Control Policies

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 IOS XE SD-WAN 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 IOS XE SD-WAN route advertisements, which are transmitted on the DTLS or TLS control connections between Cisco vSmart Controllers and Cisco IOS XE SD-WAN 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 IOS XE SD-WAN 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, which allows Cisco IOS XE SD-WAN 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 IOS XE SD-WAN 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 IOS XE SD-WAN 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 1: Control Policy Topology

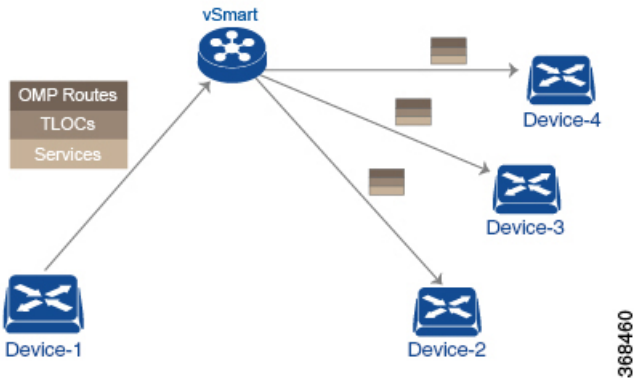

By default, no centralized control policy is provisioned. In this bare, unpolicied network, all OMP routes are placed in the Cisco vSmart Controller's route table as is, and the Cisco vSmart Controller advertises all OMP routes, as is, to all the devices in the same VPN in the network domain.

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 IOS XE SD-WAN 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 IOS XE SD-WAN 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.

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.

- Centralized Control Policy , on page 2
- Localized Control Policy, on page 37
- Device Access Policy, on page 54

Centralized Control Policy

In the Cisco IOS XE 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 IOS XE SD-WAN devices in its domain, using these routes to discover the network topology and to determine the best paths to

network destinations. The Cisco vSmart Controller distributes route information from its route table to the devices in its domain which in turn use these routes to forward data traffic through the network. The result of this architecture is that networking-wide routing decisions and routing policy are orchestrated by a central authority instead of being implemented hop by hop, by the devices in the network.

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.

Centralized control policy is provisioned and applied only on the Cisco vSmart Controller. The control policy configuration itself is never pushed to devices in the overlay network. What is pushed to the devices, using the Overlay Management Protocol (OMP), are the results of the control policy, which the devices then install in their local route tables and use for forwarding data traffic. This design means that the distribution of network-wide routes is always administered centrally, using policies designed by network administrators. These policies are always implemented by centralized Cisco vSmart Controllers, which are responsible for orchestrating the routing decisions in the Cisco IOS XE SD-WAN overlay network.

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.

Figure 2: Centralized Control Policy

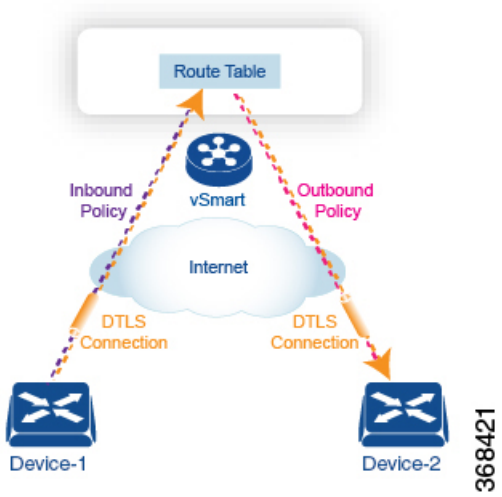

All centralized control plane traffic, including route information, is carried by OMP peering sessions that run within the secure, permanent DTLS connections between devices and the Cisco vSmart Controllers in their domain. The end points of an OMP peering session are identified by the system IDs of the devices, and the peering sessions carry the site ID, which identifies the site in which the device is located. A DTLS connection and the OMP session running over it remain active as long as the two peers are operational.

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 IOS XE SD-WAN-specific routes carried by OMP. There are three types of OMP routes:

- Cisco IOS XE 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.

The difference in these three types of routes can be viewed by using the various **show sdwan omp** operational commands when you are logged in to the CLI on a Cisco vSmart Controller or a Cisco IOS XE SD-WAN device. The **show sdwan omp routes** command displays information sorted by prefix, the **show sdwan omp services** command displays route information sorted by service, and the **show sdwan omp tlocs** command sorts route information by TLOC.

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 IOS XE SD-WAN 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 IOS XE SD-WAN 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 IOS XE SD-WAN 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 IOS XE SD-WAN 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 IOS XE SD-WAN 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 IOS XE SD-WAN 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 IOS XE SD-WAN 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

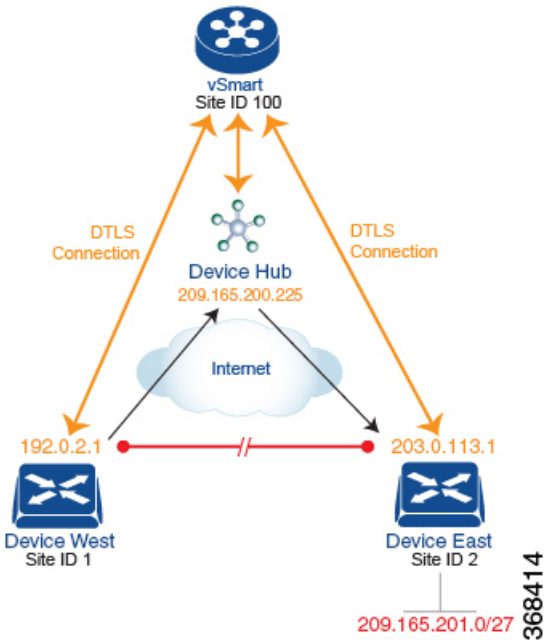

One way to minimize this overhead is to create a hub-and-spoke type of topology in which one of the devices acts as a hub site that receives the data traffic from all the spoke, or branch, devices and then redirects the traffic to the proper destination. This example shows one of the ways to create such 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 IOS XE SD-WAN device to always transit through one of the hub devices. To engineer this traffic flow, set the TLOC preference value to favor the desired hub device.

Figure 4: Traffic Engineering Topology

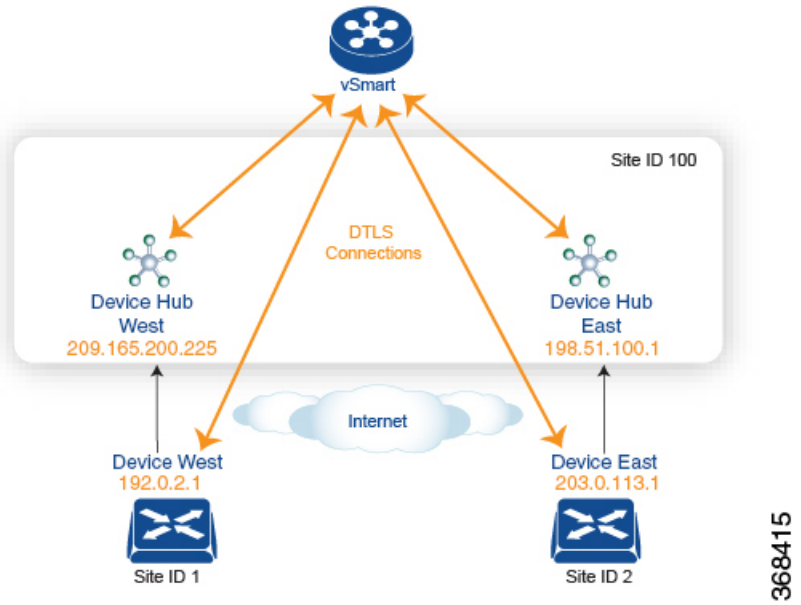

The figure shows that Site ID 100 has two hub devices, one that serves the West side of the network and a second that serves the East side. Data traffic from the Device West must be handled by the Device West hub, and similarly, data traffic from the Device East branch must go through the Device East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

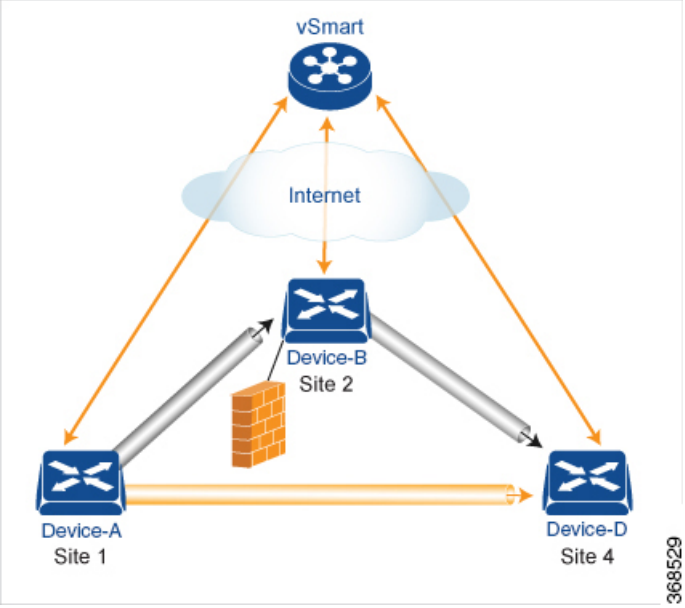

The figure Traffic Engineering 2 illustrates end-to-end path tracking. It shows that traffic from Device-A that is destined for Device-D first goes to an intermediate device, Device-B, perhaps because this intermediate device provides a service, such as a firewall. (You configure this traffic engineering 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 forwarded 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 Engineering3

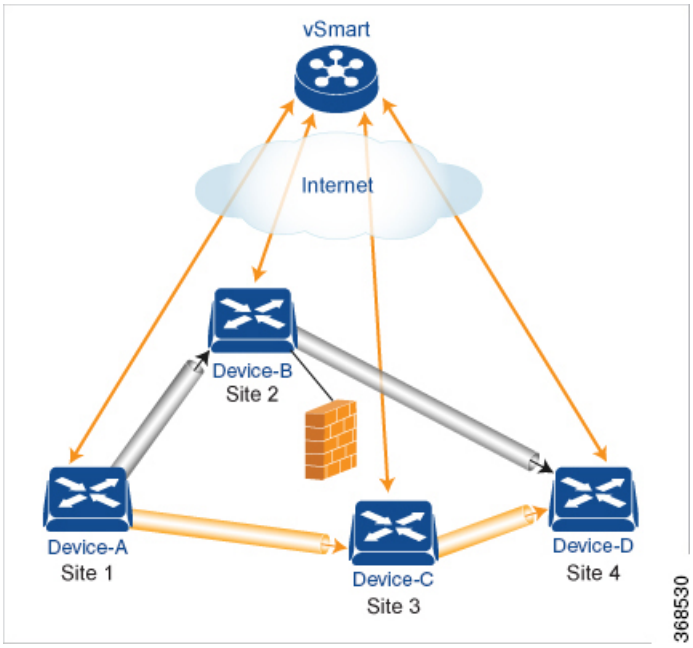

Centralized control policy, which you configure on Cisco vSmart Controllers, affects routing policy based on information in OMP routes and OMP TLOCs.

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 the Network Topology

When you first open the Configure Topology and VPN Membership screen, the **Topology** tab is selected by default.

To configure the network topology and VPN membership:

Configuration Components

A centralized control policy consists of a series of numbered (ordered) sequences of match-action pairs that are evaluated in order, from lowest sequence number to highest sequence number. When a route or TLOC matches the match conditions, the associated action or actions are taken and policy evaluation on that packets stops. Keep this process in mind as you design your policies to ensure that the desired actions are taken on the items subject to policy.

If a route or TLOC matches no parameters in any of the sequences in the policy configure, it is, by default, rejected and discarded.

The figure illustrates the configuration components for centralized control policy.

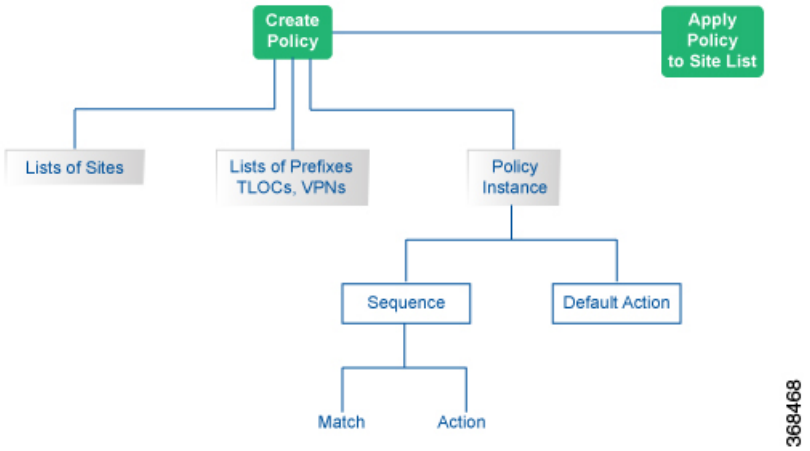

Create a Hub and Spoke Policy

Step 1 In the Add Topology drop-down, select **Hub and Spoke**.

Step 2 Enter a name for the hub-and-spoke policy.

Step 3 Enter a description for the policy.

Step 4 In the VPN List field, select the VPN list for the policy.

Step 5 In the left pane, click **Add Hub and Spoke**. A hub-and-spoke policy component containing the text string My Hub-and-Spoke is added in the left pane.

Step 6 Double-click the **My Hub-and-Spoke** text string, and enter a name for the policy component.

Step 7 In the right pane, add hub sites to the network topology:

a) Click **Add Hub Sites**.
b) In the **Site List Field**, select a site list for the policy component.
c) Click **Add**.
d) Repeat these steps to add more hub sites to the policy component.

Step 8 In the right pane, add spoke sites to the network topology:

a) Click **Add Spoke Sites**.
b) In the **Site List Field**, select a site list for the policy component.
c) Click **Add**.
d) Repeat these steps to add more spoke sites to the policy component.

Step 9 Repeat steps as needed to add more components to the hub-and-spoke policy.

Step 10 Click **Save Hub and Spoke Policy**.

Create a Policy for Mesh

Step 1 In the Add Topology drop-down, select **Mesh**.

Step 2 Enter a name for the mesh region policy component.

Step 3 Enter a description for the mesh region policy component.

Step 4 In the **VPN List** field, select the VPN list for the policy.
Step 5 Click **New Mesh Region**.
Step 6 In the **Mesh Region Nam**e field, enter a name for the individual mesh region.
Step 7 In the **Site List** field, select one or more sites to include in the mesh region.
Step 8 Repeat these steps to add more mesh regions to the policy.
Step 9 Click **Save Mesh Region**.

Custom Control (Route and TLOC)

Policy for a topology with custom route and TLOC configuration.

Step 1 In the Add Topology drop-down, select **Custom Control (Route & TLOC)**.
Step 2 Enter a name for the custom control policy component.
Step 3 Enter a description of the custom control policy component.
Step 4 Click **Sequence Type**. The Add Control Policy popup displays.
Step 5 Click **Route** or **TLOC** to create a policy of that type.
Step 6 Click **Sequence Rule**.

Custom Control (Route)

Create a policy to apply on an OMP route. By default, the Match tab is selected, displaying match condition options.

Step 1 From the **Add Custom Control Policy** screen, click **Route**.
Step 2 Click **Sequence Rule**. Match and Actions options display.
Step 3 From the Match tab, select and configure match conditions for your route.

Match Condition	Description
Color List	Select a color list to match, or click **New Color List** to create a new list: **a.** Enter a name for the Color list. **b.** From the **Select Color** drop-down menu, select the color(s) you want included in your list. **c.** Click **Save**.
OMP Tag	Enter the OMP route tag, a number between 0-4294967295.

Match Condition	Description
Origin	Select an origin for the route from the drop-down menu. Options include: • **Aggregate** • **BGP External** • **BGP Internal** • **Connected** • **OSPF** • **OSPF External 1** • **OSPF External 2** • **OSPF Intra-Area** • **Static.**
Originator	Enter the IP address of the originator of this route.
Preference	Enter the preference number for the route, a number between 0-4294967295.
Site	Select a site list from the list of options., or create a new site list: **a.** Enter a name for the Site list. **b.** Enter the Site numbers, following the example. **c.** Click **Save**.
TLOC	Select a TLOC list to match, or create a new TLOC list: **a.** Enter a name for the TLOC list. **b.** In the **TLOC IP** field, enter the IP address for the TLOC. **c.** In the **Color** drop-down menu, select the color you want to apply to the TLOC list. **d.** From the **Encap** drop-down menu, select the encapsulation type for the TLOC list. **e.** In the **Preference** field, enter the preference number for the route, a number between 0-4294967295. **f.** Optionally, click **Add TLOC** and repeat steps 1-5 to open another TLOC list. **g.** Click **Save**.

Match Condition	Description
VPN	**a.** From the **Match Conditions > VPN list** field, select a VPN list, or click **New VPN List** to create a new one: **b.** Enter a name for the VPN List. **c.** In the **VPN** field, enter the VPN numbers, for example, 100 or 200 separated by commas, or 1000-2000 by range. **d.** Click **Save**.
Prefix List	From the **Match Conditions > Prefix List** field, select a **Prefix** list, or click **New Prefix List** to create a new one: **a.** From the **Prefix List** drop-down menu, select a prefix list, or create a new one. **b.** In the **Add Prefix** field, enter the IP prefixes, or click **Import** on the right to import prefixes. **c.** Click **Save**. **Note** The **Prefix List** option is not available if you select protocol **Both** (IPv4 and IPv6).

Step 4 From the **Actions** tab, select **IPv4**, **IPv6**, or **Both**, to designate which protocol the actions should apply to. Not all of the following options are available for all protocols.

Step 5 Click **Accept** or **Reject** for the IP traffic meeting the match conditions:

Match Condition	Description
Accept	Allow traffic from the selected protocol. Click the following menu buttons to open configuration fields:
	Export To—Select a VPN list, or create a new one.
	OMP Tag—Enter the OMP route tag, a number between 0-4294967295.
	Preference—Enter the preference number for the route, a number between 0-4294967295.
	Service— Enter the following information: **Type**—Select a service type. Options are: • **Firewall** • **Intrusion Detection Prevention** • **Intrusion Detection System** • **Net Service 1** • **Net Service 2** • **Net Service 3** • **Net Service 4** • **Net Service 5** **VPN**—Enter the number of the Service VPN. **TLOC IP**—Enter the IP address of the Service TLOC. **Color**—Select a Color type from the drop-down list. **Encapsulation**—Select **IPSEC** or **GRE** as the encapsulation type. **TLOC List**—Select a service TLOC list from the drop-down menu, or create a new one.

Match Condition	Description
	TLOC Action Select an action from the drop-down menu: • **Strict**—Direct matching traffic only to the intermediate destination. With this action, if the intermediate destination is down, no traffic reaches the final destination. If you do not configure a set tloc-action action in a centralized control policy, strict is the default behavior. • **Primary**—First direct matching traffic to the intermediate destination. If that driver is not reachable, then direct it to the final destination. With this action, if the intermediate destination is down, all traffic reaches the final destination. • **Backup**—First direct matching traffic to the final destination. If that driver is not reachable, then direct it to the intermediate destination. With this action, if the source is unable to reach the final destination directly, it is possible for all traffic to reach the final destination via the intermediate destination. • **Equal Cost Multi-path**—Equally direct matching control traffic between the intermediate destination and the ultimate destination. With this action, if the intermediate destination is down, all traffic reaches the ultimate destination.
	TLOC—Enter the following information: • **TLOC List**—Select a TLOC list, or create a new one. • **TLOC IP**—Enter the IP address of the designated TLOC. • **Color**—Select a color from the available options.
	Encapsulation—Select **IPSEC** or **GRE** as the encapsulation type.
Reject	Reject traffic for the selected conditions. **a.** Select a protocol from the **Protocol** dropdown: **IPv4**, **IPv6**, or **Both**. **b.** Click **Accept** or **Reject** for the match conditions. **c.** Optionally, repeat these steps with a different protocol.

Step 6 Click **Save Match and Actions**.

Create a Custom Control (TLOC)

Create a policy to apply to a TLOC. By default, the Match tab is selected, displaying match condition options.

Step 1 From the **Add Custom Control Policy** screen, click **TLOC**.

Step 2 Click **Sequence Rule**. Match and Actions options display.

Step 3 From the Match tab, select and configure match conditions for your route.

Match Condition	Description
Carrier	Select a carrier from the drop-down list.
Color List	Select a color list from the drop-down list, or create a new one.
Domain ID	Enter a domain ID number, between 1-4294967295.
Group ID	Enter a Group ID number, between 1-4294967295.
OMP Tag	Enter an OMP tag number, between 1-4294967295.
Originator	Enter the IP address of the originator of the TLOC.
Preference	Enter a preference number for the policy, between 1-4294967295.
Site List	Select a site list from the drop-down list, create a new one, or enter a site ID in the **Site ID** field, between 1-4294967295.
TLOC	Select a TLOC from the drop-down list, or create a new one or Select a TLOC from the drop-down list, or create a new one. Enter the following values: • **TLOC IP**—Enter the IP address of the TLOC. • **Color**—Select a color list from the available options. • **Encapsulation**—Select IPSEC or GRE as the encapsulation type.

Step 4 Click **Accept** or **Reject** to apply the following match conditions to an action.

Action Condition	Description
Accept	Allow traffic from the selected protocol. Click the following menu buttons to open configuration fields: • **OMP Tag**—Enter an OMP tag number, between 1-4294967295. • **Preference**—Enter a preference number for the policy, between 1-4294967295.
Reject	Reject traffic for the selected conditions.

Import Existing Topology

Step 1 In the Add Topology drop-down, select **Import Existing Topology** to open the matching popup

Step 2 Under **Policy Type**, click the topology type you want to import:

a) **Hub and Spoke**
b) **Mesh**
c) **Custom**

Step 3 Select a policy from the field list. Cisco vManage populates this field from the available topologies for the type you select.

Step 4 Click **Import**.

Step 5 Click **Save Control Policy** to save the Route policy.

Create a VPN Membership Policy

Step 1 In the Topology bar, click **VPN Membership**. Then:

Step 2 Click **Add VPN Membership Policy**. The Update VPN Membership Policy popup displays.

Step 3 Enter a name and description for the VPN membership policy.

Step 4 In the **Site List** field, select the site list.

Step 5 In the **VPN Lists** field, select the VPN list.

Step 6 Click **Add List** to add another VPN to the VPN membership.

Step 7 Click **Save**.

Step 8 Click **Next** to move to Configure Traffic Rules in the wizard.

Configure Centralized Policy Using Cisco vManage

To configure centralized policies, use the Cisco vManage policy configuration wizard. The wizard consists of four sequential screens that guide you through the process of creating and editing policy components:

- Create Groups of Interest—Create lists that group together related items and that you call in the match or action components of a policy.
- Configure Topology—Create the network structure to which the policy applies.
- Configure Traffic Rules—Create the match and action conditions of a policy.
- Apply Policies to Sites and VPNs—Associate policy with sites and VPNs in the overlay network.

In the first three policy configuration wizard screens, you are creating policy components or blocks. In the last screen, you are applying policy blocks to sites and VPNs in the overlay network.

For a centralized policy to take effect, you must activate the policy.

Step 1: Start the Policy Configuration Wizard

1. In the Cisco vManage NMS, select the **Configure** > **Policies** screen.
2. Select the **Centralized Policy** tab.
3. Click **Add Policy**.

The policy configuration wizard appears, and the **Create Applications or Groups of Interest** screen is displayed.

Step 2: Configure Groups of Interest

In **Create Groups of Interest**, create lists of groups to use in a centralized policy:

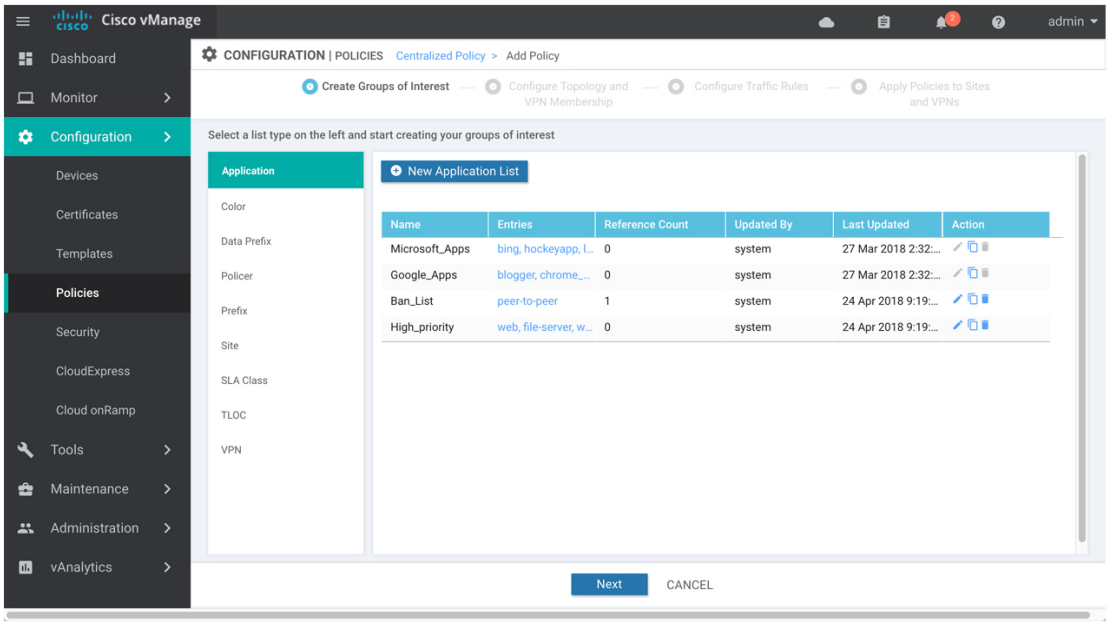

368879

1. Create new lists, as described in the following table:

Table 1:

List Type	Procedure
Color	**a.** In the left bar, click **Color**. **b.** Click **New Color List**. **c.** Enter a name for the list. **d.** From the Select Color drop-down, select the desired colors. **e.** Click **Add**.
Prefix	**a.** In the left bar, click **Prefix**. **b.** Click **New Prefix List**. **c.** Enter a name for the list. **d.** In the Add Prefix field, enter one or more data prefixes separated by commas. **e.** Click **Add**.

List Type	Procedure
Site	**a.** In the left bar, click **Site**. **b.** Click **New Site List**. **c.** Enter a name for the list. **d.** In the Add Site field, enter one or more site IDs separated by commas. **e.** Click **Add**.
TLOC	**a.** In the left bar, click **TLOC**. **b.** Click **New TLOC List**. The TLOC List popup displays. **c.** Enter a name for the list. **d.** In the TLOC IP field, enter the system IP address for the TLOC. **e.** In the Color field, select the TLOC's color. **f.** In the Encap field, select the encapsulation type. **g.** In the Preference field, optionally select a preference to associate with the TLOC. **h.** Click **Add TLOC** to add another TLOC to the list. **i.** Click **Save**.
VPN	**a.** In the left bar, click **VPN**. **b.** Click **New VPN List**. **c.** Enter a name for the list. **d.** In the Add VPN field, enter one or more VPN IDs separated by commas. **e.** Click **Add**.

2. Click **Next** to move to **Configure Topology and VPN Membership** in the wizard.

Step 3: Configure Topology and VPN Membership

When you first open the **Configure Topology and VPN Membership** screen, the **Topology** tab is selected by default:

To configure topology and VPN membership:

In the **Topology** tab, create a network topology:

Custom Control (Route & TLOC) - Centralized route control policy (for matching OMP routes)

1. In the Add Topology drop-down, select **Custom Control (Route & TLOC)**.
2. Enter a name for the control policy.

3. Enter a description for the policy.
4. In the left pane, click **Add Sequence Type**. The Add Control Policy popup displays.
5. Select **Route**. A policy component containing the text string Route is added in the left pane.
6. Double-click the **Route** text string, and enter a name for the policy component.
7. In the right pane, click **Add Sequence Rule**. The Match/Actions box opens, and Match is selected by default.
8. From the boxes under the Match box, select the desired policy match type. Then select or enter the value for that match condition. Configure additional match conditions for the sequence rule, as desired. For an explanation of the match conditions, see the OMP Route Match Attributes section in the Configuring Centralized Control Policy topic for your software release.
9. Click **Actions**. The Reject radio button is selected by default. To configure actions to perform on accepted packets, click the **Accept** radio button. Then select the action or enter a value for the action. For an explanation of the actions, see the *Action Parameters* section in the *Configuring Centralized Control Policy* topic for your software release.
10. Click **Save Match and Actions**.
11. Click **Add Sequence Rules** to configure more sequence rules, as desired. Drag and drop to re-order them.
12. Click **Add Sequence Typ**e to configure more sequences, as desired. Drag and drop to re-order them.
13. Click **Save Control Policy**.

Custom Control (Route & TLOC) - Centralized TLOC control policy (for matching TLOC routes)

1. In the Add Topology drop-down, select **Custom Control (Route & TLOC)**.
2. Enter a name for the control policy.
3. Enter a description for the policy.
4. In the left pane, click **Add Sequence Type**. The Add Control Policy popup displays.
5. Select **TLOC**. A policy component containing the text string TLOC is added in the left pane.
6. Double-click the TLOC text string, and enter a name for the policy component.
7. In the right pane, click **Add Sequence Rule**. The Match/Actions box opens, and Match is selected by default.
8. From the boxes under the Match box, select the desired policy match type. Then select or enter the value for that match condition. Configure additional match conditions for the sequence rule, as desired. For an explanation of the match conditions, see the OMP TLOC Match Attributes section in the *Configuring Centralized Control Policy* topic for your software release.
9. Click **Actions**. The Reject radio button is selected by default. To configure actions to perform on accepted packets, click the **Accept** radio button. Then select the action or enter a value for the action. For an explanation of the actions, see the *Action Parameters* section in the *Configuring Centralized Control Policy* topic for your software release.
10. Click **Save Match and Actions**.

11. Click **Add Sequence Rules** to configure more sequence rules, as desired. Drag and drop to re-order them.
12. Click **Add Sequence Type** to configure more sequences, as desired. Drag and drop to re-order them.
13. Click **Save Control Policy**.

To use an existing topology:

1. In the **Add Topology** drop-down, click **Import Existing Topology**. The Import Existing Topology popup appears.
2. Select the type of topology.
3. In the **Policy** drop-down, choose the name of the topology.
4. Click **Import**.

Click **Next** to move to **Configure Traffic Rules** in the wizard.

Click **Next** to move to **Apply Policies to Sites and VPNs** in the wizard.

Step 4: Apply Policies to Sites and VPNs

In **Apply Policies to Sites and VPNs** screen, apply a policy to sites and VPNs:

1. In the **Policy Name** field, enter a name for the policy. This field is mandatory and can contain only uppercase and lowercase letters, the digits 0 through 9, hyphens (–), and underscores (_). It cannot contain spaces or any other characters.
2. In the **Policy Description** field, enter a description of the policy. It can contain up to 2048 characters. This field is mandatory, and it can contain any characters and spaces.
3. From the **Topology** bar, choose the type of policy block. The table then lists policies that you have created for that type of policy block.
4. Associate the policy with VPNs and sites. The choice of VPNs and sites depends on the type of policy block:
   a. For a Topology policy block, click **Add New Site List** and **VPN List** or **Add New Site**. Some topology blocks might have no **Add** buttons. Choose one or more site lists, and choose one or more VPN lists. Click **Add**.
   b. For an Application-Aware Routing policy block, click **Add New Site List** and **VPN list**. Choose one or more site lists, and choose one or more VPN lists. Click **Add**.
   c. For a Traffic Data policy block, click **Add New Site List** and **VPN List**. Choose the direction for applying the policy (From Tunnel, From Service, or All), choose one or more site lists, and choose one or more VPN lists. Click **Add**.
   d. For a cflowd policy block, click **Add New Site List**. Choose one or more site lists, Click **Add**.
5. Click **Preview** to view the configured policy. The policy appears in CLI format.
6. Click **Save** Policy. The **Configuration** > **Policies** screen appears, and the policies table includes the newly created policy.

Step 5: Activate a Centralized Policy

Activating a centralized policy sends that policy to all connected Cisco vSmart controllers. To activate a centralized policy:

1. In the Cisco vManage NMS, select the **Configure** > **Policies** screen. When you first open this screen, the **Centralized Policy** tab is selected by default.
2. Choose a policy.
3. Click the **More Actions** icon to the right of the row, and click **Activate**. The Activate Policy popup appears. It lists the IP addresses of the reachable Cisco vSmart Controllers to which the policy must be applied.
4. Click **Activate**.

Structural Components for Centralized Control Policy

Following are the structural components required to configure centralized control policy. Each one is explained in more detail in the sections below.

```
policy  lists color-list list-name color color prefix-list list-name ip-prefix prefix
site-list list-name site-id site-id tloc-list list-name tloc address color color
encap encapsulation [preference value] vpn-list list-name vpn vpn-id
control-policy
policy-name
sequence
number
match
match-parameters
action reject accept export-to vpn accept set parameter
default-action (accept | reject)apply-policy site-list list-name control-policy policy-name
 (in | out)
```

Lists

Centralized control policy uses the following types of lists to group related items. In the CLI, you configure lists under the **policy lists** command hierarchy on Cisco vSmart Controllers.

List Type	Description	vManage Configuration/ CLI Configuration Command
Colors	List of one or more TLOC colors.*color* can be **3g**, **biz-internet**, **blue**, **bronze**, **custom1** through **custom3**, **default**, **gold**, **green**, **lte**, **metro-ethernet**, **mpls**, **private1** through **private6**, **public-internet**, **red**, and **silver**. To configure multiple colors in a single list, include multiple **color** options, specifying one color in each option.	**Configuration > Policies > Centralized Policy > Add Policy > Create Groups of Interest > Color** **Configuration > Policies > Custom Options > Centralized Policy > Lists > Color** **color-list** *list-name* **color** *color*

List Type	Description	vManage Configuration/ CLI Configuration Command
Prefixes	List of one or more IP prefixes. Specify the IP prefixes as follows: • *prefix*/*length*—Exactly match a single prefix–length pair. • **0.0.0.0/0**—Match any prefix–length pair. • **0.0.0.0/0 le** *length*—Match any IP prefix whose length is less than or equal to *length*. For example, **ip-prefix 0.0.0.0/0 le 16** matches all IP prefixes with lengths from /1 through /16. • **0.0.0.0/0 ge** *length*—Match any IP prefix whose length is greater than or equal to *length*. For example, **ip-prefix 0.0.0.0 ge 25** matches all IP prefixes with lengths from /25 through /32. • **0.0.0.0/0 ge** *length1* **le** *length2*, or **0.0.0.0 le** *length2* **ge** *length1*—Match any IP prefix whose length is greater than or equal to *length1* and less than or equal to *length2*. For example, **ip-prefix 0.0.0.0/0 ge 20 le 24** matches all /20, /21, /22, /23, and /24 prefixes. Also, **ip-prefix 0.0.0.0/0 le 24 ge 20** matches the same prefixes. If *length1* and *length2* are the same, a single IP prefix length is matched. For example, **ip-prefix 0.0.0.0/0 ge 24 le 24** matches only /24 prefixes. To configure multiple prefixes in a single list, include multiple **ip-prefix** options, specifying one prefix in each option.	**Configuration > Policies > Centralized Policy > Add Policy > Create Groups of Interest > Prefix** **Configuration > Policies > Custom Options > Centralized Policy > Lists > Prefix** **prefix-list** *list-name* **ip-prefix** *prefix*/*length*
Sites	List of one of more site identifiers in the overlay network. You can specify a single site identifier (such as **site-id 1**) or a range of site identifiers (such as **site-id 1-10**). To configure multiple sites in a single list, include multiple **site-id** options, specifying one site number in each option.	**Configuration > Policies > Centralized Policy > Add Policy > Create Groups of Interest > Site** **Configuration > Policies > Custom Options > Centralized Policy > Lists > Site** **site-list** *list-name* **site-id** *site-id*
TLOCs	List of one or more TLOCs in the overlay network. For each TLOC, specify its address, color, and encapsulation. *address* is the system IP address. *color* can be one of **3g**, **biz-internet**, **blue**, **bronze**, **custom1**, **custom2**, **custom3**, **default**, **gold**, **green**, **lte**, **metro-ethernet**, **mpls**, **private1** through **private6**, **public-internet**, **red**, and **silver**. *encapsulation* can be **gre** or **ipsec**. Optionally, set a preference value (from 0 to $2^{32} - 1$) to associate with the TLOC address. When you apply a TLOC list in an **action accept** condition, when multiple TLOCs are available and satisfy the match conditions, the TLOC with the lowest preference value is used. If two or more of TLOCs have the lowest preference value, traffic is sent among them in an ECMP fashion.	**Configuration > Policies > Centralized Policy > Add Policy > Create Groups of Interest > TLOC** **Configuration > Policies > Custom Options > Centralized Policy > Lists > TLOC** **tloc-list** *list-name* **tloc** *ip-address* **color** *color* **encap** (**gre** | **ipsec**) [**preference** *number*]

List Type	Description	vManage Configuration/ CLI Configuration Command
VPNs	List of one or more VPNs in the overlay network. For data policy, you can configure any VPNs except for VPN 0 and VPN 512. To configure multiple VPNs in a single list, include multiple **vpn** options, specifying one VPN number in each option. You can specify a single VPN identifier (such as **vpn 1**) or a range of VPN identifiers (such as **vpn 1-10**).	**Configuration > Policies > Centralized Policy > Add Policy > Create Groups of Interest > VPN** **Configuration > Policies > Custom Options > Centralized Policy > Lists > VPN** **vpn-list** *list-name* **vpn** *vpn-id*

Sequences

A centralized control policy contains sequences of match–action pairs. The sequences are numbered to set the order in which a route or TLOC is analyzed by the match–action pairs in the policy.

In the Cisco vManage NMS, you configure sequences from:

- **Configuration > Policies > Centralized Policy > Add Policy > Configure Traffic Rules > (Application-Aware Routing | Traffic Data | Cflowd) > Sequence Type**
- **Configuration > Policies > Custom Options > Centralized Policy > Traffic Policy > (Application-Aware Routing | Traffic Data | Cflowd) > Sequence Type**

In the CLI, you configure sequences with the **policy control-policy sequence** command.

Each sequence in a centralized control policy can contain one match condition (either for a route or for a TLOC) and one action condition.

Match Parameters

Centralized control policy can match OMP route or TLOC route attributes.

In the Cisco vManage NMS, you configure match parameters from:

- **Configuration > Policies > Centralized Policy > Add Policy > Configure Topology and VPN Membership > Add Topology > Custom Control (Route & TLOC) > Sequence Type > (Route | TLOC) > Sequence Rule > Match**
- **Configuration > Policies > Custom Options > Centralized Policy > Topology > Add Topology > Custom Control (Route & TLOC) > Sequence Type > (Route | TLOC) > Sequence Rule > Match**

In the CLI, you configure the OMP route attributes to match with the **policy control-policy sequence match route** command, and you configure the TLOC attributes to match with the **policy control-policy sequence match tloc** command.

Each sequence in a policy can contain one **match** section—either **match route** or **match tloc**.

OMP Route Match Attributes

For OMP routes (vRoutes), you can match these attributes:

Description	vManage Configuration/ CLI Configuration Command	Value or Range
Individual color.	Not available in the Cisco vManage NMS. **color** *color*	**3g**, **biz-internet**, **blue**, **bronze**, **custom1** through **custom3**,**default**, **gold**, **green**, **lte**, **metro-ethernet**, **mpls**, **private1** through **private6**, **public-internet**, **red**, and **silver**
One or more colors.	Match Color List **color-list** *list-name*	Name of a color or a **policy lists color-list** list.
Tag value associated with the route or prefix in the routing database on the device.	Match OMP Tag **omp-tag** *number*	0 through 4294967295
Protocol from which the route was learned.	Match Origin **origin** *protocol*	**bgp-external**, **bgp-internal**, **connected**, **ospf-external1**, **ospf-external2**, **ospf-inter-area**, **ospf-intra-area**, **static**
IP address from which the route was learned.	Match Originator **originator** *ip-address*	IP address
How preferred a prefix is. This is the preference value that the route or prefix has in the local site, that is, in the routing database on the device. A higher preference value is more preferred.	Match Preference **preference** *number*	0 through 255
One or more prefixes.	Match Prefix List **prefix-list** *list-name*	Name of a prefix list or a **policy lists prefix-list** list.
Individual site identifier.	Not available in Cisco vManage. **site-id** *site-id*	0 through 4294967295
One or more overlay network site identifiers.	Match Site **site-list** *list-name*	Name of a site or a **policy lists site-list** list.
Individual TLOC address.	Match TLOC **tloc** *ip-address*	IP address
One or more TLOC addresses.	Match TLOC **tloc-list** *list-name*	Name of a TLOC or a **policy lists tloc-list** list.
Individual VPN identifier.	Match VPN **vpn** *vpn-id*	0 through 65535

Description	vManage Configuration/ CLI Configuration Command	Value or Range
One or more VPN identifiers.	Match VPN **vpn-list** *list-name*	Name of a VPN or a **policy lists vpn-list** list.

TLOC Route Match Attributes

For TLOC routes, you can match these attributes:

Description	vManage Configuration/ CLI Configuration Command	Value or Range
Carrier for the control traffic.	Match Carrier **carrier** *carrier-name*	**default**, **carrier1** through **carrier8**
Individual color.	Not available in the Cisco vManage NMS. **color** *color*	**3g**, **biz-internet**, **blue**, **bronze**, **custom1** through **custom3**,**default**, **gold**, **green**, **lte**, **metro-ethernet**, **mpls**, **private1** through **private6**, **public-internet**, **red**, and **silver**
One or more colors.	Match Color List **color-list** *list-name*	See the colors above.
Domain identifier associated with a TLOC.	Match Domain ID **domain-id** *domain-id*	0 through 4294967295
Tag value associated with the TLOC route in the route table on the device.	Match OMP Tag **omp-tag** *number*	0 through 4294967295
IP address from which the route was learned.	Match Originator **originator** *ip-address*	IP address
How preferred a TLOC route is. This is the preference value that the TLOC route has in the local site, that is, in the route table on the Cisco IOS XE SD-WAN. A higher preference value is more preferred.	Match Preference **preference** *number*	0 through 255
Individual site identifier.	Match Site **site-id** *site-id*	0 through 4294967295
One or more overlay network site identifiers.	Match Site **site-list** *list-name*	Name of a **policy lists site-list** list.

Description	vManage Configuration/ CLI Configuration Command	Value or Range
Individual TLOC address.	Match TLOC **tloc** *address*	IP address
One or more TLOC addresses.	Match TLOC **tloc-list** *list-name*	Name of a **policy lists tloc-list** list.

Action Parameters

For each match condition, you configure a corresponding action to take if the route or TLOC matches.

In the Cisco vManage NMS, you configure match parameters from:

- **Configuration > Policies > Centralized Policy > Add Policy > Configure Topology and VPN Membership > Add Topology > Custom Control (Route & TLOC) > Sequence Type > (Route | TLOC) > Sequence Rule > Action**
- **Configuration > Policies > Custom Options > Centralized Policy > Topology > Add Topology > Custom Control (Route & TLOC) > Sequence Type > (Route | TLOC) > Sequence Rule > Action**

In the CLI, you configure actions with the **policy control-policy action** command.

Each sequence in a centralized control policy can contain one action condition.

In the action, you first specify whether to accept or reject a matching route or TLOC:

Description	vManage Configuration/ CLI Configuration Command	Value or Range
Accept the route. An accepted route is eligible to be modified by the additional parameters configured in the **action** portion of the policy configuration.	Click **Accept**. **accept**	—
Discard the packet.	Click **Reject**. **reject**	—

Then, for a route or TLOC that is accepted, you can configure the following actions:

Description	vManage Configuration/ CLI Configuration Command	Value or Range
Export the route the the specified VPN or list of VPNs (for a **match route** match condition only).	Click **Accept**, then action **Export To**. **export-to** (**vpn** *vpn-id* | **vpn-list** *vpn-list*)	0 through 65535 or list name.

Description	vManage Configuration/ CLI Configuration Command	Value or Range
Change the tag string in the route, prefix, or TLOC.	Click **Accept**, then action **OMP Tag**. **set omp-tag** *number*	0 through 4294967295
Change the preference value in the route, prefix, or TLOC to the specified value. A higher preference value is more preferred.	Click **Accept**, then action **Preference**. **set preference** *number*	0 through 255
Specify a service to redirect traffic to before delivering the traffic to its destination. The TLOC address or list of TLOCs identifies the TLOCs to which the traffic should be redirected to reach the service. In the case of multiple TLOCs, the traffic is load-balanced among them. The VPN identifier is where the service is located. Configure the services themselves on the Cisco IOS XE SD-WAN devices that are collocated with the service devices, using the **vpn service** configuration command.	Click **Accept**, then action Service. **set service** *service-name* (**tloc** *ip-address* | **tloc-list** *list-name*) [**vpn** *vpn-id*]	Standard services: **FW**, **IDS**, **IDP** Custom services: **netsvc1**, **netsvc2**, **netsvc3**, **netsvc4** TLOC list configured with a **policy lists tloc-list** command.
Change the TLOC address, color, and encapsulation to the specified address and color.	Click **Accept**, then action TLOC. **set tloc** *ip-address* **color** *color* [**encap** *encapsulation*]	IP address, TLOC color, and encapsulation, Color can be one of **3g**, **biz-internet**, **blue**, **bronze**, **custom1** through **custom3**,**default**, **gold**, **green**, **lte**, **metro-ethernet**, **mpls**, **private1** through **private6**, **public-internet**, **red**, and **silver**. Encapsuation can be either **gre** or **ipsec**.

Description	vManage Configuration/ CLI Configuration Command	Value or Range
Direct matching routes or TLOCs using the mechanism specified by *action*, and enable end-to-end tracking of whether the ultimate destination is reachable. Setting a TLOC action is useful when traffic is first directed, via policy, to an intermediate destination, which then forwards the traffic to its ultimate destination. For example, for traffic from vEdge-A destined for vEdge-D, a policy might direct traffic from vEdge-A first to vEdge-B (the intermediate destination), and vEdge-B then sends it to the final destination, vEdge-D. Setting the TLOC action option enables the Cisco vSmart Controller to perform end-to-end tracking of the path to the ultimate destination device. In our example, matching traffic goes from vEdge-A to vEdge-B and then, in a single hop, goes to vEdge-D. If the tunnel between vEdge-B and vEdge-D goes down, the Cisco vSmart Controller relays this information to vEdge-A, and vEdge-A removes its route to vEdge-D from its local route table. End-to-end tracking works here only because traffic goes from vEdge-B to vEdge-D in a single hop, via a single tunnel. If the traffic from vEdge-A went first to vEdge-B, then to vEdge-C, and finally to vEdge-D, the vSmart controller is unable to perform end-to-end tracking and is thus unable to keep vEdge-A informed about whether full path between it and vEdge-D is up.	Click **Accept**, then action **TLOC Action**. **set tloc-action** *action*	**ecmp**—Equally direct matching control traffic between the intermediate destination and the ultimate destination. In our example, traffic would be sent to vEdge-B (which would then send it to vEdge-D) and directly to vEdge-D. With this action, if the intermediate destination is down, all traffic reaches the ultimate destination. **primary**—First direct matching traffic to the intermediate destination. If that device is not reachable, then direct it to the final destination. In our example, traffic would first be sent to vEdge-B. If this device is down, it is sent directly to vEdge-D. With this action, if the intermediate destination is down, all traffic reaches the final destination. **backup**—First direct matching traffic to the final destination. If that device is not reachable, then direct it to the intermediate destination. In our example, traffic would first be sent directly to vEdge-D. If the vEdge-A is not able to reach vEdge-D, traffic is sent to vEdge-B, which might have an operational path to reach vEdge-D. With this action, if the source is unable to reach the final destination directly, it is possible for all traffic to reach the final destination via the intermediate destination. **strict**—Direct matching traffic only to the intermediate destination. In our example, traffic is sent only to vEdge-B, regardless of whether it is reachable. With this action, if the intermediate destination is down, no traffic reaches the final destination. If you do not configure a set tloc-action action in a centralized control policy, strict is the default behavior.
Change the TLOC address and color to those in the specified TLOC list.	Click **Accept**, then action **TLOC**. **set tloc-list** *list-name*	Name of a **policy lists tloc-list** list.

Default Action

If a route or TLOC being evaluated does not match any of the match conditions in a centralized control policy, a default action is applied to it. By default, the route or TLOC is rejected.

In the Cisco vManage NMS, you modify the default action from **Configuration > Policies > Centralized Policy > Add Policy > Configure Topology and VPN Membership > Add Topology > Custom Control (Route and TLOC) > Sequence Type > (Route | TLOC) > Sequence Rule > Default Action**.

In the CLI, you modify the default action with the **control policy default-action accept** command.

Apply Centralized Control Policy

For a centralized control policy to take effect, you apply it to a list of sites in the overlay network.

To apply a centralized policy in the Cisco vManage NMS:

1. In the Cisco vManage NMS, select the **Configure** > **Policies** screen.
2. Select a policy from the policy table.
3. Click the **More Actions** icon to the right of the row, and click **Activate**. The Activate Policy popup opens. It lists the IP addresses of the reachable Cisco vSmart Controllers to which the policy is to be applied.
4. Click **Activate**.

To apply a centralized policy in the CLI:

```
vSmart(config)# apply-policy
site-list
list-name
control-policy
policy-name (in | out)
```

You apply centralized control policy directionally:

- Inbound direction (**in**)—The policy analyzes routes and TLOCs being received from the sites in the site list before placing the routes and TLOCs into the route table on the Cisco vSmart Controller, so the specified policy actions affect the OMP routes stored in the route table.
- Outbound direction (**out**)—The policy analyzes routes and TLOCs in the Cisco vSmart Controller's route table after they are exported from the route table.

For all **control-policy** policies that you apply with **apply-policy** commands, the site IDs across all the site lists must be unique. That is, the site lists must not contain overlapping site IDs. An example of overlapping site IDs are those in the two site lists **site-list 1 site-id 1-100** and **site-list 2 site-id 70-130**. Here, sites 70 through 100 are in both lists. If you were to apply these two site lists to two different **control-policy** policies, the attempt to commit the configuration on the Cisco vSmart Controller would fail.

The same type of restriction also applies to the following types of policies:

- Application-aware routing policy (**app-route-policy**)
- Centralized data policy (**data-policy**)
- Centralized data policy used for cflowd flow monitoring (**data-policy** hat includes a **cflowd** action and **apply-policy** that includes a **cflowd-template** command)

You can, however, have overlapping site IDs for site lists that you apply for different types of policy. For example, the sites lists for **control-policy** and **data-policy** policies can have overlapping site IDs. So for the two example site lists above, **site-list 1 site-id 1-100** and **site-list 2 site-id 70-130**, you could apply one to a control policy and the other to a data policy.

Configure Centralized Policy Using CLI

To configure a centralized control policy using the CLI:

1. Create a list of overlay network sites to which the centralized control policy is to be applied (in the **apply-policy** command):

   ```
   vSmart(config)# policy
   vSmart(config-policy)# lists site-list list-name
   vSmart(config-lists-list-name)# site-id site-id
   ```

   The list can contain as many site IDs as necessary. Include one **site-id** command for each site ID. For contiguous site IDs, you can specify a range of numbers separated with a dash (–). Create additional site lists, as needed.

2. Create lists of IP prefixes, TLOCs, and VPNs as needed:

   ```
   vSmart(config)# policy lists
   vSmart(config-lists)# prefix-list list-name
   vSmart(config-lists-list-name)# ip-prefix prefix/length
   vSmart(config)# policy lists
   vSmart(config-lists)# tloc-list list-name
   vSmart(config-lists-list-name)# tloc address
                                   color
   color
                                   encap encapsulation
                                   [preference value]
   vSmart(config)# policy lists
   vSmart(config-lists)# vpn-list list-name
   vSmart(config-lists-list-name)# vpn vpn-id
   ```

   ```
   vsmart(config)# policy lists data-ipv6-prefix-list dest_ip_prefix_list
   vsmart(config-data-ipv6-prefix-list-dest_ip_prefix_list)# ipv6-prefix 2001:DB8::/32
   vsmart(config-data-ipv6-prefix-list-dest_ip_prefix_list)# commit
   Commit complete.
   ```

   ```
   vsmart(config)# policy data-policy data_policy_1 vpn-list vpn_1
   vsmart (config-sequence-100)# match destination-data-ipv6-prefix-list dest_ip_prefix_list
   vsmart (config-match)# commit
   vsmart(config-match)# exit
   vsmart(config-sequence-100)# match source-data-ipv6-prefix-list dest_ip_prefix_list
   vm9(config-match)# commit
   Commit complete.
   vm9(config-match)# end
   ```

   ```
   vsmart(config)# policy
   vsmart(config-policy)#  data-policy data_policy_1
   vsmart(config-data-policy-data_policy_1)# vpn-list vpn_1
   vsmart(config-vpn-list-vpn_1)# sequence 101
   vsmart(config-sequence-101)# match source-ipv6 2001:DB8::/32
   vsmart(config-match)# exit
   vsmart(config-sequence-101)# match destination-ipv6 2001:DB8::/32
   vsmart(config-match)#
   ```

1. Create a control policy instance:

```
vSmart(config)# policy control-policy policy-name
vSmart(config-control-policy-policy-name)#
```

2. Create a series of match–action pair sequences:

```
vSmart(config-control-policy-policy-name)# sequence
number
vSmart(config-sequence-number)#
```

The match–action pairs are evaluated in order, by sequence number, starting with the lowest numbered pair and ending when the route matches the conditions in one of the pairs. Or if no match occurs, the default action is taken (either rejecting the route or accepting it as is).

3. Define match parameters for routes and for TLOCs:

```
vSmart(config-sequence-number)# match route route-parameter
vSmart(config-sequence-number)# match tloc tloc-parameter
```

4. Define actions to take when a match occurs:

```
vSmart(config-sequence-number)# action reject
vSmart(config-sequence-number)# action accept export-to (vpn
vpn-id | vpn-list list-name)
vSmart(config-sequence-number)# action accept set omp-tag
number
```

```
vSmart(config-sequence-number)# action accept set
preference value
```

```
vSmart(config-sequence-number)# action accept set
service service-name
(tloc ip-address |
tloc-list list-name)
[vpn vpn-id]
```

```
vSmart(config-sequence-number)# action accept set tloc
ip-address
color color
[encap encapsulation]
vSmart(config-sequence-number)# action accept set tloc-action
action
```

```
vSmart(config-sequence-number)# action accept set tloc-list list-name
```

5. Create additional numbered sequences of match–action pairs within the control policy, as needed.

6. If a route does not match any of the conditions in one of the sequences, it is rejected by default. If you want nonmatching routes to be accepted, configure the default action for the policy:

```
vSmart(config-policy-name)# default-action accept
```

7. Apply the policy to one or more sites in the Cisco SD-WAN overlay network:

```
vSmart(config)# apply-policy site-list
list-name
control-policy
policy-name (in | out)
```

8. If the action you are configuring is a service, configure the required services on the Cisco IOS XE SD-WAN devices so that the Cisco vSmart Controller knows how to reach the services:

```
vsmart(config)# policy data-policy data_policy_1 vpn-list vpn_1 sequence 100
vsmart(config-sequence-100)# action accept set next-hop-ipv6 2001:DB8::/32
vsmart(config-set)#
```

Specify the VPN is which the service is located and one to four IP addresses to reach the service device or devices. If multiple devices provide the same service, the device load-balances the traffic among them. Note that the Cisco IOS XE SD-WAN device keeps track of the services, advertising them to the Cisco vSmart Controller only if the address (or one of the addresses) can be resolved locally, that is, at the device's local site, and not learned through OMP. If a previously advertised service becomes unavailable, the Cisco IOS XE SD-WAN device withdraws the service advertisement.

Centralized Control Policy Configuration Examples

This topic provides some straightforward examples of configuring centralized control policy to help you understand the configuration procedure and get an idea of how to use policy to influence traffic flow across the Cisco IOS XE SD-WAN overlay network domain.

Traffic Engineering

This example of traffic engineering forces all traffic to come to a Cisco IOS XE SD-WAN device using a device hub instead of directly.

One common way to design a domain in a Cisco IOS XE SD-WAN overlay network is to route all traffic destined for branches through a hub router, which is typically located in a data center, rather than sending the traffic directly from one Cisco IOS XE SD-WAN device to another. You can think of this as a hub-and-spoke design, where one device is acting as a hub and the devices are the spokes. With such a design, traffic between local branches travels over the IPsec connections that are established between the spoke routers and the hub routers when the devices are booted up. Using established connections means that the devices do not need to expend time and CPU cycles to establish IPsec connections with each other. If you were to imagine that this were a large network with many devices, having a full mesh of connections between each pair of routers would require a large amount of CPU from the routers. Another attribute of this design is that, from an administrative point of view, it can be simpler to institute coordinated traffic flow policies on the hub routers, both because there are fewer of them in the overlay network and because they are located in a centralized data center.

One way to direct all the device spoke router traffic to a Cisco hub router is to create a policy that changes the TLOC associated with the routes in the local network. Let’s consider the topology in the figure here:

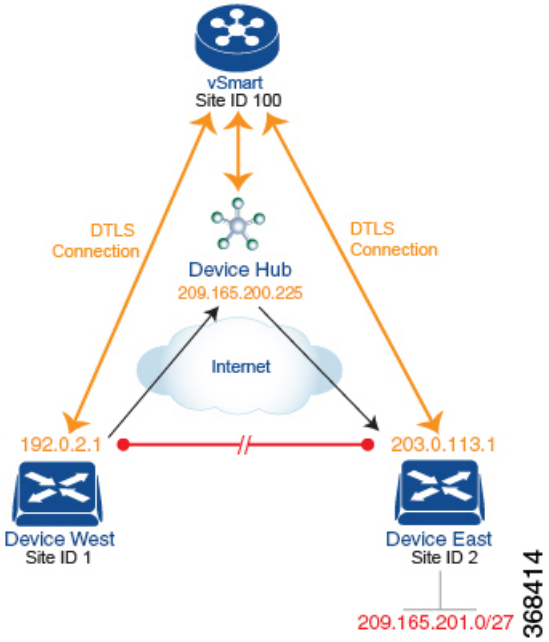

This topology has two devices in different branches:

- The Device West in site ID 1. The TLOC for this device is defined by its IP address (192.0.2.1), a color (gold), and an encapsulation (here, IPsec). We write the full TLOC address as {192.0.2.1, gold, ipsec}. The color is simply a way to identify a flow of traffic and to separate it from other flows.
- The Device East in site ID 2 has a TLOC address of {203.0.113.1, gold, ipsec}.

The devices West and East learn each other's TLOC addresses from the OMP routes distributed to them by the Cisco vSmart Controller. In this example, the Device East advertises the prefix 209.165.201.0/27 as being reachable at TLOC {203.0.113.1, gold, }. In the absence of any policy, the Device West could route traffic destined for 209.165.201.0/27 to TLOC {203.0.113.1, gold, ipsec}, which means that the Device West would be sending traffic directly to the Device East.

However, our design requires that all traffic from West to East be routed through the hub router, whose TLOC address is {209.165.200.225, gold, ipsec}, before going to the Device East. To effect this traffic flow, you define a policy that changes the route's TLOC. So, for the prefix 1209.165.201.0/27, you create a policy that changes the TLOC associated with the prefix 209.165.201.0/27 from {203.0.113.1, gold, ipsec}, which is the TLOC address of the Device East, to {209.165.200.225, gold, ipsec}, which is the TLOC address of the hub router. The result is that the OMP route for the prefix 209.165.201.0/27 that the Cisco vSmart Controller advertises to the Device West that contains the TLOC address of the hub router instead of the TLOC address of the Device East. From a traffic flow point of view, the Device West then sends all traffic destined for 209.165.201.0/27 to the hub router.

The device also learns the TLOC addresses of the West and East devices from the OMP routes advertised by the Cisco vSmart Controller. Because, devices must use these two TLOC addresses, no policy is required to control how the hub directs traffic to the devices.

Here is a policy configuration on the Cisco vSmart Controller that directs the Device West (and any other devices in the network domain) to send traffic destined to prefix 209.165.201.0/27 to TLOC 209.165.200.225, gold, which is the device:

```
policy
  lists
    prefix-list east-prefixes
      ip-prefix 209.165.201.0/27
    site-list west-sites
      site-id 1
  control-policy change-tloc
    sequence 10
      match route
        prefix-list east-prefixes
        site-id 2
      action accept
        set tloc 209.165.200.225 color gold encap ipsec
apply-policy
  site west-sites control-policy change-tloc out
```

A rough English translation of this policy is:

```
Create a list named “east-prefixes” that contains the IP prefix “209.165.201.0/27”
  Create a list named “west-sites” that contains the site-id “1”
  Define a control policy named “change-tloc”
    Create a policy sequence element that:
      Matches a prefix from list “east-prefixes”, that is, matches “209.165.201.0/27”
      AND matches a route from site-id “2”
    If a match occurs:
      Accept the route
      AND change the route’s TLOC to “209.165.200.225” with a color of "gold" and an
encapsulation of "ipsec"
  Apply the control policy “change-tloc” to OMP routes sent by the vSmart
    controller to “west-sites”, that is, to site ID 1
```

This control policy is configured on the Cisco vSmart Controller as an outbound policy, as indicated by the **out** option in the **apply-policy site** command. This option means the Cisco vSmart Controller applies the TLOC change to the OMP route after it distributes the route from its route table. The OMP route for prefix 209.165.201.0/27 that the Cisco vSmart Controller distributes to the Device West associates 209.165.201.0/27 with TLOC 209.165.200.225, gold. This is the OMP route that the Device West installs it in its route table. The end results are that when the Device West sends traffic to 209.165.201.0/27, the traffic is directed to the hub; and the Device West does not establish a DTLS tunnel directly with the Device East.

If the West side of the network had many sites instead of just one and each site had its own device, it would be straightforward to apply this same policy to all the sites. To do this, you simply add the site IDs of all the sites in the **site-list west-sites** list. This is the only change you need to make in the policy to have all the West-side sites send traffic bound for the prefix 209.165.201.0/27 through the device. For example:

```
policy
  lists
    prefix-list east-prefixes
      ip-prefix 209.165.201.0/27
    site-list west-sites
      site-id 1
      site-id 11
      site-id 12
      site-id 13
  control-policy change-tloc
    sequence 10
      match route
        prefix-list east-prefixes
        site-id 2
      action accept
        set tloc 209.165.200.225 color gold encap ipsec
apply-policy
  site west-sites control-policy change-tloc out
```

Creating Arbitrary Topologies

To provide redundancy in the hub-and-spoke-style topology discussed in the previous example, you can add a second Cisco hub to create a dual-homed hub site. The following figure shows that site ID 10 now has two Device hubs. We still want all inter-branch traffic to be routed through a device hub. However, because we now have dual-homed hubs, we want to share the data traffic between the two hub routers.

- Device Hub West, with TLOC 209.165.200.225, gold. We want all data traffic from branches on the West side of the overlay network to pass through and be processed by this device.
- Device Hub East, with TLOC 198.51.100.1, gold. Similarly, we all East-side data traffic to pass through the Device Hub East.

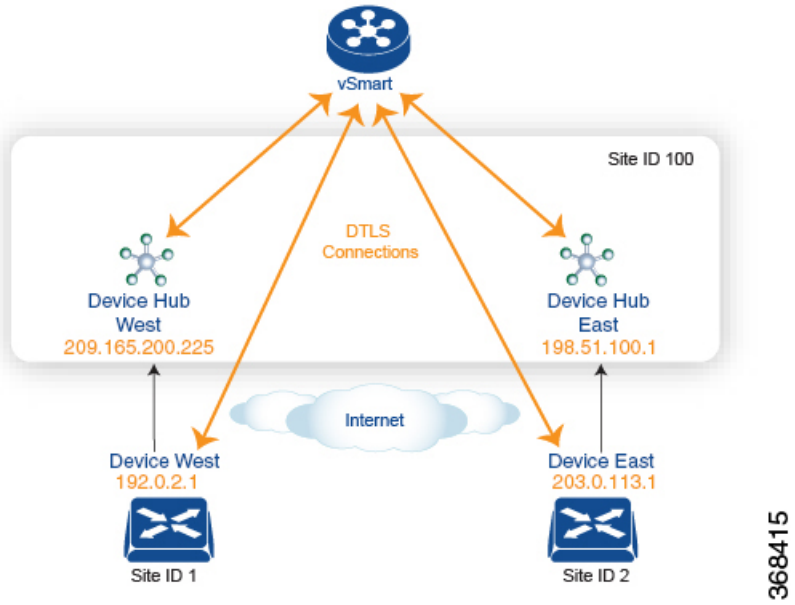

Here is a policy configuration on the Cisco vSmart Controller that would send West-side data traffic through the Cisco hub, and West and East-side traffic through the Device Hub East:

```
policy
  lists
    site-list west-sites
      site-id 1
    site-list east-sites
      site-id 2
    tloc-list west-hub-tlocs
      tloc-id 209.165.200.225 gold
    tloc-list east-hub-tlocs
      tloc-id 198.51.100.1 gold
  control-policy prefer-west-hub
    sequence 10
      match tloc
        tloc-list west-hub-tlocs
      action accept
        set preference 50
  control-policy prefer-east-hub
    sequence 10
      match tloc
        tloc-list east-hub-tlocs
      action accept
```

```
        set preference 50
apply-policy
  site west-sites control-policy prefer-west-hub out
  site east-sites control-policy prefer-east-hub out
```

Here is an explanation of this policy configuration:

Create the site lists that are required for the **apply-policy** configuration command:

- **site-list west-sites** lists all the site IDs for all the devices in the West portion of the overlay network.
- **site-list east-sites** lists the site IDs for the devices in the East portion of the network.

Create the TLOC lists that are required for the match condition in the control policy:

- **west-hub-tlocs** lists the TLOC for the Device West Hub, which we want to service traffic from the West-side device.
- **east-hub-tlocs** lists the TLOC for the Device East Hub, to service traffic from the East devices.

Define two control policies:

- **prefer-west-hub** affects OMP routes destined to TLOC 209.165.200.225, gold, which is the TLOC address of the Device West hub router. This policy modifies the preference value in the OMP route to a value of 50, which is large enough that it is likely that no other OMP routes will have a larger preference. So setting a high preference value directs traffic destined for site 100 to the Device West hub router.
- Similarly, **prefer-east-hub** sets the preference to 50 for OMP routes destined TLOC 198.51.100.1, gold, which is the TLOC address of the Device East hub router, thus directing traffic destined for site 100 site to the Device East hub 198.51.100.1 router.

Apply the control policies:

- The first line in the **apply-policy** configuration has the Cisco vSmart Controller apply the **prefer-west-hub** control policy to the sites listed in the **west-sites** list, which here is only site ID 1, so that the preference in their OMP routes destined to TLOC 209.165.200.225 is changed to 50 and traffic sent from the Device West to the hub site goes through the Device West hub router.
- The Cisco vSmart Controller applies the **prefer-east-hub** control policy to the OMP routes that it advertises to the devices in the **east-sites** list, which changes the preference to 50 for OMP routes destined to TLOC 198.51.100.1, so that traffic from the Device East goes to the Device East hub router.

Localized Control Policy

Control policy operates on the control plane traffic in the Cisco IOS XE SD-WAN overlay network, influencing the determination of routing paths through the overlay network. Localized control policy is policy that is configured on a Cisco IOS XE SD-WAN device (hence, it is local) and affects BGP and OSPF routing decisions on the site-local network that the device is part of.

In addition to participating in the overlay network, a Cisco IOS XE SD-WAN device participates in the network at its local site, where it appears to the other network devices to be simply a regular router. As such, you can provision routing protocols, such as BGP and OSPF, on the Cisco IOS XE SD-WAN device so that it can exchange route information with the local-site routers. To control and modify the routing behavior on the local network, you configure a type of control policy called route policy on the devices. Route policy

applies only to routing performed at the local branch, and it affects only the route table entries in the local device's route table.

Localized control policy, which you configure on the devices, lets you affect routing policy on the network at the local site where the device is located. This type of control policy is called route policy.

Configuration Components

A route policy 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 configured, it is, by default, rejected and discarded.

The following figure illustrates the configuration components for localized control policy.

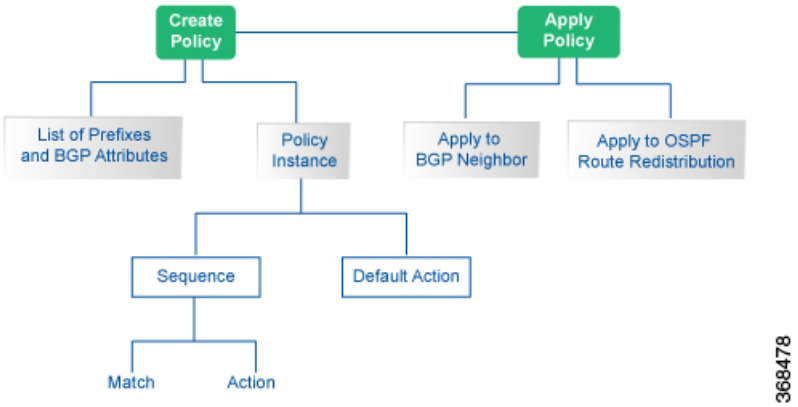

Configure Localized Control Policy Using Cisco vManage

To configure localized policies, use the Cisco vManage policy configuration wizard. The wizard is a UI policy builder that consists of five screens to configure and modify the following localized policy components:

- Groups of interest, also called lists
- Forwarding classes to use for QoS
- Access control lists (ACLs)
- Route policies
- Policy settings

You configure some or all these components depending on the specific policy you are creating. To skip a component, click the **Next** button at the bottom of the screen. To return to a component, click the **Back** button at the bottom of the screen.

Step 1: Start the Policy Configuration Wizard

To start the policy configuration wizard:

1. In the Cisco vManage NMS, select the **Configure** > **Policies** screen.
2. Select the **Localized Policy** tab.
3. Click **Add Policy**.

The policy configuration wizard opens, and the Create Groups of Interest screen is displayed.

Step 2: Configure Groups of Interest

In Create Groups of Interest, create lists of groups to use in localized policy:

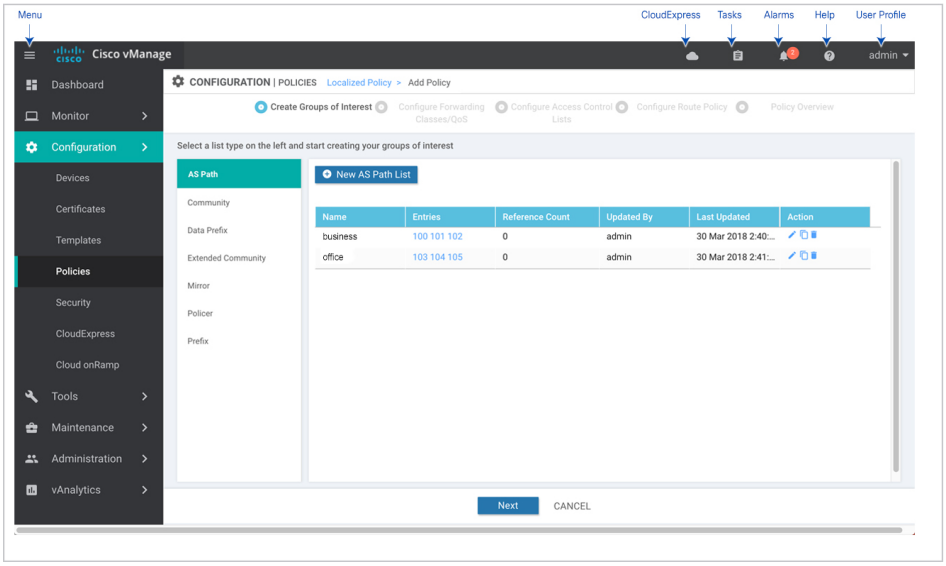

1. Create new lists, as described in the following table:

Table 2:

List Type	Procedure
AS Path	1. In the left bar, click **AS Path**. 2. Click **New AS Path List**. 3. Enter a name for the list. 4. Enter the AS path, separating AS numbers with a comma. 5. Click **Add**.

List Type	Procedure
Community	**1.** In the left bar, click **Community**. **2.** Click **New Community List**. **3.** Enter a name for the list. **4.** Enter the BGP community in the format *aa*:*nn* or as the string **internet**, **local-as**, **no-advertise**, or **no-export**, separating multiple items with a comma. For *aa*, enter a 2-byte AS number, and for *nn*, enter a 2-byte network number. **5.** Click Add.
Extended Community	**1.** In the left bar, click **Extended Community**. **2.** Click **New Extended Community List**. **3.** Enter a name for the list. **4.** Enter the BGP extended community as **rt** (*aa*:*nn* | *ip-address*), for a route target community, or **soo** (*aa*:*nn* | *ip-address*), for a route origin community, separating multiple items with a comma. For *aa,* enter a 2-byte AS number, and for *nn* enter a 2-byte network number. **5.** Click **Add**.
Mirror	**1.** In the left bar, click **TLOC**. **2.** Click **New TLOC List**. The TLOC List popup displays. **3.** Enter a name for the list. **4.** In the TLOC IP field, enter the system IP address for the TLOC. **5.** In the Color field, select the TLOC's color. **6.** In the Encap field, select the encapsulation type. **7.** In the Preference field, optionally select a preference to associate with the TLOC. **8.** Click **Add TLOC** to add another TLOC to the list. **9.** Click **Save**.
Policer	**1.** In the left bar, click **VPN**. **2.** Click **New VPN List**. **3.** Enter a name for the list. **4.** In the Add VPN field, enter one or more VPN IDs separated by commas. **5.** Click **Add**.

List Type	Procedure
Prefix	**1.** In the left bar, click **Prefix**. **2.** Click **New Prefix List**. **3.** Enter a name for the list. **4.** Enter the IP prefix in one of the following formats: • *prefix*/*length*—Exactly match a single prefix–length pair. • **0.0.0.0/0**—Match any prefix–length pair. • **0.0.0.0/0 le** *length*—Match any IP prefix whose length is less than or equal to *length*. For example, **ip-prefix 0.0.0.0/0 le 16** matches all IP prefixes with lengths from /1 through /16. • **0.0.0.0/0 ge** *length*—Match any IP prefix whose length is greater than or equal to *length*. For example, **ip-prefix 0.0.0.0 ge 25** matches all IP prefixes with lengths from /25 through /32. • **0.0.0.0/0 ge** *length1* **le** *length2*, or **0.0.0.0 le** *length2* **ge** *length1*—Match any IP prefix whose length is greater than or equal to *length1* and less than or equal to *length2*. For example, **ip-prefix 0.0.0.0/0 ge 20 le 24** matches all /20, /21, /22, /23, and /24 prefixes. Also, **ip-prefix 0.0.0.0/0 le 24 ge 20** matches the same prefixes. If *length1* and *length2* are the same, a single IP prefix length is matched. For example, **ip-prefix 0.0.0.0/0 ge 24 le 24** matches only /24 prefixes. **5.** Click **Add**.

1. Click **Next** to move to Configure Forwarding Classes/QoS in the wizard.
2. Click **Next** to move to Configure Access Control Lists in the wizard.
3. Click **Next** to move to Configure Route Policies in the wizard.

Step 3: Configure Route Policies

In Configure Route Policies, configure the routing policies:

1. In the **Add Route Policy** tab, select **Create New**.
2. Enter a name and description for the route policy.
3. In the left pane, click **Add Sequence Type**. A Route box is displayed in the left pane.
4. Double-click the **Route** box, and type a name for the route policy.
5. In the right pane, click **Add Sequence Rule** to create a single sequence in the policy. The Match tab is selected by default.
6. Click a match condition.
7. On the left, enter the values for the match condition.
8. On the right enter the action or actions to take if the policy matches.

9. Repeat Steps 6 through 8 to add match–action pairs to the route policy.
10. To rearrange match–action pairs in the route policy, in the right pane drag them to the desired position.
11. To remove a match–action pair from the route policy, click the X in the upper right of the condition.
12. Click **Save Match and Actions** to save a sequence rule.
13. To rearrange sequence rules in an route policy, in the left pane drag the rules to the desired position.
14. To copy, delete, or rename an route policy sequence rule, in the left pane, click **More Options** next to the rule's name and select the desired option.
15. If no packets match any of the route policy sequence rules, the default action is to drop the packets. To change the default action:
    a. Click **Default Action** in the left pane.
    b. Click the Pencil icon.
    c. Change the default action to Accept.
    d. Click **Save Match and Actions**.
16. Click **Next** to move to Policy Overview in the wizard.
17. Click **Preview** to view the full policy in CLI format.
18. Click **Save Policy**.

Step 4: Apply a Route Policy in a Device Template

1. In the Cisco vManage NMS, select the **Configuration > Templates** screen.
2. If you are creating a new device template:
    a. In the Device tab, click **Create Template**.
    b. From the Create Template drop-down, select **From Feature Template**.
    c. From the Device Model drop-down, select one of the devices.
    d. In the Template Name field, enter a name for the device template. This field is mandatory and can contain only uppercase and lowercase letters, the digits 0 through 9, hyphens (–), and underscores (_). It cannot contain spaces or any other characters.
    e. In the Description field, enter a description for the device template. This field is mandatory, and it can contain any characters and spaces.
    f. Continue with Step 4.
3. If you are editing an existing device template:
    a. In the Device tab, click the **More Actions** icon to the right of the desired template, and click the pencil icon.
    b. Click the **Additional Templates** tab. The screen scrolls to the Additional Templates section.
    c. From the Policy drop-down, select the name of a policy that you have configured.

4. Click the **Additional Templates** tab located directly beneath the Description field. The screen scrolls to the Additional Templates section.
5. From the Policy drop-down, select the name of the policy you configured in the above procedure.
6. To apply a route policy to BGP:
   a. Scroll to the Service VPN section.
   b. In the Service VPN drop-down, type the service VPN number (a VPN number other than 0 or 512).
   c. From Additional VPN Templates, select BGP.
   d. From the BGP drop-down, click **Create Template** or **View Template**.
   e. Select the **Neighbor** tab, click the plus sign (+), and click **More**.
   f. In Address Family, change the scope to Device Specific. Then, Click On to enable Address Family, Click On to enable Route Policy In, and specify the name of a route policy to apply to prefixes received from the neighbor, or Click On to enable Route Policy Out, and specify the name of a route policy to apply to prefixes sent to the neighbor. This name is one that you configured with a **policy route-policy** command.
   g. Click **Save** to save the neighbor configuration, and then click **Save** to save the BGP configuration.
7. To apply a route policy to routes coming from all OSPF neighbors:
   a. Scroll to the Service VPN section.
   b. In the Service VPN drop-down, type the service VPN number (a VPN number other than 0 or 512).
   c. From Additional VPN Templates, select **OSPF**.
   d. Click **Create Template** or **View Template**.
   e. Select the **Advanced** tab.
   f. In Policy Name, specify the name of a route policy to apply to incoming routes. This name is one that you configured with a **policy route-policy** command.
   g. Click **Save**.
8. To apply a route policy before redistributing routes into OSPF:
   a. Scroll to the Service VPN section.
   b. In the Service VPN drop-down, type the service VPN number (a VPN number other than 0 or 512).
   c. From Additional VPN Templates, select **OSPF**.
   d. Click **Create Template** or **View Template**.
   e. Select the **Redistribute** tab, click the plus sign (+), and select the protocol from which to redistribute routes into OSPF.
   f. Specify the name of a route policy to apply to the routes being redistributed. This name is one that you configured with a **policy route-policy** command.
   g. Click **Save**.

9. Click **Save** (for a new template) or **Update** (for an existing template).

Configure Localized Control Policy Using CLI

To configure a route policy using the CLI:

1. Create lists of prefixes, as needed:

```
Device(config)# policy
Device(config-policy)# lists
Device(config-lists)# prefix-list  list-name
Device(config-lists-list-name)# ip-prefix prefix/length
```

2. Create lists of BGP AS paths, and community and extended community attributes, as needed:

```
Device(config)# policy lists
Device(config-lists)# as-path-list list-name
Device(config-lists-list-name)# as-path path-list
Device(config)# policy lists
Device(config-lists)# community-list list-name
Device(config-lists-list-name)# community [aa:nn |
internet | local-as | no-advertise | no-export]
Device(config-lists)# ext-community-list list-name
Device(config-lists-list-name)# community [rt (aa:nn |
ip-address) | soo (aa:nn | ip-address)]
```

1. Create a route policy instance:

```
Device(config)# policy route-policy policy-name
Device(config-route-policy-policy-name)#
```

2. Create a series of match–action pair sequences:

```
Device(config-route-policy-policy-name)# sequence number
Device(config-sequence-number)#
```

   The match–action pairs are evaluated in order, by sequence number, starting with the lowest numbered pair and ending when the route matches the conditions in one of the pairs. Or if no match occurs, the default action is taken (either rejecting the route or accepting it as is).

3. Define match parameters for routes:

```
Device(config-sequence-number)#  match match-parameter
```

4. Define actions to take when a match occurs:

```
Device(config-sequence-number)# action reject
Device(config-sequence-number)# action accept set parameter
```

5. Create additional numbered sequences of match–action pairs within the router policy, as needed.
6. If a route does not match any of the conditions in one of the sequences, it is rejected by default. To accept nonmatching routes, configure the default action for the policy:

```
Device(config-policy-name)# default-action accept
```

7.

Structural Components for Localized Control Policy

Following are the structural components required to configure localized control policy. Each one is explained in more detail in the sections below.

```
policy
  lists
    as-path-list list-name
      as-path path-list
    community-list list-name
      community [aa:nn | internet | local-as | no-advertise | no-export]
    ext-community-list list-name
      community [rt (aa:nn | ip-address) | soo (aa:nn | ip-address)]
    prefix-list list-name
      ip-prefix prefix/length
  route-policy policy-name
    sequence number
      match
        match-parameters
      action
        reject
        accept
          set parameters
    default-action
      (accept | reject)
 vpn vpn-id router bgp local-as-number neighbor address
   address-family ipv4-unicast
     route-policy policy-name (in | out)
 vpn vpn-id router ospf
   route-policy policy-name in
   redistribute (bgp | connected | nat | omp | static) route-policy policy-name
```

Lists

Route policy uses the following types of lists to group related items. You configure lists under the **policy lists** command hierarchy on Cisco IOS XE SD-WAN devices.

Table 3:

List Type	Description	vManage Configuration/ CLI Configuration Command
AS paths	List of one or more BGP AS paths. You can write each AS as a single number or as a regular expression. To specify more than one AS in a single path, include the list in quotation marks (" "). To configure multiple AS paths in a single list, include multiple **as-path** options, specifying one AS path in each option.	**Configuration > Policies > Localized Policy > Add Policy > Create Groups of Interest > AS Path** **Configuration > Policies > Custom Options > Localized Policy > Lists > AS Path** **as-path-list** *list-name* **as-path** *path-list*

List Type	Description	vManage Configuration/ CLI Configuration Command
Communities	List of one of more BGP communities. In **community**, you can specify: • *aa*:*nn*: Autonomous system number and network number. Each number is a 2-byte value with a range from 1 to 65535. • **internet**: Routes in this community are advertised to the Internet community. This community comprises all BGP-speaking networking devices. • **local-as**: Routes in this community are not advertised outside the local AS. • **no-advertise**: Attach the NO_ADVERTISE community to routes. Routes in this community are not advertised to other BGP peers. • **no-export**: Attach the NO_EXPORT community to routes. Routes in this community are not advertised outside the local AS or outside a BGP confederation boundary. To configure multiple BGP communities in a single list, include multiple **community** options, specifying one community in each option.	**Configuration > Policies > Localized Policy > Add Policy > Create Groups of Interest > Community** **Configuration > Policies > Custom Options > Localized Policy > Lists > Community** **community-list** *list-name* **community** [*aa*:*nn* | **internet** | **local-as** | **no-advertise** | **no-export**]
Extended communities	List of one or more BGP extended communities. In **community**, you can specify: • **rt** (*aa*:*nn* | *ip-address*): Route target community, which is one or more routers that can receive a set of routes carried by BGP. Specify this as the autonomous system number and network number, where each number is a 2-byte value with a range from 1 to 65535, or as an IP address. • **soo** (*aa*:*nn* | *ip-address*): Route origin community, which is one or more routers that can inject a set of routes into BGP. Specify this as the autonomous system number and network number, where each number is a 2-byte value with a range from 1 to 65535, or as an IP address. To configure multiple extended BGP communities in a single list, include multiple **community** options, specifying one community in each option.	**Configuration > Policies > Localized Policy > Add Policy > Create Groups of Interest > Extended Community** **Configuration > Policies > Custom Options > Localized Policy > Lists > Extended Community** **ext-community-list** *list-name* **community** [**rt** (*aa*:*nn* | *ip-address*) | **soo** (*aa*:*nn* | *ip-address*)]

List Type	Description	vManage Configuration/ CLI Configuration Command
Prefixes	List of one or more IP prefixes. To configure multiple prefixes in a single list, include multiple **ip-prefix** options, specifying one prefix in each option. Specify the IP prefixes as follows: • *prefix*/*length*—Exactly match a single prefix–length pair. • **0.0.0.0/0**—Match any prefix–length pair. • **0.0.0.0/0 le** *length*—Match any IP prefix whose length is less than or equal to *length*. For example, **ip-prefix 0.0.0.0/0 le 16** matches all IP prefixes with lengths from /1 through /16. • **0.0.0.0/0 ge** *length*—Match any IP prefix whose length is greater than or equal to *length*. For example, **ip-prefix 0.0.0.0 ge 25** matches all IP prefixes with lengths from /25 through /32. • **0.0.0.0/0 ge** *length1* **le** *length2*, or **0.0.0.0 le** *length2* **ge** *length1*—Match any IP prefix whose length is greater than or equal to *length1* and less than or equal to *length2*. For example, **ip-prefix 0.0.0.0/0 ge 20 le 24** matches all /20, /21, /22, /23, and /24 prefixes. Also, **ip-prefix 0.0.0.0/0 le 24 ge 20** matches the same prefixes. If *length1* and *length2* are the same, a single IP prefix length is matched. For example, **ip-prefix 0.0.0.0/0 ge 24 le 24** matches only /24 prefixes.	**Configuration > Policies > Localized Policy > Add Policy > Create Groups of Interest > Prefix** **Configuration > Policies > Custom Options > Localized Policy > Lists > Prefix** **prefix-list** *list-name* **ip-prefix** *prefix*/*length*

Sequences

A localized control policy contains sequences of match–action pairs. The sequences are numbered to set the order in which a route is analyzed by the match–action pairs in the policy.

In Cisco vManage NMS, you configure sequences from:

- **Configuration > Policies > Localized Policy > Add Policy > Configure Route Policy > Sequence Type**
- **Configuration > Policies > Custom Options > Localized Policy > Route Policy > Sequence Type**

In the CLI, you configure sequences with the **route-policy sequence** command.

Each sequence in a localized control policy can contain one match condition and one action condition.

Match Parameters

In Cisco vManage NMS, you configure sequences from:

- **Configuration > Policies > Localized Policy > Add Policy > Configure Route Policy > Sequence Type > Sequence Rule > Match**

- **Configuration > Policies > Custom Options > Localized Policy > Route Policy > Sequence Type > Sequence Rule > Match**

In the CLI, you configure sequences with the **route-policy sequence match** command.

For route policy routes, you can match these attributes:

Table 4:

Description	vManage Configuration/ CLI Configuration Command	Value or Range
IP prefix or prefixes from which the route was learned	Match Address **address** *list-name*	Name of an IP prefix list
BGP AS paths	Match AS Path List **as-path** *list-name*	Name of an AS path list
BGP communities	Match Community List **community** *list-name*	Name of a BGP community list
BGP extended communities	Match Extended Community List **ext-community** *list-name*	Name of a BGP extended community list
BGP local preference	Match BGP Local Preference **local-preference** *number*	0 through 4294967295
Route metric	Match Metric **metric** *number*	0 through 4294967295
Next hop	Match Next Hop **next-hop** *list-name*	Name of an IP prefix list
OMP tag for OSPF	Match OMP Tag **omp-tag** *number*	0 through 4294967295
BGP origin code	Match Origin **origin** *origin*	**egp** (default), **igp**, **incomplete**
OSPF tag value	Match OSPF Tag **ospf-tag** *number*	0 through 4294967295
Peer address	Match Peer **peer** *address*	IP address

Action Parameters

For each match condition, you configure a corresponding action to take if the packet matches.

In Cisco vManage NMS, you configure match parameters from:

- **Configuration > Policies > Localized Policy > Add Policy > Configure Route Policy > Sequence Type > Sequence Rule > Action**
- **Configuration > Policies > Custom Options > Localized Policy > Configure Route Policy > Sequence Type > Sequence Rule > Action**

In the CLI, you configure actions with the **policy control-policy action** command.

Each sequence in a localized control policy can contain one action condition.

When a route matches the conditions in the match portion of a route policy, the route can be accepted or rejected:

Table 5:

Description	vManage Configuration/ CLI Configuration Command	Value or Range
Accept the route. An accepted route is eligible to be modified by the additional parameters configured in the **action** portion of the policy configuration.	Click Accept **accept**	—
Discard the packet.	Click Reject **reject**	—

Then, for a route that is accepted, the following actions can be configured:

Table 6:

Description	vManage Configuration/ CLI Configuration Command	Value or Range
Set the AS number in which a BGP route aggregator is located and the IP address of the route aggregator.	Click Accept, then action Aggregator **set aggregator** *as-number ip-address*	1 through 65535
Set an AS number or a series of AS numbers to exclude from the AS path or to prepend to the AS path.	Click Accept, then action AS Path **set as-path** (**exclude** | **prepend**) *as-number*	1 through 65535
Set the BGP atomic aggregate attribute.	Click Accept, then action Atomic Aggregate **set atomic-aggregate**	—
Set the BGP community value.	Click Accept, then action Community **set community** *value*	[*aa*:*nn* | **internet** | **local-as** | **no-advertise** | **no-export**]
Set the BGP local preference.	Click Accept, then action Local Preference **set local-preference** *number*	0 through 4294967295

Description	vManage Configuration/ CLI Configuration Command	Value or Range
Set the metric value.	Click Accept, then action Metric **set metric** *number*	0 through 4294967295
Set the metric type.	Click Accept, then action Metric Type **set metric-type** *type*	**type1**, **type2**
Set the next-hop address.	Click Accept, then action Next Hop **set next-hop** *ip-address*	IP address
Set the OMP tag for OSPF to use.	Click Accept, then action OMP Tag **set omp-tag** *number*	0 through 4294967295
Set the BGP origin code.	Click Accept, then action Origin **set origin** *origin*	**egp**, **igp** (default), **incomplete**
Set the IP address from which the route was learned.	Click Accept, then action Originator **set originator** *ip-address*	IP address
Set the OSPF tag value.	Click Accept, then action OSPF Tag **set ospf-tag** *number*	0 through 4294967295
Set the BGP weight.	Click Accept, then action Weight **set weight** *number*	0 through 4294967295

To display the OMP and OSPF tag values associated with a route, use the **show ip routes detail** command.

Action Parameters

For each match condition, you configure a corresponding action to take if the packet matches.

In Cisco vManage NMS, you configure match parameters from:

- **Configuration > Policies > Localized Policy > Add Policy > Configure Route Policy > Sequence Type > Sequence Rule > Action**
- **Configuration > Policies > Custom Options > Localized Policy > Configure Route Policy > Sequence Type > Sequence Rule > Action**

In the CLI, you configure actions with the **policy control-policy action** command.

Each sequence in a localized control policy can contain one action condition.

When a route matches the conditions in the match portion of a route policy, the route can be accepted or rejected:

Table 7:

Description	vManage Configuration/ CLI Configuration Command	Value or Range
Accept the route. An accepted route is eligible to be modified by the additional parameters configured in the **action** portion of the policy configuration.	Click Accept **accept**	—
Discard the packet.	Click Reject **reject**	—

Then, for a route that is accepted, the following actions can be configured:

Table 8:

Description	vManage Configuration/ CLI Configuration Command	Value or Range
Set the AS number in which a BGP route aggregator is located and the IP address of the route aggregator.	Click Accept, then action Aggregator **set aggregator** *as-number ip-address*	1 through 65535
Set an AS number or a series of AS numbers to exclude from the AS path or to prepend to the AS path.	Click Accept, then action AS Path **set as-path** (**exclude** | **prepend**) *as-number*	1 through 65535
Set the BGP atomic aggregate attribute.	Click Accept, then action Atomic Aggregate **set atomic-aggregate**	—
Set the BGP community value.	Click Accept, then action Community **set community** *value*	[*aa*:*nn* | **internet** | **local-as** | **no-advertise** | **no-export**]
Set the BGP local preference.	Click Accept, then action Local Preference **set local-preference** *number*	0 through 4294967295
Set the metric value.	Click Accept, then action Metric **set metric** *number*	0 through 4294967295
Set the metric type.	Click Accept, then action Metric Type **set metric-type** *type*	**type1**, **type2**
Set the next-hop address.	Click Accept, then action Next Hop **set next-hop** *ip-address*	IP address
Set the OMP tag for OSPF to use.	Click Accept, then action OMP Tag **set omp-tag** *number*	0 through 4294967295

Description	vManage Configuration/ CLI Configuration Command	Value or Range
Set the BGP origin code.	Click Accept, then action Origin **set origin** *origin*	**egp**, **igp** (default), **incomplete**
Set the IP address from which the route was learned.	Click Accept, then action Originator **set originator** *ip-address*	IP address
Set the OSPF tag value.	Click Accept, then action OSPF Tag **set ospf-tag** *number*	0 through 4294967295
Set the BGP weight.	Click Accept, then action Weight **set weight** *number*	0 through 4294967295

To display the OMP and OSPF tag values associated with a route, use the **show ip routes detail** command.

Default Action

If a route being evaluated does not match any of the match conditions in a localized control policy, a default action is applied to this route. By default, the route is rejected.

In Cisco vManage NMS, you modify the default action from **Configuration > Policies > Localized Policy > Add Policy > Configure Route Policy > Sequence Type > Sequence Rule > Default Action**.

In the CLI, you modify the default action with the **control policy default-action accept** command.

Apply Route Policy for BGP

For a route policy to take effect for BGP, you must apply it to an address family. Currently, the Cisco SD-WAN software supports only the IPv4 address family.

To apply a BGP route policy in the Cisco vManage NMS:

1. In the Cisco vManage NMS, select the **Configure > Templates** screen.
2. In the Device tab, click the **Create Template** drop-down and select **From Feature Template**.
3. From the Device Model drop-down, select the type of device for which you are creating the template. The Cisco vManage NMS displays all the feature templates for that device type. The required feature templates are indicated with an asterisk (*), and the remaining templates are optional. The factory-default template for each feature is selected by default.
4. In the Template Name field, enter a name for the device template. This field is mandatory and can contain only uppercase and lowercase letters, the digits 0 through 9, hyphens (-), and underscores (_). It cannot contain spaces or any other characters.
5. In the Description field, enter a description for the device template. This field is mandatory, and it can contain any characters and spaces.
6. In the Basic Information bar, click the **Service VPN** tab.
7. In the Service VPN field, select the **VPN number**.

8. In Additional VPN Templates, select **BGP**.
9. Select **Create Template**.
10. In the Basic Configuration bar, click **IPv4 Unicast Address Family**.
11. In the Address Family field, select **ipv4-unicast**.
12. In the Redistribute tab, click **New Redistribute**.
13. In the Route Policy field, enter the name of the route policy to apply to redistributed routes.
14. Click **Add**.
15. Click **Save**.

To apply a BGP route policy in the CLI:

```
Device(config)# vpn
vpn-id
router bgp
local-as-number
neighbor address
address-family ipv4-unicast route-policy
policy-name (in | out)
```

Applying the policy in the inbound direction (**in**) affects routes being received by BGP. Applying the policy in the outbound direction (**out**) affects routes being advertised by BGP.

Apply Route Policy for OSPF

For a route policy to take effect for OSPF, you can apply it to all inbound traffic.

To apply an OSPF route policy in the Cisco vManage NMS:

1. In the Cisco vManage NMS, select the **Configure > Templates** screen.
2. In the Device tab, click the **Create Template** drop-down and select From Feature Template.
3. From the Device Model drop-down, select the type of device for which you are creating the template. The Cisco vManage NMS displays all the feature templates for that device type. The required feature templates are indicated with an asterisk (*), and the remaining templates are optional. The factory-default template for each feature is selected by default.
4. In the Template Name field, enter a name for the device template. This field is mandatory and can contain only uppercase and lowercase letters, the digits 0 through 9, hyphens (-), and underscores (_). It cannot contain spaces or any other characters.
5. In the Description field, enter a description for the device template. This field is mandatory, and it can contain any characters and spaces.
6. In the Basic Information bar, click the **Service VPN** tab.
7. In the Service VPN field, select the **VPN number**.
8. In Additional VPN Templates, select **OSPF**.
9. Select **Create Template**.
10. In the Basic Configuration bar, click **Redistribute**.

11. Click **New Redistribute**.
12. In the Route Policy field, enter the name of the route policy to apply to redistributed routes.
13. Click **Add**.
14. Click **Save**.

To apply an OSPF route policy in the CLI:

```
Device(config)# vpn vpn-id
router ospf route-policy policy-name in
```

You can also apply the policy when redistributing routes into OSPF:

```
Device(config)# vpn
vpn-id
router ospf redistribute(bgp|connected|nat|omp|static)route-policy
policy-name
```

Device Access Policy

Table 9: Feature History

Feature Name	Release Information	Description
Ability to apply ACL to SNMP like on Cisco products ACL matching SSH,VTY		Access policies define rules that traffic must meet to pass through an interface. When you define rules for incoming traffic, they are applied to the traffic before any other policies are applied. Cisco SD-WAN vEdge control plane processes data traffic for local services (like SSH and SNMP) from a set of sources in a VPN. Routing packets are required to form the overlay. It is important to protect the control plane CPU from device access traffic by applying the filter.

Device Access Policy Overview

The Cisco vManage user interface is enhanced to configure device access policy on Cisco IOS XE SD-WANCisco SD-WAN devices.

The Cisco vManage control plane processes the data traffic for local services like, SSH and SNMP from a set of sources in a VPN. It is important to protect the control plane CPU from device access traffic by applying the filter to avoid unnecessary traffic.

Access policies define the rules that traffic must meet to pass through an interface. When you define rules for incoming traffic, they are applied to the traffic before any other policies are applied. You can use access policies, in routed and transparent firewall mode to control IP traffic. An access rule permits or denies traffic based on the protocol used, the source and destination IP address or network, and optionally, the users and user groups. Each incoming packet at an interface is analysed to determine if it must be forwarded or dropped based on criteria you specify. If you define access rules for the outgoing traffic, packets are also analyzed before they are allowed to leave an interface. Access policies are applied in order. That is, when the device compares a packet to the rules, it searches from top to bottom in the access policies list, and applies the policy

for the first matched rule, ignoring all subsequent rules (even if a later rule is a better match). Thus, you should place specific rules above more general rules to ensure the specific rules are not skipped.

Configure Device Access Policy Using vManage

Cisco IOS XE SD-WAN devices supports device access policy configuration to handle SNMP and SSH traffic directed towards Control Plane. Use Cisco vManage to configure destination port based on device access policy.

Note In order to allow connection to device from vManage Tools > SSH Terminal tab, create a rule to accept **Device Access Protocol** as SSH and **Source Data Prefix** as 192.168.1.5/32.

To configure localized device access control policies, use the Cisco vManage policy configuration wizard.

Configure specific or all components depending on the specific policy you are creating. To skip a component, click the **Next** button. To return to a component, click the **Back** button at the bottom of the screen.

To configure Device Access Policy:

1. In the Cisco vManage, select the **Configure > Policies** screen.
2. Select the **Localized Policy** tab.
3. From **Custom Options > Localized Policy** pane, select **Access Control Lists.**
4. Click **Add Device Access Policy** drop down list to add a device. The options are **Add IPv4 Device Access Policy** and **Add IPv6 Device Access Policy**.
5. Select **Add IPv4 Device Access Policy** from the drop-down list to add IPv4 ACL Policy. The Edit Device IPv4 ACL Policy page displays.
6. Enter the name and the description for the new policy.
7. Click **Add ACL Sequence** to add a sequence. The Device Acces Control List page displays.
8. Click **Sequence Rule**. Match and Actions options display.
9. From the **Match** pane, select and configure the following conditions for your ACL policy:

Match Condition	Description
Device Access Protocol (required)	Select a carrier from the drop-down list. For example SNMP, SSH.
Source Data Prefix	Enter the source IP address. For example, 10.0.0.0/12.
Source Port	Enter the list of source ports. The range is 0-65535.
Destination Data Prefix	Enter the destination IP address. For example, 10.0.0.0/12.
Destination VPN	Enter a VPN ID.

10. From the **Actions** tab, configure the following conditions for your ACL policy:

Action Condition	Description
Accept	
Counter Name	Enter the counter name to be accepted. The maximum length can be 20 characters.
Drop	
Counter Name	Enter the counter name to drop. The maximum length can be 20 characters.

11. Click **Save Match And Actions** to save all the conditions for ACL policy.
12. Click **Save Device Access Control List Policy** to apply the selected match conditions to an action.
13. If no packets match any of the route policy sequence rules, the **Default Action** in the left pane is to drop the packets.

Note IPv6 Prefix match is not supported on Cisco IOS XE SD-WAN devices. When you try to configure IPv6 prefix match on these devices, Cisco vManage fails to generate device configuration.

Configure Device Access Policy Using CLIs

To configure Device Access Policy:

```
Device(config)# system
Device(config-system) device-access-policy ipv4 <pol-name>
```

Configuration:

```
Device(config)# policy
Device(config-policy)policy device-access-policy <name>
    sequence 1
        match
            destination-data-prefix-list    Destination prefix list
            destination-ip                  List of destination addresses
            destination-port                List of destination ports
            dscp                            List of DSCP values
            packet-length                   Packet length
            protocol                        List of protocols
            source-data-prefix-list         Source prefix list
            source-ip                       List of source addresses
            source-port                     List of source ports
            destination-vpn                 List of VPN-ID
        action
            accept
            count                           Number of packets/bytes matching this rule
            drop
    default-action       Accept or drop

system
    device-access-policy ipv4 <pol-name>
```

Note IPv6 Prefix match is not supported on Cisco IOS XE SD-WAN devices.

The following example shows the sample configuration for Device Access Policy:

```
policy device-access-policy dev_pol
  sequence 1
   match
    destination-port 22
   !
   action drop
    count ssh_packs
   !
  !
  default-action drop
 !
 device-access-policy snmp_policy
  sequence 2
   match
    destination-port 161
   !
   action drop
    count snmp_packs
   !
  !
  default-action accept
 !
!
system
 device-access-policy ipv4 snmp_policy
!
```

Examples for ACL Statistics and Counters

To configure ACL statistics and counters using yang:

```
Yang file: Cisco-IOS-XE-acl-oper.yang

  grouping ace-oper-data {
    description
      "ACE operational data";
    leaf match-counter {
      type yang:counter64;
      description
        "Number of matches for an access list entry";
    }
  }
```

Example configuration using yang model:

```
Router#config-t
Router(config)# ip access-list extended ACL-1
Router(config-ext-nacl)# 1 permit ip 10.10.10.1 0.0.0.0 any
Router(config-ext-nacl)# 2 deny ip 20.20.0.0 0.0.255.255 any
Router(config-ext-nacl)# commit
Commit complete.

Router#
Router#
Router#request platform software system shell
Activity within this shell can jeopardize the functioning of the system.
```

```
Are you sure you want to continue? [y/n] y
[Router:/]$

[Router:/]$

 [Router:/]$


[Router:/]$

[Router:/]$ confd_cli -C -P 3010 -noaaa -g sdwan-oper

root connected from 127.0.0.1 using console on Router

Router# show access-lists access-list ACL-1
ACCESS
CONTROL
LIST     RULE  MATCH
NAME     NAME  COUNTER
------------------------
ACL-1    1     0
         2     0

Router# show access-lists access-list ACL-1 | display xml
<config xmlns="http://tail-f.com/ns/config/1.0">
  <access-lists xmlns="http://cisco.com/ns/yang/Cisco-IOS-XE-acl-oper">
  <access-list>
    <access-control-list-name>ACL-1</access-control-list-name>
    <access-list-entries>
      <access-list-entry>
        <rule-name>1</rule-name>
        <access-list-entries-oper-data>
          <match-counter>0</match-counter>
        </access-list-entries-oper-data>
      </access-list-entry>
      <access-list-entry>
        <rule-name>2</rule-name>
        <access-list-entries-oper-data>
          <match-counter>0</match-counter>
        </access-list-entries-oper-data>
      </access-list-entry>
    </access-list-entries>
  </access-list>
  </access-lists>
</config>
Router#
```

To configure ACL statistics and counters using CLI, use the command `show ip access-list [access-list-number | access-list-name]`.

Example configuration using CLI:

```
show ip access-list [access-list-number | access-list-name]

Example:
Router# show ip access-list ACL-1
Extended IP access list ACL-1
10 permit ip host 10.1.1.1 any (3 matches) 30
30 permit ip host 10.2.2.2 any (27 matches)



To clear counters in ACL stats:

clear ip access-list counters {access-list-number | access-list-name}
```

Verifying Device Access Policy Configuration

Cisco IOS XE SD-WAN devices support the following operational commands to provide information for device-access-policy. These commands provide a visual for the counters and the names of the configured device-access-policy. The two commands and the respective yang models are shown in the following sections.

Yang Model for the command **device-access-policy-counters**:

```
list device-access-policy-counters {
        tailf:info "IPv6 Device Access Policy counters";
        when "/viptela-system:system/viptela-system:personality = 'vedge'";
        tailf:callpoint device-access-policy-counters-v6; // _nfvis_exclude_line_
        key "name";
        tailf:hidden cli;

        leaf name {
          tailf:info "Device Access Policy name";
          type viptela:named-type-127;
        }
        config false;
        list device-access-policy-counter-list {
          tailf:info "Device access policy counter list";
          tailf:callpoint device-access-policy-counter-list-v6; // _nfvis_exclude_line_
          tailf:cli-no-key-completion;
          tailf:cli-suppress-show-match;
          key "counter-name";
          tailf:hidden cli;

          leaf counter-name {
            tailf:info "Counter name";
            tailf:cli-suppress-show-match;
            type viptela:named-type;
          }
          leaf packets {
            type yang:counter64;
            tailf:cli-suppress-show-match;
          }
          leaf bytes {
            type yang:counter64;
            tailf:cli-suppress-show-match;
          }
        }
```

The folllwoing example shows the policy details of a conter.

```
show policy device-access-policy-counters

                  COUNTER
NAME          NAME        PACKETS   BYTES
-------------------------------------------------------------
dev_pol       ssh_packs      -       -
snmp_policy   snmp_packs     0       0
```

Yang Model for the command **device-access-policy-names**:

```
list device-access-policy-names {
        tailf:info "IPv6 device access policy names";
        when "/viptela-system:system/viptela-system:personality = 'vedge'";
        tailf:callpoint device-access-policy-names-v6; // _nfvis_exclude_line_
        tailf:cli-no-key-completion;
        key "name";
```

```
        tailf:hidden cli;

        leaf name {
          tailf:info "Device Access Policy name";
          type viptela:named-type-127;
        }
        config false;
      }
```

The following example shows the list of configured policies:

```
show policy device-access-policy-names

NAME
-------------
dev_pol
snmp_policy
```

Verifying ACL Policy on SNMP Server

For device access policies on SNMP servers, Cisco vManage validates to block the template push on the device, if SNMP feature template is not configured.

Yang Model for the command **snmp-server community**. Following is the ACL settings sample from Cisco-IOS-XE-snmp.yang:

```
container community {
          description
            "Configure a SNMP v2c Community string and access privs";
          tailf:cli-compact-syntax;
          tailf:cli-sequence-commands;
          leaf community-string {
            tailf:cli-drop-node-name;
            type string;
          }
          container access {
            tailf:cli-drop-node-name;
            tailf:cli-flatten-container;
            leaf standard-acl {
              tailf:cli-drop-node-name;
              tailf:cli-full-command;
              type uint32 {
                range "1..99";
              }
            }
            leaf expanded-acl {
              tailf:cli-drop-node-name;
              tailf:cli-full-command;
              type uint32 {
                range "1300..1999";
              }
            }
            leaf acl-name {
              tailf:cli-drop-node-name;
              tailf:cli-full-command;
              type string;
            }
            leaf ipv6 {
              description
                "Specify IPv6 Named Access-List";
              tailf:cli-full-command;
```

```
        type string;
      }
      leaf ro {
        description
          "Read-only access with this community string";
        type empty;
      }
      leaf rw {
        description
          "Read-write access with this community string";
        type empty;
      }
    }
  }
```

Verifying ACL Policy on SSH

For device access policies on SSH servers using Virtual Teletype (VTY) lines, Cisco vManage uses all the available VTY lines in the backend and pushes policy accordingly.

Following is the ACL settings sample from Cisco-IOS-XE-line.yang:

```
// line * / access-class
    container access-class {
      description
        "Filter connections based on an IP access list";
      tailf:cli-compact-syntax;
      tailf:cli-sequence-commands;
      tailf:cli-reset-container;
      tailf:cli-flatten-container;
      list acccess-list {
        tailf:cli-drop-node-name;
        tailf:cli-compact-syntax;
        tailf:cli-reset-container;
        tailf:cli-suppress-mode;
        tailf:cli-delete-when-empty;
        key "direction";
        leaf direction {
          type enumeration {
            enum "in";
            enum "out";
          }
        }
        leaf access-list {
          tailf:cli-drop-node-name;
          tailf:cli-prefix-key;
          type ios-types:exp-acl-type;
          mandatory true;
        }
        leaf vrf-also {
          description
            "Same access list is applied for all VRFs";
          type empty;
        }
      }
    }
```