



Policy Overview

Policy influences the flow of data traffic and routing information among Cisco vEdge devices in the overlay network.


Policy comprises:

- Routing policy—which affects the flow of routing information in the network's control plane.
- Data policy—which affects the flow of data traffic in the network's data plane.

To implement enterprise-specific traffic control requirements, you create basic policies, and deploy advanced features that are activated by means of the policy configuration infrastructure.

Just as the Cisco SD-WAN overlay network architecture clearly separates the control plane from the data plane and control between centralized and localized functions, the Cisco SD-WAN policy is cleanly separated. Policies apply either to control plane or data plane traffic, and they are configured either centrally on Cisco vSmart Controllers or locally on Cisco vEdge devices. The following figure illustrates the division between control and data policy, and between centralized and local policy.

Figure 1: Policy Architecture



368464

Control and Data Policy

Control policy is the equivalent of routing protocol policy, and data policy is equivalent to what are commonly called access control lists (ACLs) and firewall filters.

Centralized and Localized Policy

The Cisco SD-WAN policy design provides a clear separation between centralized and localized policy. In short, centralized policy is provisioned on the centralized Cisco vSmart Controllers in the overlay network, and the localized policy is provisioned on Cisco vEdge devices, which sit at the network edge between a branch or enterprise site and a transport network, such as the Internet, MPLS, or metro Ethernet.

Centralized Policy

Centralized policy refers to policy provisioned on Cisco vSmart Controllers, which are the centralized controllers in the Cisco SD-WAN overlay network. Centralized policy comprises two components:

- Control policy, which affects the overlay network–wide routing of traffic
- Data policy, which affects the data traffic flow throughout the VPN segments in the network

Centralized control policy applies to the network-wide routing of traffic by affecting the information that is stored in the Cisco vSmart Controller's route table and that is advertised to the Cisco vEdge devices. The effects of centralized control policy are seen in how Cisco vEdge devices direct the overlay network's data traffic to its destination.



Note

The centralized control policy configuration itself remains on the Cisco vSmart Controller and is never pushed to local devices.

Centralized data policy applies to the flow of data traffic throughout the VPNs in the overlay network. These policies can permit and restrict access based either on a 6-tuple match (source and destination IP addresses and ports, DSCP fields, and protocol) or on VPN membership. These policies are pushed to the selected Cisco vEdge devices.

Localized Policy

Localized policy refers to a policy that is provisioned locally through the CLI on the Cisco vEdge devices, or through a Cisco vManage device template.

Localized control policy is also called as route policy, which affects (BGP and OSPF) routing behavior on the site-local network.

Localized data policy allows you to provision access lists and apply them to a specific interface or interfaces on the device. Simple access lists permit and restrict access based on a 6-tuple match (source and destination IP addresses and ports, DSCP fields, and protocol), in the same way as with centralized data policy. Access lists also allow provisioning of class of service (CoS), policing, and mirroring, which control how data traffic flows out of and in to the device's interfaces and interface queues.

The design of the Cisco SD-WAN policy distinguishes basic and advanced policies. Basic policy allows you to influence or determine basic traffic flow through the overlay network. Here, you perform standard policy tasks, such as managing the paths along which traffic is routed through the network, and permitting or blocking traffic based on the address, port, and DSCP fields in the packet's IP header. You can also control the flow of data traffic into and out of a Cisco vEdge device's interfaces, enabling features such as class of service and queuing, mirroring, and policing.

Advanced features of Cisco SD-WAN policy offer specialized policy-based network applications. Examples of these applications include the following:

- Service chaining, which redirects data traffic to shared devices in the network, such as firewall, intrusion detection and prevention (IDS), load balancer, and other devices, before the traffic is delivered to its destination. Service chaining obviates the need to have a separate device at each branch site.
- Application-aware routing, which selects the best path for traffic based on real-time network and path performance characteristics.
- Cflowd, for monitoring traffic flow.
- Converting a Cisco vEdge device into a NAT device, to allow traffic destined for the Internet or other public network can exit directly from the Cisco vEdge device.

By default, no policy of any kind is configured on Cisco vEdge devices, either on the centralized Cisco vSmart Controllers or the local Cisco vEdge devices. When control plane traffic, which distributes route information, is unpoliced:

- All route information that OMP propagates among the Cisco vEdge devices is shared, unmodified, among all Cisco vSmart Controllers and all Cisco vEdge devices in the overlay network domain.
- No BGP or OSPF route policies are in place to affect the route information that Cisco vEdge devices propagate within their local site network.

When data plane traffic is unpoliced, all data traffic is directed towards its destination based solely on the entries in the local Cisco vEdge device's route table, and all VPNs in the overlay network can exchange data traffic.

- [Policy Architecture, on page 3](#)
- [Cisco vSmart Policy Components, on page 11](#)
- [Design Cisco vSmart Controller Policy Processing and Application, on page 17](#)
- [Cisco vSmart Policy Operation, on page 18](#)
- [Configure and Execute Cisco vSmart Policies, on page 24](#)

Policy Architecture

This topic offers an orientation about the architecture of the Cisco SD-WAN policy used to implement overlay network-wide policies. These policies are called **vSmart policy** or **centralized policy**, because you configure them centrally on a Cisco vSmart Controller. Cisco vSmart policy affects the flow of both control plane traffic (routing updates carried by Overlay Management Protocol (OMP) and used by the Cisco vSmart Controllers to determine the topology and status of the overlay network) and data plane traffic (data traffic that travels between the Cisco vEdge devices across the overlay network).

With Cisco SD-WAN, you can also create routing policies on the Cisco vEdge devices. These policies are simply traditional routing policies that are associated with routing protocol (BGP or OSPF) locally on the devices. You use them in the traditional sense for controlling BGP and OSPF, for example, to affect the exchange of route information, to set route attributes, and to influence path selection.

Centralized Control Policy Architecture

In the Cisco SD-WAN network architecture, centralized control policy is handled by the Cisco vSmart Controller, which effectively is the routing engine of the network. The Cisco vSmart Controller is the centralized manager of network-wide routes, maintaining a primary route table for these routes. The Cisco vSmart Controller builds its route table based on the route information advertised by the Cisco vEdge devices in its

Centralized control policy allows you to influence the network routes advertised by the Cisco vSmart Controllers. This type of policy, which is provisioned centrally on the Cisco vSmart Controller, affects both the route information that the Cisco vSmart Controller stores in its primary route table and the route information that it distributes to the devices.

Within a network domain, the network topology map on all Cisco vSmart Controllers must be synchronized. To support this, you must configure identical policies on all the Cisco vSmart Controllers in the domain.

The diagram shows a network topology where two devices, Device-1 and Device-2, establish a DTLS connection through the Internet. A central vSmart router acts as the intermediary, connected to a Route Table. The inbound path from Device-1 to the vSmart router is labeled 'Inbound Policy' and 'DTLS Connection'. The outbound path from the vSmart router to Device-2 is labeled 'Outbound Policy' and 'DTLS Connection'. The Internet is represented by a cloud in the center.

Control policy can be applied both inbound, to the route advertisements that the Cisco vSmart Controller receives from the devices, and outbound, to advertisements that it sends to them. Inbound policy controls which routes and route information are installed in the local routing database on the Cisco vSmart Controller, and whether this information is installed as-is or is modified. Outbound control policy is applied after a route is retrieved from the routing database, but before a Cisco vSmart Controller advertises it, and affects whether the route information is advertised as-is or is modified.

Route Types

The Cisco vSmart Controller learns the network topology from OMP routes, which are Cisco SD-WAN-specific routes carried by OMP. There are three types of OMP routes:

- Cisco SD-WAN OMP routes—These routes carry prefix information that the devices learn from the routing protocols running on its local network, including routes learned from BGP and OSPF, as well as direct, connected, and static routes. OMP advertises OMP routes to the Cisco vSmart Controller by means of an OMP route SAFI (Subsequent Address Family Identifier). These routes are commonly simply called OMP routes.
- TLOC routes—These routes carry properties associated with transport locations, which are the physical points at which the devices connect to the WAN or the transport network. Properties that identify a TLOC include the IP address of the WAN interface and a color that identifies a particular traffic flow. OMP advertises TLOC routes using a TLOC SAFI.
- Service routes—These routes identify network services, such as firewalls and IDPs, that are available on the local-site network to which the devices are connected. OMP advertises these routes using a service SAFI.

Default Behavior Without Centralized Control Policy

By default, no centralized control policy is provisioned on the Cisco vSmart Controller. This results in the following route advertisement and redistribution behavior within a domain:

- All Cisco vEdge devices redistribute all the route-related prefixes that they learn from their site-local network to the Cisco vSmart Controller. This route information is carried by OMP route advertisements that are sent over the DTLS connection between the devices and the Cisco vSmart Controller. If a domain contains multiple Cisco vSmart Controllers, the devices send all OMP route advertisements to all the controllers.
- All the devices send all TLOC routes to the Cisco vSmart Controller or controllers in their domain, using OMP.
- All the devices send all service routes to advertise any network services, such as firewalls and IDPs, that are available at the local site where the device is located. Again, these are carried by OMP.
- The Cisco vSmart Controller accepts all the OMP, TLOC, and service routes that it receives from all the devices in its domain, storing the information in its route table. The Cisco vSmart Controller tracks which OMP routes, TLOCs, and services belong to which VPNs. The Cisco vSmart Controller uses all the routes to develop a topology map of the network and to determine routing paths for data traffic through the overlay network.
- The Cisco vSmart Controller redistributes all information learned from the OMP, TLOC, and service routes in a particular VPN to all the devices in the same VPN.
- The devices regularly send route updates to the Cisco vSmart Controller.
- The Cisco vSmart Controller recalculates routing paths, updates its route table, and advertises new and changed routing information to all the devices.

Behavior Changes with Centralized Control Policy

When you do not want to redistribute all route information to all Cisco vEdge devices in a domain, or when you want to modify the route information that is stored in the Cisco vSmart Controller's route table or that is

advertised by the Cisco vSmart Controller, you design and provision a centralized control policy. To activate the control policy, you apply it to specific sites in the overlay network in either the inbound or the outbound direction. The direction is with respect to the Cisco vSmart Controller. All provisioning of centralized control policy is done on the Cisco vSmart Controller.

Applying a centralized control policy in the inbound direction filters or modifies the routes being advertised by the Cisco vEdge device before they are placed in the route table on the Cisco vSmart Controller. As the first step in the process, routes are either accepted or rejected. Accepted routes are installed in the route table on the Cisco vSmart Controller either as received or as modified by the control policy. Routes that are rejected by a control policy are silently discarded.

Applying a control policy in outbound direction filters or modifies the routes that the Cisco vSmart Controller redistributes to the Cisco vEdge devices. As the first step of an outbound policy, routes are either accepted or rejected. For accepted routes, centralized control policy can modify the routes before they are distributed by the Cisco vSmart Controller. Routes that are rejected by an outbound policy are not advertised.

VPN Membership Policy

A second type of centralized data policy is VPN membership policy. It controls whether a Cisco vEdge device can participate in a particular VPN. VPN membership policy defines which VPNs of a device is allowed and which is not allowed to receive routes from.

VPN membership policy can be centralized, because it affects only the packet headers and has no impact on the choice of interface that a Cisco vEdge device uses to transmit traffic. What happens instead is that if, because of a VPN membership policy, a device is not allowed to receive routes for a particular VPN, the Cisco vSmart Controller never forwards those routes to that driver.


Examples of Modifying Traffic Flow with Centralized Control Policy

This section provides some basic examples of how you can use centralized control policies to modify the flow of data traffic through the overlay network.

Create an Arbitrary Topology

When data traffic is exchanged between two Cisco vEdge devices, if you have provisioned no control policy, the two devices establish an IPsec tunnel between them and the data traffic flows directly from one device to the next. For a network with only two devices or with just a small number of devices, establishing connections between each pair of devices is generally not been an issue. However, such a solution does not scale. In a network with hundreds or even thousands of branches, establishing a full mesh of IPsec tunnels tax the CPU resources of each device.

Figure 3: Arbitrary Topology



One way to minimize this overhead is to create a hub-and-spoke type of topology in which one of the devices acts as a hub site that receives the data traffic from all the spoke, or branch, devices and then redirects the traffic to the proper destination. This example shows one of the ways to create such a hub-and-spoke topology, which is to create a control policy that changes the address of the TLOC associated with the destination.


The figure illustrates how such a policy might work. The topology has two branch locations, West and East. When no control policy is provisioned, these two devices exchange data traffic with each other directly by creating an IPsec tunnel between them (shown by the red line). Here, the route table on the Device West contains a route to Device East with a destination TLOC of 203.0.113.1, color gold (which we write as the tuple {192.0.2.1, gold}), and Device East route table has a route to the West branch with a destination TLOC of {203.0.113.1, gold}.

To set up a hub-and-spoke-type topology here, we provision a control policy that causes the West and East devices to send all data packets destined for the other device to the hub device. (Remember that because control policy is always centralized, you provision it on the Cisco vSmart Controller.) On the Device West, the policy simply changes the destination TLOC from {203.0.113.1, gold} to {209.165.200.225, gold}, which is the TLOC of the hub device, and on the Device East, the policy changes the destination TLOC from {192.0.2.1, gold} to the hub's TLOC, {209.165.200.225, gold}. If there were other branch sites on the west and east sides of the network that exchange data traffic, you could apply these same two control policies to have them redirect all their data traffic through the hub.

Set Up Traffic Engineering

Control policy allows you to design and provision traffic engineering. In a simple case, suppose that you have two devices acting as hub devices. If you want data traffic destined to a branch Cisco vEdge device to always transit through one of the hub devices, set the TLOC preference value to favor the desired hub device.

Figure 4: Traffic Engineering Topology




368415

The figure shows that Site ID 100 has two hub devices, one that serves the West side of the network and a second that serves the East side. Data traffic from the Device West must be handled by the Device West hub, and similarly, data traffic from the Device East branch must go through the Device East hub.

To engineer this traffic flow, you provision two control policies, one for Site ID 1, where the Device West device is located, and a second one for Site ID 2. The control policy for Site ID 1 changes the TLOC for traffic destined to the Device East to {209.165.200.225, gold}, and the control policy for Site ID 2 changes the TLOC for traffic destined for Site ID 1 to {198.51.100.1, gold}. One additional effect of this traffic engineering policy is that it load-balances the traffic traveling through the two hub devices.

With such a traffic engineering policy, a route from the source device to the destination device is installed in the local route table, and traffic is sent to the destination regardless of whether the path between the source and destination devices is available. Enabling end-to-end tracking of the path to the ultimate destination allows the Cisco vSmart Controller to monitor the path from the source to the destination, and to inform the source device when that path is not available. The source device can then modify or remove the path from its route table.


Figure 5: Traffic Engineering 2



The figure Traffic Engineering 2 illustrates end-to-end path tracking. It shows that traffic from Device-A that is destined for Device-D first goes to an intermediate device, Device-B, perhaps because this intermediate device provides a service, such as a firewall. (You configure this traffic engineering with a centralized control policy that is applied to Device-A, at Site 1.) Then Device-B, which has a direct path to the ultimate destination, forwards the traffic to Device-D. So, in this example, the end-to-end path between Device-A and Device-D comprises two tunnels, one between Device-A and Device-B, and the second between Device-B and Device-D. The Cisco vSmart Controller tracks this end-to-end path, and it notifies Device-A if the portion of the path between Device-B and Device-D becomes unavailable.

As part of end-to-end path tracking, you can specify how to forward traffic from the source to the ultimate destination using an intermediate device. The default method is strict forwarding, where traffic is always sent from Device-A to Device-B, regardless of whether Device-B has a direct path to Device-D or whether the tunnel between Device-B and Device-D is up. More flexible methods forward some or all traffic directly from Device-A to Device-D. You can also set up a second intermediate device to provide a redundant path with the first intermediate device is unreachable and use an ECMP method to forward traffic between the two. The figure Traffic Engineering3 adds Device-C as a redundant intermediate device.

Figure 6: Traffic Engineering 3



Centralized control policy, which you configure on Cisco vSmart Controllers, affects routing policy based on information in OMP routes and OMP TLOCs.

This type of policy allows you to set actions for matching routes and TLOCs, including redirecting packets through network services, such as firewalls, a feature that is called service chaining.

In domains with multiple Cisco vSmart Controllers, all the controllers must have the same centralized control policy configuration to ensure that routing within the overlay network remains stable and predictable.


Configure Centralized Policy Based on Prefixes and IP Headers

A centralized data policy based on source and destination prefixes and on headers in IP packets consists of a series of numbered (ordered) sequences of match-action pair that are evaluated in order, from lowest sequence number to highest sequence number. When a packet matches one of the match conditions, the associated action is taken and policy evaluation on that packets stops. Keep this in mind as you design your policies to ensure that the desired actions are taken on the items subject to policy.

If a packet matches no parameters in any of the sequences in the policy configuration, it is dropped and discarded by default.

Configuration Components

The following figure illustrates the configuration components for a centralized data policy:



Cisco vSmart Policy Components

The Cisco vSmart policies that implement overlay network-wide policies are implemented on a Cisco vSmart Controller. Because Cisco vSmart Controllers are centralized devices, you can manage and maintain Cisco vSmart policies centrally, and you can ensure consistency in the enforcement of policies across the overlay network.

The implementation of Cisco vSmart policy is done by configuring the entire policy on the Cisco vSmart Controller. Cisco vSmart policy configuration is accomplished with three building blocks:

- Lists define the targets of policy application or matching.
- Policy definition, or policies, controls aspects of control and forwarding. There are different types of policy, including:
 - app-route-policy (for application-aware routing)
 - cflowd-template (for cflowd flow monitoring)
 - control-policy (for routing and control plane information)
 - data-policy (for data traffic)
 - vpn-membership-policy (for limiting the scope of traffic to specific VPNs)
- Policy application controls what a policy is applied towards. Policy application is site-oriented, and is defined by a specific list called a site-list.

You assemble these three building blocks to Cisco vSmart policy. More specifically, policy is the sum of one or more lists, one policy definition, and at least one policy applications, as shown in the table below.

Table 1: The Three Building Blocks of Cisco vSmart Policies

Lists		Policy Definition		Policy Application
<code>data-prefix-list</code>: List of prefixes for use with a <code>data-policy</code> <code>prefix-list</code>: List of prefixes for use with any other policy <code>site-list</code>: List of site-id:s for use in policy and <code>apply-policy</code> <code>tloc-list</code> : List of tloc:s for use in policy <code>vpn-list</code> : List of vpn:s for use in policy	+	<code>app-route-policy</code>: Used with <code>sla-classes</code> for application-aware routing <code>cflowd-template</code>: Configures the <code>cflowd</code> agents on the Cisco vEdge devices <code>control-policy</code>: Controls OMP routing control <code>data-policy</code>: Provides vpn-wide policy-based routing <code>vpn-membership-policy</code>: Controls vpn membership across nodes	+	<code>apply-policy</code>: Used with a <code>site-list</code> to determine where policies are applied
=				
Complete policy definition configured on Cisco vSmart and enforced either on Cisco vSmart or on Cisco vEdge devices.				

Lists

Lists are how you group related items so that you can reference them all together. Examples of items you put in lists are prefixes, TLOCs, VPNs, and overlay network sites. In the Cisco vSmart Controller policy, you invoke lists in two places: when you create a policy definition and when you apply a policy. Separating the definition of the related items from the definition of policy means that when you can add or remove items from a lists, you make the changes only in a single place: You do not have to make the changes through the policy definition. So if you add ten sites to your network and you want to apply an existing policy to them, you simply add the site identifiers to the site list. You can also change policy rules without having to manually modify the prefixes, VPNs, or other things that the rules apply to.

Table 2: List Types

List type	Usage
data-prefix-list	Used in <code>data-policy</code> to define prefix and upper layer ports, either individually or jointly, for traffic matching.
prefix-list	Used in <code>control-policy</code> to define prefixes for matching RIB entries.
site-list	Used in <code>control-policy</code> to match source sites, and in <code>apply-policy</code> to define sites for policy application.
tloc-list	Used in <code>control-policy</code> to define TLOCs for matching RIB entries and to apply redefined TLOCs to vRoutes.

List type	Usage
vpn-list	Used in control-policy to define prefixes for matching RIB entries, and in data-policy and app-route-policy to define VPNs for policy application.

The following configuration shows the types of Cisco vSmart Controller policy lists:

```

policy
  lists
    data-prefix-list appl
      ip-prefix 209.165.200.225/27 port 100
    !
    prefix-list pfx1
      ip-prefix 209.165.200.225/27
    !
    site-list sitel
      site-id 100
    !
    tloc-list sitel-tloc
      tloc 209.165.200.225 color mpls
    vpn-list vpn1
      vpn1
    !
  !

```

Policy Definition

The policy definition is where you create the policy rules. You specify match conditions (route-related properties for control policy and data-related fields for data policy) and actions to perform when a match occurs. A policy contains match–action pairings that are numbered and that are examined in sequential order. When a match occurs, the action is performed, and the policy analysis on that route or packet terminates. Some types of policy definitions apply only to specific VPNs.

Table 3: Policy Types

Policy type	Usage
policy-type	Can be control-policy , data-policy , or vpn-membership —dictates the type of policy. Each type has a particular syntax and a particular set of match conditions and settable actions.
vpn-list	Used by data-policy and app-route-policy to list the VPNs for which the policy is applicable.
sequence	Defines each sequential step of the policy by sequence number.
match	Decides what entity to match on in the specific policy sequence.
action	Determines the action that corresponds to the preceding match statement.

Policy type	Usage
default-action	Action to take for any entity that is not matched in any sequence of the policy. By default, the action is set to reject.

The following configuration shows the components of the Cisco vSmart Controller policy definition. These items are listed in the logical order you should use when designing policy, and this order is also how the items are displayed in the configuration, regardless of the order in which you add them to the configuration.

```

policy
  policy-type name
  vpn-list vpn-list
  sequence number
  match
    <route | tloc vpn | other>
  !
  action <accept reject drop>
  set attribute value
  !
  default-action <reject accept>
  !
  !
  !

```

Policy Application

The following are the configuration components:

Component	Usage
site-list	Determines the sites to which a given policy is applies. The direction (in out) applies only to control-policy.
policy-type	The policy type can be control-policy , data-policy , or vpn-membership —and name refer to an already configured policy to be applied to the sites specified in the site-list for the section.

For a policy definition to take effect, you associate it with sites in the overlay network.

```

apply-policy
  site-list name
  control-policy name <inout>
  !
  site-list name
  data-policy name
  vpn-membership name
  !
  !

```

Policy Example

For a complete policy, which consists of lists, policy definition, and policy application. The example illustrated below creates two lists (a site-list and a tloc-list), defines one policy (a control policy), and applies the policy to the site-list. In the figure, the items are listed as they are presented in the node configuration. In a normal configuration process, you create lists first (group together all the things you want to use), then define the

policy itself (define what things you want to do), and finally apply the policy (specify the sites that the configured policy affects).

```

apply-policy
  site-list sitel -----> Apply the defined policy towards the sites in site-list
  control-policy prefer_local out
  !
policy
  lists
  site-list sitel
  site-id 100
  tloc-list prefer_sitel ----> Define the lists required for apply-policy and for use within
  the policy
  tloc 192.0.2.1 color mols encap ipsec preference 400
  control-policy prefer_local
  sequence 10
  match route
  site-list sitele ----->Lists previously defined used within policy
  !
  action accept
  set
  tloc-list prefer_site
  !
  !
  !

```

TLOC Attributes Used in Policies

A transport location, or TLOC, defines a specific interface in the overlay network. Each TLOC consists of a set of attributes that are exchanged in OMP updates among the Cisco SD-WAN devices. Each TLOC is uniquely identified by a 3-tuple of IP address, color, and encapsulation. Other attributes can be associated with a TLOC.

The TLOC attributes listed below can be matched or set in Cisco vSmart Controller policies.

Table 4:

TLOC Attribute	Function	Application Point Set By	Application Point Modify By
Address (IP address)	system-ip address of the source device on which the interface is located.	Configuration on source device	control-policy data-policy
Carrier	Identifier of the carrier type. It primarily indicates whether the transport is public or private.	Configuration on source device	control-policy
Color	Identifier of the TLOC type.	Configuration on source device	control-policy data-policy
Domain ID	Identifier of the overlay network domain.	Configuration on source device	control-policy
Encapsulation	Tunnel encapsulation, either IPsec or GRE.	Configuration on source device	control-policy data-policy

TLOC Attribute	Function	Application Point Set By	Application Point Modify By
Originator	system-ip address of originating node.	Configuration on any originator	control-policy
Preference	OMP path-selection preference. A higher value is a more preferred path.	Configuration on source device	control-policy
Site ID	Identification for a give site. A site can have multiple nodes or TLOCs.	Configuration on source device	control-policy
Tag	Identifier of TLOC on any arbitrary basis.	Configuration on source device	control-policy

vRoute Attributes Used in Policies

A Cisco SD-WAN route, or vRoute, defines a route in the overlay network. A vRoute, which is similar to a standard IP route, has a number attributes such as TLOC and VPN. The Cisco vEdge devices exchange vRoutes in OMP updates.

The vRoutes attributes listed below can be matched or set in Cisco vSmart Controller policies.

Table 5:

vRoute Attribute	Function	Application Point Set By	Application Point Modify By
Origin	Source of the route, either BGP, OSPF, connected, static.	Source device	control-policy
Originator	Source of the update carrying the route.	Any originator	control-policy
Preference	OMP path-selection preference. A higher value is a more preferred path.	Configuration on source device or policy	control-policy
Service	Advertised service associated with the vRoute.	Configuration on source device	control-policy
Site ID	Identifier for a give site. A site can have multiple nodes or TLOCs.	Configuration on source device	control-policy
Tag	Identification on any arbitrary basis.	Configuration on source device	control-policy
TLOC	TLOC used as next hop for the vRoute.	Configuration on source device or policy	control-policy data-policy
VPN	VPN to which the vRoute belongs.	Configuration on source device or policy	control-policy data-policy

Design Cisco vSmart Controller Policy Processing and Application

Understanding how a Cisco vSmart Controller policy is processed and applied allows for proper design of policy and evaluation of how policy is implemented across the overlay network.

Policy is processed as follows:

- A policy definition consists of a numbered, ordered sequence of match–action pairings. Within each policy, the pairings are processed in sequential order, starting with the lowest number and incrementing.
- As soon as a match occurs, the matched entity is subject to the configured action of the sequence and is then no longer subject to continued processing.
- Any entity not matched in a sequence is subject to the default action for the policy. By default, this action is reject.

Cisco vSmart Controller policy is applied on a per-site-list basis, so:

- When applying policy to a site-list, you can apply only one of each type of policy. For example, you can have one control-policy and one data-policy, or one control-policy in and one control-policy out. You cannot have two data policies or two outbound control policies.
- Because a site-list is a grouping of many sites, you should be careful about including a site in more than one site-list. When the site-list includes a range of site identifiers, ensure that there is no overlap. If the same site is part of two site-lists and the same type of policy is applied to both site-lists, the policy behavior is unpredictable and possibly catastrophic.
- Control-policy is unidirectional, being applied either inbound to the vSmart controller or outbound from it. When control-policy is needed in both directions, configure two control policies.
- Data-policy is bidirectional and can be applied either to traffic received from the service side of the Cisco vEdge device, traffic received from the tunnel side, or all of these combinations.
- VPN membership policy is always applied to traffic outbound from the Cisco vSmart Controller.
- Control-policy remains on the Cisco vSmart Controller and affects routes that the controller sends and receives.
- Data-policy is sent to either the Cisco vEdge devices in the site-list. The policy is sent in OMP updates, and it affects the data traffic that the devices send and receive.
- When any node in the overlay network makes a routing decision, it uses any and all available routing information. In the overlay network, it is the Cisco vSmart Controller that distributes routing information to the Cisco vEdge device nodes.
- In a network deployment that has two or more Cisco vSmart Controllers, each controller acts independently to disseminate routing information to other Cisco vSmart Controllers and to Cisco vEdge devices in the overlay network. So, to ensure that the Cisco vSmart Controller policy has the desired effect in the overlay network, each Cisco vSmart Controller must be configured with the same policy, and the policy must be applied identically. For any given policy, you must configure the identical policy and apply it identically across all the Cisco vSmart Controllers.

Cisco vSmart Policy Operation

At a high level, control policy operates on routing information, which in the Cisco SD-WAN network is carried in OMP updates. Data policy affects data traffic, and VPN membership controls the distribution of VPN routing tables.

The basic Cisco vSmart policies are:

- Control Policy
- Data Policy
- VPN Membership

Control Policy


Control policy, which is similar to standard routing policy, operates on routes and routing information in the control plane of the overlay network. Centralized control policy, which is provisioned on the Cisco vSmart Controller, is the Cisco SD-WAN technique for customizing network-wide routing decisions that determine or influence routing paths through the overlay network. Local control policy, which is provisioned on a Cisco vEdge device, allows customization of routing decisions made by BGP and OSPF on site-local branch or enterprise networks.

The routing information that forms the basis of centralized control policy is carried in Cisco SD-WAN route advertisements, which are transmitted on the DTLS or TLS control connections between Cisco vSmart Controllers and Cisco vEdge devices. Centralized control policy determines which routes and route information are placed into the centralized route table on the Cisco vSmart Controller and which routes and route information are advertised to the Cisco vEdge devices in the overlay network. Basic centralized control policy establish traffic engineering, to set the path that traffic takes through the network. Advanced control policy supports a number of features, including service chaining, which allows Cisco vEdge devices in the overlay network to share network services, such as firewalls and load balancers.

Centralized control policy affects the OMP routes that are distributed by the Cisco vSmart Controller throughout the overlay network. The Cisco vSmart Controller learns the overlay network topology from OMP routes that are advertised by the Cisco vEdge devices over the OMP sessions inside the DTLS or TLS connections between the Cisco vSmart Controller and the devices.

Three types of OMP routes carry the information that the Cisco vSmart Controller uses to determine the network topology:

- Cisco SD-WAN OMP routes, which are similar to IP route advertisements, advertise routing information that the devices have learned from their local site and the local routing protocols (BGP and OSPF) to the Cisco vSmart Controller. These routes are also referred to as OMP routes or vRoutes.
- TLOC routes carry overlay network-specific locator properties, including the IP address of the interface that connects to the transport network, a link color, which identifies a traffic flow, and the encapsulation type. (A TLOC, or transport location, is the physical location where a Cisco vEdge device connects to a transport network. It is identified primarily by IP address, link color, and encapsulation, but a number of other properties are associated with a TLOC.)
- Service routes advertise the network services, such as firewalls, available to VPN members at the local site.

Figure 7: Control Policy Topology

By default, no centralized control policy is provisioned. In this bare, unpoliced network, all OMP routes are placed in the Cisco vSmart Controller's route table as is, and the Cisco vSmart Controller advertises all OMP routes, as is, to all the devices in the same VPN in the network domain.

By provisioning centralized control policy, you can affect which OMP routes are placed in the Cisco vSmart Controller's route table, what route information is advertised to the devices, and whether the OMP routes are modified before being put into the route table or before being advertised.

Cisco vEdge devices place all the route information learned from the Cisco vSmart Controllers, as is, into their local route tables, for use when forwarding data traffic. Because the Cisco vSmart Controller's role is to be the centralized routing system in the network, Cisco vEdge devices can never modify the OMP route information that they learn from the Cisco vSmart Controllers.

The Cisco vSmart Controller regularly receives OMP route advertisements from the devices and, after recalculating and updating the routing paths through the overlay network, it advertises new routing information to the devices.

The centralized control policy that you provision on the Cisco vSmart Controller remains on the Cisco vSmart Controller and is never downloaded to the devices. However, the routing decisions that result from centralized control policy are passed to the devices in the form of route advertisements, and so the affect of the control policy is reflected in how the devices direct data traffic to its destination.

A type of centralized control policy called service chaining allows data traffic to be routed through one or more network services, such as firewall, load balancer, and intrusion detection and prevention (IDP) devices, en route to its destination.

Localized control policy, which is provisioned locally on the devices, is called route policy. This policy is similar to the routing policies that you configure on a regular driver, allowing you to modify the BGP and OSPF routing behavior on the site-local network. Whereas centralized control policy affects the routing behavior across the entire overlay network, route policy applies only to routing at the local branch.


The Cisco SD-WAN devices periodically exchange OMP updates, which carry routing information pertaining to the overlay network. Two of the things that these updates contain are vRoute attributes and Transport Locations (TLOC) attributes.

The Cisco vSmart Controller uses these attributes from the OMP updates to determine the topology and status of the overlay network, and installs routing information about the overlay network into its route table. The controller then advertises the overlay topology to the Cisco vEdge devices in the network by sending OMP updates to them.

Control policy examines the vRoute and TLOC attributes carried in OMP updates and can modify attributes that match the policy. Any changes that results from control policy are applied directionally, either inbound or outbound.

The figure shows a control-policy named **prefer_local** that is configured on a Cisco vSmart Controller and that is applied to Site 1 (via site-list list1) and to Site 2 (via site-list list2).

Figure 8: Control Policy Topology



```
Device# apply-policy
site-list list1
control-policy prefer_local in
!
```

The upper left arrow shows that the policy is applied to Site 1—more specifically, to **site-list list1**, which contains an entry for Site 1. The command **control-policy prefer_local in** is used to apply the policy to OMP updates that are coming in to the Cisco vSmart Controller from the Cisco vEdge device, which is inbound from the perspective of the controller. The **in** keyword indicates an **inbound** policy. So, for all OMP updates that the Site 1 devices send to the Cisco vSmart Controller, the "prefer_local" control policy is applied before the updates reach the route table on the Cisco vSmart Controller. If any vRoute or TLOC attributes in an OMP update match the policy, any changes that result from the policy actions occur before the Cisco vSmart Controller installs the OMP update information into its route table.

The route table on the Cisco vSmart Controller is used to determine the topology of the overlay network. The Cisco vSmart Controller then distributes this topology information, again via OMP updates, to all the devices in the network. Because applying policy in the inbound direction influences the information available to the Cisco vSmart Controller. It determines the network topology and network reachability, modifying vRoute and TLOC attributes before they are placed in the controller's route table.

```
apply-policy
site-list list2
```

```
control-policy prefer_local out
!
```


On the right side of the figure above, the "prefer_local" policy is applied to Site 2 via the **control-policy prefer_local out** command. The **out** keyword in the command indicates an **outbound policy**, which means that the policy is applied to OMP updates that the Cisco vSmart Controller is sending to the devices at Site 2. Any changes that result from the policy occur, after the information from the Cisco vSmart Controller's route table is placed in to an OMP update and before the devices receive the update. Again, note that the direction is outbound from the perspective of the Cisco vSmart Controller.

In contrast to an inbound policy, which affects the centralized route table on the Cisco vSmart Controller and has a broad effect on the route attributes advertised to all the devices in the overlay network. A control policy applied in the outbound direction influences only the route tables on the individual devices included in the site-list.

The same control policy (the **prefer_local** policy) is applied to both the inbound and outbound OMP updates. However, the effects of applying the same policy to inbound and outbound are different. The usage shown in the figure illustrates the flexibility of the Cisco SD-WAN control policy design architecture and configuration.

Data Policy


Data policy influences the flow of data traffic traversing the network based either on fields in the IP header of packets or the router interface on which the traffic is being transmitted or received. Data traffic travels over the IPsec connections between Cisco vEdge devices, shown in purple in the adjacent figure.



The Cisco SD-WAN architecture implements two types of data policy:

- Centralized data policy controls the flow of data traffic based on the source and destination addresses and ports and DSCP fields in the packet's IP header (referred to as a 5-tuple), and based on network segmentation and VPN membership. These types of data policy are provisioned centrally, on the Cisco vSmart controller, and they affect traffic flow across the entire network.
- Localized data policy controls the flow of data traffic into and out of interfaces and interface queues on a Cisco vEdge device. This type of data policy is provisioned locally using access lists. It allows you to classify traffic and map different classes to different queues. It also allows you to mirror traffic and to police the rate at which data traffic is transmitted and received.

By default, no centralized data policy is provisioned. The result is that all prefixes within a VPN are reachable from anywhere in the VPN. Provisioning centralized data policy allows you to apply a 6-tuple filter that controls access between sources and destinations.




As with centralized control policy, you provision a centralized data policy on the Cisco vSmart Controller, and that configuration remains on the Cisco vSmart Controller. The effects of data policy are reflected in how the Cisco vEdge devices direct data traffic to its destination. Unlike control policy, however, centralized data policies are pushed to the devices in a read-only fashion. They are not added to the router's configuration file, but you can view them from the CLI on the router.

With no access lists provisioned on a Cisco vEdge device, all data traffic is transmitted at line rate and with equal importance, using one of the interface's queues. Using access lists, you can provision class of service, which allows you to classify data traffic by importance, spread it across different interface queues, and control the rate at which different classes of traffic are transmitted. You can provision policing. You can also provision packet mirroring.

Data policy examines fields in the headers of data packets, looking at the source and destination addresses and ports, and the protocol and DSCP values, and for matching packets, it can modify the next hop in a variety of ways or apply a policer to the packets. Data policy is configured and applied on the Cisco vSmart Controller, and then it is carried in OMP updates to the Cisco vEdge devices in the site-list that the policy is applied to. The match operation and any resultant actions are performed on the devices as it transmits or receives data traffic.

In the Data Policy Topology figure, a data policy named “change_next_hop” is applied to a list of sites that includes Site 3. The OMP update that the vSmart controller sends to the devices at Site 3 includes this policy definition. When the device sends or receives data traffic that matches the policy, it changes the next hop to the specified TLOC. Nonmatching traffic is forwarded to the original next-hop TLOC.

Figure 9: Data Policy Topology



In the **apply-policy** command for a data policy, specify a direction from the perspective of the device. The "all" direction in the figure applies the policy to incoming and outgoing data traffic transiting the tunnel interface. You can limit the span of the policy to only incoming traffic with a **data-policy change_next_hop from-tunnel** command or to only outgoing traffic with a **data-policy change_next_hop from-service** command.

VPN Membership Policy Operation


VPN membership policy, as the name implies, affects the VPN route tables that are distributed to particular Cisco vEdge devices. In an overlay network with no VPN membership policy, the Cisco vSmart Controller pushes the routes for all VPNs to all the devices. If your business usage model restricts participation of specific devices in particular VPNs, a VPN membership policy is used to enforce this restriction.

The figure VPN Membership Topology illustrates how VPN membership policy works. This topology has three Cisco vEdge devices:

- The Cisco vEdge devices at Sites 1 and 2 service only VPN 2.
- The Cisco vEdge devices at Site 3 services both VPN 1 and VPN 2.

In the figure, the device at Site 3 receives all route updates from the Cisco vSmart Controller, because these updates are for both VPN 1 and VPN 2. However, because the other Cisco vEdge devices service only VPN 2, it can filter the route updates sent to them, remove the routes associated with VPN 1 and sends only the ones that apply to VPN 2.

Figure 10: VPN Membership Topology



Notice that here, direction is not set when applying VPN membership policy. The Cisco vSmart Controller always applies this type of policy to the OMP updates that it sends outside to the Cisco vEdge devices.

Configure and Execute Cisco vSmart Policies

All Cisco vSmart Controller policies are configured on the Cisco vEdge devices, using a combination of policy definition and lists. All Cisco vSmart Controller policies are also applied on the Cisco vEdge devices, with a combination of apply-policy and lists. However, where the actual Cisco vSmart Controller policy executes depends on the type of policy, as shown in this figure:

Figure 11: Cisco vSmart Policy

 vSmart	Action	App-route Policy	Cflowd Template	Control Policy	Data Policy	VPN Membership Policy
	Configure	✓	✓	✓	✓	✓
	Apply	✓	✓	✓	✓	✓
	Execute			✓		✓

 Device	Action	App-route Policy	Cflowd Template	Control Policy	Data Policy	VPN Membership Policy
	Configure					
	Apply					
	Execute	✓	✓		✓	

368503

For control policy and VPN membership policy, the entire policy configuration remains on the Cisco vSmart Controller, and the actions taken as a result of routes or VPNs that match a policy are performed on the Cisco vSmart Controller.

For the other three policy types—application-aware routing, cflowd templates, and data policy—the policies are transmitted in OMP updates to the Cisco vEdge devices, and any actions taken as a result of the policies are performed on the devices.

