



Policies Configuration Guide, Cisco IOS XE SD-WAN Releases 16.11, 16.12

First Published: 2019-07-03

Last Modified: 2020-04-28

Americas Headquarters

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



CONTENTS

CHAPTER 1	What's New for Cisco SD-WAN	1
	What's New for Cisco IOS XE SD-WAN Releases 16.12.1b, 16.12.1d, and 16.12.2r	1

CHAPTER 2	Policy Basics	5
	Policy Overview	5
	Policies in Cisco vManage	7

CHAPTER 3	Cisco SD-WAN Policy Framework Basics	11
	Cisco vSmart Policy Components	11
	TLOC Attributes Used in Policies	15
	vRoute Attributes Used in Policies	16
	Design Cisco vSmart Controller Policy Processing and Application	17
	Cisco vSmart Policy Operation	18
	Control Policy	18
	Data Policy	20
	VPN Membership Policy Operation	21
	Configure and Execute Cisco vSmart Policies	22

CHAPTER 4	Control Policies	25
	Centralized Control Policy	26
	Route Types	28
	Default Behavior Without Centralized Control Policy	28
	Behavior Changes with Centralized Control Policy	29
	Examples of Modifying Traffic Flow with Centralized Control Policy	29
	Create an Arbitrary Topology	29
	Set Up Traffic Engineering	30

Configure the Network Topology	33
Configuration Components	33
Create a Hub and Spoke Policy	34
Create a Policy for Mesh	34
Custom Control (Route and TLOC)	35
Import Existing Topology	40
Create a VPN Membership Policy	41
Configure Centralized Policy Using Cisco vManage	41
Structural Components for Centralized Control Policy	46
Apply Centralized Control Policy	54
Configure Centralized Policy Using CLI	55
Centralized Control Policy Configuration Examples	57
Localized Control Policy	61
Configure Localized Control Policy Using Cisco vManage	62
Configure Localized Control Policy Using CLI	68
Structural Components for Localized Control Policy	69
Apply Route Policy for BGP	76
Apply Route Policy for OSPF	77
Device Access Policy	78
Device Access Policy Overview	78
Configure Device Access Policy Using vManage	79
Configure Device Access Policy Using CLIs	80
Examples for ACL Statistics and Counters	81
Verifying Device Access Policy Configuration	83
Verifying ACL Policy on SNMP Server	84
Verifying ACL Policy on SSH	85
<hr/>	
CHAPTER 5	Data Policies 87
Centralized Data Policy	88
Configure Centralized Data Policy Based on Prefixes and IP Headers	89
Start the Policy Configuration Wizard	90
Step 1: Create Policy Lists	90
Step 2: Configure Traffic Rules	93
Step 3: Apply Policies to Sites and VPNs	97

Step 4: Activate a Centralized Data Policy	98
Configure Centralized Data Policy Using CLI	98
Structural Components of Policy Configuration for Centralized Data Policy	100
Lists	101
VPN Lists	103
Policer Parameters	103
Sequences	104
Match Parameters	104
Action Parameters	106
Default Action	109
Apply Centralized Data Policy	110
Deep Packet Inspection	111
Configure Deep Packet Inspection Using vManage	111
Configure Deep Packet Inspection Using CLI	113
Structural Components of Policy Configuration for Deep Packet Inspection	115
Action Parameters for Configuring Deep Packet Inspection	116
Apply Centralized Data Policy for Deep Packet Inspection	118
Centralized Data Policy Configuration Examples	120
Localized Data Policy	122
Localized Data Policy for IPv4	123
Configure Localized Data Policy for IPv4 Using Cisco vManage	123
Structural Components of Configuration for Access Lists	129
Configure Localized Data Policy for IPv4 Using the CLI for Cisco IOS XE SD-WAN Devices	136
Localized Data Policy for IPv6	138
Configure Localized Data Policy for IPv6 Using vManage	139
Structural Components of Configuration for Access Lists	144
Configure Localized Data Policy for IPv6 Using the CLI	149
Localized Data Policy Configuration Examples	150
<hr/>	
CHAPTER 6	Policy Basics CLI Reference 151
<hr/>	
CHAPTER 7	Forward Error Correction 155
	Configure Forward Error Correction for a Policy 155

Monitor Forward Error Correction Tunnel Information	156
Monitor Forward Error Application Family Information	157

CHAPTER 8

Packet Duplication for Noisy Channels	159
Information about Packet Duplication	159
Configure Packet Duplication	159
Monitor Packet Duplication Per Application	160

CHAPTER 9

Integrate Cisco IOS XE SD-WAN Device with Cisco ACI	161
Guidelines to Integrate with Cisco ACI	162
Verify Cisco ACI Registration	162
SLA Classes	162
Data Prefixes	163
VPNs	163
Map Data Prefix and VPN to SLA	163
Create an App-Route-Policy	164
Map ACI Sites	164
Unmap ACI Sites	165
Delete a Controller	165

CHAPTER 10

Application-Aware Routing	167
Components of Application-Aware Routing	168
Classification of Tunnels into SLA Classes	169
Configure Application-Aware Routing	171
Configure Application-Aware Routing Using CLIs	178
Structural Components of Policy Configuration for Application-Aware Routing	180
Apply Application-Aware Routing Policy	185
Configure the Monitoring of Data Plane Tunnel Performance	187
Application-Aware Routing Policy Configuration Example	189

CHAPTER 11

Traffic Flow Monitoring with Cflowd	193
Configure Traffic Flow Monitoring on Cisco XE SD-WAN Devices	196
Configure Global Flow Visibility	196
Configure Global Application Visibility	197

Configure Cflowd Monitoring Policy	197
Display Cflowd Information	198
Configure Cflowd Traffic Flow Monitoring Using CLI	199
Structural Components of Policy Configuration for Cflowd	200
Apply and Enable Cflowd Policy	203
Cflowd Traffic Flow Monitoring Configuration Example	204

CHAPTER 12**Lawful Intercept 209**

Information About Lawful Intercept	209
Prerequisites for Lawful Intercept	212
Install Lawful Intercept using vManage	213
Lawful Intercept MIBs	213
Restrict Access to Trusted Hosts (Without Encryption)	214
Restrict Trusted Mediation Device	214
Configure Lawful Intercept	215
Configure Lawful Intercept Using CLI	215
Encrypt Lawful Intercept Traffic	216
Configure Encryption in the Device	216
Configure Lawful Intercept Encryption using CLI	217
Verify Static Tunnel with Media Device Gateway	218

CHAPTER 13**Policy Applications Using CLIs 219**



CHAPTER 1

What's New for Cisco SD-WAN



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

This chapter describes what's new in Cisco SD-WAN for each release.

- [What's New for Cisco IOS XE SD-WAN Releases 16.12.1b, 16.12.1d, and 16.12.2r, on page 1](#)

What's New for Cisco IOS XE SD-WAN Releases 16.12.1b, 16.12.1d, and 16.12.2r

This section applies to Cisco IOS XE SD-WAN devices.

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

Table 1: What's New for Cisco IOS XE SD-WAN Devices

Feature	Description
Getting Started	
API Cross-Site Request Forgery Prevention	This feature adds protection against Cross-Site Request Forgery (CSRF) that occurs when using Cisco SD-WAN REST APIs. This protection is provided by including a CSRF token with API requests. You can put requests on an allowed list so that they do not require protection if needed. See Cross-Site Request Forgery Prevention .
Systems and Interfaces	

Feature	Description
IPv6 Support for NAT64 Devices	This feature supports NAT64 to facilitate communication between IPv4 and IPv6 on Cisco IOS XE SD-WAN devices. See IPv6 Support for NAT64 Devices .
Secure Shell Authentication Using RSA Keys	This feature helps configure RSA keys by securing communication between a client and a Cisco SD-WAN server. See SSH Authentication using vManage on Cisco XE SD-WAN Devices. See Configure SSH Authentication .
DHCP option support	This feature allows DHCP server options, 43 and 191 to configure vendor-specific information in client-server exchanges. See Configure DHCP .
Communication with an UCS-E Server	This feature allows you to connect a UCS-E interface with a UCS-E server through the interface feature template. See Create a UCS-E Template .
Bridging, Routing, Segmentation, and QoS	
QoS on Subinterface	This feature enables Quality of Service (QoS) policies to be applied to individual subinterfaces. See QoS on Subinterface .
Policies	
Packet Duplication for Noisy Channels	This feature helps mitigate packet loss over noisy channels, thereby maintaining high application QoE for voice and video. See Configure and Monitor Packet Duplication .
Control Traffic Flow Using Class of Service Values	This feature lets you control the flow of traffic into and out of a Cisco device's interface based on the conditions defined in the quality of service (QoS) map. A priority field and a layer 2 class of service (CoS) were added for configuring the re-write rule. See Configure Localized Data Policy for IPv4 Using Cisco vManage .
Integration with Cisco ACI	The Cisco SD-WAN and Cisco ACI integration functionality now supports predefined SLA cloud beds. It also supports dynamically generated mappings from a data prefix-list and includes a VPN list to an SLA class that is provided by Cisco ACI. See Integration with Cisco ACI .
Encryption of Lawful Intercept Messages	This feature encrypts lawful intercept messages between a Cisco IOS XE SD-WAN device and a media device using static tunnel information. See Encryption of Lawful Intercept Messages .
Security	
High-Speed Logging for Zone-Based Firewalls	This feature allows a firewall to log records with minimum impact to packet processing. See Firewall High-Speed Logging .
Self zone policy for Zone-Based Firewalls	This feature can help define policies to impose rules on incoming and outgoing traffic. See <i>Apply Policy to a Zone Pair</i> in Use the Policy Configuration Wizard .

Feature	Description
Secure Communication Using Pairwise IPsec Keys	This feature allows private pairwise IPsec session keys to be created and installed for secure communication between IPsec devices and its peers. See IPsec Pairwise Keys Overview .
Network Optimization and High Availability	
TCP Optimization	This feature optimizes TCP data traffic by decreasing any round-trip latency and improving throughput. See TCP Optimization: Cisco XE SD-WAN Routers .
Share VNF Devices Across Service Chains	This feature lets you share Virtual Network Function (VNF) devices across service chains to improve resource utilisation and reduce resource fragmentation. See Share VNF Devices Across Service Chains .
Monitor Service Chain Health	This feature lets you configure periodic checks on the service chain data path and reports the overall status. To enable service chain health monitoring, NFVIS version 3.12.1 or later should be installed on all CSP devices in a cluster. See Monitor Service Chain Health .
Manage PNF Devices in Service Chains	This feature lets you add Physical Network Function (PNF) devices to a network, in addition to the Virtual Network function (VNF) devices. These PNF devices can be added to service chains and shared across service chains, service groups, and a cluster. Inclusion of PNF devices in the service chain can overcome the performance and scaling issues caused by using only VNF devices in a service chain. See Manage PNF Devices in Service Chains .
Devices	
Cisco 1101 Series Integrated Services Routers	Cisco SD-WAN capability can now be enabled on Cisco 1101 Series Integrated Services Routers.
Commands	
Loopback interface support for WAN (IPsec)	This feature allows you to configure a loopback transport interface on a Cisco IOS XE SD-WAN device for troubleshooting and diagnostic purposes. See the bind command.



CHAPTER 2

Policy Basics

- [Policy Overview, on page 5](#)
- [Policies in Cisco vManage, on page 7](#)

Policy Overview

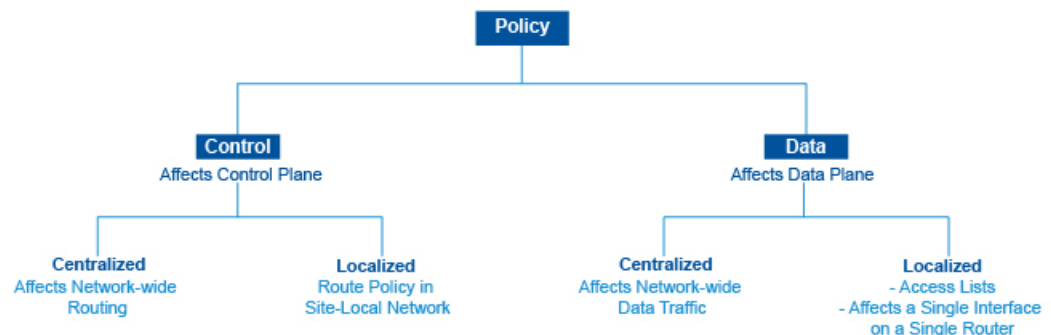
Policy influences the flow of data traffic and routing information among Cisco IOS XE SD-WAN devices in the overlay network. Policy comprises:

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

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

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

Figure 1: Policy Architecture



368464

Control and Data Policy

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

Centralized and Localized Policy

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

Centralized Policy

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

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

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



Note

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

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

Localized Policy

Localized policy refers to a policy that is provisioned locally through the CLI on the Cisco IOS XE SD-WAN devices, or through a Cisco vManage device template.

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

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

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

data traffic into and out of a Cisco IOS XE SD-WAN device's interfaces, enabling features such as class of service and queuing, and policing.

- Application-aware routing, which selects the best path for traffic based on real-time network and path performance characteristics.
- Cflowd, for monitoring traffic flow.

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

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

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

This section examines the structural components of routing and data policy in the Cisco SD-WAN overlay network.

Policies in Cisco vManage

Use the Policies screen to create and activate centralized and localized control and data policies for Cisco vSmart Controllers and Cisco IOS XE SD-WAN devices.

Figure 2: Policy Configuration

This screen allows you to perform several tasks related to Policies in Cisco vManage:

- View centralized or localized policies
- Copy, edit, or delete policies
- Create and edit policy components
- Activate and deactivate a centralized policy on Cisco vSmart controllers

Create and Manage Policies via Cisco vManage

View centralized or localized policies

To view centralized or localized policies, do the following:

1. From the **Centralized Policy** or **Localized Policy** tab, select a policy.
2. For a policy created using the UI policy builder or via CLI, click **More Actions** and click **View**. The policy created using the UI policy builder is displayed in graphical format while the policy created using the CLI method is displayed in text format.

3. For a policy created using the vManage policy configuration wizard, click **More Actions** and click **Preview**. This policy is displayed in text format.

Copy, edit and delete policies

1. To copy a policy:
 - a. From the **Centralized Policy** or **Localized Policy** tab, select a policy.
 - b. Click **More Actions** and click **Copy**.
 - c. In the Policy Copy popup window, enter the policy name and a description of the policy.



Note If you are upgrading to 18.4.4 version, Data Policy names need to be under 26 characters.

- d. Click **Copy**.
2. To edit policies created using the vManage policy configuration wizard:
 - a. Click **More Actions** and click **Edit**.
 - b. Edit the policy as needed.
 - c. Click **Save Policy Changes**.
3. To edit policies created using the CLI method:
 - a. In the **Custom Options** drop-down, click **CLI Policy**.
 - b. Click **More Actions** and click **Edit**.
 - c. Edit the policy as needed.
 - d. Click **Update**.
4. To delete policies:
 - a. From the **Centralized Policy** or **Localized Policy** tab, select a policy.
 - b. Click **More Actions** and click **Delete**.
 - c. Click **OK** to confirm deletion of the policy.

Edit or Create a Policy Component

You can create individual policy components directly and then use them or import them when you are using the policy configuration wizard:

1. In the Title bar, click the **Custom Options** drop-down.
2. For centralized policies, select the **Centralized Policy** tab and then select a policy component:
 - CLI policy—Create the policy using the command-line interface rather than the policy configuration wizard.

- Lists—Create groups of interest to import in the Group of Interest screen in the policy configuration wizard.
 - Topology—Create a hub-and-spoke, mesh, or custom topology or a VPN membership to import in the Topology screen in the policy configuration wizard.
 - Traffic Policy—Create an application-aware routing, traffic data, or cflowd policy to import in the Traffic Rules screen in the policy configuration wizard.
3. For localized policies, select the **Localized Policy** and then select a policy component:
 - CLI policy—Create the policy using the command-line interface rather than the policy configuration wizard.
 - Lists—Create groups of interest to import in the Group of Interest screen in the policy configuration wizard.
 - Forwarding Class/QoS—Create QoS mappings and rewrite rules to import in the Forwarding Classes/QoS screen in the policy configuration wizard.
 - Access Control Lists—Create ACLs of interest to import in the Configure Access Lists screen in the policy configuration wizard.
 - Route Policy—Create route policies to import in the Configure Route Policies screen in the policy configuration wizard.

Activate a Centralized Policy on Cisco vSmart Controllers

1. In the Title bar, click the **Custom Options** drop-down.
2. In the **Centralized Policy** tab, and then select a policy.
3. Click **More Actions** and click **Activate**.
4. In the **Activate Policy** popup, click **Activate** to push the policy to all reachable Cisco vSmart Controllers in the network.
5. Click **OK** to confirm activation of the policy on all Cisco vSmart Controllers.
6. To deactivate the centralized policy, select the = tab, and then select a policy.
7. 6. Click **More Actions** and click **Deactivate**.
8. In the **Deactivate Policy** popup, click **Deactivate** to confirm that you want to remove the policy from all reachable Cisco vSmart Controllers.



CHAPTER 3

Cisco SD-WAN Policy Framework Basics

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

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

- [Cisco vSmart Policy Components, on page 11](#)
- [Design Cisco vSmart Controller Policy Processing and Application, on page 17](#)
- [Cisco vSmart Policy Operation, on page 18](#)
- [Configure and Execute Cisco vSmart Policies, on page 22](#)

Cisco vSmart Policy Components

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

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

- Lists define the targets of policy application or matching.
- Policy definition, or policies, controls aspects of control and forwarding. There are different types of policy, including:
 - app-route-policy (for application-aware routing)
 - cflowd-template (for cflowd flow monitoring)
 - control-policy (for routing and control plane information)
 - data-policy (for data traffic)

- vpn-membership-policy (for limiting the scope of traffic to specific VPNs)
- Policy application controls what a policy is applied towards. Policy application is site-oriented, and is defined by a specific list called a site-list.

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

Table 2: The Three Building Blocks of Cisco vSmart Policy

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

Lists

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

Table 3: List Types

List type	Usage
data-prefix-list	Used in data-policy to define prefix and upper layer ports, either individually or jointly, for traffic matching.

List type	Usage
prefix-list	Used in control-policy to define prefixes for matching RIB entries.
site-list	Used in control-policy to match source sites, and in apply-policy to define sites for policy application.
tloc-list	Used in control-policy to define TLOCs for matching RIB entries and to apply redefined TLOCs to vRoutes.
vpn-list	Used in control-policy to define prefixes for matching RIB entries, and in data-policy and app-route-policy to define VPNs for policy application.

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

```

policy
  lists
    data-prefix-list appl
      ip-prefix 209.165.200.225/27 port 100
    !
    prefix-list pfx1
      ip-prefix 209.165.200.225/27
    !
    site-list site1
      site-id 100
    !
    tloc-list site1-tloc
      tloc 209.165.200.225 color mpls
    vpn-list vpn1
      vpn1
    !
  !

```

Policy Definition

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

Table 4: Policy Types

Policy type	Usage
policy-type	Can be control-policy , data-policy , or vpn-membership —dictates the type of policy. Each type has a particular syntax and a particular set of match conditions and settable actions.
vpn-list	Used by data-policy and app-route-policy to list the VPNs for which the policy is applicable.

Policy type	Usage
sequence	Defines each sequential step of the policy by sequence number.
match	Decides what entity to match on in the specific policy sequence.
action	Determines the action that corresponds to the preceding match statement.
default-action	Action to take for any entity that is not matched in any sequence of the policy. By default, the action is set to reject.

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

```

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

```

Policy Application

The following are the configuration components:

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

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

```

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

```

```
!
!
```

Policy Example

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

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

TLOC Attributes Used in Policies

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

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

Table 5:

TLOC Attribute	Function	Application Point Set By	Application Point Modify By
Address (IP address)	system-ip address of the source device on which the interface is located.	Configuration on source device	control-policy data-policy
Carrier	Identifier of the carrier type. It primarily indicates whether the transport is public or private.	Configuration on source device	control-policy

TLOC Attribute	Function	Application Point Set By	Application Point Modify By
Color	Identifier of the TLOC type.	Configuration on source device	control-policy data-policy
Domain ID	Identifier of the overlay network domain.	Configuration on source device	control-policy
Encapsulation	Tunnel encapsulation, either IPsec or GRE.	Configuration on source device	control-policy data-policy
Originator	system-ip address of originating node.	Configuration on any originator	control-policy
Preference	OMP path-selection preference. A higher value is a more preferred path.	Configuration on source device	control-policy
Site ID	Identification for a give site. A site can have multiple nodes or TLOCs.	Configuration on source device	control-policy
Tag	Identifier of TLOC on any arbitrary basis.	Configuration on source device	control-policy

vRoute Attributes Used in Policies

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

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

Table 6:

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

vRoute Attribute	Function	Application Point Set By	Application Point Modify By
Tag	Identification on any arbitrary basis.	Configuration on source device	control-policy
TLOC	TLOC used as next hop for the vRoute.	Configuration on source device or policy	control-policy data-policy
VPN	VPN to which the vRoute belongs.	Configuration on source device or policy	control-policy data-policy

Design Cisco vSmart Controller Policy Processing and Application

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

Policy is processed as follows:

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

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

- When applying policy to a site-list, you can apply only one of each type of policy. For example, you can have one control-policy and one data-policy, or one control-policy in and one control-policy out. You cannot have two data policies or two outbound control policies.
- Because a site-list is a grouping of many sites, you should be careful about including a site in more than one site-list. When the site-list includes a range of site identifiers, ensure that there is no overlap. If the same site is part of two site-lists and the same type of policy is applied to both site-lists, the policy behavior is unpredictable and possibly catastrophic.
- Control-policy is unidirectional, being applied either inbound to the vSmart controller or outbound from it. When control-policy is needed in both directions, configure two control policies.
- Data-policy is bidirectional and can be applied either to traffic received from the service side of the Cisco IOS XE SD-WAN device, traffic received from the tunnel side, or all of these combinations.
- VPN membership policy is always applied to traffic outbound from the Cisco vSmart Controller.
- Control-policy remains on the Cisco vSmart Controller and affects routes that the controller sends and receives.
- Data-policy is sent to either the Cisco IOS XE SD-WAN devices in the site-list. The policy is sent in OMP updates, and it affects the data traffic that the devices send and receive.

- When any node in the overlay network makes a routing decision, it uses any and all available routing information. In the overlay network, it is the Cisco vSmart Controller that distributes routing information to the Cisco IOS XE SD-WAN device nodes.
- In a network deployment that has two or more Cisco vSmart Controllers, each controller acts independently to disseminate routing information to other Cisco vSmart Controllers and to Cisco IOS XE SD-WAN devices in the overlay network. So, to ensure that the Cisco vSmart Controller policy has the desired effect in the overlay network, each Cisco vSmart Controller must be configured with the same policy, and the policy must be applied identically. For any given policy, you must configure the identical policy and apply it identically across all the Cisco vSmart Controllers.

Cisco vSmart Policy Operation

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

The basic Cisco vSmart policies are:

- Control Policy
- Data Policy
- VPN Membership

Control Policy

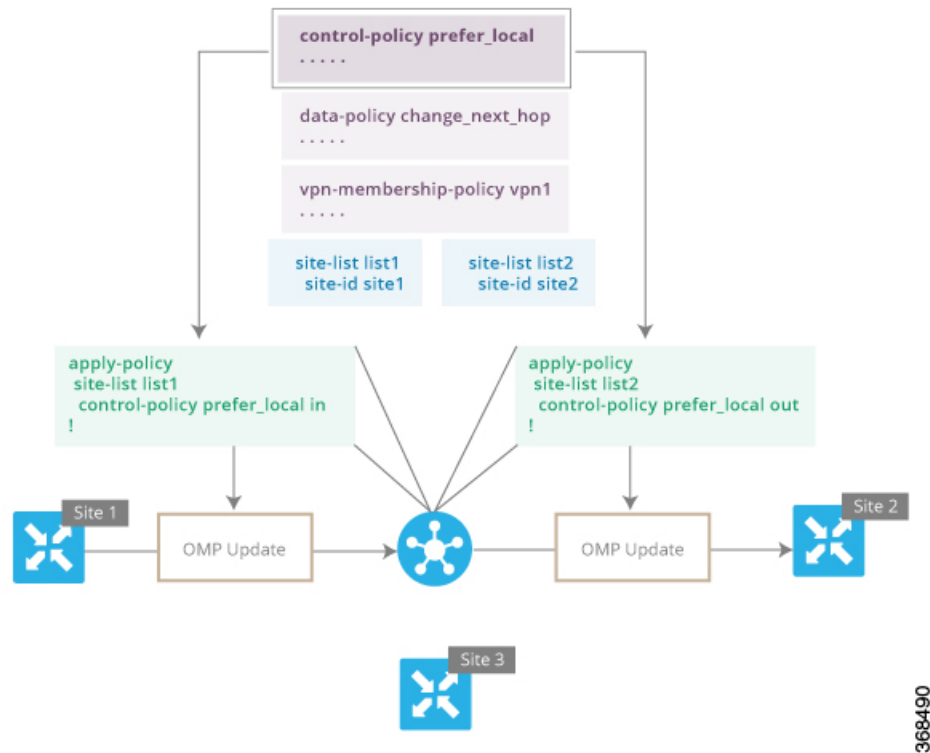
The Cisco IOS XE SD-WAN devices periodically exchange OMP updates, which carry routing information pertaining to the overlay network. Two of the things that these updates contain are vRoute attributes and Transport Locations (TLOC) attributes.

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

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

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

Figure 3: Control Policy Topology



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

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

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

```
apply-policy
site-list list2
control-policy prefer_local out
!
```

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

route table is placed in to an OMP update and before the devices receive the update. Again, note that the direction is outbound from the perspective of the Cisco vSmart Controller.

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

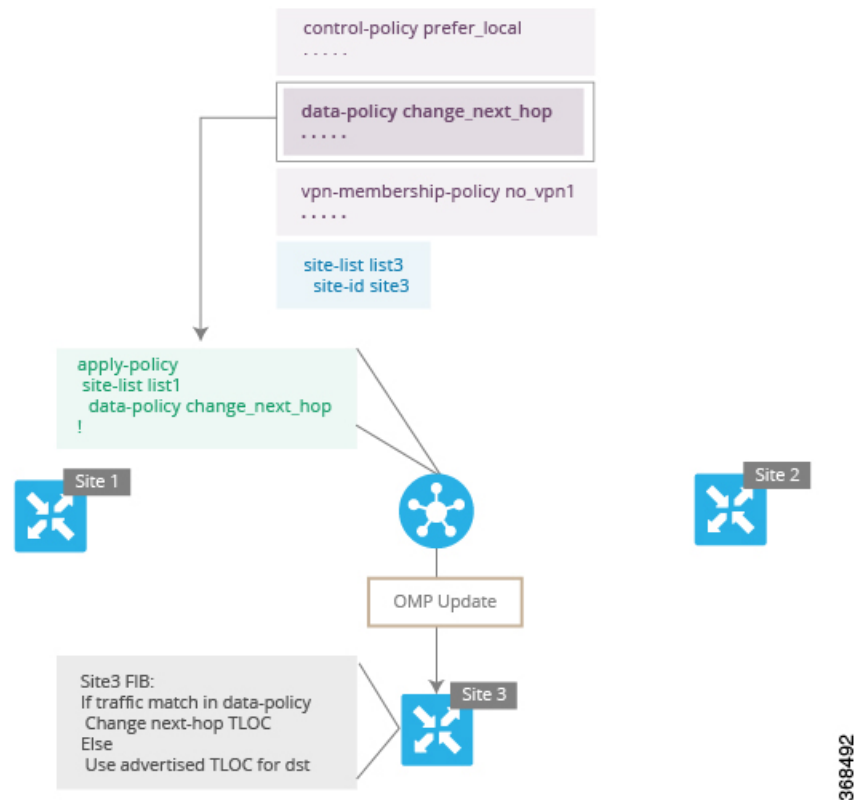
The same control policy (the **prefer_local** policy) is applied to both the inbound and outbound OMP updates. However, the effects of applying the same policy to inbound and outbound are different. The usage shown in the figure illustrates the flexibility of the Cisco IOS XE SD-WAN control policy design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

Data Policy

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

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

Figure 4: Data Policy Topology



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

VPN Membership Policy Operation

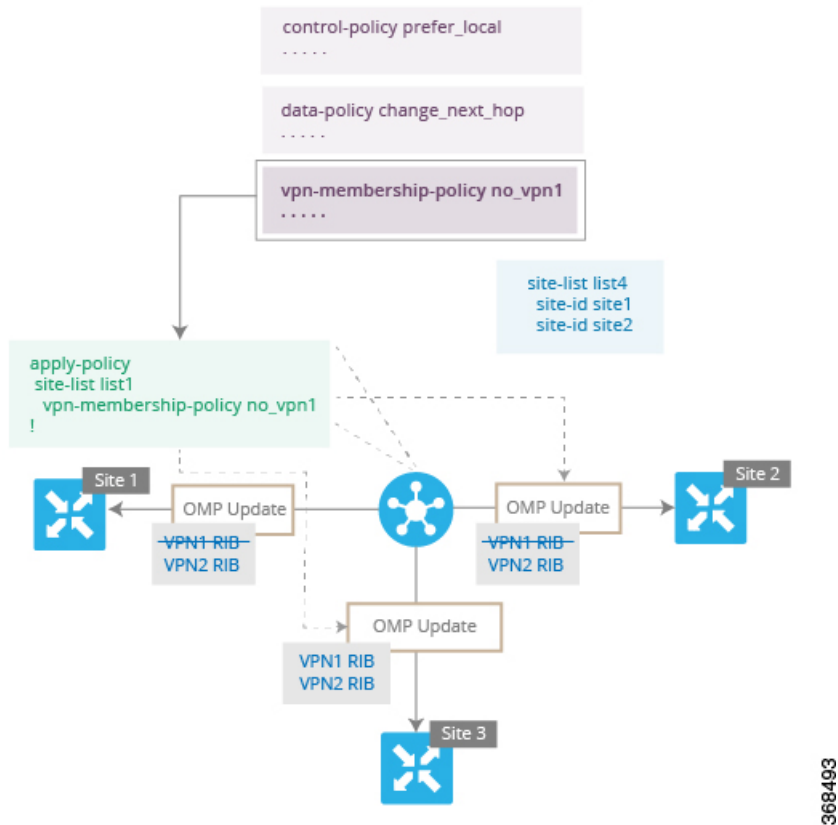
VPN membership policy, as the name implies, affects the VPN route tables that are distributed to particular Cisco IOS XE SD-WAN devices. In an overlay network with no VPN membership policy, the Cisco vSmart Controller pushes the routes for all VPNs to all the devices. If your business usage model restricts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

The figure VPN Membership Topology illustrates how VPN membership policy works. This topology has three Cisco IOS XE SD-WAN devices:

- The Cisco IOS XE SD-WAN devices at Sites 1 and 2 service only VPN 2.
- The Cisco IOS XE SD-WAN devices at Site 3 services both VPN 1 and VPN 2.

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

Figure 5: VPN Membership Topology




368493


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

Configure and Execute Cisco vSmart Policies

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

Figure 6: Cisco vSmart Policy

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

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

368503

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

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



CHAPTER 4

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 7: Control Policy Topology



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

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

Cisco 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 26
- [Localized Control Policy](#), on page 61
- [Device Access Policy](#), on page 78

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 8: 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.

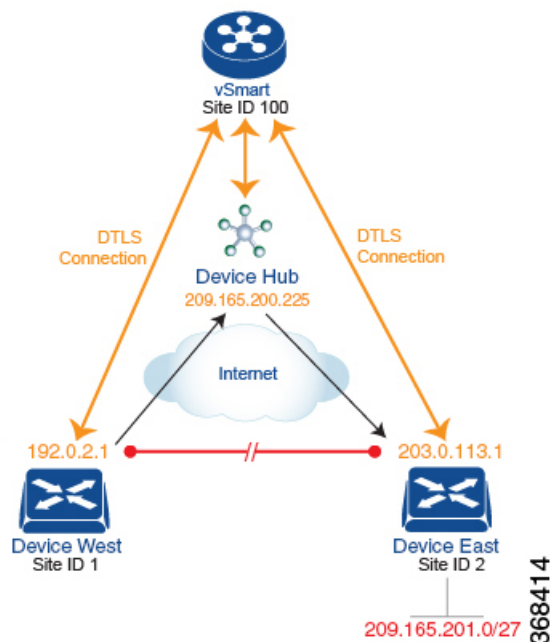
Examples of Modifying Traffic Flow with Centralized Control Policy

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 9: 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.

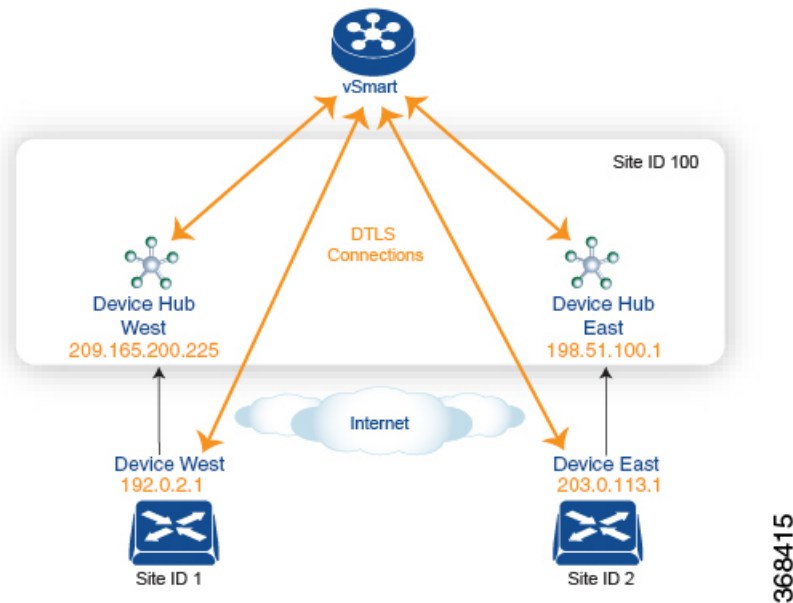
The figure illustrates how such a policy might work. The topology has two branch locations, West and East. When no control policy is provisioned, these two devices exchange data traffic with each other directly by creating an IPsec tunnel between them (shown by the red line). Here, the route table on 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 10: 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 11: Traffic Engineering 2



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

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

Figure 12: 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:
- Click **Add Hub Sites**.
 - In the **Site List Field**, select a site list for the policy component.
 - Click **Add**.
 - 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:
- Click **Add Spoke Sites**.
 - In the **Site List Field**, select a site list for the policy component.
 - Click **Add**.
 - 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 Name** 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: <ul style="list-style-type: none"> 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: <ul style="list-style-type: none"> • 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: <ol style="list-style-type: none"> 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: <ol style="list-style-type: none"> 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	<p>a. From the Match Conditions > VPN list field, select a VPN list, or click New VPN List to create a new one:</p> <p>b. Enter a name for the VPN List.</p> <p>c. In the VPN field, enter the VPN numbers, for example, 100 or 200 separated by commas, or 1000-2000 by range.</p> <p>d. Click Save.</p>
Prefix List	<p>From the Match Conditions > Prefix List field, select a Prefix list, or click New Prefix List to create a new one:</p> <p>a. From the Prefix List drop-down menu, select a prefix list, or create a new one.</p> <p>b. In the Add Prefix field, enter the IP prefixes, or click Import on the right to import prefixes.</p> <p>c. Click Save.</p> <p>Note The Prefix List option is not available if you select protocol Both (IPv4 and IPv6).</p>

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.
	<p>Service— Enter the following information:</p> <p>Type—Select a service type. Options are:</p> <ul style="list-style-type: none"> • Firewall • Intrusion Detection Prevention • Intrusion Detection System • Net Service 1 • Net Service 2 • Net Service 3 • Net Service 4 • Net Service 5 <p>VPN—Enter the number of the Service VPN.</p> <p>TLOC IP—Enter the IP address of the Service TLOC.</p> <p>Color—Select a Color type from the drop-down list.</p> <p>Encapsulation—Select IPSEC or GRE as the encapsulation type.</p> <p>TLOC List—Select a service TLOC list from the drop-down menu, or create a new one.</p>

Match Condition	Description
	<p>TLOC Action</p> <p>Select an action from the drop-down menu:</p> <ul style="list-style-type: none"> • 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.
	<p>TLOC—Enter the following information:</p> <ul style="list-style-type: none"> • 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.
	<p>Encapsulation—Select IPSEC or GRE as the encapsulation type.</p>
Reject	<p>Reject traffic for the selected conditions.</p> <ol style="list-style-type: none"> Select a protocol from the Protocol dropdown: IPv4, IPv6, or Both. Click Accept or Reject for the match conditions. 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: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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.

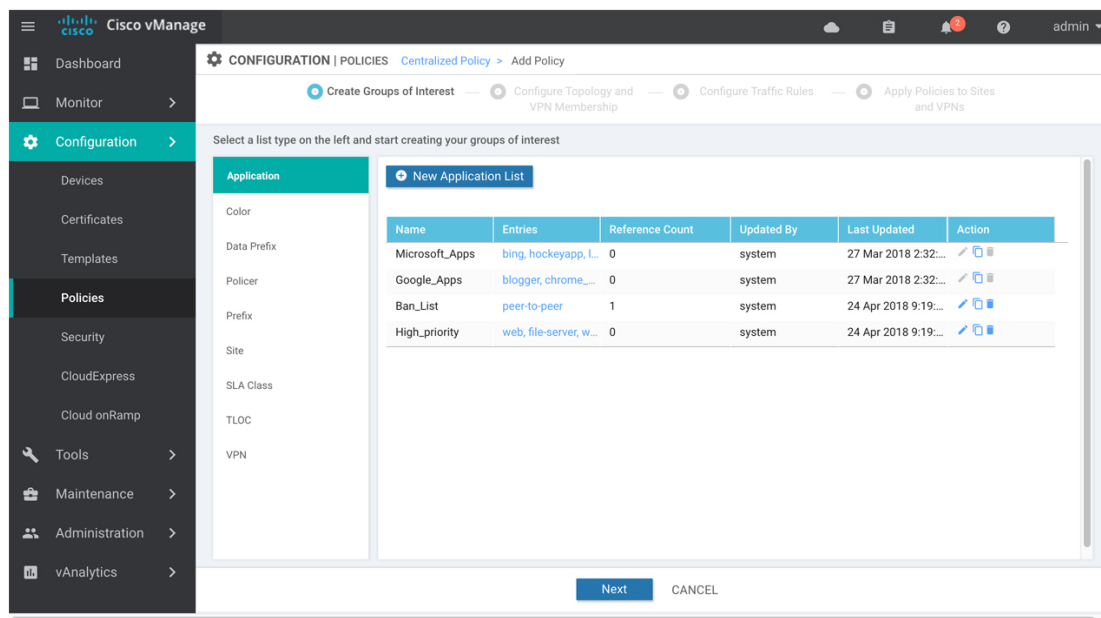
Step 1: Start the Policy Configuration Wizard

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 7:

List Type	Procedure
Color	<ol style="list-style-type: none"> In the left bar, click Color. Click New Color List. Enter a name for the list. From the Select Color drop-down, select the desired colors. Click Add.
Prefix	<ol style="list-style-type: none"> In the left bar, click Prefix. Click New Prefix List. Enter a name for the list. In the Add Prefix field, enter one or more data prefixes separated by commas. Click Add.

List Type	Procedure
Site	<ol style="list-style-type: none"> 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	<ol style="list-style-type: none"> 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	<ol style="list-style-type: none"> 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 Type** 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. <i>color</i> 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 <i>list-name</i> color <i>color</i>

List Type	Description	vManage Configuration/ CLI Configuration Command
Prefixes	<p>List of one or more IP prefixes. Specify the IP prefixes as follows:</p> <ul style="list-style-type: none"> • <i>prefix/length</i>—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 <i>length</i>. 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 <i>length</i>. 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 <i>length1</i> and less than or equal to <i>length2</i>. 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 <i>length1</i> and <i>length2</i> 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. 	<p>Configuration > Policies > Centralized Policy > Add Policy > Create Groups of Interest > Prefix</p> <p>Configuration > Policies > Custom Options > Centralized Policy > Lists > Prefix</p> <p>prefix-list list-name ip-prefix prefix/length</p>
Sites	<p>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.</p>	<p>Configuration > Policies > Centralized Policy > Add Policy > Create Groups of Interest > Site</p> <p>Configuration > Policies > Custom Options > Centralized Policy > Lists > Site</p> <p>site-list list-name site-id site-id</p>
TLOCs	<p>List of one or more TLOCs in the overlay network.</p> <p>For each TLOC, specify its address, color, and encapsulation. <i>address</i> is the system IP address. <i>color</i> 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. <i>encapsulation</i> 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.</p>	<p>Configuration > Policies > Centralized Policy > Add Policy > Create Groups of Interest > TLOC</p> <p>Configuration > Policies > Custom Options > Centralized Policy > Lists > TLOC</p> <p>tloc-list list-name tloc ip-address color color encap (gre ipsec) [preference number]</p>

List Type	Description	vManage Configuration/ CLI Configuration Command
VPNs	<p>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.</p> <p>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).</p>	<p>Configuration > Policies > Centralized Policy > Add Policy > Create Groups of Interest > VPN</p> <p>Configuration > Policies > Custom Options > Centralized Policy > Lists > VPN</p> <p>vpn-list list-name vpn vpn-id</p>

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 <i>color</i>	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 <i>list-name</i>	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 <i>number</i>	0 through 4294967295
Protocol from which the route was learned.	Match Origin origin <i>protocol</i>	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 <i>ip-address</i>	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 <i>number</i>	0 through 255
One or more prefixes.	Match Prefix List prefix-list <i>list-name</i>	Name of a prefix list or a policy lists prefix-list list.
Individual site identifier.	Not available in Cisco vManage. site-id <i>site-id</i>	0 through 4294967295
One or more overlay network site identifiers.	Match Site site-list <i>list-name</i>	Name of a site or a policy lists site-list list.
Individual TLOC address.	Match TLOC tloc <i>ip-address</i>	IP address
One or more TLOC addresses.	Match TLOC tloc-list <i>list-name</i>	Name of a TLOC or a policy lists tloc-list list.
Individual VPN identifier.	Match VPN vpn <i>vpn-id</i>	0 through 65535

Description	vManage Configuration/ CLI Configuration Command	Value or Range
One or more VPN identifiers.	Match VPN vpn-list <i>list-name</i>	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 <i>carrier-name</i>	default, carrier1 through carrier8
Individual color.	Not available in the Cisco vManage NMS. color <i>color</i>	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 <i>list-name</i>	See the colors above.
Domain identifier associated with a TLOC.	Match Domain ID domain-id <i>domain-id</i>	0 through 4294967295
Tag value associated with the TLOC route in the route table on the device.	Match OMP Tag omp-tag <i>number</i>	0 through 4294967295
IP address from which the route was learned.	Match Originator originator <i>ip-address</i>	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 <i>number</i>	0 through 255
Individual site identifier.	Match Site site-id <i>site-id</i>	0 through 4294967295
One or more overlay network site identifiers.	Match Site site-list <i>list-name</i>	Name of a policy lists site-list list.

Description	vManage Configuration/ CLI Configuration Command	Value or Range
Individual TLOC address.	Match TLOC tloc <i>address</i>	IP address
One or more TLOC addresses.	Match TLOC tloc-list <i>list-name</i>	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 (<i>vpn vpn-id vpn-list vpn-list</i>)	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
<p>Specify a service to redirect traffic to before delivering the traffic to its destination.</p> <p>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.</p> <p>The VPN identifier is where the service is located.</p> <p>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.</p>	Click Accept , then action Service . set service service-name (tloc ip-address tloc-list list-name) [vpn vpn-id]	<p>Standard services: FW, IDS, IDP</p> <p>Custom services: netsvc1, netsvc2, netsvc3, netsvc4</p> <p>TLOC list configured with a policy lists tloc-list command.</p>
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 . Encapsulation can be either gre or ipsec .

Description	vManage Configuration/ CLI Configuration Command	Value or Range
<p>Direct matching routes or TLOCs using the mechanism specified by <i>action</i>, 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.</p> <p>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.</p>	<p>Click Accept, then action TLOC Action.</p> <p>set tloc-action <i>action</i></p>	<p>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.</p> <p>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.</p> <p>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.</p> <p>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 strict action in a centralized control policy, strict is the default behavior.</p>
<p>Change the TLOC address and color to those in the specified TLOC list.</p>	<p>Click Accept, then action TLOC.</p> <p>set tloc-list <i>list-name</i></p>	<p>Name of a policy lists tloc-list list.</p>

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** that 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
                                encaps 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 in 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

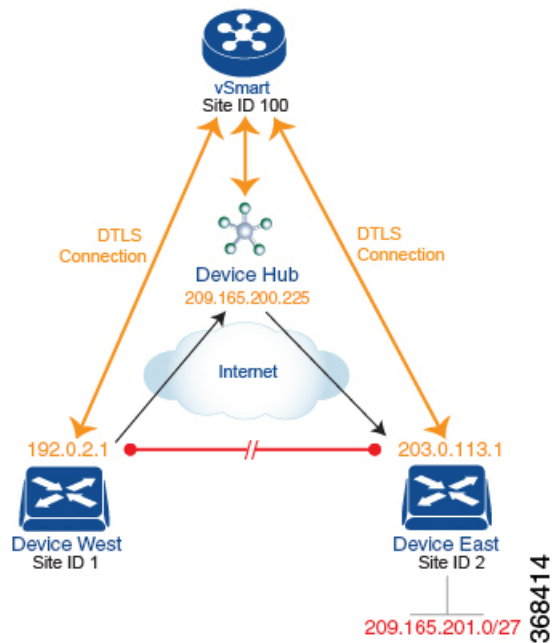
This topic provides some straightforward examples of configuring centralized control policy to help you understand the configuration procedure and get an idea of how to use policy to influence traffic flow across the Cisco IOS XE SD-WAN overlay network domain.

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 209.165.201.0/27, 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 encaps 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:

```

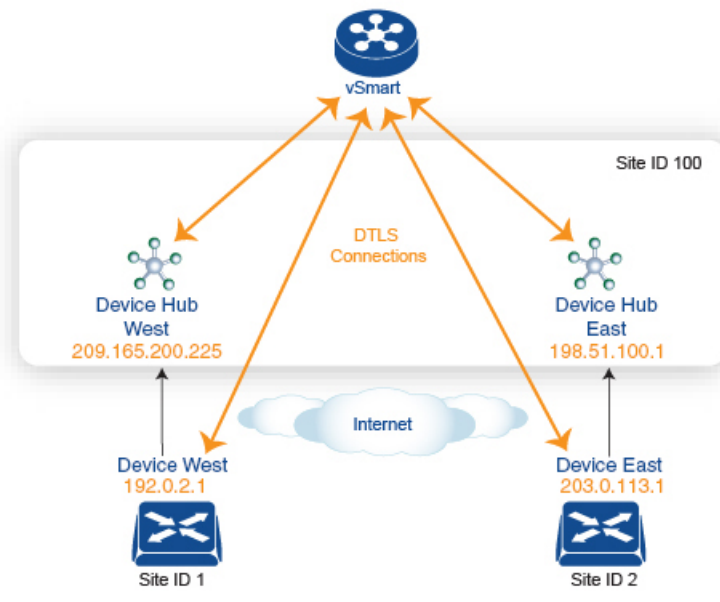
policy
  lists
    prefix-list east-prefixes
      ip-prefix 209.165.201.0/27
    site-list west-sites
      site-id 1
      site-id 11
      site-id 12
      site-id 13
  control-policy change-tloc
    sequence 10
    match route
      prefix-list east-prefixes
      site-id 2
    action accept
      set tloc 209.165.200.225 color gold encaps 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 100 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 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.

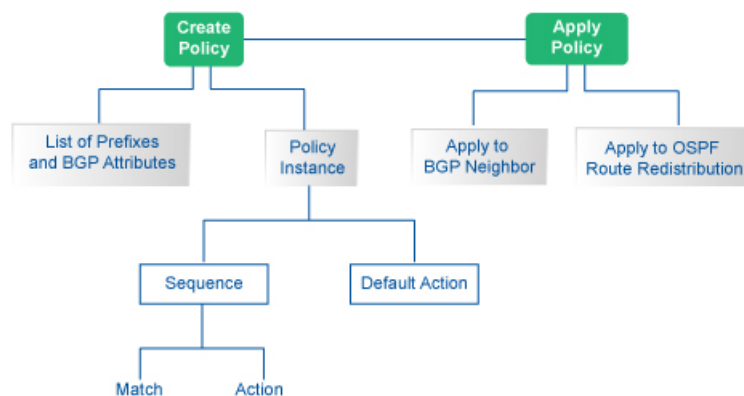
Localized control policy, which you configure on the devices, lets you affect routing policy on the network at the local site where the device is located. This type of control policy is called route policy.

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

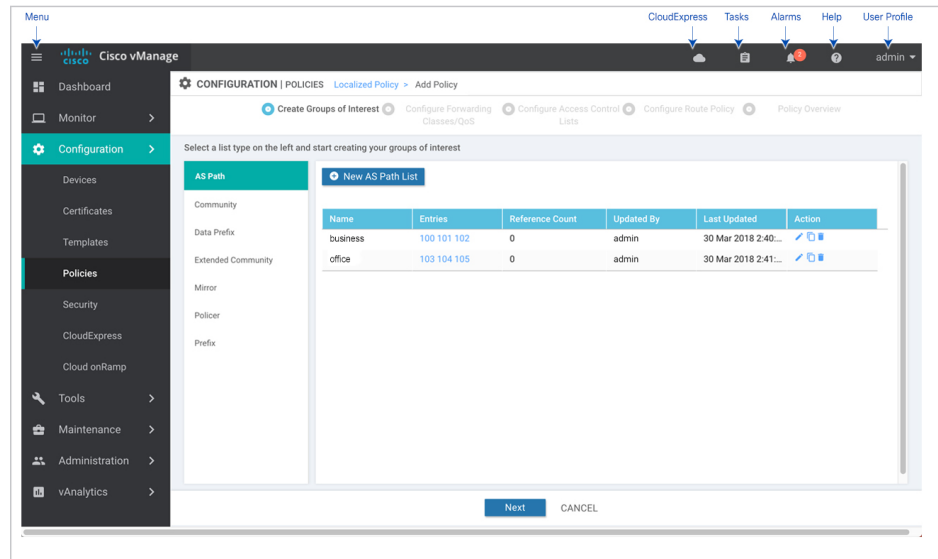
To start the policy configuration wizard:

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 8:

List Type	Procedure
AS Path	<ol style="list-style-type: none"> 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	<ol style="list-style-type: none"> 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 <i>aa:nn</i> or as the string internet, local-as, no-advertise, or no-export, separating multiple items with a comma. For <i>aa</i>, enter a 2-byte AS number, and for <i>nn</i>, enter a 2-byte network number. 5. Click Add.
Extended Community	<ol style="list-style-type: none"> 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 (<i>aa:nn ip-address</i>), for a route target community, or soo (<i>aa:nn ip-address</i>), for a route origin community, separating multiple items with a comma. For <i>aa</i>, enter a 2-byte AS number, and for <i>nn</i> enter a 2-byte network number. 5. Click Add.
Mirror	<ol style="list-style-type: none"> 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	<ol style="list-style-type: none"> 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	<ol style="list-style-type: none"> 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: <ul style="list-style-type: none"> • <i>prefix/length</i>—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 <i>length</i>. 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 <i>length</i>. 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 <i>length1</i> and less than or equal to <i>length2</i>. 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 <i>length1</i> and <i>length2</i> 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**.

- 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:

- 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
```

- 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)]
```

- Create a route policy instance:

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

- 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).

- Define match parameters for routes:

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

- Define actions to take when a match occurs:

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

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

- 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
```

-

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 9:

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	<p>List of one or more BGP communities. In community, you can specify:</p> <ul style="list-style-type: none"> • 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. 	<p>Configuration > Policies > Localized Policy > Add Policy > Create Groups of Interest > Community</p> <p>Configuration > Policies > Custom Options > Localized Policy > Lists > Community</p> <p>community-list <i>list-name</i> community [<i>aa:nn</i> internet local-as no-advertise no-export]</p>
Extended communities	<p>List of one or more BGP extended communities. In community, you can specify:</p> <ul style="list-style-type: none"> • rt (<i>aa:nn</i> <i>ip-address</i>): 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 (<i>aa:nn</i> <i>ip-address</i>): 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. 	<p>Configuration > Policies > Localized Policy > Add Policy > Create Groups of Interest > Extended Community</p> <p>Configuration > Policies > Custom Options > Localized Policy > Lists > Extended Community</p> <p>ext-community-list <i>list-name</i> community [rt (<i>aa:nn</i> <i>ip-address</i>) soo (<i>aa:nn</i> <i>ip-address</i>)]</p>

List Type	Description	vManage Configuration/ CLI Configuration Command
Prefixes	<p>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:</p> <ul style="list-style-type: none"> • <i>prefix/length</i>—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 <i>length</i>. 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 <i>length</i>. For example, ip-prefix 0.0.0.0/0 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 <i>length1</i> and less than or equal to <i>length2</i>. 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 <i>length1</i> and <i>length2</i> 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. 	<p>Configuration > Policies > Localized Policy > Add Policy > Create Groups of Interest > Prefix</p> <p>Configuration > Policies > Custom Options > Localized Policy > Lists > Prefix</p> <p>prefix-list list-name ip-prefix prefix/length</p>

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 10:

Description	vManage Configuration/ CLI Configuration Command	Value or Range
IP prefix or prefixes from which the route was learned	Match Address address <i>list-name</i>	Name of an IP prefix list
BGP AS paths	Match AS Path List as-path <i>list-name</i>	Name of an AS path list
BGP communities	Match Community List community <i>list-name</i>	Name of a BGP community list
BGP extended communities	Match Extended Community List ext-community <i>list-name</i>	Name of a BGP extended community list
BGP local preference	Match BGP Local Preference local-preference <i>number</i>	0 through 4294967295
Route metric	Match Metric metric <i>number</i>	0 through 4294967295
Next hop	Match Next Hop next-hop <i>list-name</i>	Name of an IP prefix list
OMP tag for OSPF	Match OMP Tag omp-tag <i>number</i>	0 through 4294967295
BGP origin code	Match Origin origin <i>origin</i>	egp (default), igp , incomplete
OSPF tag value	Match OSPF Tag ospf-tag <i>number</i>	0 through 4294967295
Peer address	Match Peer peer <i>address</i>	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 11:

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 12:

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	[<i>aa:nn</i> 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 <i>number</i>	0 through 4294967295
Set the metric type.	Click Accept, then action Metric Type set metric-type <i>type</i>	type1, type2
Set the next-hop address.	Click Accept, then action Next Hop set next-hop <i>ip-address</i>	IP address
Set the OMP tag for OSPF to use.	Click Accept, then action OMP Tag set omp-tag <i>number</i>	0 through 4294967295
Set the BGP origin code.	Click Accept, then action Origin set origin <i>origin</i>	egp, igp (default), incomplete
Set the IP address from which the route was learned.	Click Accept, then action Originator set originator <i>ip-address</i>	IP address
Set the OSPF tag value.	Click Accept, then action OSPF Tag set ospf-tag <i>number</i>	0 through 4294967295
Set the BGP weight.	Click Accept, then action Weight set weight <i>number</i>	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 13:

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 14:

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 <i>as-number ip-address</i>	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) <i>as-number</i>	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 <i>value</i>	[<i>aa:nn</i> internet local-as no-advertise no-export]
Set the BGP local preference.	Click Accept, then action Local Preference set local-preference <i>number</i>	0 through 4294967295
Set the metric value.	Click Accept, then action Metric set metric <i>number</i>	0 through 4294967295
Set the metric type.	Click Accept, then action Metric Type set metric-type <i>type</i>	type1, type2
Set the next-hop address.	Click Accept, then action Next Hop set next-hop <i>ip-address</i>	IP address
Set the OMP tag for OSPF to use.	Click Accept, then action OMP Tag set omp-tag <i>number</i>	0 through 4294967295

Description	vManage Configuration/ CLI Configuration Command	Value or Range
Set the BGP origin code.	Click Accept, then action Origin set origin <i>origin</i>	egp, igp (default), incomplete
Set the IP address from which the route was learned.	Click Accept, then action Originator set originator <i>ip-address</i>	IP address
Set the OSPF tag value.	Click Accept, then action OSPF Tag set ospf-tag <i>number</i>	0 through 4294967295
Set the BGP weight.	Click Accept, then action Weight set weight <i>number</i>	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 15: 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-WAN Cisco 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 following example shows the policy details of a counter.

show policy device-access-policy-counters

NAME	COUNTER		
	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 access-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;
    }
}
}


```




CHAPTER 5

Data Policies

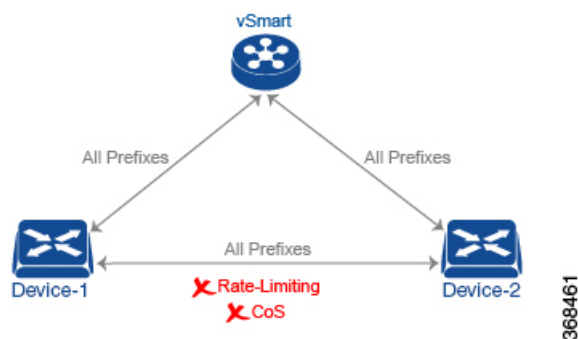
Data policy influences the flow of data traffic traversing the network based either on fields in the IP header of packets or the router interface on which the traffic is being transmitted or received. Data traffic travels over the IPsec connections between Cisco IOS XE SD-WAN devices, shown in purple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



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

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

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



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

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

- [Centralized Data Policy, on page 88](#)
- [Localized Data Policy, on page 122](#)

Centralized Data Policy

Centralized data policy is policy that is configured on a Cisco vSmart Controller (hence, it is centralized) and that affects data traffic being transmitted between the routers on the Cisco SD-WAN overlay network.

Centralized Data Policy Overview

Data policy operates on the data plane in the Cisco IOS XE SD-WAN overlay network and affects how data traffic is sent among Cisco IOS XE SD-WAN devices in the network. The Cisco IOS XE SD-WAN architecture defines two types of data policy, centralized data policy, which controls the flow of data traffic based on the IP header fields in the data packets and based on network segmentation, and localized data policy, which controls the flow of data traffic into and out of interfaces and interface queues on the devices.

Centralized data policy is applied to packets that originate from a specific sender, or source address, for instance, from a workstation in a local site that is sending voice, data, or other traffic, and it controls which destinations within a VPN the traffic can reach. Data policy is applied to data traffic based on a 6-tuple of fields in the packet's IP header: source IP address, source port, destination IP address, destination port, DSCP, and protocol.

As with control policy, data policy is provisioned centrally on a Cisco vSmart Controller and is applied only on the Cisco vSmart Controller controller. The data policy itself is never pushed to the devices in the network. What is pushed to the Cisco IOS XE SD-WAN devices, via OMP and based on the site ID, are the results of the data policy; hence, the effects of the policy are reflected on the devices. Normally, the data policy on a Cisco IOS XE SD-WAN device acts as the data policy for the entire site that sits behind the device. Data policy that comes from the Cisco vSmart Controller is always implicitly applied in the inbound direction.

Data policy can be applied to data traffic based on the packet header fields, such as the prefix, port, protocol, and DSCP value, and they can also be applied based on the VPN in the overlay network to which the traffic flows.

Data Policy Based on Packet Header Fields

Policy decisions affecting data traffic can be based on the packet header fields, specifically, on the source and destination IP prefixes, the source and destination IP ports, the protocol, and the DSCP.

This type of policy is often used to modify traffic flow in the network. Here are some examples of the types of control that can be effected with centralized data policy:

- Which set of sources are allowed to send traffic to any destination outside the local site. For example, local sources that are rejected by such a data policy can communicate only with hosts on the local network.
- Which set of sources are allowed to send traffic to a specific set of destinations outside the local site. For example, local sources that match this type of data policy can send voice traffic over one path and data traffic over another.
- Which source addresses and source ports are allowed to send traffic to any destination outside the local site or to a specific port at a specific destination.

Deep Packet Inspection

In addition to examining the network- and transport-layer headers in data packets, centralized data policy can be used to examine the application information in the data packets' payload. This deep packet inspection offers control over how data packets from specific applications or application families are forwarded across the network, allowing you to assign the traffic to be carried by specific tunnels. To control the traffic flow of specific application traffic based on the traffic loss or latency properties on a tunnel, use application-aware routing.

To base policy decisions on source and destination prefixes and on the headers in the IP data packets, you use centralized data policy, which you configure with the **policy data-policy** command. The Cisco vSmart Controller pushes this type of data policy to the Cisco IOS XE SD-WAN devices. In domains with multiple Cisco vSmart Controllers, all the controllers must have the same centralized data policy configuration to ensure that traffic flow within the overlay network remains synchronized.

To base policy decisions on the application information in the packet payload, you use centralized data policy to perform deep packet inspection. You configure this by creating lists of applications with the **policy lists app-list** command and then calling these lists in a **policy data-policy** command.

To configure the VPNs that Cisco IOS XE SD-WAN devices are allowed to receive routes from, you use centralized data policy, which you configure with the **policy vpn-membership** command. VPN membership policy affects which routes the Cisco vSmart Controller sends to the devices. The policy itself remains on the Cisco vSmart Controller and is not pushed to the devices.

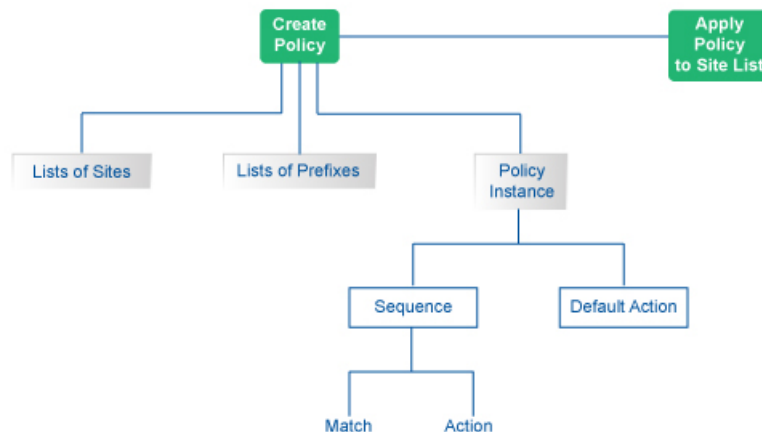
Configure Centralized Data Policy Based on Prefixes and IP Headers

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

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

Configuration Components

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



To configure centralized data 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 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 data policy to take effect, you must activate the policy.

This section provides general procedures for configuring centralized data policy on Cisco vSmart Controllers. Centralized data policy can be used for different purposes, which are described in the sections that follow.

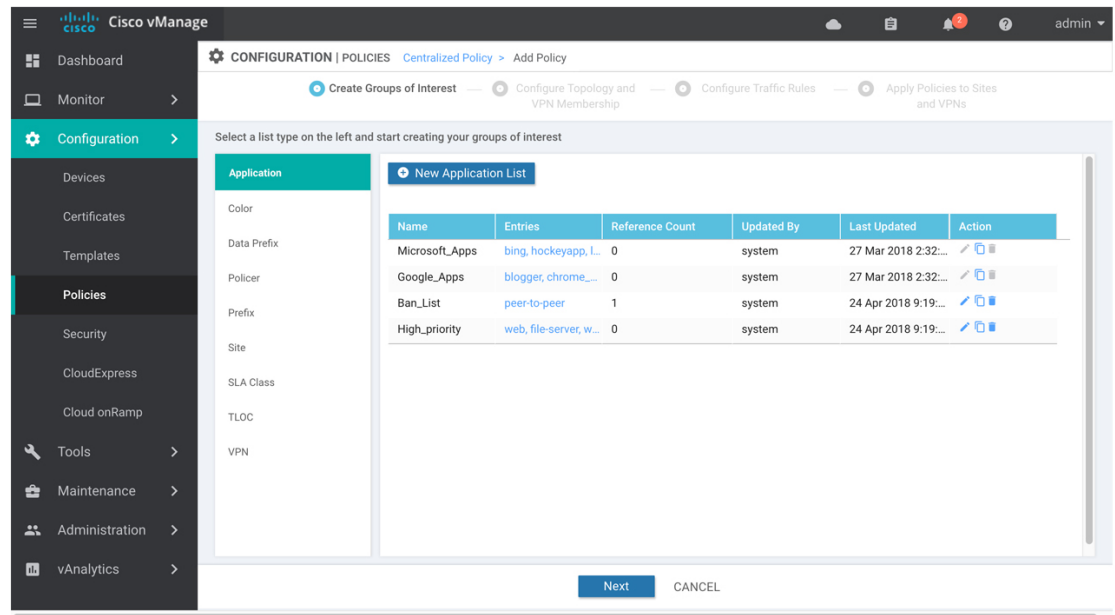
Start the Policy Configuration Wizard

To start the policy configuration wizard:

-
- Step 1** In the Cisco vManage NMS, select the **Configure > Policies** screen.
 - Step 2** Select the **Centralized Policy** tab.
 - Step 3** Click **Add Policy**. The policy configuration wizard opens, and the Create Groups of Interest screen displays.
-

Step 1: Create Policy Lists

You can create lists of groups to use in centralized policy.



368879

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

List Type	Procedure
Application	<ol style="list-style-type: none"> In the left bar, click Application. Click New Application List. Enter a name for the list. Click either the Application or Application Family button. From the Select drop-down, select the desired applications or application families. Click Add. <p>Two application lists are preconfigured. You cannot edit or delete these lists.</p> <ul style="list-style-type: none"> Google_Apps—Includes Google applications, such as gmail, Google maps, and YouTube. To display a full list of Google applications, click the list in the Entries column. Microsoft_Apps—Includes Microsoft applications, such as Excel, Skype, and Xbox. To display a full list of Microsoft applications, click the list in the Entries column.

List Type	Procedure
Data Prefix	<ol style="list-style-type: none"> a. In the left bar, click Data Prefix. b. Click New Data Prefix List. c. Enter a name for the list. d. Select either IPv4 or IPv6. e. In the Add Data Prefix field, enter one or more data prefixes separated by commas. f. Click Add.
Policer	<ol style="list-style-type: none"> a. In the left bar, click Policer. b. Click New Policer List. c. Enter a name for the list. d. Define the policing parameters: <ol style="list-style-type: none"> 1. In the Burst field, enter the maximum traffic burst size, a value from 15,000 to 10,000,000 bytes. 2. In the Exceed field, select the action to take when the burst size or traffic rate is exceeded. It can be drop, which sets the packet loss priority (PLP) to low. 3. In the Rate field, enter the maximum traffic rate, a value from 0 through $2^{64} - 1$ bits per second (bps). e. Click Add.
Prefix	<ol style="list-style-type: none"> 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.
Site	<ol style="list-style-type: none"> 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.

List Type	Procedure
SLA Class	<ol style="list-style-type: none"> a. In the left bar, click SLA Class. b. Click New SLA Class List. c. Enter a name for the list. d. Define the SLA class parameters: <ol style="list-style-type: none"> 1. In the Loss field, enter the maximum packet loss on the connection, a value from 0 through 100 percent. 2. In the Latency field, enter the maximum packet latency on the connection, a value from 0 through 1,000 milliseconds. 3. In the Jitter field, enter the maximum jitter on the connection, a value from 1 through 1,000 milliseconds. e. Click Add.
TLOC	<ol style="list-style-type: none"> 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	<ol style="list-style-type: none"> 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.

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

Step 2: Configure Traffic Rules

When you first open the Traffic Rules screen, the Application-Aware Routing tab is selected by default. To configure traffic rules for deep packet inspection, see [Deep Packet Inspection, on page 111](#).

To configure traffic rules for centralized data policy:

- Step 1** Click the **Traffic Data** tab.
- Step 2** Click the **Add Policy** drop-down.
- Step 3** Click **Create New**. The Add Data Policy screen displays.
- Step 4** Enter a name and description for the data policy.
- Step 5** In the right pane, click **Sequence Type**. The Add Data Policy popup opens.
- Step 6** Select the type of data policy you want to create. Choices are: **Application Firewall**, **QoS**, **Traffic Engineering**, and **Custom**.
- Step 7** A policy sequence containing the text string **Application Firewall**, **QoS**, **Traffic Engineering**, or **Custom** is added in the left pane
- Step 8** Double-click the text string, and enter a name for the policy sequence. The name you type is displayed both in the Sequence Type list in the left pane and in the right pane.
- Step 9** In the right pane, click **Sequence Rule**. The Match/Action box opens, and Match is selected by default. The available policy match conditions are listed below the box.
- Step 10** For QoS and Traffic Engineering data policies: From the **Protocol** drop-down list, select **IPv4** to apply the policy only to IPv4 address families, **IPv6** to apply the policy only to IPv6 address families, or **Both** to apply the policy IPv4 and IPv6 address families.
- Step 11** To select one or more Match conditions, click its box and set the values as described in the following table. Note that not all match conditions are available for all policy sequence types.

Match Condition	Procedure	IPv4 Fields	IPv6 Fields
None (match all packets)	Do not specify any match conditions.		
Applications /Application Family List	<ol style="list-style-type: none"> a. In the Match conditions, click Applications/Application Family List. b. In the drop-down, select the application family. c. To create an application list: <ol style="list-style-type: none"> 1. Click New Application List. 2. Enter a name for the list. 3. Click Application to create a list of individual applications. Click Application Family to create a list of related applications. 4. In the Select Application drop-down, select the desired applications or application families. 5. Click Save. 	app-list	

Match Condition	Procedure	IPv4 Fields	IPv6 Fields
Destination Data Prefix	<p>a. In the Match conditions, click Destination Data Prefix.</p> <p>b. To match a list of destination prefixes, select the list from the drop-down.</p> <p>c. To match an individual destination prefix, enter the prefix in the Destination: IP Prefix field.</p>	source/ destination-data-prefix-list	source/ destination-data-prefix-list
Destination Port	<p>a. In the Match conditions, click Destination Port.</p> <p>b. In the Destination: Port field, enter the port number. Specify a single port number, a list of port numbers (with numbers separated by a space), or a range of port numbers (with the two numbers separated with a hyphen [-]).</p>	src/dst ip	src/dst ip
DNS Application List	<p>Add an application list to enable split DNS.</p> <p>a. In the Match conditions, click DNS Application List.</p> <p>b. In the drop-down, select the application family.</p>	dns-app-list	
DNS	<p>Add an application list to process split DNS.</p> <p>a. In the Match conditions, click DNS.</p> <p>b. In the drop-down, select Request to process DNS requests for the DNS applications, and select Response to process DNS responses for the applications.</p>	dns-request dns-response	
DSCP	<p>a. In the Match conditions, click DSCP.</p> <p>b. In the DSCP field, type the DSCP value, a number from 0 through 63.</p>	dscp	dscp
Packet Length	<p>a. In the Match conditions, click Packet Length.</p> <p>b. In the Packet Length field, type the length, a value from 0 through 65535.</p>	packet-len	packet-len
PLP	<p>a. In the Match conditions, click PLP to set the Packet Loss Priority.</p> <p>b. In the PLP drop-down, select Low or High. To set the PLP to high, apply a policer that includes the exceed remark option.</p>		
Protocol	<p>a. In the Match conditions, click Protocol.</p> <p>b. In the Protocol field, type the Internet Protocol number, a number from 0 through 255.</p>	protocol	protocol/next header

Match Condition	Procedure	IPv4 Fields	IPv6 Fields
Source Data Prefix	<p>a. In the Match conditions, click Source Data Prefix.</p> <p>b. To match a list of source prefixes, select the list from the drop-down.</p> <p>c. To match an individual source prefix, enter the prefix in the Source field.</p>	source/ destination-data-prefix-list	source /destination-data-prefix-list
Source Port	<p>a. In the Match conditions, click Source Port.</p> <p>b. In the Source field, enter the port number. Specify a single port number, a list of port numbers (with numbers separated by a space), or a range of port numbers (with the two numbers separated with a hyphen [-]).</p>	ports	ports
TCP	<p>a. In the Match conditions, click TCP.</p> <p>b. In the TCP field, syn is the only option available.</p>	tcp flag	

Step 12 To select actions to take on matching data traffic, click the **Actions** box.

Step 13 To drop matching traffic, click **Drop**. The available policy actions are listed to the right of the button.

Step 14 To accept matching traffic, click **Accept**. The available policy actions are listed to the right of the button.

Step 15 Set the policy action as described in the following table. Note that not all actions are available for all match conditions

Match Condition	Description	Procedure
Counter	Count matching data packets.	<p>a. In the Action conditions, click Counter.</p> <p>b. In the Counter Name field, enter the name of the file in which to store packet counters.</p>
DSCP	Assign a DSCP value to matching data packets.	<p>a. In the Action conditions, click DSCP.</p> <p>b. In the DSCP field, type the DSCP value, a number from 0 through 63.</p>
Forwarding Class	Assign a forwarding class to matching data packets.	<p>a. In the Match conditions, click Forwarding Class.</p> <p>b. In the Forwarding Class field, type the class value, which can be up to 32 characters long.</p>

Match Condition	Description	Procedure
Policer	Apply a policer to matching data packets.	<ol style="list-style-type: none"> a. In the Match conditions, click Policer. b. In the Policer drop-down field, select the name of a policer.
Loss Correction	<p>Apply loss correction to matching data packets.</p> <p>Forward Error Correction (FEC) recovers lost packets on a link by sending redundant data, enabling the receiver to correct errors without the need to request retransmission of data.</p> <p>FEC is supported only for IPSEC tunnels, it is not supported for GRE tunnels.</p> <ul style="list-style-type: none"> • FEC Adaptive – Corresponding packets are subjected to FEC only if the tunnels that they go through have been deemed unreliable based on measured loss. Adaptive FEC starts to work at 2% packet loss; this value is hard-coded and is not configurable. • FEC Always – Corresponding packets are always subjected to FEC. • Packet Duplication – Sends duplicate packets over a single tunnel. If more than one tunnel is available, duplicated packets will be sent over the tunnel with the best parameters. 	<ol style="list-style-type: none"> a. In the Match conditions, click Loss Correction. b. In the Loss Correction field, select FEC Adaptive, FEC Always, or Packet Duplication.
Click Save Match and Actions .		

- Step 16** Create additional sequence rules as desired. Drag and drop to re-arrange them.
- Step 17** Click **Save Data Policy**.
- Step 18** Click **Next** to move to Apply Policies to Sites and VPNs in the wizard.

Step 3: Apply Policies to Sites and VPNs

In Apply Policies to Sites and VPNs, apply a policy to overlay network sites and VPNs.

- Step 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.
- Step 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.
- Step 3** From the Topology bar, select the tab that corresponds to the type of policy block—**Topology**, **Application-Aware Routing**, **Traffic Data**, or **Cflowd**. The table then lists policies that you have created for that type of policy block.
- Step 4** Associate the policy with VPNs and sites. The choice of VPNs and sites depends on the type of policy block:

Step 4: Activate a Centralized Data Policy

- 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. Select one or more site lists, and select one or more VPN lists. Click **Add**.
- b) For an **Application-Aware Routing** policy block, click **Add New Site List and VPN list**. Select one or more site lists, and select one or more VPN lists. Click **Add**.
- c) For a **Traffic Data** policy block, click **Add New Site List and VPN List**. Select the direction for applying the policy (**From Tunnel**, **From Service**, or **All**), select one or more site lists, and select one or more VPN lists. Click **Add**.
- d) For a **cflowd** policy block, click **Add New Site List**. Select one or more site lists, Click **Add**.

Step 5 Click **Preview** to view the configured policy. The policy is displayed in CLI format.

Step 6 Click **Save Policy**. The **Configuration > Policies** screen appears, and the policies table includes the newly created policy.

Step 4: Activate a Centralized Data Policy

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

Step 1 In the Cisco vManage NMS, select the **Configure > Policies** screen.

Step 2 Select a policy from the policy table.

Step 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.

Step 4 Click **Activate**.

Configure Centralized Data Policy Using CLI

Following are the high-level steps for configuring a VPN membership data policy:

1. Create a list of overlay network sites to which the VPN membership 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 and VPNs, as needed:

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

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:DB8:19::1
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:DB8:19::1
vsmart(config-match)# exit
vsmart(config-sequence-101)# match destination-ipv6 2001:DB8:19::1
vsmart(config-match)#

```

3. Create lists of TLOCs, as needed.

```

vSmart(config)# policy
vSmart(config-policy)# lists tloc-list list-name
vSmart(config-lists-list-name)# tloc ip-address color color encapsulation
[preference number]

```

4. Define policing parameters, as needed:

```

vSmart(config-policy)# policer policer-name
vSmart(config-policer)# rate bandwidth
vSmart(config-policer)# burst bytes
vSmart(config-policer)# exceed action

```

5. Create a data policy instance and associate it with a list of VPNs:

```

vSmart(config)# policy data-policy policy-name
vSmart(config-data-policy-policy-name)# vpn-list list-name

```

6. Create a series of match–pair sequences:

```

vSmart(config-vpn-list)# 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).

7. Define match parameters for packets:

```

vSmart(config-sequence-number)# matchparameters

```

8. Define actions to take when a match occurs:

```

vSmart(config-sequence-number)# action (accept | drop) [count counter-name] [log]
[tcp-optimization]
vSmart(config-sequence-number)# action accept nat [pool number] [use-vpn 0]
vSmart(config-sequence-number)# action accept redirect-dns (host | ip-address)
vSmart(config-sequence-number)# action accept set parameters

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:DB8:19::1
vsmart(config-set)#

```

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

10. If a route does not match any of the conditions in one of the sequences, it is rejected by default. To accept nonmatching prefixed, configure the default action for the policy:

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

11. Apply the policy to one or more sites in the overlay network:

```
vSmart(config)# apply-policy site-list list-name data-policy policy-name (all
| from-service | from-tunnel)
```

Structural Components of Policy Configuration for Centralized Data Policy

The following commands are the structural components required to configure VPN membership policy. Each one is explained in more detail in the sections that follow.

```
policy
 lists
  app-list list-name
    (app applications | app-family application-families)
  data-prefix-list list-name
  ip-prefix prefix
  site-list list-name
  site-id site-id
  tloc-list list-name
  tloc ip-address color color encaps encapsulation [preference value]
  vpn-list list-name
  vpn vpn-id
 policer policer-name
  burst bytes
  exceed action
  rate bandwidth
 data-policy policy-name
 vpn-list list-name
  sequence number
  match
    app-list list-name
    destination-data-prefix-list list-name
    destination-ip prefix/length
    destination-port port-numbers
    dscp number
    dns-app-list list-name
    dns (request | response)
    packet-length number
    protocol number
    source-data-prefix-list list-name
    source-ip prefix/length
    source-port port-numbers
    tcp flag
  action
    cflowd (not available for deep packet inspection)
    count counter-name
    drop
    log
    redirect-dns (dns-ip-address | host)
    tcp-optimization
    accept
      nat [pool number] [use-vpn 0]
      set
        dscp number
        forwarding-class class
        local-tloc color color [encap encapsulation] [restrict]
        next-hop ip-address
        policer policer-name

    tloc ip-address color color [encap encapsulation]
```

```

tloc-list list-name
  vpn vpn-id
  default-action
    (accept | drop)
apply-policy site-list list-name
  data-policy policy-name (all | from-service | from-tunnel)

```

Lists

Centralized data policy for deep packet inspection 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.

- **Configuration > Policies > Centralized Policy > Add Policy > Create Groups of Interest**
- **Configuration > Policies > Custom Options > Lists.**

In the CLI, you configure lists under the **policy lists** command hierarchy on Cisco vSmart Controllers.

List Type	Description	vManage / CLI Command
Applications and application families	<p>List of one or more applications or application families running on the subnets connected to the device.</p> <ul style="list-style-type: none"> • <i>application-names</i> can be the names of one or more applications. The Cisco IOS XE SD-WAN devices supports about 2300 different applications. To list the supported applications, use the ? in the CLI. • <i>application-families</i> can be one or more of the following: antivirus, application-service, audio_video, authentication, behavioral, compression, database, encrypted, erp, file-server, file-transfer, forum, game, instant-messaging, mail, microsoft-office, middleware, network-management, network-service, peer-to-peer, printer, routing, security-service, standard, telephony, terminal, thin-client, tunneling, wap, web, and webmail. 	<p>Configuration > Policies > Centralized Policy > Add Policy > Create Groups of Interest > Application</p> <p>or</p> <p>Configuration > Policies > Centralized Policy > Lists > Application</p> <p>app-list list-name</p> <p>(app applications app-family application-families)</p>

List Type	Description	vManage / CLI Command
Prefixes	List of one or more IP prefixes.	Configuration > Policies > Centralized Policy > Add Policy > Create Groups of Interest > Prefix or Configuration > Policies > Custom Options > Centralized Policy > Lists > Prefix prefix-list <i>list-name</i> ip-prefix <i>prefix/length</i>
Sites	List of one or 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).	Configuration > Policies > Centralized Policy > Add Policy > Create Groups of Interest > Site or Configuration > Policies > Custom Options > Centralized Policy > Lists > Site site-list <i>list-name</i> site-id <i>site-id</i>
TLOCs	<p>List of one or more TLOCs in the overlay network.</p> <p>For each TLOC, specify its address, color, and encapsulation. <i>address</i> is the system IP address. <i>color</i> can be one of 3g, biz-internet, blue, bronze, custom1, custom2, custom3, default, gold, green, lte, metro-ethernet, mpls, mpls-restricted, private1 through private6, public-internet, red, and silver. <i>encapsulation</i> can be gre or ipsec.</p> <p>Optionally, set a preference value (from 0 to 232 – 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.</p>	Configuration > Policies > Centralized Policy > Add Policy > Create Groups of Interest > TLOC or Configuration > Policies > Custom Options > Centralized Policy > Lists > Site tloc-list <i>list-name</i> tloc <i>ip-address color color encap encapsulation</i> [preference number]

List Type	Description	vManage / CLI Command
VPNs	<p>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.</p> <p>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).</p>	<p>Configuration > Policies > Centralized Policy > Add Policy > Create Groups of Interest > VPN</p> <p>or</p> <p>Configuration > Policies > Custom Options > Centralized Policy > Lists > VPN</p> <p>vpn-list <i>list-name</i></p> <p>vpn <i>vpn-id</i></p>

VPN Lists

Each centralized data policy is associated with a VPN list. You configure VPN lists with the **policy data-policy vpn-list** command. The list you specify must be one that you created with a VPN Group of Interest or List in the Cisco vManage policy configuration wizard or with the **policy lists vpn-list** command.

For centralized data policy, you can include any VPNs except for VPN 0 and VPN 512. VPN 0 is reserved for control traffic, so never carries any data traffic, and VPN 512 is reserved for out-of-band network management, so also never carries any data traffic. Note that while the CLI allows you to include these two VPNs in a data policy configuration, the policy is not applied to these two VPNs.

Policer Parameters

To configure policing parameters, create a policer that specifies the maximum bandwidth and burst rate for traffic on an interface, and how to handle traffic that exceeds these values.

In Cisco vManage NMS, you configure policer parameters from:

- **Configuration > Policies > Centralized Policy > Add Policy > Create Groups of Interest > Policer**
- **Configuration > Policies > Custom Options > Centralized Policy > Lists > Policer**

In the CLI, you configure policer parameters as follows:

```
vSmart(config)# policy policer policer-name
vSmart(config-policer)# rate bps
vSmart(config-policer)# burst bytes
vSmart(config-policer)# exceed action
```

rate is the maximum traffic rate. It can be a value from 0 through 264 – 1 bits per second.

burst is the maximum traffic burst size. It can be a value from 15000 to 1000000 bytes.

exceed is the action to take when the burst size or traffic rate is exceeded. *action* can be **drop** (the default) or **remark**. The **drop** action is equivalent to setting the packet loss priority (PLP) bit to low. The **remark** action sets the PLP bit to high. In centralized data policy, access lists, and application-aware routing policy, you can match the PLP with the **match plp** option.

Sequences

Each VPN list consists of sequences of match–action pairs. The sequences are numbered to set the order in which data traffic 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 data-policy vpn-list sequence` command.

Each sequence can contain one match condition and one action condition.

Match Parameters

Centralized data policy can match IP prefixes and fields in the IP headers, as well as applications. You can also enable split DNS.

In Cisco vManage NMS, you configure match parameters from:

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

Each sequence in a policy can contain one match condition.

For data policy, you can match these parameters:

Description	vManage Configuration/CLI Configuration Command	Value or Range
Match all packets	Omit Match Omit <code>match</code> command	—
Applications or application families	Match Applications/Application Family List <code>app-list list-name</code>	Name of an application list or an <code>app-list</code> list
Group of destination prefixes	Match Destination Data Prefix <code>destination-data-prefix-list list-name</code>	Name of a data prefix list or a <code>data-prefix-list</code> list
Individual destination prefix	Match Destination Data Prefix <code>destination-ip prefix/length</code>	IP prefix and prefix length

Description	vManage Configuration/CLI Configuration Command	Value or Range
Destination port number	Match Destination Port destination-port <i>number</i>	0 through 65535; specify a single port number, a list of port numbers (with numbers separated by a space), or a range of port numbers (with the two numbers separated with a hyphen [-])
Enable split DNS, to resolve and process DNS requests and responses on an application-by-application basis	Match DNS Application List dns-app-list <i>list-name</i>	Name of an app-list list. This list specifies the applications whose DNS requests are processed.
Specify the direction in which to process DNS packets	Match DNS dns (request response)	To process DNS requests sent by the applications (for outbound DNS queries), specify dns request . To process DNS responses returned from DNS servers to the applications, specify dns response .
DSCP value	Match DSCP dscp <i>number</i>	0 through 63
Packet length	Match Packet Length packet-length <i>number</i>	0 through 65535; specify a single length, a list of lengths (with numbers separated by a space), or a range of lengths (with the two numbers separated with a hyphen [-])
Packet loss priority (PLP)	Match PLP plp	(high low) By default, packets have a PLP value of low . To set the PLP value to high , apply a policer that includes the exceed remark option.
Internet protocol number	Match Protocol protocol <i>number</i>	0 through 255
Group of source prefixes	Match Source Data Prefix source-data-prefix-list <i>list-name</i>	Name of a data prefix or a data-prefix-list list
Individual source prefix	Match Source Data Prefix source-ip <i>prefix/length</i>	IP prefix and prefix length
Source port number	Match Source Port source-port <i>address</i>	0 through 65535; specify a single port number, a list of port numbers (with numbers separated by a space), or a range of port numbers (with the two numbers separated with a hyphen [-])

Description	vManage Configuration/CLI Configuration Command	Value or Range
TCP flag	<code>tcp flag</code>	<code>syn</code>

Action Parameters

Table 16: Feature History

Feature Name	Release Information	Description
Path Preference Support for Cisco IOS XE SD-WAN Devices	Cisco IOS XE Release Amsterdam 17.2.1r	This feature extends to Cisco IOS XE SD-WAN devices, support for selecting one or more local transport locators (TLOCs) for a policy action.

When data traffic matches the conditions in the match portion of a centralized data policy, the packet can be accepted or dropped, and it can be counted. Then, you can associate parameters with accepted packets.

In Cisco vManage NMS, you configure match parameters from:

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

In the CLI, you configure the action parameters with the `policy data-policy vpn-list sequence action` command.

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

In the action, you first specify whether to accept or drop a matching data packet, and whether to count it:

Description	vManage Configuration/CLI Configuration Parameter	Value or Range
Accept the packet. An accepted packet is eligible to be modified by the additional parameters configured in the action portion of the policy configuration.	Click Accept . accept	—
Enable cflowd traffic monitoring.	Click Accept , then action Cflowd cflowd	—
Count the accepted or dropped packets.	Action Counter Click Accept , then action Counter count counter-name	Name of a counter. Use the show policy access-lists counters command on the Cisco IOS XE SD-WAN device.
Discard the packet. This is the default action.	Click Drop drop	—

Description	vManage Configuration/CLI Configuration Parameter	Value or Range
Redirect DNS requests to a particular DNS server. Redirecting requests is optional, but if you do so, you must specify both actions.	Click Accept , then action Redirect DNS redirect-dns host redirect-dns ip-address	For an inbound policy, redirect-dns host allows the DNS response to be correctly forwarded back to the requesting service VPN. For an outbound policy, specify the IP address of the DNS server.
Fine-tune TCP to decrease round-trip latency and improve throughput for matching TCP traffic.	Click Accept , then action TCP Optimization tcp-optimization	—



Note On Cisco IOS XE routers, all the ongoing optimized flows are dropped when the TCP Optimization is removed.

Then, for a packet that is accepted, the following parameters can be configured:

Description	vManage	CLI Configuration Parameter	Value or Range
Enable cflowd traffic monitoring.	Click Accept , then action Cflowd .	cflowd	—
Direct matching traffic to the NAT functionality so that it can be redirected directly to the Internet or other external destination.	Click Accept , then action NAT Pool or NAT VPN .	nat [pool number] [use-vpn 0]	—
DSCP value.	Click Accept , then action DSCP .	set dscp value	0 through 63
Forwarding class.	Click Accept , then action Forwarding Class .	set forwarding-class value	Name of forwarding class

Description	vManage	CLI Configuration Parameter	Value or Range
<p>Direct matching packets to a TLOC that matches the color and encapsulation</p> <p>By default, if the TLOC is not available, traffic is forwarded using an alternate TLOC.</p>	Click Accept , then action Local TLOC .	set local-tloc color <i>color</i> [encap <i>encapsulation</i>]	<p><i>color</i> can be:</p> <p>3g, biz-internet, blue, bronze, custom1, custom2,</p>
<p>Direct matching packets to one of the TLOCs in the list if the TLOC matches the color and encapsulation</p> <p>By default, if the TLOC is not available, traffic is forwarded using an alternate TLOC. To drop traffic if a TLOC is unavailable, include the restrict option.</p>	Click Accept , then action Local TLOC	set local-tloc-list color <i>color</i> encap <i>encapsulation</i> [restrict]	<p>custom3, default, gold, green lte, metro-ethernet, mpls, private1 through private6, public-internet, red, and silver.</p> <p>By default, <i>encapsulation</i> is ipsec. It can also be gre.</p>
Set the next hop to which the packet should be forwarded.	Click Accept , then action Next Hop .	set next-hop <i>ip-address</i>	IP address
Apply a policer.	Click Accept , then action Policer .	set policer <i>policer-name</i>	Name of policer configured with a policy policer command.
<p>Specify a service to redirect traffic to before delivering the traffic to its destination.</p> <p>The TLOC address or list of TLOCs identifies the remote 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.</p> <p>The VPN identifier is where the service is located.</p> <p>Configure the services themselves on the Cisco IOS XE SD-WAN devices that are collocated with the service devices, using the vpn service command.</p>	Click Accept , then action Service .	set service <i>service-name</i> [tloc <i>ip-address</i> tloc-list <i>list-name</i>] [vpn <i>vpn-id</i>]	<p>Standard services: FW, IDS, IDP</p> <p>Custom services: netsvc1, netsvc2, netsvc3, netsvc4</p> <p>TLOC list is configured with a policy lists tloc-list list.</p>

Description	vManage	CLI Configuration Parameter	Value or Range
Direct traffic to a remote TLOC that matches the IP address, color, and encapsulation.	Click Accept , then action TLOC .	set tloc address color color [encap encapsulation]	TLOC address, color, and encapsulation
Direct traffic to one of the remote TLOCs in the TLOC list if it matches the IP address, color, and encapsulation of one of the TLOCs in the list. If a preference value is configured for the matching TLOC, that value is assigned to the traffic.	Click Accept , then action TLOC .	set tloc-list list-name	Name of a policy lists tloc-list list
Set the VPN that the packet is part of.	Click Accept , then action VPN .	set vpn vpn-id	0 through 65530

The following table describes the IPv4 and IPv6 actions.

IPv4 Actions	IPv6 Actions
drop, dscp, next-hop (from-service only)/vpn, count, forwarding class, policer (only in interface ACL), App-route SLA (only)	
App-route preferred color, app-route sla strict, cflowd, nat, redirect-dns	
	drop, dscp, next-hop/vpn, count, forwarding class, policer (only in interface ACL) App-route SLA (only), App-route preferred color, app-route sla strict
policer (DataPolicy), tcp-optimization, fec-always,	policer (DataPolicy)
tloc, tloc-list (set tloc, set tloc-list)	tloc, tloc-list (set tloc, set tloc-list)
App-Route backup-preferred color, local-tloc, local-tloc-list	App-Route backup-preferred color, local-tloc, local-tloc-list

Default Action

If a data packet being evaluated does not match any of the match conditions in a data policy, a default action is applied to the packet. By default, the data packet is dropped

In the Cisco vManage NMS, you modify the default action from:

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

- **Configuration > Policies > Custom Options > Centralized Policy > Traffic Policy > (Application-Aware Routing | Traffic Data | Cflowd) > Sequence Type > Sequence Rule > Default Action.**

In the CLI, you modify the default action with the `policy data-policy vpn-list default-action accept` command.

Apply Centralized Data Policy

For a centralized data 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 data-policy policy-name (all | from-service
| from-tunnel)
```

By default, data policy applies to all data traffic passing through the device: the policy evaluates all data traffic going from the local site (that is, from the service side of the router) into the tunnel interface, and it evaluates all traffic entering to the local site through the tunnel interface. You can explicitly configure this behavior by including the **all** option. To have the data policy apply only to traffic coming from the service site and exiting from the local site through the tunnel interface, include the **from-service** option. To have the policy apply only to traffic entering from the tunnel interface and traveling to the service site, include the **from-tunnel** option. You can apply different data policies in each of the two traffic directions.

For all **data-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 **data-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 control policy (**control-policy**)
- Centralized data policy used for cflowd flow monitoring (**data-policy** that 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.

As soon as you successfully activate the configuration by issuing a **commit** command, the Cisco vSmart Controller pushes the data policy to the devices located in the specified sites. To view the policy as configured on the Cisco vSmart Controllers, use the **show running-config** command on the Cisco vSmart Controller:

```
vSmart# show running-config policy
vSmart# show running-config apply-policy
```

To view the policy that has been pushed to the Cisco IOS XE SD-WAN device, use the **show sdwan policy from-vsmart** command on the Cisco IOS XE SD-WAN device.

```
Device# show sdwan policy from-vsmart
```

Deep Packet Inspection

You configure deep packet inspection using a standard centralized data policy. You define the applications of interest in a vManage policy list or with **policy lists app-list** CLI command, and you call these lists in the match portion of the data policy. You can control the path of the application traffic through the network by defining, in the **action** portion of the data policy, the local TLOC or the remote TLOC, or for strict control, you can define both.

Configure Deep Packet Inspection Using vManage

To configure a centralized data policy for deep packet inspection, use the vManage policy configuration wizard. Use the wizard to create and edit deep packet inspection policy components:

- Configure groups of interest (lists) to group related items to be called in the centralized data policy.
- Configure traffic rules.
- Apply the policy.

Step 1: Start the Policy Configuration Wizard

To start the policy configuration wizard:

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

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

Step 2: Create Groups of Interest

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

To configure groups of interest for deep packet inspection:

1. In the left pane, select the type of list. For centralized data policy for deep packet inspection, you can use Application, Site, and VPN lists.
2. To create a new list, click New List.

To modify an existing list, click the More Actions icon to the right of the desired list, and click the pencil icon.

3. In the List Name field, enter a name for the list. 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.
4. In the field below the List Name field, enter the desired values for the list. For some lists you type the desired values, and for others you select from a drop-down.
5. Click Add (for a new list) or Save (for an existing list).
6. Click Next to move to the Configure Topology and VPN Membership screen.
7. Click Next to move the Configure Traffic Rules in the wizard.

Step 3: Configure Traffic Rules

When you first open the Traffic Rules screen, the Application-Aware Routing tab is selected by default:

To configure traffic rules for deep packet inspection policy:

1. In the Application-Aware Routing bar, click Traffic Data.
2. To create a new centralized data policy, click Add Policy.
To modify an existing policy, click the More Actions icon to the right of the desired policy, and click the pencil icon.
3. If data traffic does not match any of the conditions in one of the sequences, it is dropped by default. If you want nonmatching routes to be accepted, click the pencil icon in the Default Action, click Accept, and click Save Match And Actions.
4. To create a match–action sequence for data traffic:
 - a. Click Sequence Type.
 - b. To create a match–action rule, click Sequence Rule. The Match button is selected by default.
 - c. Click the desired Match button, and enter the desired values in Match Conditions. For some conditions, you type the desired values, and for others you select from a drop-down.
 - d. Click the Actions button. The default action is Reject. To accept matching packets, click the Accept radio button. Then click the desired action, and enter the desired values for Actions.
 - e. Click Save Match and Actions.
 - f. Create additional Sequence Rules or Sequence Types, as needed.
5. To rename a Sequence Type, double-click its name in the right pane, and type the new name. The name also changes in the right pane.
6. To re-order sequence rules and types, drag and drop them them.
7. Click Save.
8. Click Next to move to the Apply Policies to Sites and VPNs in the wizard.

Step 4: Apply Policies to Sites and VPNs

1. In Apply Policies to Sites and VPNs, apply a policy to overlay network sties and VPNs:

2. 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.
3. 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.
4. From the Topology bar, select the Application-Aware Routing tab. The table then lists policies that you have created for that type of policy block.
5. Click Add New Site List and VPN List or Add New Site. Some topology blocks might have no Add buttons. Select one or more site lists, and select one or more VPN lists. Click Add.
6. Click Preview to view the configured policy. The policy is displayed in CLI format.
7. Click Save Policy. The Configuration > Policies screen opens, and the policies table includes the newly created policy.

Step 5: Activate a Centralized Data Policy

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

1. In 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 vSmart controllers to which the policy is to be applied.
4. Click Activate.

Configure Deep Packet Inspection Using CLI

Following are the high-level steps for configuring a centralized data policy to use for deep packet inspection:

1. Create a list of overlay network sites to which the data 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 applications and application families that are to be subject to the data policy. Each list can contain one or more application names, or one or more application families. A single list cannot contain both applications and application families.

```
vSmart(config)# policy lists
vSmart(config-lists)# app-list list-name
vSmart(config-app-list)# app application-name
```

```
vSmart(config)# policy lists
vSmart(config-lists)# app-list list-name
vSmart(config-applist)# app-family family-name
```

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

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

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

4. Create lists of TLOCs, as needed:

```
vSmart(config)# policy
vSmart(config-policy)# lists tloc-list list-name
vSmart(config-lists-list-name)# tloc ip-address color color encap encapsulation
[preference number]
```

5. Define policing parameters, as needed:

```
vSmart(config-policy)# policer policer-name
vSmart(config-policer)# rate bandwidth
vSmart(config-policer)# burst bytes
vSmart(config-policer)# exceed action
```

6. Create a data policy instance and associate it with a list of VPNs:

```
vSmart(config)# policy data-policy policy-name
vSmart(config-data-policy-policy-name)# vpn-list list-name
```

7. Create a series of match–pair sequences:

```
vSmart(config-vpn-list)# 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).

8. Define match parameters based on applications:

```
vSmart(config-sequence-number)# match app-list list-name
```

9. Define additional match parameters for data packets:

```
vSmart(config-sequence-number)# match parameters
```

10. Define actions to take when a match occurs:

```
vSmart(config-sequence-number)# action (accept | drop) [count]
```

11. For packets that are accepted, define the actions to take. To control the tunnel over which the packets travels, define the remote or local TLOC, or for strict control over the tunnel path, set both:

```
vSmart(config-action)# set tloc ip-address color color encap encapsulation
vSmart(config-action)# set tloc-list list-name
vSmart(config-action)# set local-tloc color color encap encapsulation
vSmart(config-action)# set local-tloc-list color color encap encapsulation [restrict]
```

12. Define additional actions to take.**13. Create additional numbered sequences of match–action pairs within the data policy, as needed.****14. If a route does not match any of the conditions in one of the sequences, it is rejected by default. If you want nonmatching prefixes to be accepted, configure the default action for the policy:**

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

15. Apply the policy to one or more sites in the overlay network:

```
vSmart(config)# apply-policy site-list list-name data-policy policy-name (all |
from-service | from-tunnel)
```

To enable the infrastructure for deep packet inspection on the vEdge routers, include the following command in the configuration on the routers:

```
vEdge(config)# policy app-visibility
```

Structural Components of Policy Configuration for Deep Packet Inspection

Following are the structural components required to configure centralized data policy for deep packet inspection. Each one is explained in more detail in the sections below.

On the vSmart controller:

```
policy
  lists
    app-list list-name
      (app applications | app-family application-families)
    data-prefix-list list-name
      ip-prefix prefix
    site-list list-name
      site-id site-id
    tloc-list list-name
      tloc ip-address color color encap encapsulation [preference value]
    vpn-list list-name
      vpn vpn-id
  policer policer-name
    burst bytes
    exceed action
    rate bps
  data-policy policy-name
    vpn-list list-name
      sequence number
    match
      app-list list-name
      destination-data-prefix-list list-name
      destination-ip ip-addresses
      destination-port port-numbers
      dscp number
      packet-length number
      protocol protocol
      source-data-prefix-list list-name
      source-ip ip-addresses
      source-port port-numbers
      tcp flag
    action
      drop
      count counter-name
      log
      accept
        nat [pool number] [use-vpn 0]
        set
          dscp number
          forwarding-class class
          local-tloc color color [encap encapsulation] [restrict]
          next-hop ip-address
          policer policer-name
          service service-name local [restrict] [vpn vpn-id]
          service service-name (tloc ip-address | tloc-list list-name) [vpn vpn-id]
          tloc ip-address color color encap encapsulation
          tloc-list list-name
          vpn vpn-id
      default-action
```

```
(accept | drop)
apply-policy site-list list-name
  data-policy policy-name (all | from-service | from-tunnel)
```

On the vEdge router:

```
policy
  app-visibility
```

Action Parameters for Configuring Deep Packet Inspection

When data traffic matches the conditions in the match portion of a centralized data policy, the packet can be accepted or dropped, and it can be counted. Then, you can associate parameters with accepted packets.

In vManage NMS, you configure match parameters from:

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

In the CLI, you configure the action parameters under the **policy data-policy vpn-list sequence action** command.

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

In the action, you first specify whether to accept or drop a matching data packet, and whether to count it:

Description	vManage Configuration/CLI Configuration Parameter	Value or Range
Accept the packet. An accepted packet is eligible to be modified by the additional parameters configured in the action portion of the policy configuration.	Click Accept . accept	—
Count the accepted or dropped packets.	Action Counter Click Accept , then action Counter count counter-name	Name of a counter. Use the show policy access-lists counters command on the Cisco device.
Discard the packet. This is the default action.	Click Drop . drop	—

To view the packet logs, use the **show app log flows** and **show log** commands.

Then, for a packet that is accepted, the following parameters can be configured. Note that you cannot use DPI with either cflowd or NAT.

Description	vManage	CLI Configuration Parameter	Value or Range
DSCP value.	Click Accept , then action DSCP .	set dscp value	0 through 63

Description	vManage	CLI Configuration Parameter	Value or Range
Forwarding class.	Click Accept , then action Forwarding Class .	set forwarding-class <i>value</i>	Name of forwarding class
Direct matching packets to a TLOC that matches the color and encapsulation By default, if the TLOC is not available, traffic is forwarded using an alternate TLOC.	Click Accept , then action Local TLOC .	set local-tloc color <i>color</i> [encap <i>encapsulation</i>]	<i>color</i> can be: 3g, biz-internet, blue, bronze, custom1, custom2, custom3, default, gold,
Direct matching packets to one of the TLOCs in the list if the TLOC matches the color and encapsulation By default, if the TLOC is not available, traffic is forwarded using an alternate TLOC. To drop traffic if a TLOC is unavailable, include the restrict option.	Click Accept , then action Local TLOC	set local-tloc-list color <i>color</i> encap <i>encapsulation</i> [restrict]	green lte, metro-ethernet, mpls, private1 through private6, public-internet, red, and silver. By default, <i>encapsulation</i> is ipsec . It can also be gre .
Set the next hop to which the packet should be forwarded.	Click Accept , then action Next Hop .	set next-hop <i>ip-address</i>	IP address
Apply a policer.	Click Accept , then action Policer .	set policer <i>policer-name</i>	Name of policer configured with a policy policer command.
Direct matching packets to the name service, before delivering the traffic to its ultimate destination. The TLOC address or list of TLOCs identifies the remote 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 vEdge routers that are collocated with the service devices, using the vpn service configuration command.	Click Accept , then action Service .	set service <i>service-name</i> [tloc <i>ip-address</i> tloc-list <i>list-name</i>] [vpn <i>vpn-id</i>]	Standard services: FW, IDS, IDP Custom services: netsvc1, netsvc2, netsvc3, netsvc4 TLOC list is configured with a policy lists tloc-list list.

Description	vManage	CLI Configuration Parameter	Value or Range
Direct matching packets to the named service that is reachable via a GRE tunnel whose source is in the transport VPN (VPN 0). If the GRE tunnel used to reach the service is down, packet routing falls back to using standard routing. To drop packets when a GRE tunnel to the service is unreachable, include the restrict option. In the service VPN, you must also advertise the service using the service command. You configure the GRE interface or interfaces in the transport VPN (VPN 0).	Click Accept , then action Service .	set service <i>service-name</i> [tloc <i>ip-address</i> tloc-list <i>list-name</i>] [vpn <i>vpn-id</i>]	Standard services: FW, IDS, IDP Custom services: netsvc1, netsvc2,netsvc3, netsvc4
Direct traffic to a remote TLOC. The TLOC is defined by its IP address, color, and encapsulation.	Click Accept , then action TLOC .	set local-tloc color <i>color</i> [encap <i>encapsulation</i>]	TLOC address, color, and encapsulation
Direct traffic to one of the remote TLOCs in the TLOC list.	Click Accept , then action TLOC .	set tloc-list <i>list-name</i>	Name of a policy lists tloc-list list
Set the VPN that the packet is part of.	Click Accept , then action VPN .	set vpn <i>vpn-id</i>	0 through 65530

Default Action

If a data packet being evaluated does not match any of the match conditions in a data policy, a default action is applied to the packet. By default, the data packet is dropped.

In vManage NMS, you modify the default action from Configuration > Policies > Centralized Policy > Add Policy > Configure Traffic Rules > Application-Aware Routing > Sequence Type > Sequence Rule > Default Action.

In the CLI, you modify the default action with the **policy data-policy vpn-list default-action accept** command.

Apply Centralized Data Policy for Deep Packet Inspection

For a deep packet inspection centralized data policy to take effect, you apply it to a list of sites in the overlay network.

To apply a centralized policy in vManage NMS:

1. In 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 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 data-policy policy-name (all | from-service  
| from-tunnel)
```

By default, data policy applies to all data traffic passing through the vEdge router: the policy evaluates all data traffic going from the local site (that is, from the service side of the router) into the tunnel interface, and it evaluates all traffic entering to the local site through the tunnel interface. You can explicitly configure this behavior by including the **all** option. To have the data policy apply only to policy exiting from the local site, include the **from-service** option. To have the policy apply only to incoming traffic, include the **from-tunnel** option.

You cannot apply the same type of policy to site lists that contain overlapping site IDs. That is, all data policies cannot have overlapping site lists among themselves. If you accidentally misconfigure overlapping site lists, the attempt to commit the configuration on the vSmart controller fails.

As soon as you successfully activate the configuration by issuing a **commit** command, the vSmart controller pushes the data policy to the vEdge routers located in the specified sites. To view the policy as configured on the vSmart controller, use the **show running-config** command on the vSmart controller:

```
vSmart# show running-config policy  
vSmart# ;show running-config apply-policy
```

To view the policy that has been pushed to the vEdge router, use the **show policy from-vsmart** command on the vEdge router.

```
vEdge# show policy from-vsmart
```

Monitor Running Applications

To enable the deep packet inspection infrastructure on the vEdge routers, you must enable application visibility on the routers:

```
vEdge(config)# policy app-visibility
```

To display information about the running applications, use the **show app dpi supported-applications**, **show app dpi applications**, and **show app dpi flows** commands on the router.

View DPI Applications Using vManage

You can view the list of all the application-aware applications supported by the SD-WAN software on the router using the following steps:

1. In the Cisco vManage, select the **Monitor > Network** screen.
2. From the **WAN-Edge** pane, select the **Device** that supports DPI. The vManage Control Connections page displays.
3. In the left pane, select **Real Time** to view the device details.
4. From the **Device Options** drop-down, choose **DPI Applications** to view the list of applications running on the device.
5. From the **Device Options** drop-down, choose **DPI Supported Applications** to view the list of applications that are supported on the device.

Centralized Data Policy Configuration Examples

This topic provides some examples of configuring centralized data policy to influence traffic flow across the Cisco IOS XE SD-WAN domain and to configure a Cisco IOS XE SD-WAN device to be an Internet exit point.

General Centralized Data Policy Example

This section shows a general example of a centralized data policy to illustrate that you configure centralized data policy on a Cisco vSmart Controller and that after you commit the configuration, the policy itself is pushed to the required Cisco IOS XE SD-WAN devices.

Here we configure a simple data policy on the Cisco vSmart Controller vm9:

```
vm9# show running-config policy
policy
data-policy test-data-policy
vpn-list test-vpn-list
sequence 10
match
  destination-ip 209.165.201.0/27
!
action drop
count test-counter
!
!
default-action drop
!
!
lists
vpn-list test-vpn-list
vpn 1
!
site-list test-site-list
site-id 500
!
!
!
```

Then, apply this policy to the site list named **test-site-list**, which includes site 500:

```
vm9# show sdwan running-config apply-policy
apply-policy
site-list test-site-list
data-policy test-data-policy
!
!
```

Immediately after you activate the configuration on the Cisco vSmart Controller, it pushes the policy configuration to the Cisco IOS XE SD-WAN devices in site 500. One of these devices is vm5, where you can see that the policy has been received:

```
vm5# show sdwan policy from-vsmart
policy-from-vsmart
data-policy test-data-policy
vpn-list test-vpn-list
sequence 10
match
  destination-ip 209.165.201.0/27
!
action drop
count test-counter
```



```

!
!
default-action drop
!
!
lists
vpn-list test-vpn-list
vpn 1
!
!
!

```

Control Access

This example shows a data policy that limits the type of packets that a source can send to a specific destination. Here, the host at source address 192.0.2.1 in site 100 and VPN 100 can send only TCP traffic to the destination host at 203.0.113.1. This policy also specifies the next hop for the TCP traffic sent by 192.0.2.1, setting it to be TLOC 209.165.200.225, color gold. All other traffic is accepted as a result of the **default-action** statement.

```

policy
lists
site-list north
site-id 100
vpn-list vpn-north
vpn 100
!
data-policy tcp-only
vpn-list vpn-north
sequence 10
match
source-ip 192.0.2.1/32
destination-ip 203.0.113.1/32
protocol tcp
action accept
set tloc 209.165.200.225 gold
!
default-action accept
!
!
apply-policy
site north data-policy tcp-only

```

Restrict Traffic

This examples illustrates how to disallow certain types of data traffic from being sent from between VPNs. This policy drops data traffic on port 25, which carries SMTP mail traffic, that originates in 209.165.201.0/27. However, the policy accepts all other data traffic, including non-SMTP traffic from 209.165.201.0/27.

```

policy
lists
data-prefix-list north-ones
ip-prefix 209.165.201.0/27
port 25
vpn-list all-vpns
vpn 1
vpn 2
site-list north
site-id 100
!
data-policy no-mail
vpn-list all-vpns
sequence 10

```

```

        match
        source-data-prefix-list north-ones
        action drop
    !
    default-action accept
!
!
apply-policy
site north data-policy no-mail

```

Localized Data Policy

Data policy operates on the data plane in the Cisco IOS XE SD-WAN overlay network and affects how data traffic is sent among the Cisco IOS XE SD-WAN devices in the network. The Cisco SD-WAN architecture defines two types of data policy, centralized data policy, which controls the flow of data traffic based on the IP header fields in the data packets and based on network segmentation, and localized data policy, which controls the flow of data traffic into and out of interfaces and interface queues on a Cisco IOS XE SD-WAN device.

Localized data policy, so called because it is provisioned on the local Cisco IOS XE SD-WAN device, is applied on a specific router interface and affects how a specific interface handles the data traffic that it is transmitting and receiving. Localized data policy is also referred to as access lists (ACLs). With access lists, you can provision class of service (CoS), classifying data packets and prioritizing the transmission properties for different classes. You can configure policing.

For IPv4, you can configure QoS actions.

You can apply IPv4 access lists in any VPN on the router, and you can create access lists that act on unicast and multicast traffic. You can apply IPv6 access lists only to tunnel interfaces in the transport VPN (VPN 0).

You can apply access lists either in the outbound or inbound direction on the interface. Applying an IPv4 ACL in the outbound direction affects data packets traveling from the local service-side network into the IPsec tunnel toward the remote service-side network. Applying an IPv4 ACL in the inbound direction affects data packets exiting from the IPsec tunnel and being received by the local Cisco IOS XE SD-WAN device. For IPv6, an outbound ACL is applied to traffic being transmitted by the router, and an inbound ACL is applied to received traffic.

Explicit and Implicit Access Lists

Access lists that you configure using localized data policy are called *explicit* ACLs. You can apply explicit ACLs in any VPN on the router.

Router tunnel interfaces also have *implicit ACLs*, which are also referred to as *services*. Some of these are present by default on the tunnel interface, and they are in effect unless you disable them. Through configuration, you can also enable other implicit ACLs. On Cisco IOS XE SD-WAN devices, the following services are enabled by default: DHCP (for DHCPv4 and DHCPv6), DNS, and ICMP. You can also enable services for BGP, Netconf, NTP, OSPF, SSHD, and STUN.

Perform QoS Actions

With access lists, you can provision quality of service (QoS) which allows you to classify data traffic by importance, spread it across different interface queues, and control the rate at which different classes of traffic are transmitted. See Forwarding and QoS Overview.

Localized Data Policy for IPv4

This topic provides procedures for configuring IPv4 localized data policy. This type of data policy is called access lists, or ACLs. You can provision simple access lists that filter traffic based on IP header fields. You also use access lists to apply QoS, and policing to data packets. You can create access lists that act on unicast and multicast traffic.

In Cisco vManage NMS, you configure localized data policy from the **Configuration > Policies** screen, using a policy configuration wizard. In the CLI, you configure these policies on the Cisco IOS XE SD-WAN device.

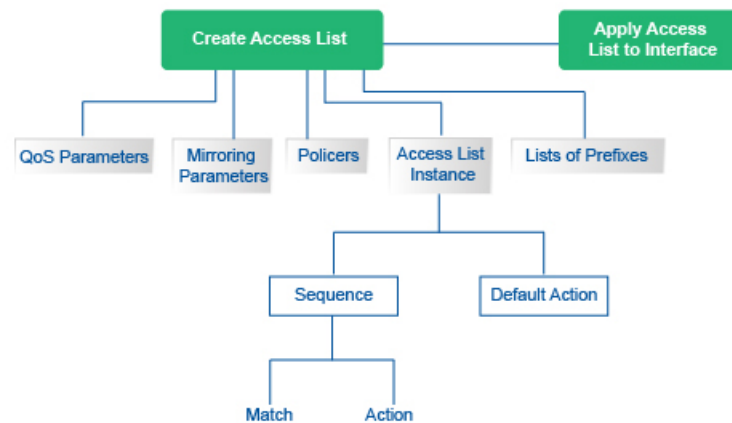
Configuration Components

An access list consists of a sequence of match–action pairs 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 packet stops. Keep this in mind as you design your policies to ensure that the desired actions are taken on the items subject to policy.

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

The following figure illustrates the configuration components for access lists.

Figure 13: Configuration Components



368486

Configure Localized Data Policy for IPv4 Using Cisco vManage

Table 17: Feature History

Feature Name	Release Information	Description
Control Traffic Flow Using Class of Service Values	Cisco IOS XE SD-WAN Release 16.12.1b	This feature lets you control the flow of traffic into and out of a Cisco IOS XE SD-WAN device interface based on the conditions defined in the quality of service (QoS) map. A priority field and a layer 2 class of service (CoS) were added for configuring the re-write rule.

To configure IPv4 localized policy, use the Cisco vManage policy configuration wizard. The wizard is a UI policy builder that consists of five screens to configure IPv4 localized policy components:

- Groups of Interest, also called lists—Create data prefix lists and policer parameters that group together related items and that you call in the match or action components of a policy.
- Forwarding Classes—Define forwarding classes and rewrite rules to use for QoS.
- Access Control Lists—Define the match and action conditions of ACLs.
- Route Policies—Define the match and action conditions of route policies.
- Policy Settings—Define additional policy settings, including Cloud QoS 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 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: Create Groups of Interest

In the Create Groups of interest screen create lists to use in the localized data policy:

Name	Entries	Reference Count	Updated By	Last Updated	Action
business	100 101 102	0	admin	30 Mar 2018 2:40...	Edit Delete
office	103 104 105	0	admin	30 Mar 2018 2:41...	Edit Delete

1. Create news lists of groups, as described in the following table:

Table 18:

List Type	Procedure
Data Prefix	<ol style="list-style-type: none"> 1. In the left bar, click Data Prefix. 2. Click New Data Prefix List. 3. Enter a name for the list. 4. Enter one or more IP prefixes. 5. Click Add.
Mirror	<ol style="list-style-type: none"> 1. In the left bar, click Mirror. 2. Click New Mirror List. The Mirror List popup displays. 3. Enter a name for the list. 4. In the Remote Destination IP field, enter the IP address of the destination to which to mirror the packets. 5. In the Source IP field, enter the IP address of the source of the packets to mirror. 6. Click Save.
Policer	<ol style="list-style-type: none"> 1. In the left bar, click Policer. 2. Click New Policer List. 3. Enter a name for the list. 4. In the Burst field, enter maximum traffic burst size. It can be a value from 15000 to 10000000 bytes. 5. In the Exceed field, select the action to take when the burst size or traffic rate is exceeded. Select Drop (the default) to set the packet loss priority (PLP) to low. Select Remark to set the PLP to high. 6. In the Rate field, enter the maximum traffic rate. It can be value from 0 through $2^{64} - 1$ bps 7. Click Add.

1. Click **Next** to move to Configure Forwarding Classes/QoS in the wizard.

Step 3: Configure Forwarding Classes for QoS

When you first open the Forwarding Classes/QoS screen, the **QoS** tab is selected by default:

To configure forwarding classes for use by QoS:

1. To create a new QoS mapping:
 - a. In the QoS tab, click the **Add QoS** drop-down.
 - b. Select **Create New**.

- c. Enter a name and description for the QoS mapping.
 - d. Click **Add Queue**. The Add Queue popup displays.
 - e. Select the queue number from the Queue drop-down.
 - f. Select the maximum bandwidth and buffer percentages, and the scheduling and drop types. Enter the forwarding class.
 - g. Click **Save**.
2. To import an existing QoS mapping:
 - a. In the QoS tab, click the **Add QoS** drop-down.
 - b. Select **Import Existing**.
 - c. Select a QoS mapping.
 - d. Click **Import**.
3. To view or copy a QoS mapping or to remove the mapping from the localized policy, click the **More Actions** icon to the right of the row, and select the desired action.
4. To configure policy rewrite rules for the QoS mapping:
 - a. In the QoS tab, click the **Add Rewrite Policy** drop-down..
 - b. Select **Create New**.
 - c. Enter a name and description for the rewrite rule.
 - d. Click **Add Rewrite Rule**. The Add Rule popup displays.
 - e. Select a class from the Class drop-down.
 - f. Select the priority (**Low** or **High**) from the Priority drop-down.
Low priority is supported only for Cisco IOS XE SD-WAN devices.
 - g. Enter the DSCP value (0 through 63) in the DSCP field.
 - h. Enter the class of service (CoS) value (0 through 7) in the Layer 2 Class of Service field.
 - i. Click **Save**.
5. To import an existing rewrite rule:
 - a. In the QoS tab, click the **Add Rewrite Policy** drop-down..
 - b. Select **Import Existing**.
 - c. Select a rewrite rule.
 - d. Click **Import**.
6. Click **Next** to move to Configure Access Lists in the wizard.

Step 4: Configure ACLs

1. In the Configure Access Control Lists screen, configure ACLs.
2. To create a new IPv4 ACL, click the **Add Access Control List Policy** drop-down. Then select **Add IPv4 ACL Policy**:
3. Enter a name and description for the ACL.
4. In the left pane, click **Add ACL Sequence**. An Access Control List box is displayed in the left pane.
5. Double-click the **Access Control List** box, and type a name for the ACL.
6. In the right pane, click **Add Sequence Rule** to create a single sequence in the ACL. The Match tab is selected by default.
7. Click a match condition.
8. On the left, enter the values for the match condition.
9. On the right enter the action or actions to take if the policy matches.
10. Repeat Steps 6 through 8 to add match–action pairs to the ACL.
11. To rearrange match–action pairs in the ACL, in the right pane drag them to the desired position.
12. To remove a match–action pair from the ACL, click the **X** in the upper right of the condition.
13. Click **Save Match and Actions** to save a sequence rule.
14. To rearrange sequence rules in an ACL, in the left pane drag the rules to the desired position.
15. To copy, delete, or rename an ACL sequence rule, in the left pane, click **More Options** next to the rule's name and select the desired option.
16. If no packets match any of the ACL 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**.
17. Click **Next** to move to Configure Route Policy in the wizard.
18. Click **Next** to move to the Policy Overview screen.

Step 5: Configure Policy Settings

In Policy Overview, configure policy settings:

1. Enter a name and description for the ACL.
2. To enable cflowd visibility so that a Cisco IOS XE SD-WAN device can perform traffic flow monitoring on traffic coming to the router from the LAN, click **Netflow**.

3. To enable application visibility so that a Cisco IOS XE SD-WAN device can monitor and track the applications running on the LAN, click **Application**.
4. To enable QoS scheduling and shaping for traffic that a Cisco IOS XE SD-WAN device receives from transport-side interfaces, click **Cloud QoS**.
5. To enable QoS scheduling and shaping for traffic that a Cisco IOS XE SD-WAN device receives from service-side interfaces, click **Cloud QoS Service Side**.
6. To log the headers of all packets that are dropped because they do not match a service configured by an Allow Service parameter on a tunnel interface, click **Implicit ACL Logging**.
7. To configure how often packets flows are logged, click **Log Frequency**. Packet flows are those that match an access list (ACL), a cflowd flow, or an application-aware routing flow.
8. Click **Preview** to view the full policy in CLI format.
9. Click **Save Policy**.

Step 6: Apply a Localized Data Policy in a Device Template

1. In 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 Cisco IOS XE SD-WAN 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. Click **Create** (for a new template) or **Update** (for an existing template).

Structural Components of Configuration for Access Lists

Following are the structural components required to configure access lists, shown as they appear in the CLI and when you click **Preview** in the Cisco vManage localized policy configuration wizard. Each component is explained in the sections below.

```

policy
  lists
    data-prefix-list list-name
    ip-prefix prefix/length
  class-map
    class class map map
  cloud-qos
  cloud-qos-service-side
  implicit-acl-logging
  log-frequency number
  qos-scheduler scheduler-name
    class class-name
    bandwidth-percent percentage
    buffer-percent percentage
    drops drop-type
    scheduling (llq | wrr)
  qos-map map-name
    qos-scheduler scheduler-name
  rewrite-rule rule-name
    class class-name priority dscp dscp-value layer-2-cos number
  mirror mirror-name
    remote-dest ip-address source ip-address
  policer policer-name
    rate bandwidth
    burst bytes
    exceed action
  access-list list-name
    sequence number
      match
        match-parameters
      action
        drop
          count counter-name
          log
        accept
          class class-name
          count counter-name
          log
          mirror mirror-name
          policer policer-name
          set dscp value
          set next-hop ipv4-address
      default-action
        (accept | drop)
  vpn vpn-id
    interface interface-name
      access-list list-name (in | out)
      policer policer-name (in | out)
      rewrite-rule rule-name

```

Lists

Access lists use prefix lists to group related prefixes.

In the Cisco vManage NMS, you configure prefix lists from:

- **Configuration > Policies > Localized Policy > Add Policy > Create Groups of Interest**

- **Configuration > Policies > Custom Options > Localized Policy > Lists > Data Prefix**

In the CLI, you configure lists under the **policy lists** command hierarchy on Cisco IOS XE SD-WAN devices.

Table 19:

List Type	Description	vManage Configuration/ CLI Configuration Command
Data prefixes	List of one or more IP prefixes. You can specify both unicast and multicast addresses. To configure multiple prefixes in a single list, include multiple ip-prefix options, specifying one prefix in each option.	Configuration > Policies > Localized Policy > Add Policy > Create Groups of Interest > Data Prefix > New Data Prefix List Configuration > Policies > Custom Options > Localized Policy > Lists > Data Prefix > New Data Prefix List data-prefix-list <i>list-name</i> ip-prefix <i>prefix/length</i>

QoS Parameters

In Cisco vManage NMS, you configure QoS parameters:

- **Configuration > Policies > Localized Policy > Add Policy > Create Groups of Interest > Class Map, or Configuration > Policies > Custom Options > Localized Policy > Lists > Class Map**
- **Configuration > Policies > Localized Policy > Add Policy > Configuring Forwarding Classes/QoS, or Configuration > Policies > Custom Options > Localized Policy > Configuring Forwarding Classes/QoS**
- **Configuration > Policies > Localized Policy > Add Policy > Policy Overview, or Configuration > Policies > Custom Options > Localized Policy > Policy Overview**

This section explains how to configure QoS parameters from the CLI.

To configure QoS parameters on a device, first define a classification. In Cisco vManage NMS:

```
Device(config)# policy class-map class class-name queue number
```

class-name is the name of the class. It can be a text string from 1 through 32 characters long.

For hardware, each interface has eight queues, numbered from 0 through 7. Queue 0 is reserved for low-latency queuing (LLQ), so any class that is mapped to queue 0 must be configured to use LLQ. The default scheduling method for all is weighted round-robin (WRR).

For Cisco IOS XE SD-WAN devices, each interface has eight queues, numbered from 0 through 7. Queue 0 is reserved for control traffic, and queues 1, 2, 3, 4, 5, 6 and 7 are available for data traffic. The scheduling method for all eight queues is WRR. LLQ is not supported.

To configure QoS parameters on a Cisco IOS XE SD-WAN device, you must enable QoS scheduling and shaping. To enable QoS parameters for traffic that the Cisco IOS XE SD-WAN device receives from transport-side interfaces:

```
Device(config)# policy cloud-qos
```

To enable QoS parameters for traffic that the Cisco IOS XE SD-WAN device receives from service-side interfaces:

```
Device(config)# policy cloud-qos-service-side
```

Next, configure scheduling:

```
Device(config)# policy qos-scheduler scheduler-name
Device(config-qos-scheduler)# class percentage
Device(config-qos-scheduler)# buffer-percent percentage
Device(config-qos-scheduler)# drops (red-drop | tail-drop)
Device(config-qos-scheduler)# scheduling (llq | wrr)
```

scheduler-name is the name of the QoS scheduler. It can be a text string from 1 through 32 characters long.

class-name is the name of the forwarding class and can be a text string from 1 through 32 characters long. The common class names correspond to the per-hop behaviors AF (assured forwarding), BE (best effort), and EF (expedited forwarding).

The bandwidth percentage is the percentage of the interface's bandwidth to allocate to the forwarding class. The sum of the bandwidth on all forwarding classes on an interface should not exceed 100 percent.

The buffer percentage is the percentage of the interface's buffering capacity to allocate to the forwarding class. The sum of the buffering capacity of all forwarding classes on an interface should not exceed 100 percent.

Packets that exceed the bandwidth or buffer percentage are dropped either randomly, using random early detection (**red-drop**), or from the end of the queue (**tail-drop**). Low-latency queuing (LLQ) cannot use random early detection.

The algorithm to schedule interface queues can be either low-latency queuing (**llq**) or weighted round-robin (**wrr**).

Then, assign the scheduler to a QoS map:

```
Device(config-policy)# qos-map map-name qos-scheduler scheduler-name
```

map-name is the name of the QoS map, and *scheduler-name* is the name of the scheduler you configured above. Each name can be a text string from 1 through 32 characters long.

Finally, to configure a rewrite rule to overwrite the DSCP field of a packet's outer IP header:

```
Device(config)# policy rewrite-rule rule-name class class-name loss-priority
dscp dscp-value layer-2-cos number
```

rule-name is the name of the rewrite rule. It can be a text string from 1 through 32 characters long.

class-name is the name of a class you configured with the **qos-scheduler class** command. The packet loss priority (PLP) can be either **high** or **low**. To have a DSCP value overwrite the DSCP field of the packet's outer IP header, set a value from 0 through 63. To include an 802.1p marking in the packet, specify a number from 0 through 7.

Policer Parameters

To configure policing parameters, create a policer that specifies the maximum bandwidth and burst rate for traffic on an interface, and how to handle traffic that exceeds these values.

In Cisco vManage NMS, you configure policer parameters from:

- **Configuration > Policies > Centralized Policy > Add Policy > Create Groups of Interest > Policer**
- **Configuration > Policies > Custom Options > Centralized Policy > Lists > Policer**

In the CLI, you configure policer parameters as follows:

```
vSmart(config)# policy policer policer-name
vSmart(config-policer)# rate bps
vSmart(config-policer)# burst bytes
vSmart(config-policer)# exceed action
```

rate is the maximum traffic rate. It can be a value from 0 through 264 – 1 bits per second.

burst is the maximum traffic burst size. It can be a value from 15000 to 1000000 bytes.

exceed is the action to take when the burst size or traffic rate is exceeded. *action* can be **drop** (the default) or **remark**. The **drop** action is equivalent to setting the packet loss priority (PLP) bit to low. The **remark** action sets the PLP bit to high. In centralized data policy, access lists, and application-aware routing policy, you can match the PLP with the **match plp** option.

Sequences

An access list contains sequences of match–action pairs. The sequences are numbered to set the order in which a packet is analyzed by the match–action pairs in the access lists.

In Cisco vManage NMS, you configure sequences from:

- **Configuration > Policies > Localized Policy > Add Policy > Configure Access Control Lists > Add Access Control List Policy > Add ACL Sequence**
- **Configuration > Policies > Custom Options > Localized Policy > Access Control List Policy > Add Access Control List Policy > Add ACL Sequence**

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

Each sequence in an access list can contain one match condition and one action condition.

Match Parameters

Access lists can match IP prefixes and fields in the IP headers.

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

- **Configuration > Policies > Localized Policy > Add Policy > Configure Access Control Lists > Add Access Control List Policy > Add ACL Sequence > Add Sequence Rule > Match**
- **Configuration > Policies > Custom Options > Localized Policy > Access Control List Policy > Add Access Control List Policy > Add ACL Sequence > Add Sequence Rule > Match**

In the CLI, you configure the match parameters with the **policy access-list sequence match** command.

Each sequence in an access-list must contain one match condition.

For access lists, you can match these parameters:

Table 20:

Description	vManage Configuration/ CLI Configuration Command	Value or Range
Classification map	Match Class class <i>class-name</i>	Name of a class defined with a policy class-map command.
Group of destination prefixes	Match Destination Data Prefix destination-data-prefix-list <i>list-name</i>	Name of a data-prefix-list list.
Individual destination prefix	Not available in vManage NMS destination-ip <i>prefix/length</i>	IP prefix and prefix length

Description	vManage Configuration/ CLI Configuration Command	Value or Range
Destination port number	Match Destination Port destination-port <i>number</i>	0 through 65535; specify a single port number, a list of port numbers (with numbers separated by a space), or a range of port numbers (with the two numbers separated with a hyphen [-])
DSCP value	Match DSCP dscp <i>number</i>	0 through 63
Internet Protocol number	Match Protocol protocol <i>number</i>	0 through 255
Packet length	Match Packet Length packet-length <i>number</i>	Length of the packet. <i>number</i> can be from 0 through 65535. Specify a single length, a list of lengths (with numbers separated by a space), or a range of lengths (with the two numbers separated with a hyphen [-])
Group of source prefixes	Match Source Data Prefix source-data-prefix-list <i>list-name</i>	Name of a data-prefix-list list.
Packet loss priority (PLP)	Match PLP plp	(high low) By default, packets have a PLP value of low . To set the PLP value to high , apply a policer that includes the exceed remark option.
Individual source prefix	Match Source Data Prefix source-ip <i>prefix/length</i>	IP prefix and prefix length
Source port number	Match Source Port source-port <i>address</i>	0 through 65535; specify a single port number, a list of port numbers (with numbers separated by a space), or a range of port numbers (with the two numbers separated with a hyphen [-])
TCP flag	Match TCP tcp <i>flag</i>	syn

Action Parameters

When a packet matches the conditions in the match portion of an access list, the packet can be accepted or dropped, and it can be counted. Then, you can classify, mirror, or police accepted packets.

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

- **Configuration > Policies > Localized Policy > Add Policy > Configure Access Control Lists > Add Access Control List Policy > Add ACL Sequence > Add Sequence Rule > Action**
- **Configuration > Policies > Custom Options > Localized Policy > Access Control List Policy > Add Access Control List Policy > Add ACL Sequence > Add Sequence Rule > Action**

In the CLI, you configure the actions parameters with the **policy access-list sequence action** command. Each sequence in an access list can contain one action condition.

In the action, you first specify whether to accept or drop a matching data packet, and whether to count it:

Table 21:

Description	vManage Configuration/ CLI Configuration Parameter	Value or Range
Accept the packet. An accepted packet is eligible to be modified by the additional parameters configured in the action portion of the access list.	Click Accept accept	—
Count the accepted or dropped packets.	Action Counter Click Accept, then action Counter count <i>counter-name</i>	Name of a counter. To display counter information, use the show policy access-lists counters command on the Cisco IOS XE SD-WAN device.
Discard the packet. This is the default action.	Click Drop drop	—

For a packet that is accepted, the following actions can be configured:

Table 22:

Description	vManage Configuration/ CLI Configuration Parameter	Value or Range
Classify the packet.	Click Accept, then Class class <i>class-name</i>	Name of a QoS class defined with a policy class-map command.
Mirror the packet.	Click Accept, then Mirror List mirror <i>mirror-name</i>	Name of mirror defined with a policy mirror command.
Police the packet.	Click Accept, then Policer policer <i>policer-name</i>	Name of a policer defined with a policy policer command.
Packet's DSCP value.	Click Accept, then DSCP set dscp <i>value</i>	0 through 63.
Next-hop address.	Click Accept, then Next Hop set next-hop <i>ipv4-address</i>	IPv4 address.

Default Action

If a packet being evaluated does not match any of the match conditions in a access list, a default action is applied to this packet. By default, the packet is dropped.

In the Cisco vManage NMS, you modify the default action from:

- **Configuration > Policies > Localized Policy > Add Policy > Configure Access Control Lists > Default Action**
- **Configuration > Policies > Custom Options > Localized Policy > Access Control List Policy > Default Action**

In the CLI, you modify this behavior with the **access-list default-action accept** command.

Apply Access Lists

For an access list to take effect, you must apply it to an interface.

In the Cisco vManage NMS, you apply the access list in one of these interface feature configuration templates:

- **Configuration > Templates > VPN Interface Bridge**
- **Configuration > Templates > VPN Interface Cellular**
- **Configuration > Templates > VPN Interface Ethernet**
- **Configuration > Templates > VPN Interface GRE**
- **Configuration > Templates > VPN Interface PPP**
- **Configuration > Templates > VPN Interface PPP Ethernet**

In the CLI, you apply the access list as follows:

```
Device(config)# vpn vpn-id interface interface-name
Device(config-interface)# access-list list-name (in|out)
```

Applying the policy in the inbound direction (**in**) affects prefixes being received on the interface. Applying it in the outbound direction (**out**) affects prefixes being transmitted on the interface.

For an access list that applies QoS classification, apply any DSCP rewrite rules to the same interface to which you apply the access list:

```
Device(config)# vpn vpn-id interface interface-name rewrite-rule rule-name
```

Note that you can also apply a policer directly to an interface, which has the effect of policing all packets transiting the interface, rather than policing only the selected packets that match the access list. You can apply the policer to either inbound or outbound packets:

```
Device(config)# vpn vpn-id interface interface-name
Device(config-interface)# policer
policer-name (in|out) interface-name
```

Explicit and Implicit Access Lists

Access lists that you configure through localized data policy using the **policy access-list** command are called *explicit ACLs*. You can apply explicit ACLs to any interface in any VPN on the device.

The device's tunnel interfaces in VPN 0 also have *implicit ACLs*, which are also referred to as *services*. Some services are enabled by default on the tunnel interface, and are in effect unless you disable them. Through configuration, you can also enable other services. You configure and modify implicit ACLs with the **allow-service** command:

```
Device(config)# vpn 0
Device(config-vpn)# interface interface-name
Device(config-interface)# tunnel-interface
```

```
Device(config-tunnel-interface)# allow-service service-name
Device(config-tunnel-interface)# no allow-service service-name
```

On Cisco IOS XE SD-WAN devices, the following services are enabled by default: DHCP (for DHCPv4 and DHCPv6), DNS, and ICMP. These three services allow the tunnel interface to accept DHCP, DNS, and ICMP packets. You can also enable services for BGP, Netconf, NTP, OSPF, SSHD, and STUN.

**Note**

If a connection is initiated from a device, and if NAT is enabled on the device (for example, Direct Internet Access (DIA) is configured), return traffic is allowed by the NAT entry even if the implicit ACL has been configured as **no allow-service**. You can still block this traffic with an explicit ACL.

Do not confuse an explicit ACL with an IOS XE ACL. An IOS XE ACL does not interact with a Cisco SD-WAN explicit and an implicit ACL and cannot override an implicit ACL or explicit ACL. IOS XE ACLs are executed later in the order of traffic processing operations.

When data traffic matches both an explicit ACL and an implicit ACL, how the packets are handled depends on the ACL configuration. Specifically, it depends on:

- Whether the implicit ACL is configured as allow (**allow-service** *allow-service*) or deny (**no allow-service** *service-name*). Allowing a service in an implicit ACL is the same as specifying the **accept** action in an explicit ACL, and a service that is not allowed in an implicit ACL is the same as specifying the **drop** action in an explicit ACL
- Whether, in an explicit ACL, the **accept** or **deny** action is configured in a policy sequence or in the default action.

The following table explains how traffic matching both an implicit and an explicit ACL is handled:

Table 23:

Implicit ACL	Explicit ACL: Sequence	Explicit ACL: Default	Result
Allow (accept)	Deny (drop)	—	Deny (drop)
Allow (accept)	—	Deny (drop)	Allow (accept)
Deny (drop)	Allow (accept)	—	Allow (accept)
Deny (drop)	—	Allow (accept)	Deny (drop)

Configure Localized Data Policy for IPv4 Using the CLI for Cisco IOS XE SD-WAN Devices

Following are the high-level steps for configuring an access list using the CLI on Cisco IOS XE SD-WAN devices:

1. Create lists of IP prefixes, as needed:

```
Device(config)# policy lists data-prefix-list ipv4_prefix_list
Device(config-data-prefix-list-ipv4_prefix_list)
# ip-prefix 192.168.0.3/24
```


2. For QoS, configure the **class-map ios**:

```
Device(config)# class-map match-any class1
Device(config)# match qos-group 1
class-map match-any class6
match qos-group 6
class-map match-any class7
match qos-group 7
class-map match-any class4
match qos-group 4
class-map match-any class5
match qos-group 5
class-map match-any class2
match qos-group 2
class-map match-any class3
match qos-group 3
class-map match-any class1
match qos-group 1
end
```



Note queue2 is optional here since we are using **class-default**.

3. For QoS, define rewrite rules to overwrite the DSCP field of a packet's outer IP header, if desired:

```
Device(config)# policy rewrite-rule rule1
Device(config-rewrite-rule-rule1)# class class1 low dscp 3
Device(config-rewrite-rule-rule1)# class class2 high dscp 4
Will be a table to map class-id → QoS-Group, QID, DSCP, Discard-Class
```

4. For QoS, map each forwarding class to an output queue, configure a QoS scheduler for each forwarding class, and group the QoS schedulers into a QoS map:

```
Device(config)# policy class-map class class1 queue 1
<0..7>[1]
```

5. For QoS map configuration, merge with interface shaping configuration, if shaping is configured.

If shaping is not configured, you can apply the **policy-map** generated for the **qos-map**.

```
Device(config)# policy-map qos_map_for_data_policy
<name:string>
Device(config-pmap)# class class1<name:string>
Device(config-pmap-c)# bandwidth<percentage>
Device(config-pmap-c)# random-detect
```

6. Configure a WAN interface without a shaping configuration:

```
Device(config)# policy-map qos_map_for_data_policy <name:string>
Device(config-pmap)# class class1<name:string>
Device(config-pmap-c)# bandwidth<percentage>
Device(config-pmap-c)# random-detect
```

7. Configure a WAN interface with a shaping configuration:

```
Device(config)# policy-map shaping_interface
Device(config-pmap)# class class-default
Device(config-pmap-c)# shape average 100000000(rate-in-bps)
Device(config-pmap-c)# service-policy qos_map_for_data_policy
```

8. Associate a **service-policy** to a Cisco IOS XE SD-WAN device:

```
Device(config)# sdwan interface GigabitEthernet 1
Device(config-if)# rewrite-rule rule1
Device(config-if)# service-policy output qos_map_for_data_policy
```

9. Define policing parameters:

```
Device(config)# policy policer policer_On_gige
Device(config-policer-policer_On_gige)# rate ?
Description: Bandwidth for 1g interfaces: <8..1000000000>bps; for 10g interfaces:
<8..10000000000>bps
Possible completions:<0..2^64-1>
Device(config-policer-policer_On_gige)# burst
Description: Burst rate, in bytes
Possible completions:<15000..10000000>
Device(config-policer-policer_On_gige)# exceed drop
```

10. Associate an access list set to policer:

```
Device(config)# policy access-list ipv4_acl
Device(config-access-list-ipv4_acl)# sequence 100
Device(config-sequence-100)# match dscp 10
Device(config-match)# exit
Device(config-sequence-100)# action accept
Device(config-sequence-100)# action count dscp_10_count
Device(config-sequence-100)# policer policer_On_gige
Device(config-sequence-100)# action drop
vm5(config-action)#
```

11. Associate an access list to a LAN or a WAN interface:

```
Device(config)# sdwan interface GigabitEthernet5
Device(config-interface-GigabitEthernet5)# access-list ipv4_acl
Device(config-interface-GigabitEthernet5)# commit
```

Localized Data Policy for IPv6

This topic provides procedures for configuring IPv6 localized data policy. This type of data policy is called access lists, or ACLs. You can provision simple access lists that filter traffic based on IP header fields. You also use access lists to apply policing to data packets.

For IPv6, you can apply access lists only to interfaces in the transport VPN, VPN 0.

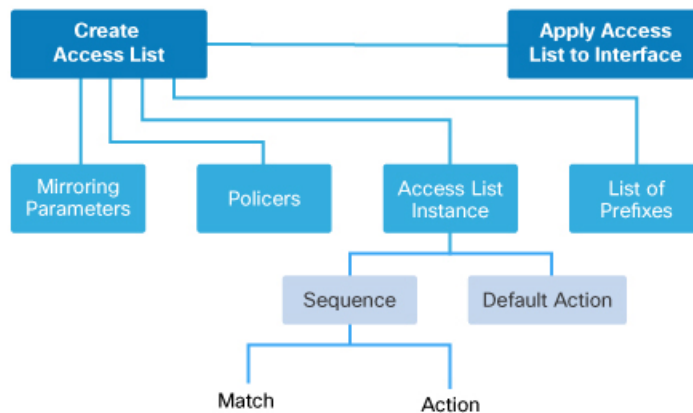
In Cisco vManage NMS, you configure localized data policy from the **Configuration > Policies** screen, using a policy configuration wizard. In the CLI you configure these policies on the Cisco IOS XE SD-WAN device.

Configuration Components

An access list consists of a sequence of match–action pairs 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 packet stops. Keep this in mind as you design your policies to ensure that the desired actions are taken on the items subject to policy.

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

The following figure illustrates the configuration components for IPv6 access lists:



368534

Configure Localized Data Policy for IPv6 Using vManage

To configure IPv6 localized data policy, use the Cisco vManage policy configuration wizard. The wizard is a UI policy builder that consists of five screens, and you use four of them to configure IPv6 localized policy components:

- Groups of Interest, also called *lists*—Create data prefix lists and policer parameters that group together related items and that you call in the match or action components of a policy.
- Access Control Lists—Define the match and action conditions of ACLs.
- Route Policies—Define the match and action conditions of route policies.
- Policy Settings—Define additional 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.

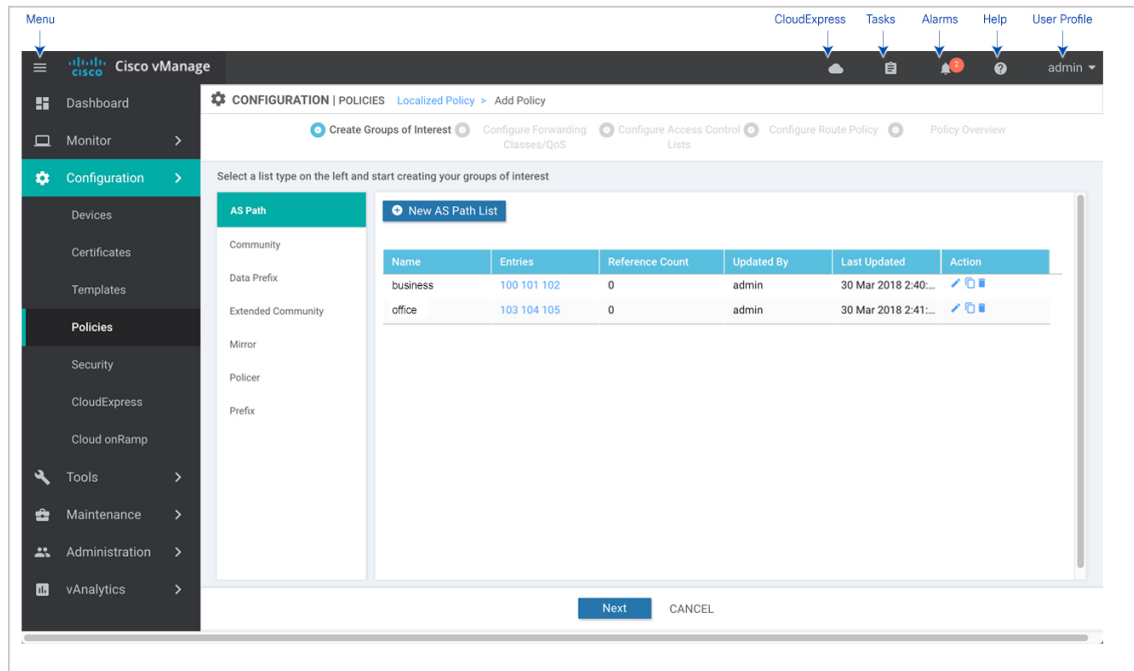
Step 1: Start the Policy Configuration Wizard

To start the policy configuration wizard:

1. In Cisco vManage NMS, select the **Configure > Policies** screen. When you first open this screen, the Centralized Policy tab is selected by default.
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: Create Groups of Interest

In the Create Groups of interest screen create lists to use in the localized data policy:



368398

1. Create news lists of groups, as described in the following table:

List Type	Procedure
AS Path	<p>Permit or deny prefixes from certain autonomous systems.</p> <ol style="list-style-type: none"> In the left bar, click AS Path. Enter a name for the list. For Cisco IOS XE SD-WAN devices: Enter a number from 1 to 500. Set the preference value for the list in the Add AS Path field.
Community	<ol style="list-style-type: none"> In the left bar, click Community. Click New Community List. Enter a name for the list. In the Add Community field, enter one or more data prefixes separated by commas. Click Add.

List Type	Procedure
Data Prefix	<ol style="list-style-type: none"> a. In the left bar, click Data Prefix. b. Click New Data Prefix List. c. Enter a name for the list. d. In the Internet Protocol field, click IPv4 or IPv6. e. In the Add Data prefix field, enter one or more data prefixes separated by commas. f. Click Add.
Extended Community	<ol style="list-style-type: none"> a. In the left bar, click Extended Community. b. Click New Extended Community List. c. Enter a name for the list. d. In the Add Extended Community field, enter one or more data prefixes separated by commas. e. Click Add.
Class Map	<p>Map a class name to an interface queue number.</p> <ol style="list-style-type: none"> a. In the left bar, click Class Map. b. Click New Class List. The Class List popup displays. c. Enter a name for the list. The class name can be a text string from 1 to 32 characters long. d. Select a queue number between 0 and 7 from the Queue drop-down menu. e. Click Save.
Mirror	<p>Define the remote destination for mirrored packets, and define the source of the packets.</p> <ol style="list-style-type: none"> a. In the left bar, click Mirror. b. Click New Mirror List. c. Enter a name for the list. d. Enter the Remote Destination IP address in the left field, where the mirrored traffic should be routed. e. Enter the Source IP address of the mirrored traffic in the right field. f. Click Add.

List Type	Procedure
Policer	<ol style="list-style-type: none"> a. In the left bar, click Policer. b. Click New Policer List. c. Enter a name for the list. d. Define the policing parameters: <ol style="list-style-type: none"> 1. In the Burst field, enter the maximum traffic burst size, a value from 15,000 to 10,000,000 bytes. 2. In the Exceed field, select the action to take when the burst size or traffic rate is exceeded. It can be drop, which sets the packet loss priority (PLP) to low, or remark, which sets the PLP to high. 3. In the Rate field, enter the maximum traffic rate, a value from 0 through 264 – 1 bits per second (bps). e. Click Add.
Prefix	<ol style="list-style-type: none"> a. In the left bar, click Prefix. b. Click New Prefix List. c. Enter a name for the list. d. Click either IPv4 or IPv6. e. Under Add Prefix, enter the prefix for the list. (An example is displayed.) Optionally, click the green Import link on the right-hand side to import a prefix list. f. Click Add.

2. Click **Next** to move to Configure Forwarding Classes/QoS in the wizard. For IPv6 localized data policy, you cannot configure QoS.
3. Click **Next** to move to Configure Access Lists in the wizard.

Step 3: Configure ACLs

1. In the Configure Access Control Lists screen, click **Add Access Control List Policy**, and choose **Add IPv6 ACL Policy** from the drop-down.
2. Enter a name and description for the ACL.
3. From the left column, click **Add ACL Sequence**.
4. Click **Sequence Rule** to open the ACL match/action sequence menu.
5. Click a match condition. See [Match Parameters](#) for a full description of these options.
6. On the left side, enter the values for the match condition.

7. On the right side, enter the action or actions to take if the policy matches. See [Action Parameters](#) for a full description of these options.
8. Repeat Steps 3 through 7 to add match–action pairs to the ACL.
9. To rearrange match–action pairs in the ACL, drag them to the desired position in the right pane.
10. To remove a match–action pair from the ACL, click the X in the upper right of the condition.
11. Click **Save Match and Actions** to save a sequence rule.
12. To copy, delete, or rename an ACL sequence rule, in the left pane, click the **More Options** menu (three dots) next to the rule's name and select the desired option.
13. If no packets match any of the ACL 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**.
14. Click **Next** to move to Configure Route Policy in the wizard.
15. Click **Next** to move to the Policy Overview screen.

Step 4: Configure Policy Settings

In Policy Overview, configure policy settings:

1. Enter a name and description for the ACL.
2. Under **Policy Settings**, select one of the following policy options:

Policy Settings Options	Description
Netflow	
Application	
Cloud QoS	
Cloud QoS Service side	
Implicit ACL Logging	Log the headers of all packets that are dropped because they do not match a service configured by an Allow Service parameter on a tunnel interface.

3. Click **Preview** to view the full policy in CLI format.
4. Click **Save Policy**.

Step 5: Apply a Localized Data Policy in a Device Template

1. In 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 a Cisco IOS XE SD-WAN device.
 - 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. Click **Create** (for a new template) or **Update** (for an existing template).

Structural Components of Configuration for Access Lists

Following are the structural components required to configure access lists. Each one is explained in more detail in the sections below.

```

policy
  remote-dest ip-address source ip-address
  policer policer-name
    rate bandwidth
    burst bytes
    exceed action
policy ipv6
  access-list list-name
    sequence number
    match match-parameters
    action
      drop
      count counter-name
      log
      accept
      class class-name
      mirror mirror-name

```



```

        policer policer-name
        default-action (accept | drop)
vpn vpn-id
        interface interface-name
        ipv6 access-list list-name (in | out)

```

Policer Parameters

To configure policing parameters, create a policer that specifies the maximum bandwidth and burst rate for traffic on an interface, and how to handle traffic that exceeds these values.

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

- **Configuration > Policies > Localized Policy > Add Policy > Create Groups of Interest > Policer > New Policer List**
- **Configuration > Policies > Custom Options > Localized Policy > Lists > Policer > New Policer List**

In the CLI, you configure policer parameters as follows:

```

Device(config)# policy policer policer-name
Device(config-policer)# rate bps
Device(config-policer)# burst bytes
Device(config-policer)# exceed action

```

- **rate** is the maximum traffic rate. It can be a value from 0 through $2^{64} - 1$ bits per second.
- **burst** is the maximum traffic burst size. It can be a value from 15000 to 1000000 bytes.
- **exceed** is the action to take when the burst size or traffic rate is exceeded. *action* can be **drop** (the default) or **remark**. The **drop** action is equivalent to setting the packet loss priority (PLP) bit to low. The **remark** action sets the PLP bit to high. In centralized data policy, access lists, and application-aware routing policy, you can match the PLP with the **match plp** option.

Sequences

Sequences

An access list contains sequences of match–action pairs. The sequences are numbered to set the order in which a packet is analyzed by the match–action pairs in the access lists.

In the Cisco vManage NMS, you configure sequences from:

- **Configuration > Policies > Localized Policy > Add Policy > Configure Access Control Lists > Add Access Control List Policy > Add ACL Sequence**
- **Configuration > Policies > Custom Options > Localized Policy > Access Control List Policy > Add Access Control List Policy > Add ACL Sequence**

In the CLI, you configure sequences with the **policy ipv6 access-list sequence** command.

Each sequence in an access list can contain one match condition and one action condition.

Match Parameters

Access lists can match IP prefixes and fields in the IP headers.

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

- **Configuration > Policies > Localized Policy > Add Policy > Configure Access Control Lists > Add Access Control List Policy > Add ACL Sequence > Add Sequence Rule > Match**
- **Configuration > Policies > Custom Options > Localized Policy > Access Control List Policy > Add Access Control List Policy > Add ACL Sequence > Add Sequence Rule > Match**

In the CLI, you configure the match parameters with the **policy ipv6 access-list sequence match** command. Each sequence in an access list must contain one match condition.

For access lists, you can match these parameters:

Description	vManage Match Tab / CLI Command	Value or Range
Enter a Destination port number.	Destination Port destination-port <i>number</i>	0 through 65535; specify a single port number, a list of port numbers (with numbers separated by a space), or a range of port numbers (with the two numbers separated with a hyphen [-])
Select the Next Header protocol.	Protocol next-header <i>number</i>	0 through 255, corresponding to an Internet Protocol number
Specify the packet length	Packet Length packet-length <i>number</i>	Length of the packet. <i>number</i> can be from 0 through 65535. Specify a single length, a list of lengths (with numbers separated by a space), or a range of lengths (with the two numbers separated with a hyphen [-])
Specify the packet loss priority (PLP)	PLP plp	(high low) By default, packets have a PLP value of low . To set the PLP value to high , apply a policer that includes the exceed remark option.
Select a Source data prefix list		
Enter a Source port number	Source Port source-port <i>address</i>	0 through 65535; specify a single port number, a list of port numbers (with numbers separated by a space), or a range of port numbers (with the two numbers separated with a hyphen [-])
Enter a Destination Data Prefix		
TCP	TCP tcp <i>flag</i>	syn
Set the packet's DSCP value	Class set class <i>value</i>	0 through 63
Traffic class	Traffic Class traffic-class <i>value</i>	0 through 63

Action Parameters

When a packet matches the conditions in the match portion of an access list, the packet can be accepted or dropped, and it can be counted. Then, you can classify, mirror, or police accepted packets.

In Cisco vManage NMS, you configure match parameters from:

- **Configuration > Policies > Localized Policy > Add Policy > Configure Access Control Lists > Add Access Control List Policy > Add ACL Sequence > Add Sequence Rule > Action**
- **Configuration > Policies > Custom Options > Localized Policy > Access Control List Policy > Add Access Control List Policy > Add ACL Sequence > Add Sequence Rule > Action**

In the CLI, you configure the actions parameters with the **policy ipv6 access-list sequence action** command.

Each sequence in an access list can contain one action condition.

In the action, you first specify whether to accept or drop a matching data packet, and whether to count it:

For a packet that is accepted, the following actions can be configured:

Description	vManage Action Tab / CLI Command	Value or Range
Accept the packet. An accepted packet is eligible to be modified by the additional parameters configured in the action portion of the access list.	Click Accept . accept	—
Count the accepted or dropped packets.	Counter Name count <i>counter-name</i>	Name of a counter. To display counter information, use the show ipv6 policy access-lists counters command on the Cisco IOS XE SD-WAN device.
Designate the next hop router.	Next Hop	
Traffic Class	set traffic-class <i>value</i>	0-63
Mirror the packet.	Mirror List mirror <i>mirror-name</i>	Name of mirror defined with a policy mirror command.
Set the packet's DSCP value.	Class	0 through 63
Police the packet.	Policer policer <i>policer-name</i>	Name of a policer defined with a policy policer command.
Discard the packet. This is the default action.	Click Drop . drop	—

Default Action

If a packet being evaluated does not match any of the match conditions in a access list, a default action is applied to this packet. By default, the packet is dropped.

In the Cisco vManage NMS, you modify the default action from:

- **Configuration > Policies > Localized Policy > Add Policy > Configure Access Control Lists > Default Action**
- **Configuration > Policies > Custom Options > Localized Policy > Access Control List Policy > Default Action**

In the CLI, you modify this behavior with the `access-list ipv6 default-action accept` command.

Apply Access Lists

For an access list to take effect, you must apply it to a tunnel interface in VPN 0.

In the Cisco vManage NMS, you apply the access list in one of the interface feature configuration templates.

In the CLI, you apply the access list as follows:

```
vEdge(config)# vpn 0 interface interface-name
vEdge(config-interface)# ipv6 access-list list-name (in | out)
```

Applying the policy in the inbound direction (**in**) affects prefixes being received on the interface. Applying it in the outbound direction (**out**) affects prefixes being transmitted on the interface.

Explicit and Implicit Access Lists

Access lists that you configure through localized data policy using the `policy access-list` command are called *explicit* ACLs. You can apply explicit ACLs to any interface in any VPN on the router.

The router's tunnel interfaces in VPN 0 also have implicit ACLs, which are also referred to as *services*. Some services are enabled by default on the tunnel interface, and are in effect unless you disable them. Through configuration, you can also enable other services. You configure and modify implicit ACLs with the `allow-service` command:

```
Device(config)# vpn 0
Device(config-vpn)# interface interface-name
Device(config-interface)# tunnel-interface
Device(config-tunnel-interface)# allow-service service-name
Device(config-tunnel-interface)# no allow-service service-name
```

On Cisco IOS XE SD-WAN devices, the following services are enabled by default: DHCP (for DHCPv4 and DHCPv6), DNS, and ICMP. These three services allow the tunnel interface to accept DHCP, DNS, and ICMP packets. You can also enable services for BGP, Netconf, NTP, OSPF, SSHD, and STUN.



Note

If a connection is initiated from a device, and if NAT is enabled on the device (for example, Direct Internet Access (DIA) is configured), return traffic is allowed by the NAT entry even if the implicit ACL has been configured as `no allow-service`. You can still block this traffic with an explicit ACL.

Do not confuse an explicit ACL with an IOS XE ACL. An IOS XE ACL does not interact with a Cisco SD-WAN explicit and an implicit ACL and cannot override an implicit ACL or explicit ACL. IOS XE ACLs are executed later in the order of traffic processing operations.

When data traffic matches both an explicit ACL and an implicit ACL, how the packets are handled depends on the ACL configuration. Specifically, it depends on:

- Whether the implicit ACL is configured as allow (`allow-service allow-service`) or deny (`no allow-service service-name`). Allowing a service in an implicit ACL is the same as specifying the

accept action in an explicit ACL, and a service that is not allowed in an implicit ACL is the same as specifying the **drop** action in an explicit ACL.

- Whether, in an explicit ACL, the **accept** or **deny** action is configured in a policy sequence or in the default action.

The following table explains how traffic matching both for an implicit and an explicit ACL is handled:

Implicit ACL	Explicit ACL: Sequence	Explicit ACL: Default	Result
Allow (accept)	Deny (drop)	—	Deny (drop)
Allow (accept)	—	Deny (drop)	Allow (accept)
Deny (drop)	Allow (accept)	—	Allow (accept)
Deny (drop)	—	Allow (accept)	Deny (drop)

Configure Localized Data Policy for IPv6 Using the CLI

Following are the high-level steps for configuring an access list using the CLI:

1. Define policing parameters:

```
Device(config)# policy policer policer_On_gige
Device (config-policer-policer_On_gige)# rate ?
Description: Bandwidth for lg interfaces: <8..1000000000>bps;for 10g interfaces:
<8..10000000000>bps Possible completions: <0..2^64-1>
Device(config-policer-policer_On_gige)# burst
Description: Burst rate, in bytes Possible completions:<15000..10000000>
Device(config-policer-policer_On_gige)# exceed drop
```

2. Create an access list instance:

```
Device (config)# policy ipv6 access-list ipv6_access_list
```

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

```
Device(config-access-list-ipv6_access_list)# sequence 100
```

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).

4. Define match parameters for packets:

```
Device(config-sequence-100)# match traffic-class 10
Device(config-match)# exit
```

5. Define actions to take when a match occurs:

```
Device(config-sequence-100)# action accept count traffic_class10_count
Device(config-sequence-100)# action drop
Device(config-sequence-100)# action accept class class1
Device(config-sequence-100)# action accept policer policer_On_gige
```

6. Create additional numbered sequences of match–action pairs within the access list, as needed.

7. If a packet does not match any of the conditions in one of the sequences, it is rejected by default. If you want nonmatching packets to be accepted, configure the default action for the access list:
8. Apply the access list to an interface:

```
Device(config)# sdwan interface GigabitEthernet5
Device(config-interface-GigabitEthernet5)
# ipv6 access-list ipv6_access_list in
Device(config-interface-GigabitEthernet5)
# commit
```

Applying the access list in the inbound direction (**in**) affects packets being received on the interface. Applying it in the outbound direction (**out**) affects packets being transmitted on the interface.

Localized Data Policy Configuration Examples

This topic provides some straightforward examples of configuring localized data policy to help you get an idea of how to use policy to influence traffic flow across the Cisco SD-WAN domain. Localized data policy, also known as access lists, is configured directly on the local Cisco vEdge devices.

QoS

You can configure quality of service (QoS) to classify data packets and control how traffic flows out of and in to the interfaces on a Cisco vEdge device and on the interface queues. For examples of how to configure a QoS policy, see Forwarding and QoS Configuration Examples.



CHAPTER 6

Policy Basics CLI Reference

CLI commands for configuring and monitoring policy.

Centralized Control Policy Command Hierarchy

Configure on Cisco vSmart Controllers only.

```
policy
  lists
    color-list list-name
      color color
    prefix-list list-name
      ip-prefix prefix/length
    site-list list-name
      site-id site-id
    tloc-list list-name
      tloc address color color encap encapsulation [preference value weight value]
    vpn-list list-name
      vpn vpn-id

policy
  control-policy policy-name
  default-action action
  sequence number
  match
    route
      color color
      color-list list-name
      omp-tag number
      origin protocol
      originator ip-address
      preference number
      prefix-list list-name
      site-id site-id
      site-list list-name
      tloc address
      tloc-list list-name
      vpn vpn-id
      vpn-list list-name
    tloc
      carrier carrier-name
      color color
      color-list list-name
      domain-id domain-id
      group-id group-id
      omp-tag number
      originator ip-address
      preference number
```

```

        site-id site-id
        site-list list-name
        tloc address
        tloc-list list-name
    action
        reject
        accept
            export-to (vpn vpn-id | vpn-list list-name)
            set
                omp-tag number
                preference value
                service service-name (tloc ip-address | tloc-list list-name) [vpn vpn-id]
                tloc-action action
                tloc-list list-name

    apply-policy
        site-list list-name control-policy policy-name (in | out)

```

Centralized Data Policy Command Hierarchy

Configure on Cisco vSmart Controllers only.

```

policy
  lists
    app-list list-name
      (app applications | app-family application-families)
    data-prefix-list list-name
      ip-prefix prefix/length
    site-list list-name
      site-id site-id
    tloc-list list-name
      tloc ip-address color color encap encapsulation [preference value weight value]
    vpn-list list-name
      vpn vpn-id

  policy
    data-policy policy-name
    vpn-list list-name
    default-action action
    sequence number
    match
      app-list list-name
      destination-data-prefix-list list-name
      destination-ip prefix/length
      destination-port number
      dns (request | response)
      dns-app-list list-name
      dscp number
      packet-length number
      plp (high | low)
      protocol number
      source-data-prefix-list list-name
      source-ip prefix/length
      source-port number
      tcp flag
    action
      cflowd
      count counter-name
      drop
      log
      tcp-optimization
      accept
        nat [pool number] [use-vpn-0]
        redirect-dns (host | ip-address)
        set

```



```

        dscp number
        forwarding-class class
        local-tloc color color [encap encapsulation]
        local-tloc-list color color [encap encapsulation] [restrict]
        next-hop ip-address
        policer policer-name
        service service-name local [restrict] [vpn vpn-id]
        service service-name [tloc ip-address | tloc-list list-name] [vpn vpn-id]
        tloc ip-address color color [encap encapsulation]
        tloc-list list-name
        vpn vpn-id
    vpn-membership policy-name
    default-action action
    sequence number
    match
        vpn vpn-id
        vpn-list list-name
    action
        (accept | reject)

apply-policy
    site-list list-name data-policy policy-name (all | from-service | from-tunnel)
    site-list list-name vpn-membership policy-name

```

Operational Commands

show running-config



CHAPTER 7

Forward Error Correction

Forward Error Correction (FEC) is a mechanism to recover lost packets on a link by sending extra “parity” packets for every group (N) of packets. As long as the receiver receives a subset of packets in the group (at-least N-1) and the parity packet, up to a single lost packet in the group can be recovered. FEC is supported on Cisco IOS XE SD-WAN devices.

Table 24: Feature History

Release	Description
Cisco IOS XE SD-WAN Release 16.11.x	Feature introduced. Forward Error Correction (FEC) is a mechanism to recover lost packets on a link by sending extra “parity” packets for every group (N) of packets.

- [Configure Forward Error Correction for a Policy, on page 155](#)
- [Monitor Forward Error Correction Tunnel Information, on page 156](#)
- [Monitor Forward Error Application Family Information, on page 157](#)

Configure Forward Error Correction for a Policy

- Step 1** Select **Configuration > Policies**.
- Step 2** Select **Centralized Policy** at the top of the page and then click **Add Policy**.
- Step 3** Click **Next** to select Configure Traffic Rules
- Step 4** Select **Traffic Data**, and from the Add Policy drop-down menu select click **Create New**.
- Step 5** Click **Sequence Type** in the left panel.
- Step 6** From the Add Data Policy pop-up menu, select **QoS**.
- Step 7** Click **Sequence Rule**.
- Step 8** In the **Applications/Application Family List/Data Prefix**, Select one or more applications or lists..
- Step 9** Click **Actions** and select **Loss Correction**.
- Step 10** In the Actions area, select one of the following:

- **FEC Adaptive**—Only send FEC information only when the system detects a packet loss.

Note The **FEC Adaptive** option is not supported on Cisco IOS XE SD-WAN devices.

Note FEC adaptive only works when the **app-route** interval is set at least twice that of the BFD Hello packet interval.

- **FEC Always**—Always send FEC information with every transmission
- **Packet Duplication** check box—Duplicates packets through secondary links to reduce packet loss if one link goes down

Step 11 Click **Save Match and Actions**.

Step 12 Click **Save Data Policy**.

Step 13 Click **Next** and take these actions to create a Centralized Policy:

- a) Enter a Name and Description.
- b) Select **Traffic Data Policy**.
- c) Choose VPNs/site list for the policy.
- d) Save the policy.

Monitor Forward Error Correction Tunnel Information

Step 1 Select **Monitor > Network**.

Step 2 Select a device group.

Step 3 In the left panel, click **Tunnel**, which displays under WAN

The WAN tunnel information includes the following:

- A graph that shows the total tunnel loss for the selected tunnels.
- A graph that shows the FEC loss recovery rate for the selected tunnels. The system calculates this rate by dividing the total number of reconstructed packets by the total number of lost packets on FEC:
- A table that provides the following information for each tunnel endpoint:
 - Name of the tunnel endpoint
 - Communications protocol that the endpoint uses
 - State of the endpoint
 - Jitter, in ms, on the endpoint
 - Packet loss percentage for the endpoint
 - FEC loss recovery percentage for the endpoint
 - Latency, in ms, on the endpoint
 - Total bytes transmitted from the endpoint
 - Total bytes received by the endpoint

- Application usage link
-

Monitor Forward Error Application Family Information

Step 1 Select **Monitor** > **Network**.

Step 2 Select a device group.

Step 3 In the left panel, click **DPI Applications**, which displays under **Applications**.

The FEC Recovery Rate application information includes the following:

- A graph for which you can select any of the following perspectives:
 - Application Usage—Usage of various types of traffic for the selected application families, in KB.
 - FEC Recovery Rate—FEC loss recovery rate for the selected application families. The system calculates this rate by dividing the total number of reconstructed packets by the total number of lost FEC-enabled packets.
 - A table that provides the following for each application family:
 - Name of the application family.
 - Packet Delivery Performance for the application family.
 - Note** If you need to see the packet delivery performance for the selected application family, ensure that packet duplication is enabled. Packet delivery performance is calculated based on the formula as displayed in the Cisco vManage tooltip for the **Packet Delivery Performance** column.
 - Traffic usage, in KB, MB, or GB for the selected application family.
-



CHAPTER 8

Packet Duplication for Noisy Channels

Table 25: Feature History

Feature Name	Release Information	Description
Packet Duplication for Noisy Channels	Cisco IOS XE SD-WAN Release 16.12.1b	This feature helps mitigate packet loss over noisy channels, thereby maintaining high application QoE for voice and video.

- [Information about Packet Duplication, on page 159](#)
- [Configure Packet Duplication, on page 159](#)
- [Monitor Packet Duplication Per Application, on page 160](#)

Information about Packet Duplication

Cisco IOS XE SD-WAN devices use packet duplication to overcome packet loss.

Packet duplication sends copies of packets on alternate available paths to reach Cisco IOS XE SD-WAN devices. If one of the packets is lost, a copy of the packet is forwarded to the server. Receiving Cisco IOS XE SD-WAN devices discard copies of the packet and forward one packet to the server.

Packet duplication is suitable for edges with multiple access links. Once packet duplication is configured and pushed to your device, you can see the tunnel packet duplication statistics.

Configure Packet Duplication

- Step 1** Select **Configuration > Policies**.
- Step 2** Select **Centralized Policy** at the top of the page and then click **Add Policy**.
- Step 3** Click **Next** twice to select Configure Traffic Rules.
- Step 4** Select **Traffic Data**, and from the Add Policy drop-down menu, select click **Create New**.
- Step 5** Click **Sequence Type** in the left panel.
- Step 6** From the Add Data Policy pop-up menu, select **QoS**.
- Step 7** Click **Sequence Rule**.
- Step 8** In the **Applications/Application Family List/Data Prefix**, Select one or more applications or lists..

- Step 9** Click **Actions** and select **Loss Correction**.
- Step 10** In the Actions area, select the **Pack Duplication** option to enable the packet duplication feature.
- **FEC Adaptive**—Only send Forward Error Correction (FEC) information when the system detects a packet loss.
 - **FEC Always**—Always send FEC information with every transmission.
 - **None**—Use when no loss protection is needed.
 - **Packet Duplication**—Enable when packets need to be duplicated and sent on the next available links to reduce packet loss.
- Step 11** Click **Save Match and Actions**.
- Step 12** Click **Save Data Policy**.
- Step 13** Click **Next** and take these actions to create a Centralized Policy:
- Enter a Name and a Description.
 - Select **Traffic Data Policy**.
 - Choose **VPNs/site list** for the policy.
 - Save the policy.
-

Monitor Packet Duplication Per Application

- Step 1** Select **Monitor > Network**.
- Step 2** Select a device group.
- Step 3** In the left panel, click **Applications**.
- Step 4** On the Application usage tab, select the application family of interest, and click on the Application family listed.
- Step 5** If packet duplication is enabled for any application, vManage displays Packet Delivery Performance as GOOD, MODERATE, or POOR or the field displays as N/A.
- Step 6** GOOD and MODERATE performance is a clickable link. When clicking on the link, the status pops up a window.
- Step 7** On the pop-up window, you see Application, Packet Delivery Performance, Overall for the Application, Average Drop Rate, and Overall for the Application information. The time slot graph represents the packets transmitted with different available link colors and the overall performance calculated when packet duplication is enabled.
- Step 8** If you hover over the time slot, you can see the performance status and the average drop rate for each link.
-



CHAPTER 9

Integrate Cisco IOS XE SD-WAN Device with Cisco ACI

Table 26: Feature History

Feature Name	Release Information	Description
Integration with Cisco ACI	Cisco IOS XE SD-WAN Release 16.12.1b	The Cisco IOS XE SD-WAN and Cisco ACI integration functionality now supports predefined SLA cloud beds. It also supports dynamically generated mappings from a data prefix-list and includes a VPN list to an SLA class that is provided by Cisco ACI.

Cisco ACI release 4.1(1) adds support for WAN SLA policies. This feature enables tenant administrators to apply preconfigured policies to specify the levels of packet loss, jitter, and latency for tenant traffic over the WAN. When a WAN SLA policy is applied to tenant traffic, the Cisco APIC sends the configured policies to a Cisco vSmart Controller. The Cisco vSmart Controller, which is configured in Cisco ACI as an external device manager that provides Cisco IOS XE SD-WAN capabilities, chooses the best possible WAN link that meets the loss, jitter, and latency parameters specified in the SLA policy.

The WAN SLA policies are applied to tenant traffic through contracts.

As an example of where this feature can be useful, consider a deployment in which branches connect to a data center over a WAN via multiple transport technologies, such as MPLS, Internet, and 4G. In such deployments, there can be multiple paths between the branches and data centers. This feature provides optimized path selection in these situations based on application groups and SLA.

- [Guidelines to Integrate with Cisco ACI, on page 162](#)
- [Verify Cisco ACI Registration, on page 162](#)
- [SLA Classes, on page 162](#)
- [Data Prefixes, on page 163](#)
- [VPNs, on page 163](#)
- [Map Data Prefix and VPN to SLA, on page 163](#)
- [Create an App-Route-Policy, on page 164](#)
- [Map ACI Sites, on page 164](#)
- [Unmap ACI Sites, on page 165](#)
- [Delete a Controller, on page 165](#)

Guidelines to Integrate with Cisco ACI

The general steps that you perform in Cisco vManage to configure the integration are:

1. Verify that Cisco ACI has registered the desired controller as a partner with a Cisco vSmart Controller, as described in the procedure, [Verify Cisco ACI Registration](#).
2. Attach devices to the Cisco vSmart Controller, as described in the Map ACI Sites section.

The following guidelines apply when integrating Cisco vManage with Cisco ACI:

- Only new Cisco IOS XE SD-WAN deployments support this integration.
- Make sure that any devices to which the Cisco APIC sends policies do not have any application-aware routing policies configured for them.
- Make sure each device to which the Cisco APIC sends policies has an attached template.
- Before you begin the integration, use the CLI policy builder to create a centralized policy and activate it by using the Cisco vManage policy builder.
- Before you apply WAN SLA policies, establish a connection between the Cisco vSmart Controller and the Cisco APIC. For instructions, see [Cisco ACI and Cisco IOS XE SD-WAN Integration](#).
- Before you attach devices, configure Cisco ACI for this integration.

Verify Cisco ACI Registration

After you configure Cisco ACI for integration with Cisco vManage, perform the following steps in the Cisco vManage to verify that Cisco ACI has registered the desired controller as a Cisco vManage partner:

Step 1 In Cisco vManage, select **Administration** > **Integration Management**.

The Integration Management page displays.

Step 2 On the Integration Management page, verify that ACI Partner Registration appears in the Description for the controller to which the Cisco APIC is to send policies.

SLA Classes

Cisco vManage provides preconfigured SLA classes for use with the ACI integration. These SLA classes are available automatically and cannot be modified or deleted.

To view these SLA classes, follow these steps:

1. In Cisco vManage, select **Configuration** > **Policies**.
2. From the Custom Options drop-down menu, select **Lists**.

3. Select **SLA Class** from the type list on the left.

The following SLA classes are available:

- Business Normal—Designed for normal business operations
- Voice—Designed for voice operations
- Business Critical—Designed for critical business operations that require low packet loss and latency
- Business High—Designed for highly important business operations

Data Prefixes

Cisco ACI creates data prefix lists that are required for integration and updates these lists dynamically as required. You do not need to configure the data prefixes in Cisco vManage.

To view these data prefixes, follow these steps:

1. In Cisco vManage, select **Configuration > Policies**.
2. From the Custom Options drop-down menu, select **Lists**.
3. Select **Data Prefix** from the type list on the left.

Because Cisco ACI provides these data prefixes automatically, the information in this list can vary. To make sure you are viewing current information, refresh the page occasionally.

VPNs

Cisco ACI creates VPNs that are required for integration and sends them to Cisco vManage. These VPNs become available in Cisco vManage automatically. You do not need to configure the VPNs in Cisco vManage.

To view these VPNs, follow these steps:

-
- Step 1** In Cisco vManage, select **Configuration > Policies**.
 - Step 2** From the Custom Options drop-down menu, select **Lists**.
 - Step 3** Select **VPN** from the type list on the left.
-

Map Data Prefix and VPN to SLA

After Cisco ACI establishes a mapping from a data prefix list and a VPN list to an SLA class, Cisco ACI sends the mapping to Cisco vManage. You can view these mappings in Cisco vManage on the page where you configure the app route policy.

Create an App-Route-Policy

After Cisco ACI maps a data prefix and a VPN to an SLA class list, you can create an app-rout-policy to define sequence rules for the Cisco ACI integration.

To create an app-route-policy, follow these steps:

-
- Step 1** In Cisco vManage, select **Configuration > Policies**.
 - Step 2** Click the **More Actions** icon at the right of a row that contains a centralized policy, and then click **Edit**.
 - Step 3** Select the **Traffic Rules** tab.
 - Step 4** Select **Add Policy > Create New**.
 - Step 5** Click **ACI Sequence Rules**.
 - Step 6** From the VPN drop-down, choose a VPN ID. Cisco vManage displays a list of data prefixes and SLA classes that are mapped to this VPN. (These mappings were sent by Cisco ACI.)
 - Step 7** Check the box to the left of the data prefix and SLA class that you want to include with the policy, and then click **Import**.
 - Step 8** Enter a name for the policy in the Name field and a description of the policy in the Description field, and then click **Save Application Aware Routing Policy**. Cisco vManage creates the policy.
 - Step 9** To apply a site list and a VPN list to the policy, select the **Policy Application** tab, then select **Application-Aware Routing**, and then click **New Site Lists and VPN List**.
 - Step 10** Select a site list and a VPN list for the policy.
 - Step 11** Add sequence rules to the policy as needed.
 - Step 12** Click **Save Policy Changes**.
-

Map ACI Sites

Mapping ACI sites designates the controller devices to which the policies from Cisco APIC apply.

Before you begin, review the guidelines in the [Guidelines to Integrate with Cisco ACI](#) section.

To attach devices to a controller, follow these steps:

-
- Step 1** In Cisco vManage, select **Administration > Integration Management**.
The Integration Management page displays.
 - Step 2** Click the **More Actions** icon to the right of the row for the applicable site and select **Attach Devices**.
 - Step 3** In the Available Devices column on the left, select a group and search for one or more devices, select a device from the list, or click **Select All**.
 - Step 4** Click the arrow pointing right to move the device to the Selected Devices column on the right.
To remove devices from the Selected Devices column, in that column select a group and search for one or more devices, select a device from the list, or click **Select All**, and then click the arrow pointing left.

Step 5 Click **Attach**.

Unmap ACI Sites

Unmapping ACI sites stops Cisco APIC policies from being applied to the unmapped devices.

To detach devices from a controller, follow these steps:

Step 1 In Cisco vManage, select **Administration > Integration Management**.

The Integration Management page displays.

Step 2 Click the **More Actions** icon to the right of the row for the applicable site and select **Detach Devices**.

Step 3 In the Available Devices column on the left, select a group and search for one or more devices, select a device from the list, or click **Select All**.

Step 4 Click the arrow pointing right to move the device to the Selected Devices column on the right.

To remove devices from the Selected Devices column, in that column select a group and search for one or more devices, select a device from the list, or click **Select All**, and then click the arrow pointing left.

Step 5 Click **Detach**.

Delete a Controller

If you want to remove a controller as a partner with Cisco ACI, we recommend that you remove its registration by using Cisco ACI instead of deleting it in Cisco vManage. Deleting an ACI partner from Cisco vManage automatically deletes the data prefixes and VPNs that Cisco ACI created for the partner.

Before you begin, remove from policy definitions and data prefix lists and VPN lists that ACI created and make sure that these lists are not referenced from any policy.

Step 1 In Cisco vManage, select **Administration > Integration Management**.

The Integration Management page displays.

Step 2 Detach all devices that are attached to the controller.

For instructions, see the Detach Devices from a Controller section.

Step 3 Click the **More Actions** icon to the right of the row for the applicable site and select **Delete Controller**.




CHAPTER 10

Application-Aware Routing

Application-aware routing tracks network and path characteristics of the data plane tunnels between Cisco IOS XE SD-WAN devices and uses the collected information to compute optimal paths for data traffic. These characteristics include packet loss, latency, and jitter, and the load, cost and bandwidth of a link. The ability to consider factors in path selection other than those used by standard routing protocols—such as route prefixes, metrics, link-state information, and route removal on the Cisco IOS XE SD-WAN device—offers a number of advantages to an enterprise:

- In normal network operation, the path taken by application data traffic through the network can be optimized, by directing it to WAN links that support the required levels of packet loss, latency, and jitter defined in an application's SLA.
- In the face of network brownouts or soft failures, performance degradation can be minimized. The tracking of network and path conditions by application-aware routing in real time can quickly reveal performance issues, and it automatically activates strategies that redirect data traffic to the best available path. As the network recovers from the soft failure conditions, application-aware routing automatically readjusts the data traffic paths.
- Network costs can be reduced because data traffic can be more efficiently load-balanced.
- Application performance can be increased without the need for WAN upgrades.



368471

Each Cisco IOS XE SD-WAN device supports up to eight TLOCs, allowing a single Cisco IOS XE SD-WAN device to connect to up to eight different WAN networks. This capability allows path customization for application traffic that has different needs in terms of packet loss and latency.

- [Components of Application-Aware Routing, on page 168](#)
- [Classification of Tunnels into SLA Classes, on page 169](#)
- [Configure Application-Aware Routing, on page 171](#)
- [Configure Application-Aware Routing Using CLIs, on page 178](#)
- [Structural Components of Policy Configuration for Application-Aware Routing, on page 180](#)
- [Apply Application-Aware Routing Policy, on page 185](#)
- [Configure the Monitoring of Data Plane Tunnel Performance, on page 187](#)
- [Application-Aware Routing Policy Configuration Example, on page 189](#)

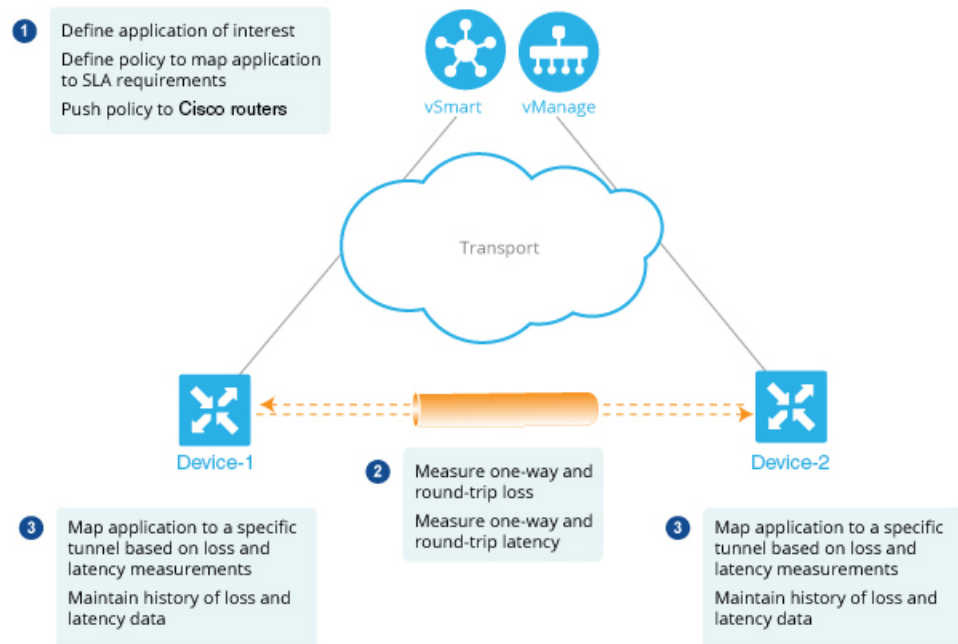
Components of Application-Aware Routing

The Cisco IOS XE SD-WAN Application-Aware Routing solution consists of three elements:

- **Identification**—You define the application of interest, and then you create a centralized data policy that maps the application to specific SLA requirements. You single out data traffic of interest by matching on the Layer 3 and Layer 4 headers in the packets, including source and destination prefixes and ports, protocol, and DSCP field. As with all centralized data policies, you configure them on a Cisco vSmart Controller, which then passes them to the appropriate Cisco IOS XE SD-WAN devices.
- **Monitoring and measuring**—The Cisco IOS XE SD-WAN software uses BFD packets to continuously monitor the data traffic on the data plane tunnels between devices, and periodically measures the performance characteristics of the tunnel. To gauge performance, the Cisco IOS XE SD-WAN device looks for traffic loss on the tunnel, and it measures latency by looking at the one-way and round-trip

times of traffic traveling over the tunnel. These measurements might indicate a suboptimal data traffic conditions.

- **Mapping application traffic to a specific transport tunnel**—The final step is to map an application's data traffic to the data plane tunnel that provides the desired performance for the application. The mapping decision is based on two criteria: the best-path criteria computed from measurements performed on the WAN connections and on the constraints specified in a policy specific to application-aware routing.



To create data policy based on the Layer 7 application itself, use configure deep packet inspection with a centralized data policy. With deep packet inspection, you can direct traffic to a specific tunnel, based on the remote TLOC, the remote TLOC, or both. You cannot direct traffic to tunnels based on SLA classes.

Classification of Tunnels into SLA Classes

The process of classifying tunnels into one or more SLA classes for application-aware routing has three parts:

- Measure loss, latency, and jitter information for the tunnel.
- Calculate the average loss, latency, and jitter for the tunnel.
- Determine the SLA classification of the tunnel.

Measure Loss, Latency, and Jitter

When a data plane tunnel in the overlay network is established, a BFD session automatically starts on the tunnel. In the overlay network, each tunnel is identified with a color that identifies a specific link between a local TLOC and a remote TLOC. The BFD session monitors the liveness of the tunnel by periodically sending Hello packets to detect whether the link is operational. Application-aware routing uses the BFD Hello packets to measure the loss, latency, and jitter on the links.

By default, the BFD Hello packet interval is 1 second. This interval is user-configurable (with the **bfd color interval** command). Note that the BFD Hello packet interval is configurable per tunnel.

Calculate Average Loss, Latency, and Jitter

BFD periodically polls all the tunnels on the Cisco IOS XE SD-WAN devices to collect packet latency, loss, jitter, and other statistics for use by application-aware routing. At each poll interval, application-aware routing calculates the average loss, latency, and jitter for each tunnel, and then calculates or recalculates each tunnel's SLA. Each poll interval is also called a "bucket."

By default, the poll interval is 10 minutes. With the default BFD Hello packet interval at 1 second, this means that information from about 600 BFD Hello packets is used in one poll interval to calculate loss, latency, and jitter for the tunnel. The poll interval is user-configurable (with the **bfd app-route poll-interval** command). Note that the application-aware routing poll interval is configurable per Cisco IOS XE SD-WAN device; that is, it applies to all tunnels originating on a device.

Reducing the poll interval without reducing the BFD Hello packet interval may affect the quality of the loss, latency, and jitter calculation. For example, setting the poll interval to 10 seconds when the BFD Hello packet interval is 1 second means that only 10 Hello packets are used to calculate the loss, latency, and jitter for the tunnel.

The loss, latency, and jitter information from each poll interval is preserved for six poll intervals. At the seventh poll interval, the information from the earliest polling interval is discarded to make way for the latest information. In this way, application-aware routing maintains a sliding window of tunnel loss, latency, and jitter information.

The number of poll intervals (6) is not user-configurable. Each poll interval is identified by an index number (0 through 5) in the output of the **show app-route statistics** command.

Determine the SLA Classification

To determine the SLA classification of a tunnel, application-aware routing uses the loss, latency, and jitter information from the latest poll intervals. The number of poll intervals used is determined by a multiplier. By default, the multiplier is 6, so the information from all the poll intervals (specifically, from the last six poll intervals) is used to determine the classification. For the default poll interval of 10 minutes and the default multiplier of 6, the loss, latency, and jitter information collected over the last hour is considered when classifying the SLA of each tunnel. These default values have to be chosen to provide damping of sorts, as a way to prevent frequent reclassification (flapping) of the tunnel.

The multiplier is user-configurable (with the **bfd app-route multiplier** command). Note that the application-aware routing multiplier is configurable per Cisco IOS XE SD-WAN device; that is, it applies to all tunnels originating on a device.

If there is a need to react quickly to changes in tunnel characteristics, you can reduce the multiplier all the way down to 1. With a multiplier of 1, only the latest poll interval loss and latency values are used to determine whether this tunnel can satisfy one or more SLA criteria.

Based on the measurement and calculation of tunnel loss and latency, each tunnel may satisfy one or more user-configured SLA classes. For example, a tunnel with a mean loss of 0 packets and mean latency of 10 milliseconds would satisfy a class that has been defined with a maximum packet loss of 5 and a minimum latency of 20 milliseconds, and it would also satisfy a class that has been defined with a maximum packet loss of 0 and minimum latency of 15 milliseconds.

Regardless of how quickly a tunnel is reclassified, the loss, latency, and jitter information is measured and calculated continuously. You can configure how quickly application-aware routing reacts to changes by modifying the poll interval and multiplier.

Configure Application-Aware Routing

This topic provides general procedures for configuring application-aware routing. Application-aware routing policy affects only traffic that is flowing from the service side (the local/WAN side) to the tunnel (WAN) side of the Cisco IOS XE SD-WAN device.

An application-aware routing policy matches applications with an SLA, that is, with the data plane tunnel performance characteristics that are necessary to transmit the applications' data traffic. The primary purpose of application-aware routing policy is to optimize the path for data traffic being transmitted by Cisco IOS XE SD-WAN devices.

An application-aware routing policy is a type of centralized data policy: you configure it on the vSmart controller, and the controller automatically pushes it to the affected Cisco IOS XE SD-WAN devices. As with any policy, an application-aware routing 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 data packet matches one of the match conditions, an SLA action is applied to the packet to determine the data plane tunnel to use to transmit the packet. If a packet matches no parameters in any of the policy sequences, and if no default SLA class is configured, the packet is accepted and forwarded with no consideration of SLA. Because application-aware routing policy accepts nonmatching traffic by default, it is considered to be a positive policy. Other types of policies in the Cisco IOS XE SD-WAN software are negative policies, because by default they drop nonmatching traffic.

General Cisco vManage Configuration Procedure

To configure application-aware routing policy, 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 Applications or 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 application-aware routing policy to take effect, you must activate the policy.

Step 1: Start the Policy Configuration Wizard

To start the policy configuration wizard:

1. In Cisco vManage NMS, select the **Configure** > **Policies** screen. When you first open this screen, the Centralized Policy tab is selected by default.
2. Click **Add Policy**.

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

Step 2: Create Applications or Groups of Interest

To create lists of applications or groups to use in centralized policy:

1. Create new lists of groups, as described:
 - Application
 - a. In the left bar, click **Application**.
 - b. Click **New Application List**.
 - c. Enter a name for the list.
 - d. Click either the **Application** or **Application Family** button.
 - e. From the Select drop-down, select the desired applications or application families.
 - f. 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**.
 - 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**.
 - SLA Class
 - a. In the left bar, click **SLA Class**.
 - b. Click **New SLA Class List**.
 - c. Enter a name for the list.
 - d. Define the SLA class parameters:
 1. In the Loss field, enter the maximum packet loss on the connection, a value from 0 through 100 percent.
 2. In the Latency field, enter the maximum packet latency on the connection, a value from 0 through 1,000 milliseconds.
 3. In the Jitter field, enter the maximum jitter on the connection, a value from 1 through 1,000 milliseconds.

- e. Click **Add**.
 - 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 in the wizard. When you first open this screen, the Topology tab is selected by default.

Step 3: Configure the Network Topology

To configure the network topology:

1. In the Topology tab, create a network topology
Hub and Spoke - Policy for a topology with one or more central hub sites and with spokes connected to a hub.
 - a. In the Add Topology drop-down, select **Hub and Spoke**.
 - b. Enter a name for the hub-and-spoke policy.
 - c. Enter a description for the policy.
 - d. In the VPN List field, select the VPN list for the policy.
 - e. 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.
 - f. Double-click the **My Hub-and-Spoke** text string, and enter a name for the policy component.
 - g. In the right pane, add hub sites to the network topology:
 1. Click **Add Hub Sites**.
 2. In the Site List Field, select a site list for the policy component.
 3. Click **Add**.
 4. Repeat Steps 7a, 7b, and 7c to add more hub sites to the policy component.
 - h. In the right pane, add spoke sites to the network topology:
 1. Click **Add Spoke Sites**.
 2. In the Site List Field, select a site list for the policy component.
 3. Click **Add**.
 4. Repeat Steps 8a, 8b, and 8c to add more spoke sites to the policy component.

- i. Repeat Steps 5 through 8 to add more components to the hub-and-spoke policy.
- j. Click **Save Hub and Spoke Policy**.

Mesh - Partial-mesh or full-mesh region

- a. In the Add Topology drop-down, select **Mesh**.
 - b. Enter a name for the mesh region policy component.
 - c. Enter a description for the mesh region policy component.
 - d. In the VPN List field, select the VPN list for the policy.
 - e. Click **New Mesh Region**.
 - f. In the Mesh Region Name field, enter a name for the individual mesh region.
 - g. In the Site List field, select one or more sites to include in the mesh region.
 - h. Repeat Steps 5 through 7 to add more mesh regions to the policy.
 - i. Click **Save Mesh Region**.
2. To use an existing topology:
 - a. In the Add Topology drop-down, click **Import Existing Topology**. The Import Existing Topology popup displays.
 - b. Select the type of topology.
 - c. In the Policy drop-down, select the name of the topology.
 - d. Click **Import**.
 3. Click **Next** to move to Configure Traffic Rules in the wizard. When you first open this screen, the Application-Aware Routing tab is selected by default.

Step 4: Configure Traffic Rules

To configure traffic rules for application-aware routing policy:

1. In the Application-Aware Routing bar, select the **Application-Aware Routing** tab.
2. Click the **Add Policy** drop-down.
3. Select **Create New**, and in the left pane, click **Sequence Type**. A policy sequence containing the text string App Route is added in the left pane.
4. Double-click the App Route text string, and enter a name for the policy sequence. The name you type is displayed both in the Sequence Type list in the left pane and in the right pane.
5. In the right pane, click **Sequence Rule**. The Match/Action box opens, and Match is selected by default. The available policy match conditions are listed below the box.
6. To select one or more Match conditions, click its box and set the values as described in the following table:

Table 27:

Match Condition	Procedure
None (match all packets)	Do not specify any match conditions.
Applications/Application Family List	<ol style="list-style-type: none"> a. In the Match conditions, click Applications/Application Family List. b. In the drop-down, select the application family. c. To create an application list: <ol style="list-style-type: none"> 1. Click New Application List. 2. Enter a name for the list. 3. Click the Application button to create a list of individual applications. Click the Application Family button to create a list of related applications. 4. In the Select Application drop-down, select the desired applications or application families. 5. Click Save.
Destination Data Prefix	<ol style="list-style-type: none"> a. In the Match conditions, click Destination Data Prefix. b. To match a list of destination prefixes, select the list from the drop-down. c. To match an individual destination prefix, type the prefix in the Destination box.
Destination Port	<ol style="list-style-type: none"> a. In the Match conditions, click Destination Port. b. In the Destination field, enter the port number. Specify a single port number, a list of port numbers (with numbers separated by a space), or a range of port numbers (with the two numbers separated with a hyphen [-]).
DNS Application List (to enable split DNS)	<ol style="list-style-type: none"> a. In the Match conditions, click DNS Application List. b. In the drop-down, select the application family.
DNS (to enable split DNS)	<ol style="list-style-type: none"> a. In the Match conditions, click DNS. b. In the drop-down, select Request to process DNS requests for the DNS applications, and select Response to process DNS responses for the applications.
DSCP	<ol style="list-style-type: none"> a. In the Match conditions, click DSCP. b. In the DSCP field, type the DSCP value, a number from 0 through 63.

PLP	<ol style="list-style-type: none"> a. In the Match conditions, click PLP. b. In the PLP drop-down, select Low or High. To set the PLP to high, apply a policer that includes the exceed remark option.
Protocol	<ol style="list-style-type: none"> a. In the Match conditions, click Protocol. b. In the Protocol field, type the Internet Protocol number, a number from 0 through 255.
Source Data Prefix	<ol style="list-style-type: none"> a. In the Match conditions, click Source Data Prefix. b. To match a list of source prefixes, select the list from the drop-down. c. To match an individual source prefix, type the prefix in the Source box.
Source Port	<ol style="list-style-type: none"> a. In the Match conditions, click Source Port. b. In the Source field, enter the port number. Specify a single port number, a list of port numbers (with numbers separated by a space), or a range of port numbers (with the two numbers separated with a hyphen [-]).

7. To select actions to take on matching data traffic, click the Actions box. The available policy actions are listed below the box.
8. Set the policy action for a **Counter** match condition. Count matching data packets.
 - a. In the Action conditions, click **Counter**.
 - b. In the Counter Name field, enter the name of the file in which to store packet counters.
9. Set the policy action for a **Log** match condition. Place a sampled set of packets that match the SLA class rule into system logging (syslog) files. In addition to logging the packet headers, a syslog message is generated the first time a packet header is logged and then every 5 minutes thereafter, as long as the flow is active.
 - a. In the Action conditions, click **Log** to enable logging.
10. Set the policy action for a **SLA Class List** match condition. For the SLA class, all matching data traffic is directed to a tunnel whose performance matches the SLA parameters defined in the class. The software first tries to send the traffic through a tunnel that matches the SLA. If a single tunnel matches the SLA, data traffic is sent through that tunnel. If two or more tunnels match, traffic is distributed among them. If no tunnel matches the SLA, data traffic is sent through one of the available tunnels.
 - a. In the Action conditions, click **SLA Class List**.
 - b. In the SLA Class drop-down, select one or more SLA classes.
 - c. Optionally, in the Preferred Color drop-down, select the color of the data plane tunnel or tunnels to prefer. Traffic is load-balanced across all tunnels. If no tunnels match the SLA, data traffic is sent through any available tunnel. That is, color preference is a loose matching, not a strict matching.
 - d. Click **Strict** to perform strict matching of the SLA class. If no data plane tunnel is available that satisfies the SLA criteria, traffic is dropped.

11. Click **Save Match and Actions**.
12. Create additional sequence rules as desired. Drag and drop to re-arrange them.
13. Create additional sequence types as desired. Drag and drop to re-arrange them.
14. Click **Save Application-Aware Routing Policy**.
15. Click **Next** to move to Apply Policies to Sites and VPNs in the wizard.

Step 5: Apply Policies to Sites and VPNs

In the last screen of the policy configuration wizard, you associate the policy blocks that you created on the previous three screens with VPNs and with sites in the overlay network.

To apply a policy block to sites and VPNs in the overlay network:

1. If you are already in the policy configuration wizard, skip to Step 6. Otherwise, in Cisco vManage NMS, select the **Configure > Policies** screen. When you first open this screen, the Centralized Policy tab is selected by default.
2. Click **Add Policy**. The policy configuration wizard opens, and the Create Applications or Groups of Interest screen is displayed.
3. Click **Next**. The Network Topology screen opens, and in the Topology bar, the Topology tab is selected by default.
4. Click **Next**. The Configure Traffic Rules screen opens, and in the Application-Aware Routing bar, the Application-Aware Routing tab is selected by default.
5. Click **Next**. The Apply Policies to Sites and VPNs screen opens.
6. 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.
7. 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.
8. From the Topology bar, select the type of policy block. The table then lists policies that you have created for that type of policy block.
9. Click **Add New Site List** and **VPN list**. Select one or more site lists, and select one or more VPN lists. Click **Add**.
10. Click **Preview** to view the configured policy. The policy is displayed in CLI format.
11. Click **Save Policy**. The **Configuration > Policies** screen opens, and the policies table includes the newly created policy.

Step 6: Activate an Application-Aware Routing Policy

Activating an application-aware routing policy sends that policy to all connected Cisco vSmart Controllers. To activate a policy:

1. In Cisco vManage NMS, select the **Configure > Policies** screen. When you first open this screen, the Centralized Policy tab is selected by default.

2. Select a policy.
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**.

Configure Application-Aware Routing Using CLIs

Following are the high-level steps for configuring an application-aware routing policy:

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

```
vSmart(config)# policy
vSmart(config-policy)# lists site-list list-name
vSmart(config-site-list)# 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 SLA classes and traffic characteristics to apply to matching application data traffic:

```
vSmart(config)# policy sla-class sla-class-name
vSmart(config-sla-class)# jitter milliseconds
vSmart(config-sla-class)# latency milliseconds
vSmart(config-sla-class)# loss percentage
```

3. Create lists of applications, IP prefixes, and VPNs to use in identifying application traffic of interest (in the **match** section of the policy definition):

```
vSmart(config)# policy lists
vSmart(config-lists)# app-list list-name
vSmart(config-app-list)# (app application-name | app-family family-name)

vSmart(config-lists)# prefix-list list-name
vSmart(config-prefix-list)# ip-prefix prefix/length

vSmart(config-lists)# vpn-list list-name
vSmart(config-vpn-list)# vpn vpn-id
```

4. Create an application-aware routing policy instance and associate it with a list of VPNs:

```
vSmart(config)# policy app-route-policy policy-name
vSmart(config-app-route-policy)# vpn-list list-name
```

5. Within the policy, create one or more numbered sequence of match–action pairs, where the match parameters define the data traffic and applications of interest and the action parameters specify the SLA class to apply if a match occurs.

- a. Create a sequence:

```
vSmart(config-app-route-policy)# sequence number
```

- b. Define match parameters for data packets:

```
vSmart(config-sequence)# match parameters
```

- c. Define the action to take if a match occurs:

```
vSmart(config-sequence) # action sla-class sla-class-name [strict]
vSmart(config-sequence) # action sla-class sla-class-name [strict] preferred-color
colors
```

The first two **action** options direct matching data traffic to a tunnel interface that meets the SLA characteristics in the specified SLA class:

- **sla-class** *sla-class-name*—When you specify an SLA class with no additional parameters, data traffic that matches the SLA is forwarded as long as one tunnel interface is available. The software first tries to send the traffic through a tunnel that matches the SLA. If a single tunnel matches the SLA, data traffic is sent through that tunnel. If two or more tunnels match, traffic is distributed among them. If no tunnel matches the SLA, data traffic is sent through one of the available tunnels.
- **sla-class** *sla-class-name* **preferred-color** *color*—To set a specific tunnel to use when data traffic matches an SLA class, include the **preferred-color** option, specifying the color of the preferred tunnel. If more than one tunnel matches the SLA, traffic is sent to the preferred tunnel. If a tunnel of the preferred color is not available, traffic is sent through any tunnel that matches the SLA class. If no tunnel matches the SLA, data traffic is sent through any available tunnel. In this sense, color preference is considered to be a loose matching, not a strict matching, because data traffic is always forwarded, whether a tunnel of the preferred color is available or not.
- **sla-class** *sla-class-name* **preferred-color** *colors*—To set multiple tunnels to use when data traffic matches an SLA class, include the **preferred-color** option, specifying two or more tunnel colors. Traffic is load-balanced across all tunnels.

If no tunnel matches the SLA, data traffic is sent through any available tunnel. In this sense, color preference is considered to be a loose matching, not a strict matching, because data traffic is always forwarded, whether a tunnel of the preferred color is available or not. When no tunnel matches the SLA, you can choose how to handle the data traffic:

- **strict**—Drop the data traffic.

d. Count the packets or bytes that match the policy:

```
vSmart(config-sequence) # action count counter-name
```

e. Place a sampled set of packets that match the SLA class rull into syslog files:

```
vSmart(config-sequence) # action log
```

f. The match–action pairs within a policy are evaluated in numerical order, based on the sequence number, starting with the lowest number. If a match occurs, the corresponding action is taken and policy evaluation stops.

6. If a packet does not match any of the conditions in one of the sequences, a default action is taken. For application-aware routing policy, the default action is to accept nonmatching traffic and forward it with no consideration of SLA. You can configure the default action so that SLA parameters are applied to nonmatching packets:

```
vSmart(config-policy-name) # default-action sla-class sla-class-name
```

7. Apply the policy to a site list:

```
vSmart(config) # apply-policy site-list list-name app-route-policy policy-name
```

Structural Components of Policy Configuration for Application-Aware Routing

Here are the structural components required to configure application-aware routing policy. Each one is explained in more detail in the sections below.

```

policy
  lists
    app-list list-name
      (app application-name | app-family application-family)
    prefix-list list-name
      ip-prefix prefix
    site-list list-name
      site-id site-id
    vpn-list list-name
      vpn-id vpn-id
  log-frequency number
  sla-class sla-class-name
    jitter milliseconds
    latency milliseconds
    loss percentage
  app-route-policy policy-name
    vpn-list list-name
      sequence number
      match
        match-parameters
      action
        count counter-name
        log
        sla-class sla-class-name [strict] [preferred-color colors]
      default-action
        sla-class sla-class-name
  apply-policy site-list list-name
    app-route-policy policy-name

```

Lists

Application-aware routing policy uses the following types of lists to group related items. You configure these lists under the **policy lists** command hierarchy on Cisco vSmart Controllers.

Table 28:

List Type	Description	Command
Applications and application families	<p>List of one or more applications or application families running on the subnets connected to the Cisco IOS XE SD-WAN device. Each app-list can contain either applications or application families, but you cannot mix the two. To configure multiple applications or application families in a single list, include multiple app or app-family options, specifying one application or application family in each app or app-family option.</p> <ul style="list-style-type: none"> • <i>application-name</i> is the name of an application. The Cisco IOS XE SD-WAN device supports about 2300 different applications. • <i>application-family</i> is the name of an application family. It can be one of the following: antivirus, application-service, audio_video, authentication, behavioral, compression, database, encrypted, erp, file-server, file-transfer, forum, game, instant-messaging, mail, microsoft-office, middleware, network-management, network-service, peer-to-peer, printer, routing, security-service, standard, telephony, terminal, thin-client, tunneling, wap, web, and webmail. 	<pre>app-list list-name (app application-name app-family application-family)</pre>
Data prefixes	<p>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.</p>	<pre>data-prefix-list list-name ip-prefix prefix/length</pre>

List Type	Description	Command
Sites	List of one or more site identifiers in the overlay network. To configure multiple sites in a single list, include multiple site-id options, specifying one site number in each option. You can specify a single site identifier (such as site-id 1) or a range of site identifiers (such as site-id 1-10).	site-list <i>list-name</i> site-id <i>site-id</i>
VPNs	List of one or more VPNs in the overlay network. 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-id 1) or a range of VPN identifiers (such as vpn-id 1-10).	vpn-list <i>list-name</i> vpn <i>vpn-id</i>

In the Cisco vSmart Controller configuration, you can create multiple iterations of each type of list. For example, it is common to create multiple site lists and multiple VPN lists so that you can apply data policy to different sites and different customer VPNs across the network.

When you create multiple iterations of a type of list (for example, when you create multiple VPN lists), you can include the same values or overlapping values in more than one of these list. You can do this either on purpose, to meet the design needs of your network, or you can do this accidentally, which might occur when you use ranges to specify values. (You can use ranges to specify data prefixes, site identifiers, and VPNs.) Here are two examples of lists that are configured with ranges and that contain overlapping values:

- **vpn-list list-1 vpn 1-10**
vpn-list list-2 vpn 6-8
- **site-list list-1 site 1-10**
site-list list-2 site 5-15

When you configure data policies that contain lists with overlapping values, or when you apply data policies, you must ensure that the lists included in the policies, or included when applying the policies, do not contain overlapping values. To do this, you must manually audit your configurations. The Cisco IOS XE SD-WAN configuration software performs no validation on the contents of lists, on the data policies themselves, or on how the policies are applied to ensure that there are no overlapping values.

If you configure or apply data policies that contain lists with overlapping values to the same site, one policy is applied and the others are ignored. Which policy is applied is a function of the internal behavior of Cisco IOS XE SD-WAN device when it processes the configuration. This decision is not under user control, so the outcome is not predictable.

VPN Lists

Each application-aware policy instance is associated with a VPN list. You configure VPN lists with the **policy app-route-policy vpn-list** command. The VPN list you specify must be one that you created with a **policy lists vpn-list** command.

Sequences

Within each VPN list, an application-aware policy contains sequences of match–action pairs. The sequences are numbered to set the order in which data traffic is analyzed by the match–action pairs in the policy. You configure sequences with the **policy app-aware-policy vpn-list sequence** command.

Each sequence in an application-aware policy can contain one **match** command and one **action** command.

Match Parameters

Application-aware routing policy can match IP prefixes and fields in the IP headers. You configure the match parameters with the **match** command under the **policy app-route-policy vpn-list sequence** command hierarchy on Cisco vSmart Controllers.

You can match these parameters:

Table 29:

Description	Command	Value or Range
Match all packets	Omit match command	—
Applications or application families	app-list <i>list-name</i>	Name of an app-list list
Group of destination prefixes	destination-data-prefix-list <i>list-name</i>	Name of a data-prefix-list list
Individual destination prefix	destination-ip <i>prefix/length</i>	IP prefix and prefix length
Destination port number	destination-port <i>number</i>	0 through 65535. Specify a single port number, a list of port numbers (with numbers separated by a space), or a range of port numbers (with the two numbers separated with a hyphen [-]).
DSCP value	dscp <i>number</i>	0 through 63
Internet Protocol number	protocol <i>number</i>	0 through 255
Packet loss priority (PLP)	plp	(high low) By default, packets have a PLP value of low . To set the PLP value to high , apply a policer that includes the exceed remark option.
Group of source prefixes	source-data-prefix-list <i>list-name</i>	Name of a data-prefix-list list
Individual source prefix	source-ip <i>prefix/length</i>	IP prefix and prefix length
Source port number	source-port <i>number</i>	0 through 65535; enter a single port number, a list of port numbers (with numbers separated by a space), or a range of port numbers (with the two numbers separated with a hyphen [-])

Description	Command	Value or Range
Split DNS, to resolve and process DNS requests on an application-by-application basis	dns-app-list <i>list-name</i> dns (request response)	Name of an app-list list. This list specifies the applications whose DNS requests are processed. To process DNS requests sent by the applications (for outbound DNS queries), specify dns request . To process DNS responses returned from DNS servers to the applications, specify dns response .

Action Parameters

When data traffic matches the match parameters, the specified action is applied to it. For application-aware routing policy, the action is to apply an SLA class. The SLA class defines the maximum packet latency or maximum packet loss, or both, that the application allows on the data plane tunnel used to transmit its data. The Cisco SD-WAN software examines the recently measured performance characteristics of the data plane tunnels and directs the data traffic to the WAN connection that meets the specified SLA.

The following actions can be configured:

Table 30:

Description	Command	Value or Range
Count matching data packets.	action count <i>counter-name</i>	Name of a counter.
SLA class to match. All matching data traffic is directed to a tunnel whose performance matches the SLA parameters defined in the class. The software first tries to send the traffic through a tunnel that matches the SLA. If a single tunnel matches the SLA, data traffic is sent through that tunnel. If two or more tunnels match, traffic is distributed among them. If no tunnel matches the SLA, data traffic is sent through one of the available tunnels.	action sla-class <i>sla-class-name</i>	SLA class name defined in policy sla-class command
Group of data plane tunnel colors to prefer when an SLA class match occurs. Traffic is load-balanced across all tunnels. If no tunnels match the SLA, data traffic is sent through any available tunnel. That is, color preference is a loose matching, not a strict matching.	action sla-class <i>sla-class-name</i> preferred-color <i>colors</i>	SLA class name defined in policy sla-class command and one of the supported tunnel colors.

Description	Command	Value or Range
Strict matching of the SLA class. If no data plane tunnel is available that satisfies the SLA criteria, traffic is dropped. Note that for policy configured with this option, data traffic that matches the match conditions is dropped until the application-aware routing path is established.	action sla-class <i>sla-class-name</i> strict action sla-class <i>sla-class-name</i> preferred-color <i>color</i> strict action sla-class <i>sla-class-name</i> preferred-color <i>colors</i> strict	SLA class name defined in policy sla-class command

If more than one data plane tunnel satisfies an SLA class criteria, the Cisco IOS XE SD-WAN device selects one of them by performing load-balancing across the equal paths.

Default Action

A policy's default action defines how to handle packets that match none of the match conditions. For application-aware routing policy, if you do not configure a default action, all data packets are accepted and transmitted based on normal routing decisions, with no consideration of SLA.

To modify this behavior, include the **default-action sla-class** *sla-class-name* command in the policy, specifying the name of an SLA class you defined in the **policy sla-class** command.

When you apply an SLA class in a policy's default action, you cannot specify the **strict** option.

If no data plane tunnel satisfies the SLA class in the default action, the Cisco IOS XE SD-WAN device selects one of the available tunnels by performing load-balancing across equal paths.

Apply Application-Aware Routing Policy

For an application-aware route policy to take effect, you apply it to a list of sites in the overlay network:

```
vSmart(config)# apply-policy site-list list-name app-route-policy policy-name
```

When you apply the policy, you do not specify a direction (either inbound or outbound). Application-aware routing policy affects only the outbound traffic on the Cisco IOS XE SD-WAN devices.

For all **app-route-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 **app-route-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:

- Centralized control policy (**control-policy**)
- Centralized data policy (**data-policy**)
- Centralized data policy used for cflowd flow monitoring (**data-policy** that 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 **app-route-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.

As soon as you successfully activate the configuration on the Cisco vSmart Controller by issuing a **commit** command, the controller pushes the application-aware routing policy to the Cisco IOS XE SD-WAN devices at the specified sites.

To view the policy configured on the Cisco vSmart Controller, use the **show running-config** command on the controller.

To view the policy that the Cisco vSmart Controller has pushed to the device, issue the **show policy from-vsmart** command on the router.

To display flow information for the application-aware applications running on the device, issue the **show app dpi flows** command on the router.

How Application-Aware Routing Policy Is Applied in Combination with Other Data Policies

If you configure a Cisco IOS XE SD-WAN device with application-aware routing policy and with other policies, the policies are applied to data traffic sequentially.

On a Cisco IOS XE SD-WAN device, you can configure the following types of data policy:

- Centralized data policy. You configure this policy on the Cisco vSmart Controller, and the policy is passed to the device. You define the configuration with the **policy data-policy configuration** command, and you apply it with the **apply-policy site-list data-policy**, or **apply-policy site-list vpn-membership** command.
- Localized data policy, which is commonly called access lists. You configure access lists on the device with the **policy access-list** configuration command. You apply them, within a VPN, to an incoming interface with the **vpn interface access-list in** configuration command or to an outgoing interface with the **vpn interface access-list out** command.
- Application-aware routing policy. Application-aware routing policy affects only traffic that is flowing from the service side (the local/LAN side) to the tunnel (WAN) side of the Cisco IOS XE SD-WAN device. You configure application-aware routing policy on the Cisco vSmart Controller with the **policy app-route-policy** configuration command, and you apply it with the **apply-policy site-list app-route-policy** command. When you commit the configuration, the policy is passed to the appropriate devices. Then, matching data traffic on the device is processed in accordance with the configured SLA conditions. Any data traffic that is not dropped as a result of this policy is passed to the data policy for evaluation. If the data traffic does not match and if no default action is configured, transmit it without SLA consideration.

You can apply only one data policy and one application-aware routing policy to a single site in the overlay network. When you define and apply multiple site lists in a configuration, you must ensure that a single data policy or a single application-aware routing policy is not applied to more than one site. The CLI does not check for this circumstance, and the **validate** configuration command does not detect whether multiple policies of the same type are applied to a single site.

For data traffic flowing from the service side of the router to the WAN side of the router, policy evaluation of the traffic evaluation occurs in the following order:

1. Apply the input access list on the LAN interface. Any data traffic that is not dropped as a result of this access list is passed to the application-aware routing policy for evaluation.

2. Apply the application-aware routing policy. Any data traffic that is not dropped as a result of this policy is passed to the data policy for evaluation. If the data traffic does not match and if no default action is configured, transmit it without SLA consideration.
3. Apply the centralized data policy. Any data traffic that is not dropped as a result of the input access list is passed to the output access list for evaluation.
4. Apply the output access list on the WAN interface. Any data traffic that is not dropped as a result of the output access list is transmitted out the WAN interface.

For data traffic coming from the WAN through the router and into the service-side LAN, the policy evaluation of the traffic occurs in the following order:

1. Apply the input access list on the WAN interface. Any data traffic that is not dropped as a result of the input access list is passed to the data policy for evaluation.
2. Apply the data policy. Any data traffic that is not dropped as a result of the input access list is passed to the output access list for evaluation.
3. Apply the output access list on the LAN interface. Any data traffic that is not dropped as a result of the output access list is transmitted out the LAN interface, towards its destination at the local site.

As mentioned above, application-aware routing policy affects only traffic that is flowing from the service side (the local/LAN side) to the tunnel (WAN) side of the Cisco IOS XE SD-WAN device, so data traffic inbound from the WAN is processed only by access lists and data policy.

Configure the Monitoring of Data Plane Tunnel Performance

The Bidirectional Forwarding Detection (BFD) protocol runs over all data plane tunnels between Cisco IOS XE SD-WAN devices, monitoring the liveness, and network and path characteristics of the tunnels. Application-aware routing uses the information gathered by BFD to determine the transmission performance of the tunnels. Performance is reported in terms of packet latency and packet loss on the tunnel.

BFD sends Hello packets periodically to test the liveness of a data plane tunnel and to check for faults on the tunnel. These Hello packets provide a measurement of packet loss and packet latency on the tunnel. The Cisco IOS XE SD-WAN device records the packet loss and latency statistics over a sliding window of time. BFD keeps track of the six most recent sliding windows of statistics, placing each set of statistics in a separate bucket. If you configure an application-aware routing policy for the device, it is these statistics that the router uses to determine whether a data plane tunnel's performance matches the requirements of the policy's SLA.

The following parameters determine the size of the sliding window:

Parameters	Default	Configuration Command	Range
BFD Hello packet interval	1 second	bfd color <i>color</i> hello-interval <i>seconds</i>	1 through 65535 seconds
Polling interval for application-aware routing	10 minutes (600,000 milliseconds)	bfd app-route poll-interval <i>milliseconds</i>	1 through 4,294,967 ($2^{32} - 1$) milliseconds
Multiplier for application-aware routing	6	bfd app-route multiplier <i>number</i>	1 through 6

Let us use the default values for these parameters to explain how application-aware routing works:

- For each sliding window time period, application-aware routing sees 600 BFD Hello packets (BFD Hello interval x polling interval: 1 second x 600 seconds = 600 Hello packets). These packets provide measurements of packet loss and latency on the data plane tunnels.
- Application-aware routing retains the statistics for 1 hour (polling interval x multiplier: 10 minutes x 6 = 60 minutes).
- The statistics are placed in six separate buckets, indexed with the numbers 0 through 5. Bucket 0 contains the latest statistics, and bucket 5 the oldest. Every 10 minutes, the newest statistics are placed in bucket 0, the statistics in bucket 5 are discarded, and the remaining statistics move into the next bucket.
- Every 60 minutes (every hour), application-aware routing acts on the loss and latency statistics. It calculates the mean of the loss and latency of all the buckets in all the sliding windows and compares this value to the specified SLAs for the tunnel. If the calculated value satisfies the SLA, application-aware routing does nothing. If the value does not satisfy the SLA, application-aware routing calculates a new tunnel.
- Application-aware routing uses the values in all six buckets to calculate the mean loss and latency for a data tunnel. This is because the multiplier is 6. While application-aware always retains six buckets of data, the multiplier determines how many it actually uses to calculate the loss and latency. For example, if the multiplier is 3, buckets 0, 1, and 2 are used.

Because these default values take action only every hour, they work well for a stable network. To capture network failures more quickly so that application-aware routing can calculate new tunnels more often, adjust the values of these three parameters. For example, if you change just the polling interval to 1 minute (60,000 milliseconds), application-aware routing reviews the tunnel performance characteristics every minute, but it performs its loss and latency calculations based on only 60 Hello packets. It may take more than 1 minute for application-aware routing to reset the tunnel if it calculates that a new tunnel is needed.

To display statistics for each data plane tunnel, use the **show sdwan app-route stats** command:

```
Device# show sdwan app-route stats
```

SRC IP	DST IP	PROTO	SRC PORT	DST PORT	MEAN LOSS	MEAN LATENCY	INDEX	TOTAL PACKETS	LOSS	AVERAGE LATENCY	AVERAGE JITTER	TX DATA PKTS	RX DATA PKTS	
193.0.2.1	194.0.2.1	ipsec	12346	12346	0	22	0	596	0	2	0	0		
								1	596	0	21	2	0	0
								2	596	0	21	2	0	0
								3	597	1	21	2	0	0
								4	596	0	21	2	0	0
193.0.2.1	194.0.2.1	ipsec	12346	12346	0	24	0	596	0	24	3	0		
								1	596	0	25	3	0	0
								2	596	0	25	3	0	0
								3	596	0	24	3	0	0
								4	596	0	24	3	0	0
193.0.2.1	194.0.2.1	ipsec	12346	34083	0	21	0	596	0	21	3	0		
								1	596	0	22	3	0	0
								2	596	0	22	3	0	0
								3	596	0	21	3	0	0
								4	596	0	21	3	0	0
193.0.2.1	194.0.2.1	ipsec	12346	36464	0	23	0	596	0	23	3	0		
								1	596	0	23	3	0	0
								2	596	0	24	3	0	0
								3	596	0	23	4	0	0
								4	596	0	23	4	0	0

```
5      596      0      23      4      0      0
```

To display the next-hop information for an IP packet that a device sends out a service side interface, use the **show policy service-path** command. To view the similar information for packets that the router sends out a WAN transport tunnel interface, use the **show policy tunnel-path** command.

Enable Application Visibility on Cisco IOS XE SD-WAN Devices

You can enable application visibility directly on Cisco IOS XE SD-WAN devices, without configuring application-aware routing policy so that you can monitor all the applications running in all VPNs in the LAN. To do this, configure application visibility on the router:

```
vEdge(config)# policy app-visibility
```

To monitor the applications, use the **show app dpi applications** and **show app dpi supported-applications** commands on the device.

Application-Aware Routing Policy Configuration Example

This topic shows a straightforward example of configuring application-aware routing policy. This example defines a policy that applies to ICMP traffic, directing it to links with latency of 50 milliseconds or less when such links are available.

Configuration Steps

You configure application-aware routing policy on a Cisco vSmart Controller. The configuration consists of the following high-level components:

- Definition of the application (or applications)
- Definition of SLA parameters
- Definition of sites, prefixes, and VPNs
- Application-aware routing policy itself
- Specification of overlay network sites to which the policy is applied

The order in which you configure these components is immaterial from the point of view of the CLI. However, from an architectural design point of view, a logical order is to first define all the parameters that are invoked in the application-aware routing policy itself or that are used to apply the policy to various sites in the overlay network. Then, you specify the application-aware routing policy itself and the network sites to which you want to apply the policy.

Here is the procedure for configuring this application-aware routing policy on a Cisco vSmart Controller:

1. Define the SLA parameters to apply to matching ICMP traffic. In our example, we want to direct ICMP traffic to links that have a latency of 50 milliseconds or less:

```
vSmart# config
vSmart(config)# policy sla-class test_sla_class latency 50
vSmart(config-sla-class-test_sla_class)#
```

2. Define the site and VPN lists to which we want to apply the application-aware routing policy:

```
vSmart(config-sla-class-test_sla_class)# exit
vSmart(config-sla-class-test_sla_class)# lists vpn-list vpn_1_list vpn 1
```

```
vSmart(config-vpn-list-vpn_1_list)# exit
vSmart(config-lists)# site-list site_500 site-id 500
vSmart(config-site-list-site_500)#
```

3. Configure the application-aware routing policy. Note that in this example, we apply the policy to the application in two different ways: In sequences 1, 2, and 3, we specify the protocol number (protocol 1 is ICMP, protocol 6 is TCP, and protocol 17 is UDP).

```
vSmart(config-site-list-site_500)# exit
vSmart(config-lists)# exit
vSmart(config-policy)# app-route-policy test_app_route_policy
vSmart(config-app-route-policy-test_app_route_policy)# vpn-list vpn_1_list
vSmart(config-vpn-list-vpn_1_list)# sequence 1 match protocol 6
vSmart(config-match)# exit
vSmart(config-sequence-1)# action sla-class test_sla_class strict
vSmart(config-sequence-1)# exit
vSmart(config-vpn-list-vpn_1_list)# sequence 2 match protocol 17
vSmart(config-match)# exit
vSmart(config-sequence-2)# action sla-class test_sla_class
vSmart(config-sequence-2)# exit
vSmart(config-vpn-list-vpn_1_list)# sequence 3 match protocol 1
vSmart(config-match)# exit
vSmart(config-sequence-3)# action sla-class test_sla_class strict
vSmart(config-sequence-3)# exit
vSmart(config-sequence-4)#
```

4. Apply the policy to the desired sites in the Cisco IOS XE SD-WAN overlay network:

```
vSmart(config-sequence-4)# top
vSmart(config)# apply-policy site-list site_500 app-route-policy test_app_route_policy
```

5. Display the configuration changes:

```
vSmart(config-site-list-site_500)# top
vSmart(config)# show config
```

6. Validate that the configuration contains no errors:

```
vSmart(config)# validate
Validation complete
```

7. Activate the configuration:

```
vSmart(config)# commit
Commit complete.
```

8. Exit from configuration mode:

```
vSmart(config)# exit
vSmart#
```

Full Example Configuration

Putting all the pieces of the configuration together gives this configuration:

```
vSmart# show running-config policy
policy
sla-class test_sla_class
latency 50
!
app-route-policy test_app_route_policy
vpn-list vpn_1_list
sequence 1
match
```

```
        protocol 6
        !
        action sla-class test_sla_class strict
        !
sequence 2
  match
    protocol 17
    !
    action sla-class test_sla_class
    !
sequence 3
  match
    protocol 1
    !
    action sla-class test_sla_class strict
    !
  !
!
lists
  vpn-list vpn_1_list
  vpn 1
  !
  site-list site_500
  site-id 500
  !
  site-list site_600
  site-id 600
  !
!
!
apply-policy
  site-list site_500
  app-route-policy test_app_route_policy
  !
!
```




CHAPTER 11

Traffic Flow Monitoring with Cflowd


Cflowd monitors traffic flowing through Cisco IOS XE SD-WAN devices in the overlay network and exports flow information to a collector, where it can be processed by an IPFIX analyzer. For a traffic flow, cflowd periodically sends template reports to flow collector. These reports contain information about the flows and the data is extracted from the payload of these reports.

You can create a cflowd-template that defines the location of cflowd collectors, how often sets of sampled flows are sent to the collectors, and how often the template is sent to the collectors (on Cisco vSmart Controllers only). You can configure a maximum of four cflowd collectors per Cisco IOS XE SD-WAN device. To have a cflowd-template take effect, apply it with the appropriate data policy.

You must configure at least one cflowd-template, but it need not contain any parameters. With no parameters, the data flow cache on the nodes is managed using default settings, and no flow export occurs.

Cflowd traffic flow monitoring is equivalent to Flexible Netflow (FNF).

The cflowd software implements cflowd version 10, as specified in *RFC 7011* and *RFC 7012*. Cflowd version 10 is also called the IP Flow Information Export (IPFIX) protocol.



Cflowd performs 1:1 sampling. Information about all flows is aggregated in the cflowd records; flows are not sampled. Cisco IOS XE SD-WAN devices do not cache any of the records that are exported to a collector.

Cisco IOS XE SD-WAN device IPFIX Information Elements Exported to the Collector

The Cisco IOS XE SD-WAN device cflowd software exports the following IPFIX information elements to the cflowd collector. Fields vary depending on the release that you are on. Common fields are exported to Cisco vManage and external exporters. Feature fields are exported only to Cisco vManage.

Before Cisco XE SD-WAN Release 17.2, Flexible NetFlow (FNF) exports all fields to external collectors and vManage. Starting from Cisco XE SD-WAN Release 17.2, FNF export the elements (that are marked

yes) in the following table to both external collectors and vManage. Other fields like “drop cause id” are for specific features and these fields are exported only to vManage, but not to external collector.

Information Element	Element ID	Exported to External Collector	Description	Data Type	Data Type Semantics	Units or Range
sourceIPv4Address	8	Yes	IPv4 source address in the IP packet header.	ipv4Address (4 bytes)	default	—
destinationIPv4Address	12	Yes	IPv4 destination address in the IP packet header.	IPv4Address (4 bytes)	default	—
ingressInterface	10	Yes	Index of the IP interface where packets of this flow are being received.	unsigned32 (4 bytes)	identifier	—
ipDiffServCodePoint	195	Yes	Value of a Differentiated Services Code Point (DSCP) encoded in the Differentiated Services field. This field spans the most significant 6 bits of the IPv4 TOS field.	unsigned8 (1 byte)	identifier	0 through 63
protocolIdentifier	4	Yes	Value of the protocol number in the Protocol field of the IP packet header. The protocol number identifies the IP packet payload type. Protocol numbers are defined in the IANA Protocol Numbers registry.	unsigned8 (1 byte)	identifier	—
sourceTransportPort	7	Yes	Source port identifier in the transport header. For the transport protocols UDP, TCP, and SCTP, this is the destination port number given in the respective header. For GRE and IPsec flows, the value of this field is 0.	unsigned16 (2 bytes)	identifier	—
destinationTransportPort	11	Yes	Destination port identifier in the transport header. For the transport protocols UDP, TCP, and SCTP, this is the destination port number given in the respective header.	unsigned16 (2 bytes)	identifier	—
tcpControlBits	6	Yes	TCP control bits observed for the packets of this flow. This information is encoded as a bit field; each TCP control bit has a bit in this set. The bit is set to 1 if any observed packet of this flow has the corresponding TCP control bit set to 1. Otherwise, the bit is set to 0. For values of this field, see the <i>IANA IPFIX web page</i> .	unsigned8 (1 byte)	identifier	—

Information Element	Element ID	Exported to External Collector	Description	Data Type	Data Type Semantics	Units or Range
flowEndReason	136	Yes	Reason for the flow termination. For values of this field, see the <i>IANA IPFIX web page</i> .	unsigned8 (1 byte)	identifier	—
ingressoverlaysessionid	12432	Yes	A 32-bit identifier for input overlay session id.	unsigned32 (4 bytes)	identifier	—
VPN Identifier	Enterprise specific	Yes	Cisco IOS XE SD-WAN device VPN identifier. The device uses the enterprise ID for VIP_IANA_ENUM or 41916, and the VPN element ID is 4321.	unsigned32 (4 bytes)	identifier	0 through 65535
connection id long	12441	Yes	A 64-bit identifier for a connection between client and server.	Unsigned64 (8 bytes)	identifier	—
application id	95	Yes	A 32 bit identifier for an application name	unsigned32 (4 bytes)	identifier	—
egressInterface	14	Yes	Index of the IP interface where packets of this flow are being sent.	unsigned32 (4 bytes)	default	—
egressoverlaysessionid	12433	Yes	A 32-bit identifier for output overlay session id.	unsigned32 (4 bytes)	identifier	—
sdwan qos-queue-id	12446	No	Queue index for QoS.	unsigned8 (1 byte)	identifier	—
drop cause id	12442	No	A 16-bit identifier for a drop cause name.	unsigned16 (2 bytes)	identifier	—
counter bytes sdwan dropped long	12443	No	Total number of dropped octets in incoming packets for this flow at the observation point since initialization or re-initialization of the metering process for the observation point. The count includes the IP heads and the IP payload.	unsigned64 (8 bytes)	totalCounter	Octets
sdwan sla-not-met	12444	No	A Boolean to indicate if required SLA is met or not.	unsigned8 (1 byte)	identifier	—
sdwan preferred-color-not-met	12445	No	A Boolean to indicate if preferred color is met or not.	unsigned8 (1 byte)	identifier	—
counter packets sdwan dropped long	42329	No	Total number of dropped packets in incoming packets for this flow at the observation point since initialization or re-initialization of the metering process for the observation point.	unsigned64 (8 bytes)	totalCounter	Packets

Information Element	Element ID	Exported to External Collector	Description	Data Type	Data Type Semantics	Units or Range
octetDeltaCount	1	Yes	Number of octets since the previous report in incoming packets for this flow at the observation point. This number includes IP headers and IP payload.	unsigned64 (8 bytes)	deltaCounter	Octets
packetDeltaCount	2	Yes	Number of incoming packets since the previous report for this flow at this observation point.	unsigned64 (8 bytes)	deltaCounter	Packets
flowStartMilliseconds	152	Yes	Absolute timestamp of the first packet of this flow.	dateTime-MilliSeconds (8 bytes)	—	—
flowEndMilliseconds	153	Yes	Absolute timestamp of the last packet of this flow.	dateTime-MilliSeconds (8 bytes)	—	—

- [Configure Traffic Flow Monitoring on Cisco XE SD-WAN Devices, on page 196](#)
- [Configure Cflowd Traffic Flow Monitoring Using CLI, on page 199](#)
- [Structural Components of Policy Configuration for Cflowd, on page 200](#)
- [Apply and Enable Cflowd Policy, on page 203](#)
- [Cflowd Traffic Flow Monitoring Configuration Example, on page 204](#)

Configure Traffic Flow Monitoring on Cisco XE SD-WAN Devices

This topic provides the procedure for configuring cflowd traffic flow monitoring on Cisco IOS XE SD-WAN devices. Cflowd traffic flow monitoring uses Flexible Netflow (FNF) to export traffic data. To configure cflowd monitoring, follow these steps:

1. Configure global flow visibility.
2. Configure cflowd monitoring policy.

Configure Global Flow Visibility

To enable cflowd visibility globally on all Cisco IOS XE SD-WAN devices so that you can perform traffic flowing monitoring on traffic coming to the router from all VPNs in the LAN.

In Cisco vManage NMS

1. Select the **Configuration > Policies** screen.
2. Select the **Localized Policy** tab.
3. Click **Add Policy**.
4. Click **Next** to display the Configure Policy Setting screen.

5. Click **Netflow**.

From the CLI

```
Device# config-transaction
Device(config)# policy flow-visibility
Device(config-policy)# commit
Commit complete.
Device(config-policy)# end
Device#
```



Note The **policy app-visibility** command also enables global flow visibility by enabling **nbar** to get the application name.

Configure Global Application Visibility

To enable cflowd visibility globally on all Cisco IOS XE SD-WAN devices so that you can perform traffic flowing monitoring on traffic coming to the router from all VPNs in the LAN.

The difference between **flow-visibility** and **app-visibility** is that **app-visibility** enables **nbar** to see each application of the flows coming to the router from all VPNs in the LAN.

In Cisco vManage NMS

1. Select the **Configuration > Policies** screen.
2. Select the **Localized Policy** tab.
3. Click **Add Policy**.
4. Click **Next** to display the Configure Policy Setting screen.
5. Click **Application**.

From the CLI

```
Device# config-transaction
Device(config)# policy app-visibility
Device(config-policy)# commit
Commit complete.
Device(config-policy)# end
Device#
```

Configure Cflowd Monitoring Policy

To configure policy for cflowd traffic flow monitoring, 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:

1. Create Applications or Groups of Interest—Create lists that group together related items and that you call in the match or action components of a policy.
2. Configure Topology—Create the network structure to which the policy applies.

3. Configure Traffic Rules—Create the match and action conditions of a policy.
4. 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 the cflowd policy to take effect, you must activate the policy.

For details of the Cisco vManage configuration procedure, see *Configuring Cflowd Traffic Flow Monitoring*.

From the CLI on the Cisco vSmart Controller that is controlling the Cisco IOS XE SD-WAN device:

1. Configure a cflowd template to specify flow visibility and flow sampling parameters:

```
vSmart(config)# policy cflowd-template template-name
vSmart(config-cflowd-template)# flow-active-timeout seconds
vSmart(config-cflowd-template)# flow-inactive-timeout seconds
vSmart(config-cflowd-template)# flow-sampling-interval number
vSmart(config-cflowd-template)# template-refresh seconds
```

2. Configure a flow collector:

```
vSmart(config-cflowd-template)# collector vpn vpn-id address
ip-address port port-number transport transport-type
source-interface interface-name
```



Note Cisco IOS XE SD-WAN devices only support UDP collector. Irrespective of which transport protocol is configured, the collector functionality on Cisco IOS XE SD-WAN device is always UDP.

3. Configure a data policy that defines traffic match parameters and that includes the action **cflowd**:

```
vSmart(config)# policy data-policy policy-name
vSmart(config-data-policy)# sequence number
vSmart(config-sequence)# match match-parameters
vSmart(config-sequence)# action cflowd
vSmart(config-data-policy)# default-action accept
```

4. Create lists of sites in the overlay network that contain the Cisco IOS XE SD-WAN devices to which you want to apply the traffic flow monitoring policy. To include multiple site in the list, configure multiple **vpn vpn-id** commands.

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

5. Apply the data policy to the sites in the overlay network that contain the Cisco IOS XE SD-WAN devices:

```
vSmart(config)# apply-policy site-list list-name
vSmart(config-site-list)# data-policy policy-name
vSmart(config-site-list)# cflowd-template template-name
```

Display Cflowd Information

To display cflowd information, use the following commands on the Cisco IOS XE SD-WAN device.

- show sdwan app-fwd cflowd collector

- show sdwan app-fwd cflowd flow-count
- show sdwan app-fwd cflowd flows [vpn *vpn-id*] format table
- show sdwan app-fwd cflowd statistics
- show sdwan app-fwd cflowd template [name *template-name*]
- show sdwan app-fwd cflowd flows format table

Configure Cflowd Traffic Flow Monitoring Using CLI

Following are the high-level steps for configuring a cflowd centralized data policy to perform traffic monitoring and to export traffic flows to a collector:

1. Create a list of overlay network sites to which the cflowd centralized data 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 a list of VPN for which the cflowd centralized data policy is to be configured (in the **policy data-policy** command):

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

3. Create lists of IP prefixes, as needed:

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

4. Configure a cflowd template, and optionally, configure template parameters, including the location of the cflowd collector, the flow export timers, and the flow sampling interval:

```
vSmart(config)# policy cflowd-template template-name
vSmart(config-cflowd-template-template-name)# collector vpn vpn-id address ip-address
port port-number transport-type (transport_tcp | transport_udp) source-interface
interface-name
vSmart(config-cflowd-template-template-name)# flow-active-timeout seconds
vSmart(config-cflowd-template-template-name)# flow-inactive-timeout seconds
vSmart(config-cflowd-template-template-name)# template-refresh seconds

vSmart(config)# policy cflowd-template cflowd_server
    flow-active-timeout 60
    flow-inactive-timeout 30
    template-refresh 80
```

You must configure a cflowd template, but it need not contain any parameters. With no parameters, the data flow cache on router is managed using default settings, and no flow export occurs. You can configure one cflowd template per router, and it can export to a maximum of four collectors.

By default, an actively flowing data set is exported to the collector every 600 seconds (10 minutes), a data set for a flow on which no traffic is flowing is sent every 60 seconds (1 minute), and the cflowd template record fields (the three timer values) are sent to the collector every 90 seconds.

Also by default, a new flow is created immediately after an existing flow has ended. If you modify the configuration of the template record fields, the changes take effect only on flows that are created after the configuration change has been propagated to the router. Because an existing flow continues indefinitely, to have configuration changes take effect, clear the flow with the **clear app cflowd flows** command.



Note On Cisco IOS XE SD-WAN devices, flow-active-timeout is fixed as 60 seconds. If a flow-inactive-timeout is fixed as 10 seconds, it cannot be configured.

5. Create a data policy instance and associate it with a list of VPNs:

```
vSmart(config)# policy data-policy policy-name
vSmart(config-data-policy-policy-name)# vpn-list list-name
```

6. Create a sequence to contain a single match–action pair:

```
vSmart(config-vpn-list-list-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. If no match occurs, the default action is taken.

7. Define match parameters for the data packets:

```
vSmart(config-sequence-number)# match parameters
```

8. In the action, enable cflowd:

```
vSmart(config-sequence-number)# action cflowd
```

9. In the action, count or log data packets:

```
vSmart(config-sequence-number)# action count counter-name
vSmart(config-sequence-number)# action log
```

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

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

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

12. Apply the policy and the cflowd template to one or more sites in the overlay network:

```
vSmart(config)# apply-policy site-list list-name data-policy policy-name
vSmart(config)# apply-policy site-list list-name cflowd-template template-name
```

Structural Components of Policy Configuration for Cflowd

Here are the structural components required to configure cflowd on a Cisco vSmart Controller. Each component is explained in more detail in the sections below.


```

policy
  lists
    prefix-list list-name
      ip-prefix prefix
    site-list list-name
      site-id site-id
    vpn-list list-name
      vpn vpn-id
  log-frequency number
  cflowd-template template-name
    collector vpn vpn-id address ip-address port port-number transport transport-type
  source-interface interface-name
  flow-active-timeout seconds
  flow-inactive-timeout seconds
  flow-sampling-interval number
  template-refresh seconds
  data-policy policy-name
  vpn-list list-name
    sequence number
    match
      match-parameters
    action
      cflowd
      count counter-name
      drop
      log
    default-action
      (accept | drop)
  apply-policy site-list list-name
    data-policy policy-name
    cflowd-template template-name

```

Lists

Centralized data policy uses the following types of lists to group related items. You configure lists under the **policy lists** command hierarchy on Cisco vSmart Controllers.

Table 31:

List Type	Description	Command
Data 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.	data-prefix-list list-name ip-prefix prefix/length
Sites	List of one or more site identifiers in the overlay network. To configure multiple sites in a single list, include multiple site-id options, specifying one site number in each option. You can specify a single site identifier (such as site-id 1) or a range of site identifiers (such as site-id 1-10).	site-list list-name site-id site-id
VPNs	List of one or more VPNs in the overlay network. 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).	vpn-list list-name vpn vpn-id

Cflowd Templates

For each cflowd data policy, you must create a template that defines the location of the flow collector:

```
vSmart(config)# policy cflowd-template template-name
```

The template can specify cflowd parameters or it can be empty. With no parameters, the data flow cache on vEdge nodes is managed using default settings, and no flow export occurs.

In the cflowd template, you can define the location of the flow collection:

```
vSmart# (config-cflowd-template-template-name)
vSamrt# collector vpn vpn-id address
ip-address port port-number
transport transport-type source-interface
interface-name
```

You can configure one cflowd template per Cisco vEdge device, and it can export to a maximum of four collectors.

You can configure flow export timers:

```
vSmart(config)# policy cflowd-template template-name
vSmart(config-cflowd-template-template-name)# flow-active-timeout seconds
vSmart(config-cflowd-template-template-name)# flow-inactive-timeout seconds
vSmart(config-cflowd-template-template-name)# template-refresh seconds
```

By default, an actively flowing data set is exported to the collector every 600 seconds (10 minutes), a data set for a flow on which no traffic is flowing is sent every 60 seconds (1 minute), and the cflowd template record fields are sent to the collector every 90 seconds. For flow sampling, by default, a new flow is started immediately after an existing flow ends.

For a single Cisco IOS XE SD-WAN device, you can configure a maximum of four collectors.

Data Policy Instance

For each centralized data policy, you create a named container for that policy with a **policy data-policy** *policy-name* command. For a single Cisco IOS XE SD-WAN device, you can configure a maximum of four cflowd policies.

VPN Lists

Each centralized data policy instance applies to the VPNs contained in a VPN list. Within the policy, you specify the VPN list with the **policy data-policy vpn-list** *list-name* command. The list name must be one that you created with a **policy lists vpn-list** *list-name* command.

Sequences

Within each VPN list, a centralized data policy contains sequences of match–action pairs. The sequences are numbered to set the order in which data traffic is analyzed by the match–action pairs in the policy. You configure sequences with the **policy data-policy vpn-list sequence** command.

Each sequence in a centralized data policy can contain one **match** command and one **action** command.

Match Parameters

Centralized data policy can match IP prefixes and fields in the IP headers. You configure the match parameters under the **policy data-policy vpn-list sequence match** command.

For data policy, you can match these parameters:

Table 32:

Description	Command	Value or Range
Group of destination prefixes	destination-data-prefix-list <i>list-name</i>	Name of a data-prefix-list list.
Individual destination prefix	destination-ip <i>prefix/length</i>	IP prefix and prefix length
Destination port number	destination-port <i>number</i>	0 through 65535
DSCP value	dscp <i>number</i>	0 through 63
Internet Protocol number	protocol <i>number</i>	0 through 255
Group of source prefixes	source-data-prefix-list <i>list-name</i>	Name of a data-prefix-list list
Individual source prefix	source-ip <i>prefix/length</i>	IP prefix and prefix length
Source port number	source-port <i>address</i>	0 through 255

Action Parameters

When data traffic matches the conditions in the match portion of a centralized data policy, the packet can be accepted or rejected, and you can configure a counter for the accepted or rejected packets. You configure the action parameters under the **policy data-policy vpn-list sequence action** command.

Table 33:

Description	Command	Value or Range
Count the accepted or dropped packets.	count <i>counter-name</i>	Name of a counter. To display counter information, use the show policy access-lists counters command on the Cisco vEdge device.
Enable cflowd.	cflowd	—

For a packet that is accepted, configure the parameter **cflowd** to enable packet collection.

Default Action

If a data packet being evaluated does not match any of the match conditions in a control policy, a default action is applied to this route. By default, the route is rejected. To modify this behavior, include the **policy data-policy vpn-list default-action accept** command.

Apply and Enable Cflowd Policy

For a centralized data policy to take effect, you must apply it to a list of sites in the overlay network:

```
vSmart(config)# apply-policy site-list list-name data-policy policy-name
```

To activate the cflowd template, associate it with the data policy:

```
vSmart(config)# apply-policy cflowd-template template-name
```

For all **data-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 **data-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 control policy (**control-policy**)
- Centralized data policy (**data-policy**)

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.

As soon as you successfully activate the configuration by issuing a **commit** command, the Cisco vSmart Controller pushes the data policy to the Cisco IOS XE SD-WAN devices located in the specified sites. To view the policy as configured on the Cisco vSmart Controller, use the **show running-config** command on the Cisco vSmart Controller. To view the policy that has been pushed to the device, use the **show policy from-vsmart** command on the device.

To display the centralized data policy as configured on the Cisco vSmart Controller, use the **show running-config** command:

```
vSmart# show running-config policy
vSmart# show running-config apply-policy
```

To display the centralized data policy that has been pushed to the Cisco IOS XE SD-WAN device, issue the **show omp data-policy** command on the device:

```
Device# show sdwan policy from-vsmart
```

Enable Cflowd Visibility on Cisco IOS XE SD-WAN device Devices

You can enable cflowd visibility directly on Cisco IOS XE SD-WAN devices, without configuring a data policy, so that you can perform traffic flow monitoring on traffic coming to the router from all VPNs in the LAN. To do this, configure cflowd visibility on the device:

```
Device(config)# policy flow-visibility
```

To monitor the applications, use the **show app cflowd flows** and **show app cflowd statistics** commands on the device.

Cflowd Traffic Flow Monitoring Configuration Example

This topic shows a straightforward example of configuring traffic flow monitoring.

Configuration Steps

You enable cflowd traffic monitoring with a centralized data policy, so all configuration is done on a Cisco vSmart Controller. The following example procedure monitors all TCP traffic, sending it to a single collector:

1. Create a cflowd template to define the location of the collector and to modify cflowd timers:

```
vsmart(config)# policy cflowd-template test-cflowd-template
vsmart(config-cflowd-template-test-cflowd-template)# collector vpn 1 address 172.16.155.15
port 13322 transport transport_udp
vsmart(config-cflowd-template-test-cflowd-template)# flow-inactive-timeout 60
vsmart(config-cflowd-template-test-cflowd-template)# template-refresh 90
```

2. Create a list of VPNs whose traffic you want to monitor:

```
vsmart(config)# policy lists vpn-list vpn_1 vpn 1
```

3. Create a list of sites to apply the data policy to:

```
vsmart(config)# policy lists site-list cflowd-sites site-id 400,500,600
```

4. Configure the data policy itself:

```
vsmart(config)# policy data-policy test-cflowd-policy
vsmart(config-data-policy-test-cflowd-policy)# vpn-list vpn_1
vsmart(config-vpn-list-vpn_1)# sequence 1
vsmart(config-sequence-1)# match protocol 6
vsmart(config-match)# exit
vsmart(config-sequence-1)# action accept cflowd
vsmart(config-action)# exit
vsmart(config-sequence-1)# exit
vsmart(config-vpn-list-vpn_1)# default-action accept
```

5. Apply the policy and the cflowd template to sites in the overlay network:

```
vsmart(config)# apply-policy site-list cflowd-sites data-policy test-cflowd-policy
Device(config-site-list-cflowd-sites)# cflowd-template test-cflowd-template
```

6. Activate the data policy:

```
vsmart(config-site-list-cflowd-sites)# validate
Validation complete
vsmart(config-site-list-cflowd-sites)# commit
Commit complete.
vsmart(config-site-list-cflowd-sites)# exit configuration-mode
```

Full Example Configuration

Here is what the full example cflowd configuration looks like:

```
vsmart(config)# show configuration
apply-policy
site-list cflowd-sites
data-policy test-cflowd-policy
cflowd-template test-cflowd-template
!
!
policy
data-policy test-cflowd-policy
vpn-list vpn_1
sequence 1
match
protocol 6
!
action accept
cflowd
!
!
default-action accept
!
```

```

!
cflowd-template test-cflowd-template
  flow-inactive-timeout 60
  template-refresh      90
  collector vpn 1 address 192.0.2.1 port 13322 transport transport_udp
!
lists
  vpn-list vpn_1
    vpn 1
  !
  site-list cflowd-sites
    site-id 400,500,600
  !
!
!
!

```

Check the Cflowd Configuration

After you activate the cflowd configuration on the Cisco vSmart Controller, you can check it with the **show running-config policy** and **show running-config apply-policy** commands on the Cisco vSmart Controller. In addition, the configuration is immediately pushed down to the Cisco IOS XE SD-WAN devices at the affected sites.

You can view the pushed cflowd template with the **show sdwan policy from-vsmart cflowd** command. Here is the output from a device at site 500:

```

Device# show sdwan policy from-vsmart cflowd-template
from-vsmart cflowd-template test-cflowd-template
  flow-active-timeout 30
  flow-inactive-timeout 60
  template-refresh 90
  collector vpn 1 address 192.0.2.1 port 13322 transport transport_udp

```

You can view all the pushed policy components with the **show sdwan policy from-vsmart** command:

```

Device# show sdwan policy from-vsmart
from-vsmart data-policy test-cflowd-policy
  vpn-list vpn_1
    sequence 1
      match
        protocol 6
      action accept
        cflowd
      default-action accept
from-vsmart cflowd-template test-cflowd-template
  flow-active-timeout 30
  flow-inactive-timeout 60
  template-refresh 90
  collector vpn 1 address 192.0.2.1 port 13322 transport transport_udp
from-vsmart lists vpn-list vpn_1
  vpn 1

```

Check the Flows

On the Cisco IOS XE SD-WAN devices affected by the cflowd data policy, various commands let you check the status of the cflowd flows.

```

Device# show sdwan app-fwd cflowd statistics

  data_packets           :      0
  template_packets      :      0
  total-packets         :      0

```

```
flow-refresh           :      123
flow-ageout            :      117
flow-end-detected      :          0
flow-end-forced        :          0
```




CHAPTER 12

Lawful Intercept

The Lawful Intercept (LI) feature supports service providers in meeting the requirements of law enforcement agencies (LEA) to provide electronic surveillance as authorized by a judicial or administrative order. The surveillance is performed using wiretaps to intercept Voice-over-Internet protocol (VoIP) or data traffic going through the edge routers. The LEA delivers a request for a wiretap to the target's service provider, who is responsible for intercepting data communication to and from the individual using IP sessions. A user session is tapped using either the Source and Destination IP addresses, or VRF name, which is translated to a vrf-tableid value within the router.

Table 34: Feature History

Feature Name	Release Information	Description
Encryption of Lawful Intercept Messages	Cisco IOS XE SD-WAN Release 16.12.1b	This feature encrypts lawful intercept messages between a Cisco IOS XE SD-WAN device and a media device using static tunnel information.

- [Information About Lawful Intercept, on page 209](#)
- [Prerequisites for Lawful Intercept, on page 212](#)
- [Install Lawful Intercept using vManage, on page 213](#)
- [Lawful Intercept MIBs, on page 213](#)
- [Restrict Access to Trusted Hosts \(Without Encryption\), on page 214](#)
- [Restrict Trusted Mediation Device, on page 214](#)
- [Configure Lawful Intercept, on page 215](#)
- [Configure Lawful Intercept Using CLI, on page 215](#)
- [Encrypt Lawful Intercept Traffic , on page 216](#)
- [Verify Static Tunnel with Media Device Gateway, on page 218](#)

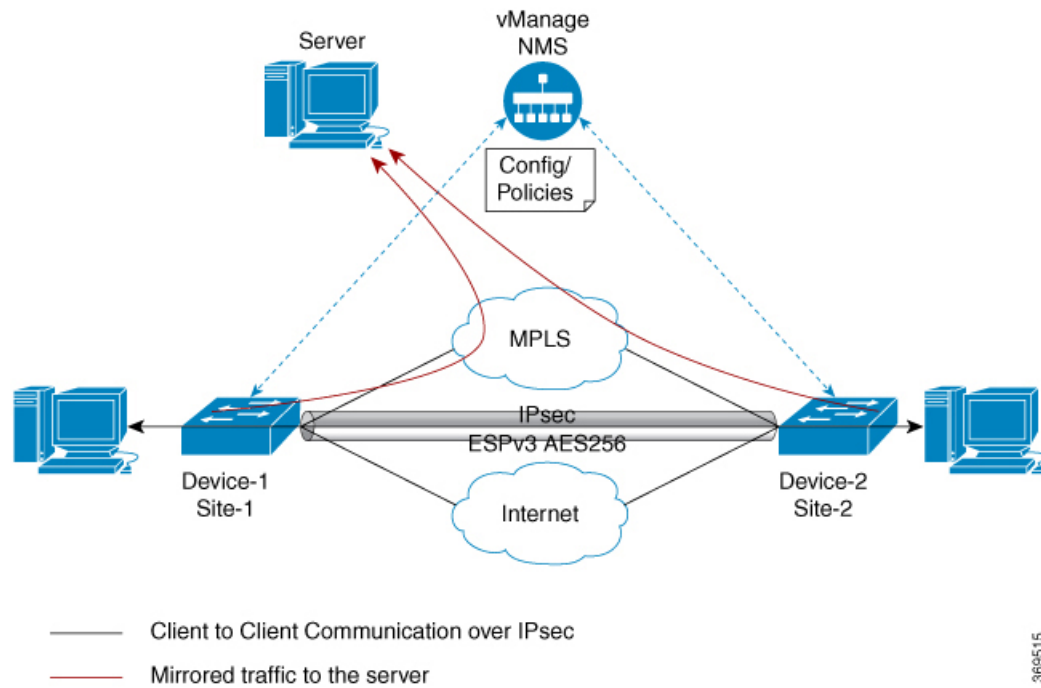
Information About Lawful Intercept

Lawful intercept is a process that enables a Law Enforcement Agency (LEA) to perform electronic surveillance on an individual (a target) as authorized by a judicial or administrative order. To facilitate the lawful intercept process, certain legislation and regulations require service providers (SPs) and Internet service providers (ISPs) to implement their networks to explicitly support authorized electronic surveillance.

Lawful Intercept Process

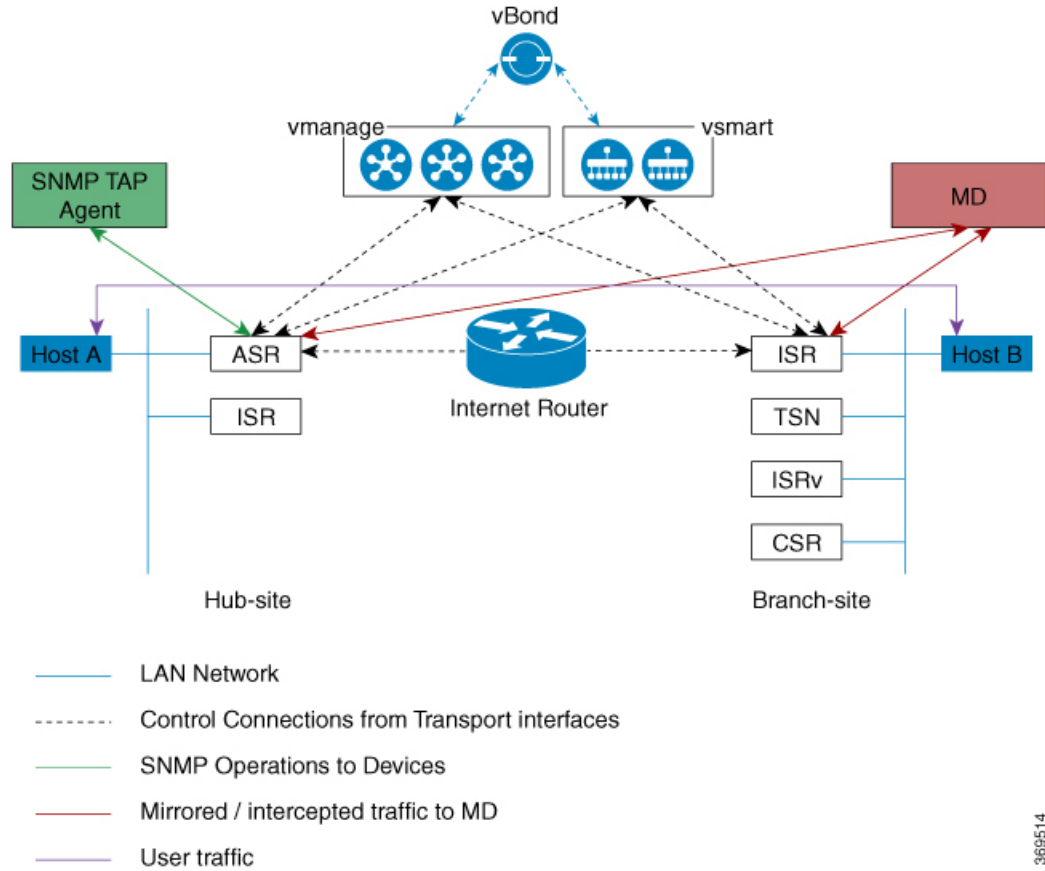
When triggering a lawful intercept for communications from Site A to Site B, the edge platform duplicates the traffic and sends an unencrypted copy of the traffic to a target server, which is hosted in the customer network designed for Lawful Intercept. The Cisco vManage NMS ensures that Cisco vManage users (non-Lawful Intercept users) who have access to Site A and Site B for any information, are unaware of the duplicated flow of information.

Figure 14: Cisco-SD-WAN Lawful Intercept Workflow



369515

Figure 15: Cisco SD-WAN Lawful Intercept Process



369514

Licence-based Lawful Intercept

Cisco SD-WAN solution is a term-based licensed feature. This feature license enables the Cisco vManage component of the Cisco SD-WAN solution and allows the customer to access the Lawful Intercept function. Once the Lawful Intercept license is enabled on the solution, Cisco vManage provides a new privilege in the Manage Users menu of the Cisco vManage UI. By default, this privilege is available to all admin users. In addition, administrators can assign the Lawful Intercept privilege to any other user.

Any user with Lawful Intercept privilege would be able to enable Lawful Intercept function on an edge device in the WAN network. All changes made by any user with Lawful Intercept function would be audit logged and changes will be recorded just like any other change made by any user in the system.

After acquiring a court order or warrant to perform surveillance, any user with Lawful Intercept privilege will be able to make Lawful Intercept related changes on sites with a warrant.

1. Install license for Lawful Intercept on Cisco vManage.
2. Create an lawful intercept admin (liadmin) user on Cisco vManage. The **liadmin** user must be associated with the user group, Basic.
3. Login to Cisco vManage as **liadmin** user and configure Lawful Intercept specific templates.
4. Cisco vManage automatically pushes templates to all Cisco IOS XE SD-WAN devices with Lawful Intercept compatible images.

5. Configuration is pushed to device from Cisco vManage using the following:
 - a. SNMP TAP MIB configuration
 - b. SNMP Access list (li-acl keyword)
 - c. MD List
6. SNMP SET is sent to device to achieve the following goals:
 - a. To setup and activate MD entry on Cisco IOS XE SD-WAN devices.
 - b. To setup and activate stream to be intercepted.
 - c. To activate or deactivate intercept
7. Mediation Device receives the intercepted or mirrored traffic.

VRF-Aware Lawful Intercept

VRF Aware Lawful Intercept is the ability to provision a Lawful Intercept wiretap on IPv4 data in a particular VPN. This feature allows a LEA to lawfully intercept targeted data within that VPN. Only IPv4 data within that VPN is subject to the VRF-based Lawful Intercept tap.

To provision a VPN-based IPv4 tap, the LI administrative function (running on the mediation device) uses the CISCO-IP-TAP-MIB to identify the name of the VRF table that the targeted VPN uses. The VRF name is used to select the VPN interfaces on which to enable LI in order to execute the tap. The device determines which traffic to intercept and which mediation device to send the intercepted packets based on the VRF name (along with the source and destination address, source and destination port, and protocol).

Prerequisites for Lawful Intercept

Access to the Cisco Lawful Intercept MIB view should be restricted to the mediation device and to system administrators who need to be aware of lawful intercepts on the router. To access the MIB, users must have level-15 access rights on the router.

For the router to communicate with the mediation device to execute a lawful intercept, the following configuration requirements must be met:

- The domain name for both the router and the mediation device must be registered in the Domain Name System (DNS). In DNS, the router IP address is typically the address of the FastEthernet0/0/0 interface on the router.
- The mediation device must have an access function (AF) and an access function provisioning interface (AFPI).
- You must add the mediation device to the Simple Network Management Protocol (SNMP) user group that has access to the CISCO-TAP2-MIB view. Specify the username of the mediation device as the user to add to the group.
 - When you add the mediation device as a CISCO-TAP2-MIB user, you can include the mediation device's authorization password if you want. The password must be at least eight characters in length.

- You must configure SNMP service in vManage using the VPN Interface Ethernet screen of Feature Template. See VPN Interface Ethernet section in Templates topic.

Install Lawful Intercept using vManage



Note The following process must be repeated for every Cisco vManage node.

1. Connect to a Cisco vManage device as administrator
2. Request tools license

```
vm12# tools license request
Your org-name is: XYZ Inc
Your license-request challenge is:
Uwk3u4Vwkl8n632fKDIpKDEFkzfeJlhFQPOHobpvewmed0U83LQDgajO7GnmCIgA
```

3. Contact Cisco Support to generate the license using the output of Step 2.
4. Run the install file command and reboot:

```
vm12# tools license install file license.lic
License installed. Please reboot to activate.
vm12# reboot
Are you sure you want to reboot? [yes,no] yes
```

```
Broadcast message from root@vm12 (somewhere) (Tue Jan 22 17:07:47 2019):
Tue Jan 22 17:07:47 UTC 2019: The system is going down for reboot NOW!
Connection to 10.0.1.32 closed.
tester@vip-vmanage-dev-109:~$
```

5. Verify if the Lawful Intercept license is installed successfully, using the following command:

```
vm12# show system status
LI License Enabled True
```

6. Create lawful intercept admin user using Cisco vManage.
7. Login to Cisco vManage using the lawful intercept admin credentials.



Note Use the **tools license remove-all** command to remove all licenses after reboot. You will not be able to re-install the previous license.

Lawful Intercept MIBs

Due to its sensitive nature, the Cisco Lawful Intercept MIBs are only available in software images that support the Lawful Intercept feature.

These MIBs are not accessible through the Network Management Software MIBs [Support page](#).

Restricting Access to the Lawful Intercept MIBs

Only the mediation device and users who need to know about lawful intercepts must be allowed to access the Lawful Intercept MIBs. To restrict access to these MIBs, you must perform the following actions:

1. Create a view that includes the Cisco Lawful Intercept MIBs.
2. Create an SNMP user group that has read-and-write access to the view. Only users assigned to this user group can access information in the MIBs.
3. Add users to the Cisco LI user groups to define who can access the MIBs and any information related to lawful intercepts. Be sure to add the mediation device as a user in this group; otherwise, the router cannot perform lawful intercepts.



Note Detail MD5 authentication key generation algorithm is defined at <https://tools.ietf.org/html/rfc3414#appendix-A.2.1>

Restrict Access to Trusted Hosts (Without Encryption)

SNMPv3 provides support for both security models and security levels. A security model is an authentication strategy that is set up for a user and the group in which the user resides. A security level is the permitted level of security within a security model. A combination of a security model and a security level will determine the security mechanism employed when handling an SNMP packet.

Additionally, the SNMP support for the Named Access Lists feature adds support for standard named access control lists (ACLs) to several SNMP commands.

To configure a new SNMP group or a table that maps SNMP users to SNMP views, use the **snmp-server** command in global configuration mode.

In the following example, the access list named 99 allows SNMP traffic only from 10.1.1.1 to access Cisco IOS XE SD-WAN devices. This access list is then applied to the SNMP user, testuser.

```
access-list 99 permit ip host 10.1.1.1
snmp-server user testuser INTERCEPT_GROUP v3 encrypted auth sha
testPassword1 priv aes testPassword2 access 99
```

SNMP traffic is only allowed from WAN interface (gigabitEthernet 1).

```
control-plane host
management-interface gigabitEthernet 1 allow snmp
```

Restrict Trusted Mediation Device

In the following example, the **md-list** command allows an SNMP request **config MD** in the subnet 10.3.3.0/24.

When a Cisco IOS XE SD-WAN device receives an SNMP request to create a mediation device, it first checks the Mediation Device List configuration information.

If the IP address of the mediation device is not in the configured Mediation Device List, the Mediation Device entry is not active.

```
md-list 10.3.3.0 255.255.255.0
```



Note You can configure up to a maximum of eight Mediation Device List subnets.

Configure Lawful Intercept

The following are the two components for Lawful Intercept vManage configuration:

- Lawful Intercept SNMP template – This template provisions the configuration for the following:
 - SNMPv3 group for lawful intercept – The group name is INTERCEPT_GROUP by default.
 - SNMPv3 users for lawful intercept – All users are restricted by an access list by default.
 - SNMPv3 view is configured by default. The view included Cisco TAP MIBs.
 - The following TAP MIBs are configured:
 - ciscoIpTapMIB
 - ciscoTap2MIB
 - ifIndex
 - ifDescr
- Lawful intercept access list template – The access list template provides configuration for the following:
 - Mediation Device-List configuration – Provides option to configure up to 8 subnets.
 - SNMP access-list – provides option to configure up to 8 subnets or host addresses, and a wildcard mask.

Configure Lawful Intercept Using CLI

```
control-plane host
management-interface GigabitEthernet0/0/0 allow ftp ssh snmp
management-interface GigabitEthernet0/0/1 allow ftp ssh snmp
!
!
md-list 10.101.0.0 255.255.255.0
md-list 10.102.0.10 255.255.255.255
md-list 10.103.0.0 255.255.255.0
md-list 10.104.0.4 255.255.255.255
md-list 10.105.0.0 255.255.255.0
md-list 10.106.0.0 255.255.255.0
md-list 10.107.0.7 255.255.255.255
md-list 10.108.0.0 255.255.0.0
!
ip access-list standard li-acl
permit 174.16.50.254
```

Example: Enabling Mediation Device Access Lawful Intercept MIBs

The following example shows how to enable the mediation device to access the lawful intercept MIBs. It creates an SNMP view (tapV) that includes four LI MIBs (CISCO-TAP2-MIB, CISCO-IP-TAP-MIB, CISCO-802-TAP-MIB, and CISCO-USER-CONNECTION-TAP-MIB). It also creates a user group that has read, write, and notify access to MIBs in the tapV view.

```
snmp-server enable trap
snmp-server engineID local 766D616E6167652Dac10ff31
snmp-server group INTERCEPT_GROUP v3 noauth read INTERCEPT_VIEW write INTERCEPT_VIEW notify
SNG_VIEW
snmp-server user UItestuser1 INTERCEPT_GROUP v3 encrypted auth md5
DA:B2:36:03:6A:5C:D0:6D:F6:D8:9C:5E:56:77:AD:43 priv aes 128
DA:B2:36:03:6A:5C:D0:6D:F6:D8:9C:5E:56:77:AD:43 access li-acl
snmp-server user UItestuser2 INTERCEPT_GROUP v3 encrypted auth md5
D2:01:1E:47:D8:9E:3E:B5:58:CD:90:0F:49:FC:36:56 priv aes 128
CF:32:C4:3E:34:27:3F:4A:D8:18:A7:19:E5:04:A7:DF access li-acl
!
snmp-server engineID local 766D616E6167652DAC10FF31
snmp-server group INTERCEPT_GROUP v3 noauth read INTERCEPT_VIEW write INTERCEPT_VIEW notify
SNG_VIEW
snmp-server view INTERCEPT_VIEW ciscoIpTapMIB included
snmp-server view INTERCEPT_VIEW ciscoTap2MIB included
snmp-server view INTERCEPT_VIEW ifIndex included
snmp-server view INTERCEPT_VIEW ifDescr included
```

Encrypt Lawful Intercept Traffic

Encryption of intercepted traffic between the router (the content Intercept Access Point (IAP)) and the Mediation Device (MD) is recommended.

The following is the required configuration:

- Configuring encryption in the router, and either an encryption client in the MD or a router associated with the MD to decrypt the traffic.
- Restricting access to trusted hosts.
- Configuring the VPN client.

Configure Encryption in the Device

To configure encryption, configure Authentication, Authorization, and Accounting (AAA) parameters. The following example shows how to configure the parameters:

```
aaa authentication login userauthen local
username <username> password 0 <password>
```

In CISCO-TAP2-MIB, the source interface must be the tunnel interface of the Cisco IOS XE SD-WAN devices and the destination address must be IP address of the mediation device.

Configure Lawful Intercept Encryption using CLI

In the following example, an IPSec tunnel is configured between Cisco IOS XE SD-WAN device and Media Device Gateway. Media Device Gateway terminates IPSec tunnel and adds a route to Media Device list through the IPSec Tunnel.

In CISCO-TAP2-MIB, source interface is the tunnel interface of the Cisco IOS XE SD-WAN devices; destination address is the IP address of the media device.

```
crypto ikev2 diagnose error 1000
crypto ikev2 keyring ikev2_keyring
peer mypeer
address 0.0.0.0 0.0.0.0
pre-shared-key cisco123
devic gateway
!
crypto ikev2 profile ikev2_profile
authentication local pre-share
authentication remote pre-share
dpd 10 3 on-demand
lifetime 14400
keyring local ikev2_keyring
match identity remote address 0.0.0.0 0.0.0.0
!
crypto ikev2 proposal default
encryption aes-cbc-256
group 14 16 19 2 20 21
integrity sha256 sha384 sha512
!
crypto ipsec profile ipsec_profile
set ikev2-profile ikev2_profile
set pfs group16
set transform-set tfs
set security-association lifetime seconds 7200
set security-association replay window-size 256
!
crypto ipsec transform-set tfs esp-gcm 256
mode tunnel
!
interface Tunnel100
no shutdown
ip address 10.2.2.1 255.255.255.0
tunnel source GigabitEthernet1
tunnel destination 10.124.19.57
tunnel mode ipsec ipv4
tunnel protection ipsec profile ipsec_profile

ip route 10.3.3.0 255.255.255.0 Tunnel100
```

□ pre-shared key should be same on media

□ tunnel address

□ Cisco XE SD-WAN WAN interface

□ Media Device Gateway address

□ route MD list traffic through IPSec Tunnel

Use the following configuration to configure media gateway to terminate IPSec tunnel:

```
crypto ikev2 proposal default
encryption aes-cbc-256
integrity sha384 sha512 sha256
group 20 16 19 14 21 2
!
crypto ikev2 keyring ikev2_keyring
peer mypeer
address 0.0.0.0 0.0.0.0
pre-shared-key cisco123
!
crypto ikev2 profile ikev2-profile
match identity remote address 0.0.0.0 0.0.0.0
authentication remote pre-share
```

□ pre-shared key, should be same on cEdge

```
authentication local pre-share
keyring local ikev2_keyring
lifetime 14400
dpd 10 3 on-demand
crypto ipsec transform-set tfs esp-gcm 256
mode tunnel
crypto ipsec profile ipsec_profile
set security-association lifetime seconds 7200
set security-association replay window-size 256
set transform-set tfs
set pfs group16
set ikev2-profile ikev2_profile
!
interface Tunnel100
ip address 10.2.2.2 255.255.255.0
tunnel source GigabitEthernet1
tunnel mode ipsec ipv4
tunnel destination 10.74.5.213
tunnel protection ipsec profile ipsec_profile
!
```

□ Tunnel address
□ MD GW phy interface
□ cEdge wan interface

Verify Static Tunnel with Media Device Gateway

The IPsec tunnel between the Cisco IOS XE SD-WAN device and the Media Device gateway is static and is always in the UP state.

Use the following commands to verify static tunnel configuration with the Media Device gateway:

- **show crypto session detail**
- **show crypto ipsec sa**



CHAPTER 13

Policy Applications Using CLIs

CLI commands for configuring and monitoring policy applications.

Application-Aware Routing Command Hierarchy

Configure and apply the policy on Cisco vSmart Controllers:

```
policy
  lists
    app-list list-name
      (app application-name | app-family application-family)
    data-prefix-list list-name
      ip-prefix prefix/length
    site-list list-name
      site-id site-id
    vpn-list list-name
      vpn vpn-id
    sla-class sla-class-name
      jitter milliseconds
      latency milliseconds
      loss percentage

policy
  app-route-policy policy-name
  vpn-list list-name
  default-action sla-class sla-class-name
  sequence number
  match
    app-id app-id-name
    app-list list-name
    destination-data-prefix-list list-name
    destination-ip prefix/length
    destination-port number
    dns (request | response)
    dns-app-list list-name
    dscp number
    plp (high | low)
    protocol number
    source-data-prefix-list list-name
    source-ip prefix/length
    source-port number
  action
    backup-sla-preferred-color colors
    count
    log
    sla-class sla-class-name [strict] [preferred-color colors]
```

```

apply-policy site-list list-name
app-route-policy policy-name

```

Cflowd Traffic Flow Monitoring Command Hierarchy

Configure on Cisco vSmart Controllers only:

```

policy
  lists
    prefix-list list-name
      ip-prefix prefix/length
    site-list list-name
      site-id site-id
    vpn-list list-name
      vpn vpn-id
  cflowd-template template-name
    collector vpn vpn-id address ip-address port port-number transport transport-type
    flow-active-timeout seconds
    flow-inactive-timeout seconds
    flow-sampling-interval number
    template-refresh seconds

policy
  data-policy policy-name vpn-list list-name
  default-action action
  sequence number
  match
    destination-data-prefix-list list-name
    destination-ip prefix/length
    destination-port number
    dscp number
    protocol number
    source-data-prefix-list list-name
    source-ip prefix/length
    source-port number
  action
    count counter-name
    drop
    accept
      cflowd

apply-policy
  site-list list-name
  data-policy policy-name direction
  cflowd-template template-name

```

Local Internet Exit Command Hierarchy

Configure and apply a centralized data policy on the Cisco vSmart Controller:

```

policy
  lists
    prefix-list list-name
      ip-prefix prefix/length
    site-list list-name
      site-id site-id
    vpn-list list-name
      vpn vpn-id
  cflowd-template template-name
    collector vpn vpn-id address ip-address port port-number
    flow-active-timeout seconds
    flow-inactive-timeout seconds
    template-refresh seconds

```

```

policy
  data-policy policy-name vpn-list list-name
  default-action action
  sequence number
  match
    destination-data-prefix-list list-name
    destination-ip prefix/length
    destination-port number
    dscp number
    protocol number
    source-data-prefix-list list-name
    source-ip prefix/length
    source-port number
  action
    count counter-name
    drop
    accept
      nat use-vpn 0

apply-policy
  site-list list-name
  data-policy policy-name direction

```

Zone-Based Firewalls

```

policy
  lists
    prefix-list list-name
      ip-prefix prefix/length
  tcp-syn-flood-limit number
  zone (destination-zone-name | source-zone-name)
    vpn vpn-id
  zone-to-no-zone-internet (allow | deny)
  zone-pair pair-name
    source-zone source-zone-name
    destination-zone destination-zone-name
  zone-policy policy-name
  zone-based-policy policy-name
  default-action action
  sequence number
  match
    destination-data-prefix-list list-name
    destination-ip prefix/length
    destination-port number
    protocol number
    source-data-prefix-list list-name
    source-ip prefix-length
    source-port number
  action
    drop
    inspect
    log
    pass

```

