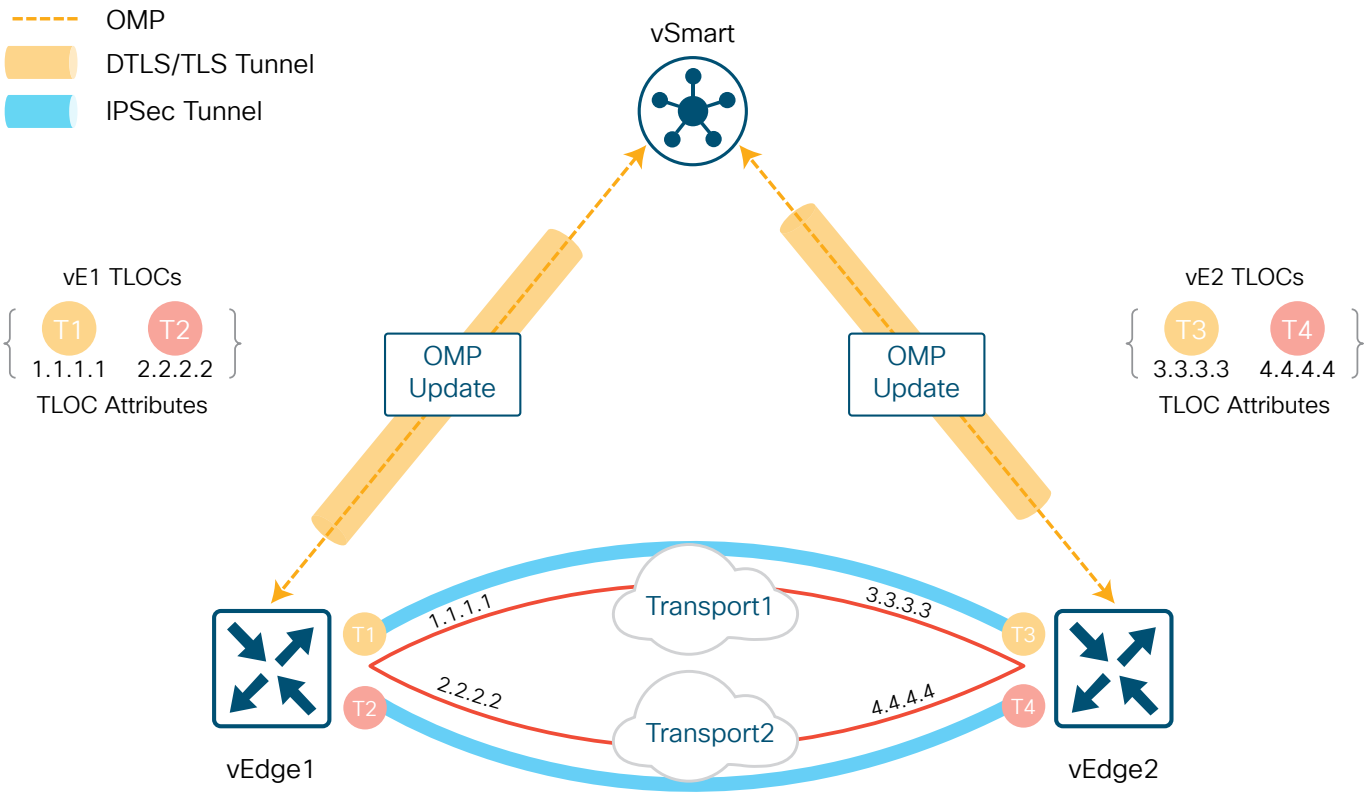

Virtual Private Networks (VPNs)

In the SD-WAN overlay, Virtual Private Networks (VPNs) provide segmentation, much like Virtual Routing and Forwarding instances (VRFs) that many are already familiar with. Each VPN is isolated from one another and each have their own forwarding table. An interface or subinterface is explicitly configured under a single VPN and cannot be part of more than one VPN. Labels are used in OMP route attributes and in the packet encapsulation, which identifies the VPN a packet belongs to.

The VPN number is a four-byte integer with a value from 0 to 65530. There are two VPNs present by default in the vEdge devices and controllers, VPN 0 and VPN 512.

- VPN 0 is the transport VPN. It contains the interfaces that connect to the WAN transports. Secure DTLS/TLS connections to the vSmart or between vSmart and vBond controllers are initiated from this VPN. Static or default routes or a dynamic routing protocol needs to be configured inside this VPN in order to get appropriate next-hop information so the control plane can be established and IPSec tunnels can connect to remote sites.
- VPN 512 is the management VPN. It carries the out-of-band management traffic to and from the Cisco SD-WAN devices. This VPN is not carried across the overlay network.

In addition to the default VPNs that are already defined, one or more service-side VPNs need to be created that will contain interfaces that will connect to the local-site network and carry user data traffic. These VPNs can be enabled for features such as OSPF or BGP, Virtual Router Redundancy Protocol (VRRP), QoS, traffic shaping, or policing. User traffic can be directed over the IPSec tunnels to other sites by redistributing OMP routes received from the vSmart controllers at the site into the service-side VPN routing protocol. In turn, routes from the local site can be advertised to other sites by advertising the service VPN routes into the OMP routing protocol, which will be sent to the vSmart controllers and redistributed to the other vEdge routers in the network.

The following figure demonstrates VPNs on a vEdge router. The interfaces, ge0/2 and ge0/0, are part of the transport VPN; ge0/4 and ge0/5 are part of the service VPN, which is attached to the local network at the site; and the mgmt0 port is part of VPN512. Note that while physical interfaces are displayed in the diagram, the interfaces in the transport and service VPNs could be subinterfaces instead.

Figure 4. VPNs on a vEdge router

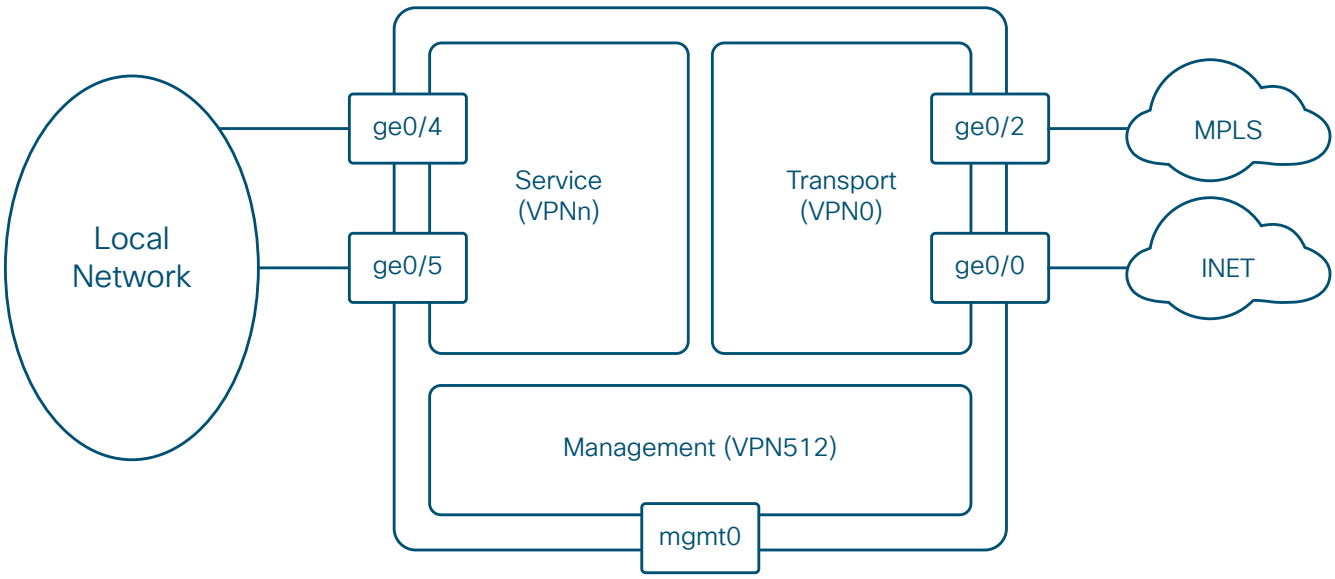

TLOC extension

A very common network setup in a site with two vEdge routers is for each vEdge router to be connected to just one transport as shown in Figure 4. There are links between the vEdge routers, which allow each vEdge router to access the opposite transport through a TLOC-extension interface on the neighboring vEdge router. In the figure below, vEdge1 connects directly to the biz-internet transport and uses the TLOC extension interface on vEdge2 to connect to the public-internet transport. In turn, vEdge2 connects directly to the public-internet transport and uses the TLOC extension interface on vEdge1 to connect to the biz-internet transport. TLOC extensions can be separate physical interfaces or subinterfaces.

Figure 5. TLOC extension

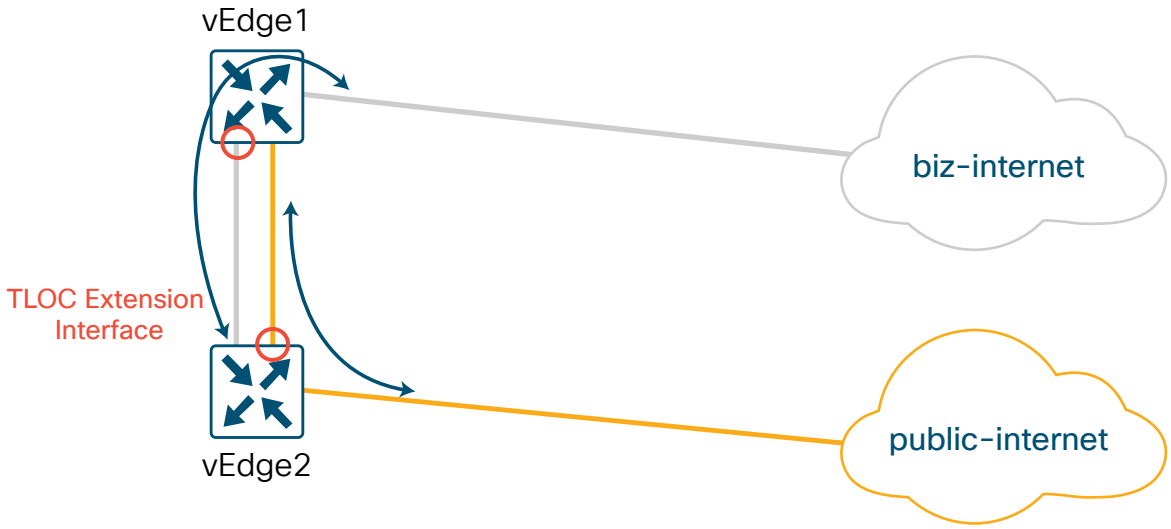

When you configure the TLOC extension interface, you configure it in VPN 0, assign it an IP address, and then specify the WAN interface to which it is bound. In Figure 6, vEdge1's TLOC extension interface is ge0/7 and is bound to the MPLS transport through ge0/2. vEdge2's TLOC extension interface is ge0/7 and is bound to the INET transport through ge0/4.

Figure 6. TLOC extension

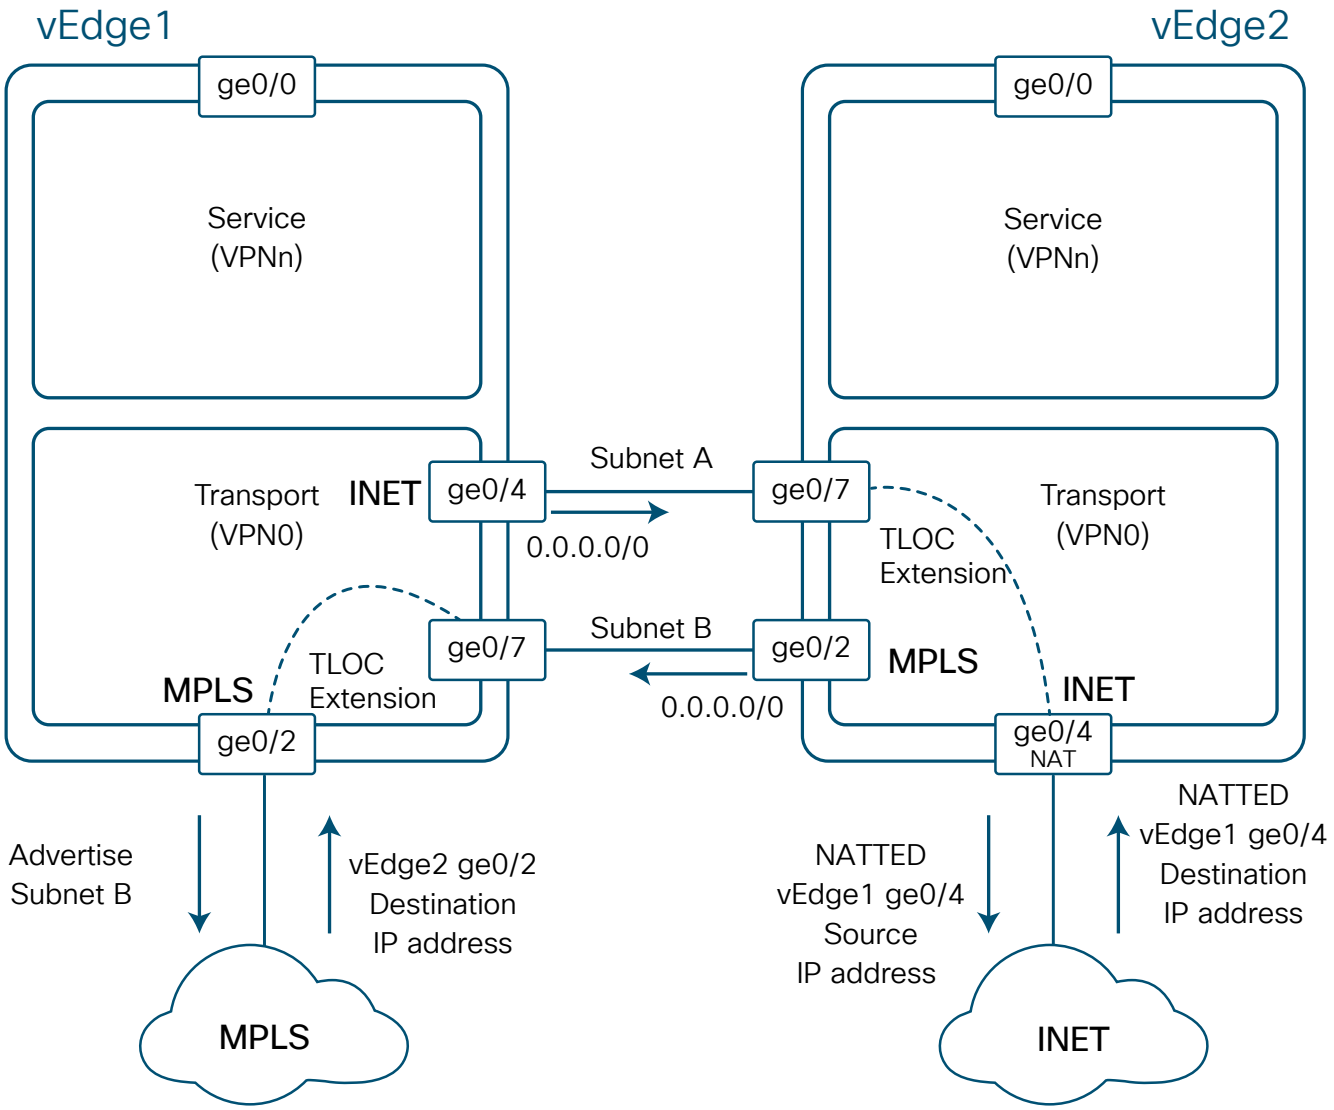

Some routing considerations need to take place in order for controller reachability to occur and for IPSec tunnels to come up with other sites over the TLOC extension interfaces. To reach the INET transport, vEdge1's INET interface should be configured with a default route pointing to vEdge2's ge0/7 IP address. If subnet A is in a private address space, then NAT should be configured on vEdge2's ge0/4 transport interface to ensure traffic can be routed back. To reach the MPLS transport, vEdge2's MPLS interface should be configured with a default route pointing to vEdge1's ge0/7 IP address. To ensure traffic can be routed back to the TLOC extension interface, a routing protocol (typically BGP) can be run in the transport VPN of vEdge1 to advertise subnet B so that the MPLS provider has a route to subnet B through vEdge1.

Bringing the vEdge into the overlay

In order to join the overlay network, a vEdge router needs to establish a secure connection to the vManage so that it can receive a full configuration, and it needs to establish a secure connection with the vSmart controller so that it can participate in the overlay network. The discovery of the vManage and vSmart happens automatically and is accomplished by first establishing a secure connection to the vBond orchestrator.

The following figure shows the sequence of events that occurs when bringing the vEdge router into the overlay.

Figure 7. TLOC extension

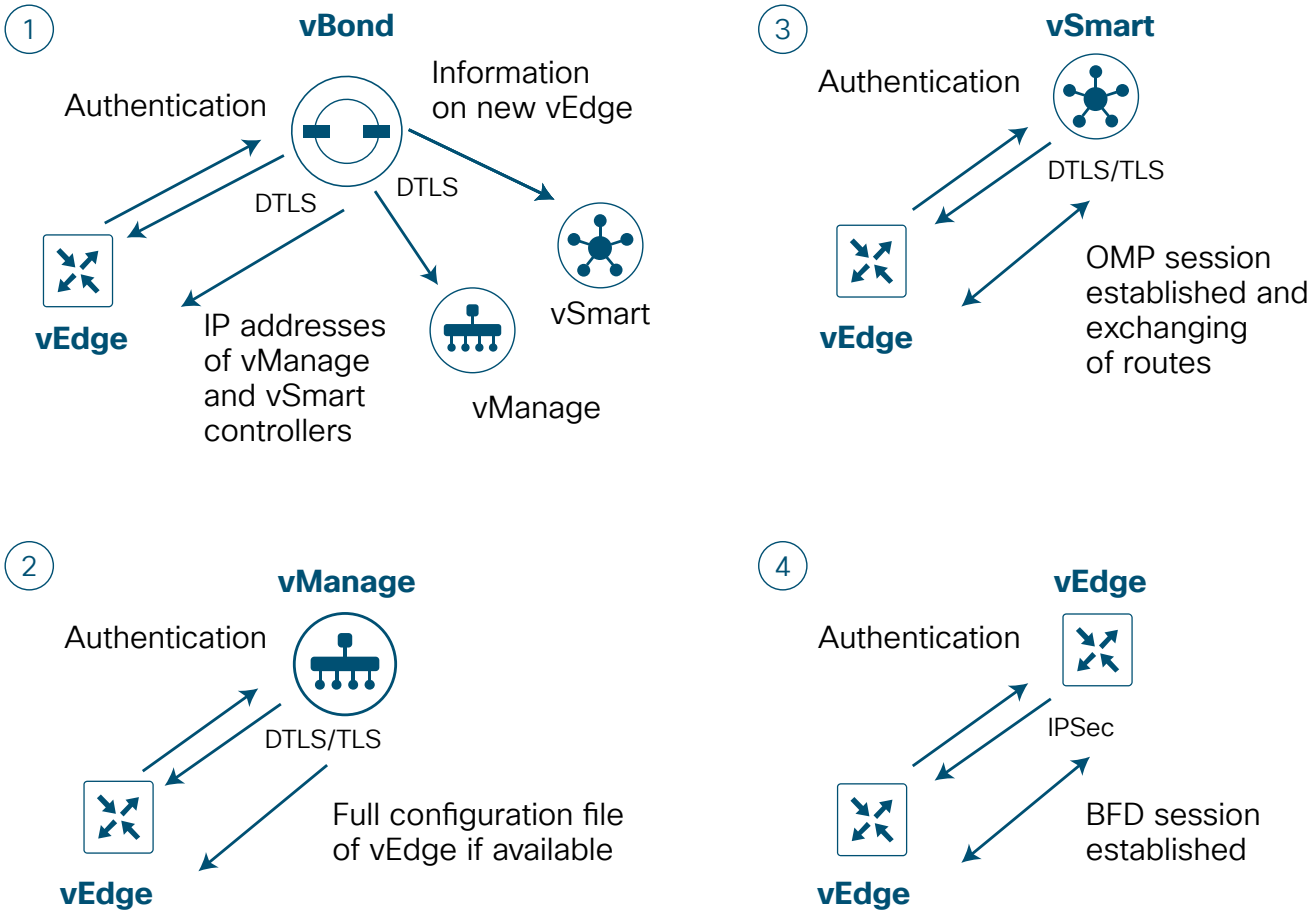

1. Through a minimal bootstrap configuration or through the Zero-Touch Provisioning (ZTP) process, the vEdge router will first attempt to authenticate with the vBond orchestrator through an encrypted DTLS connection. Once authenticated, the vBond orchestrator sends the vEdge router the IP addresses of the vManage Network Management System (NMS) and the vSmart controllers. The vBond orchestrator also informs the vSmart controllers and vManage of the new vEdge router wanting to join the domain.
2. The vEdge router begins establishing secure DTLS or TLS sessions with the vManage and the vSmart controllers and tears down the session with the vBond orchestrator. Once the vEdge router authenticates with the vManage NMS, the vManage with push the full configuration to the vEdge router if available.

3. The vEdge router attempts to establish DTLS/TLS connections to the vSmart controllers over each transport link. When it authenticates to a vSmart controller, it will establish an OMP session and then learn the routes, including prefixes, TLOCs, and service routes, encryption keys, and policies.
4. The vEdge router will attempt to establish an IPSec tunnel to TLOCs over each transport. A TLOC on a private transport color attempts to connect to TLOCs on both public and private colors, and a TLOC on a public color tries to connect to other TLOCs on public colors by default. The restrict keyword on the tunnel will only build tunnels between TLOCs of the same color. BFD will then run over these established connections.

See https://sdwan-docs.cisco.com/Product_Documentation/Getting_Started/Viptela_Overlay_Network_Bringup/01Bringup_Sequence_of_Events for additional details on the vEdge router and controller connection establishment.

See https://sdwan-docs.cisco.com/Product_Documentation/Software_Features/Release_18.1/05Security/01Security_Overview/Data_Plane_Security_Overview for information on data plane security.

Bootstrapping the vEdge router

Two ways to get a vEdge router up and running on the network is by establishing a console to it and configuring a few configuration lines, or by using ZTP, where you can plug the vEdge router into the network and power it on and it will be provisioned automatically.

With the bootstrap configuration method, the idea is to configure the minimum network connectivity and the minimum identifying information along with the vBond orchestrator IP address or hostname. The vEdge router will attempt to connect to the vBond orchestrator and discover the other network controllers from there. In order for you to bring up the vEdge router successfully, there are a few things that need to be present on the vEdge:

- Configure an IP address and gateway address on an interface connected to the network, or alternatively, configure Dynamic Host Configuration Protocol (DHCP) in order to obtain an IP address and gateway address dynamically. The vEdge should be able to reach the vBond through the network.
- Configure the vBond IP address or hostname. If you configure a hostname, the vEdge router needs to be able to reach a DNS server in order to resolve it. You do this by configuring a DNS server address under VPN 0.
- Configure the organization name, system IP address, and site ID. Optionally, configure the host name.

Tech tip

In addition to the above requirements, the vEdge router needs to have a valid certificate installed, but certificates are already installed on hardware-based vEdge routers at the factory. The system clock also needs to reflect accurate time because of the certificate authentication, and can be set manually or through Network Time Protocol (NTP) if need be, but rarely does this need to be addressed when onboarding new devices.

Zero-Touch Provisioning (ZTP) process

ZTP is an automatic provisioning procedure which starts when the vEdge router is powered up for the first time. The vEdge will attempt to connect to a ZTP server with the hostname ztp.viptela.com, where it will get its vBond orchestrator information. Once the vBond orchestrator information is obtained, it can then subsequently make connections to the vManage and vSmart controllers in order to get its full configuration and join the overlay network.

Figure 8. Zero-touch provisioning for a vEdge appliance

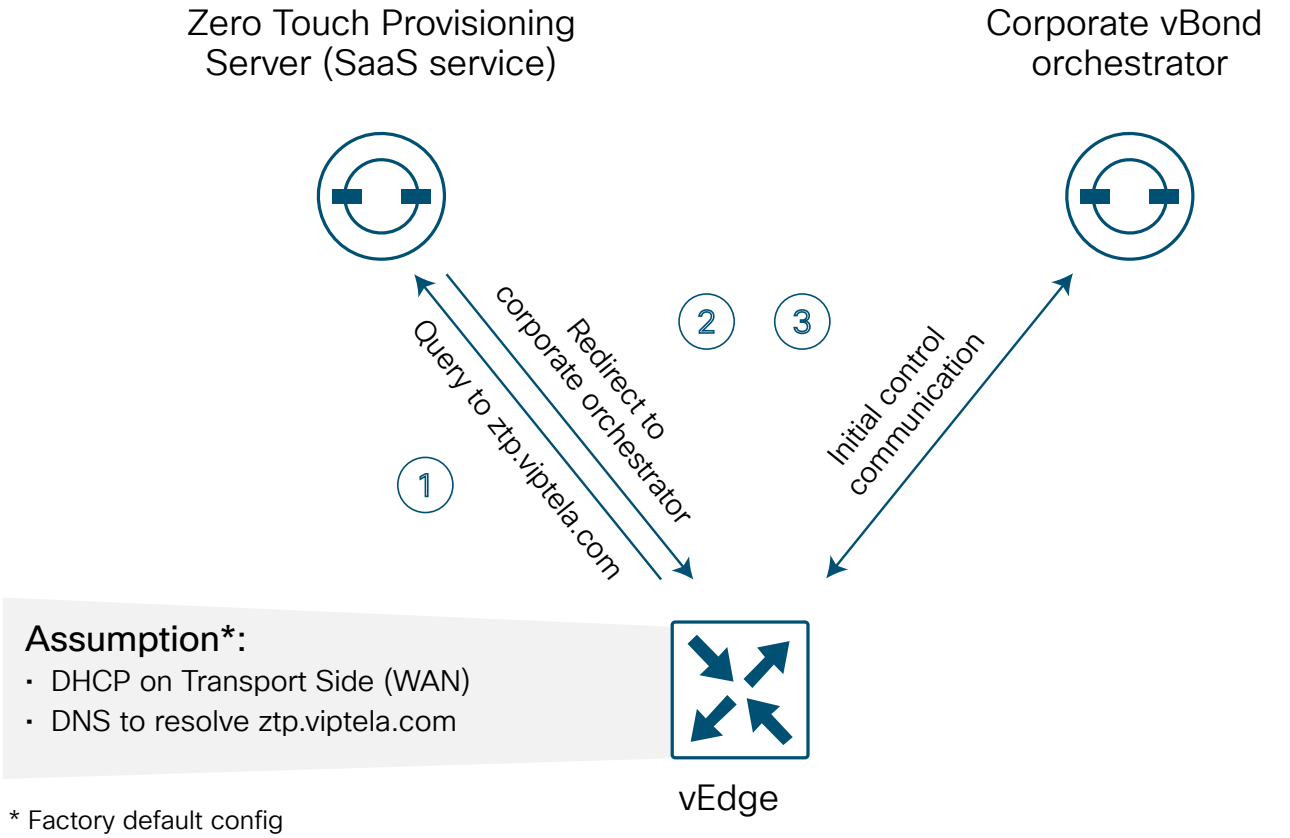

* Factory default config

There are a few requirements for ZTP provisioning:

- With the hardware vEdge appliances, only certain ports are pre-configured by default to be a DHCP client interface and can be used for ZTP. The following table outlines the ports that must be plugged into the network for ZTP to work.

Table 1. vEdge ZTP interfaces

vEdge model	Interface
vEdge 5000	ge0/0 (for network modules in slot 0)
vEdge 2000	ge2/0
vEdge 1000	ge0/0
vEdge 100b	ge0/4
vEdge 100m	ge0/4
vEdge 100wm	ge0/4, cellular0

- The gateway router for the vEdge router in the network should have reachability to public DNS servers and be able to reach ztp.viptela.com.
- In vManage, there must be a device configuration template for the vEdge router attached to the vEdge device. The system IP address and site ID need to be included in this device template in order for the process to work. The ZTP process will not succeed without this.

See https://sdwan-docs.cisco.com/Product_Documentation/Getting_Started/Viptela_Overlay_Network_Bringup/07Deploy_the_vEdge_Routers/08Prepare_vEdge_Routers_for_ZTP for additional information as well as information on ZTP with wireless routers.

Controller connections

The secure sessions between the vEdge routers and the controllers (and between controllers), by default are DTLS, which is User Datagram Protocol (UDP)-based. The default base source port is 12346. The vEdge may use port hopping where the devices try different source ports when trying to establish connections to each other in case the connection attempt on the first port fails. The vEdge will increment the port by 20 and try ports 12366, 12386, 12406, and 12426 before returning back to 12346. Port hopping is configured by default on a vEdge router, but you can disable it globally or on a per-tunnel-interface basis. It is recommended to run port-hopping at the branches, but disable this feature in the controllers, and data center, regional hub, or a place where aggregate traffic exists. Control connections on vManage with multiple vCPUs will have a different base port for each vCPU core.

For vEdge routers that sit behind the same NAT device and share a public IP address, you do not want each vEdge to attempt to connect to the same controller using the same public IP and port number. In this case, you can configure an offset to the base port number of 12346, so the port attempts will be unique among the vEdge routers. A port offset of 1 will cause the vEdge to use the base port of 12347, and then port-hop with ports 12367, 12387, 12407, and 12427. Port offsets need to be explicitly configured, and by default, the port offset is 0.

Alternatively, you can use TLS to connect to the vManage and vSmart controllers, which is TCP-based instead of UDP-based. vBond controller connections will always use DTLS, however. TCP ports will originate on the vEdge from a random port number destined to the base port of 23456, and control connections with multiple vCPUs will have a different base port for each vCPU core, similar to the DTLS case.

IPSec tunnels and BFD from a vEdge router to another vEdge router use UDP with similar ports as defined by DTLS.

Ensure that any firewalls in the network allow communication between vEdge routers and to controllers. Ensure that they are configured to allow return traffic as well. The following table is a summary of the ports used, assuming the controllers are configured to not use port-hopping.

Table 2. DTLS, TLS, and IPSec ports for vEdge and controller connections

Source device	Source port	Destination device	Destination port
vEdge (DTLS)	UDP 12346+n, 12366+n, 12386+n, 12406+n, and 12426+n, where n=0–19 and represents the configured offset	vSmart/vBond controller (no port hopping enabled)	UDP 12346
vEdge (DTLS)	UDP 12346+n, 12366+n, 12386+n, 12406+n, and 12426+n, where n=0–19 and represents the configured offset	vManage	vCPU1 = UDP 12346 vCPU2 = UDP 12446 vCPU3 = UDP 12546 vCPU4 = UDP 12646 vCPU5 = UDP 12746 vCPU6 = UDP 12846 vCPU7 = UDP 12946 vCPU8 = UDP 13046
vEdge (TLS)	TCP random port number > 1024	vSmart	TCP 23456
vEdge (TLS)	TCP random port number > 1024	vManage	vCPU1 = TCP 23456 vCPU2 = TCP 23556 vCPU3 = TCP 23656 vCPU4 = TCP 23756 vCPU5 = TCP 23856 vCPU6 = TCP 23956 vCPU7 = TCP 24056 vCPU8 = TCP 24156
vEdge (IPSec)	UDP 12346+n, 12366+n, 12386+n, 12406+n, and 12426+n, where n=0–19 and represents the configured offset	vEdge	UDP 12346+n, 12366+n, 12386+n, 12406+n, and 12426+n, where n=0–19 and represents the configured offset

Additional NAT considerations

In addition to potential issues of multiple vEdge routers behind a single NAT address for controller connections, there are considerations for vEdge-to-vEdge tunnel traffic. Though several types of NAT are supported with vEdge routers, if full mesh traffic is desired, take care to ensure at least one side of the vEdge tunnel can initiate a connection inbound to a second vEdge if there is a firewall in the path. It is recommended to configure full-cone, or 1-to-1 NAT at the data center or hub site so that, regardless of what NAT type is running at the branch (restricted-cone, port-restricted cone, or symmetric NAT), the branch can connect into the hub site with an

IPSec tunnel at a minimum without issue. Two sites with firewalls running symmetric NAT will have issues forming a tunnel connection, as this NAT translates the source of each side to a random port number, and traffic cannot be initiated from the outside. Symmetric NAT configured at one site requires full-cone NAT or a public IP with no NAT on the other site in order to establish a direct IPSec tunnel between them. Sites which cannot connect directly should be set up to reach each other through the data center.

Configuration templates

Configurations and policies apply to vEdge routers and vSmart controllers which enable traffic to flow between the data center and the branch or between branches. An administrator can enable configurations and policies through the Command-Line Interface (CLI) using console or Secure Shell (SSH) on the vEdge device, or remotely through the vManage GUI.

To configure a vEdge device or controller on the network using the vManage GUI, an administrator applies a device template to a vEdge router or multiple vEdge routers. These templates can be CLI-based or feature-based. While you can create CLI-based templates, we recommend feature-based templates because they are modular, more scalable, and less error-prone. Each device template is made up of several feature templates that describe the interface configurations, tunnel configurations, and local routing behavior.

Tech tip

In order to apply centralized policy to the network, the vSmart controllers must be managed by the vManage. You accomplish this by attaching a CLI or feature-based device template to them.

Templates are extremely flexible, and there are a number of approaches to putting templates together. You can choose to have more variables inside your template, which will result in less feature templates, or you can have less variables but more feature templates. For example, you can choose to enable NAT as a variable or a global value. You can create one interface feature template and choose to enable or disable NAT through a variable, or you can create two different feature templates, one with NAT disabled and one with NAT enabled, and choose the most appropriate feature template to use, depending on the device template. In any case, you should add a detailed description of each feature and device template in detail in the GUI and create very descriptive variable names so that it is very clear what each template and variable is.

When designing configuration templates, it is helpful to think about how operations may interact with the templates on a day-to-day basis. It might be useful to use variables for interface names so that interfaces can be moved for troubleshooting purposes, without having to create new feature templates to accomplish it. It also might be helpful to create variables for states of interfaces and routing protocols for troubleshooting reasons, such as allowing the disabling of an interface or a BGP neighbor by just changing a variable.

Device templates

Device templates are specific to only one vEdge model type, but you may need to create multiple device templates of the same model type due to their location and function in the network. Each device template references a series of feature templates which makes up the entire configuration of the device. A device template configuration cannot be shared between vEdge models, but a feature template can span across several model types and be used by different device templates.

Figure 9 illustrates device template components. The device template is made up of feature templates grouped into the following sections:

- Basic information – This section includes system, logging, AAA, OMP, BFD, security, archive, and NTP feature templates.
- Transport and management VPN – This section includes the templates used to configure VPN 0 and VPN 512, which includes BGP, OSPF, VPN interface, VPN interface cellular, VPN interface GRE, and VPN interface PPP feature templates.
- Service VPN – This section includes the templates used to configure the service VPNs, which contains the BGP, IGMP, Multicast, OSPF, PIM, VPN interface, VPN interface bridge, VPN interface GRE, VPN interface IPSec, VPN interface Natpool, and DHCP server feature templates.
- Additional templates – This section includes banner, Simple Network Management Protocol (SNMP), bridge, localized policy, and cellular feature templates.

Figure 9. Device template

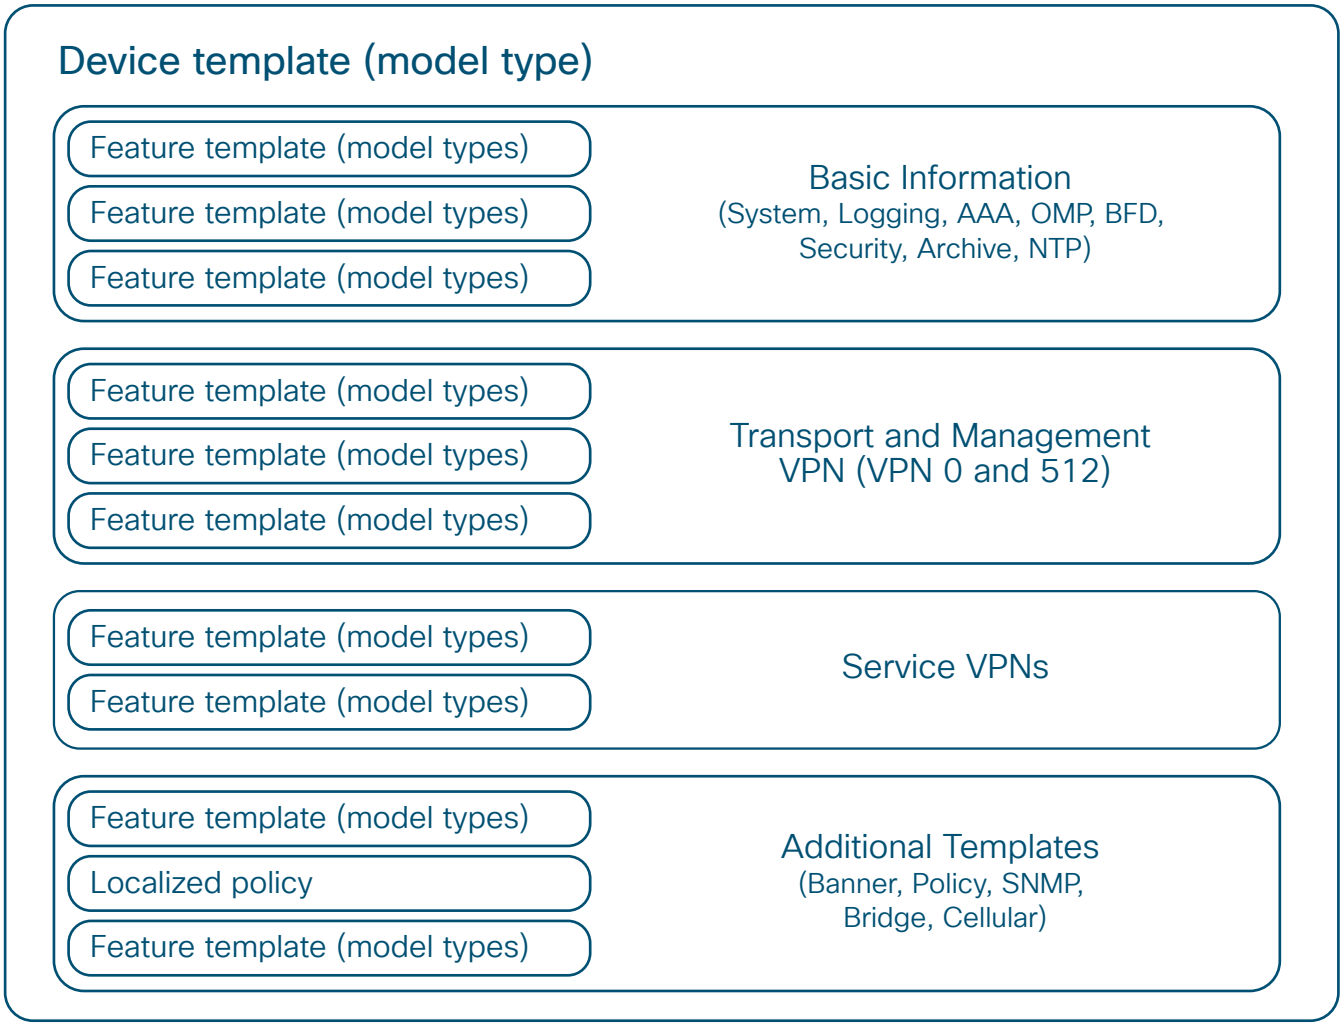

Feature templates

Following is a brief description of some of the different feature templates and a subset of the information each will allow you to configure.

- System – Configure basic system information, such as site ID, system IP, time zone, hostname, device groups, GPS coordinates, port hopping, and port offset.
- Logging – Configure logging to disk and/or to a remote logging server.
- AAA – Specify the authentication method and order and configure Radius, TACACs, or local authentication, including local user groups with different read/write permissions.
- BFD – Specify the BFD app-route multiplier and poll interval and specify the hello and BFD multiplier for each transport.
- OMP – Change the graceful restart timers and advertisement timers and hold timers; change the number of paths advertised; configure an AS overlay number; choose which local protocols will be advertised into OMP; and change the number of equal-cost paths installed in the vEdge router.
- Security – Change the rekey time, anti-replay window, and authentication types for IPSec.
- Archive (optional) – Archive the full running configuration onto a file server within a time period specified.
- NTP (optional) – Configure NTP servers and authentication if required.
- VPN – Change the ECMP hash, add DNS servers, advertise protocols (BGP, static, connected, OSPF external) from the VPN into OMP, and add IPv4 or v6 static routes, service routes, and GRE routes.
- BGP (optional) – Configure the AS number, router ID, distance, maximum paths, neighbors, redistribution of protocols into BGP, hold time, and keepalive timers.
- OSPF (optional) – Configure router ID, distance, areas, OSPF interfaces, reference bandwidth, default information originate, metrics, metric type, and SPF timers.
- VPN Interface configuration – Configure an interface name, the status of the interface, static or dynamic IPv4 and v6 addressing, DHCP helper, NAT, VRRP, shaping, QoS, ingress/egress Access Control List (ACL) for IPv4 and 6, policing, static Address Resolution Protocol (ARP), 802.1x, duplex, MAC address, IP Maximum Transmission Unit (MTU), Transmission Control Protocol Maximum Segment Size (TCP MSS), TLOC extension, and more. In the case of the transport VPN, configure tunnel, transport color, allowed protocols for the interface, encapsulation, preference, weight, and more.
- VPN interface bridge (optional) – Configure layer 3 characteristics of a bridge interface, including IPv4 address, DHCP helper, ACLs, VRRP, MTU, and TCP MSS.
- DHCP server (optional) – Configure DHCP server characteristics, such as address pool, lease time, static leases, domain name, default gateway, DNS servers, and TFTP servers.
- Banner (optional) – Configure the login banner or message-of-the-day banner.
- Policy (optional) – Attach a localized policy.
- SNMP (optional) – Configure SNMP parameters, including SNMP device name and location, SNMP version, views, and communities, and trap groups.
- Bridge (optional) – Define layer 2 characteristics of a bridge, including the VLAN ID, MAC address aging, maximum MAC addresses, and physical interfaces for the bridge.

Routing protocol templates, such as BGP or OSPF, and VPN interface templates are configured under a VPN. DHCP server feature templates are configured under a VPN interface.

Configuring parameters

An administrator uses vManage to configure device and feature templates, specifying variables where needed since templates can apply to multiple vEdge devices that have unique settings.

When configuring values of parameters inside of feature templates, there is often a drop-down box that gives you three different types of values:

- Global – When you specify a global value, you specify the desired value, either by typing the value into a text box, selecting a choice from a radio button, or selecting a value from a drop-down box. Whatever value you select will be applied to all devices the feature template is applied to.
- Device-specific – When you specify a device-specific value, you will create a variable name. The value for this variable will be defined when the device template is applied.
- Default – When you specify a default value, a default value will be applied to all devices the feature template is applied to. If there is a specific value, it will appear in a textbox in grey scale.

In the illustration below, Timezone is shown as a global, device-specific, or default value. A variable name is entered when specifying the device-specific value.

Figure 10. Feature template parameter value types

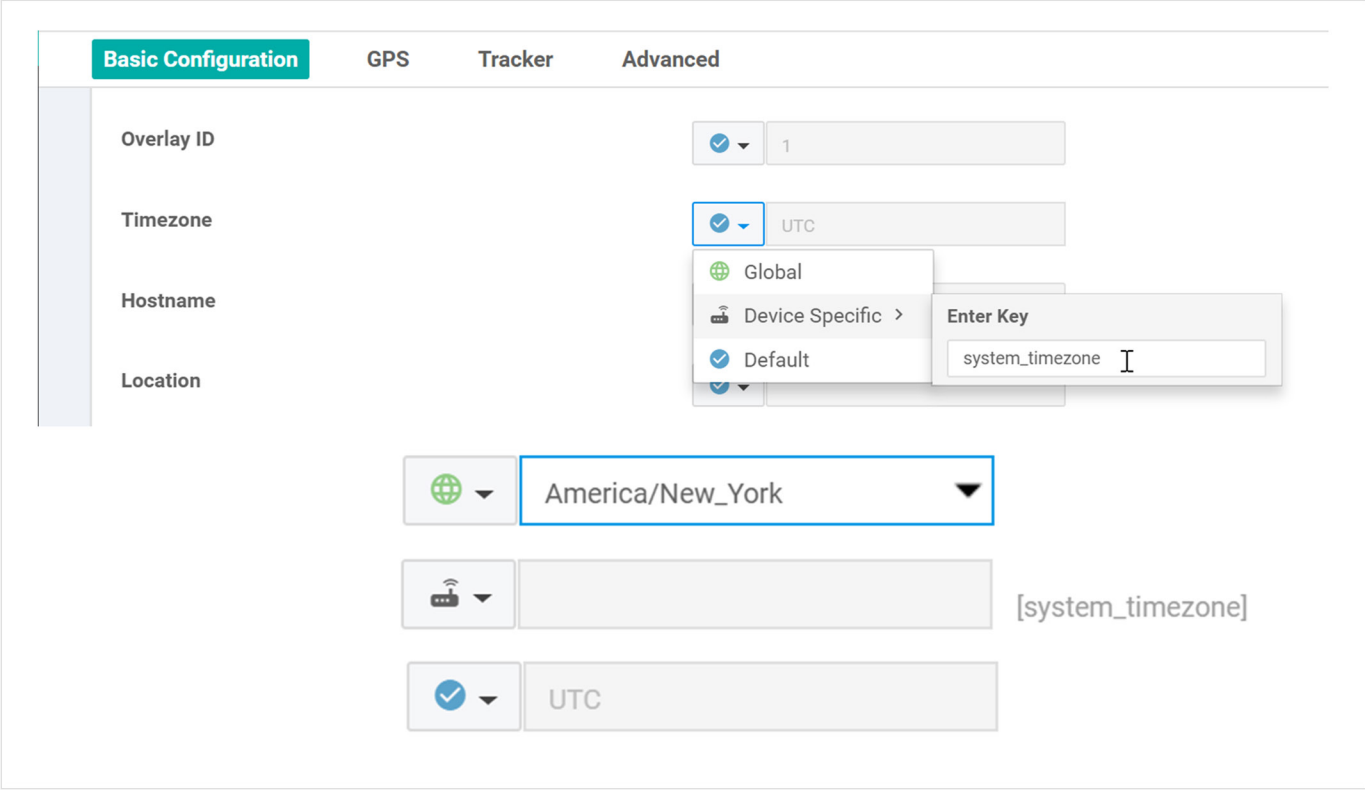

Tech tip

Within each feature template, you can use the same variable name for two different parameter values, but they will be treated like two separate variables. Descriptive and unique variable names are important so that it's clear what values need to be entered when the device template is applied to a device. Variables with the same name in different templates are also different variables and you cannot share them across templates.

Deploying device templates

Once feature templates are configured, the device template configuration is completed by referencing the desired feature template in each configuration category (system, AAA, BFD, VPN, VPN interface, etc.). Once a device template is configured, it can be attached to a specific vEdge device. Once attached, you will be required to fill in the values for any variables in the template for each vEdge the template will apply to before the configuration can be deployed. You can enter values through the vManage GUI directly, or by filling out a .csv file that can be uploaded. The .csv file method allows you to deploy a large number of vEdge routers quickly and more easily. vManage will then modify the configuration of the targeted vEdge devices in the database and then push out the entire configuration to the intended vEdge routers on the network.

When making an update to a feature or device template, the application will happen immediately if there are devices attached to those templates. If the configuration gets pushed out and if there is an error, such as an incorrect value format or a reference to a loopback interface that doesn't exist, the template configuration rolls back to its previous state before the edit.

Policies

Policies are an important part of the Cisco SD-WAN Solution and are used to influence the flow of data traffic among the vEdge routers in the overlay network. Policies apply either to control plane or data plane traffic and are configured either centrally on vSmart controllers (centralized policy) or locally (localized policy) on vEdge routers.

Centralized control policies operate on the routing and TLOC information and allow for customizing routing decisions and determining routing paths through the overlay network. These policies can be used in configuring traffic engineering, path affinity, service insertion, and different types of VPN topologies (full-mesh, hub-and-spoke, regional mesh, etc.). Another centralized control policy is application-aware routing, which selects the optimal path based on real-time path performance characteristics for different traffic types. Localized control policies allow you to affect routing policy at a local site, specifically through OSPF or BGP route maps and prefix lists.

Data policies influence the flow of data traffic through the network based on fields in the IP packet headers and VPN membership. Centralized data policies can be used in configuring application firewalls, service chaining, traffic engineering, quality of service (QoS), and Cflowd. Localized data policies allow you to configure how data traffic is handled at a specific site, such as ACLs, QoS, mirroring, and policing. Some centralized data policy may affect handling on the vEdge itself, as in the case of app-route policies or a QoS classification policy. In these cases, the configuration is still downloaded directly to the vSmart controllers, but any policy information that needs to be conveyed to the vEdge routers is communicated through OMP.

Figure 11. Centralized and localized policies

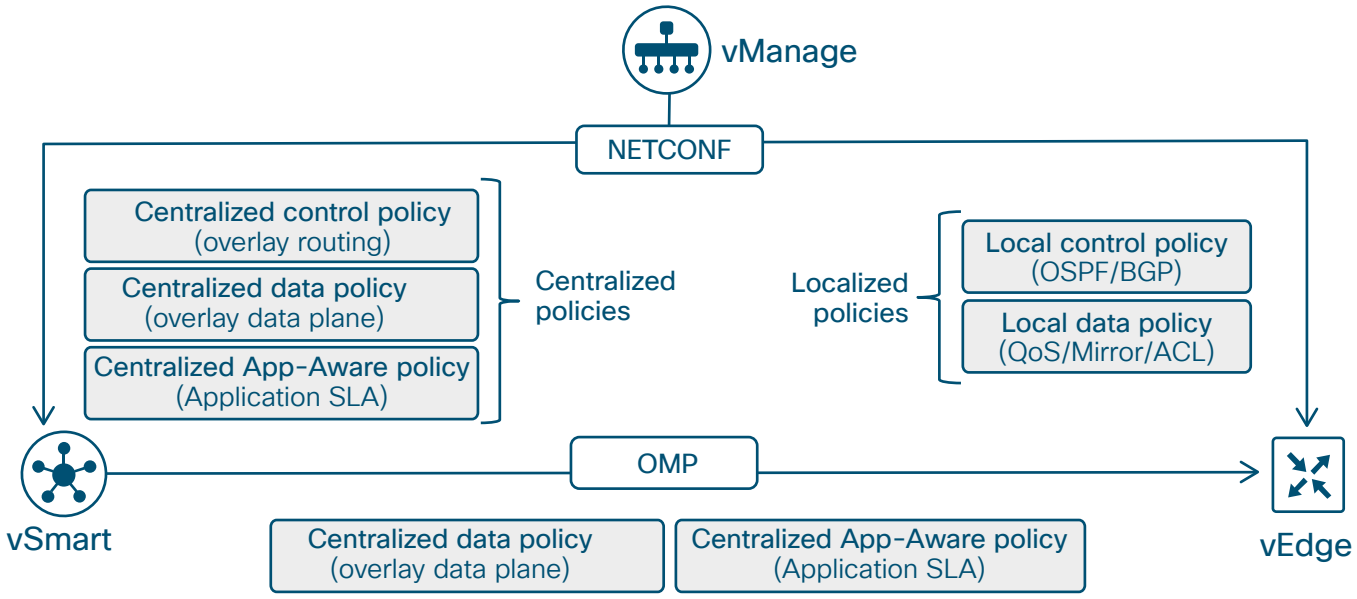

Configuring localized policy

There are three steps for applying localized policy:

1. In the vManage GUI, create the localized policy under **Configuration>Policies** and select the **Localized** Policy tab. Before Release 18.2, the policy is added as a CLI policy. Starting in Release 18.2, a policy configuration wizard was created to assist with policy creation.
2. In the device template, under the **Additional Templates** section next to **Policy**, reference the name of the localized policy.
3. Reference any policy components, like route policies and prefix lists, inside the feature templates.

When you are creating a device template and referencing a feature template that already has a route policy or prefix list or another localized policy component configured in it, you must have a policy name referenced in the device template before you can create or update the device template. If a device is already attached to an existing device template, you must first attach a localized policy to the device template before referencing any localized policy elements within the feature templates that are associated with that device template.

You can only apply one localized policy to a vEdge device. Within this policy, you will create both control and data policies components; prefix-lists, route-policies, as-path lists, community-lists, QoS class-maps, qos-map policies, mirror and policing policies, rewrite-rule policies, and access lists will all be included in this one localized policy.

See https://sdwan-docs.cisco.com/Product_Documentation/Software_Features/Release_17.2/06Policy_Basics/03Localized_Control_Policy/Configuring_Localized_Control_Policy for information on configuring localized control policy and https://sdwan-docs.cisco.com/Product_Documentation/Software_Features/Release_17.2/06Policy_Basics/05Localized_Data_Policy/Configuring_Localized_Data_Policy_for_IPv4 for information on configuring localized data policy.

Configuring centralized policy

When configuring centralized policy in the vManage GUI, there are three main components:

- Lists – Lists are used to group related items so you can reference them as a group. They are used when applying policy or used in matching or actions within the policy definitions. You can create lists for applications, color, data prefixes, policers, prefixes, sites, SLA classes, TLOCs, and VPNs. Data prefixes are used in data policies to define data prefixes, and prefixes are used in control polices to match on route prefixes.
- Policy definition – The policy definitions control the aspects of control and forwarding. Within the policy definition is where you create policy rules, specifying a series of match-action pairs which are examined in sequential order. There are several types of policy definitions: app-route policy, cflowd-template, control-policy, data-policy, and a vpn-membership policy.
- Policy application – The policy is applied to a site list.

There are several different types of policy definitions:

- App-route policy – Allows you to create an application-aware routing policy which tracks path characteristics such as loss, latency, and jitter. Traffic is put into different SLA categories (loss, delay, and jitter), and traffic is directed to different paths depending on the abilities to meet the SLA categories.
- Cflowd template – Allows you to enable cflowd, which sends sampled network data flows to collectors.
- Control policy – Operates on the control plane traffic and influences the routing paths in the network.
- Data policy – Influences the flow of data traffic based on the fields in the IP packet header.
- VPN membership policy – Can restrict participation in VPNs on vEdge routers and the population of their route tables.

Control policy examines the routes and TLOC attributes in the routing information and modifies attributes that match the policy. This policy is unidirectional and can be applied to a site list in an inbound or outbound direction. The direction is from the perspective of the vSmart controller. A policy applied to a site list in the inbound direction means that policy would affect routes coming from the sites on the site list and actions would be applied on the receive side of the vSmart controller. A policy applied to a site list in the outbound direction means the policy would affect routes going to the sites on the site list and actions would be applied to the sending side of the vSmart controller.

Figure 12. Applying centralized policy

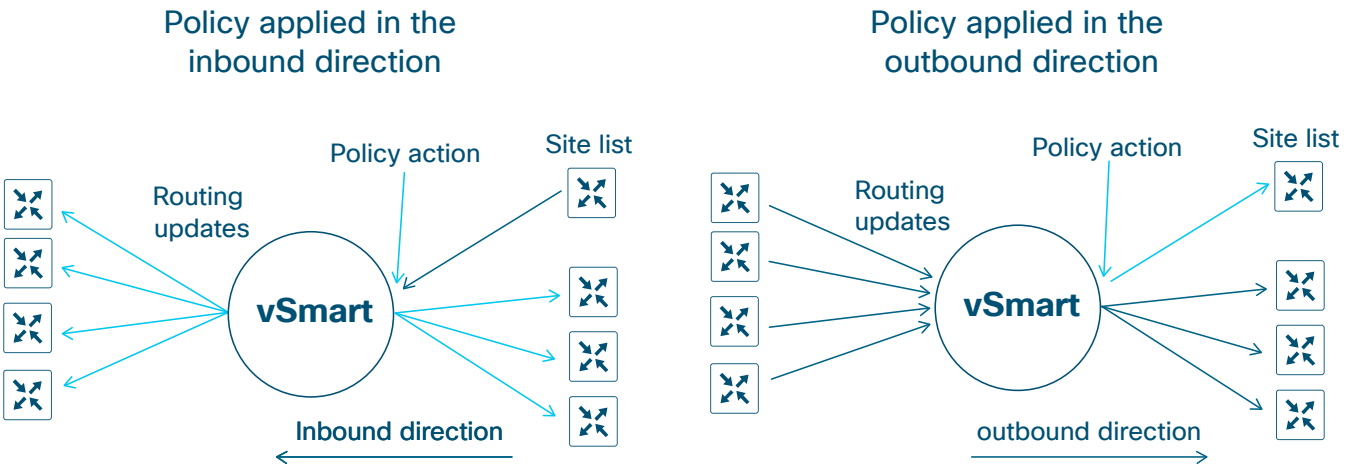

No direction is set with app-route polices – this policy is sent to the vEdge router via OMP and applied to the vEdge as traffic moves in the direction from LAN to WAN. No direction is set with cFlowd and VPN policy as well. Data policy, however, is directional from the perspective of the vEdge. You can apply this either from-service, from-tunnel, or all.

Note that you can create several centralized polices within the vManage GUI, but only one can be activated at a time on any particular vSmart controller. Inside the centralized policy, you will be able to create several different policy definitions that make up the centralized policy, for example, app-route, cflowd, control, data, and vpn-membership policies. Note that with a given site list, you are restricted to one of each type of policy, but you can have a different control policy in each direction (inbound and outbound). When creating site ID lists for the purpose of applying policy definitions, you must not overlap site IDs in different lists.

Order of operations

Following is the order of operations on a packet as it traverses from service VPN to transport VPN on a vEdge router:

1. Local policy/configuration – includes QoS classification, policer, and marking
2. Centralized application-aware routing policy
3. Centralized data policy – includes QoS classification, policer, marking, and path selection
4. Routing/forwarding
5. Scheduling and queueing
6. Local policy shaping and ACL – includes shaping, re-marking, and policer

Figure 13. Policy order of operations on a vEdge router

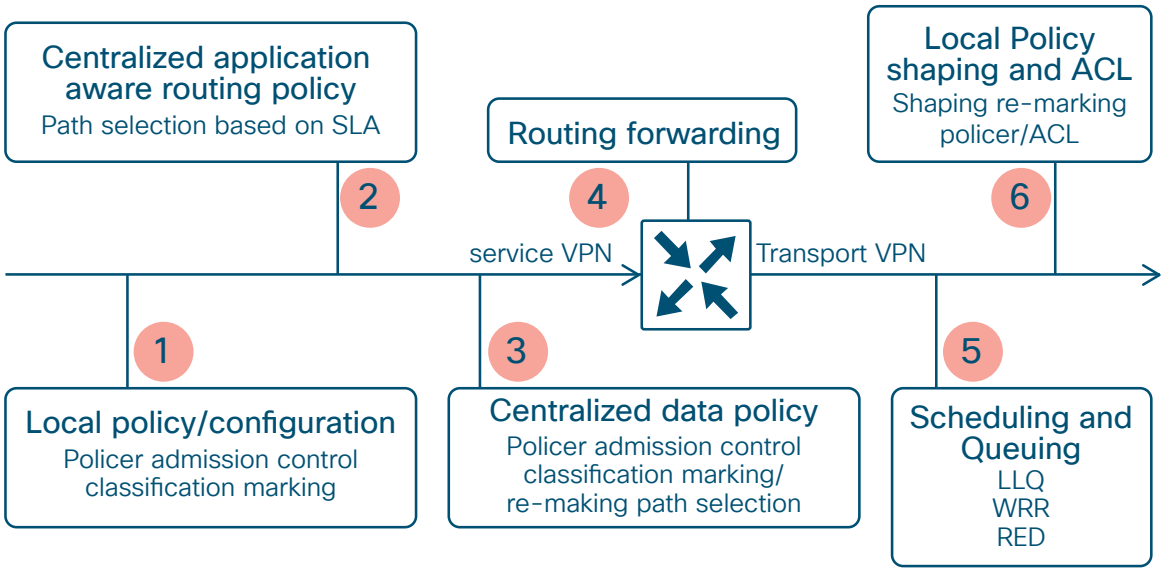

Packet flow through the vEdge router (from service interface to WAN-Transport interface)

From the ordering, it's possible for a centralized data policy to overwrite the actions of a local data policy configuration, and it's also possible for a centralized data policy to influence the path selection that is different than what was chosen as part of the application-aware routing policy. Keep this information in mind as you define the policies for the network.

Traffic symmetry for DPI

Application-aware routing uses Deep Packet Inspection (DPI) for matching on applications within the policy. In order for DPI on a vEdge router to be able to classify most application traffic, it is important that the vEdge router sees network traffic in both directions. In dual-vEdge sites without any policy enabled, equal cost paths exist over each transport and to each vEdge router, and network traffic is hashed depending on fields in the IP header. Traffic is unlikely to always be forwarded to the same vEdge router in both the LAN-to-WAN direction and the WAN-to-LAN direction. To maintain symmetric traffic, it is recommended to set up routing so that traffic prefers one vEdge over another at dual-vEdge router sites.

To ensure symmetry, traffic needs to prefer one router in both directions, from the LAN to the WAN and from the WAN to the LAN. There are different ways to accomplish this.

To influence traffic in the LAN to WAN direction:

- For VRRP, use VRRP priority to prefer one vEdge router over the other.
- For OSPF, use the cost metric, configured either on the interface of the neighboring switch itself or through a route policy on the vEdge router that modifies the metric of routes redistributed from OMP to OSPF.
- For BGP, use a route policy and set AS path prepend or Multi-Exit Discriminator (MED) on routes redistributed from OMP to BGP.

Figure 14. Influencing traffic in the LAN-to-WAN direction

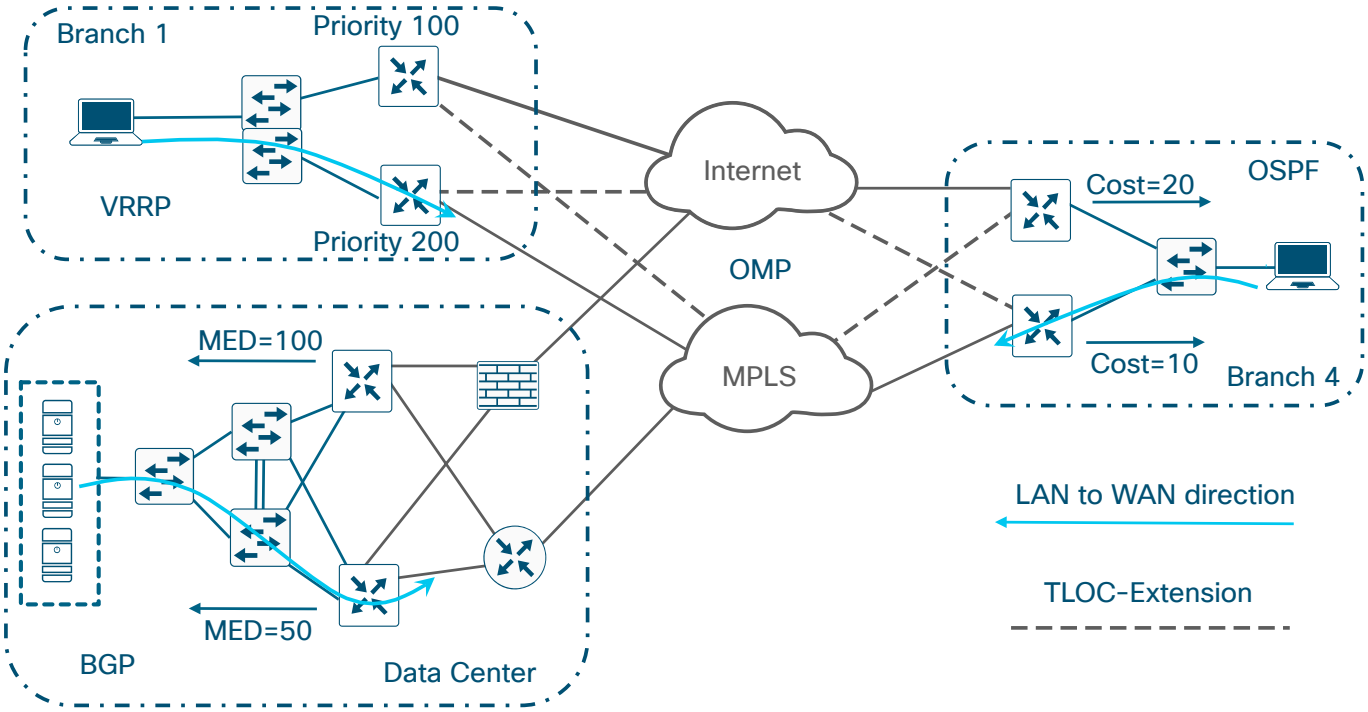

To influence traffic in the WAN-to-LAN direction over the overlay, you can influence an OMP attribute or set the IPSec tunnel preference. When BGP or OSPF is redistributed into OMP, the MED setting for BGP and the cost for OSPF is automatically translated into the OMP origin metric, which is used in the decision making for picking the best route.

Some common methods to influence traffic for the WAN-to-LAN direction:

- For BGP, use a route-policy and set MED (metric) on routes inbound from the LAN BGP neighbors
- For OSPF, use vEdge router interface cost to set the metric on routes coming into the LAN interface
- For any vEdge router, use IPSec tunnel preference to influence which is the preferred vEdge through the WAN overlay

Figure 15. Influencing traffic in the WAN-to-LAN direction

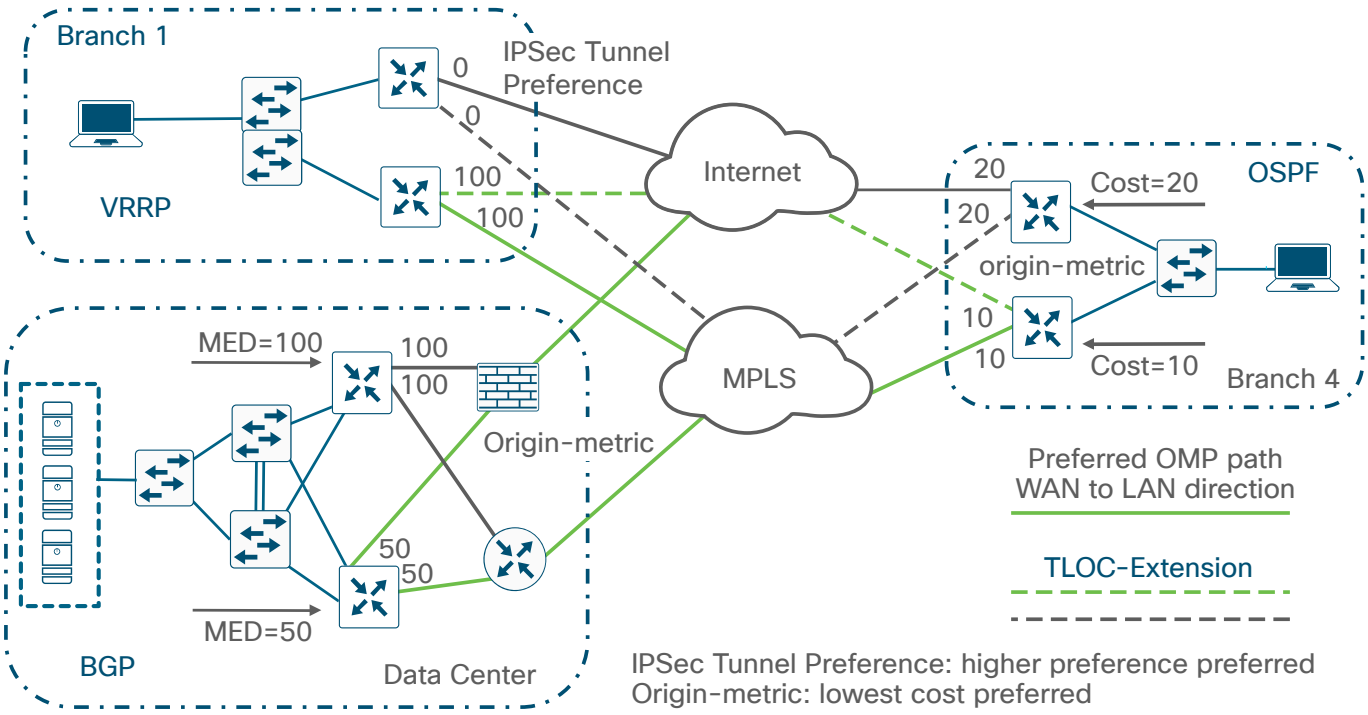

Quality of Service (QoS)

QoS is frequently deployed on the WAN transport, as bandwidth is less there, compared to the LAN side. QoS can only be applied to physical interfaces and not subinterfaces.

Each vEdge router physical interface has eight queues, labeled 0-7. Queue 0 uses Low Latency Queueing (LLQ), while queues 1-7 use Weighted Round Robin (WRR) for scheduling. Queue 0 (using LLQ) is a strict-priority traffic queue, meaning delay-sensitive traffic that is assigned to this queue is transmitted before packets in other queues. In addition, tail-drop is the only congestion avoidance algorithm used for this queue, and the queue is strictly policed.

By default, control traffic and BFD traffic (both marked as DSCP 48 decimal) use queue 0, while user traffic uses queue 2.

Figure 16. Quality of service queues for physical vEdge hardware

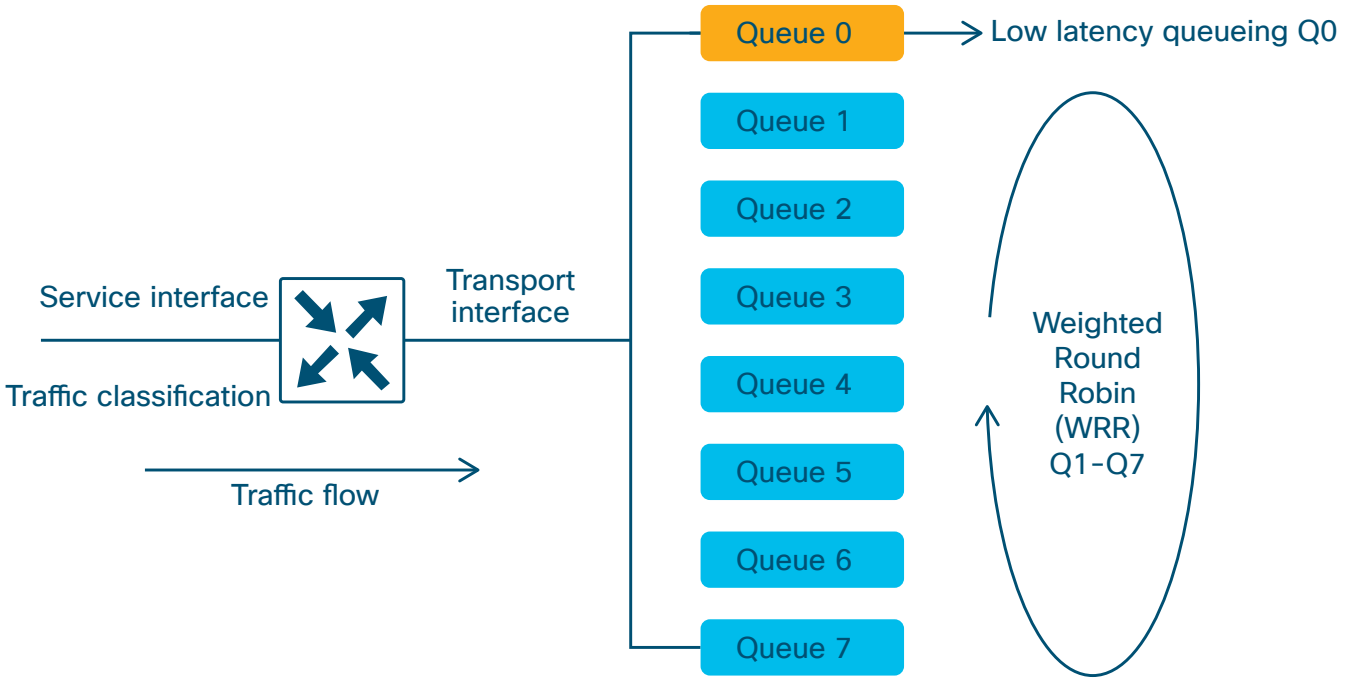

There are a few steps to take when configuring QoS on vEdge router hardware:

1. Map QoS forwarding classes to output queues. There are eight queues, 0 through 7. Queue 0 is reserved for control traffic and low-latency queueing (LLQ) traffic. If you assign traffic queue 0, it must be configured for LLQ scheduling. This is configured through localized policy.
2. Configure the QoS scheduler for each forwarding class. This assigns the scheduling method (LLQ or WRR), bandwidth percentage, buffer percentage, and drop algorithm (Random Early Detection [RED] or tail drop) to each forwarding class. The bandwidth and buffer percentages must add up to 100 percent. This is configured through localized policy.
3. Create a QoS map, where all of the QoS schedulers are grouped. This is configured through localized policy.
4. For vEdge 5000 routers (and vEdge cloud routers), use the cloud-qos command to enable QoS scheduling and shaping for the transport-side tunnel interfaces. Use the cloud-qos-service-side command to enable QoS scheduling and shaping for the service-side interfaces. This is configured within the localized policy.
5. Create a re-write policy (optional). This policy changes the DSCP value in the tunnel header and allows you to map to a smaller number of DSCP values that the service provider might support in the provider cloud. This is configured through localized policy.
6. Define an access list to match traffic and assign to forwarding classes. This can be configured either through centralized policy as a traffic data policy or through localized policy.
7. Apply the access list to an interface. It is typically applied as an inbound list on the service-side interface. This is configured either through centralized policy by specifying direction when the policy is applied (typically, the from-service option is used), or this can be configured by specifying the access list created in the localized policy in the VPN Ethernet template.
8. Apply the QoS map and, optionally, the re-write policy, to an egress interface. This is configured through the VPN interface Ethernet template of the desired interfaces in VPN 0.

Deployment planning

It is important to plan out your SD-WAN deployment carefully, as to make it easier for configuration, day-to-day operations, and maintenance. Following are some considerations.

Port numbering

It is recommended to have a port-numbering scheme that is consistent throughout the network. Consistency assists in easier configuration and troubleshooting.

In addition, the default factory configuration of a vEdge router specifies certain ports in VPN 0 for DHCP so the vEdge can automatically obtain a DHCP address, resolve DNS, and communicate with the ZTP server. So, if you utilize ZTP, be sure this port has reachability to the DHCP and DNS servers by connecting them to the most appropriate place in the network.

System IP

System IP is a persistent, system-level IPv4 address that uniquely identifies the device independently of any interface addresses. It acts much like a router ID, so it doesn't need to be advertised or known by the underlay. A best practice, however, is to advertise this system IP address in the service VPN and use it as a source IP address for SNMP and logging, making it easier to correlate network events with vManage information. A system IP address is required to be configured in order for a vEdge router to be authenticated by the controllers and brought into the overlay network.

A logical scheme for your system IP addresses is recommended to make sites more easily recognizable.

Site ID

A site ID is a unique identifier of a site in the SD-WAN overlay network with a numeric value 1 through 4294967295. This ID must be the same for all of the vEdge devices that reside at the same site. A site could be a data center, a branch office, a campus, or something similar. A site ID is required to be configured in order for a vEdge router to be authenticated by the controllers and brought into the overlay network. By default, IPSec tunnels are not formed between vEdge routers within the same site.

A site ID scheme should be chosen carefully, as this makes it easier to apply policy. When you apply policy, you apply policy to a list or range of site IDs (ex. 100,200-299), and there is no wildcard support.

Although there are several different ways to organize a site ID scheme, the following table provides an example of a scheme that uses nine digits.

Table 3. Nine-digit site ID example

Digit	Representation	Examples
1	Country/continent	1=North America, 2=Europe, 3=APAC
2	Region	1=US West, 2=US East, 3=Canada West, 4=Canada East
3-6	Site type	0000-0099=Hub locations, 1000-1999=Type 1 sites, 2000-2999=Type 2 sites, 3000-3999 = Type 3 sites, 4000-4999=Type 4 sites, 5000-9999 = future use
7-9	Store/site/branch number, or any other ID specifier	001, 002, 003

Grouping according to geography is helpful in cases where you might want to prefer a regional data center over another for centralized Internet access or for connectivity to hubs in other countries and regions.

Site types should be created according to types of policies applied in order to make applying policy easier. When a new site is created, just creating a site ID that falls into the matching range of a policy will automatically cause the policy to be applied to it. Some examples of how you may want to group branches according to type include:

- Branches that use a centrally-located firewall or another centrally-located service
- Branches that use Direct Internet Access
- Lower versus higher bandwidth sites since you may want different topologies for each. Low-bandwidth sites could use a hub-and-spoke topology to save bandwidth while higher bandwidth sites use a full-mesh topology.
- Different SLA and transport requirements, such as using MPLS for critical traffic, voice, and video while everything else traverses the Internet circuit, and perhaps some sites using MPLS for voice only, while everything else traverses the Internet circuit.

Obviously, you can have overlapping types, but the idea is to put them in categories that makes it easier to apply policy from a configuration perspective. It helps to think about the requirements and policies required before assigning site IDs.

Please use the feedback form to send comments and suggestions about this guide.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

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