



Cisco Catalyst SD-WAN Policies Configuration Guide, Cisco IOS XE Catalyst SD-WAN Release 17.x

First Published: 2020-03-17

Last Modified: 2024-08-27

Americas Headquarters

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

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

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

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

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

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

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

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

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

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

© 2019–2024 Cisco Systems, Inc. All rights reserved.



CONTENTS

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

CHAPTER 2	What's New in Cisco IOS XE (SD-WAN)	3
------------------	--	----------

CHAPTER 3	Policy Overview	5
	Policy Architecture	7
	Centralized Control Policy Architecture	7
	Route Types	8
	Default Behavior Without Centralized Control Policy	9
	Behavior Changes with Centralized Control Policy	9
	Examples of Modifying Traffic Flow with Centralized Control Policy	10
	Configure Centralized Policy Based on Prefixes and IP Headers	14
	Cisco Catalyst SD-WAN Controller Policy Components	15
	TLOC Attributes Used in Policies	19
	Cisco Catalyst SD-WAN Route Attributes Used in Policies	19
	Design Cisco Catalyst SD-WAN Controller Policy Processing and Application	20
	Cisco Catalyst SD-WAN Controller Policy Operation	21
	Control Policy	22
	Data Policy	24
	VPN Membership Policy Operation	26
	Configure and Execute Cisco SD-WAN Controller Policies	27

CHAPTER 4	Centralized Policy	29
	Overview of Centralized Policies	29
	Types of Centralized Policies	29
	Configure Centralized Policies Using Cisco SD-WAN Manager	30

Start the Policy Configuration Wizard	30
Configure Groups of Interest for Centralized Policy	31
Integrating WAN Insight (WANI) into Cisco SD-WAN Manager	38
Predictive Path Recommendations	39
Configure Topology and VPN Membership	40
Import Existing Topology	42
Create a VPN Membership Policy	42
Configure Traffic Rules	43
Match Parameters - Control Policy	48
Match Parameters - Data Policy	51
Action Parameters - Control Policy	57
Action Parameters - Data Policy	58
Apply Policies to Sites and VPNs	62
NAT Fallback on Cisco IOS XE Catalyst SD-WAN Devices	63
Activate a Centralized Policy	65
Configure Centralized Policies Using the CLI	66
Centralized Policies Configuration Examples	70

CHAPTER 5
Localized Policy 81

Overview of Localized Policies	81
Types of Localized Policies	81
Configure Localized Policy Using Cisco SD-WAN Manager	83
Start the Policy Configuration Wizard	83
Configure Groups of Interest for Localized Policy	83
Configure Forwarding Classes/QoS	86
Configure ACLs	88
Explicit and Implicit Access Lists	89
Configure Route Policies	90
Match Parameters	91
Action Parameters	93
Configure Policy Settings	94
Apply Localized Policy in a Device Template	95
Activate a Localized Policy	96
Configure Localized Policy for IPv4 Using the CLI	97

Configure Localized Policy for IPv6 Using the CLI	99
Localized Data Policy Configuration Examples	100
QoS For Router Generated Cisco SD-WAN Manager Traffic	101
Information About QoS For Router-Generated Cisco SD-WAN Manager Traffic	101
Restrictions For QoS For Router Generated Cisco SD-WAN Manager Traffic	102
Configure QoS for Router Generated Cisco SD-WAN Manager Traffic Using a CLI Template	102
Verify QoS for Router Generated Cisco SD-WAN Manager Traffic Using CLI	103
Troubleshooting QoS For Router Generated Cisco SD-WAN Manager Traffic	104

CHAPTER 6**Redirect DNS in a Service-Side VPN 105**

Information About Redirect DNS in a Service-Side VPN	105
Restrictions for Redirect DNS in a Service-Side VPN	106
Use Cases for Redirect DNS in a Service-Side VPN	106
Configure Redirect DNS in a Service-Side VPN	108
Configure Redirect DNS in a Service-Side VPN Using the CLI	110
Verify Redirect DNS in a Service-Side VPN	112
Configuration Examples for Redirect DNS	112

CHAPTER 7**Default AAR and QoS Policies 115**

Information About Default AAR and QoS Policies	115
Benefits of Default AAR and QoS Policies	116
Prerequisites for Default AAR and QoS Policies	116
Restrictions for Default AAR and QoS Policies	117
Supported Devices for Default AAR and QoS Policies	117
Use Cases for Default AAR and QoS Policies	117
Configure Default AAR and QoS Policies Using Cisco SD-WAN Manager	117
Monitor Default AAR and QoS Policies	122

CHAPTER 8**Device Access Policy 123**

Device Access Policy Overview	123
Configure Device Access Policy Using Cisco SD-WAN Manager	124
Configure Device Access Policy Using the CLI	125
Examples for ACL Statistics and Counters	126
Verifying ACL Policy on an SNMP Server	127

Verifying ACL Policy on SSH 129

CHAPTER 9

Cisco Catalyst SD-WAN Application Intelligence Engine Flow 131

Cisco Catalyst SD-WAN Application Intelligence Engine Flow Overview 131

Configure Cisco Catalyst SD-WAN Application Intelligence Engine Flow Using Cisco SD-WAN Manager 132

Apply Centralized Policy for SD-WAN Application Intelligence Engine Flow 132

Monitor Running Applications 132

View SAIE Applications 133

Action Parameters for Configuring SD-WAN Application Intelligence Engine Flow 133

Configure SD-WAN Application Intelligence Engine Flow Using the CLI 136

CHAPTER 10

Application-Aware Routing 139

Information About Application-Aware Routing 139

Application-Aware Routing Support for Multicast Protocols 140

Restrictions for Multicast Protocols 140

Components of Application-Aware Routing 141

SLA Classes 142

Classification of Tunnels into SLA Classes 145

Measure Loss, Latency, and Jitter 145

Calculate Average Loss, Latency, and Jitter 146

Determine SLA Classification 146

Per-Class Application-Aware Routing 147

Per-Class Application-Aware Routing Overview 147

Application Probe Class 147

Default DSCP Values 148

Configure Application-Aware Routing 148

Configure Application-Aware Routing Policies Using Cisco SD-WAN Manager 149

Configure Best Tunnel Path 149

Best Tunnel Path Overview 149

Recommendation for the Best Tunnel Path 150

Configure Variance for Best Tunnel Path 150

Verify Configuration of Variance for Best Tunnel Path 151

Configure SLA Class 152

Configure Traffic Rules	153
Default Action of Application-Aware Routing Policy	158
Configure Application Probe Class through Cisco Catalyst SD-WAN Manager	158
Add App-Probe-Class to an SLA Class	159
Configure Default DSCP on Cisco BFD Template	159
Apply Policies to Sites and VPNs	160
How Application-Aware Routing Policy is Applied in Combination with Other Data Policies	161
Activate an Application-Aware Routing Policy	162
Monitor Data Plane Tunnel Performance	163
Enable Application Visibility on Cisco IOS XE Catalyst SD-WAN Devices	164
Configure Application-Aware Routing Using CLIs	164
Configure Application Probe Class Using CLI	166
Application-Aware Routing Policy Configuration Example	167

CHAPTER 11**Enhanced Application-Aware Routing 173**

Information About Enhanced Application-Aware Routing	173
Overview of Enhanced Application-Aware Routing	174
PfR Measurements	174
Application-Aware Routing Design and Measurements	175
Benefits of Enhanced Application-Aware Routing	175
Guidelines of Enhanced Application-Aware Routing	176
Compatibility With Cisco IOS XE Catalyst SD-WAN devices Not Running Enhanced Application-Aware Routing	176
Supported Devices for Enhanced Application-Aware Routing	177
Restrictions for Enhanced Application-Aware Routing	177
Prerequisites for Enhanced Application-Aware Routing	177
Configure Enhanced Application-Aware Routing	177
Configure Enhanced Application-Aware Routing Using a Feature Template in Cisco Catalyst SD-WAN Manager	177
Configure Enhanced Application-Aware Routing Using a Configuration Group in Cisco Catalyst SD-WAN Manager	178
Configure Enhanced Application-Aware Routing Using a CLI Template	178
Verify the Enhanced Application-Aware Routing Configuration	179
Monitor Enhanced Application-Aware Routing Using Cisco Catalyst SD-WAN Manager	180

Troubleshooting Enhanced Application-Aware Routing 181

CHAPTER 12

Traffic Flow Monitoring 183

Traffic Flow Monitoring 183

Information About Traffic Flow Monitoring 185

Traffic Flow Monitoring with Cflowd Overview 185

IPFIX Information Elements for Cisco IOS XE Catalyst SD-WAN Devices 186

Flexible Netflow for VPN0 Interface 190

Limitations of Flexible Netflow on VPN0 Interface 191

Flexible NetFlow Export Spreading 192

Flexible NetFlow Export of BFD Metrics 193

How the Export of BFD Metrics Works 193

Cflowd Traffic Flow Monitoring with SAIE Flows 193

Benefits of Cflowd Traffic Flow Monitoring with SAIE Flows 194

Prerequisites for Cflowd Traffic Flow Monitoring with SAIE Flows 194

Restrictions for Cflowd Traffic Flow Monitoring with SAIE Flows 194

Information About Configuring a Maximum FNF Record Rate for Aggregated Data 194

Restrictions for Traffic Flow Monitoring 195

Restrictions for Enabling Collect Loopback in Flow Telemetry When Using Loopbacks as TLOCs 195

Configure Traffic Flow Monitoring 195

Configure Traffic Flow Monitoring on Cisco IOS XE Catalyst SD-WAN Devices 195

Configure Global Flow Visibility 195

Configure Global Application Visibility 197

Configure Cflowd Monitoring Policy 199

View Cflowd Information 201

Configure Cflowd Traffic Flow Monitoring Using the CLI 202

Configure Flexible Netflow on VPN0 Interface 204

Configure Flexible NetFlow with Export of BFD Metrics Using the CLI 204

Configuration Examples for Flexible NetFlow Export of BFD Metrics 205

Apply and Enable Cflowd Policy 206

Cflowd Traffic Flow Monitoring Configuration Examples 207

Configure the Maximum FNF Record Rate for Aggregated Data, Using CLI Commands 212

Verify Traffic Flow Monitoring 213

Verify Collect Loopback 213

Verify Interface Binding on the Device	215
Verify Flexible Netflow Configuration on VPN0 Interface	216
Verify Flexible NetFlow Configuration with Export of BFD Metrics	219

CHAPTER 13	Application Performance Monitor	221
	Overview of Application Performance Monitor	221
	Limitations and Restrictions	223
	Configure Application Performance Monitor	224
	Verify Performance Monitoring Configuration	225

CHAPTER 14	Enhanced Policy Based Routing	237
	Overview of ePBR	237
	Configure ePBR	239
	Monitor ePBR	242

CHAPTER 15	Forward Error Correction	245
	Supported Devices for Forward Error Correction	245
	Configure Forward Error Correction for a Policy	245
	Monitor Forward Error Correction Tunnel Information	246
	Monitor Forward Error Application Family Information	247
	Monitor Forward Error Correction Status Using the CLI	247

CHAPTER 16	Packet Duplication	249
	Information about Packet Duplication	249
	Configure Packet Duplication Using Centralized Policy	250
	Configure Packet Duplication Using Policy Groups	251
	Configure Underlay Fragmentation Using Cisco SD-WAN Manager	252
	Restrictions for Packet Duplication	252
	Monitor Packet Duplication Statistics for a Device	252
	Monitor Tunnel Information for a Device	252

CHAPTER 17	Policy Configuration Tagging	253
	Supported Devices for Policy Configuration Tagging	254

Restrictions for Policy Configuration Tagging	254
Information About Policy Configuration Tagging	254
Benefits of Policy Configuration Tagging	257
Configure Policy Configuration Tagging Using a CLI Template	258
Verify Tag-Instances Configuration Using the CLI	260

CHAPTER 18	Integrate Cisco IOS XE Catalyst SD-WAN Device with Cisco ACI	263
	Guidelines to Integrate with Cisco ACI	264
	Verify Cisco ACI Registration	264
	SLA Classes	264
	Data Prefixes	265
	VPNs	265
	Map Data Prefix and VPN to SLA	265
	Create an App-Route-Policy	265
	Map ACI Sites	266
	Unmap ACI Sites	267
	Delete a Controller	267

CHAPTER 19	Custom Applications	269
	Information About Custom Applications	269
	Restrictions for Custom Applications	271
	Configure Custom Applications Using Cisco SD-WAN Manager	272
	Verify Custom Applications	273

CHAPTER 20	Service Insertion	275
	Information About Service Insertion	276
	Restrictions for Service Insertion	280
	Use Cases for Service Insertion	280
	Configure Service Insertion	281
	Configure Service Chain Actions in a Data Policy	282
	Traffic Steering to a Service Chain	283
	Traffic Steering Using a Control Policy	283
	Traffic Steering Using a Data Policy	284
	Traffic Steering Using an Interface Access Control List	285

Path Preference	286
Share Service Chains Across User VPN	286
Separate Interfaces for Transmitted and Received Traffic	287
Service Chaining Trusted and Untrusted Traffic	287
Service Chain Between Two Routers	288
Configure Fall Back and Restrict Behavior for Traffic Through a Service Chain	288
Interfaces for Attaching Services in a Service Chain to a Router	289
Service Chaining with Software Defined Cloud Interconnect Bring Your Own Service	289
Configure Service Insertion Using a CLI Template	290

CHAPTER 21 **Service Chaining** 291

Configure Service Chaining	294
Service Chaining Configuration Examples	295
Monitor Service Chaining	303

CHAPTER 22 **Lawful Intercept** 307

Information About Lawful Intercept	307
Prerequisites for Lawful Intercept	310
Install Lawful Intercept using Cisco Catalyst SD-WAN Manager	311
Lawful Intercept MIBs	312
Restrict Access to Trusted Hosts (Without Encryption)	312
Restrict Trusted Mediation Device	313
Configure Lawful Intercept	313
Configure Lawful Intercept Using CLI	313
Encrypt Lawful Intercept Traffic	314
Configure Encryption in the Device	315
Configure Lawful Intercept Encryption using CLI	315
Verify Static Tunnel with Media Device Gateway	316

CHAPTER 23 **Lawful Intercept 2.0** 317

Information About Lawful Intercept 2.0	318
Prerequisites for Cisco Catalyst SD-WAN Lawful Intercept 2.0	320
Benefits of Cisco Catalyst SD-WAN Lawful Intercept 2.0	320
Configure Lawful Intercept 2.0 Workflow	320

Create a Lawful Intercept Administrator	320
Create a Lawful Intercept API User	321
Create an Intercept	321
Retrieve an Intercept	323
Troubleshooting Cisco SD-WAN Controller for Lawful Intercept from Cisco SD-WAN Manager	323

CHAPTER 24	Troubleshoot Cisco Catalyst SD-WAN Policies	325
	Overview	325
	Support Articles	325
	Feedback Request	326
	Disclaimer and Caution	326



CHAPTER 1

Read Me First



Note To achieve simplification and consistency, the Cisco SD-WAN solution has been rebranded as Cisco Catalyst SD-WAN. In addition, from Cisco IOS XE SD-WAN Release 17.12.1a and Cisco Catalyst SD-WAN Release 20.12.1, the following component changes are applicable: **Cisco vManage to Cisco Catalyst SD-WAN Manager**, **Cisco vAnalytics to Cisco Catalyst SD-WAN Analytics**, **Cisco vBond to Cisco Catalyst SD-WAN Validator**, **Cisco vSmart to Cisco Catalyst SD-WAN Controller**, and **Cisco Controllers to Cisco Catalyst SD-WAN Control Components**. See the latest Release Notes for a comprehensive list of all the component brand name changes. While we transition to the new names, some inconsistencies might be present in the documentation set because of a phased approach to the user interface updates of the software product.

Related References

- [Cisco Catalyst SD-WAN Control Components Compatibility Matrix and Server Recommendations](#)
- [Cisco Catalyst SD-WAN Device Compatibility](#)

User Documentation

- [User Documentation for Cisco IOS XE Catalyst SD-WAN Release 17](#)

Communications, Services, and Additional Information

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

Documentation Feedback

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



CHAPTER 2

What's New in Cisco IOS XE (SD-WAN)

[What's New in Cisco IOS XE Catalyst SD-WAN Release 17.x](#)



CHAPTER 3

Policy Overview

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

Figure 1: Policy Architecture



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 Catalyst SD-WAN policy design provides a clear separation between centralized and localized policy. In short, centralized policy is provisioned on the centralized Cisco SD-WAN Controllers in the overlay network, and the localized policy is provisioned on Cisco IOS XE Catalyst 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 SD-WAN Controllers, which are the centralized controllers in the Cisco Catalyst 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 SD-WAN Controller's route table and that is advertised to the Cisco IOS XE Catalyst SD-WAN devices. The effects of centralized control policy are seen in how Cisco IOS XE Catalyst SD-WAN devices direct the overlay network's data traffic to its destination.



Note The centralized control policy configuration itself remains on the Cisco SD-WAN 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 IOS XE Catalyst SD-WAN devices.

Localized Policy

Localized policy refers to a policy that is provisioned locally through the CLI on the Cisco IOS XE Catalyst SD-WAN devices, or through a Cisco SD-WAN Manager 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 Catalyst 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 Catalyst 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 Catalyst SD-WAN devices, either on the centralized Cisco SD-WAN Controllers or the local Cisco IOS XE Catalyst SD-WAN devices. When control plane traffic, which distributes route information, is unpoliced:

- All route information that OMP propagates among the Cisco IOS XE Catalyst SD-WAN devices is shared, unmodified, among all Cisco SD-WAN Controllers and all Cisco IOS XE Catalyst 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 Catalyst 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 Catalyst SD-WAN device's route table, and all VPNs in the overlay network can exchange data traffic.

- [Policy Architecture, on page 7](#)
- [Cisco Catalyst SD-WAN Controller Policy Components, on page 15](#)
- [Design Cisco Catalyst SD-WAN Controller Policy Processing and Application, on page 20](#)
- [Cisco Cisco Catalyst SD-WAN Controller Policy Operation, on page 21](#)
- [Configure and Execute Cisco SD-WAN Controller Policies, on page 27](#)

Policy Architecture

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

With Cisco Catalyst SD-WAN, you can also create routing policies on the Cisco IOS XE Catalyst 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.

Centralized Control Policy Architecture


In the Cisco IOS XE Catalyst SD-WAN network architecture, centralized control policy is handled by the Cisco SD-WAN Controller, which effectively is the routing engine of the network. The Cisco SD-WAN Controller is the centralized manager of network-wide routes, maintaining a primary route table for these routes. The Cisco SD-WAN Controller builds its route table based on the route information advertised by the Cisco IOS XE Catalyst 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 SD-WAN 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 SD-WAN Controllers. This type of policy, which is provisioned centrally on the Cisco SD-WAN Controller, affects both the route information that the Cisco SD-WAN 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 SD-WAN 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 SD-WAN Controllers, which are responsible for orchestrating the routing decisions in the Cisco IOS XE Catalyst SD-WAN overlay network.

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

Figure 2: Centralized Control Policy



All centralized control plane traffic, including route information, is carried by OMP peering sessions that run within the secure, permanent DTLS connections between devices and the Cisco SD-WAN 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 SD-WAN 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 SD-WAN 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 SD-WAN Controller advertises it, and affects whether the route information is advertised as-is or is modified.

Route Types

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

- Cisco IOS XE Catalyst 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 SD-WAN 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 SD-WAN Controller or a Cisco IOS XE Catalyst 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 SD-WAN Controller. This results in the following route advertisement and redistribution behavior within a domain:

- All Cisco IOS XE Catalyst SD-WAN devices redistribute all the route-related prefixes that they learn from their site-local network to the Cisco SD-WAN Controller. This route information is carried by OMP route advertisements that are sent over the DTLS connection between the devices and the Cisco SD-WAN Controller. If a domain contains multiple Cisco SD-WAN Controllers, the devices send all OMP route advertisements to all the controllers.
- All the devices send all TLOC routes to the Cisco SD-WAN 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 SD-WAN 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 SD-WAN Controller tracks which OMP routes, TLOCs, and services belong to which VPNs. The Cisco SD-WAN 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 SD-WAN 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 SD-WAN Controller.
- The Cisco SD-WAN 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 Catalyst SD-WAN devices in a domain, or when you want to modify the route information that is stored in the Cisco Catalyst SD-WAN Controller's route table or that is advertised by the Cisco Catalyst SD-WAN 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 Catalyst

SD-WAN Controller. All provisioning of centralized control policy is done on the Cisco Catalyst SD-WAN Controller.

Applying a centralized control policy in the inbound direction filters or modifies the routes being advertised by the Cisco IOS XE Catalyst SD-WAN device before they are placed in the route table on the Cisco Catalyst SD-WAN 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 Catalyst SD-WAN 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 Catalyst SD-WAN Controller redistributes to the Cisco IOS XE Catalyst 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 Catalyst SD-WAN 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 Catalyst 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 Catalyst 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 Catalyst SD-WAN 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 Catalyst SD-WAN devices, if you have provisioned no control policy, the two devices establish an IPsec tunnel between them and the data traffic flows directly from one device to the next. For a network with only two devices or with just a small number of devices, establishing connections between each pair of devices is generally not been an issue. However, such a solution does not scale. In a network with hundreds or even thousands of branches, establishing a full mesh of IPsec tunnels tax the CPU resources of each device.

Figure 3: Arbitrary Topology



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


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

To set up a hub-and-spoke-type topology here, we provision a control policy that causes the West and East devices to send all data packets destined for the other device to the hub device. (Remember that because control policy is always centralized, you provision it on the Cisco Catalyst SD-WAN 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 Catalyst SD-WAN device to always transit through one of the hub devices, set the TLOC preference value to favor the desired hub device.

Figure 4: Traffic Engineering Topology




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 Catalyst SD-WAN Controller to monitor the path from the source to the destination, and to inform the source device when that path is not available. The source device can then modify or remove the path from its route table.


Figure 5: Traffic Engineering 2



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

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

Figure 6: Traffic Engineering 3



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

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


Configure Centralized Policy Based on Prefixes and IP Headers

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

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

Configuration Components

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



Cisco Catalyst SD-WAN Controller Policy Components

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

The implementation of Cisco SD-WAN Controller policy is done by configuring the entire policy on the Cisco Catalyst SD-WAN Control Components. Cisco SD-WAN Controller 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 SD-WAN Controller policy. More specifically, policy is the sum of one or more lists, one policy definition, and at least one policy applications, as shown in the table below.

Table 1: The Three Building Blocks of Cisco SD-WAN Controller Policies

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

Lists	Policy Definition	Policy Application
=		
Complete policy definition configured on Cisco SD-WAN Controller and enforced either on Cisco SD-WAN Controller or on Cisco IOS XE Catalyst 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 SD-WAN Controller policy, you invoke lists in two places: when you create a policy definition and when you apply a policy. Separating the definition of the related items from the definition of policy means that when you can add or remove items from a lists, you make the changes only in a single place: You do not have to make the changes through the policy definition. So if you add ten sites to your network and you want to apply an existing policy to them, you simply add the site identifiers to the site list. You can also change policy rules without having to manually modify the prefixes, VPNs, or other things that the rules apply to.

Table 2: List Types

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

```

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

```

!
!

```

Policy Definition

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

Table 3: Policy Types

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

Table 4:

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

Cisco Catalyst SD-WAN Route Attributes Used in Policies

A Cisco Catalyst SD-WAN route, defines a route in the overlay network and is similar to a standard IP route, has a TLOC and VPN attributes. The Cisco IOS XE Catalyst SD-WAN devices exchange routes in OMP updates.

The routes attributes listed below can be matched or set in Cisco SD-WAN Controller policies.

Table 5:

Route 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 route.	Configuration on source device	control-policy
Site ID	Identifier for a give site. A site can have multiple nodes or TLOCs.	Configuration on source device	control-policy
Tag	Identification on any arbitrary basis.	Configuration on source device	control-policy
TLOC	TLOC used as next hop for the route.	Configuration on source device or policy	control-policy data-policy
VPN	VPN to which the route belongs.	Configuration on source device or policy	control-policy data-policy

Design Cisco Catalyst SD-WAN Controller Policy Processing and Application

Understanding how a Cisco SD-WAN 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 SD-WAN 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 Cisco SD-WAN 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 Catalyst 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 SD-WAN Controller.
- Control-policy remains on the Cisco SD-WAN Controller and affects routes that the controller sends and receives.
- Data-policy is sent to either the Cisco IOS XE Catalyst 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 Catalyst SD-WAN Controller that distributes routing information to the Cisco IOS XE Catalyst SD-WAN device nodes.
- In a network deployment that has two or more Cisco Catalyst SD-WAN Controllers, each controller acts independently to disseminate routing information to other Cisco SD-WAN Controllers and to Cisco IOS XE Catalyst SD-WAN devices in the overlay network. So, to ensure that the Cisco SD-WAN Controller policy has the desired effect in the overlay network, each Cisco SD-WAN 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 SD-WAN Controllers.



Note When you deploy a policy, the deployment status is updated only for 30 minutes, which is the timeout limit for policies. After the timeout period, the deployment task status is not monitored. If you are deploying a bigger policy with more number of lines, and if it takes more than 30 minutes, the task status will not be monitored.

Cisco Cisco Catalyst SD-WAN Controller Policy Operation

At a high level, control policy operates on routing information, which in the Cisco IOS XE Catalyst 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 SD-WAN Controller policies are:

- Control Policy
- Data Policy
- VPN Membership

Control Policy

Control policy, which is similar to standard routing policy, operates on routes and routing information in the control plane of the overlay network. Centralized control policy, which is provisioned on the Cisco SD-WAN Controller, is the Cisco Catalyst 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 Catalyst SD-WAN device, allows customization of routing decisions made by BGP and OSPF on site-local branch or enterprise networks.

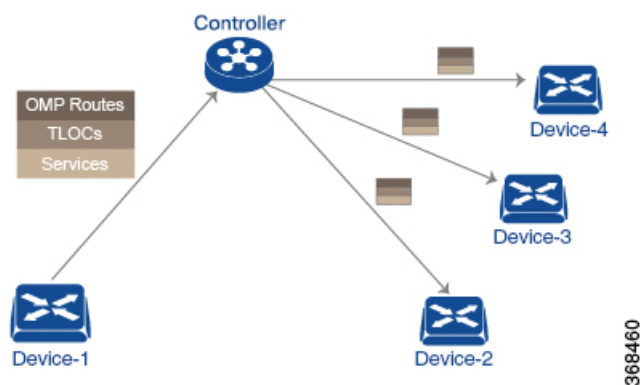
The routing information that forms the basis of centralized control policy is carried in Cisco IOS XE Catalyst SD-WAN route advertisements, which are transmitted on the DTLS or TLS control connections between Cisco SD-WAN Controllers and Cisco IOS XE Catalyst SD-WAN devices. Centralized control policy determines which routes and route information are placed into the centralized route table on the Cisco SD-WAN Controller and which routes and route information are advertised to the Cisco IOS XE Catalyst 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 Catalyst 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 SD-WAN Controller throughout the overlay network. The Cisco SD-WAN Controller learns the overlay network topology from OMP routes that are advertised by the Cisco IOS XE Catalyst SD-WAN devices over the OMP sessions inside the DTLS or TLS connections between the Cisco SD-WAN Controller and the devices.

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

- Cisco Catalyst 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 SD-WAN Controller. These routes are also referred to as OMP routes or Routes.
- 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 Catalyst 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 SD-WAN Controller's route table as is, and the Cisco SD-WAN 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 SD-WAN 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 Catalyst SD-WAN devices place all the route information learned from the Cisco SD-WAN Controllers, as is, into their local route tables, for use when forwarding data traffic. Because the Cisco SD-WAN Controller's role is to be the centralized routing system in the network, Cisco IOS XE Catalyst SD-WAN devices can never modify the OMP route information that they learn from the Cisco SD-WAN Controllers.

The Cisco SD-WAN 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 SD-WAN Controller remains on the Cisco SD-WAN 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.

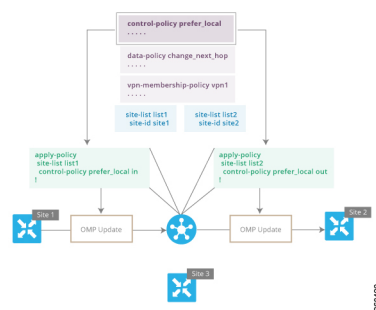
The Cisco IOS XE Catalyst SD-WAN devices periodically exchange OMP updates, which carry routing information pertaining to the overlay network. Two of the things that these updates contain are Route attributes and Transport Locations (TLOC) attributes.

The Cisco SD-WAN 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 Catalyst SD-WAN devices in the network by sending OMP updates to them.

Control policy examines the Route 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 SD-WAN Controller and that is applied to Site 1 (via site-list list1) and to Site 2 (via site-list list2).

Figure 8: Control Policy Topology



```
Device# apply-policy
site-list list1
```

```
control-policy prefer_local in
!
```

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

The route table on the Cisco SD-WAN Controller is used to determine the topology of the overlay network. The Cisco SD-WAN 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 SD-WAN Controller. It determines the network topology and network reachability, modifying Route 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 SD-WAN Controller is sending to the devices at Site 2. Any changes that result from the policy occur, after the information from the Cisco SD-WAN 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 SD-WAN Controller.

In contrast to an inbound policy, which affects the centralized route table on the Cisco SD-WAN 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 Catalyst SD-WAN control policy design architecture and configuration.

Data Policy


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



The Cisco IOS XE Catalyst 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 SD-WAN 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 Catalyst 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 a centralized data policy on the Cisco SD-WAN Controller, and that configuration remains on the Cisco SD-WAN Controller. The effects of data policy are reflected in how the Cisco IOS XE Catalyst 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.

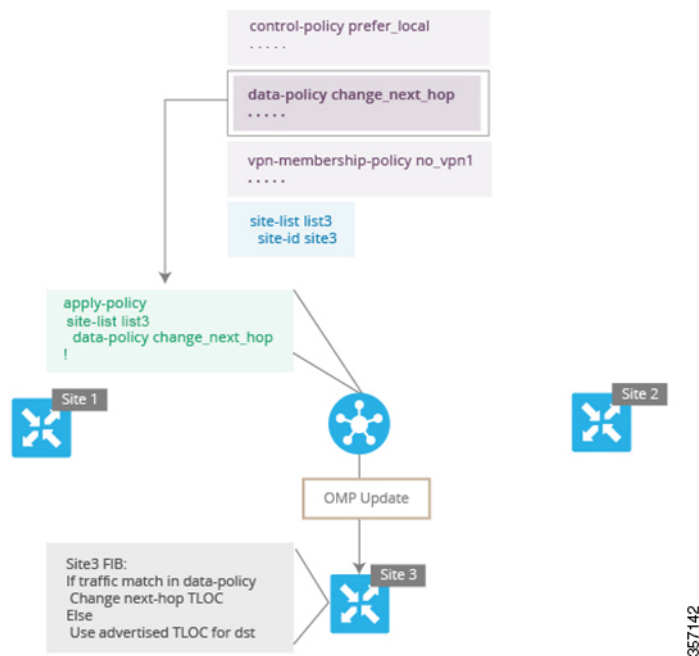
With no access lists provisioned on a Cisco IOS XE Catalyst SD-WAN device, all data traffic is transmitted at line rate and with equal importance, using one of the interface's queues. Using access lists, you can provision class of service, which allows you to classify data traffic by importance, spread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

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 SD-WAN Controller, and then it is carried in OMP updates to the Cisco IOS XE Catalyst 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 Cisco SD-WAN 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. Non-matching traffic is forwarded to the original next-hop TLOC.

Figure 9: Data Policy Topology



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

VPN Membership Policy Operation

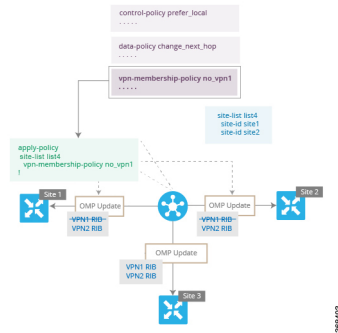
VPN membership policy, as the name implies, affects the VPN route tables that are distributed to particular Cisco IOS XE Catalyst SD-WAN devices. In an overlay network with no VPN membership policy, the Cisco Catalyst SD-WAN Controller pushes the routes for all VPNs to all the devices. If your business 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 Catalyst SD-WAN devices:

- The Cisco IOS XE Catalyst SD-WAN devices at Sites 1 and 2 service only VPN 2.
- The Cisco IOS XE Catalyst 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 SD-WAN Controller, because these updates are for both VPN 1 and VPN 2. However, because the other Cisco IOS XE Catalyst 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 10: VPN Membership Topology





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

Configure and Execute Cisco SD-WAN Controller Policies

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

Figure 11: Cisco SD-WAN Controller Policy

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

368503

For control policy and VPN membership policy, the entire policy configuration remains on the Cisco SD-WAN Controller, and the actions taken as a result of routes or VPNs that match a policy are performed on the Cisco SD-WAN 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 Catalyst SD-WAN devices, and any actions taken as a result of the policies are performed on the devices.



CHAPTER 4

Centralized Policy

The topics in this section provide overview information about the different types of centralized policies, the components of centralized policies, and how to configure centralized policies using Cisco SD-WAN Manager or the CLI.

- [Overview of Centralized Policies, on page 29](#)
- [Configure Centralized Policies Using Cisco SD-WAN Manager, on page 30](#)
- [Configure Centralized Policies Using the CLI, on page 66](#)
- [Centralized Policies Configuration Examples, on page 70](#)

Overview of Centralized Policies

Centralized policies refer to policies that are provisioned on Cisco SD-WAN Controllers, which are the centralized controllers in the Cisco Catalyst SD-WAN overlay network.

Types of Centralized Policies

Centralized Control Policy

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



Note The centralized control policy configuration itself remains on the Cisco Catalyst SD-WAN Controller and is never pushed to local devices.

Centralized Data Policy

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 IOS XE Catalyst SD-WAN devices.

Centralized 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 a 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.

Configure Centralized Policies Using Cisco SD-WAN Manager

To configure a centralized policy, use the Cisco SD-WAN Manager policy configuration wizard. The wizard consists of the following operations 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 and VPN Membership:** 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 the policy with sites and VPNs in the overlay network.
- **Activate the centralized policy.**

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

To configure centralized policies using Cisco SD-WAN Manager, use the steps identified in the procedures that follow this section.

Start the Policy Configuration Wizard

To start the policy configuration wizard:

1. From the Cisco SD-WAN Manager menu, choose **Configuration > Policies**.
2. Click **Centralized Policy**.
3. Click **Add Policy**.

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

Configure Groups of Interest for Centralized Policy

In **Create Groups of Interest**, create new groups of list types as described in the following sections to use in a centralized policy:

Configure Application

1. In the groups of interest list, click **Application** list type.
2. Click **New Application List**.
3. Enter a name for the list.
4. Choose either **Application** or **Application Family**.

Application can be the names of one or more applications, such as **Third Party Control**, **ABC News**, **Microsoft Teams**, and so on. The Cisco IOS XE Catalyst SD-WAN devices support about 2300 different applications. To list the supported applications, use the **?** in the CLI.

Application Family 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**.

5. In the **Select** drop-down, in the 'Search' filter, select the required applications or application families.
6. Click **Add**.

A few application lists are preconfigured. You cannot edit or delete these lists.

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.

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.

Configure Color

1. In the groups of interest list, click **Color**.
2. Click **New Color List**.
3. Enter a name for the list.
4. In the **Select Color** drop-down, in the 'Search' filter select the required colors.

Colors 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.

5. Click **Add**.

To configure multiple colors in a single list, you can select multiple colors from the drop-down.

Configure Community

Table 6: Feature History

Feature Name	Release Information	Description
Ability to Match and Set Communities	Cisco SD-WAN Release 20.5.1 Cisco IOS XE Catalyst SD-WAN Release 17.5.1a Cisco vManage Release 20.5.1	This feature lets you match and set communities using a control policy. Control policies are defined and applied on Cisco IOS XE Catalyst SD-WAN device devices to manipulate communities. With this feature, you can match and assign single or multiple BGP community tags to your prefixes based on which routing policies can be manipulated.

A community list is used to create groups of communities to use in a match clause of a route map. A community list can be used to control which routes are accepted, preferred, distributed, or advertised. You can also use a community list to set, append, or modify the communities of a route.

- In the group of interest list, click **Community**.
- Click **New Community List**.
- Enter a name for the community list.
- Choose either **Standard** or **Expanded**.
 - Standard community lists are used to specify communities and community numbers.
 - Expanded community lists are used to filter communities using a regular expression. Regular expressions are used to specify patterns to match community attributes.
- In the **Add Community** field, enter one or more data prefixes separated by commas in any of the following formats:
 - aa:nn**: Autonomous System (AS) 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 number.
 - no-advertise**: Attaches the NO_ADVERTISE community to routes. Routes in this community are not advertised to other BGP peers.
 - no-export**: Attaches 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.
- Click **Add**.

Configure Data Prefix

- In the **Groups of Interest** list, click **Data Prefix**.

2. Click **New Data Prefix List**.
3. Enter a name for the list.
4. Choose either **IPv4** or **IPv6**.
5. In the **Add Data Prefix** field, enter one or more data prefixes separated by commas.
6. Click **Add**.

Configure Policer

1. In the groups of interest list, click **Policer**.
2. Click **New Policer List**.
3. Enter a name for the list.
4. Define the policing parameters:
 - a. In the **Burst** field, enter the maximum traffic burst size, a value from 15,000 to 10,000,000 bytes.
 - b. 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**.
You can use the **remark** action to set the packet loss priority (PLP) to **high**.
 - c. In the **Rate** field, enter the maximum traffic rate, a value from 0 through $2^{64} - 1$ bits per second (bps).
5. Click **Add**.



Note For ACL or data policy, even if the same policer object is used in different sequences, each sequence is policed separately. For example, consider that policer rate is 1000000 or 1Mb. If the same policer is set for multiple sequences then each sequence polices 1Mb traffic. If there are different policers configured for difference sequences, then each policer polices traffic based on the respective policing rate.

Configure Prefix

1. In the groups of interest list, click **Prefix**.
2. Click **New Prefix List**.
3. Enter a name for the list.
4. In the **Add Prefix** field, enter one or more data prefixes separated by commas.
5. Click **Add**.

Configure Site

1. In the groups of interest list, click **Site**.
2. Click **New Site List**.
3. Enter a name for the list.

4. In the **Add Site** field, enter one or more site IDs separated by commas.
For example, 100 or 200 separated by commas or in the range, 1- 4294967295.
5. Click **Add**.

Configure App Probe Class

1. In the groups of interest list, click **App Probe Class**.
2. Click **New App Probe Class**.
3. Enter the probe class name in the **Probe Class Name** field.
4. Select the required forwarding class from the **Forwarding Class** drop-down list.
5. In the **Entries** pane, select the appropriate color from the **Color** drop-down list and enter the **DSCP** value.
You can add more entries if needed by clicking on the + symbol.
6. Click **Save**.

Configure SLA Class

1. In the groups of interest list, click **SLA Class**.
2. Click **New SLA Class List**.
3. Enter a name for the list.
4. Define the SLA class parameters:
 - a. In the **Loss** field, enter the maximum packet loss on the connection, a value from 0 through 100 percent.
 - b. In the **Latency** field, enter the maximum packet latency on the connection, a value from 0 through 1,000 milliseconds.
 - c. In the **Jitter** field, enter the maximum jitter on the connection, a value from 1 through 1,000 milliseconds.
 - d. Select the required app probe class from the **App Probe Class** drop-down list.
5. (Optional) Select the **Fallback Best Tunnel** checkbox to enable the best tunnel criteria.
This optional field is available from Cisco IOS XE Catalyst SD-WAN Release 17.5.1a to pick the best path or color from the available colors when SLA is not met. When this option is selected, you can choose the required criteria from the drop-down. The criteria are a combination of one or more of losses, latency, and, jitter values.
6. Select the **Criteria** from the drop-down list. The available criteria are:
 - Latency
 - Loss
 - Jitter
 - Latency, Loss

- Latency, Jitter
- Loss, Latency
- Loss, Jitter
- Jitter, Latency
- Jitter, Loss
- Latency, Loss, Jitter
- Latency, Jitter, Loss
- Loss, Latency, Jitter
- Loss, Jitter, Latency
- Jitter, Latency, Loss
- Jitter, Loss, Latency

7. Enter the **Loss Variance (%)**, **Latency Variance (ms)**, and the **Jitter Variance (ms)** for the selected criteria.
8. Click **Add**.

Configure TLOC

1. In the groups of interest list, 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.
The range is 0 to 4294967295.
8. Click **Add TLOC** to add another TLOC to the list.
9. Click **Save**.



Note To use the `set tloc` and `set tloc-list` commands, you must use the `set-vpn` command.

For each TLOC, specify its address, color, and encapsulation. 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 highest preference value is used. If two or more TLOCs have the highest preference value, traffic is sent among them in an ECMP fashion.

If IPsec preference is set on the local preferred color for an edge router, the local TLOC and the color does not overlap with the centralized policy configured with local color preference. The edge router with local TLOC preference takes the precedence. In this case, the preferred TLOC configured in centralized policy is not considered.

Configure VPN

1. In the groups of interest list, 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.
For example, 100 or 200 separated by commas or in the range, 1- 65530.
5. Click **Add**.

Configure Region

Minimum release: Cisco vManage Release 20.7.1

To configure a list of regions for Multi-Region Fabric (formerly Hierarchical SD-WAN), ensure that Multi-Region Fabric is enabled in **Administration > Settings**.

1. In the groups of interest list, click **Region**.
2. Click **New Region List**.
3. In the **Region List Name** field, enter a name for the region list.
4. In the **Add Region** field, enter one or more regions, separated by commas, or enter a range.
For example, specify regions 1, 3 with commas, or a range 1-4.
5. Click **Add**.

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

Configure Preferred Color Group

Table 7: Feature History

Feature Name	Release Information	Description
Tiered Transport Preference in Application-aware Routing and Data Policy	Cisco IOS XE Catalyst SD-WAN Release 17.9.1a Cisco vManage Release 20.9.1	This feature adds support for ranking of Application Aware Routing (AAR) preferred and backup preferred colors. You can configure up to three levels of priority based on the color or path preference on a Cisco IOS XE Catalyst SD-WAN device.

You can configure the order of transport preference to choose the preference order for forwarding traffic.

The **Preferred Color Group** is supported only on overlay traffic but not on DIA traffic.

1. In the groups of interest list, click **Preferred Color Group**.

2. Click **New Preferred Color Group**.
3. In the **Preferred Color Group Name** field, enter a name for the preferred color group.
4. In the **Primary Colors** pane, do the following:
 - a. Choose the color preference from the **Color Preference** drop-down list.
 - b. Choose the path preference from the **Path Preference** drop-down list.

Field	Description
Preferred Color Group Name	Enter a name of the preferred color group.
Color Preference	<p>Choose the color preference from the drop-down list. The options are:</p> <ul style="list-style-type: none"> • default • 3g • biz-internet • blue • bronze • custom1 • custom2, and so on <p>You can select multiple colors.</p>
Path Preference	<p>Choose the path preference from the drop-down list. The options are:</p> <ul style="list-style-type: none"> • Direct Path: Use only a direct path between the source and the destination devices. • Multi Hop Path: In a Multi-Region Fabric network, use a multi-hop path, which includes the core region, between the source and destination devices, even if a direct path is available. • All Paths: Use any path between the source and destination devices. <p>Note This option is equivalent to not configuring path preference at all. If you are applying the policy to a non-Multi-Region Fabric network, use this option.</p>

5. In the **Secondary Colors** pane, do the following:
 - a. Choose the color preference in the **Color Preference** drop-down list.
 - b. Choose the path preference from the **Path Preference** drop-down list.
6. In the **Tertiary Colors** pane, do the following:
 - a. Choose the color preference from the **Color Preference** drop-down list.

- b. Choose the path preference from the **Path Preference** drop-down list.

7. Click **Add**.

The following guidelines are helpful when configuring the ranking for colors:

- Primary preference is mandatory, and at each priority level, at least one preference path or color is mandatory. Both can also be configured.
- More than one color can be configured as a preference.
- If path preference is not configured, all paths are constrained by the preferred colors that are available.
- If color preference is not configured within the constraint of the path preference, then all the colors are available.
- The preferences apply in order of priority to determine the path or color for forwarding traffic.

When the primary, secondary, and tertiary colors are down, packets are not dropped. The traffic falls back to the usual routing preference to choose if any other colors are up.

Integrating WAN Insight (WANI) into Cisco SD-WAN Manager

Table 8: Feature History

Feature Name	Release Information	Description
WAN Insight Policy Automation	Cisco IOS XE Catalyst SD-WAN Release 17.12.1a Cisco Catalyst SD-WAN Manager Release 20.12.1	With this feature, you can apply the recommendations that are available on Cisco SD-WAN Analytics to Cisco SD-WAN Manager AAR policy and view the applied recommendations on Cisco SD-WAN Manager.

Cisco SD-WAN Analytics is a cloud-based analytics service for Cisco Catalyst SD-WAN offering comprehensive insights into application and network performance. The analytics service is available with Cisco DNA Advantage and Cisco DNA Premier software subscriptions. Cisco SD-WAN Analytics collects and stores metadata about traffic flows in its cloud storage and produces analytics based on this collected data. Predictive Path Analytics generates recommendations for path based on long term insights. These recommendations need to be converted into policy created manually on Cisco SD-WAN Manager and then applied to the network.

The Predictive Path Recommendations feature allows you to apply active recommendations to the actionable centralized AAR policy to influence the forwarding decisions in the Cisco Catalyst SD-WAN network. The recommendations are applied as a part of the AAR policy and then pushed to Cisco SD-WAN Controller. The Predictive Path Recommendations are applied to the SD-WAN network as TLOC preferences in AAR policies.

For more information about using Predictive Path Recommendations, see [Predictive Path Recommendations](#).

Apply Predictive Path Recommendations

When there are predictive path recommendations in Cisco SD-WAN Analytics, perform the following steps to apply the recommendations to the Application-Aware routing policies:

1. In the Cisco SD-WAN Manager menu, click the bell icon at the top-right corner. The **Notifications** pane is displayed with active alarms.
2. If there are any **Active Recommendations** in the **Notifications** pane, click on the site to view the recommendations. Alternatively, you can view from the Cisco SD-WAN Manager menu, click **Analytics > Predictive Networks**.
3. Click **Active Recommendations**, and then click **Apply**.
4. In the **Apply Predictive Path Recommendations** window, click **Proceed to Apply** to apply new recommendations.

You can review the applied recommendations in the Cisco SD-WAN Manager generated configs and push the recommendations to Cisco SD-WAN Controller.

Points to Consider

- Cisco SD-WAN Manager pulls recommendations when you log in. If you want to update the recommendations, refresh the page or log in again.
- Cisco SD-WAN Manager support recommendations for application lists which are associated with some AAR policy only. If AAR Policy does not exist for a given application list, the recommendations are not valid and policy processing is not done.
- WAN Insights generates recommendations for standard App Groups even when the AAR Policy is not defined. However, the policy automation is not done since AAR policy is not defined.
- When for the same site and application list, if WANI generates a terminate for a recommendation which is applied and also generates another recommendation, the recommendations are applied based on the preferences.
- Application of WANI recommendations for Cloud OnRamp for SaaS is not supported.

Predictive Path Recommendations

WAN Insights (WANI) allows you to track the performance of your current network setup and tune your policies and paths to achieve the best user experience. Predictive path recommendations influence AAR policy TLOC preferences.

WAN Insights is a predictive network optimization tool that uses a statistical model to examine historical data from Cisco Catalyst SD-WAN, in order to find the best paths for application traffic. WANI analyzes the telemetry data exported during application traffic flows, and then generates long-term recommendations for paths that would reduce the probability of experiencing an SLA violation (for example, low-quality performance).

Predictive network associates some SLA with each application list that is defined in the AAR policy in order to detect SLA violations for the applications. This is used to calculate a probability of SLA violation on a given site and TLOC and generates recommendations.

For more information about configuring group of interest for data policies, see [Configure Groups of Interest for Centralized Policy](#).

Configure Topology and VPN Membership

When you first open the **Configure Topology and VPN Membership** window, the **Topology** window is displayed by default.

To configure topology and VPN membership:

Hub-and-Spoke

1. In the **Add Topology** drop-down, select **Hub-and-Spoke**.
2. Enter a name for the hub-and-spoke policy.
3. Enter a description for the policy.
4. In the **VPN List** field, select the VPN list for the policy.
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.
6. Double-click the **My Hub-and-Spoke** text string, and enter a name for the policy component
7. In the right pane, add hub sites to the network topology:
 - a. Click **Add Hub Sites**.
 - b. In the **Site List** field, select a site list for the policy component.
 - c. Click **Add**.
 - d. Repeat these steps to add more hub sites to the policy component.
8. In the right pane, add spoke sites to the network topology:
 - a. Click **Add Spoke Sites**.
 - b. In the **Site List Field**, select a site list for the policy component.
 - c. Click **Add**.
 - d. Repeat these steps to add more spoke sites to the policy component.
9. Repeat steps as needed to add more components to the hub-and-spoke policy.
10. Click **Save Hub-and-Spoke Policy**.

Mesh

1. In the **Add Topology** drop-down, select **Mesh**.
2. Enter a name for the mesh region policy component.
3. Enter a description for the mesh region policy component.
4. In the **VPN List** field, select the VPN list for the policy.
5. Click **New Mesh Region**.
6. In the **Mesh Region Name** field, enter a name for the individual mesh region.
7. In the **Site List** field, select one or more sites to include in the mesh region.

8. Click **Add**.
9. Repeat these steps to add more mesh regions to the policy.
10. Click **Save Mesh 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 **Sequence Type**. The **Add Custom 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 **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.
9. Click **Actions**. The **Reject** option is selected by default. To configure actions to perform on accepted packets, click the **Accept** option. Then select the action or enter a value for the action.
10. Click **Save Match and Actions**.
11. Click **Sequence Rule** to configure more sequence rules, as desired. Drag and drop to re-order them.
12. Click **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 **Sequence Type**. The **Add Custom 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 **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.
9. Click **Actions**. The **Reject** option is selected by default. To configure actions to perform on accepted packets, click the **Accept** option. Then select the action or enter a value for the action.
10. Click **Save Match and Actions**.
11. Click **Sequence Rule** to configure more sequence rules, as desired. Drag and drop to re-order them.

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

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.



Note Sequence can have either **match app-list** or **dns-app-list** configured for a policy, but not both. Configuring both **match app-list** and **dns-app-list** for a policy is not supported.

NAT DIA fallback and DNS redirection are not supported at the same time in data policy.

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

Default Action

If a selected route or TLOC 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.

If a selected data packet 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.

Import 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. For **Policy Type**, choose the name of the topology you want to import.
4. In the **Policy** drop-down, select a policy to import.



Note The policy configuration wizard does not let you import an already configured policy as in other instances of centralized policies (data, control, or application-aware routing). The policy must be configured in its entirety.

5. Click **Import**.

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

Create a VPN Membership Policy

1. In the **Specify your network topology** area, click **VPN Membership**.
2. Click **Add VPN Membership Policy**.



Note You can add only one VPN membership at a time, therefore all site lists and VPN lists must be included in a single policy.

The **Add VPN Membership Policy** popup displays.

3. Enter a name and description for the VPN membership policy.
4. In the **Site List** field, select the site list.
5. In the **VPN Lists** field, select the VPN list.
6. Click **Add List** to add another VPN to the VPN membership.
7. Click **Save**.
8. Click **Next** to move to **Configure Traffic Rules** in the wizard.

Configure Traffic Rules

Table 9: Feature History

Feature Name	Release Information	Description
Policy Matching with ICMP Message	Cisco IOS XE Release 17.4.1 Cisco vManage Release 20.4.1	This feature provides support for a new match condition that you can use to specify a list of ICMP messages for centralized data policies, localized data policies, and Application-Aware Routing policies.

When you first open the **Configure Traffic Rules** window, **Application-Aware Routing** is selected by default.

You can also view already created AAR routing policies listed in the page. It provides various information related to the policies such as the Name of the policy, Type, Mode, Description, Update By, and Last Updated details.



Note You can refer to the Mode column for the security status details of the policy. The status helps to differentiate whether the policy is used in unified security or not. The mode status is applicable only for security policies and not relevant to any centralized or localized policies.

For more information on configuring traffic rules for the Cisco Catalyst SD-WAN Application Intelligence Engine (SAIE) flow, see [Cisco Catalyst SD-WAN Application Intelligence Engine Flow](#).



Note In Cisco vManage Release 20.7.1 and earlier releases, the SAIE flow is called the deep packet inspection (DPI) flow.

To configure traffic rules for a centralized data policy:

1. Click **Traffic Data**.
2. Click the **Add Policy** drop-down.
3. Click **Create New**. The **Add Data Policy** window displays.
4. Enter a name and a description for the data policy.
5. In the right pane, click **Sequence Type**. The **Add Data Policy** popup opens.
6. Select the type of data policy you want to create, **Application Firewall**, **QoS**, **Traffic Engineering**, or **Custom**.



Note If you want to configure multiple types of data policies for the same match condition, you need to configure a custom policy.

7. A policy sequence containing the text string **Application**, **Firewall**, **QoS**, **Traffic Engineering**, or **Custom** is added in the left pane.
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.
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.

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 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. <p>This match condition is available for IPv6 traffic from Cisco IOS XE Catalyst SD-WAN Release 17.9.1a and Cisco vManage Release 20.9.1.</p>
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, enter the prefix in the Destination: IP Prefix field.

Match Condition	Procedure
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>
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> <p>This match condition is available for IPv6 traffic from Cisco IOS XE Catalyst SD-WAN Release 17.9.1a and Cisco vManage Release 20.9.1.</p>
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>
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>
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>
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>
ICMP Message	<p>To match ICMP messages, in the Protocol field, set the Internet Protocol Number to 1, or 58, or both.</p> <p>Note This field is available from Cisco IOS XE Release 17.4.1, Cisco vManage Release 20.4.1.</p>
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>

Match Condition	Procedure
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>
TCP	<p>a. In the Match conditions, click TCP.</p> <p>b. In the TCP field, syn is the only option available.</p>

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 to IPv4 and IPv6 address families.
11. To select one or more **Match** conditions, click its box and set the values as described.



Note Not all match conditions are available for all policy sequence types.

12. To select actions to take on matching data traffic, click the **Actions** box.
13. To drop matching traffic, click **Drop**. The available policy actions are listed in the right side.
14. To accept matching traffic, click **Accept**. The available policy actions are listed in the right side.
15. Set the policy action as described.



Note Not all actions are available for all match conditions.



Note If IPv4 packet contains non-initial fragment of UDP or TCP datagram, it has no L4 ports information available because there is no UDP or TCP header. For such fragments destination-port or source-port match is ignored.

In the following example, all the UDP packets to destination port 161 and any other IPv4 packets having protocol ID field in IPv4 header set to 17 with IPv4 header having fragment-offset set will be dropped.

```

policy
app-visibility
access-list SDWAN_101
sequence 100
match
destination-port 161
protocol          17
!
action drop
!
!
```

Action Condition	Description	Procedure
Counter	Count matching data packets.	<ol style="list-style-type: none"> 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.
DSCP	Assign a DSCP value to matching data packets.	<ol style="list-style-type: none"> a. In the Action conditions, click DSCP. b. In the DSCP field, type the DSCP value, a number from 0 through 63.
Forwarding Class	Assign a forwarding class to matching data packets.	<ol style="list-style-type: none"> a. In the Match conditions, click Forwarding Class. b. In the Forwarding Class field, type the class value, which can be up to 32 characters long.
Log	<p>Minimum release: Cisco vManage Release 20.11.1 and Cisco IOS XE Release 17.11.1a</p> <p>Click Log to enable logging.</p> <p>When (DP, AAR or ACL) data policy packets are configured with log action, logs generated and logged to syslog. Due to the global log-rate-limit, not all logs are logged. 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.</p>	<ol style="list-style-type: none"> a. In the Action conditions, click Log to enable logging.
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.

Action Condition	Description	Procedure
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. <p>If you choose FEC Adaptive, an additional field, Loss Threshold, displays that allows you to specify the packet loss threshold for automatically enabling FEC.</p> <p>Adaptive FEC starts to work at 2% packet loss; this value is configurable.</p> <p>You can specify a loss threshold of 1 to 5%. The default packet loss threshold is 2%.</p> <ul style="list-style-type: none"> • 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"> In the Match conditions, click Loss Correction. In the Loss Correction field, select FEC Adaptive, FEC Always, or Packet Duplication.
Click Save Match and Actions .		

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

Match Parameters - Control Policy

For OMP and TLOC routes , you can match the following attributes:

Match Condition	Description
Color List	One or more colors. The available colors are: 3g, biz-internet, blue, bronze, custom1,custom2, custom3, default, gold, green, lte, metro-ethernet, mpls, private1 through private6, public-internet, red and silver.

Match Condition	Description
Community List	<p>List of one or more BGP communities. In the Community List field, you can specify:</p> <ul style="list-style-type: none"> • aa:nn: AS 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.
Types	<p>Specifies the community type. Choose Standard to specify communities and community numbers or, Expanded to filter communities using a regular expression. Regular expressions are used to specify patterns to match community attributes.</p>
Criteria OR	<p>Compares each regex string in the community list against the community string of the route.</p> <p>The OR condition is applicable across multiple community lists and is valid for all devices.</p> <p>Starting from Cisco IOS XE Catalyst SD-WAN Release 17.5.1a, the Community Types and Criteria fields are available.</p>
OMP Tag	<p>Tag value associated with the route or prefix in the routing database on the device.</p> <p>The range is 0 through 4294967295.</p>
Origin	<p>Protocol from which the route was learned.</p>
Originator	<p>IP address from which the route was learned.</p>

Match Condition	Description
<p>Path Type</p>	<p>In a Hierarchical SD-WAN architecture, match a route by its path type, which can be one of the following:</p> <ul style="list-style-type: none"> • Hierarchical Path: A route that includes hops from an access region to a border router, through region 0, to another border router, then to an edge router in a different access region • Direct Path: A direct path route from one edge router to another edge router. • Transport Gateway Path: A route that is re-originated by a router that has transport gateway functionality enabled. <p>Note This option is available beginning with Cisco vManage Release 20.8.1.</p>
<p>Preference</p>	<p>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. The range is 0 through 255.</p>
<p>Prefix List</p>	<p>One or more prefixes. Specifies the name of a prefix list.</p>
<p>Not available in Cisco SD-WAN Manager.</p>	<p>Individual site identifier. The range is 0 through 4294967295.</p>
<p>Site</p>	<p>One or more overlay network site identifiers.</p>
<p>Region</p>	<p>Region defined for Hierarchical SD-WAN. The range is 1 to 63.</p> <p>Note This option is available beginning with Cisco vManage Release 20.7.1.</p>
<p>Role</p>	<p>In a Hierarchical SD-WAN architecture, match by the device type, which can be Border Router or Edge Router.</p> <p>Note This option is available beginning with Cisco vManage Release 20.8.1.</p>
<p>TLOC</p>	<p>Individual TLOC address.</p> <p>Note To use the <code>set tloc</code> and <code>set tloc-list</code> commands, you must use the <code>set-vpn</code> command.</p>

Match Condition	Description
VPN	Individual VPN identifier. The range is 0 through 65535.
Carrier	Carrier for the control traffic. Values are: default, carrier1 through carrier8.
Domain ID	Domain identifier associated with a TLOC. The range is 0 through 4294967295.
OMP Tag	Tag value associated with the TLOC route in the route table on the device. The range is 0 through 4294967295.
Site	Individual site contributor or more overlay network site identifiers.. The range is 0 through 4294967295.

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.

Match Parameters - Data Policy

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

Each sequence in a policy can contain one or more match conditions.

Table 10:

Match Condition	Description
Omit	Match all packets.
Applications/Application Family List	Applications or application families. This match condition is available for IPv6 traffic from Cisco IOS XE Release 17.9.1a and Cisco vManage Release 20.9.1.
Destination Data Prefix	Group of destination prefixes, IP prefix and prefix length. The range is 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 [-]).

Match Condition	Description
Destination Region	<p>Choose one of the following:</p> <ul style="list-style-type: none"> • Primary: Match traffic if the destination device is in the same primary region (also called access region) as the source. This traffic reaches the destination using a multi-hop path, through the core region. • Secondary: Match traffic if the destination device is not in the same primary region as the source but is within the same secondary region as the source. This traffic can reach the destination using a direct tunnel, as described for secondary regions. • Other: Match traffic if the destination device is not in the same primary region or secondary region as the source. This traffic requires a multi-hop path from the source to the destination. <p>Note Minimum releases: Cisco vManage Release 20.9.1, Cisco IOS XE Catalyst SD-WAN Release 17.9.1a</p>
DNS Application List	<p>Enables split DNS, to resolve and process DNS requests and responses on an application-by-application basis. Name of an app-list list . This list specifies the applications whose DNS requests are processed.</p> <p>This match condition is available for IPv6 traffic from Cisco IOS XE Release 17.9.1a and Cisco vManage Release 20.9.1.</p>
DNS	<p>Specify the direction in which to process DNS packets. 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.</p>
DSCP	<p>Specifies the DSCP value.</p>
Packet length	<p>Specifies the packet length. The range is 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 [-]).</p>
Packet Loss Priority (PLP)	<p>Specifies the packet loss priority. 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.</p>
Protocol	<p>Specifies Internet protocol number. The range is 0 through 255.</p>
ICMP Message	<p>For Protocol IPv4 when you enter a Protocol value as 1, the ICMP Message field displays where you can select an ICMP message to apply to the data policy. Likewise, the ICMP Message field displays for Protocol IPv6 when you enter a Protocol value as 58.</p> <p>When you select Protocol as Both, the ICMP Message or ICMPv6 Message field displays.</p> <p>Note This field is available from Cisco IOS XE Release 17.4.1, Cisco vManage Release 20.4.1.</p>
Source Data Prefix	<p>Specifies the group of source prefixes or an individual source prefix.</p>
Source Port	<p>Specifies the source port number. The range is 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 [-]).</p>
TCP Flag	<p>Specifies the TCP flag, syn.</p>

Match Condition	Description
Traffic To	In a Multi-Region Fabric architecture, match border router traffic flowing to the access region that the border router is serving, the core region, or a service VPN. Note Minimum release: Cisco vManage Release 20.8.1



Note If IPv4 packet contains non-initial fragment of UDP or TCP datagram, it has no L4 ports information available because there is no UDP or TCP header. For such fragments destination-port or source-port match is ignored.

In the following example, all the UDP packets to destination port 161 and any other IPv4 packets having protocol ID field in IPv4 header set to 17 with IPv4 header having fragment-offset set will be dropped.

```

policy
app-visibility
access-list SDWAN_101
sequence 100
match
destination-port 161
protocol          17
!
action drop
!
!
```

Table 11: ICMP Message Types/Codes and Corresponding Enumeration Values

Type	Code	Enumeration
0	0	echo-reply

3		unreachable
	0	net-unreachable
	1	host-unreachable
	2	protocol-unreachable
	3	port-unreachable
	4	packet-too-big
	5	source-route-failed
	6	network-unknown
	7	host-unknown
	8	host-isolated
	9	dod-net-prohibited
	10	dod-host-prohibited
	11	net-tos-unreachable
	12	host-tos-unreachable
	13	administratively-prohibited
	14	host-precedence-unreachable
15	precedence-unreachable	
5		redirect
	0	net-redirect
	1	host-redirect
	2	net-tos-redirect
	3	host-tos-redirect
8	0	echo
9	0	router-advertisement
10	0	router-solicitation
	11	time-exceeded
	0	ttl-exceeded
11	1	reassembly-timeout
	12	parameter-problem
	0	general-parameter-problem
12	1	option-missing
	2	no-room-for-option
	13	0

14	0	timestamp-reply
40	0	photuris
42	0	extended-echo
43		extended-echo-reply
	0	echo-reply-no-error
	1	malformed-query
	2	interface-error
	3	table-entry-error
	4	multiple-interface-match

Table 12: ICMPv6 Message Types/Codes and Corresponding Enumeration Values

Type	Code	Enumeration
1		unreachable
	0	no-route
	1	no-admin
	2	beyond-scope
	3	destination-unreachable
	4	port-unreachable
	5	source-policy
	6	reject-route
	7	source-route-header
2	0	packet-too-big
3		time-exceeded
	0	hop-limit
	1	reassembly-timeout
4		parameter-problem
	0	Header
	1	next-header
	2	parameter-option
128	0	echo-request
129	0	echo-reply
130	0	mld-query
131	0	mld-report

132	0	mld-reduction
133	0	router-solicitation
134	0	router-advertisement
135	0	nd-ns
136	0	nd-na
137	0	redirect
138		router-renumbering
	0	renum-command
	1	renum-result
	255	renum-seq-number
139		ni-query
	0	ni-query-v6-address
	1	ni-query-name
	2	ni-query-v4-address
140		ni-response
	0	ni-response-success
	1	ni-response-refuse
	2	ni-response-qtype-unknown
141	0	ind-solicitation
142	0	ind-advertisement
143		mldv2-report
144	0	dhaad-request
145	0	dhaad-reply
146	0	mpd-solicitation
147	0	mpd-advertisement
148	0	cp-solicitation
149	0	cp-advertisement
151	0	mr-advertisement
152	0	mr-solicitation
153	0	mr-termination
155	0	rpl-control

Action Parameters - Control Policy

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

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:

Table 13:

Description	Cisco SD-WAN Manager
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 .
Discard the packet.	Click Reject .

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

Action Condition	Description
Export To	Export the route to the specified VPN or list of VPNs (for a match route match condition only). The range is 0 through 65535 or list name.
OMP Tag	Change the tag string in the route, prefix, or TLOC. The range is 0 through 4294967295.
Preference	Change the preference value in the route, prefix, or TLOC to the specified value. A higher preference value is more preferred. The range is 0 through 255.
Service	Specify a service to redirect traffic to before delivering the traffic to its destination. The TLOC address or list of TLOCs identifies the TLOCs to which the traffic should be redirected to reach the service. In the case of multiple TLOCs, the traffic is load-balanced among them. The VPN identifier is where the service is located. Standard services: FW, IDS, IDP Custom services: netsvc1, netsvc2, netsvc3, netsvc4 Configure the services themselves on the Cisco IOS XE Catalyst SD-WAN devices that are collocated with the service devices, using the vpn service configuration command.
TLOC	Change the TLOC address, color, and encapsulation to the specified address and color. 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 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 highest preference value is used. If two or more of TLOCs have the highest preference value, traffic is sent among them in an ECMP fashion.

Action Condition	Description
TLOC Action	<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.</p> <p>Setting the TLOC action option enables the Cisco Catalyst SD-WAN Controller to perform end-to-end tracking of the path to the ultimate destination device.</p>



Note The **preference** command controls the preference for directing inbound and outbound traffic to a tunnel. The preference can be a value from 0 through 4294967295 (232 – 1), and the default value is 0. A higher value is preferred over a lower value.

When a Cisco vEdge device has two or more tunnels, if all the TLOCs have the same preference and no policy is applied that affects traffic flow, all the TLOCs are advertised into OMP. When the router transmits or receives traffic, it distributes traffic flows evenly among the tunnels, using ECMP.

Action Parameters - Data Policy

Table 14: Feature History

Feature Name	Release Information	Description
Path Preference Support for Cisco IOS XE Catalyst SD-WAN Devices	Cisco IOS XE Catalyst SD-WAN Release 17.2.1r	This feature extends to Cisco IOS XE Catalyst SD-WAN devices, support for selecting one or more local transport locators (TLOCs) for a policy action.
Traffic Redirection to SIG Using Data Policy	Cisco IOS XE Release 17.4.1 Cisco vManage Release 20.4.1	With this feature, while creating a data policy, you can define an application list along with other match criteria and redirect the application traffic to a Secure Internet Gateway (SIG).
Next Hop Action Enhancement in Data Policies	Cisco IOS XE Catalyst SD-WAN Release 17.5.1a Cisco vManage Release 20.5.1	This feature enhances match action conditions in a centralized data policy for parity with the features configured on Cisco IOS XE Catalyst SD-WAN devices. When you are setting up next-hop-loose action, this feature helps to redirect application traffic to an available route when next-hop address is not available.
Traffic Redirection to SIG Using Data Policy: Fallback to Routing	Cisco IOS XE Catalyst SD-WAN Release 17.8.1a Cisco vManage Release 20.8.1	With this feature, you can configure internet-bound traffic to be routed through the Cisco Catalyst SD-WAN overlay, as a fallback mechanism, when all SIG tunnels are down.

Feature Name	Release Information	Description
Log Action for both Localized and Centralized Data Policies	Cisco IOS XE Catalyst SD-WAN Release 17.11.1a Cisco vManage Release 20.11.1	This feature enables you to set a log action parameter for data policy, application route policy, and localized policy while configuring data policies on Cisco IOS XE Catalyst SD-WAN devices. The log parameter allows packets to get logged and generate syslog messages. Logs are exported to an external syslog server every five minutes when a flow is active. You can control policy logs as per the configured rate using the command policy log-rate-limit .
Remote Preferred Color in Data Policy	Cisco IOS XE Catalyst SD-WAN Release 17.15.1a Cisco Catalyst SD-WAN Manager Release 20.15.1	You can set a remote preferred color in the data policy to control traffic routing based on the SLA criteria. See Configure Traffic Rules, on page 153 for more information.

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

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:

Action Condition	Description
Click Accept	Accepts the packet. An accepted packet is eligible to be modified by the additional parameters configured in the action portion of the policy configuration.
Cflowd	Enables cflowd traffic monitoring.
Counter	Counts the accepted or dropped packets. Specifies the name of a counter. Use the show policy access-lists counters command on the Cisco IOS XE Catalyst SD-WAN device.
Click Drop	Discards the packet. This is the default action.
Log	<p>Minimum release: Cisco IOS XE Catalyst SD-WAN Release 17.11.1a and Cisco vManage Release 20.11.1</p> <p>Click Log to enable logging.</p> <p>When (DP, AAR or ACL) data policy packets are configured with log action, logs generated and logged to syslog. Due to the global log-rate-limit, not all logs are logged. 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.</p> <p>For information on policy log-rate-limit CLI, see policy log-rate-limit command in the Cisco Catalyst SD-WAN Qualified Command Reference Guide.</p>

Action Condition	Description
Redirect DNS	<p>Redirects DNS requests to a particular DNS server. Redirecting requests is optional, but if you do so, you must specify both actions.</p> <p>For an inbound policy, redirect-dns host allows the DNS response to be correctly forwarded back to the requesting service VPN.</p> <p>For an outbound policy, specify the IP address of the DNS server.</p> <p>Note When you upgrade to releases later than Cisco IOS XE Catalyst SD-WAN Release 17.7.1a, you must configure redirect DNS through nat use-vpn 0 to redirect DNS to Direct Internet Interface (DIA).</p> <p>Note You can set only local TLOC preferences with redirect-dns as actions on the same sequence, but not remote TLOC.</p> <p>Note You cannot configure Redirect DNS and SIG at the same time. NAT DIA fallback and DNS redirection are not supported at the same time in data policy.</p>
TCP Optimization	Fine-tune TCP to decrease round-trip latency and improve throughput for matching TCP traffic.
Secure Internet Gateway	<p>Redirect application traffic to a SIG.</p> <p>Note Before you apply a data policy for redirecting application traffic to a SIG, you must have configured the SIG tunnels.</p> <p>For more information on configuring Automatic SIG tunnels, see Automatic Tunnels. For more information on configuring Manual SIG tunnels, see Manual Tunnels.</p> <p>Check the Fallback to Routing check box to route internet-bound traffic through the Cisco SD-WAN overlay when all SIG tunnels are down. This option is introduced in Cisco IOS XE Catalyst SD-WAN Release 17.8.1a and Cisco vManage Release 20.8.1.</p>



Note On Cisco IOS XE Catalyst SD-WAN devices, 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:

Action Condition	Description
Cflowd	Enables cflowd traffic monitoring.
NAT Pool or NAT VPN	Enables NAT functionality, so that traffic can be redirected directly to the internet or other external destination. You can configure up to 31 (1–31) NAT pools per router.

Action Condition	Description
DSCP	DSCP value. The range is 0 through 63.
Forwarding Class	Name of the forwarding class.
Local TLOC	<p>Enables sending packets to one of the TLOCs that matches the color and encapsulation. The available colors are: 3g, biz-internet, blue, bronze, custom1,custom2, custom3, default, gold, green, lte, metro-ethernet, mpls, private1 through private6, public-internet, red and silver.</p> <p>The encapsulation options are: ipsec and gre.</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> <p>By default, encapsulation is ipsec.</p>
Next Hop	<p>Sets the next hop IP address to which the packet should be forwarded.</p> <p>Note Starting from Cisco IOS XE Catalyst SD-WAN Release 17.5.1a and Cisco vManage Release 20.5.1, the Use Default Route when Next Hop is not available field is available next to the Next Hop action parameter. This option is available only when the sequence type is Traffic Engineering or Custom, and the protocol is either IPv4 or IPv6, but not both.</p>
Policer	Applies a policer. Specifies the name of policer configured with the policy policer command.
Service	<p>Specifies 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>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> <p>Configure the services themselves on the Cisco IOS XE Catalyst SD-WAN devices that are collocated with the service devices, using the vpn service command.</p>
TLOC	Direct traffic to a remote TLOC that 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 VPN .	Set the VPN that the packet is part of. The range is 0 through 65530.



Note Data policies are applicable on locally generated packets, including routing protocol packets, when the match conditions are generic.

Example configuration:

```
sequence 21
  match
    source-ip 10.0.0.0/8
  action accept
```

In such situations, it may be necessary to add a sequence in the data policy to escape the routing protocol packets. For example to skip OSPF, use the following configuration:

```
sequence 20
  match
    source-ip 10.0.0.0/8
    protocol 89
  action accept
sequence 21
  match
    source-ip 10.0.0.0/8
  action accept
```

The following table describes the IPv4 and IPv6 actions.

Table 15:

IPv4 Actions	IPv6 Actions
drop, dscp, next-hop (from-service only)/vpn, count, forwarding class, policer (only in interface ACL), App-route SLA (only)	N/A
App-route preferred color, app-route sla strict, cflowd, nat, redirect-dns	N/A
N/A	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

Apply Policies to Sites and VPNs

In the **Apply Policies to Sites and VPNs** page, 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. 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 **New Site List**, **Inbound Site List**, **Outbound Site List**, or **VPN List**. 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 **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 **New Site List and VPN List**. Choose the direction for applying the policy (**From Service**, **From Tunnel**, or **All**), choose one or more site lists, and choose one or more VPN lists. Click **Add**.
 - d. For a cflowd policy block, click **New Site List**. Choose one or more site lists, and click **Add**.
4. Click **Preview** to view the configured policy. The policy appears in CLI format.
5. Click **Save Policy**. The **Configuration > Policies** page appears, and the policies table includes the newly created policy.

NAT Fallback on Cisco IOS XE Catalyst SD-WAN Devices

	Release Information	
NAT Fallback on Cisco IOS XE Catalyst SD-WAN Devices	Cisco IOS XE Release 17.3.2 Cisco vManage Release 20.3.2	Cisco IOS XE Catalyst SD-WAN devices support the NAT fallback feature for Direct Internet Access (DIA). The NAT fallback feature provides a routing-based mechanism for all traffic that is sent to the DIA route to use an alternative route when required. With this release, fallback is supported on the service and tunnel side.



Note To use Cisco SD-WAN Manager to configure NAT DIA fallback, Cisco SD-WAN Manager must manage your Cisco Catalyst SD-WAN Controller.

To enable NAT fallback using Cisco SD-WAN Manager, create and configure a data policy by doing the following:

1. From the Cisco SD-WAN Manager menu, choose **Configuration > Policies**.

2. From the **Custom** options drop-down, under **Centralized Policy**, select **Traffic Policy**.
3. Click **Traffic Data**.
4. From the **Add Policy** drop-down, click **Create New**.
5. Click **Sequence Type** and select **Custom**.
6. Click (+) **Sequence Rule** to create a new sequence rule.
7. After adding match conditions, click **Actions** and click **Accept**.
8. Click **NAT VPN** and select the **Fallback** checkbox.
9. Click **Save and Match Actions**.
10. Click **Save Data Policy**.

Edit your existing centralized policy and import the policy:

1. Click **Centralized Policy** and for the required centralized policy, click ... and select **Edit**.
2. Click **Traffic Rules** and select **Traffic Data**.
3. From the **Add Policy** drop-down, select **Import Existing**.
4. Select the NAT policy that you created from the **Policy** drop-down.
5. Click **Policy Application** and select **Traffic Data**.
6. Click + **New Site List and VPN List**.
7. Select the direction, VPN, and site as required.
8. Click **Add**.
9. Click **Save Policy Changes**.
10. Click to select **VPN**, and **Site** from the drop-down.



Note Policy configured for the **from-tunnel** traffic is also applied to the return DIA (Underlay) traffic apart from the return traffic coming over the tunnel. If none of the sequences in that policy match, it matches the default sequence in that policy.



Note NAT DIA fallback and DNS redirection are not supported at the same time in data policy.
The following NAT fallback actions/commands are now supported:

- Action: `nat fallback`
- When applying a policy: `direction from-tunnel`

Activate a Centralized Policy

Activating a centralized policy sends that policy to all connected Cisco SD-WAN Controllers. To activate a centralized policy:

1. From the Cisco SD-WAN Manager menu, choose **Configuration > Policies**. **Centralized Policy** is selected and displayed by default.
2. For the required policy, click ... and select **Activate**. The **Activate Policy** popup appears. It lists the IP addresses of the reachable Cisco SD-WAN Controllers to which the policy must be applied.
3. Click **Activate**.

View Centralized Policies

To view centralized policies:

1. From the **Centralized Policy**, select a policy.
2. For a policy created using the UI policy builder or using the CLI, click ... and select **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 Cisco SD-WAN Manager policy configuration wizard, click ... and select **Preview**. This policy is displayed in text format.

Copy, Edit, and Delete Policies

To copy a policy:

1. From the **Centralized Policy**, select a policy.
2. For the desired policy, click ... and select **Copy**.
3. In the Policy Copy popup window, enter the policy name and a description of the policy.



Note Starting with the Cisco IOS XE Release 17.2, 127 characters are supported for policy names for the following policy types:

- Central route policy
- Local route policy
- Local Access Control IOST (ACL)
- Local IPv6 ACL
- Central data policy
- Central app route policy
- QoS map
- Rewrite rule

All other policy names support 32 characters.

4. Click **Copy**.

To edit policies created using the Cisco SD-WAN Manager policy configuration wizard:

1. For the desired policy, click ... and select **Edit**.
2. Edit the policy as needed.
3. Click **Save Policy Changes**.

To edit policies created using the CLI method:

1. In the **Custom Options** drop-down, click **CLI Policy**.
2. For the desired policy, click ... and select **Edit**.
3. Edit the policy as needed.
4. Click **Update**.

To delete policies:

1. From the **Centralized Policy**, select a policy.
2. For the desired policy, click ... and select **Delete**.
3. Click **OK** to confirm deletion of the policy.

Configure Centralized Policies Using the CLI

To configure a centralized control policy using the CLI:

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

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

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

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

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

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

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

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

3. Create a control policy instance:

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

4. 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).

5. Define match parameters for routes and for TLOCs:

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

6. 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
```

7. Create additional numbered sequences of match–action pairs within the control policy, as needed.
8. 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
```

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

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

10. If the action you are configuring is a service, configure the required services on the Cisco IOS XE Catalyst SD-WAN devices so that the Cisco Catalyst SD-WAN 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 Catalyst SD-WAN device keeps track of the services, advertising them to the Cisco Catalyst SD-WAN 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 Catalyst SD-WAN device withdraws the service advertisement.

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)# match parameters
```

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
```

- 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)
```

Centralized Policies Configuration Examples

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

General Centralized Policy Example

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

Here we configure a simple data policy on the Cisco Catalyst SD-WAN 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 Catalyst SD-WAN Controller, it pushes the policy configuration to the Cisco IOS XE Catalyst 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 198.51.100.1/32
  protocol tcp
  action accept
  set tloc 203.0.113.1 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
```

Traffic Engineering

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

One common way to design a domain in a Cisco IOS XE Catalyst 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 Catalyst 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 Catalyst SD-WAN 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 Catalyst SD-WAN 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 Catalyst SD-WAN 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 Catalyst SD-WAN 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 Catalyst SD-WAN Controller as an outbound policy, as indicated by the **out** option in the `apply-policy site` command. This option means the Cisco Catalyst SD-WAN 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 Catalyst SD-WAN 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 Catalyst SD-WAN Controller that would send West-side data traffic through the Cisco hub, and West and East-side traffic through the Device Hub East:

```

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

```

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

```

Here is an explanation of this policy configuration:

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

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

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

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

Define two control policies:

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

Apply the control policies:

- The first line in the **apply-policy** configuration has the Cisco Catalyst SD-WAN 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 Catalyst SD-WAN 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.

Community Example

This example displays the configuration for centralized control policy for community lists.

```

policy
  lists
    expanded-community-list test
      community 0:110* 100:[7-9]+
      community 0:110* 11:*

    community-list test-com
      community 0:1
      community 0:2

  control-policy test
    sequence 10
    match route
      expanded-community-list test

```



```

action accept
set
  community 100:2 100:3
additive

```

This example displays the configuration for standard community lists.

```

Standard Community list

route : 0:1234 0:11 0:12

community-list
  community 0:100
  community 0:1234
  community 0:101
*MATCH*

route : 0:1234 0:11 0:12
community-list
  community 0:100
  community 0:5678
  community 0:101
*NO MATCH*

```

This example displays the configuration for expanded community lists. OR match compares each regex string in the community list against the route's community string.

```

Expanded Community list
route - 0:1234 0:5678
expanded-community-list:
  community 0:110* 11:
  community 0:110* 100:[7-9]+
  community 0:12[3-7]+
*MATCH*

route - 0:1234 0:5678
expanded-community-list:
  community 0:111*
  community 0:110* 11:*
*NO MATCH*

```

EXACT match input strings need to have communities in sorted order. Sorts it by byte value and add the meta characters for start and end of string.

```

route - 0:1234 0:5678
expanded-community-list:
community ^0:1234 0:5678$
*MATCH*

```

AND match input strings need to have communities in sorted order. Add '.' to blindly match between the sorted communities.

```

route - 0:0 0:1234 0:5678 0:9789 0:9800 0:9900 0:9999 1:10
expanded-community-list:
  community 0:1234 .+ 0:9900 .+
*MATCH*

```

SIG Data Policy Fallback

From Cisco IOS XE Catalyst SD-WAN Release 17.8.1a and Cisco vManage Release 20.8.1, you can use the **sig-action fallback-to-routing** command to configure internet-bound traffic to be routed through the Cisco

Catalyst SD-WAN overlay when all SIG tunnels are down. The following example shows the configuration of this fallback mechanism.

```
data-policy _VPN10_SIG_Fall_Back
  vpn-list VPN10
    sequence 1
      match
        app-list Google_Apps
        source-ip 0.0.0.0/0
      !
      action accept
        sig
        sig-action fallback-to-routing
      !
    !
  default-action drop
```

Ranking Color Preference Example

```
policy lists
  preferred-color-group GROUP1_COLORS
    primary-preference
      color-preference biz-internet
      path-preference direct-tunnel
    !
    secondary-preference
      color-preference mpls
      path-preference multi-hop-path
    !
    tertiary-preference
      color-preference lte
    !
  !
  preferred-color-group GROUP2_COLORS
    primary-preference
      color-preference mpls
    !
    secondary-preference
      color-preference biz-internet
    !
  !
  preferred-color-group GROUP3_COLORS
    primary-preference
      color-preference mpls biz-internet lte
    !
```

Data Policy for IPv6 Applications Example

```
policy
  data-policy _VPN1_Data-Policy-For-Ipv6-Traffic
    vpn-list VPN1
      sequence 1
        match
          app-list Msft-0365
          source-ipv6 0::0/0
        !
        action accept
      !
    !
  default-action drop
  !
lists
```

```
app-list Msft-0365
  app ms-office-web-apps
!
site-list SITE-100
  site-id 100
!
vpn-list VPN1
  vpn 1
!
!
!
apply-policy
  site-list SITE-100
  data-policy _VPN1_Data-Policy-For-Ipv6-Traffic all
!
!
```




CHAPTER 5

Localized Policy

The topics in this section provide overview information about the different types of localized policies, the components of localized policies, and how to configure localized policies using Cisco SD-WAN Manager or the CLI.

- [Overview of Localized Policies, on page 81](#)
- [Configure Localized Policy Using Cisco SD-WAN Manager , on page 83](#)
- [Configure Localized Policy for IPv4 Using the CLI, on page 97](#)
- [Configure Localized Policy for IPv6 Using the CLI, on page 99](#)
- [Localized Data Policy Configuration Examples, on page 100](#)
- [QoS For Router Generated Cisco SD-WAN Manager Traffic, on page 101](#)
- [Information About QoS For Router-Generated Cisco SD-WAN Manager Traffic, on page 101](#)
- [Restrictions For QoS For Router Generated Cisco SD-WAN Manager Traffic, on page 102](#)
- [Configure QoS for Router Generated Cisco SD-WAN Manager Traffic Using a CLI Template, on page 102](#)
- [Verify QoS for Router Generated Cisco SD-WAN Manager Traffic Using CLI, on page 103](#)
- [Troubleshooting QoS For Router Generated Cisco SD-WAN Manager Traffic, on page 104](#)

Overview of Localized Policies

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

Types of Localized Policies

Localized Control Policy

Control policy operates on the control plane traffic in the Cisco IOS XE Catalyst 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 Catalyst 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 Catalyst 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 Catalyst 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. 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.

Localized Data Policy

Data policy operates on the data plane in the Cisco IOS XE Catalyst SD-WAN overlay network and affects how data traffic is sent among the Cisco IOS XE Catalyst SD-WAN devices in the network. The Cisco Catalyst 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 Catalyst SD-WAN device.

Localized data policy, so called because it is provisioned on the local Cisco IOS XE Catalyst 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 and provision packet mirroring.

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 Catalyst 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 Catalyst 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.

Mirror Data Packets

Once packets are classified, you can configure access lists to send a copy of data packets seen on a Cisco vEdge device to a specified destination on another network device. The Cisco IOS XE Catalyst SD-WAN devices support 1:1 mirroring; that is, a copy of every packet is sent to the alternate destination.

Configure Localized Policy Using Cisco SD-WAN Manager

To configure localized policies, use the Cisco SD-WAN Manager policy configuration wizard. The wizard is a UI policy builder that consists of five windows 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 **Next** at the bottom of the window. To return to a component, click **Back** at the bottom of the window.

To configure localized policies using Cisco SD-WAN Manager, use the steps identified in the procedures that follow this section.

Start the Policy Configuration Wizard

To start the policy configuration wizard:

1. From the Cisco SD-WAN Manager menu, choose **Configuration > Policies**.
2. Select **Localized Policy**.
3. Click **Add Policy**.

The **Create Groups of Interest** page is displayed.

Configure Groups of Interest for Localized Policy

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

In Create Groups of Interest, create new groups of list types as described in the following sections to use in a localized policy:

Configure As Path

1. In the group of interest list, 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**.

AS Path list specifies 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 separated by commas. To configure multiple AS paths in a single list, include multiple **as-path** options, specifying one AS path in each option.

Configure Community

A community list is used to create groups of communities to use in a match clause of a route map. A community list can be used to control which routes are accepted, preferred, distributed, or advertised. You can also use a community list to set, append, or modify the communities of a route.

1. In the group of interest list, click **Community**.
2. Click **New Community List**.
3. Enter a name for the community list.
4. In the **Add Community** field, enter one or more data prefixes separated by commas in any of the following formats:
 - **aa:nn**: Autonomous System (AS) 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 number.
 - **no-advertise**: Attaches the NO_ADVERTISE community to routes. Routes in this community are not advertised to other BGP peers.
 - **no-export**: Attaches 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.
5. Click **Add**.

Configure Data Prefix

1. In the **Group of Interest** list, 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**.

A data prefix list specifies 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.

Configure Extended Community

1. In the group of interest list, click **Extended Community**.
2. Click **New Extended Community List**.
3. Enter a name for the list.
4. Enter the BGP extended community in the following formats:
 - **rt** (*aa:nn | ip-address*): Route target community, which is one or more routers that can receive a set of routes carried by BGP. Specify this as the AS number and network number, where each number is a 2-byte value with a range from 1 to 65535, or as an IP address.
 - **soo** (*aa:nn | ip-address*): Route origin community, which is one or more routers that can inject a set of routes into BGP. Specify this as the AS 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.
5. Click **Add**.

Configure Class Map

1. In the group of interest list, click **Class Map**.
2. Click **New Class List**.
3. Enter a name for the class.
4. Select a required queue from the **Queue** drop-down list.
5. Click **Save**.

Configure Mirror

1. In the group of interest list, 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 for 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 **Add**.

To configure mirroring parameters, define the remote destination to which to mirror the packets, and define the source of the packets. Mirroring applies to unicast traffic only. It does not apply to multicast traffic.

Configure Policer

1. In the group of interest list, click **Policer**.
2. Click **New Policer List**.
3. Enter a name for the list.
4. In the **Burst (bps)** 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 (bps)** field, enter the maximum traffic rate. It can be value from 8 through 2^{64} bps (8 through 100000000000).
7. Click **Add**.

Configure Prefix

1. In the group of interest list, click **Prefix**.
2. Click **New Prefix List**.
3. Enter a name for the list.
4. In the **Internet Protocol** field, click either **IPv4** or **IPv6**.
5. 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.
6. Click **Add**.

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

Configure Forwarding Classes/QoS

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

QoS Map

To create a new QoS mapping:

1. In **QoS**, click the **Add QoS Map** drop-down.
2. Select **Create New**.
3. Enter a name and description for the QoS mapping.
4. Click **Add Queue**. The **Add Queue** popup appears.
5. Select the queue number from the **Queue** drop-down.
6. Select the maximum bandwidth and buffer percentages, and the scheduling and drop types.
7. Enter the **Forwarding Class**.

8. Click Save Queue.

To import an existing QoS mapping:

1. In **QoS**, click the **Add QoS Map** drop-down.
2. Select **Import Existing**. The **Import Existing Application QoS Map Policy** popup displays.
3. Select a **QoS Map** policy.
4. Click **Import**.

To view or copy a QoS mapping or to remove the mapping from the localized policy, click ... and select the desired action.

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

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

Policy Rewrite

To configure policy rewrite rules for the QoS mapping:

1. In **Policy Rewrite**, click the **Add Rewrite Policy** drop-down.
2. Select **Create New**.
3. Enter a name and description for the rewrite rule.
4. Click **Add Rewrite Rule**. The **Add Rule** popup appears.
5. Select a class from the **Class** drop-down.
6. Select the priority (**Low** or **High**) from the Priority drop-down.
Low priority is supported only for Cisco IOS XE Catalyst SD-WAN devices.
7. Enter the DSCP value (0 through 63) in the **DSCP** field.
8. Enter the class of service (CoS) value (0 through 7) in the **Layer 2 Class of Service** field.
9. Click **Save Rule**.

To import an existing rewrite rule:

1. In **QoS**, click the **Add Rewrite Policy** drop-down..
2. Select **Import Existing**. The **Import Existing Policy Rewrite** popup appears.

3. Select a rewrite rule policy.
4. Click **Import**.

Click **Next** to move to **Configure Access Lists** page.

Configure ACLs

1. In the **Configure Access Control Lists** page, configure ACLs.
2. To create a new ACL, click the **Add Access Control List Policy** drop-down. Select one from the following options:
 - **Add IPv4 ACL Policy**: Configure IPv4 ACL policy.
 - **Add IPv6 ACL Policy**: Configure IPv6 ACL policy.
 - **Import Existing**: Import existing ACL policy.
3. If you click **Add IPv4 ACL Policy**, the **Add IPv4 ACL Policy** page appears.
or
If you click **Add IPv6 ACL Policy**, the **Add IPv6 ACL Policy** page appears.
4. Enter a name and description for the ACL in the **ACL Policy** page.
5. In the left pane, click **Add ACL Sequence**. An **Access Control List** box is displayed in the left pane.
6. Double-click the **Access Control List** box, and type a name for the ACL.
7. In the right pane, click **Add Sequence Rule** to create a single sequence in the ACL. **Match** is selected by default.
8. Click a match condition.
9. On the left, enter the values for the match condition.
 - a. 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 ... next to the rule's name and select the desired option.

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. To change the default action:

1. Click **Default Action** in the left pane.
2. Click the **Pencil** icon.
3. Change the default action to **Accept**.
4. Click **Save Match and Actions**.
5. Click **Save Access Control List Policy**.

To configure **Device Access Policy**, see [Device Access Policy](#).

Click **Next** to move to Configure Route Policy page.

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 Catalyst 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 a Cisco IOS XE ACL. A Cisco IOS XE ACL does not interact with a Cisco Catalyst SD-WAN explicit and an implicit ACL and cannot override an implicit ACL or explicit ACL. Cisco 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 service-name**) 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 16:

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 Route Policies

In **Configure Route Policies**, configure the routing policies:

1. In **Add Route Policy**, 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. **Match** is selected by default.
6. Select a desired protocol from the **Protocol** drop-down list. The options are: IPv4, IPv6, or both.
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 route policy.
11. To rearrange match–action pairs in the route policy, in the right pane drag them to the desired position.
12. To remove a match–action pair from the route policy, 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 route policy, in the left pane drag the rules to the desired position.
15. To copy, delete, or rename the route policy sequence rule, in the left pane, click ... next to the rule's name and select the desired option.
16. 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**.
17. Click **Save Route Policy**.
 18. Click **Next** to move to **Policy Overview** page.

Match Parameters

Access List Parameters

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

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.

Match class in ACL is not supported. You can use rewrite policy to configure DSCP values.

For access lists, you can match these parameters:

Match Condition	Description
Class	Name of a class defined with a policy class-map command.
Destination Data Prefix	Name of a data-prefix-list list.
Destination Port	Specifies 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 [-]). The range is 0 through 65535.
DSCP	Specifies the DSCP value. The range is 0 through 63.
Protocol	Specifies the internet protocol number. The range is 0 through 255.
ICMP Message	<p>When you select a Protocol value as 1 the ICMP Message field displays where you can select an ICMP message to apply to the data policy.</p> <p>When you select a Next Header value as 58 the ICMP Message field displays where you can select an ICMP message to apply to the data policy.</p> <p>Note This field is available from Cisco IOS XE Release 17.4.1, Cisco vManage Release 20.4.1.</p> <p>For icmp-msg and icmp6-msg message types, refer to the ICMP Message Types/Codes and Corresponding Enumeration Values table in the Centralized chapter.</p>
Packet Length	Specifies the length of the packet. The range 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 [-]).
Source Data Prefix	Specifies the name of a data-prefix-list list.
PLP	Specifies the Packet Loss Priority (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.

Match Condition	Description
Source Port	Specifies 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 [-]). The range is 0 through 65535.
TCP	syn

Route Policy Parameters

For route policies, you can match these parameters:

Match Condition	Description
Address	Specifies the name of a Prefix-List list.
AS Path List	Specifies one or more BGP AS path lists. You can write each AS as a single number or as a regular expression. To specify more than one AS number in a single path, include the list in quotation marks (" "). To configure multiple AS numbers in a single list, include multiple AS Path options, specifying one AS path in each option.
Community List	List of one or more BGP communities. In Community List , you can specify: <ul style="list-style-type: none"> • aa:nn: AS 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.
Extended Community List	Specifies the list of one or more BGP extended communities. In community , you can specify: <ul style="list-style-type: none"> • rt (aa:nn ip-address): Route target community, which is one or more routers that can receive a set of routes carried by BGP. Specify this as the AS number and network number, where each number is a 2-byte value with a range from 1 to 65535, or as an IP address. • soo (aa:nn ip-address): Route origin community, which is one or more routers that can inject a set of routes into BGP. Specify this as the AS 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.
BGP Local Preference	Specifies the BGP local preference number. The range is 0 through 4294967295.
Metric	Specifies the route metric value. The range is 0 through 4294967295.

Match Condition	Description
Next Hop	Specifies the name of an IP prefix list.
OMP Tag	Specifies the OMP tag number. The range is 0 through 4294967295.
Origin	Specifies the BGP origin code. The options are: EGP (default), IGP, Incomplete.
OSPF Tag	Specifies the OSPF tag number. The range is 0 through 4294967295.
Peer	Specifies the peer IP address.

Action Parameters

Access List 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 CLI, you configure the action 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:

Action Condition	Description
Accept	Accepts the packet. An accepted packet is eligible to be modified by the additional parameters configured in the action portion of the access list.
Counter	Name of a counter. To display counter information, use the show policy access-lists counters command on the Cisco IOS XE Catalyst SD-WAN device.
Drop	Discards the packet. This is the default action.

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

Description	Value or Range
Class	Specifies the name of a QoS class. It can also be defined with a policy class-map command.
Mirror List	Specifies the name of mirror . It is defined with a policy mirror command.
Policer	Specifies the name of a policer defined with a policy policer command.
DSCP	Specifies the packet's DSCP value. The range is 0 through 63.
Next Hop	Specifies the IPv4 address. It sets the next hop IP address to which the packet should be forwarded. Note Starting from Cisco vManage Release 20.5.1 and Cisco IOS XE Catalyst SD-WAN Release 17.5.1a, Use Default Route when Next Hop is not available field is available next to Next Hop action parameter.

Route Policy Parameters

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:

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

Description	Value or Range
Aggregator	Set the AS number in which a BGP route aggregator is located and the IP address of the route aggregator. The range is 1 through 65535.
As Path	Sets an AS number or a series of AS numbers to exclude from the AS path or to prepend to the AS path. The range is 1 through 65535.
Atomic Aggregate	Sets the BGP atomic aggregate attribute.
Community	Sets the BGP community value. Starting from Cisco IOS XE Catalyst SD-WAN Release 17.5.1a, the Community Additive option field is available. Additive option appends the communities to the existing communities of the route.
Local Preference	Sets the BGP local preference. The range is 0 through 4294967295.
Metric	Sets the metric value. The range is 0 through 4294967295.
Metric Type	Sets the metric type. The options are type1 or type2.
Next Hop	Sets the IPv4 address. It sets the next hop IP address to which the packet should be forwarded. Note Starting from Cisco vManage Release 20.5.1 and Cisco IOS XE Catalyst SD-WAN Release 17.5.1a, Use Default Route when Next Hop is not available field is available next to Next Hop action parameter.
OMP Tag	Sets the OMP tag for OSPF to use. The range is 0 through 4294967295.
Origin	Sets the BGP origin code. The options are: EGP (default), IGP, Incomplete.
Originator	Sets the IP address from which the route was learned.
OSPF Tag	Sets the OSPF tag value. The range is 0 through 4294967295.
Weight	Sets the BGP weight. The range is 0 through 4294967295.

Configure Policy Settings

In **Policy Overview**, configure the policy settings:

1. In the **Enter name and description for your localized master policy** pane, enter name and description for the policy.

2. In the **Policy Settings** pane, select the policy application checkboxes that you want to configure. The options are:
 - **Netflow**: Perform traffic flow monitoring on IPv4 traffic.
 - **Netflow IPv6**: Perform traffic flow monitoring on IPv6 traffic.
 - **Application**: Track and monitor IPv4 applications.
 - **Application IPv6**: Track and monitor IPv6 applications.
 - **Cloud QoS**: Enable QoS scheduling.
 - **Cloud QoS Service Side**: Enable QoS scheduling on the service side.
 - **Implicit ACL Logging**: Log the headers of all the packets that are dropped because they do not match a service perform traffic flow monitoring.
3. 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.
4. Click **Preview** to view the full policy in CLI format.
5. Click **Save Policy**.

Apply Localized Policy in a Device Template

1. From the Cisco SD-WAN Manager menu, choose **Configuration > Templates**.
2. If you are creating a new device template:
 - a. Click **Device Templates**.



Note In Cisco vManage Release 20.7.1 and earlier releases, **Device Template** is titled as **Device**.

- 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 Catalyst 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. Click **Device Templates**, and for the desired template, click ... and select **Edit**.



Note In Cisco vManage Release 20.7.1 and earlier releases, **Device Template** is titled as **Device**.

- b. Click **Additional Templates**. 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 **Additional Templates** 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).

Activate a Localized Policy

1. Click **Localized Policy**, and select a policy.
2. For the desired policy, click ... and select **Activate**.
3. In the **Activate Policy** popup, click **Activate** to push the policy to all reachable Cisco SD-WAN Controllers in the network.
4. Click **OK** to confirm activation of the policy on all Cisco SD-WAN Controllers.
5. To deactivate the localized policy, select =, and then select a policy.
6. For the desired policy, click ... and select **Deactivate**.
7. In the **Deactivate Policy** popup, click **Deactivate** to confirm that you want to remove the policy from all reachable Cisco SD-WAN Controllers.

View Localized Policies

To view localized policies:

1. Click **Localized Policy**, and select a policy.
2. For a policy created using the UI policy builder or using the CLI, click ... and select **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 Cisco SD-WAN Manager policy configuration wizard, click ... and select **Preview**. This policy is displayed in text format.

Copy, Edit, and Delete Policies

To copy a policy:

1. Click **Localized Policy**, and select a policy.
2. For the desired policy, click ... and select **Copy**.
3. In the Policy Copy popup window, enter the policy name and a description of the policy.



Note Starting with the Cisco IOS XE Release 17.2, 127 characters are supported for policy names for the following policy types:

- Central route policy
- Local route policy
- Local Access Control List (ACL)
- Local IPv6 ACL
- Central data policy
- Central app route policy
- QoS map
- Rewrite rule

All other policy names support 32 characters.

4. Click **Copy**.

To edit policies created using the Cisco SD-WAN Manager policy configuration wizard:

1. For the desired policy, click ... and select **Edit**.
2. Edit the policy as needed.
3. Click **Save Policy Changes**.

To edit policies created using the CLI method:

1. From the **Custom Options** drop-down, under Localized Policy, select **CLI Policy**.
2. For the desired policy, click ... and select **Edit**.
3. Edit the policy as needed.
4. Click **Update**.

To delete policies:

1. Click **Localized Policy**, and select a policy.
2. For the desired policy, click ... and select **Delete**.
3. Click **OK** to confirm deletion of the policy.

Configure Localized Policy for IPv4 Using the CLI

Following are the high-level steps for configuring an access list using the CLI on Cisco IOS XE Catalyst 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
```

- 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**.

- 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

- 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]
```

- 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
```

- 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
```

- 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
```

- Associate a **service-policy** to a Cisco IOS XE Catalyst 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
```

- 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
```

- 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)#
```

- 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
```

Configure Localized Policy for IPv6 Using the CLI

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

- 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
```

- Create an access list instance:

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

- 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 Catalyst 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.

ICMP Message Example

This example displays the configuration for localized data policy for ICMP messages.

```
policy
access-list acl_1
sequence 100
match
protocol 1
icmp-msg administratively-prohibited
!
action accept
count administratively-prohibited
!
!
```


QoS For Router Generated Cisco SD-WAN Manager Traffic

Table 17: Feature History

Feature Name	Release Information	Description
QoS for Router Generated Cisco SD-WAN Manager Traffic	Cisco IOS XE Catalyst SD-WAN Release 17.11.1a	This feature helps you to prioritize or queue router-generated Cisco SD-WAN Manager traffic based on your specific requirements. Use QoS policies and class maps to route Cisco SD-WAN Manager traffic through a queue of your choice.

Information About QoS For Router-Generated Cisco SD-WAN Manager Traffic

Quality of Service (QoS) is a technique used to manage and prioritize network traffic to ensure that certain types of traffic are given priority over others. QoS is particularly important for router-generated Cisco SD-WAN Manager traffic, which is used for managing and monitoring network devices. For more information see, [Forwarding and QoS](#).

You can prioritize or queue router-generated traffic based on your specific requirements. The prioritization can be achieved through the use of QoS policies and class maps.

Use the following steps to put router-generated traffic into the queue of your choice:

1. Define a class map using a CLI template: Identifies the type of traffic you want to prioritize. In this case, you create a class map to identify the router-generated traffic to queue.
2. Define a policy map using a CLI template: Defines the actions that you want to take on the traffic identified in the class map. Create a policy map that assigns a priority or places the router-generated traffic into a specific queue.

Benefits of QoS For Router Generated Cisco SD-WAN Manager Traffic

- Improved network performance: By prioritizing critical router-generated traffic over less important traffic, ensure that your network management functions operate smoothly and monitor and control network devices effectively.
- Better user experience: Queuing router-generated traffic helps preventing congestion on the network and ensure that user-generated traffic does not negatively impact network management functions. The queuing can result in a better user experience.
- Increased network availability: Reduces the risk of network downtime caused by network management issues. This improves network availability and reduce the impact of any network issues on your business operations.

- Simplified network management: Simplifies network management and reduces the need for manual intervention. The simplification can save time and reduce the risk of human error.
- Efficient use of network resources: QoS policies and class maps allow you to allocate network resources efficiently, ensuring that critical router-generated traffic flow efficiently, minimizing the impact on other network traffic.

Restrictions For QoS For Router Generated Cisco SD-WAN Manager Traffic

- The QoS for router generated Cisco SD-WAN Manager traffic feature is supported only on Cisco IOS XE Catalyst SD-WAN devices.
- Configuring QoS for router generated Cisco SD-WAN Manager traffic is possible only using a CLI template.
- With this feature, you can prioritize, using a queue, only for the traffic that devices generate for Cisco SD-WAN Manager. Other data and management plane traffic continue to take Queue 0 by default.

Configure QoS for Router Generated Cisco SD-WAN Manager Traffic Using a CLI Template

For more information about using CLI templates, see [CLI Add-On Feature Templates](#) and [CLI Templates](#).



Note By default, CLI templates execute commands in global config mode.

Define a Class Map and Map to a Queue Number

1. Using a localized policy, define a class-map and map the class-map to a queue number :


```
policy class-map class Queue_1 queue 2
```
2. Commit the changes.

Here's the complete configuration example for defining a class map and mapping it to a queue number:

```
config-t
policy class-map class Queue_1 queue 2
!
```

Enable QoS For Router Generated Cisco SD-WAN Manager Traffic

This section provides example CLI configurations to enable QoS for router generated Cisco SD-WAN Manager traffic:

1. Enter config-policy mode:


```
policy
```

2. Use a forwarding class and use the class map that you mapped to a queue that you want to prioritize:
vmanage-forwarding-class *queue_name*
3. Commit the changes.

QoS for router generated Cisco SD-WAN Manager traffic is enabled.

Here's the complete configuration example for enabling QoS for router generated Cisco SD-WAN Manager traffic:

```
config-t
policy
vmanage-forwarding-class Queue_1
!
```

Verify QoS for Router Generated Cisco SD-WAN Manager Traffic Using CLI

The following is sample output from the **show policy-map interface** command using the **GigabitEthernet 1** keyword:

```
Device# show policy-map interface GigabitEthernet 1

Service-policy output: shape_GigabitEthernet1

Class-map: class-default (match-any)
  8619 packets, 5056404 bytes
  5 minute offered rate 113000 bps, drop rate 0000 bps
Match: any
Queueing
queue limit 64 packets
(queue depth/total drops/no-buffer drops) 0/0/0
(pkts output/bytes output) 8619/5056404
shape (average) cir 4200000, bc 16800, be 16800
target shape rate 4200000

Service-policy : qosmap

queue stats for all priority classes:
Queueing
priority level 1
queue limit 512 packets
(queue depth/total drops/no-buffer drops) 0/0/0
(pkts output/bytes output) 565/95064

Class-map: Queue0 (match-any)
  565 packets, 95064 bytes
  5 minute offered rate 4000 bps, drop rate 0000 bps
Match: qos-group 0
police:
  rate 30 %
  rate 1260000 bps, burst 39375 bytes
  conformed 565 packets, 95064 bytes; actions:
  transmit
  exceeded 0 packets, 0 bytes; actions:
  drop
  conformed 4000 bps, exceeded 0000 bps
Priority: Strict, b/w exceed drops: 0
```

```

Priority Level: 1

Class-map: Queue_1 (match-any)
 8050 packets, 4961100 bytes ----->
 5 minute offered rate 111000 bps, drop rate 0000 bps
 Match: qos-group 1
 Queueing
  queue limit 64 packets
  (queue depth/total drops/no-buffer drops) 0/0/0
  (pkts output/bytes output) 8050/4961100
 bandwidth remaining ratio 10

Class-map: Queue_2 (match-any)
 4 packets, 240 bytes
 5 minute offered rate 0000 bps, drop rate 0000 bps
 Match: qos-group 2
 Queueing
  queue limit 64 packets
  (queue depth/total drops/no-buffer drops) 0/0/0
  (pkts output/bytes output) 4/240
 bandwidth remaining ratio 10

```

In this example, **Class-map** for the respective queues displays the number, size, and the rate of packet transfer from the router to the destination. You can see a change in the Queue_1 and keep track of the packet transfer.

Troubleshooting QoS For Router Generated Cisco SD-WAN Manager Traffic

Problem

Unable to commit changes using the CLI

Possible Causes

There could be typos or incorrect queue names entered while committing the changes. For example, if you type queuee 2 instead of queue 2, the following error is displayed: Aborted: illegal reference 'policy vmanage-traffic-forwarding-class'

Solution

Enter the right queue name that you want the Cisco SD-WAN Manager traffic from the router to flow through.



CHAPTER 6

Redirect DNS in a Service-Side VPN

Table 18: Feature History

Feature Name	Release Information	Description
Redirect DNS in a Service-Side VPN	Cisco IOS XE Catalyst SD-WAN Release 17.8.1a Cisco vManage Release 20.8.1	This feature allows you to configure a Cisco IOS XE Catalyst SD-WAN device to respond to Domain Name System (DNS) queries using proxy servers. This feature adds support for DNS proxy for service-side VPN hosts and DNS redirects inside the service VPNs.

- [Information About Redirect DNS in a Service-Side VPN, on page 105](#)
- [Restrictions for Redirect DNS in a Service-Side VPN, on page 106](#)
- [Use Cases for Redirect DNS in a Service-Side VPN, on page 106](#)
- [Configure Redirect DNS in a Service-Side VPN, on page 108](#)
- [Configure Redirect DNS in a Service-Side VPN Using the CLI, on page 110](#)
- [Verify Redirect DNS in a Service-Side VPN, on page 112](#)
- [Configuration Examples for Redirect DNS, on page 112](#)

Information About Redirect DNS in a Service-Side VPN

The Redirect DNS feature enables Cisco IOS XE Catalyst SD-WAN devices to respond to DNS queries using a specific configuration and associated host table cache that are selected based on certain characteristics of the queries. In a redirect DNS environment, multiple DNS databases can be configured on the device. The Cisco Catalyst SD-WAN software can be configured to choose one of the DNS name server configurations whenever the device responds to a DNS query, by forwarding or resolving the query. Prior to Cisco IOS XE Catalyst SD-WAN Release 17.8.1a, redirect DNS is supported only through NAT Direct Internet Access (DIA) path.

When an application-aware routing policy allows a Cisco IOS XE Catalyst SD-WAN device to send application traffic to a service VPN and receive application traffic from a service VPN, the device performs a DNS lookup to determine the path to reach the application server. If the router does not have a connection to the internet, it sends DNS queries to an edge device that has such a connection, and that device determines how to reach a server for that application.



Note In a network in which the device that is connected to the internet is in a geographically distant data center, the resolved DNS address points to a server that is also geographically distant from the site where the service VPN is located.

Because you can configure a Cisco IOS XE Catalyst SD-WAN device to be an internet exit point, it is possible for any router to reach the internet directly to perform DNS lookups.

You can configure redirect DNS with either a centralized data policy or, if you want to apply SLA criteria to the data traffic, you can use application-aware routing policy.


Restrictions for Redirect DNS in a Service-Side VPN

- A redirect DNS request is not accepted without NAT configuration if the request is from the same VPN with the same port from a different host.
- If you configure DNS server IP address using NAT, it cannot be changed through the data policy.
- DNS fragmented packets and self-generated DNS are not supported.
- DNS requests from the overlay tunnel are not supported.
- Redirect DNS is supported only on IPv4 traffic, and not on IPv6 traffic.
- DNS requests through User Datagram Protocol (UDP) are supported. However, requests from Transmission Control Protocol (TCP) are not supported.

Use Cases for Redirect DNS in a Service-Side VPN

Unconditional Redirect DNS

In unconditional redirect DNS (scenario A), a host sends all the DNS requests to a local edge router, and the local edge router redirects the DNS request to an enterprise DNS server in the data center (which is available only using a service-side VPN) and acts as a DNS forwarder. A use case for this feature redirects statically configured IP addresses for printers to an enterprise DNS server in a data center. In this use case, all the legacy printers are statically configured with an IP address of a local router as DNS server, which acts as DNS forwarder to forward all the DNS requests from printers.


Figure 12: Unconditional Redirect DNS**Conditional Redirect DNS**

In conditional redirect DNS (scenario B), a host uses a service provider (SP) or managed service provider (MSP) DNS by default. For known applications that use an Cisco Catalyst SD-WAN Application Intelligence Engine (SAIE) or custom applications, for example, *.google.com, the DNS request is forwarded to the enterprise DNS server using a Cisco Catalyst SD-WAN overlay network. All the other DNS requests are sent to the SP or MSP DNS server.



Note In Cisco vManage Release 20.7.1 and earlier releases, SAIE is called deep packet inspection (DPI).

Figure 13: Conditional Redirect DNS



Configure Redirect DNS in a Service-Side VPN

1. From the Cisco SD-WAN Manager menu, choose **Configuration > Policies**.
2. From the **Custom Options** drop-down list, choose **Traffic Policy** from the **Centralized Policy** menu.
3. Click **Traffic Data** to create a traffic data policy.
4. From the **Add Policy** drop-down list, choose **Create New**.
5. In the **Name** and **Description**, enter a name and a description for the data policy.
6. Click **Sequence Type**.
The **Add Data Policy** dialog box is displayed.
7. Choose the type of data policy that you want to create—**Application Firewall**, **QoS**, **Service Chaining**, **Traffic Engineering**, or **Custom**.
A policy sequence containing the selected type of data policy is added in the left pane.
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.
9. Click **Sequence Rule**. The **Match/Action** dialog box is displayed, where **Match** is selected by default. The available policy match conditions are listed in the menu.
10. From the **Protocol** drop-down list, choose **IPv4** to apply the policy only to IPv4 address families.
11. To choose one or more **Match** conditions, click the fields and set the values as described.



Note Not all match conditions are available for all policy sequence types.

12. To select the actions to take on matching data traffic, click the **Actions** menu.
13. To drop matching traffic, click **Drop**.
The available policy actions are listed on the right side.
14. To accept matching traffic, click **Accept**.
The available policy actions are listed on the right side.
15. In the **Actions** menu, choose **Redirect DNS** to configure redirect DNS.
16. In the **Redirect DNS** condition field, enter the **IP Address** and click **Save Match and Actions**.
17. Click **Save Data Policy**.

Match Condition	Procedure
None (match all the packets)	Do not specify any match conditions.
Applications / Application Family List / Custom Applications	<ol style="list-style-type: none"> 1. In the Match conditions menu, click Applications/Application Family List. 2. From the drop-down list, choose the application family. 3. To create an application list: <ol style="list-style-type: none"> a. Click New Application List. b. Enter a name for the list. c. Click Application to create a list of individual applications. Click Application Family to create a list of related applications. d. From the Select Application drop-down list, choose the corresponding applications or application families. e. Click Save.
DNS Application List	Add an application list to enable split DNS: <ol style="list-style-type: none"> 1. In the Match conditions menu, click DNS Application List. 2. From the drop-down list, choose the application family.
DNS	Add an application list to process split DNS: <ol style="list-style-type: none"> 1. In the Match conditions menu, click DNS. 2. From the drop-down list, choose Request to process DNS requests for the DNS applications.

Match Condition	Procedure
Destination Data Prefix	<ol style="list-style-type: none"> 1. In the Match conditions menu, click Destination Data Prefix. 2. To match a list of destination prefixes, from the Data Prefix drop-down list, choose a list. 3. To match an individual destination prefix, enter the prefix in the Destination: IP Prefix field.
Destination Port	<ol style="list-style-type: none"> 1. In the Match conditions menu, click Destination Port. 2. 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 two numbers separated with a hyphen [-]).
DSCP	<ol style="list-style-type: none"> 1. In the Match conditions menu, click DSCP. 2. In the DSCP field, enter the DSCP value—a number from 0 through 63.
Packet Length	<ol style="list-style-type: none"> 1. In the Match conditions menu, click Packet Length. 2. In the Packet Length field, enter the length—a value from 0 through 65535.
PLP	<ol style="list-style-type: none"> 1. In the Match conditions menu, click PLP to set the Packet Loss Priority. 2. From the PLP drop-down list, choose Low or High.
Protocol	<ol style="list-style-type: none"> 1. In the Match conditions menu, click Protocol. 2. In the Protocol field, enter the Internet Protocol number—a number from 0 through 255.
Source Data Prefix	<ol style="list-style-type: none"> 1. In the Match conditions menu, click Source Data Prefix. 2. To match a list of source prefixes, from the Source Data Prefix List drop-down list, choose a data prefix list. 3. To match an individual source prefix, enter the prefix in the Source field.
Source Port	<ol style="list-style-type: none"> 1. In the Match conditions menu, click Source Port. 2. 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 [-]).

Configure Redirect DNS in a Service-Side VPN Using the CLI

The following steps show the minimum policy components required to enable redirect DNS with a centralized data policy:

1. Create a list of overlay network sites to which the centralized control policy is to be applied:

```
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 an en dash (–). Create additional site lists, as needed.

2. Create lists of applications or application families for which you want to enable redirect DNS. Refer to these lists in the **match** section of the data policy.

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

3. Create list VPNs to which the redirect DNS policy is to be applied:

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

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

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

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

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

6. Process the DNS server resolution for the applications or application families contained in an application list. For the *list-name* argument, specify the list name.

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

7. Configure the match–action pair sequence to process DNS requests (for outbound data traffic) or responses (for inbound data traffic):

```
vSmart(config-sequence)# match dns (request | response)
```

8. By default, the DNS servers configured in the VPN in which the policy is applied are used to process DNS lookups for the applications. You can direct the DNS requests to a particular DNS server. For a data policy condition that applies to outbound traffic (from the service network), configure the IP address of the DNS server:

```
vSmart(config-sequence)# action accept redirect-dns ip-address
```

For a data policy condition that applies to inbound traffic (from the tunnel), include the following action so that the DNS response can be correctly forwarded back to the service VPN:

```
vSmart(config-sequence)# action accept redirect-dns host
```

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

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

Verify Redirect DNS in a Service-Side VPN

The following is a sample output from the `show sdwan policy from-vsmart` command that shows how to verify the redirect DNS configuration:

```
vSmart# show sdwan policy from-vsmart
from-vsmart data-policy vpn1_dns-redirect-prefer-lte
direction from-service
vpn-list vpn1
sequence 1
match
source-ip 10.0.0.0/0
dns request
action accept
count          gdns2_-396115821
redirect-dns 10.255.255.254
default-action accept
from-vsmart lists vpn-list vpn1
vpn 1
```

Configuration Examples for Redirect DNS

Unconditional DNS Redirect

The following example shows how to configure an unconditional DNS redirect, where all the DNS requests are matched:

```
policy
data-policy rdns
vpn-list vpn10
sequence 10
match
source-ip 0.0.0.0/0
dns request
!
action
redirect-dns 209.165.200.225
!
default-action accept
!
!
!
apply-policy
site-list siteA
data-policy rdns from-service
```

Conditional DNS Redirect

The following example shows how to configure a conditional DNS redirect, where a selective DNS request is defined using an app list:

```
policy
data-policy rdns
vpn-list vpn10
sequence 10
```

```
match
  source-ip 10.0.0.0/8
  dns      request
  dns-app-list YouTube
  !
action
  redirect-dns 209.165.200.225
  !
default-action accept
!
!
!
!
apply-policy
site-list siteA
data-policy rdns from-service
```




CHAPTER 7

Default AAR and QoS Policies

Table 19: Feature History

Feature Name	Release Information	Description
Configure Default AAR and QoS Policies	Cisco IOS XE Catalyst SD-WAN Release 17.7.1a Cisco vManage Release 20.7.1	This feature enables you to efficiently configure default application-aware routing (AAR), data, and quality of service (QoS) policies for Cisco IOS XE Catalyst SD-WAN devices. The feature provides a step-by-step workflow for categorizing the business relevance, path preference, and other parameters for network applications, and applying those preferences as traffic policy.

- [Information About Default AAR and QoS Policies, on page 115](#)
- [Prerequisites for Default AAR and QoS Policies, on page 116](#)
- [Restrictions for Default AAR and QoS Policies, on page 117](#)
- [Supported Devices for Default AAR and QoS Policies, on page 117](#)
- [Use Cases for Default AAR and QoS Policies, on page 117](#)
- [Configure Default AAR and QoS Policies Using Cisco SD-WAN Manager, on page 117](#)
- [Monitor Default AAR and QoS Policies, on page 122](#)

Information About Default AAR and QoS Policies

It is often helpful to create an AAR policy, a data policy, and a QoS policy for devices in a network. These policies route and prioritize traffic for best performance. When creating these policies, it is helpful to distinguish among the applications producing network traffic, based on the likely business relevance of the applications, and to give higher priority to business-relevant applications.

Cisco SD-WAN Manager provides an efficient workflow to help you create a default set of AAR, data, and QoS policies to apply to devices in the network. The workflow presents a set of more than 1000 applications that can be identified by network-based application recognition (NBAR), an application recognition technology built into Cisco IOS XE Catalyst SD-WAN devices. The workflow groups the applications into one of three business-relevance categories:

- Business-relevant: Likely to be important for business operations, for example, Webex software.
- Business-irrelevant: Unlikely to be important for business operations, for example, gaming software.

- Default: No determination of relevance to business operations.

Within each of the business-relevance categories, the workflow groups the applications into application lists, such as broadcast video, multimedia conferencing, VoIP telephony, and so on.

Using the workflow, you can accept the predefined categorization of each application's business relevance or you can customize the categorization of specific applications by moving them from one of the business-relevance categories to another. For example, if, by default, the workflow predefines a specific application as business-irrelevant, but that application is important for your business operations, then you can recategorize the application as Business-relevant.

The workflow provides a step-by-step procedure for configuring the business relevance, path preference, and service level agreement (SLA) category.

After you complete the workflow, Cisco SD-WAN Manager produces a default set of the following:

- AAR policy
- QoS policy
- Data policy

After you attach these policies to a centralized policy, you can apply these default policies to Cisco IOS XE Catalyst SD-WAN devices in the network.

Background Information About NBAR

NBAR is an application recognition technology included in Cisco IOS XE Catalyst SD-WAN devices. NBAR uses a set of application definitions called protocols to identify and categorize traffic. One of the categories that it assigns to traffic is the business-relevance attribute. The values of this attribute are Business-relevant, Business-irrelevant, and Default. In developing protocols to identify applications, Cisco estimates whether an application is likely to be important for typical business operations, and assigns a business-relevance value to the application. The default AAR and QoS policy feature uses the business-relevance categorization provided by NBAR.

Benefits of Default AAR and QoS Policies

- Manage and customize bandwidth allocations.
- Prioritize applications based on their relevance to your business.

Prerequisites for Default AAR and QoS Policies

- Knowledge about the relevant applications.
- Familiarity with the SLAs and QoS markings to prioritize traffic.

Restrictions for Default AAR and QoS Policies

- When you customize a business-relevant application group, you cannot move all the applications from that group to another section. Application groups of business-relevant section need to have at least one application in them.
- Default AAR and QoS policies do not support IPv6 addressing.

Supported Devices for Default AAR and QoS Policies

- Cisco 1000 Series Integrated Services Routers (ISR1100-4G and ISR1100-6G)
- Cisco 4000 Series Integrated Services Routers (ISR44xx)
- Cisco Catalyst 8000V Edge Software
- Cisco Catalyst 8300 Series Edge Platforms
- Cisco Catalyst 8500 Series Edge Platforms
- Cisco C1100 Series Integrated Services Router

Use Cases for Default AAR and QoS Policies

If you are setting up a Cisco Catalyst SD-WAN network and want to apply an AAR and a QoS policy to all the devices in a network, use this feature to create and deploy these policies quickly.

Configure Default AAR and QoS Policies Using Cisco SD-WAN Manager

Follow these steps to configure default AAR, data, and QoS policies using Cisco SD-WAN Manager:

1. From the Cisco SD-WAN Manager menu, choose **Configuration > Policies**.
2. Click **Add Default AAR & QoS**.
The **Process Overview** page is displayed.
3. Click **Next**.
The **Recommended Settings based on your selection** page is displayed.
4. Based on the requirements of your network, move the applications between the **Business Relevant**, **Default**, and **Business Irrelevant** groups.



Note When customizing the categorization of applications as Business-relevant, Business-irrelevant, or Default, you can only move individual applications from one category to another. You cannot move an entire group from one category to another.

5. Click **Next**.

On the **Path Preferences (optional)** page, choose the **Preferred** and **Preferred Backup** transports for each traffic class.

6. Click **Next**.

The **App Route Policy Service Level Agreement (SLA) Class** page is displayed.

This page shows the default settings for **Loss**, **Latency**, and **Jitter** values for each traffic class. If necessary, customize **Loss**, **Latency**, and **Jitter** values for each traffic class.

7. Click **Next**.

The **Enterprise to Service Provider Class Mapping** page is displayed.

a. Select a service provider class option, based on how you want to customize bandwidth for different queues. For further details on QoS queues, refer to the section **Mapping of Application Lists to Queues**

b. If necessary, customize the bandwidth percentage values for each queues.

8. Click **Next**.

The **Define prefixes for the default policies and applications lists** page is displayed.

For each policy, enter a prefix name and description.

9. Click **Next**.

The **Summary** page is displayed. On this page, you can view the details for each configuration.

You can click **Edit** to edit the options that appeared earlier in the workflow. Clicking edit returns you to the relevant page.

10. Click **Configure**.

Cisco SD-WAN Manager creates the AAR, data, and QoS policies and indicates when the process is complete.

The following table describes the workflow steps or actions and their respective effects:

Table 20: Workflow Steps and Effects

Workflow Step	Affects the Following
Recommended Settings based on your selection	AAR and data policies
Path Preferences (optional)	AAR policies

Workflow Step	Affects the Following
App Route Policy Service Level Agreement (SLA) Class: <ul style="list-style-type: none"> • Loss • Latency • Jitter 	AAR policies
Enterprise to Service Provider Class Mapping	Data and QoS policies
Define prefixes for the default policies and applications	AAR, data, QoS policies, forwarding classes, application lists, SLA class lists

11. To view the policy, click **View Your Created Policy**.



Note To apply the default AAR and QoS policies to the devices in the network, create a centralized policy that attaches the AAR and data policies to the required site lists. To apply the QoS policy to the Cisco IOS XE Catalyst SD-WAN devices, attach it to a localized policy through device templates.

Mapping of Application Lists to Queues

The following lists show each service provider class option, the queues in each option, and the application lists included in each queue. The application lists are named here as they appear on the Path Preferences page in this workflow.

4 QoS class

- Voice
 - Internetwork control
 - VoIP telephony
- Mission critical
 - Broadcast video
 - Multimedia conferencing
 - Real-Time interactive
 - Multimedia streaming
- Business data
 - Signaling
 - Transactional data
 - Network management
 - Bulk data

- Default
 - Best effort
 - Scavenger

5 QoS class

- Voice
 - Internetwork control
 - VoIP telephony
- Mission critical
 - Broadcast video
 - Multimedia conferencing
 - Real-Time interactive
 - Multimedia streaming
- Business data
 - Signaling
 - Transactional data
 - Network management
 - Bulk data
- General data
 - Scavenger
- Default
 - Best effort

6 QoS class

- Voice
 - Internetwork control
 - VoIP telephony
- Video
 - Broadcast video
 - Multimedia conferencing
 - Real-Time interactive
- Mission Critical

- Multimedia streaming
- Business data
 - Signaling
 - Transactional data
 - Network management
 - Bulk data
- General data
 - Scavenger
- Default
 - Best effort

8 QoS class

- Voice
 - VoIP telephony
- Net-ctrl-mgmt
 - Internetwork control
- Interactive video
 - Multimedia conferencing
 - Real-Time interactive
- Streaming video
 - Broadcast video
 - Multimedia streaming
- Call signaling
 - Signaling
- Critical data
 - Transactional data
 - Network management
 - Bulk data
- Scavengers
 - Scavenger

- Default
 - Best effort

Monitor Default AAR and QoS Policies

Monitor Default AAR Policies

1. From the Cisco SD-WAN Manager menu, choose **Configuration > Policies**.
2. Click **Custom Options**.
3. Choose **Traffic Policy** from **Centralized Policy**.
4. Click **Application Aware Routing**.
A list of AAR policies is displayed.
5. Click **Traffic Data**.
A list of traffic data policies is displayed.

Monitor QoS Policies

1. From the Cisco SD-WAN Manager menu, choose **Configuration > Policies**.
2. Click **Custom Options**.
3. Choose **Forwarding Class/QoS** from **Localized Policy**.
4. Click **QoS Map**.
A list of QoS policies is displayed.



Note To verify QoS polices, refer to [Verify QoS Policy](#).



CHAPTER 8

Device Access Policy

Table 21: Feature History

Feature Name	Release Information	Description
Device Access Policy on SNMP and SSH	Cisco IOS XE Catalyst SD-WAN Release 17.2.1r	This feature defines 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. The control plane of a Cisco IOS XE Catalyst SD-WAN device processes the data traffic for local services (like SSH and SNMP) from a set of sources. Routing packets are required to form the overlay.

- [Device Access Policy Overview, on page 123](#)
- [Configure Device Access Policy Using Cisco SD-WAN Manager, on page 124](#)
- [Configure Device Access Policy Using the CLI, on page 125](#)
- [Examples for ACL Statistics and Counters, on page 126](#)
- [Verifying ACL Policy on an SNMP Server, on page 127](#)
- [Verifying ACL Policy on SSH, on page 129](#)

Device Access Policy Overview

Starting from Cisco IOS XE SD-WAN Release 17.2.1r, the Cisco SD-WAN Manager user interface is enhanced to configure device access policy on all the Cisco IOS XE Catalyst SD-WAN devices.

The control plane of Cisco IOS XE Catalyst SD-WAN devices process the data traffic for local services like, SSH and SNMP, from a set of sources. It is important to protect the CPU from device access traffic by applying the filter to avoid malicious 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 analyzed 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 Cisco SD-WAN Manager

Cisco IOS XE Catalyst SD-WAN devices support device access policy configuration to handle SNMP and SSH traffic directed towards the control plane. Use Cisco SD-WAN Manager to configure destination ports based on the device access policy.



Note In order to allow connections to devices from **Tools > SSH Terminal** in Cisco SD-WAN Manager, 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 SD-WAN Manager 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 a device access policy:

1. From the Cisco SD-WAN Manager menu, choose **Configuration > Policies**.
2. Click **Localized Policy** and from the **Custom Options** drop-down, under **Localized Policy**, select **Access Control Lists**.
3. From the **Add Device Access Policy** drop-down list, select **Add IPv4 Device Access Policy** or **Add IPv6 Device Access Policy** option to add a device.



Note Beginning with Cisco IOS XE Catalyst SD-WAN Release 17.9.1a and Cisco vManage Release 20.9.1, if you configure an IPv4 or an IPv6 device access policy with no policy sequences and only a default action of **Accept** or **Drop**, the device access policy creates an SSH and an SNMP configuration. You can now create a device access policy with only a default action and with no policy sequences to create a device configuration or a Cisco SD-WAN Manager configuration for both SSH and SNMP.

If you do not create an SNMP server configuration, the SNMP configuration created by the device access policy is unused.

Prior to Cisco IOS XE Catalyst SD-WAN Release 17.9.1a and Cisco vManage Release 20.9.1, if you configured a device access policy with only a default action of **Accept** or **Drop** and with no policy sequences, the device access policy would not create a device configuration or a Cisco SD-WAN Manager configuration.

4. Select **Add IPv4 Device Access Policy** from the drop-down list to add an **IPv4 ACL Policy**. The edit **Device IPv4 ACL Policy** page appears.
5. Enter the name and the description for the new policy.
6. Click **Add ACL Sequence** to add a sequence. The **Device Access Control List** page is displayed.
7. Click **Sequence Rule**. **Match** and **Actions** options are displayed.
8. Click **Match**, 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.
VPN	Enter the VPN ID. The range is 0-65536.

9. Click **Actions**, 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.

10. Click **Save Match And Actions** to save all the conditions for the ACL policy.
11. Click **Save Device Access Control List Policy** to apply the selected match conditions to an action.
12. If no packets match, then 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 Catalyst SD-WAN devices. When you try to configure IPv6 prefix matches on these devices, Cisco SD-WAN Manager fails to generate device configuration.

Configure Device Access Policy Using the CLI

Configuration:

```
ip access-list standard snmp-acl
 1 permit 10.0.1.12 255.255.255.0
 11 deny any
!

snmp-server community private view v2 ro snmp-acl

ip access-list extended ssh-acl
 1 permit tcp host 10.0.1.12 any eq 22
 11 deny tcp any any eq 22
!
```

```

line vty 0 4
 access-class ssh-acl in vrf-also
 !

```



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

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# 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 display ACL statistics and counters using the CLI, use the command:

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

Example statistics output using the CLI:

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

```
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 ACL Policy on an SNMP Server

Starting from the Cisco IOS XE Catalyst SD-WAN Release 17.2.1r release, Cisco IOS XE Catalyst SD-WAN devices support the device-access-policy feature on SNMP servers. In case of SNMP, Cisco SD-WAN Manager validates to block the template push on the device if the SNMP feature template is not configured.



Note In case of SNMP, the destination data prefix list is not applicable for Cisco IOS XE Catalyst SD-WAN devices. If you apply the localized policy with SNMP configuration for a device, then the destination data prefix will be ignored.

Configuration:

```
snmp-server community private view v2 ro snmp-acl
```

Yang model for the command **snmp-server community**. Following is the ACL settings sample from the yang model:

```
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;
  }
}
}

```

Following is the sample test log for snmp-server ACL settings:

```
Device# sh sdwan ver
16.12.1
```

```
Device# config-t
```

```
admin connected from 127.0.0.1 using console on the device
Device(config)# snmp-server community TEST_1 RO 80
Device(config)# end
Uncommitted changes found, commit them? [yes/no/CANCEL] yes
Commit complete.
```

```
Device#
*Mar 13 21:17:19.377: %SYS-5-CONFIG_P: Configured programmatically by process
session_id_for_dmi_vty_100001 from console as NETCONF on vty31266
*Mar 13 21:17:19.377: %DMI-5-CONFIG_I: RO/0: nedd: Configured from NETCONF/RESTCONF by
admin, transaction-id 518
```

```
Device#
Device# sh sdwan run | i snmp
snmp-server community TEST_1 RO 80
```

```
Device# sh sdwan run | i snmp
snmp-server community TEST_1 RO 80
Device#
```

```
admin connected from 127.0.0.1 using console on the device
Device(config)# snmp-server community TEST_V6 ipv6 acl-name-1
Device(config)# end
Uncommitted changes found, commit them? [yes/no/CANCEL] yes
Commit complete.
Device#
```

```
*Mar 13 21:18:10.040: %SYS-5-CONFIG_P: Configured programmatically by process
session_id_for_dmi_vty_100001 from console as NETCONF on vty31266
*Mar 13 21:18:10.041: %DMI-5-CONFIG_I: RO/0: nedd: Configured from NETCONF/RESTCONF by
```

```

admin, transaction-id 535

Device#
Device# sh sdwan run | i snmp
snmp-server community TEST_1 RO 80
snmp-server community TEST_V6 ipv6 acl-name-1
Device#
Device# sh run | i snmp
snmp-server community TEST_1 RO 80
snmp-server community TEST_V6 RO ipv6 acl-name-1
Device#

```

Verifying ACL Policy on SSH

Starting from the Cisco IOS XE Catalyst SD-WAN Release 17.2.1r release, the Cisco IOS XE Catalyst SD-WAN devices support device-access-policy features on SSH servers using Virtual Teletype (VTY) lines. Cisco SD-WAN Manager uses all the available VTY lines in the backend and pushes the policy accordingly.

Configuration:

```

line vty 0 4
  access-class ssh-acl in vrf-also
!
```

Following is the ACL settings sample from the yang model:

```

// 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;
  }
}
}

```

Following is the sample test log for line-server ACL settings:

```
Device# config-transaction

admin connected from 127.0.0.1 using console on Device
Device(config)# line vty 0 4
Device(config-line)# access-class acl_1 in vrf-also
Device(config-line)# transport input ssh
Device(config-line)# end
Uncommitted changes found, commit them? [yes/no/CANCEL] yes
Commit complete.
Device#
*May 24 20:51:02.994: %SYS-5-CONFIG_P: Configured programmatically by process
iosp_vty_100001_dmi_nesd from console as NETCONF on vty31266
*May 24 20:51:02.995: %DMI-5-CONFIG_I: R0/0: nesd: Configured from NETCONF/RESTCONF by
admin, transaction-id 227
Device#
Device#
Device# sh sdwan run | sec vty
Error: Licensing infrastructure is NOT initialized.
Error: Licensing infrastructure is NOT initialized.
line vty 0 4
  access-class acl_1 in vrf-also
  login local
  transport input ssh
line vty 5 80
  login local
  transport input ssh
Device#
Device# sh run | sec vty
Error: Licensing infrastructure is NOT initialized.
Error: Licensing infrastructure is NOT initialized.
line vty 0 4
  access-class acl_1 in vrf-also
  exec-timeout 0 0
  password 7 11051807
  login local
  transport preferred none
  transport input ssh
line vty 5 80
  login local
  transport input ssh
```



CHAPTER 9

Cisco Catalyst SD-WAN Application Intelligence Engine Flow

The topics in this section provide overview information about the Cisco Catalyst SD-WAN Application Intelligence Engine (SAIE) flow, and how to configure the flow using Cisco SD-WAN Manager or the CLI.

- [Cisco Catalyst SD-WAN Application Intelligence Engine Flow Overview, on page 131](#)
- [Configure Cisco Catalyst SD-WAN Application Intelligence Engine Flow Using Cisco SD-WAN Manager, on page 132](#)
- [Configure SD-WAN Application Intelligence Engine Flow Using the CLI, on page 136](#)

Cisco Catalyst SD-WAN Application Intelligence Engine Flow Overview

The Cisco Catalyst SD-WAN Application Intelligence Engine (SAIE) flow provides the ability to look into the packet past the basic header information. The SAIE flow determines the contents of a particular packet, and then either records that information for statistical purposes or performs an action on the packet.



Note In Cisco vManage Release 20.7.1 and earlier releases, the SAIE flow is called the deep packet inspection (DPI) flow.

Benefits include increased visibility into the network traffic, which enables network operators to understand usage patterns and to correlate network performance information along with providing usage base billing or even acceptable usage monitoring. The SAIE flow can also reduce the overall costs on the network.

You can configure the SAIE flow using a centralized data policy. You define the applications of interest in a Cisco SD-WAN Manager policy list or with the **policy lists app-list** CLI command, and you call these lists in a **policy data-policy** command. 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.

The following list of protocols are not supported in SAIE flow:

- Open Shortest Path First (OSPF)
- Border Gateway Protocol (BGP)

- Internet Control Message Protocol (ICMP)
- Bidirectional Forwarding Detection (BFD)

Configure Cisco Catalyst SD-WAN Application Intelligence Engine Flow Using Cisco SD-WAN Manager

To configure the Cisco Catalyst SD-WAN Application Intelligence Engine (SAIE) flow, use the Cisco SD-WAN Manager policy configuration wizard. The wizard consists of the following 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. For configuration details, see [Configure Groups of Interest](#).
- Configure Traffic Rules—Create the match and action conditions of a policy. For configuration details, see [Configure Traffic Rules](#).
- Apply Policies to Sites and VPNs—Associate policy with sites and VPNs in the overlay network.

Apply Centralized Policy for SD-WAN Application Intelligence Engine Flow

To ensure that a centralized data policy for the SD-WAN Application Intelligence Engine (SAIE) flow takes effect, you must apply it to a list of sites in the overlay network.

To apply a centralized policy in Cisco SD-WAN Manager, see *Configure Centralized Policy Using Cisco SD-WAN Manager*.

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 Cisco Catalyst SD-WAN Controller: 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 Cisco Catalyst SD-WAN Controller fails.

Monitor Running Applications

To enable the SD-WAN Application Intelligence Engine (SAIE) infrastructure on Cisco vEdge devices, you must enable application visibility on the devices:



Note In Cisco vManage Release 20.7.x and earlier releases, the SAIE flow is called the deep packet inspection (DPI) flow.

```
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 device.

View SAIE Applications

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

1. From the Cisco SD-WAN Manager menu, choose **Monitor > Devices**.
Cisco vManage Release 20.6.x and earlier: From the Cisco SD-WAN Manager menu, choose **Monitor > Network**.
2. Click **WAN-Edge**, select the **Device** that supports the SD-WAN Application Intelligence Engine (SAIE) flow. The Cisco SD-WAN Manager Control Connections page is displayed.
3. In the left pane, select **Real Time** to view the device details.
4. From the **Device Options** drop-down, choose **SAIE Applications** to view the list of applications running on the device.
5. From the **Device Options** drop-down, choose **SAIE Supported Applications** to view the list of applications that are supported on the device.

Action Parameters for Configuring SD-WAN Application Intelligence Engine Flow

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.

From the Cisco SD-WAN Manager menu, you can 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:

Table 22:

Description	Cisco SD-WAN Manager	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 policy configuration.	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. 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 flow** and **show log** commands.

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

Table 23:

Description	Cisco SD-WAN Manager	CLI Command	Value or Range
DSCP value.	Click Accept , then action DSCP .	set dscp <i>value</i>	0 through 63
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, green lte, metro-ethernet mpls, private1 through private6, public-internet, red, and silver.
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]	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.

Description	Cisco SD-WAN Manager	CLI Command	Value or Range
<p>Direct matching packets to the name service, before delivering the traffic to its ultimate 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 devices that are collocated with the service devices, using the vpn service configuration 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>
<p>Direct matching packets to the named service that is reachable using 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).</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>Direct traffic to a remote TLOC. The TLOC is defined by its IP address, color, and encapsulation.</p>	Click Accept , then action TLOC .	set local-tloc color <i>color</i> [encap <i>encapsulation</i>]	TLOC address, color, and encapsulation
<p>Direct traffic to one of the remote TLOCs in the TLOC list.</p>	Click Accept , then action TLOC .	set tloc-list <i>list-name</i>	Name of a policy lists tloc-list list
<p>Set the VPN that the packet is part of.</p>	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.

From the Cisco SD-WAN Manager menu, 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.

Configure SD-WAN Application Intelligence Engine Flow Using the CLI

Following are the high-level steps for configuring a centralized data policy for the SD-WAN Application Intelligence Engine (SAIE) flow.



Note In Cisco vManage Release 20.7.x and earlier releases, the SAIE flow is called the deep packet inspection (DPI) flow.

1. Create a list of overlay network sites to which the data policy is to be applied using 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)
```

Use the following show commands for visibility in to traffic classification:

- show app dpi flows
- show support dpi flows active detail
- show app dpi application
- show support dpi flows expired detail
- show support dpi statistics



CHAPTER 10


Application-Aware Routing

- [Information About Application-Aware Routing, on page 139](#)
- [Configure Application-Aware Routing, on page 148](#)
- [Configure Application-Aware Routing Using CLIs, on page 164](#)
- [Configure Application Probe Class Using CLI, on page 166](#)
- [Application-Aware Routing Policy Configuration Example, on page 167](#)

Information About Application-Aware Routing

Application-aware routing tracks network and path characteristics of the data plane tunnels between Cisco IOS XE Catalyst SD-WAN devices and uses the collected information to compute optimal paths for data traffic. These characteristics include packet loss, latency, and jitter. 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 Catalyst 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.



Each Cisco IOS XE Catalyst SD-WAN device supports up to eight TLOCs, allowing a single Cisco IOS XE Catalyst 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.

Application-Aware Routing Support for Multicast Protocols

Table 24: Feature History

Feature	Release Information	Description
Application-Aware Routing Policy Support for Multicast	Cisco IOS XE Catalyst SD-WAN Release 17.3.1a Cisco vManage Release 20.3.1	This feature enables support for configuring application-aware routing policy for multicast traffic on Cisco IOS XE Catalyst SD-WAN devices based on source and destination, protocol matching and SLA requirement.

Starting from Cisco IOS XE Catalyst SD-WAN Release 17.3.1a, application-aware routing supports overlay multicast traffic on Cisco IOS XE Catalyst SD-WAN devices. In older releases, an application-route policy is supported only for unicast traffic.

The Cisco IOS XE Catalyst SD-WAN devices classify the multicast traffic based on the group address and sets the SLA class. The group address can be source IP, destination IP, source prefixes, and destination prefixes. In the forwarding plane, any traffic for group address must use only those TLOC paths that meet the SLA requirement. You can perform the path selection for a group based on the preferred color, backup color, or the default action.

Restrictions for Multicast Protocols

Network-Based Application Recognition (NBAR) using the Cisco Catalyst SD-WAN Application Intelligence Engine (SAIE) flow is not supported for multicast.




Note In Cisco vManage Release 20.7.1 and earlier releases, the SAIE flow is called the deep packet inspection (DPI) flow.

Components of Application-Aware Routing

The Cisco IOS XE Catalyst 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 Catalyst SD-WAN Controller, which then passes them to the appropriate Cisco IOS XE Catalyst SD-WAN devices.
- **Monitoring and measuring**—The Cisco IOS XE Catalyst 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 Catalyst 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 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.



368472

To create a data policy based on the Layer 7 application itself, configure the Cisco Catalyst SD-WAN Application Intelligence Engine (SAIE) flow with a centralized data policy. With the SAIE flow, 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.



Note In Cisco vManage Release 20.7.1 and earlier releases, the SAIE flow is called the deep packet inspection (DPI) flow.

SLA Classes

Table 25: Feature History

Feature	Release Information	Description
Support for SLA Classes	Cisco IOS XE Catalyst SD-WAN Release 17.2.1r	This feature allows you to configure up to a maximum of eight SLA classes on Cisco SD-WAN Controller. Using this feature, you can configure additional options in an application-aware routing policy.
Support for six SLA Classes per Policy	Cisco IOS XE Catalyst SD-WAN Release 17.3.1a Cisco vManage Release 20.3.1	This feature allows you to configure up to six SLA classes per policy on Cisco IOS XE Catalyst SD-WAN devices. This enhancement allows additional options in an application-aware routing policy.
SLA Class Support Enhancement	Cisco IOS XE Catalyst SD-WAN Release 17.6.1a Cisco vManage Release 20.6.1	This feature is an enhancement to support up to 16 SLA classes on Cisco IOS XE Catalyst SD-WAN devices.
Application Aware Routing and Data Policy SLA Preferred Colors	Cisco IOS XE Catalyst SD-WAN Release 17.6.1a Cisco vManage Release 20.6.1	This feature provides different behaviors to choose preferred colors based on the SLA requirements when both application-aware routing policy and data policies are configured.

A service-level agreement (SLA) determines actions taken in application-aware routing. The SLA class defines the maximum jitter, maximum latency, maximum packet loss, or a combination of these values for data plane tunnels in Cisco IOS XE Catalyst SD-WAN devices. Each data plane tunnel comprises a local transport locators (TLOC) and a remote TLOC pair. You can configure SLA classes under the **policy sla-class** command hierarchy on Cisco SD-WAN Controllers. From Cisco IOS XE Catalyst SD-WAN Release 17.2.1r, you can configure a maximum of eight SLA classes on Cisco SD-WAN Validator. However, you can define only four unique SLA classes in an application-aware route policy. In releases earlier than Cisco IOS XE Catalyst SD-WAN Release 17.2.1r, you can configure a maximum of four SLA classes.

Starting from Cisco IOS XE Catalyst SD-WAN Release 17.3.1a, you can configure up to six SLA classes per policy on the Cisco IOS XE Catalyst SD-WAN devices.

You can configure the following parameters in an SLA class.

Table 26: SLA Components

Description	Command	Value or Range
Maximum acceptable packet jitter on the data plane tunnel	jitter <i>milliseconds</i>	1–1000 milliseconds
Maximum acceptable packet latency on the data plane tunnel.	latency <i>milliseconds</i>	1–1000 milliseconds
Maximum acceptable packet loss on the data plane tunnel.	loss <i>percentage</i>	1–100 percent

Starting from Cisco IOS XE Catalyst SD-WAN Release 17.15.1a, the threshold values for SLA class list are as follows:

Table 27: Threshold Values for SLA Class Lists

SLA Class	Loss %	Latency (ms)	Jitter (ms)
Voice-And-Video	2	300	60
Transactional Data	1	200	200
Bulk data	5	500	500
Default	5	500	500

SLA Support Enhancement

From Cisco IOS XE Catalyst SD-WAN Release 17.6.1a and Cisco vManage Release 20.6.1, you can configure more than six SLA classes per policy on Cisco IOS XE Catalyst SD-WAN devices.

Cisco IOS XE Catalyst SD-WAN devices need 16 GB RAM or more to support upto 16 SLA classes.

This feature enhancement increases the number of SLA classes supported on Cisco SD-WAN Controller and SD-WAN Edge devices. With the increase in the SLA class support, you can align SLA classes to IP Virtual Private Networks (IP-VPN) on Multi-Protocol Label Switching (MPLS) networks for transporting traffic to a global network.

The SLA enhancement helps in multitenancy, where you can push different SLA classes for different tenants. The multitenancy feature requires the Cisco SD-WAN Controller to support more than eight SLA classes. To allocate SLA classes to different tenants, the global limit for policies must be 64.



Note You cannot configure the default SLA. The default SLA is configured in all the devices to forward traffic when no user-defined SLA is met.

Table 28: Maximum SLA Classes Supported on Cisco IOS XE Catalyst SD-WAN Devices

Supported Platforms and Models	User-configurable SLA Classes prior to Cisco IOS XE Catalyst SD-WAN Release 17.6.1a (+1 Default SLA Class)	User-configurable SLA Classes from Cisco IOS XE Catalyst SD-WAN Release 17.6.1a (+1 Default SLA Class)
ASR 1001 HX -16GB • vedge-ASR-1001-HX	6	15
ASR 1002 X -16GB • vedge-ASR-1002-X	6	15
ASR 1002 HX -16GB • vedge-ASR-1002-HX	6	15
ASR 1001 X -16GB • vedge-ASR-1001-X	6	15
ISR 4451 X • vedge-ISR-4451-X	6	7
ISR 4431 • vedge-ISR-4431	6	7
Catalyst 8300 Edge Platforms • vedge-C8300-2N2S-6G • vedge-C8300-2N2S-4G2X • vedge-C8300-1N1S-6G • vedge-C8300-1N1S-4G2X • vedge-C8300-1N1S-6T • vedge-C8300-1N1S-4T2X • vedge-C8300-2N2S-6T • vedge-C8300-2N2S-4T2X	NA	7
Catalyst 8500 Edge platforms -16GB • vedge-C8500L-8S4X • vedge-C8500-12X4QC • vedge-C8500-12X	NA	15

Supported Platforms and Models	User-configurable SLA Classes prior to Cisco IOS XE Catalyst SD-WAN Release 17.6.1a (+1 Default SLA Class)	User-configurable SLA Classes from Cisco IOS XE Catalyst SD-WAN Release 17.6.1a (+1 Default SLA Class)
Any other Cisco IOS XE Catalyst SD-WAN devices (C11xx, ISR1100, and CSR1000v)	6	6

SLA-Preferred Colors

From Cisco IOS XE Catalyst SD-WAN Release 17.6.1a, when you configure both application-aware routing policy and data policy, and if data flow matches the app-route and data policy sequences, the following expected behaviors occur:

- If the preferred colors that you configure in application-aware routing meet the SLA requirements, and these preferred colors have some colors that are common with data policy, the common preferred colors are chosen over others for forwarding. (Prior to Cisco IOS XE Catalyst SD-WAN Release 17.6.1a, the data policy-preferred colors were forwarded and the application-aware routing policy preferences were ignored.)
- If preferred colors in application-aware routing do not meet the SLA, but there are colors that are common with the data policy, and these colors meet the SLA in application-aware routing, then these colors take precedence and are chosen for forwarding.
- If no tunnels or colors meet the SLA in application-aware routing, the data policy takes precedence and is chosen for forwarding. If the data policy has preferred colors, these colors are chosen. Otherwise, load balance occurs across all the colors in the data policy.

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 Catalyst 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 Catalyst 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 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 Catalyst 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.

Per-Class Application-Aware Routing

Table 29: Feature History

Feature Name	Release Information	Description
Per-Class Application-Aware Routing	Cisco IOS XE Catalyst SD-WAN Release 17.4.1a Cisco vManage Release 20.4.1	This feature enhances the capabilities of directing traffic to next-hop addresses based on the service level agreement (SLA) definitions. These SLA definitions along with the policy to match and classify traffic types can be used to direct traffic over specific Cisco Catalyst SD-WAN tunnels. The SLA definition comprises of values of loss, latency, and jitter, which are measured using the Bidirectional Forwarding Detection (BFD) channel that exists between two transport locators (TLOCs).

Per-Class Application-Aware Routing Overview

The SLA definition comprises of values of loss, latency, and jitter, which are measured using the BFD channel that exists between two TLOCs. These values collectively represent the status of the network and the BFD link. The BFD control messages are sent with a high priority Differentiated Services Code Point (DSCP) marking of 48.

The SLA metrics based on the high priority packet does not reflect the priority that is received by the actual data that flows through the edge device. The data, depending on the application class, can have different DSCP values in the network. Therefore, a more accurate representation of the loss, latency, and jitter for the traffic profiles is required for the networks to use such measurements to direct traffic types to the right tunnels.

Application-aware routing uses policies that constrain paths that can be used for forwarding the application. These constraints are usually expressed in terms of SLA classes that contain loss, latency, and jitter requirements that must be met. This requires that these metrics be measured on all the paths to the destination of the traffic using active probing or by passive monitoring.

Active probing methods include generation of synthetic traffic that is injected along with real traffic. The expectation is that the probes and the real traffic is forwarded in the same way. BFD probing, ICMP, periodic HTTP requests and IP SLA measurements are some examples of active probing mechanisms. The Cisco Catalyst SD-WAN solution uses BFD based probes for active measurements. Passive monitoring methods rely on the Cisco Catalyst SD-WAN Application Intelligence Engine (SAIE) flow and monitoring actual traffic. For example, RTP/TCP traffic is monitored for loss, latency, and jitter.



Note In Cisco vManage Release 20.7.1 and earlier releases, the SAIE flow is called the deep packet inspection (DPI) flow.

Application Probe Class

An application probe class (app-probe-class) comprises of a forwarding class, color, and DSCP. This defines the marking per color of applications that are forwarded. The color or DSCP mapping is local to a Cisco SD-WAN network site. However, a few colors and the DSCP mapping for a color does not change per site. The forwarding class determines the QoS queue in which the BFD echo request is queued at the egress tunnel

port. This is applicable only for BFD echo request packets. The packet-loss-priority for BFD packets is fixed to low. When BFD packets are sent with SLA class, they use the same DSCP value. When BFD packets are sent with app-probe-class along with SLA class, the BFD packets are sent for each SLA app-probe-class separately in a round-robin manner.



Note When the application route policy is applied at a site, only the colors relevant to the site are used. Since six SLA classes are supported on Cisco IOS XE Catalyst SD-WAN devices, the device correspondingly supports up to six app-probe-classes.

Default DSCP Values

The default DSCP value that is used in the DSCP control traffic is 48. However, there is a provision to change the default value along with the option to configure on the edge devices. All the network service providers may not necessarily use DSCP 48.

The BFD packet having the default DSCP can also be used for other features such as PMTU. A change in the default DSCP means that the other features are affected by the new default DSCP value. Therefore, we recommend that you configure the highest priority DSCP marking that the service provider provides (usually 48, but can be different based on the SLA agreement of the service provider). The color level overrides the global level default DSCP marking.

Configure Application-Aware Routing

Table 30: Feature History

Feature Name	Release Information	Description
Application-Aware Routing for IPv6	Cisco IOS XE Catalyst SD-WAN Release 17.9.1a Cisco vManage Release 20.9.1	This feature enables you to configure application-aware routing (AAR) policies to operate with IPv6 application traffic.

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 Catalyst 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 Catalyst SD-WAN devices.

An application-aware routing policy is a type of centralized data policy: you configure it on the Cisco SD-WAN Controller, and the controller automatically pushes it to the affected Cisco IOS XE Catalyst 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 SLA class is configured for the default-action, the packet is accepted and forwarded with no consideration of SLA. Because application-aware routing policy accepts nonmatching traffic by default,

it is considered as a positive policy. Other types of policies in the Cisco IOS XE Catalyst SD-WAN software are negative policies, because by default they drop nonmatching traffic.

Starting from Cisco IOS XE Catalyst SD-WAN Release 17.9.1a and Cisco vManage Release 20.9.1, you can configure AAR and data policies to control IPv6 traffic based on match application or app-list criteria.

Prior to Cisco IOS XE Catalyst SD-WAN Release 17.9.1a, IPv6 traffic did not have capability to match the IPv6 traffic based on Application name or application list to steer IPv6 traffic based on the desired intent.

Configure Application-Aware Routing Policies Using Cisco SD-WAN Manager

To configure application-aware routing policy, use the Cisco SD-WAN Manager policy configuration wizard. For Centralized Policy configuration details, see [Configure Centralized Policies](#). The wizard consists of four sequential windows 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. For configuration details, see [Configure Groups of Interest](#).
- **Configure Topology:** Create the network structure to which the policy applies. For topology configuration details, see [Configure Topology and VPN Membership](#).
- **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 windows, you are creating policy components or blocks. In the last window, you are applying policy blocks to sites and VPNs in the overlay network.

For an application-aware routing policy to take effect, you must activate the policy.

Configure Best Tunnel Path

Table 31: Feature History

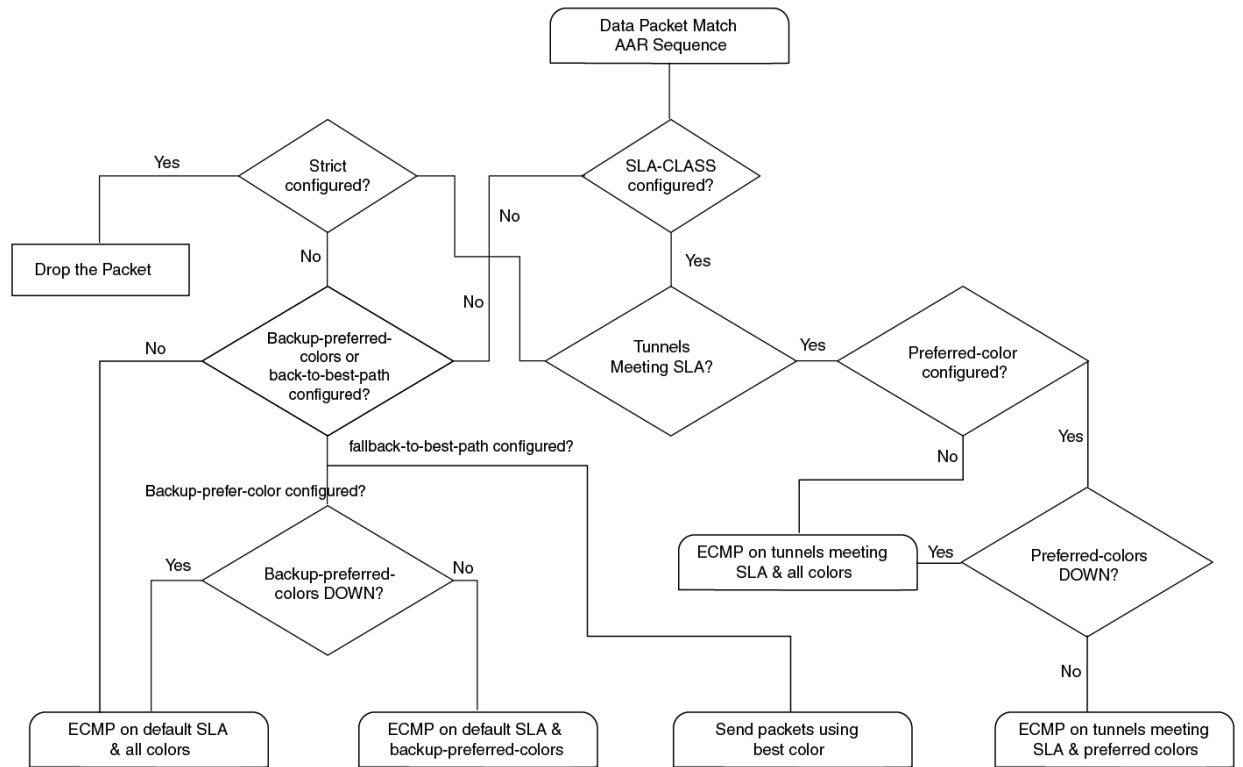
Feature Name	Release Information	Description
Best of the Worst (BOW) Tunnel Selection	Cisco IOS XE Catalyst SD-WAN Release 17.5.1a Cisco vManage Release 20.5.1	This feature introduces a new policy action fallback-to-best-path to pick the best path or color out of the available colors. When the data traffic does not meet any of the SLA class requirements, this feature allows you to select the best tunnel path criteria sequence using the Fallback Best Tunnel option under each SLA class to avoid packet loss.

Best Tunnel Path Overview

To avoid data packet loss and to configure the best application-aware routing tunnel selection when a SLA is not met, you can configure the following policy actions:

- **backup-preferred-color**
- **fallback-to-best-path**

Figure 14: Flow Chart for Application-Aware Routing Tunnel Selection



Recommendation for the Best Tunnel Path

- Configure the **fallback-to-best-path** policy action in Cisco SD-WAN Manager when configuring a SLA class.
- Configure the **backup-preferred-color** policy action in Cisco SD-WAN Manager when configuring traffic rules.

Configure Variance for Best Tunnel Path

Cisco SD-WAN Manager uses best of worst (BOW) to find a best tunnel when no tunnel meets any of the SLA class requirements.

Assume that the required latency is 100 ms to meet the SLA class requirements and tunnel T1 has 110 ms. Tunnel T2 has 111 ms and tunnel T3 has 112 ms.

As per the BOW logic, the best tunnel is T1. T2 and T3 are equally the best tunnels, with only a difference of a few ms.

You configure variance in Cisco SD-WAN Manager when configuring an SLA class. Variance accommodates small deviations as part of the best tunnel selection.

For more information, see [Configure SLA Class](#).

Example: Without Variance Configured

At time t1: T1 has 100 ms, T2 has 101 ms, and T3 has 102 ms

At time t2: T1 has 101 ms, T2 has 100 ms, and T3 has 102 ms

At time t3: T1 has 101 ms, T2 has 112 ms, and T3 has 100 ms

At time t1, the best tunnel changes from T1 to T2, and for time t2, the best tunnel changes from T2 to T3. Because variance is not configured, this leads to data path reprogramming and changes to the data traffic paths.

Assume instead that you configure variance to dampen a small deviation in ms.

For example, you configure variance as 5 ms, which means that the best tunnel SLA = 100 ms. The range is from 100 ms to 105 ms.

Example: With Variance Configured

BOW(t1) = {T1, T2, T3}

BOW(t2) = {T1, T2, T3}

BOW(t3) = {T1, T2, T3}

With variance configured, there is no data path reprogramming required or changes to data traffic paths.

Verify Configuration of Variance for Best Tunnel Path**Example for Latency Variance**

```
Device# show sdwan policy from-vsmart
from-vsmart sla-class video
latency                100
jitter                 150
  fallback-best-tunnel  latency
```

Tunnel T1: Latency: 110 msec, Loss: 0%, Jitter: 200 msec

Tunnel T2: Latency: 115 msec, Loss: 0%, Jitter: 200 msec

Tunnel T3: Latency: 120 msec, Loss: 0%, Jitter: 200 msec

Without latency variance, the best tunnel is T1.

With latency variance configured as 10 ms, T1, T2, and T3 are the best tunnels.

The range is from 110 ms to 120 ms.

The best latency + variance is 110 ms + 10 ms.

Use the following formula to find the best tunnel selection for latency variance:

(best_latency, best_latency + latency_variance)

Example for Jitter Variance

```
Device# show sdwan policy from-vsmart
from-vsmart sla-class video
latency                100
jitter                 150
  fallback-best-tunnel  jitter
```

Tunnel T1: Latency: 90 msec, Loss: 0%, Jitter: 160 msec

```
Tunnel T2: Latency: 80 msec, Loss: 0%, Jitter: 200 msec
Tunnel T3: Latency: 70 msec, Loss: 0%, Jitter: 152 msec
```

Without jitter variance, the best tunnel is T3.

With jitter variance configured as 10 ms, T1 and T3 are the best tunnels.

The range is from 152 ms to 162 ms.

The best jitter + variance is 152 ms + 10 ms.

Use the following formula to find the best tunnel selection for jitter variance:

(best_jitter, best_jitter + jitter_variance)

Example for Loss Variance

```
Device# show sdwan policy from-vsmart
from-vsmart sla-class video
latency                100
jitter                 1
  fallback-best-tunnel  loss
```

```
Tunnel T1: Latency: 110 msec, Loss: 2%, Jitter: 200 msec
Tunnel T2: Latency: 115 msec, Loss: 3%, Jitter: 200 msec
Tunnel T3: Latency: 120 msec, Loss: 4%, Jitter: 200 msec
```

Without loss variance, the best tunnel is T1.

With loss variance configured as 1%, T1 and T2 are the best tunnels.

The range is from 2% to 3%.

The best loss + variance is 2%.

Use the following formula to find the best tunnel selection for loss variance:

(best_loss, best_loss + loss_variance)

Configure SLA Class

1. From the Cisco SD-WAN Manager menu, select **Configuration** > **Policies**. Centralized Policy is selected and displayed by default.
2. Click **Add Policy**.
3. In the create groups of interest page, from the left pane, click **SLA Class**, and then click **New SLA Class List**.
4. In the **SLA Class List Name** field, enter a name for SLA class list.
5. Define the SLA class parameters:
 - a. In the **Loss** field, enter the maximum packet loss on the connection, a value from 0 through 100 percent.
 - b. In the **Latency** field, enter the maximum packet latency on the connection, a value from 1 through 1,000 milliseconds.
 - c. In the **Jitter** field, enter the maximum jitter on the connection, a value from 1 through 1,000 milliseconds.

- d. Choose the required app probe class from the **App Probe Class** drop-down list.
6. (Optional) Check the **Fallback Best Tunnel** check box to enable the best tunnel criteria.

This optional field is available from Cisco IOS XE Catalyst SD-WAN Release 17.5.1a to pick the best path or color from the available colors when a SLA is not met. When this option is selected, you can choose the required criteria from the drop-down. The criteria are a combination of one or more of loss, latency, and jitter values.
7. Select the **Criteria** from the drop-down. The available criteria are:
 - None
 - Latency
 - Loss
 - Jitter
 - Latency, Loss
 - Latency, Jitter
 - Loss, Latency
 - Loss, Jitter
 - Jitter, Latency
 - Jitter, Loss
 - Latency, Loss, Jitter
 - Latency, Jitter, Loss
 - Loss, Latency, Jitter
 - Loss, Jitter, Latency
 - Jitter, Latency, Loss
 - Jitter, Loss, Latency
8. (Optional) Enter the **Loss Variance (%)**, **Latency Variance (ms)**, and the **Jitter Variance (ms)** for the selected criteria.

For more information, see [Configure Variance for Best Tunnel Path](#).
9. Click **Add**.

Configure Traffic Rules

To configure an application-aware routing policy:

1. Click **Application Aware Routing**.
2. From the **Add Policy** drop-down list, choose **Create New**.
3. 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. You can copy, delete, or rename a policy sequence. The name you enter 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/Actions** dialog box opens, and **Match** is selected by default. The available policy match conditions are listed below the dialog box.
6. In the **Protocol** drop-down list, choose one of the following option:
 - **IPv4**
 - **IPv6**
 - **Both**



Note Depending on which protocol that you choose, the **Match** or **Actions** conditions may be different.

7. Click and choose one or more **Match** conditions. Set the values as described in the following table:

Table 32: Match Conditions

Match Condition	Procedure
None (match all packets)	Do not specify any match conditions.
Application/Application Family List	Click Application/Application Family List and choose an application list. This match condition is available for IPv6 traffic from Cisco IOS XE Release 17.9.1a and Cisco vManage Release 20.9.1.
Cloud SaaS Application List	Cisco SD-WAN Manager provides a list of several cloud applications that Cisco Catalyst SD-WAN Cloud OnRamp for SaaS can use to determine the best path selection for each SaaS application. For more information on Cisco Catalyst SD-WAN Cloud OnRamp for SaaS, see the <i>Cisco Catalyst SD-WAN Cloud OnRamp Configuration Guide, Cisco IOS XE Release 17.x</i> . Note Cloud SaaS Application List displays as a match condition if you specify IPv4 as the Protocol option. In the drop-down list, choose a SaaS application from the drop-down list.
DNS Application List	In the drop-down list, select an application family. This match condition is available for IPv6 traffic from Cisco IOS XE Release 17.9.1a and Cisco vManage Release 20.9.1.
Destination Data Prefix	To match a list of destination prefixes, choose the list from the drop-down list. To match an individual destination prefix, type the prefix in the Destination dialog box.

Destination Region	<p>You can use Destination Region in a Cisco Catalyst SD-WAN network using Cisco Catalyst SD-WAN Multi-Region Fabric.</p> <p>Choose one of the following options from the drop-down list:</p> <ul style="list-style-type: none"> • Primary: Match traffic if the destination site is in the same primary region (also called access region) as the source. • Secondary: Match traffic if the destination site is not in the same primary region as the source but is within the same secondary region as the source. This traffic can reach the destination using a direct tunnel, as described for secondary regions. • Other: Match traffic if the destination site is not in the same primary region or secondary region as the source. This traffic requires a multi-hop path from the source to the destination. <p>For more information on how to configure Multi-Region Fabric, see the <i>Cisco Catalyst SD-WAN Multi-Region Fabric (also Hierarchical SD-WAN) Configuration Guide</i>.</p>
Destination Port	<p>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>
Traffic To	<p>When creating a data policy or an application-aware policy for a border router for Multi-Region Fabric, you can use match criteria to match traffic flowing to the access region, the core region, or a service VPN.</p>
DNS (to enable split DNS)	<p>In the drop-down list, choose Request to process DNS requests for the DNS applications, and choose Response to process DNS responses for the applications.</p>
DSCP	<p>Type the DSCP value, a number from 0 through 63.</p>
PLP	<p>Choose Low or High. To set the PLP to High, apply a policer that includes the exceed remark option.</p>
Protocol	<p>Type the internet protocol number, a number from 0 through 255.</p>
ICMP Message	<p>For Protocol (IPv4), when you select a value as 1 in the Protocol field in the Match Conditions section, the ICMP Message field displays where you can select an ICMP message to apply to the data policy.</p> <p>For Protocol (IPv6), when you select a value as 58 in the Protocol field in the Match Conditions section, the ICMP Message field displays where you can select an ICMP message to apply to the data policy.</p> <p>Note This field is available from Cisco IOS XE Release 17.4.1 or Cisco SD-WAN Release 20.4.1, and also Cisco vManage Release 20.4.1.</p> <p>When Protocol is selected as Both, the ICMP Message or ICMPv6 Message field displays.</p>

Source Data Prefix	To match a list of source prefixes, choose the list from the drop-down list. To match an individual source prefix, enter the prefix in the Source field.
Source Port	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 [-]).

- To select actions for the matched data traffic, click **Actions**. Set the values as described in the following table:

Table 33: Actions

Action	Procedure
Backup SLA Preferred Color	Set the policy action for a Backup SLA Preferred Color match condition. When no tunnel matches the SLA, direct the data traffic to a specific tunnel. Data traffic is sent out the configured tunnel if that tunnel interface is available. If that tunnel interface is not available, traffic is sent out to another available tunnel. You can specify one or more colors. The backup SLA preferred color is a loose matching condition, not a strict matching condition.
Counter	Set the policy action for a Counter match condition. Click Counter . In the Counter Name field, enter the name of the file in which to store packet counters.
Log	You can 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 five minutes thereafter, as long as the flow is active. Click Log to enable logging.

Action	Procedure
SLA Class List	<p>Set the policy action for an 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 device 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.</p> <p>Click SLA Class List.</p> <p>In the SLA Class drop-down list, choose one or more SLA classes.</p> <p>Optionally, when the Preferred Color is not selected, you can choose the preferred color group from the Preferred Color Group drop-down list. Select the preferred color group of the data plane tunnel or tunnels to prefer. You can configure up to three levels of priority based on the color or path preference. This field is available from Cisco IOS XE Catalyst SD-WAN Release 17.9.1a and Cisco vManage Release 20.9.1.</p> <p>Optionally, in the Preferred Color drop-down list, choose the color of the data plane tunnel or tunnels to prefer. Traffic is load-balanced across all the 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.</p> <p>Click Strict/Drop to perform strict matching of the SLA class. If no data plane tunnel is available that satisfies the SLA criteria, traffic is dropped.</p> <p>Set a Remote Preferred Color in the AAR policy to control traffic routing based on the application list. You can add multiple remote preferred colors in the AAR policy.</p> <p>Use the Restrict to Remote Color to restrict the tunnel to preferred TLOCs. With Restrict to Remote Color option, the traffic drops when the SLA is not met with the preferred remote color.</p> <p>Click Fallback to best path to select the best available tunnel to avoid a packet drop.</p> <p>Note The Fallback to best path option is available from Cisco IOS XE Catalyst SD-WAN Release 17.5.1a and Cisco SD-WAN Release 20.5.1.</p> <p>You can select the Fallback to best path action only when the Fallback Best Tunnel option is enabled while defining a SLA class. If the Fallback Best Tunnel option is not enabled, then the following error message displays in Cisco SD-WAN Manager:</p> <pre>SLA Class selected, does not have Fallback Best Tunnel enabled. Please change the SLA class or change to Strict/Drop.</pre> <p>Click Load Balance to load balance traffic across all the tunnels.</p>
Cloud SLA	<p>Cloud SLA enables traffic to use the best path selection with Cisco Catalyst SD-WAN Cloud OnRamp for SaaS.</p> <p>Click Cloud SLA.</p>

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

Default Action of Application-Aware Routing Policy

The default action of the policy 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 Catalyst SD-WAN device selects one of the available tunnels by performing load-balancing across equal paths.

Expected behavior when data flow matches both AAR and data policies:

1. When data policy local TLOC action is configured, the **App-route preferred-color** and **backup-preferred-color** actions are ignored.
2. The **sla-class** and **sla-strict** actions are retained from the application routing configuration.
3. The data policy TLOC takes precedence.

When there is a **local-tloc-list** action that has multiple options, choose the local-TLOC that meets SLA.

- If no **local-tloc** meets SLA, then choose equal-cost multi-path routing (ECMP) for the traffic over the **local-tloc-list**.
- If none of the **local-tloc** is up, then choose a TLOC that is up.
- If none of the **local-tloc** is up and the DP is configured in restrict mode, then drop the traffic.

Configure Application Probe Class through Cisco Catalyst SD-WAN Manager

1. From the Cisco SD-WAN Manager menu, choose **Configuration > Policies**.
2. In **Centralized Policy**, click **Add Policy**. The **Create Groups of Interest** page appears.
3. Choose the list type **App Probe Class** from the left navigation panel to create your groups of interest.
4. Click **New App Probe Class**.
5. Enter the probe class name in the **Probe Class Name** field.
6. Choose the required forwarding class from the **Forwarding Class** drop-down list.

If there are no forwarding classes, then create a class from the **Class Map** list page under the **Localized Policy Lists** in the **Custom Options** menu.

To create a forwarding class:

- a. In the **Custom Options** drop-down, choose **Lists** from the Localized Policy options.
- b. In the Define Lists window, choose the list type **Class Map** from the left navigation panel.
- c. Click **New Class List** to create a new list.
- d. Enter **Class** and choose the **Queue** from the drop-down list.

- e. Click **Save**.
7. In the **Entries** pane, choose the appropriate color from the **Color** drop-down list and enter the **DSCP** value.
Click + sign, to add more entries as required.
8. Click **Save**.

Add App-Probe-Class to an SLA Class

1. From the left pane, select **SLA Class**.
2. Click **New SLA Class List**.
3. In the **SLA Class List Name** field, enter a name for SLA class list.
4. Enter the required **Loss (%)**, **Latency (ms)**, and **Jitter (ms)**.
5. Choose the required app probe class from the **App Probe Class** drop-down list.
6. Click **Add**.

The new SLA Class created with loss, latency, jitter, and app probe class is added to the table.

Configure Default DSCP on Cisco BFD Template

1. From the Cisco SD-WAN Manager menu, select **Configuration > Templates**.
2. Click **Feature Templates**.



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

3. Click **Add Template**.
4. Select a device from the device list in the left pane.
5. In the right pane, select the BFD template listed under Basic Information.
6. Enter **Template Name** and **Description** in the respective fields.
7. In the **Basic Configuration** pane, enter **Multiplier** and **Poll Interval (milliseconds)**.
8. In the **Default DSCP value for BFD Packets** field, enter the required device specific value or choose the default value for DSCP.
9. (Optional) In the **Color** pane, choose the required color from the drop-down list.
10. Enter the required **Hello Interval (milliseconds)** and **Multiplier**.
11. Choose the **Path MTU Discovery** value.
12. Enter the **BFD Default DSCP value for tloc color**.
13. Click **Add**.

The default DSCP and color values are configured on the BFD template.

Apply Policies to Sites and VPNs

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

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

1. From the Cisco SD-WAN Manager menu, choose **Configuration > Policies**. Centralized Policy is selected and displayed by default.
2. Click **Add Policy**. The Create Applications or Groups of Interest page is displayed.
3. Click **Next**. The Network Topology window opens, and in the Topology bar, Topology is selected by default.
4. Click **Next**. The Configure Traffic Rules window opens, and in the Application-Aware Routing bar, Application-Aware Routing is selected by default.
5. Click **Next**. The Apply Policies to Sites and VPNs window 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** page appears, and the policies table includes the newly created 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 Catalyst 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 Catalyst SD-WAN 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 Catalyst SD-WAN Controller by issuing a **commit** command, the controller pushes the application-aware routing policy to the Cisco IOS XE Catalyst SD-WAN devices at the specified sites.

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

To view the policy that the Cisco Catalyst SD-WAN 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 Catalyst 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 Catalyst SD-WAN device, you can configure the following types of data policy:

- Centralized data policy. You configure this policy on the Cisco Catalyst SD-WAN 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 Catalyst SD-WAN device. You configure application-aware routing policy on the Cisco Catalyst SD-WAN 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 Catalyst SD-WAN device, so data traffic inbound from the WAN is processed only by access lists and data policy.



Note When both application-aware routing and data policies are configured, if the data policy rules that contain actions such as redirect DNS, NextHop, secure internet gateway, NAT VPN, or service, the traffic which matches those rules will skip AAR policy even though the traffic also matches rules defined in the AAR policy. Data policy actions override AAR rules.

Activate an Application-Aware Routing Policy

To activate a policy:

1. From the Cisco SD-WAN Manager menu, choose **Configuration > Policies**. **Centralized Policy** is selected and displayed by default.
2. For the desired policy, click **...** and select **Activate**. The Activate Policy popup opens. It lists the IP addresses of the reachable Cisco SD-WAN Controllers to which the policy is to be applied.
3. Click **Activate**.

When you activate an application-aware routing policy, the policy is sent to all the connected Cisco SD-WAN Controllers.

Monitor Data Plane Tunnel Performance

The Bidirectional Forwarding Detection (BFD) protocol runs over all data plane tunnels between Cisco IOS XE Catalyst 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 Catalyst 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	
192.0.2.1	192.0.2.254	ipsec	12346	12346	0	22	0	596	0	21	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
5	596	0	29	4	0	0								
192.0.2.1	192.0.2.254	ipsec	12346	12346	0	24	0	596	0	24	3	0	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
5	596	0	24	3	0	0								
192.0.2.1	192.0.2.254	ipsec	12346	34083	0	21	0	596	0	21	3	0	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
5	596	0	21	3	0	0								
192.0.2.1	192.0.2.254	ipsec	12346	36464	0	23	0	596	0	23	3	0	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 Catalyst SD-WAN Devices

You can enable application visibility directly on Cisco IOS XE Catalyst 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.

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
vSmart(config-sla-class)# app-probe-class app-probe-class
vSmart(config-sla-class)# fallback-best-tunnel criteria latency loss jitter
```

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
vSmart(config-sequence)# <userinput>action backup-sla-preferred-color</userinput>
<varname>colors</varname>
```

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.
 - **backup-sla-preferred-color** *colors*—Direct the data traffic to a specific tunnel. Data traffic is sent out the configured tunnel if that tunnel interface is available; if that tunnel is unavailable, traffic is sent out another available tunnel. You can specify one or more colors. As with the **preferred-color** option, the backup SLA preferred color is loose matching. In a single **action** configuration, you cannot include both the **strict** and **backup-sla-preferred-color** options.
- 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 rule 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
```

Configure Application Probe Class Using CLI

Configure app-probe-class, real-time-video and map them with the SLA class as shown in the following example:

```
Device(config)# app-probe-class real-time-video
Device(config)# forwarding-class videofc
Device(config)# color mpls dscp 34
Device(config)# color biz-internet dscp 40
```

```
Device(config)# color lte dscp 0

Device(config)# sla-class streamsla
Device(config)# latency 20
Device(config)# loss 10
Device(config)# app-probe-class real-time-video
```

Configure the default value for DSCP using BFD template as shown:

```
Device(config)# bfd default-dscp 50
Device(config)# bfd color mpls 15
```

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.

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

- Definition of the application (or applications)
- Definition of App Probe Class (Optional)
- 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 Catalyst SD-WAN 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)#
```

- 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)#
```

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

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

- Display the configuration changes:

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

- Validate that the configuration contains no errors:

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

- Activate the configuration:

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

- Exit from configuration mode:

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

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
!
!

```

The following example defines the multicast protocol:

```

policy
!
sla-class SLA_BEST_EFFORT
  jitter 900
!
sla-class SLA_BUSINESS_CRITICAL
  loss 1
  latency 250
  jitter 300
!
sla-class SLA_BUSINESS_DATA
  loss 3
  latency 400
  jitter 500
!
sla-class SLA_REALTIME
  loss 2
  latency 300
  jitter 60
!
app-route-policy policy_multicast
  vpn-list multicast-vpn-list
  sequence 10
  match
    source-ip 10.0.0.0/8
    destination-ip 10.255.255.254/8
  !
  action
    count mc-counter-10
    sla-class SLA_BUSINESS_CRITICAL
  !
!
sequence 15

```

```

match
  source-ip      172.16.0.0/12
  destination-ip 172.31.255.254/12
  !
action
  count mc-counter-15
  sla-class SLA_BEST_EFFORT
  !
!
sequence 20
match
  destination-ip 192.168.0.1
  !
action
  count mc-counter-20
  sla-class SLA_BUSINESS_CRITICAL
  !
!
sequence 25
match
  protocol      17
  !
action
  count mc-counter-25
  sla-class SLA_REALTIME
  !
!
sequence 30
match
  source-ip      192.168.0.0/16
  destination-ip 192.168.255.254
  protocol      17
  !
action
  count mc-counter-30
  sla-class SLA_BUSINESS_DATA preferred-color lte
  !
!
default-action sla-class SLA_BEST_EFFORT
!
sequence 35
match
  source-ip      10.0.0.0/8
  destination-ip 10.255.255.254/8
  protocol      17
  !
action
  count mc-counter-35
  sla-class SLA_BUSINESS_DATA preferred-color lte
  backup-sla-preferred-color 3g
  !
!
lists
vpn-list multicast-vpn-list
  vpn 1
  vpn 60
  vpn 4001-4010
  vpn 65501-65510
  !
site-list multicast-site-list
  site-id 1100
  site-id 500
  site-id 600
  !

```

```

!
!
apply-policy
  site-list multicast-site-list
  app-route-policy policy_multicast
!
!

```

The following example defines remote color preference:

```

vSmart# show running-config policy
policy
  sla-class SLA1
  latency 100
!
app-route-policy AAR1
  vpn-list vpn1
  sequence 1
  match
    destination-ip 10.1.1.0/24
  !
  action
    sla-class SLA1 preferred-color mpls lte
    sla-class remote-preference
    remote-color mpls lte
    remote-color-restrict

```

Ranking Color Preference Example

```

app-route-policy SAMPLE _AAR
  vpn-list ONE
  sequence 10
  match
    dscp 46
  !
  action
    sla VOICE_SLA strict preferred-color-group GROUP2_COLORS
  !
!
sequence 20
  match
    dscp 34
  !
  action
    sla VOICE_SLA preferred-color-group GROUP1_COLORS
  !
!
sequence 30
  match
    dscp 28
  !
  action
    sla VOICE_SLA preferred-color-group GROUP3_COLORS
  !
!
!
policy lists
  preferred-color-group GROUP1_COLORS
  primary-preference
  color-preference biz-internet
  path-preference direct-tunnel
  !
  secondary-preference

```

```

    color-preference mpls
    path-preference multi-hop-path
    !
    tertiary-preference
    color-preference lte
    !
    !
    preferred-color-group GROUP2_COLORS
    primary-preference
    color-preference mpls
    !
    secondary-preference
    color-preference biz-internet
    !
    !
    preferred-color-group GROUP3_COLORS
    primary-preference
    color-preference mpls biz-internet lte
    !

```



Note You can configure path-preference option only if you enable the Multi-Region Fabric option in Cisco SD-WAN Manager.

AAR Policy for IPv6 Applications Example

```

policy
  sla-class Default
    jitter 100
    latency 300
    loss 25
  !
  app-route-policy _VPN1_AAR-Policy-for-IPv6-Traffic
  vpn-list VPN1
  sequence 1
  match
    app-list Msft-0365
  !
  action
    sla-class Default preferred-color public-internet
  !
  !
!
lists
  app-list Msft-0365
  app ms-office-web-apps
  !
  site-list SITE-100
  site-id 100
  !
  vpn-list VPN1
  vpn 1
  !
!
!
apply-policy
  site-list SITE-100
  app-route-policy _VPN1_AAR-Policy-for-IPv6-Traffic
  !
!

```




CHAPTER 11

Enhanced Application-Aware Routing

Table 34: Feature History

Feature Name	Release Information	Description
Enhanced Application-Aware Routing	<p>Cisco IOS XE Catalyst SD-WAN Release 17.12.1a and later.</p> <p>Cisco Catalyst SD-WAN Manager Release 20.12.1 and later.</p>	<p>Without enhanced application-aware routing enabled, Cisco IOS XE Catalyst SD-WAN devices require several minutes to switch traffic from one network path to another to meet SLA requirements when the loss, latency, and jitter exceed specific threshold values.</p> <p>Enabling enhanced application-aware routing speeds the detection of tunnel performance issues. This enables Cisco IOS XE Catalyst SD-WAN devices to redirect traffic away from tunnels that do not meet SLA requirements.</p>

- [Information About Enhanced Application-Aware Routing, on page 173](#)
- [Supported Devices for Enhanced Application-Aware Routing, on page 177](#)
- [Restrictions for Enhanced Application-Aware Routing, on page 177](#)
- [Prerequisites for Enhanced Application-Aware Routing, on page 177](#)
- [Configure Enhanced Application-Aware Routing, on page 177](#)
- [Verify the Enhanced Application-Aware Routing Configuration, on page 179](#)
- [Monitor Enhanced Application-Aware Routing Using Cisco Catalyst SD-WAN Manager, on page 180](#)
- [Troubleshooting Enhanced Application-Aware Routing, on page 181](#)

Information About Enhanced Application-Aware Routing

Without enhanced application-aware routing enabled, Cisco IOS XE Catalyst SD-WAN devices require several minutes to switch traffic from one network path to another to meet SLA requirements when the loss, latency, and jitter exceed specific threshold values. Enabling enhanced application-aware routing speeds the detection

of tunnel performance issues. This enables Cisco IOS XE Catalyst SD-WAN device to redirect traffic away from tunnels that do not meet SLA requirements.

Overview of Enhanced Application-Aware Routing

BFD (Bidirectional Forwarding Detection) detects link failure conditions and gathers performance routing data (PfR), including loss, latency, and jitter information of Cisco Catalyst SD-WAN tunnels (both IPsec and GRE). Each BFD hello packet collects the following information:

Latency: RTT (Round trip time) between BFD echo request and reply.

Jitter: The variation in the delay of packet arrival times in a network. It is a measure of the irregularity in the timing of data packets as they are transmitted and received.

Loss: Number of echo requests that fail to receive a reply.

By default, with a BFD hello timer of 1 second, one sample of PfR data is collected every second. This PfR data is collected over the duration of the poll interval (default 10 minutes). During the poll interval, the average of each statistic is computed. To determine dynamic path decisions based on the thresholds specified in application-aware routing SLAs, a default multiplier of 6 is employed to review multiple averages of the poll-interval. A poll interval average refers to the average time duration between consecutive polling or measurement events in a network monitoring or performance measurement system. The poll interval average provides an indication of how frequently the system collects data or samples network metrics over a specific time-period.

Convergence time refers to the amount of time it takes for the network to recover and resume normal operations after a failure or disruption. However, the default convergence time for detection of slowly degrading WAN circuits is between 10 minutes and 1 hour. Even with the lowest recommended poll-interval of 2 minutes and 6 intervals, the convergence time is between 2 minutes and 12 minutes. Setting a very low poll interval can result in false positives of PfR and traffic instability due to insufficient sample data for loss, latency, and jitter measurements.

PfR Measurements

Table 35: PfR Measurements

Metric	Source	Description
Loss	BFD	<p>Measured as loss of BFD packet at 1pps or one packet in <code>n_app_probe_class (n-apc) sec</code></p> <p>If the application probe class (APC) configuration is not set, the loss of BFD packets occurs at a rate of 1 packet per second (1pps). With the APC configuration, the loss is reduced to 1 packet in N seconds.</p> <p>For more information see, Application Probe Class.</p>

Metric	Source	Description
Latency	BFD	RTT measurements 1 pps or one packet in n-apc sec Without the application probe class (APC) configuration, the loss of RTT packets occurs at a rate of 1 packet per second (1pps). With the APC configuration, the loss is reduced to 1 packet in N seconds.
Jitter	BFD	Variation in RTT

Application-Aware Routing Design and Measurements

- The default BFD hello-interval is 1 sec, and the app-route/SLA poll-interval is 10 mins:

The BFD hello-interval refers to the frequency at which BFD (Bidirectional Forwarding Detection) protocol sends hello packets to detect the liveness of a network path. By default, the hello-interval is set to 1 second. On the other hand, the app-route/SLA poll-interval determines how frequently the network monitoring system collects data or measures network metrics related to application routes or Service Level Agreements (SLAs). The default poll-interval for app-route/SLA is set to 10 minutes.

- By default, the system calculates to 60 minutes using 1 pps x 600 sec x 6 buckets:

Refers to the calculation of a default value for the poll-interval in minutes. It calculates the interval by multiplying 1 packet per second (pps) by 600 seconds (10 minutes) and then multiplying the result by 6 buckets. The resulting value is 60 minutes, which is the default poll-interval.

- Experts suggest using a poll-interval of 120 seconds (2 minutes) and a multiplier of 5, which results in a 10-minute interval. This recommendation is often followed to achieve a specific monitoring frequency.
- Reducing the poll-interval/multiplier helps improve detection time but may lead to false positives with a small number of samples for PfR metrics:

Decreasing the poll-interval and/or the multiplier can enhance the speed at which network performance issues are detected. However, reducing these values may also increase the likelihood of false positives, which is that the system may incorrectly identify issues due to a small number of data samples. The detection time and the accuracy of PfR (Performance Routing) metrics must be balanced.

- The only option is to improve the measurement accuracy at a faster rate by reducing the BFD Hello interval:

To achieve a faster and more accurate measurement of network performance, the recommended approach is to decrease the BFD hello-interval. Network path liveness refers to the condition of the connectivity and availability of network paths. By reducing the interval at which hello packets are exchanged, the liveness of network paths can be detected more frequently, leading to improved measurement accuracy.

Benefits of Enhanced Application-Aware Routing

1. Improved the PfR metrics (loss/latency/jitter) measurements by introducing inline data that allows for more accurate and detailed measurements of these metrics. Inline data refers to the traffic that is processed and inspected directly at the edge of the network, within the Cisco IOS XE Catalyst SD-WAN devices.

Instead of routing all the traffic to a central location for analysis and security checks, inline data allows for real-time inspection and decision-making at the network edge.

2. Quick Enhanced-App-Route Detection and SLA Enforcement, which involves reducing the PfR poll-interval to a very low value (minimum of 10 seconds). This allows the Cisco IOS XE Catalyst SD-WAN devices to quickly detect any slow degradation of circuits. If a circuit fails to meet the SLA threshold, the tunnels are swiftly switched out from SLA forwarding to ensure efficient and reliable network performance. SLA (Service Level Agreement) forwarding refers to the capability of the Cisco Catalyst SD-WAN solution to dynamically route network traffic based on predefined performance criteria or SLAs.
3. The speed of SLA switch-over is improved.
4. SLA Dampening is introduced for a smoother transition to SLA forwarding. Before implementing SLA forwarding again, the tunnel goes through a process called dampening, which helps prevent disruptions and instabilities. This ensures a smooth transition back to SLA, minimizing any negative effects on network performance.
5. Enhancements are made to measure loss, latency, and jitter.

Guidelines of Enhanced Application-Aware Routing

- Both GRE and IPSEC tunnels are supported.
- All existing TLOCs and WAN interface types, including physical, sub interface, loopback bind, dialer, and LTE interfaces, are supported.
- TLOC Extension tunnels are supported.
- Both IPv4 and IPv6 underlay tunnels are supported.
- SLA update and switchover occur at a minimum interval of 10 seconds.
- Tunnel scale is not impacted, with minimal impact on memory and performance.
- Support is provided with and without app-probe class configuration in SLA classes.
- SLA dampening is supported.

Compatibility With Cisco IOS XE Catalyst SD-WAN devices Not Running Enhanced Application-Aware Routing

1. In the following scenario:
 - On the local side: The Cisco IOS XE Catalyst SD-WAN device is upgraded to Cisco IOS XE Catalyst SD-WAN Release 17.12.1a and later and has EAAR (Enhanced Application-Aware Routing) enabled.
 - On the remote side: The Cisco IOS XE Catalyst SD-WAN device is not upgraded to Cisco IOS XE Catalyst SD-WAN Release 17.12.1a and the EAAR is not enabled.

Then the system will fall back to using BFD based measurements where support compatibility with older releases and disabled features are present.

2. If both the local and remote sides are using Cisco IOS XE Catalyst SD-WAN Release 17.12.1a but the EAAR feature is not enabled, the system will revert to using BFD based measurements.



Note The EAAR feature is disabled by default to support existing deployments.

Supported Devices for Enhanced Application-Aware Routing

Cisco IOS XE Catalyst SD-WAN devices

Restrictions for Enhanced Application-Aware Routing

- The branch device on which you enable this feature does not support loopback unbind mode. The loopback unbind mode refers to a network interface configuration in which the loopback device is disconnected from the network stack.
- There is no per-queue measurement for GRE tunnels. Per queue measurement is used to monitor and analyze network traffic on a per-queue basis. It involves measuring and collecting various metrics and statistics for each individual queue in a network device or system. A queue is a buffer where packets are stored before they are transmitted or processed.

Prerequisites for Enhanced Application-Aware Routing

To enable application-aware routing on a Cisco IOS XE Catalyst SD-WAN device, enable enhanced application-aware routing on both the Cisco IOS XE Catalyst SD-WAN devices.

Configure Enhanced Application-Aware Routing

The procedures in this section describe how to deploy the enhanced app-aware routing configurations from Cisco Catalyst SD-WAN Manager to Cisco IOS XE Catalyst SD-WAN devices.

Configure Enhanced Application-Aware Routing Using a Feature Template in Cisco Catalyst SD-WAN Manager

1. From the **Cisco SD-WAN Manager** menu, choose **Configuration > Templates**.
2. Click **Feature Templates**.
3. Click **Add Template**.
4. Choose a device and click the **Cisco System** template under **Basic Information**.
5. In the **Enhanced App-Aware Routing** field, click **Global** from the drop-down list and choose one of the following modes:

Mode	EAAR Poll Interval	EAAR Poll Multiplier	EAAR Poll Window	SLA Dampening Multiplier	SLA Dampening Window
Aggressive	10s	6	10s - 60s	120	20 mins
Moderate	60s	5	60s - 300s	40	40 mins
Conservative	300s	6	300s - 1800s	12	60 mins



Note You can configure the enhanced application aware routing (EAAR) poll interval, poll multiplier, and SLA dampening multiplier only through CLI template.

6. Click **Save**.

Configure Enhanced Application-Aware Routing Using a Configuration Group in Cisco Catalyst SD-WAN Manager

1. From the **Cisco SD-WAN Manager** menu, choose **Configuration > Configuration Groups**.
2. Choose a configuration group. Under **Actions** click **Edit**.
3. Under **Feature Profiles** click **System Profile**.
4. Choose **basic** and under **Actions** click **Edit Feature**.
5. In the **Edit Basic Feature** page, use the **Enhanced App-Route** field and choose one of the modes as follows:

Mode	EAAR Poll Interval	EAAR Poll Multiplier	EAAR Poll Window	SLA Dampening Multiplier	SLA Dampening Window
Aggressive	10s	6	10s - 60s	120	20 mins
Moderate	60s	5	60s - 300s	40	40 mins
Conservative	300s	6	300s - 1800s	12	60 mins

6. Click **Save**.

Configure Enhanced Application-Aware Routing Using a CLI Template

For more information about using CLI templates, see [CLI Add-On Feature Templates](#) and [CLI Templates](#). By default, CLI templates execute commands in global configuration mode.

1. Enable enhanced PfR measurements for SLA enforcement.

```
bfd enhanced-app-route enable
```

Enabling the application-aware routing feature on a Cisco IOS XE Catalyst SD-WAN device requires you to enable the PfR CLI on both the remote Cisco IOS XE Catalyst SD-WAN device and the local Cisco IOS XE Catalyst SD-WAN device.

This feature involves two steps:

- a. The remote Cisco IOS XE Catalyst SD-WAN device must provide loss statistics to the local Cisco IOS XE Catalyst SD-WAN device.
 - b. The local Cisco IOS XE Catalyst SD-WAN device then utilizes these metrics to enforce Service Level Agreements (SLAs).
2. When the enhanced application aware Pfr is enabled, the default poll-interval of 10 seconds and multiplier of 6 is used for SLA enforcement and switchover. To modify these settings, use the following configuration options:

```
bfd enhanced-app-route pfr-poll-interval
```

```
bfd enhanced-app-route pfr-multiplier <number>
```

The aggressive mode setting for app route pfr multiplier is 6 by default. It is 5 for the moderate mode.

3. Configure the SLA dampening time. This is the waiting time before returning the tunnel to SLA buckets after meeting the SLA. The default time is 120 seconds. Enable the SLA dampening when the enhanced PFR is enabled.

```
bfd sla-dampening enable
```

```
bfd sla-dampening multiplier <number>
```

The aggressive mode setting for dampening multiplier is 120 by default.

Verify the Enhanced Application-Aware Routing Configuration

To verify the enhanced application-routing configuration and display the configured params for EAAR use the **show sdwan app-route params** command.

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

```
*EAAR = Enhanced Application-Aware Routing
Config:                :Enabled
Poll interval:         :10000
Poll multiplier:       :6
```

```
App route
Poll interval:         :600000
Poll multiplier:       :6
```

```
SLA dampening
Config:                :Enabled
Multiplier:           :120
```

You can use the **show sdwan bfd sessions alt** command to highlight the flags for EAAR.

```
Device# show sdwan bfd sessions alt
```

```
*Sus = Suspend
*GREinUDP = GREinUDP encap
*EAAR = Enhanced Application-Aware Routing
*NA = Flag Not Set
```

SYSTEM IP	SITE ID	STATE	DST PUBLIC		ENCAP	SOURCE TLOC		REMOTE TLOC	SOURCE IP
			COLOR PORT	COLOR		BFD-LD	FLAGS		
172.16.0.0	100	up	lte			lte			
10.0.0.0			10.0.0.1		12367	ipsec	20013	NA	
0:07:48:38									
172.16.0.1	100	up	lte			lte			
10.0.0.0			10.0.0.1		12377	ipsec	20014	NA	
0:07:48:39									
172.16.0.0	400	up	lte			lte			
10.0.0.0			10.0.0.1		12366	ipsec	20015	NA	
0:07:48:39									
172.16.0.1	500	up	lte			lte			
10.0.0.0			10.0.0.1		12366	ipsec	20016	EAAR	
0:07:48:39									

You can use **show sdwan app-route stats summary** command to display the app-route (PfR) stats details for each tunnel, across different intervals of measurements, for every configured APC.

Device# show sdwan app-route stats summary

```
app-route statistics 10.0.0.0 10.0.0.0 ipsec 12366 12367
remote-system-ip      172.16.0.0
local-color           lte
remote-color          lte
sla-class-index       0,1,2,3
fallback-sla-class-index None
enhanced-app-route    Enabled
sla-dampening-index   4,5
app-probe-class-list  None
mean-loss             0
mean-latency          0
mean-jitter           0
```

TX INDEX	TOTAL		LATENCY	JITTER	AVERAGE	AVERAGE	TX DATA	RX DATA	IPV6
	PACKETS	LOSS			PKTS	PKTS	DATA	PKTS	DATA
0	664	0	0	0	0	0	0	0	0
	0	0				0			
1	663	0	0	0	0	0	0	0	0
	0	0				0			
2	666	0	0	0	0	0	0	0	0
	0	0				0			
3	664	0	0	0	0	0	0	0	0
	0	0				0			
4	662	0	0	0	0	0	0	0	0
	0	0				0			
5	664	0	0	0	0	0	0	0	0
	0	0				0			

Monitor Enhanced Application-Aware Routing Using Cisco Catalyst SD-WAN Manager

1. From the Cisco Catalyst SD-WAN Manager menu, choose **Monitor > Devices**.
2. Under **Devices**, choose a device.



3. Click **Real Time** in the left pane.
4. In the **Device Options** field, choose **App Routes Statistics**.

EAAR-BR-SITE700 | Site Name 700 Device Model: C8000v ⓘ

Device Options:

Filter ▾

Search

Total Rows: 48  

ocol	Source Port	Destination Port	Remote System Ip	Local Color	Remote Color	Enhanced App Route	Slas Dampening Index
	12346	12346		public-internet	public-internet	Enabled	None
	12346	12346		public-internet	public-internet	Enabled	None
	12346	12346		public-internet	public-internet	Enabled	None
	12346	12346		public-internet	public-internet	Enabled	None
	12346	12346		public-internet	public-internet	Enabled	None
	12346	12346		public-internet	public-internet	Enabled	None
	12346	12346		public-internet	public-internet	Enabled	None

Troubleshooting Enhanced Application-Aware Routing

From the device:

```
Device# show sdwan run | include enhanced-app-route
```

```
bfd enhanced-app-route enable
bfd enhanced-app-route pfr-poll-interval 10000
bfd enhanced-app-route pfr-multiplier 6
```

```
show sdwan run | inc sla-dampening
bfd sla-dampening enable
bfd sla-dampening multiplier 12
```

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

```
Enhanced app route
  Config:                :Enabled <<< Enhanced app-aware routing enabled
  Poll interval:         :10000
  Poll multiplier:       :6
App route
  Poll interval:         :600000
  Poll multiplier:       :6
SLA dampening
  Config:                :Enabled
  Multiplier:            :120
```

```
Device# show platform hardware qfp active feature sdwan datapath pathmon
summary
```

Src IP	Dst IP	Src Port	Dst Port	Encap	Uidb	Bfd Discrim	PathMon
10.0.0.0	10.0.0.1	12346	12366	IPSEC	65527	20003	in/out

Device# show sdwan bfd sessions alt

*Sus = Suspend

*GREinUDP = GREinUDP encap

*EAAR = Enhanced Application-Aware Routing

*NA = Flag Not Set

SYSTEM IP	DST PUBLIC		STATE	SOURCE TLOC	REMOTE TLOC		SOURCE IP
	SITE ID	IP		COLOR	DST PUBLIC	COLOR	
UPTIME				PORT		BFD-LD	FLAGS
172.16.0.0	100	10.0.0.1	down	privatel 12367	lte ipsec	20011	10.0.0.0 EAAR
NA							
172.16.0.1	500	10.0.0.1	down	privatel 12366	3g ipsec	20013	10.0.0.0 EAAR
NA							
172.16.0.0	600	10.0.0.1	down	privatel 12366	3g ipsec	20007	10.0.0.0 EAAR
NA							

Device# show sdwan app-route stats remote-system-ip 172.16.0.0 app-route statistics 10.0.0.0 10.0.0.1 ipsec 12366 12366

```

remote-system-ip      172.16.0.0
local-color           privatel
remote-color          3g
sla-class-index       0
fallback-sla-class-index None
enhanced-app-route    Enabled
sla-dampening-index   None

```



CHAPTER 12

Traffic Flow Monitoring

- [Traffic Flow Monitoring, on page 183](#)
- [Information About Traffic Flow Monitoring, on page 185](#)
- [Restrictions for Traffic Flow Monitoring, on page 195](#)
- [Configure Traffic Flow Monitoring, on page 195](#)
- [Verify Traffic Flow Monitoring, on page 213](#)

Traffic Flow Monitoring

Table 36: Feature History

Feature Name	Release Information	Description
Flexible NetFlow Support for IPv6 and Cache Size Modification	Cisco IOS XE Catalyst SD-WAN Release 17.4.1a Cisco vManage Release 20.4.1	This feature enables export of packets to an external collector over an IPv6 transport on Cisco IOS XE Catalyst SD-WAN devices and provides the visibility of IPv6 network traffic. If you want to monitor IPv4 and IPv6 traffic together, this feature enables you to modify the cache size on the data plane. Cisco Flexible NetFlow (FNF) is a technology that provides customized visibility into network traffic. In Cisco Catalyst SD-WAN, FNF enables exporting data to Cisco SD-WAN Manager which makes it easy for the customers to monitor and improve their network.
Log Packets Dropped by Implicit ACL	Cisco IOS XE Catalyst SD-WAN Release 17.5.1a Cisco vManage Release 20.5.1	You can now enable or disable logging of dropped packets in case of a link failure. You can also configure how often the packet flows are logged.

Feature Name	Release Information	Description
Flexible NetFlow Enhancement	Cisco IOS XE Catalyst SD-WAN Release 17.6.1a Cisco vManage Release 20.6.1	This feature enhances Flexible NetFlow to collect type of service (ToS), sampler ID, and remarked DSCP values in NetFlow records. This enhancement provides the flexibility to define flow record fields to customize flow records by defining flow record fields. The ToS and remarked DSCP fields are supported only on IPv4 records. However, the sampler ID field is supported for both IPv4 and IPv6 records.
Flexible NetFlow for VPN0 Interface	Cisco IOS XE Catalyst SD-WAN Release 17.7.1a Cisco vManage Release 20.7.1	This feature supports NetFlow on VPN0 interfaces. Flexible NetFlow acts as a security tool, enables export of data to Cisco SD-WAN Manager, detects attacks on devices, and monitors traffic.
Flexible NetFlow Export Spreading	Cisco IOS XE Catalyst SD-WAN Release 17.9.1a Cisco Catalyst SD-WAN Control Components Release 20.9.x Cisco vManage Release 20.9.1	This feature enables export spreading to prevent export storms that occur when a burst of packets are sent to external collector. The export of the previous interval is spread during the current interval to prevent export storms. When NetFlow packets are sent over a low-bandwidth circuit, the export spreading functionality is enabled to avoid packet drops.
Flexible NetFlow Export of BFD Metrics	Cisco IOS XE Catalyst SD-WAN Release 17.10.1a Cisco Catalyst SD-WAN Control Components Release 20.10.1	With this feature, you can export Bidirectional Forwarding Detection (BFD) metrics to an external collector for generating BFD metrics of loss, latency, and jitter. This feature provides enhanced monitoring and faster collection of network state data. After you enable export of BFD metrics, configure an export interval for exporting the BFD metrics.
Real-Time Device Options for Monitoring Cflowd and SAIE Flows	Cisco IOS XE Catalyst SD-WAN Release 17.10.1a Cisco vManage Release 20.10.1	With this feature, you can apply filters for monitoring specific Cflowd and Cisco Catalyst SD-WAN Application Intelligence Engine (SAIE) applications or application families running within a VPN on the selected Cisco IOS XE Catalyst SD-WAN device. Real-time device options for monitoring Cflowd and SAIE flows are available on Cisco vEdge devices. This release provides support for real-time device options for monitoring Cflowd and SAIE applications on Cisco IOS XE Catalyst SD-WAN devices.
Enhancements to Flexible NetFlow for Cisco SD-WAN Analytics	Cisco IOS XE Catalyst SD-WAN Release 17.12.1a Cisco Catalyst SD-WAN Manager Release 20.12.1	This feature introduces logging enhancements to Cisco Flexible NetFlow for IPv4 and IPv6 flow records in Cisco SD-WAN Analytics. The output of the show flow record command has been enhanced for these records.

Feature Name	Release Information	Description
Flow Telemetry Enhancement When Using Loopbacks as TLOCs.	Cisco IOS XE Catalyst SD-WAN Release 17.12.1a Cisco Catalyst SD-WAN Manager Release 20.12.1	When you configure a loopback interface as an ingress or egress transport interface, this feature enables you to collect loopback instead of physical interface in FNF records. This feature is supported for IPv4 and IPv6. Updated the show command show sdwan control local-properties wan-interface-list to display the binding relationship between the loopback and physical interfaces. A new column Bind Interface is added to the existing option, Monitor > Devices > Real Time (choose the device option, Control WAN Interface Information) in Cisco SD-WAN Manager to display the binding relationship between the loopback and physical interfaces.
Configure a Maximum FNF Record Rate for Aggregated Traffic Data	Cisco IOS XE Catalyst SD-WAN Release 17.14.1a Cisco Catalyst SD-WAN Control Components Release 20.14.1	For a device, you can configure a maximum rate (records per minute) for sending Flexible NetFlow (FNF) records of aggregated traffic data. This can reduce the performance demands on a device, and may be helpful when there is a large number of applications producing network traffic.

Information About Traffic Flow Monitoring

The following sections describe traffic flow monitoring.

Traffic Flow Monitoring with Cflowd Overview

Cflowd is a flow analysis tool, used for analyzing Flexible NetFlow (FNF) traffic data. It monitors traffic flowing through Cisco IOS XE Catalyst SD-WAN devices in the overlay network and exports flow information to a collector, where it can be processed by an IP Flow Information Export (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 SD-WAN Controllers and on Cisco SD-WAN Manager). You can configure a maximum of four Cflowd collectors per Cisco IOS XE Catalyst 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 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 Catalyst SD-WAN devices do not cache any of the records that are exported to a collector.



Note NetFlow on Secure Internet Gateway (SIG) tunnels is not supported on Cisco IOS XE Catalyst SD-WAN devices.

Cflowd and SNMP Comparison

Cflowd monitors service side traffic. Cflowd mainly monitors traffic from LAN to WAN, WAN to LAN, LAN to LAN and DIA. If you use Cflowd and SNMP to monitor traffic of LAN interface (input or output), then packets and bytes should be similar. The difference of bytes in SNMP starts from L2 header, but Cflowd starts from L3 header. However, if we use Cflowd and SNMP to monitor traffic of WAN interface (input or output), then packets or bytes are unlikely to be the same. All the traffic of WAN interfaces is not service side traffic. For example, Cflowd does not monitor BFD traffic, but SNMP does. The packets or bytes of Cflowd and SNMP traffic are not the same.

IPFIX Information Elements for Cisco IOS XE Catalyst SD-WAN Devices

The Cisco Catalyst SD-WAN Cflowd software exports the following IP Flow Information Export (IPFIX) information elements to the Cflowd collector. Fields vary depending on the release that you are on. Common fields are exported to Cisco SD-WAN Manager and external exporters. Feature fields are exported only to Cisco SD-WAN Manager.

Before Cisco IOS XE Catalyst SD-WAN Release 17.2.1r, Flexible NetFlow exports all fields to external collectors and Cisco SD-WAN Manager. Starting from Cisco IOS XE Catalyst SD-WAN Release 17.2.1r, FNF exports the elements (that are marked yes) in the following table to both external collectors and Cisco SD-WAN Manager. Other fields like **drop cause id** are for specific features and these fields are exported only to Cisco SD-WAN Manager, but not to an 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	—
sourceIPv6Address	27	Yes	IPv6 source address in the IP packet header.	ipv6Address (16 bytes)	default	—
destinationIPv4Address	12	Yes	IPv4 destination address in the IP packet header.	IPv4Address (4 bytes)	default	—

Information Element	Element ID	Exported to External Collector	Description	Data Type	Data Type Semantics	Units or Range
destinationIPv6Address	28	Yes	IPv6 destination address in the IP packet header.	ipv6Address (16 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 Catalyst 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	—

Information Element	Element ID	Exported to External Collector	Description	Data Type	Data Type Semantics	Units or Range
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
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)	—	—
ip tos	5	Yes	The Type of Service field in the IP header.	unsigned8 (1 byte)	identifier	8 bits
dscp output	98	Yes	Value of a 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
flow sampler	48	Yes	A set of properties that are defined in a Netflow sampler map that are applied to at least one physical interface	unsigned8 (1 byte)	identifier	—
bfd avg latency	45296	Yes	Calculation of the Bidirectional Forwarding Detection (BFD) average latency for each tunnel	unsigned64 (8 bytes)	identifier	—
bfd avg loss	45295	Yes	Calculation of the BFD average loss for each tunnel	unsigned64 (8 bytes)	identifier	—
bfd avg jitter	45297	Yes	Calculation of the BFD average jitter for each tunnel	unsigned64 (8 bytes)	identifier	—
bfd rx cnt	45299	Yes	Count of received BFD packets	unsigned64 (8 bytes)	deltaCounter	—

Information Element	Element ID	Exported to External Collector	Description	Data Type	Data Type Semantics	Units or Range
bfd tx cnt	45300	Yes	Count of transmitted BFD packets	unsigned64 (8 bytes)	deltaCounter	—
bfd rx octets	45304	Yes	Count of received BFD octets	unsigned64 (8 bytes)	deltaCounter	—
bfd tx octets	45305	Yes	Count of transmitted BFD octets	unsigned64 (8 bytes)	deltaCounter	—
application_CATEGORY	12232	Yes	Application category name, first level categorization for each application tag	variable length	identifier	—
application_SUB_CATEGORY	12233	Yes	Application sub category name, second level categorization for each application tag	variable length	identifier	—
application_GROUP	12234	Yes	Application group name, groups multiple app tags that belong to the same application	variable length	identifier	—
application traffic-class	12243	Yes	Application traffic-class according to SRND model	variable length	identifier	—
application business-relevance	12244	Yes	Application business-relevance	variable length	identifier	—

Flexible Netflow for VPN0 Interface

From Cisco IOS XE Catalyst SD-WAN Release 17.7.1a, you can enable FNF for bidirectional traffic visibility on a VPN0 interface of a Cisco IOS XE Catalyst SD-WAN device.

Netflow provides statistics on packets flowing through the device and helps to identify the tunnel or service VPNs. Flexible Netflow on VPN0 provides visibility for all the traffic (both ingress and egress) hitting VPN0 on Cisco IOS XE SD-WAN devices.

A profile is a predefined set of traffic that you can enable or disable for a context. You can create an Easy Performance Monitor (ezPM) profile that provides an express method of provisioning monitors. This new mechanism adds functionality and does not affect the existing methods for provisioning monitors. As part of this feature, you can create **sdwan-fnf** profile to monitor traffic passing through netflow VPN0 configuration.

A context represents a performance monitor policy map that is attached to an interface in ingress and egress directions. A context contains the information about the traffic-monitor that has to be enabled. When a context is attached to an interface, two policy-maps are created, one each in ingress and egress directions. Depending on the direction specified in the traffic monitor, the policy-maps are attached in that direction and the traffic is monitored. You can modify the context to override pre-defined directions.

You can create multiple contexts based on a single profile with different traffic monitors, different exporters, and different parameters for every selected traffic monitor. An ezPM context can be attached to multiple interfaces. Only one context can be attached to an interface.

Table 37: Flexible Netflow Components

	Cisco Catalyst SD-WAN Flexible Netflow	Cisco SD-WAN Flexible Netflow VPN0 from Cisco vM Release 20.7.1
Configuration	Localized Policy: app-visibility or flow-visibility Centralized policy: cflowd policy Supported on both Cisco SD-WAN Manager feature template and CLI template/	Define Flexible Netflow VPN0 monitor using command performance monitor context xxx profile sdwan-fnf on VPN0 interface. Supported on CLI template and add-on CLI feature template Cisco SD-WAN Manager.
Interface	Cisco Catalyst SD-WAN tunnel interface and service VPN interface	VPN0 interface except Cisco Catalyst SD-WAN tunnel interface
Flow Records	Fixed records by default. Supports dynamic monitoring for records such as, FEC, packet duplication, SSL proxy and so on. Also supports collecting type of service (ToS), sampler ID and remarked DSCP values for centralized policies.	Fixed records. You cannot modify or add new fields.
Flow Direction	Supports only ingress flows	Supports both ingress and egress by default.
NBAR for APP	Network-based Application recognition (NBAR) is enabled only when app-visibility is defined.	NBAR is enabled by default.
Exporter	JSON file to Cisco SD-WAN Manager and IPFIX to external collector	Can't export to Cisco SD-WAN Manager IPFIX to external collectors

Limitations of Flexible Netflow on VPN0 Interface

- Flexible Netflow on VPN0 is not supported on Cisco Catalyst SD-WAN tunnel and Cisco Catalyst SD-WAN VPN interfaces.
- The FNF record for VPN0 traffic is a fixed record and cannot be modified.
- Cisco Catalyst SD-WAN VPN0 flow entries are reported to external collectors defined in CLI configuration and not to Cisco SD-WAN Manager.
- Cisco Catalyst SD-WAN BFD and Cisco Catalyst SD-WAN control connections such as OMP, Netconf, and SSH are encapsulated by Datagram Transport Layer Security (DTLS) or Transport Layer Security (TLS) tunnels. FNF reports on only the DTLS traffic and not the encapsulated protocol packets.
- When FNF is configured for a VPN0 WAN interface,
 - For ingress flows (WAN > Cisco Catalyst SD-WAN-tunnel > LAN) - the output interface is reported as NULL.
 - For egress flows (LAN > Cisco Catalyst SD-WAN-tunnel > WAN) - input interface is reported as WAN interface (Cisco Catalyst SD-WAN underlay tunnels).

- VPN0 monitor supports only IPv4 and IPv6 protocols.
- For routing protocols, such as OSPF, BGP, only egress traffic is supported. Ingress OSPF and BGP traffic is treated as high priority packets.
- Only loopback interfaces are supported as source interfaces for Cflowd flow export.
- FNF records only the original DSCP values when the packets are sent to the external collector. FNF supports only ingress flows.

Flexible NetFlow Export Spreading

Minimum releases: Cisco IOS XE Catalyst SD-WAN Release 17.9.1a and Cisco vManage Release 20.9.1

Enable Flexible NetFlow export spreading on Cisco IOS XE Catalyst SD-WAN devices. The export-spreading feature spreads out the export of records in the monitor cache over a time interval to improve collector performance. In the case of a synchronized cache, all network devices export records in the monitor cache at the same time. If multiple network devices are configured with the same monitor interval and synchronized cache, the collector may receive all records from all devices at the same time, which can impact the collector performance. Set the time interval for export spreading to spread out the export over a time interval.

To ensure that the collector performance is not affected, export records at a specified time interval, spreading the exporting of records evenly over the cache timeout.

Configure FNF exports using option and data templates. Use the options templates to configure system level attributes. Use the data templates to configure flow records and corresponding data.

When you enable export-spread, configure the following three spread intervals:

- **app-tables:** application-table, application-attributes option template
- **tloc-tables:** tunnel-tloc-table option template

The bfd-metric-table introduced in Cisco IOS XE Catalyst SD-WAN Release 17.10.1a and Cisco vManage Release 20.10.1 belongs to the tloc-table category.

- **other-tables:** other option templates

The following is an example of how a spreading interval works.

- When an app-table is configured with ten application-attributes or application-table, the option template packets are sent in ten seconds for all the attributes evenly.
- The default interval is one second. So, with export-spreading, one large traffic burst of ten seconds is spread into ten smaller bursts of one second each.

Flexible NetFlow option template packets are sent as a burst regularly as set by the timeout option. With export spread interval, instead of sending the option template packets as bursts, the packets are spread across the timeout and export-spread interval.

In Cisco vManage Release 20.8.1 and earlier releases, after every 60 secs option template packets are sent as a burst. For example, if there are 1000 packets, it enqueues all the 1000 packets at the end of 60 secs which causes packet drops.

When you configure export spreading, if there are 1000 packets to be sent at the end of 60 secs, then 100 packets are sent in 10 secs at the rate of 100 packets and avoids the export bursts. If no export spread is specified, the default behavior is immediate export.

When you upgrade from a previous version which doesn't support export spreading, the default value for spreading in a Cflowd template is disabled.

Flexible NetFlow Export of BFD Metrics

Minimum releases: Cisco IOS XE Catalyst SD-WAN Release 17.10.1a and Cisco Catalyst SD-WAN Control Components Release 20.10.1

With the Flexible Netflow (FNF) export of BFD metrics feature, you can export BFD telemetry data to an external FNF collector to analyze the average jitter, average latency, and loss per tunnel. Jitter and latency are measured in units of microseconds. Loss is measured in units of one hundredth of one percent, 0.01%. This feature provides enhanced monitoring and faster collection of network state data.

A new option template, `bfd-metric-table`, is added for export of BFD metrics.

Configure export of BFD metrics on Cisco IOS XE Catalyst SD-WAN devices using a Cisco SD-WAN Manager feature template or using the CLI from a Cisco SD-WAN Controller. For more information on configuring export of BFD metrics using Cisco SD-WAN Manager feature templates, see [Configure Cflowd Monitoring Policy](#). For more information on configuring export of BFD metrics using the CLI, see [Configure Flexible Netflow with Export of BFD Metrics Using the CLI](#).

How the Export of BFD Metrics Works

Minimum releases: Cisco IOS XE Catalyst SD-WAN Release 17.10.1a and Cisco Catalyst SD-WAN Control Components Release 20.10.1

A Cisco IOS XE Catalyst SD-WAN device is responsible for sending IP Flow Information Export (IPFIX) packets to an external collector. After you configure the BFD export interval on the Cisco SD-WAN Controllers or on Cisco SD-WAN Manager, the Forwarding Table Manager (FTM) generates the source metrics.

- Example 1:

If you reboot a Cisco IOS XE Catalyst SD-WAN device, the device exports the BFD metrics according to the BFD export interval that you configured. At this point, the FTM does not have any data for exporting. As a consequence, all of the fields, except for the TLOC TABLE OVERLAY SESSION ID field, contain the following invalid value:

0xFFFFFFFF

Example 2:

- The FTM interval for sending data is greater than the BFD export interval. In this situation, data may end up getting exported twice, while the FTM sends data only once. Consequently, there is no new data received from the FTM. The BFD metrics and timestamps are the same as for the last packet.

For an example of BFD telemetry data that is sent to an external collector, see [Configuration Examples for Flexible Netflow Export of BFD Metrics](#).

Cflowd Traffic Flow Monitoring with SAIE Flows

Minimum releases: Cisco IOS XE Catalyst SD-WAN Release 17.10.1a and Cisco vManage Release 20.10.1

With this feature, you can choose two Cisco SD-WAN Manager real-time device options for monitoring both Cflowd flows and SAIE flows.

For more information on SAIE flows, see the [SD-WAN Application Intelligence Engine Flow](#) chapter.

With this feature, you can apply filters for displaying specific applications or application families running within a VPN on the selected Cisco IOS XE Catalyst SD-WAN device.

For more information on the device-filtering options for Cflowd and SAIE flows, see the [Devices and Controllers](#) chapter in the *Cisco Catalyst SD-WAN Monitor and Maintain Configuration Guide*.

Benefits of Cflowd Traffic Flow Monitoring with SAIE Flows

- Provides increased visibility of the network traffic, enabling network operators to analyze network usage and improve network performance
- Provides real-time monitoring of Cisco IOS XE Catalyst SD-WAN devices
- Provides parity with Cisco SD-WAN Manager real-time device options on Cisco IOS XE Catalyst SD-WAN devices

Prerequisites for Cflowd Traffic Flow Monitoring with SAIE Flows

Minimum releases: Cisco IOS XE Catalyst SD-WAN Release 17.10.1a and Cisco vManage Release 20.10.1

Configure application and flow visibility prior to viewing the Cflowd with SAIE flow device options.

For more information on configuring application flow visibility, see [Configure Global Application Visibility, on page 197](#).

For more information on configuring global flow visibility, see [Configure Global Flow Visibility, on page 195](#).

Restrictions for Cflowd Traffic Flow Monitoring with SAIE Flows

Minimum releases: Cisco IOS XE Catalyst SD-WAN Release 17.10.1a and Cisco vManage Release 20.10.1

- Cisco SD-WAN Manager can only display 4001 Cflowd records at a time.
- If two different users attempt to access the same query from the same device at the same time, the Cisco IOS XE Catalyst SD-WAN device processes only the first request. The second user must resend their request because the first request gets timed out.
- Search filters for Cflowd with SAIE are matched against the fetched 4001 Cflowd flow records.
- Enter the full name of the application or the application family for the search filter to return a valid result.

For example, if you want to search for the **netbios-dgm** application, and you enter **netbios** for **Application** or **Application Family**, you won't receive the correct result.

Information About Configuring a Maximum FNF Record Rate for Aggregated Data

Minimum supported releases: Cisco IOS XE Catalyst SD-WAN Release 17.14.1a, Cisco Catalyst SD-WAN Control Components Release 20.14.1

Raw and Aggregated Traffic Flow Data

When traffic flow visibility is enabled (see [Configure Global Flow Visibility](#)), devices in the network send raw and aggregated traffic flow data to Cisco SD-WAN Manager.

To aggregate flow data, routers use 4-tuples of flow data (containing VPN ID, application name, ingress interface of the flow, and egress interface of the flow) as a key for consolidating the raw data of multiple flows. The router consolidates each flow for which the 4-tuple is identical into a single aggregated FNF record.

Cisco SD-WAN Manager uses the aggregated data to provide a high-level view of network traffic flow information. The aggregated data shows the network applications that are producing traffic, but is less granular than the full traffic flow data. It does not provide source and destination addresses, or source and destination ports for traffic flows.

For a detailed view of traffic flows, use functions such as On Demand Troubleshooting. For information about On Demand Troubleshooting, see [On-Demand Troubleshooting](#).

Maximum FNF Record Rate

You can configure a maximum rate (records per minute) of aggregated traffic data FNF records that a device can send to reduce the performance demands (CPU and memory) on the device. This may be helpful when there is a large number of applications producing network traffic. For information about configuring this, see [Configure the Maximum FNF Record Rate for Aggregated Data, Using CLI Commands, on page 212](#).

Restrictions for Traffic Flow Monitoring

The following sections describe notes, limitations, and restrictions related to traffic flow monitoring.

Restrictions for Enabling Collect Loopback in Flow Telemetry When Using Loopbacks as TLOCs

- Supports configuration only through the Cisco Catalyst SD-WAN Controller CLI or Cisco SD-WAN Manager CLI-template. Feature template is not supported for this release.
- Collect loopback in FNF VPN0 interfaces is not supported.
- Collect loopback in the Decidated Internet Access (DIA) scenario, is not supported.
- Multi-tenant scenario is not supported.

Configure Traffic Flow Monitoring

The following sections provide information about configuring traffic flow monitoring.

Configure Traffic Flow Monitoring on Cisco IOS XE Catalyst SD-WAN Devices

Cflowd traffic flow monitoring uses Flexible NetFlow (FNF) to export traffic data. Perform the following steps to configure Cflowd monitoring:

Configure Global Flow Visibility

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

1. From the Cisco SD-WAN Manager menu, choose **Configuration > Policies**.
2. Click **Localized Policy**.
3. Click **Add Policy**.
4. Click **Next** to advance through the wizard pages until you reach **Policy Overview** and then the **Policy Settings** page.
5. Enter **Policy Name** and **Policy Description**.
6. Check the **Netflow** check box to enable flow visibility for IPv4 traffic.
7. Check the **Netflow IPv6** check box to enable flow visibility for IPv6 traffic.



Note Enable flow visibility for IPv4 and IPv6 traffic before configuring Cflowd traffic flows with SAIE visibility. For more information on monitoring Cflowd and SAIE flows, see the [Devices and Controllers](#) chapter of the *Cisco Catalyst SD-WAN Monitor and Maintain Configuration Guide*.

8. Check **Implicit ACL Logging** to configure your Cisco IOS XE Catalyst SD-WAN device to log dropped packets in the traffic.
With this configuration, you have visibility of the packets dropped by implicit access control lists (ACL) in case of a link failure in the system.
9. Enter **Log Frequency**.
Log frequency determines how often packet flows are logged. Maximum value is 2147483647. It is rounded down to the nearest power of 2. For example, for 1000, the logging frequency is 512. Thus, every 512th packet in the flow is logged.
10. Enter **FNF IPv4 Max Cache Entries** to configure FNF cache size for IPv4 traffic.
For example, enter 100 to configure FNF cache for IPv4/IPv6 traffic as shown in the following example.
11. Enter **FNF IPv6 Max Cache Entries** to configure FNF cache size for IPv6 traffic.
For example, enter 100 to configure FNF cache for IPv4/IPv6 traffic as shown in the following example.



Note The minimum cache size value is 16. The maximum of total cache size (IPv4 cache + IPv6 cache) should not exceed the limit for each platform. If cache size is not defined and the platform is not in the list, then default maximum cache entries is 200k.

The maximum cache entries is the maximum concurrent flows that Cflowd can monitor. The maximum cache entries vary on different platforms. For more information, contact [Cisco Support](#).

The following example shows the flow-visibility configuration for both IPv4 and IPv6:

```
policy
  flow-visibility
  implicit-acl-logging
  log-frequency 1000
  flow-visibility-ipv6
```



```
ip visibility cache entries 100
ipv6 visibility cache entries 100
```

While running `policy flow-visibility` or `app-visibility` to enable the FNF monitor, you may see the following warning message displaying a GLOBAL memory allocation failure. This log is triggered by enabling FNF monitoring (`policy flow-visibility` or `app-visibility`) with a large cache size.

```
Jul  4 01:45:00.255: %CPPEXMEM-3-NOMEM: F0/0: cpp_cp_svr: QFP: 0, GLOBAL memory allocation
of 90120448 bytes by FNF failed
Jul  4 01:45:00.258: %CPPEXMEM-3-TOPUSER: F0/0: cpp_cp_svr: QFP: 0, Top User: CPR STILE
EXMEM GRAPH, Allocations: 877, Type: GLOBAL
Jul  4 01:45:00.258: %CPPEXMEM-3-TOPUSER: F0/0: cpp_cp_svr: QFP: 0, Top User: SBC, Bytes
Allocated: 53850112, Type: GLOBAL
```

The warning message does not necessarily indicate a flow monitor application failure. The warning message can indicate internal steps that FNF uses for applying memory from the External Memory Manager (EXMEM) infrastructure.

Use the `show platform hardware qfp active active classification feature-manager exmem-usage` command to display the EXMEM memory usage for various clients.

```
Device# show platform hardware qfp active active classification feature-manager exmem-usage
```

```
EXMEM Usage Information
```

```
Total exmem used by CACE: 39668
```

Client	Id	Total VMR	Total Usage	Total%	Alloc	Free
acl	0	11	2456	6	88	84
qos	2	205	31512	79	7	5
fw	4	8	892	2	2	1
obj-group	39	82	4808	12	5	2

To ensure that the FNF monitor is enabled successfully, use the `show flow monitor monitor-name` command to check the status (allocated or not allocated) of a flow monitor.

```
Device# show flow monitor sdwan_flow_monitor
```

```
Flow Monitor sdwan_flow_monitor:
```

```
Description:    monitor flows for vManage and external collectors
Flow Record:    sdwan_flow_record-003
Flow Exporter:  sdwan_flow_exporter_1
                sdwan_flow_exporter_0
```

```
Cache:
```

```
Type:          normal (Platform cache)
Status:        allocated
Size:          250000 entries
Inactive Timeout: 10 secs
Active Timeout: 60 secs
```

```
Trans end aging: off
```

```
SUCCESS
```

```
Status:        allocated
```

```
FAILURE
```

```
Status:        not allocated
```

Configure Global Application Visibility

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

The `app-visibility` enables `nbar` to see each application of the flows coming to the router from all VPNs in the LAN. If `app-visibility` or `app-visibility-ipv6` is defined, then `nbar` is enabled globally for both IPv4 and IPv6 flows.

1. From the Cisco SD-WAN Manager menu, choose **Configuration > Policies**.
2. Click **Localized Policy**.
3. Click **Add Policy**.
4. Click **Next** to advance through the wizard pages until you reach **Policy Overview** and then the **Policy Settings** page.
5. Enter **Policy Name** and **Policy Description**.
6. Check the **Application** check box to enable application visibility for IPv4 traffic.
7. Check the **Application IPv6** check box to enable application visibility for IPv6 traffic.



Note Enable application visibility for IPv4 and IPv6 traffic before configuring Cflowd traffic flows with SAIE visibility.

For more information on monitoring Cflowd and SAIE flows, see the [Devices and Controllers](#) chapter of the *Cisco Catalyst SD-WAN Monitor and Maintain Configuration Guide*.

8. Enter **FNF IPv4 Max Cache Entries** to configure FNF cache size for IPv4 traffic.
For example, enter 100 to configure FNF cache size for IPv4 traffic as shown in the following example.
9. Enter **FNF IPv6 Max Cache Entries** to configure FNF cache size for IPv6 traffic.
For example, enter 100 to configure FNF cache size for IPv6 traffic as shown in the following example.

The following example shows the application visibility configuration for both IPv4 and IPv6:

```
policy
 app-visibility

 app-visibility-ipv6
 ip visibility cache entries 100
 ipv6 visibility cache entries 100
!
```



Note The `policy app-visibility` command also enables global flow visibility by enabling `nbar` to get the application name.



Note If you configure Cflowd global `flow-visibility`, but you do not configure Cflowd `app-visibility`, the exported application to Cisco SD-WAN Manager returns a result of unknown. The same application exported to an external collector using the IPFIX analyzer may contain an incorrect application name.

If you want to retain the application name, define Cflowd `app-visibility` to avoid this issue.

Configure Cflowd Monitoring Policy

To configure a policy for Cflowd traffic flow monitoring, use the Cisco SD-WAN Manager policy configuration wizard. The wizard consists of four sequential pages that guide you through the process of creating and editing policy components:

1. **Create Applications or Groups of Interest:** Create lists that group related items together 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 a policy with sites and VPNs in the overlay network.

In the first three policy configuration wizard pages, create policy components or blocks. In the last page, apply policy blocks to sites and VPNs in the overlay network. For the Cflowd policy to take effect, activate the policy.

1. From the Cisco SD-WAN Manager menu, choose **Configuration > Policies**.
2. Click **Custom Options**.
3. Under **Centralized Policy**, click **Traffic Policy**.
4. Click **Cflowd**.
5. Click **Add Policy** and then click **Create New**.
6. Enter the **Name** and **Description** for the policy.
7. In the **Cflowd Template** section, enter **Active Flow Timeout**.
8. In the **Inactive Flow Timeout** field, enter the timeout range.
9. In the **Flow Refresh** field, enter the range.
10. In the **Sampling Interval** field, enter the sample duration.
11. In the **Protocol** drop-down list, choose an option from the drop-down list.

Starting from Cisco IOS XE Catalyst SD-WAN Release 17.6.1a and Cisco vManage Release 20.6.1, the **Advanced Settings** field displays when you choose **IPv4** or **Both** from the options.

12. Under the **Advanced Settings**, do the following to collect additional IPv4 flow records:
 - Check the **TOS** check box.
 - Check the **Re-marked DSCP** check box.
13. Under the **Collector List**, click **New Collector**. You can configure up to four collectors.
 - a. In the **VPN ID** field, enter the number of the VPN in which the collector is located.
 - b. In the **IP Address** field, enter the IP address of the collector.
 - c. In the **Port** field, enter the collector port number.
 - d. In the **Transport Protocol** drop-down list, choose the transport type to use to reach the collector.
 - e. In the **Source Interface** field, enter the name of the interface to use to send flows to the collector.

- f. In the **Export Spreading** field, click the **Enable** or **Disable** radio button.

Starting from Cisco IOS XE Catalyst SD-WAN Release 17.9.1a and Cisco vManage Release 20.9.1, the **Export Spreading** field is available to prevent export storms that occur due to the creation of a synchronized cache. The export of the previous interval is spread during the current interval to prevent export storms.

- g. In the **BFD Metrics Exporting** field, click the **Enable** or **Disable** radio button.

Starting from Cisco IOS XE Catalyst SD-WAN Release 17.10.1a and Cisco vManage Release 20.10.1, the **BFD Metrics Exporting** field is available for collecting BFD metrics of loss, jitter, and latency.

- h. In the **Exporting Interval** field, enter the interval in seconds for sending BFD metrics.

Starting from Cisco IOS XE Catalyst SD-WAN Release 17.10.1a and Cisco vManage Release 20.10.1, the **Exporting Interval** field is available for specifying the export interval for BFD metrics.

Once you enable BFD metrics exporting, you can see the **Exporting Interval** field.

The **Exporting Interval** field controls the intervals by which BFD metrics are sent.

The default BFD export interval is 600 seconds.

Field	Description
Cflowd Policy Name	Enter a name for the Cflowd policy.
Description	Enter a description for the Cflowd policy.
Active Flow Timeout	Enter an active flow timeout value. The range is 30 to 3600 seconds. Active flow timeout is the time interval Netflow records are exported for long lived flows.
Inactive Flow Timeout	Enter an inactive flow timeout value. The range is 1 to 3600 seconds. Inactive flow timeout is the time interval that flows are not active for a period of time (For example, 15 seconds) that is exported from the flow cache.
Flow Refresh	Enter the interval for sending Cflowd records to an external collector. The range is 60 through 86400 seconds.
Sampling Interval	Enter the sample duration. The range is 1 through 65536 seconds. Sampling interval is the time duration taken to collect one of the sample in packets.
Protocol	Choose the traffic protocol type from the drop-down list. The options are: IPv4 , IPv6 , or Both . The default protocol is IPv4 .
TOS	Check the TOS check box. This indicates the type of field in the IPv4 header.

Field	Description
Re-marked DSCP	Check the Re-marked DSCP check box. This indicates the traffic output specified by the remarked data policy.
VPN ID	Enter the VPN ID. The range is 0 through 65536.
IP Address	Enter the IP address of the collector.
Port	Enter the port number of the collector. The range is from 1024 through 65535.
Transport Protocol	Choose the transport type from the drop-down list to reach the collector. The options are: TCP or UDP .
Source Interface	Choose the source interface from the drop-down list.
Export Spreading	Click the Enable or Disable radio button to configure export spreading. The default is Disable .
BFD Metrics Exporting	Click the Enable or Disable radio button to configure export of Bidirectional Forwarding Detection (BFD) metrics. The default is Disable .
Exporting Interval	Enter the export interval in seconds for sending the BFD metrics to an external collector. Enter an integer value. This field is displayed only if you enable BFD metrics export. The default BFD export interval is 600 seconds.

14. Click **Save Cflowd Policy**.

View Cflowd Information

To view Cflowd information, use the following commands on the Cisco IOS XE Catalyst 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



Note Starting from Cisco IOS XE Catalyst SD-WAN Release 17.15.1a, the preceding show commands retrieve up to 4000 flow records for each monitor (IPv4 and IPv6) from the cflowd database. The flow records exceeding 4000 are not shown.

The following sample output displays Cflowd information:

```
Device# show sdwan app-fwd cflowd flows
Generating output, this might take time, please wait ...
app-fwd cflowd flows vpn 1 src-ip 10.2.2.11 dest-ip 10.20.24.17 src-port 0 dest-port 2048
dscp 63 ip-prot 1
tcp-cntrl-bits          0
icmp-opcode            2048
total-pkts             6
total-bytes            600
start-time              "Fri May 14 02:57:23 2021"
egress-intf-name       GigabitEthernet5
ingress-intf-name      GigabitEthernet1
application            unknown
family                 network-service
drop-cause              "No Drop"
drop-octets            0
drop-packets           0
sla-not-met            0
color-not-met          0
queue-id               2
tos                    255
dscp-output            63
sampler-id             3
fec-d-pkts             0
fec-r-pkts             0
pkt-dup-d-pkts-orig    0
pkt-dup-d-pkts-dup     0
pkt-dup-r-pkts         0
pkt-cxp-d-pkts         0
traffic-category       0
```

For more information on Cflowd flows, see the [show sdwan app-fwd cflowd flows](#) command page.

Configure Cflowd Traffic Flow Monitoring Using the CLI

From the CLI on the Cisco SD-WAN Controller that is controlling the Cisco IOS XE Catalyst 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
vSmart(config-cflowd-template)# protocol ipv4|ipv6|Both
```



Note On Cisco IOS XE Catalyst SD-WAN devices, a flow-active-timeout is fixed as 60 seconds. If a flow-inactive-timeout is fixed as 10 seconds. The **flow-active-timeout** and **flow-inactive-timeout** value that is configured on Cisco SD-WAN Controller or Cisco SD-WAN Manager do not take effect on Cisco IOS XE Catalyst SD-WAN devices.

- To collect TOS, DSCP output and TLOC loopback in flow monitor:

Starting Cisco Catalyst SD-WAN Manager Release 20.12.1, when you configure a loopback interface as an ingress or egress transport interface, this feature enables you to collect loopback instead of physical interface in FNF records. This feature is supported for IPv4 and IPv6.

```
vSmart(config-cflowd-template)# customized-ipv4-record-fields
vsmart(config-customized-ipv4-record-fields)# collect-tos
vsmart(config-customized-ipv4-record-fields)# collect-dscp-output
vSmart(config-cflowd-template)# collect-tloc-loopback
```

- 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
export-spread
enable
app-tables app-tables
tloc-tables tloc-tables
other-tablesother-tables
```



Note You can configure app-tables, tloc-tables, and other-tables options only using Cisco SD-WAN Controllers.



Note Cisco IOS XE Catalyst SD-WAN devices only support UDP collector. Irrespective of the transport protocol that is configured, UDP is the default collector for Cisco IOS XE Catalyst SD-WAN devices.

- 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
```

- Create lists of sites in the overlay network that contain the Cisco IOS XE Catalyst 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
```

- Apply the data policy to the sites in the overlay network that contain the Cisco IOS XE Catalyst 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
```

Configure Flexible Netflow on VPN0 Interface

You can enable FNF on a VPN0 interface using a CLI template or the CLI add-on template. The ezPM profile helps in creating a new profile to carry all the Netflow VPN0 monitor configuration. On selecting a profile and specifying a few parameters, ezPM provides the remaining provisioning information. A profile is a pre-defined set of traffic monitors that can be enabled or disabled for a context. You can configure Easy Performance Monitor (ezPM) and enable FNF as follows.

```
Device# config-transaction
Device(config)# performance monitor context <monitor_name> profile <sdwan-fnf> traffic-monitor
<all> [ipv4/ipv6]
Device(config-perf-mon)# exporter destination <destination address> source <source interface>
transport udp vrf <vrf-name> port <port-number> dscp <dscp>
```

The following example shows how to configure a performance monitor context using the sdwan-fnf profile. This configuration enables monitoring of traffic metrics. Here, 10.1.1.1 is the IP address of the third-party collector, GigabitEthernet5 is the source interface, and 4739 is the listening port of the third-party collector.

```
Device# config-transaction
Device(config)# performance monitor context <monitor_name> profile sdwan-fnf traffic-monitor
all [ipv4/ipv6]
Device(config-perf-mon)# exporter destination <10.1.1.1> source <GigabitEthernet5> transport
udp vrf <vrf1> port <4739> dscp <1>
```

Configure Flexible NetFlow with Export of BFD Metrics Using the CLI

Minimum releases: Cisco IOS XE Catalyst SD-WAN Release 17.10.1a and Cisco Catalyst SD-WAN Control Components Release 20.10.1

From the CLI on the Cisco SD-WAN Controller that is controlling the Cisco IOS XE Catalyst SD-WAN device, enter the following commands depending on if you want to enable or disable the export of BFD metrics using a data policy:

1. Enable the export of BFD metrics.

```
policy
  cflowd-template template-name
  collector vpn vpn-id address ip-address port port transport transport
  source-interface interface
  bfd-metrics-export
  export-interval export-interval
```

The default BFD export interval is 600 seconds. BFD export interval is independent of a Cflowd template refresh. The BFD export interval only controls the interval for sending data from the bfd-metrics-export table. For the tunnel-tloc table, the BFD export interval uses the minimum value between the BFD export interval and the Cflowd template refresh as the interval to send data.

2. Disable the export of BFD metrics.

```
policy
  cflowd-template template-name
```



```

collector vpn vpn-id address ip-address port port transport transport
source-interface interface
no bfd-metrics-export

```

Here is a complete configuration example for enabling BFD metrics export.

```

policy
cflowd-template fnf
  template-refresh 600
  collector vpn 0 address 10.0.100.1 port 4739 transport transport_udp
  bfd-metrics-export
  export-interval 30
  !
  !
  !
lists
  site-list 500
  site-id 500
  !
  !
!
apply-policy
  site-list 500
  cflowd-template fnf
  !
!

```

Configuration Examples for Flexible NetFlow Export of BFD Metrics

Minimum releases: Cisco IOS XE Catalyst SD-WAN Release 17.10.1a and Cisco Catalyst SD-WAN Control Components Release 20.10.1

The following example shows a centralized policy configuration with export of BFD metrics enabled:

```

Device# show sdwan policy from-vsmart
from-vsmart cflowd-template fnf
flow-active-timeout 600
flow-inactive-timeout 60
template-refresh 600
flow-sampling-interval 1
protocol ipv4
customized-ipv4-record-fields
no collect-tos
no collect-dscp-output
collector vpn 0 address 10.0.100.1 port 4739 transport transport_udp
bfd-metrics-export
export-interval 600

```

The following example shows FNF BFD telemetry data with average jitter, average latency, and loss metrics:

```

{ 'Data_Template': 'Data_Flow',
  'ObservationDomainId': 6,
  'Version': 10,
  'arrive_time': 1658807309.2496994,
  'dfs_tfs_length': 200,
  'export_dfs_tfs_templates_list_dict': { 'FlowSequence': 3354,
                                          'Flowset_id': '258',
                                          'Flowset_length': 200,
                                          'Length': 286,
                                          'ObservationDomainId': 6,
                                          'TimeStamp': 1658807269,
                                          'Version': 10,

```

```

'flow': [ { 'bfd_avg_jitter': 1000,
           'bfd_avg_latency': 1000,
           'bfd_loss': 15,
           'bfd_pfr_update_ts': 1658806692155,
           'bfd_rx_cnt': 0,
           'bfd_tx_cnt': 0,
           'ipDiffServCodePoint': 48,
           'tloc_table_overlay_session_id': 10},
          ...
        ]},
'flow_length': 4,
'flow_time': 1658807269,
'flowset_id': '258',
'header': { 'FlowSequence': 3354,
           'Length': 286,
           'ObservationDomainId': 6,
           'TimeStamp': 1658807269,
           'Version': 10},
'host': '10.0.100.15',
'ipfix_length': 286,
'packet_number': 2,
'template_id': '258'}

```

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 apply these two site lists to two different **data-policy** policies, the attempt to commit the configuration on the Cisco Catalyst SD-WAN 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.

After you successfully activate the configuration by issuing a **commit** command, the Cisco Catalyst SD-WAN Controller pushes the data policy to the Cisco IOS XE Catalyst SD-WAN devices located in the specified sites. To view the policy as configured on the Cisco Catalyst SD-WAN Controller, use the **show running-config** command in the Cisco Catalyst SD-WAN 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 Catalyst SD-WAN 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 Catalyst 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 Catalyst SD-WAN devices

You can enable Cflowd visibility directly on Cisco IOS XE Catalyst 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 Examples

This topic shows a complete example of configuring traffic flow monitoring.

Configuration Steps

Enable Cflowd traffic monitoring with a centralized data policy, so all configuration is done on a Cisco Catalyst SD-WAN 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.

```
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
```

Example Configuration

Here is a complete example of a Cflowd configuration:

```
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 192.168.0.1 protocol ipv4 port 13322 transport transport_udp
!
lists
  vpn-list vpn_1
  vpn 1
  !
  site-list cflowd-sites
  site-id 400,500,600
  !
!
!
```

The following sample output from the **show sdwan run policy** command displays the configuration for IPv4 and IPv6 application visibility and flow visibility for Cflowd with SAIE flows:

```
Device# show sdwan run policy
policy
  app-visibility
  app-visibility-ipv6
  flow-visibility
  flow-visibility-ipv6
```

Verify Cflowd Configuration

To verify the Cflowd configuration after activating it on the Cisco Catalyst SD-WAN Controller, use the **show running-config policy** and **show running-config apply-policy** commands.

The following is a sample output from the **show sdwan policy from-vsmart cflowd-template** command:

```
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
  flow-sampling-interval 1
  protocol ipv4/ipv6/both
  customized-ipv4-record-fields
    collect-tos
    collect-dscp-output

collector vpn 1 address 192.0.2.1 protocol ipv4 port 13322 transport transport_udp
```

The following is a sample output from 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
  protocol ipv4/ipv6/both
  template-refresh 90
  customized-ipv4-record-fields
    collect-tos
    collect-dscp-output
  collector vpn 1 address 192.0.2.1 port 13322 transport transport_udp
from-vsmart lists vpn-list vpn_1
  vpn 1
```

Starting from Cisco IOS XE Catalyst SD-WAN Release 17.12.1a, the cflowd commands have been enhanced for both IPv4 and IPv6 flow records.

The following is the sample output from the **show flow record** command where it has been enhanced by the addition of a new field **collect connection initiator** which specifies the direction of flow.

```
Device# show flow record sdwan_flow_record-xxx
```

IPv4 flow record:

```
flow record sdwan_flow_record-1666223692122679:
  Description:      flow and application visibility records
  No. of users:    1
  Total field space: 102 bytes
  Fields:
    match ipv4 protocol
    match ipv4 source address
    match ipv4 destination address
    match transport source-port
    match transport destination-port
    match routing vrf service
    collect ipv4 dscp
```

```

collect transport tcp flags
collect interface input
collect interface output
collect counter bytes long
collect counter packets long
collect timestamp absolute first
collect timestamp absolute last
collect application name
collect flow end-reason
collect connection initiator
collect overlay session id input
collect overlay session id output
collect connection id long
collect drop cause id
collect counter bytes sdwan dropped long
collect sdwan sla-not-met
collect sdwan preferred-color-not-met
collect sdwan qos-queue-id
collect counter packets sdwan dropped long

```

IPv6 flow format:

```

flow record sdwan_flow_record_ipv6-1667963213662363:
  Description:          flow and application visibility records
  No. of users:        1
  Total field space:   125 bytes
  Fields:
    match ipv6 protocol
    match ipv6 source address
    match ipv6 destination address
    match transport source-port
    match transport destination-port
    match routing vrf service
    collect ipv6 dscp
    collect transport tcp flags
    collect interface input
    collect interface output
    collect counter bytes long
    collect counter packets long
    collect timestamp absolute first
    collect timestamp absolute last
    collect application name
    collect flow end-reason
    collect connection initiator
    collect overlay session id input
    collect overlay session id output
    collect connection id long
    collect drop cause id
    collect counter bytes sdwan dropped long
    collect sdwan sla-not-met
    collect sdwan preferred-color-not-met
    collect sdwan qos-queue-id
    collect counter packets sdwan dropped long

```

The following is the enhanced sample output from the **show flow monitor *monitor-name* cache** command where a new field `connection initiator` indicating flow direction has been added in the output. The `connection initiator` field can have one of these values - `initiator` for client to server traffic flow, `reverse` for server to client and `unknown` when the direction of traffic flow is not known.

```

Device# show flow monitor sdwan_flow_monitor cache
Cache type: Normal (Platform cache)
Cache size: 128000
Current entries: 4
High Watermark: 5

```

```

Flows added: 6
Flows aged: 2
- Inactive timeout ( 10 secs) 2
IPV4 SOURCE ADDRESS: 10.20.24.110
IPV4 DESTINATION ADDRESS: 10.20.25.110
TRNS SOURCE PORT: 40254
TRNS DESTINATION PORT: 443
IP VPN ID: 1
IP PROTOCOL: 6
tcp flags: 0x02
interface input: Gi5
interface output: Gi1
counter bytes long: 3966871
counter packets long: 52886
timestamp abs first: 02:07:45.739
timestamp abs last: 02:08:01.840
flow end reason: Not determined
connection initiator: Initiator
interface overlay session id input: 0
interface overlay session id output: 4
connection connection id long: 0xD8F051F000203A22

```

Check the Flows

On the Cisco IOS XE Catalyst 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

```

FNF IPv6 Configuration Example for IPv6 traffic

The following example shows the centralized policy configuration with Cflowd for IPv6 traffic:

```

policy
data-policy _vpn_1_accept_cflowd_vpn_1
vpn-list vpn_1
sequence 102
match
source-ipv6          2001:DB8:0:/32
destination-ipv6    2001:DB8:1:/32
!
action accept
count cflowd_ipv6_1187157291
cflowd
!
!
default-action accept
!
!
cflowd-template cflowd_server
flow-active-timeout 60
flow-inactive-timeout 30
protocol            ipv6
!

```

```

lists
  vpn-list vpn_1
  vpn 1
  site-list vedge1
  site-id 500
!

apply-policy
  site-list vedge1
  data-policy _vpn_1_accept_cflowd_vpn_1 all
  cflowd-template cflowd_server

```

FNF Export Spread Configuration Example

The following example shows the configuration for export spreading:

```

Device# show sdwan policy from-vsmart
from-vsmart cflowd-template cflowd
flow-active-timeout 600
flow-inactive-timeout 60
template-refresh 60
flow-sampling-interval 1
protocol ipv4
customized-ipv4-record-fields
no collect-tos
no collect-dscp-output
collector vpn 0 address 10.0.100.1 port 4739 transport transport_udp
export-spread
  app-tables 20
  tloc-tables 10
  other-tables 5

```

Configure the Maximum FNF Record Rate for Aggregated Data, Using CLI Commands

Minimum supported releases: Cisco IOS XE Catalyst SD-WAN Release 17.14.1a, Cisco Catalyst SD-WAN Control Components Release 20.14.1

Before You Begin

For more information about using CLI templates, see [CLI Add-On Feature Templates](#) and [CLI Templates](#). By default, CLI templates execute commands in global configuration mode.

Configure the Maximum FNF Record Rate

Configure the maximum rate (FNF records per minute) for a device to send aggregated traffic data to Cisco SD-WAN Manager.

```
policy app-agg-node max-records-per-minute
```

Example

The following configures a device to send a maximum of 1000 FNF records per minute of aggregated traffic data.

```
policy app-agg-node 1000
```


Example

The following restores a device to the default value of sending a maximum of 10000 FNF records per minute of aggregated traffic data.

```
no policy app-agg-node
```

Verify Traffic Flow Monitoring

The following sections provide information about verifying traffic flow monitoring.

Verify Collect Loopback

You can verify the ingress and egress interface output using the following command.

show sdwan app-fwd cflowd flows

The following is a sample output from the **show sdwan app-fwd cflowd flows** using the **flows** keyword.

```
Device#show sdwan app-fwd cflowd flows
app-fwd cflowd flows vpn 1 src-ip 10.10.15.12 dest-ip 10.20.15.12 src-port 0 dest-port 0
dscp 0 ip-protocol 1
tcp-ctrl-bits          24
icmp-opcode           0
total-pkts            5
total-bytes           500
start-time            "Tue Jun 27 09:21:09 2023"
egress-intf-name      Loopback1
ingress-intf-name     GigabitEthernet5
application           ping
family                network-service
drop-cause            "No Drop"
drop-octets           0
drop-packets          0
sla-not-met           0
color-not-met         0
queue-id              2
initiator             2
tos                   0
dscp-output           0
sampler-id            0
fec-d-pkts            0
fec-r-pkts            0
pkt-dup-d-pkts-orig   0
pkt-dup-d-pkts-dup    0
pkt-dup-r-pkts        0
pkt-cxp-d-pkts        0
category              0
service-area          0
cxp-path-type         0
region-id             0
ssl-read-bytes        0
ssl-written-bytes     0
ssl-en-read-bytes     0
ssl-en-written-bytes  0
ssl-de-read-bytes     0
ssl-de-written-bytes  0
ssl-service-type      0
ssl-traffic-type      0
ssl-policy-action     0
```

```

appqoe-action          0
appqoe-sn-ip          0.0.0.0
appqoe-pass-reason    0
appqoe-dre-input-bytes 0
appqoe-dre-input-packets 0
appqoe-flags          0
    
```

You can verify the ingress and egress interface output using the following command.

show sdwan app-fwd cflowd table

The following is a sample output from the **show sdwan app-fwd cflowd table** using the **table** keyword.

```

show sdwan app-fwd cflowd flows table
PKT  PKT  PKT  PKT
SSL
SSL
TCP
SLA  COLOR
CXP
SSL  SSL
EN  DRE
DRE
ICMP  TOTAL  TOTAL
DSCP  SAMPLER  D  R  PKTS  PKTS  R  D
READ  WRITTEN  READ  WRITTEN  READ  WRITTEN  SERVICE  TRAFFIC  POLICY  APPQOE  PATH  REGION
PASS  INPUT  INPUT  APPQOE
VPN  SRC  IP  DEST  IP
OPCODE  PKTS  BYTES  START  TIME  NAME  PORT  PORT  DSCP  PROTO  BITS
APPLICATION  FAMILY  DROP  CAUSE  OCTETS  PACKETS  MET  MET  ID  INITIATOR
TOS  OUTPUT  ID  PKTS  PKTS  ORIG  DUP  PKTS  PKTS  CATEGORY  AREA  TYPE  ID
BYTES  BYTES  BYTES  BYTES  BYTES  BYTES  BYTES  TYPE  TYPE  ACTION  ACTION  SN
IP  REASON  BYTES  PACKETS  FLAGS
-----
1  10.10.15.11  10.20.20.10  0  0  0  1  24  0
5  500  Tue Jun 27 09:21:06 2023  Loopback1  GigabitEthernet5 ping
network-service  No Drop  0  0  0  0  0
0  0  0  0  0  0  0  0  0  0  0  0  0.0.0.0
0  0  0  0
0  10.0.5.5  10.0.15.10  58048  22  4  6  24
41  1752  Tue Jun 27 09:21:06 2023  internal0/0/rp:0  GigabitEthernet9 unknown
network-service  No Drop  0  0  0  0  0  0  0  0  0  0  0  0.0.0.0
0  0  0  0  0  0  0  0  0  0  0  0  0.0.0.0
0  0  0  0
1  10.10.15.11  10.20.20.10  0  2048  0  1  24
2048  5  500  Tue Jun 27 09:21:06 2023  GigabitEthernet5  Loopback1 ping
network-service  No Drop  0  0  0  0  0  0  0  0  0  0  0  0.0.0.0
0  0  0  0  0  0  0  0  0  0  0  0  0.0.0.0
0  0  0  0
1  10.10.15.11  10.5.10.15  0  2048  0  1  31
2048  20  960  Tue Jun 27 09:21:06 2023  Null  GigabitEthernet5 ping
network-service  Ipv4NoRoute  960  20  0  0  2  2  0  0  0  0  0.0.0.0
0  0  0  0  0  0  0  0  0  0  0  0  0.0.0.0
0  0  0  0
1  10.10.15.11  10.20.20.10  50920  4739  0  17  31  0
473  524768  Tue Jun 27 09:21:06 2023  GigabitEthernet5  internal0/0/rp:0 ipfix
network-management  No Drop  0  0  0  0  0  0  0  0  0  0  0  0.0.0.0
0  0  0  0  0  0  0  0  0  0  0  0  0.0.0.0
0  0  0  0  0  0  0  0  0  0  0  0  0.0.0.0
    
```

```

0      0      0      0
0  10.0.5.10      10.0.5.10      22      58048      48      6      24
0      39      3020      Tue Jun 27 09:21:05 2023 GigabitEthernet9 internal0/0/rp:0 ssh
      terminal      No Drop      0      0      0      0      0      0      0      0      0
0      0      0      0      0      0      0      0      0      0      0      0.0.0.0
0      0      0      0
1  10.10.15.11      10.20.20.10      0      771      48      1      31
771      8      4192      Tue Jun 27 09:21:05 2023 internal0/0/rp:0 GigabitEthernet5 icmp
      network-service      No Drop      0      0      0      0      0      0      0      0
0      0      0      0      0      0      0      0      0      0      0      0.0.0.0
0      0      0      0
1  fe40::6044:ff:feb7:c2db ff01::1:ff00:10      0      34560      0      58      0
34560      6      432      Tue Jun 27 09:20:41 2023 internal0/0/rp:0 GigabitEthernet5 ipv6-icmp
      network-service      No Drop      0      0      0      0      0      0      0      0
0      0      0      0      0      0      0      0      0      0      0      0.0.0.0
0      0      0      0
1  10:20:20::10      fe40::6024:ff:feb6:c1db      0      34816      56      58      0
34816      4      288      Tue Jun 27 09:20:41 2023 GigabitEthernet5 internal0/0/rp:0 ipv6-icmp
      network-service      No Drop      0      0      0      0      0      0      0      0
0      0      0      0      0      0      0      0      0      0      0      0.0.0.0
0      0      0      0

```

Verify Interface Binding on the Device

You can verify the interface binding on the device using the following command.

show sdwan control local-properties wan-interface-list

The following is a sample output from the **show sdwan control local-properties wan-interface-list** using the **wan-interface-list** keyword.

The command displays:

- The physical interface bound to the loopback WAN interface in bind mode.
- Unbind for loopback WAN interface in unbind mode.
- N/A for any other cases.

```
Device#show sdwan control local-properties wan-interface-list
```

```
NAT TYPE: E -- indicates End-point independent mapping
```

```
A -- indicates Address-port dependent mapping
```

```
N -- indicates Not learned
```

```
Note: Requires minimum two vbonds to learn the NAT type
```

	PRIVATE	PUBLIC	PUBLIC PRIVATE	PRIVATE
MAX RESTRICT/ INTERFACE	LAST IPv4	SPI TIME PORT IPv4	NAT VM STATE	BIND IPv6
CNTRL CONTROL/ PORT	VS/VM COLOR LR/LB CONNECTION	REMAINING	TYPE CON REG	INTERFACE
PRF IDs				STUN
GigabitEthernet1	10.0.10.10	12346 10.0.10.10	::	
0:01:14:20 N	5 Default N/A	up 2	no/yes/no	No/No 0:20:20:27

```

GigabitEthernet4          10.0.10.10      12346 10.0.10.10    ::
      12346 2/0 blue          up 2 no/yes/no No/No 0:20:20:27
0:01:14:20 N 5 Default N/A
Loopback1                 1.1.1.1         12366 1.1.1.1       ::
      12366 2/0 custom1       up 2 no/yes/no No/No 0:20:20:27
0:01:14:20 N 5 Default GigabitEthernet1
Loopback2                 2.2.2.2         12406 2.2.2.2       ::
      12406 2/0 custom2       up 2 no/yes/no No/No 0:20:20:27
0:01:14:20 N 5 Default Unbind

```

Verify Flexible Netflow Configuration on VPN0 Interface

View Flexible Netflow Record Configuration Summary

You can verify FNF record configuration using the following command.

```
Device# show flow record <monitor-context-name>
```



Note The monitor name is used as temp0 in the following examples.

The following sample output displays the information about IPv4 traffic flow records using ezPM profile.

```

Device# show flow record temp0-sdwan-fnf-vpn0-monitor_ipv4
flow record temp0-sdwan-fnf-vpn0-monitor_ipv4:
  Description:          ezPM record
  No. of users:         1
  Total field space:    66 bytes
  Fields:
    match ipv4 dscp
    match ipv4 protocol
    match ipv4 source address
    match ipv4 destination address
    match transport source-port
    match transport destination-port
    match flow direction
    collect routing next-hop address ipv4
    collect transport tcp flags
    collect interface input
    collect interface output
    collect flow sampler
    collect counter bytes long
    collect counter packets long
    collect timestamp absolute first
    collect timestamp absolute last
    collect application name
    collect flow end-reason

```

The following sample output displays the information about IPv6 traffic flow records using ezPM profile.

```

Device# show flow record temp0-sdwan-fnf-vpn0-monitor_ipv6

flow record temp0-sdwan-fnf-vpn0-monitor_ipv6:
  Description:          ezPM record
  No. of users:         1
  Total field space:    102 bytes
  Fields:
    match ipv6 dscp

```

```

match ipv6 protocol
match ipv6 source address
match ipv6 destination address
match transport source-port
match transport destination-port
match flow direction
collect routing next-hop address ipv6
collect transport tcp flags
collect interface input
collect interface output
collect flow sampler
collect counter bytes long
collect counter packets long
collect timestamp absolute first
collect timestamp absolute last
collect application name
collect flow end-reason

```

The following sample output displays the monitor information about IPv4 traffic netflow configuration using ezPM profile.

```

Device# show flow monitor temp0-sdwan-fnf-vpn0-monitor_ipv4
Flow Monitor temp0-sdwan-fnf-vpn0-monitor_ipv4:
  Description:      ezPM monitor
  Flow Record:     temp0-sdwan-fnf-vpn0-monitor_ipv4
  Cache:
    Type:           normal (Platform cache)
    Status:         allocated
    Size:           5000 entries
    Inactive Timeout: 10 secs
    Active Timeout: 60 secs

    Trans end aging:  off

```

The following sample output displays the monitor information about IPv6 traffic netflow configuration using ezPM profile.

```

Device# show flow monitor temp0-sdwan-fnf-vpn0-monitor_ipv6
Flow Monitor temp0-sdwan-fnf-vpn0-monitor_ipv6:
  Description:      ezPM monitor
  Flow Record:     temp0-sdwan-fnf-vpn0-monitor_ipv6
  Cache:
    Type:           normal (Platform cache)
    Status:         allocated
    Size:           5000 entries
    Inactive Timeout: 10 secs
    Active Timeout: 60 secs

    Trans end aging:  off

```

View Flow Record Cache

The following sample output displays flow record cache for the specified monitor, in this case, temp0-sdwan-fnf-vpn0-monitor_ipv4.

```

Device# show flow monitor temp0-sdwan-fnf-vpn0-monitor_ipv4 cache
Cache type:           Normal (Platform cache)
Cache size:          5000
Current entries:     14

```

```

High Watermark:                               14

Flows added:                                  170
Flows aged:                                   156
  - Active timeout      (    60 secs)         156

IPV4 SOURCE ADDRESS:      10.0.0.0
IPV4 DESTINATION ADDRESS: 10.255.255.254
TRNS SOURCE PORT:         0
TRNS DESTINATION PORT:    0
FLOW DIRECTION:           Input
IP DSCP:                   0x00
IP PROTOCOL:              1
ipv4 next hop address:    10.0.0.1
tcp flags:                 0x00
interface input:          Gi1
interface output:         Gi2
flow sampler id:          0
counter bytes long:       840
counter packets long:     10
timestamp abs first:      02:55:24.359
timestamp abs last:       02:55:33.446
flow end reason:         Not determined
application name:         layer7 ping
.....

```

The following sample output displays flow record cache for the specified IPv6 monitor, temp0-sdwan-fnf-vpn0-monitor_ipv6.

```

Device# show flow monitor temp0-sdwan-fnf-vpn0-monitor_ipv6 cache
Cache type:                               Normal (Platform cache)
Cache size:                               5000
Current entries:                           6
High Watermark:                            6

Flows added:                               10
Flows aged:                                 4
  - Inactive timeout      (    10 secs)     4

IPV6 SOURCE ADDRESS:      2001:DB8::/32
IPV6 DESTINATION ADDRESS: 2001:DB8::1
TRNS SOURCE PORT:         0
TRNS DESTINATION PORT:    32768
FLOW DIRECTION:           Output
IP DSCP:                   0x00
IP PROTOCOL:              58
ipv6 next hop address:    2001:DB8:1::1
tcp flags:                 0x00
interface input:          Gi2
interface output:         Gi1
flow sampler id:          0
counter bytes long:       2912
counter packets long:     28
timestamp abs first:      02:57:06.025
timestamp abs last:       02:57:33.378
flow end reason:         Not determined
application name:         prot ipv6-icmp

```

The following sample output displays the flow exporter details.

```

Device# show flow exporter temp0
Flow Exporter temp0:
  Description:           performance monitor context temp0 exporter
  Export protocol:       IPFIX (Version 10)

```

```

Transport Configuration:
  Destination type:      IP
  Destination IP address: 10.0.0.1
  VRF label:            1
  Source IP address:    10.0.0.0
  Source Interface:     GigabitEthernet5
  Transport Protocol:   UDP
  Destination Port:     4739
  Source Port:          51242
  DSCP:                 0x1
  TTL:                  255
  Output Features:      Used
Export template data timeout:      300
Options Configuration:
  interface-table (timeout 300 seconds) (active)
  vrf-table (timeout 300 seconds) (active)
  sampler-table (timeout 300 seconds) (active)
  application-table (timeout 300 seconds) (active)
  application-attributes (timeout 300 seconds) (active)

```

Verify Flexible NetFlow Configuration with Export of BFD Metrics

Minimum releases: Cisco IOS XE Catalyst SD-WAN Release 17.10.1a and Cisco Catalyst SD-WAN Control Components Release 20.10.1

The following sample output from the **show flow exporter** command displays the configuration of each flow exporter:

```

Device# show flow exporter
...
Flow Exporter sdwan_flow_exporter_1:
  Description:          export flow records to collector
  Export protocol:      IPFIX (Version 10)
  Transport Configuration:
    Destination type:   IP
    Destination IP address: 10.0.100.1
    Source IP address:  10.0.100.15
    Transport Protocol:  UDP
    Destination Port:   4739
    Source Port:        54177
    DSCP:                0x0
    TTL:                 255
    MTU:                 1280
    Output Features:     Used
  Options Configuration:
    interface-table (timeout 600 seconds) (active)
    tunnel-tloc-table (timeout 600 seconds) (active)
    bfd-metrics-table (timeout 600 seconds) (active)

```

The following sample output from the **show flow exporter statistics** command displays the client-sent statistics of each flow exporter:

```

Device# show flow exporter statistics
...
Flow Exporter sdwan_flow_exporter_1:
  Packet send statistics (last cleared 3d05h ago):
    Successfully sent:      1433                (907666 bytes)

  Client send statistics:
    Client: Option options interface-table
      Records added:        6552
      - sent:                6552

```

```

Bytes added:          694512
- sent:              694512

Client: Option options tunnel-tloc-table
Records added:       1916
- sent:              1916
Bytes added:         99632
- sent:              99632

Client: Flow Monitor sdwan_flow_monitor
Records added:       0
Bytes added:         0

Client: Option options bfd-metrics-table
Records added:       4
- sent:              4
Bytes added:         196
- sent:              196

```

The following sample output from the **show flow exporter templates** command displays the details for each template:

```
Device# show flow exporter templates
```

```
...
```

```
Client: Option options tunnel-tloc-table
Exporter Format: IPFIX (Version 10)
Template ID      : 257
Source ID       : 6
Record Size     : 52
Template layout

```

Field	ID	Ent.ID	Offset	Size
TLOC TABLE OVERLAY SESSION ID	12435	9	0	4
tloc local color	12437	9	4	16
tloc remote color	12439	9	20	16
tloc tunnel protocol	12440	9	36	8
tloc local system ip address	12436	9	44	4
tloc remote system ip address	12438	9	48	4

```
Client: Option options bfd-metrics-table
Exporter Format: IPFIX (Version 10)
Template ID      : 262
Source ID       : 6
Record Size     : 49
Template layout

```

Field	ID	Ent.ID	Offset	Size
TLOC TABLE OVERLAY SESSION ID	12435	9	0	4
IP DSCP	195		4	1
bfd loss	12527	9	5	4
bfd pfr update ts	12530	9	9	8
bfd avg latency	12528	9	17	8
bfd avg jitter	12529	9	25	8
bfd rx cnt	12531	9	33	8
bfd tx cnt	12532	9	41	8



CHAPTER 13

Application Performance Monitor

Table 38: Feature History

Feature Name	Release Information	Description
Application Performance Monitor	Cisco IOS XE Catalyst SD-WAN Release 17.5.1a Cisco vManage Release 20.5.1	This feature provides an express method for configuring an intent-based performance monitor with the help of predefined monitoring profiles. Configure this feature using the CLI Add-on feature template in Cisco SD-WAN Manager.

- [Overview of Application Performance Monitor, on page 221](#)
- [Limitations and Restrictions, on page 223](#)
- [Configure Application Performance Monitor, on page 224](#)
- [Verify Performance Monitoring Configuration, on page 225](#)

Overview of Application Performance Monitor

The Application Performance Monitor feature is a simplified framework that enables you to configure intent-based performance monitors. With this feature, you can view real-time, end-to-end application performance filtered by client segments, network segments, and server segments. This information helps you optimize application performance.

An application performance monitor is a predefined configuration that is used to collect performance metrics for specific traffic.

Key Concepts in Application Performance Monitoring

Monitoring Profile: A profile is a predefined set of traffic monitors that can be enabled or disabled for a context. As part of this feature, the `sdwan-performance` profile has been enhanced to include Application Response Time (ART) and media monitors to monitor traffic passing through Cisco Catalyst SD-WAN tunnel interfaces. The `sdwan-performance` profile has a dedicated policy to filter traffic based on your intent.

When you choose the `sdwan-performance` profile, the related configuration is generated and applied automatically.

Context: A context represents a performance monitor policy map that is attached to an interface for ingress and egress traffic. A context contains information about a traffic monitor that has to be enabled. When a

context is attached to an interface, two policy-maps are created, one each for ingress and egress traffic. Depending on the direction specified in the traffic monitor, the policy maps are attached in that direction and the traffic is monitored.



Note A context can be attached to multiple interfaces. Only one context can be attached to an interface. You can modify the context only when it is not attached to an interface.

Traffic Monitoring Specifications: You can choose to filter performance metrics using classification and sampler.


- **Classification:** Classification is a filter that defines the traffic that should be monitored for specified applications. This filter reduces the load on the device and performance collectors because they only need to monitor performance for specific applications.
- **Sampler:** A sampler monitors random traffic flows, based on the sampling rate specified, rather than all the flows. Enabling the sampler reduces scaling and performance impact when the scale of traffic is large.

Features and Benefits

- ART can be monitored for TCP flows. Some of the parameters that can be monitored are—server network delay, client network delay, and application delay.
- Jitter can be monitored for Real-time Transport Protocol (RTP) audio and video traffic.
- Information about input and output interfaces and local and remote TLOCs can be collected for every flow that matches the performance monitor.
- Performance monitor can be configured on all WAN tunnel interfaces or specific WAN tunnel interfaces using CLI commands.
- Global performance sampler is supported. The sampler allows you to monitor random flows based on the sampling rate configured, rather than the entire traffic, therefore, reducing performance and scaling overhead.

How Application Performance Monitor Works

Figure 15: Performance Monitoring Workflow



In this image, performance monitor has been applied globally (on all tunnel interfaces). You also have the option to enable it on specific interfaces. Performance is monitored for traffic going out of, and coming into the WAN tunnel interfaces. Based on the exporter parameters defined in the context that is initiated from the monitoring profile, the metrics that are collected are sent to the third-party collector that is defined. You can then view details of the application or media that you are monitoring using various show commands.

Limitations and Restrictions

- Performance monitoring is only supported on IPv4 traffic. IPv6 traffic is not supported.
- Once a performance monitor is applied to a device, the configuration cannot be modified and reapplied to the device. Follow these steps to make any modifications to performance monitor configuration:
 1. Edit the CLI Add-on feature template or device CLI template to remove the **performance monitor apply** command from the template. Update the device CLI template or the device template to which the CLI Add-on feature template is attached.
 2. Edit the **performance monitor context** in the CLI Add-on feature template, and apply the performance monitor again using the **performance monitor apply** command. Update the device template to which the CLI Add-on feature template is attached.

Alternatively, configure a new context based on the same monitoring profile, and remove the previous context configuration.

- App visibility must be enabled in a policy to be able to set the connector initiator value appropriately.

Configure Application Performance Monitor

You can enable application performance monitor globally (on all WAN tunnel interfaces) or on specific WAN tunnel interfaces. You can also enable performance monitoring for ART, or media monitors, or both.

To configure application performance monitoring using Cisco SD-WAN Manager, [create a CLI add-on feature template and attach it to the device template](#).

Enable Performance Monitor Globally

The following example shows how to configure a performance monitor context using the `sdwan-performance` profile. This configuration enables monitoring of traffic metrics for ART and media, and applies the configuration to all SD-WAN tunnel interfaces. Here, 10.0.1.128 is the IP address of the third-party collector, GigabitEthernet9 is the source interface, and 2055 is the listening port of the third-party collector.

```
performance monitor context CISCO-APP-MONITOR profile sdwan-performance
  exporter destination 10.0.1.128 source GigabitEthernet9 port 2055
  traffic-monitor application-response-time
  traffic-monitor media
!
performance monitor apply CISCO-APP-MONITOR sdwan-tunnel
```

Enable Performance Monitor on a Specific Interface

The following example shows how to configure a performance monitor context using the `sdwan-performance` profile. This configuration enables monitoring of traffic metrics for ART and media, and applies it to a specific tunnel interface, in this case, Tunnel1. Here, 10.0.1.128 is the IP address of the third-party collector, GigabitEthernet9 is the source interface, and 2055 is the listening port of the third-party collector.

```
performance monitor context CISCO-APP-MONITOR profile sdwan-performance
  exporter destination 10.0.1.128 source GigabitEthernet9 port 2055
  traffic-monitor application-response-time
  traffic-monitor media
!
interface Tunnel1
  performance monitor context CISCO-APP-MONITOR
```

Specify Additional Monitoring Filters and Sampling Rate

The following example shows how to enable specific type of traffic to be monitored. In this case, the match protocol of `rtp-audio` is defined in the class map named `match-audio`. This class is then referenced in **traffic-monitor media class-and** `match-audio` so that `rtp-audio` traffic is specifically monitored. Alternatively, you can use the keyword **class-and**. In such a case, the customized class map replaces the default class map, which is automatically created when you enable the `sdwan-performance` profile.

In this example, performance monitor is applied globally, which means that it is applied on all Cisco Catalyst SD-WAN tunnel interfaces. The sampling rate of 10 indicates that one in 10 flows is monitored. Sampling rate 100 indicates that one in 100 flows is monitored.

```
class-map match-any match-audio
  match protocol rtp-audio
!
performance monitor context CISCO-APP-MONITOR profile sdwan-performancekeyword
  exporter destination 10.75.212.84 source GigabitEthernet0/0/0 port 2055
```

```

traffic-monitor application-response-time
traffic-monitor media class-and (or class-replace) match-audio
!
performance monitor apply CISCO-APP-MONITOR sdwan-tunnel
performance monitor sampling-rate 10

```

Verify Performance Monitoring Configuration

View Performance Monitor Configuration Summary

The following sample out displays the information about traffic monitors that are enabled and the interfaces to which they are applied.

```
Device# show performance monitor context CISCO-MONITOR summary
```

```

=====
|                               CISCO-MONITOR                               |
=====
Description: User defined

Based on profile: sdwan-performance

Coarse-grain NBAR based profile

Configured traffic monitors
=====
application-response-time:
media: class-and match_audio

Attached to Interfaces
=====

Tunnell

```

The following sample out displays operational information about the third-party exporters that are attached to the specified context.

```
Device# show performance monitor context CISCO-MONITOR exporter
```

```

=====
|                               Exporters information of context CISCO-MONITOR                               |
=====

```

```

Flow Exporter 175_SDWAN-1:
  Description:                performance monitor context CISCO-MONITOR exporter
  Export protocol:            IPFIX (Version 10)
  Transport Configuration:
    Destination type:        IP
    Destination IP address:  10.75.212.84
    Source IP address:       10.74.28.19
    Source Interface:        GigabitEthernet0/0/0
    Transport Protocol:      UDP
    Destination Port:        2055
    Source Port:             63494
    DSCP:                    0x0
    TTL:                     255
    Output Features:         Used
  Options Configuration:
    interface-table (timeout 600 seconds) (active)
    sampler-table (timeout 600 seconds) (active)
    application-table (timeout 600 seconds) (active)
    sub-application-table (timeout 600 seconds) (active)
    application-attributes (timeout 600 seconds) (active)
    tunnel-tloc-table (timeout 600 seconds) (active)
Flow Exporter 175_SDWAN-1:
  Packet send statistics (last cleared 04:13:19 ago):
    Successfully sent:        10270                (13709142 bytes)

  Client send statistics:
    Client: Option options interface-table
      Records added:          312
        - sent:              312
      Bytes added:           31824
        - sent:              31824

```

Client: Option options sampler-table

```
Records added:          28
- sent:                 28
Bytes added:            1344
- sent:                 1344
```

Client: Option options application-name

```
Records added:          38766
- sent:                 38766
Bytes added:            3217578
- sent:                 3217578
```

Client: Option sub-application-table

```
Records added:          858
- sent:                 858
Bytes added:            144144
- sent:                 144144
```

Client: Option options application-attributes

```
Records added:          38038
- sent:                 38038
Bytes added:            9813804
- sent:                 9813804
```

Client: Option options tunnel-tloc-table

```
Records added:          26
- sent:                 26
Bytes added:            1352
- sent:                 1352
```

Client: MMA EXPORTER GROUP MMA-EXP-1

```
Records added:          0
```

```
Bytes added:          0
```

```
Client: Flow Monitor 175_SDWAN-art_ipv4
```

```
Records added:       0
```

```
Bytes added:         0
```

For more information, see the [show performance monitor context](#) command page.

View Flow Record Cache

The following sample output displays flow record cache for the specified monitor, in this case, CISCO-MONITOR-art_ipv4 .

```
Device# show performance monitor cache
```

```
Monitor: CISCO-MONITOR
```

```
Data Collection Monitor:
```

```
Cache type:          Synchronized (Platform cache)
Cache size:          4000
Current entries:     0

Flows added:        0
Flows aged:         0
Synchronized timeout (secs): 60
```

```
Monitor: CISCO-MONITOR-art_ipv4
```

```
Data Collection Monitor:
```

```
Cache type:          Synchronized (Platform cache)
Cache size:          11250
```



```

Current entries:                0

Flows added:                    0

Flows aged:                    0

Synchronized timeout (secs):   60

```

For more information, see the [show performance monitor cache](#) command page.

View Performance Monitor Templates

The following sample output displays flow exporter template information for the specified monitor.

```
Device# show flow exporter CISCO-MONITOR templates
```

```
Flow Exporter CISCO-MONITOR:
```

```
Client: Option options sampler-table
```

```
Exporter Format: IPFIX (Version 10)
```

```
Template ID      : 257
```

```
Source ID       : 6
```

```
Record Size     : 48
```

```
Template layout
```

Field	ID	Ent.ID	Offset	Size
FLOW SAMPLER	48		0	4
flow sampler name	84		4	41
flow sampler algorithm export	49		45	1
flow sampler interval	50		46	2

```
Client: Option options application-name
```

```
Exporter Format: IPFIX (Version 10)
```

```
Template ID      : 258
```

```
Source ID       : 6
```

```
Record Size     : 83
```

```
Template layout
```

Field	ID	Ent.ID	Offset	Size
APPLICATION ID	95		0	4
application name	96		4	24
application description	94		28	55

Client: Option sub-application-table

Exporter Format: IPFIX (Version 10)

Template ID : 259

Source ID : 6

Record Size : 168

Template layout

Field	ID	Ent.ID	Offset	Size
APPLICATION ID	95		0	4
SUB APPLICATION TAG	97		4	4
sub application name	109		8	80
sub application description	110		88	80

Client: Option options application-attributes

Exporter Format: IPFIX (Version 10)

Template ID : 260

Source ID : 6

Record Size : 258

Template layout

Field	ID	Ent.ID	Offset	Size
-------	----	--------	--------	------

APPLICATION ID	95	0	4
application category name	12232	9	32
application sub category name	12233	9	32
application group name	12234	9	32
application traffic-class	12243	9	32
application business-relevance	12244	9	32
p2p technology	288	164	10
tunnel technology	289	174	10
encrypted technology	290	184	10
application set name	12231	9	32
application family name	12230	9	32

Client: Option options tunnel-tloc-table

Exporter Format: IPFIX (Version 10)

Template ID : 261

Source ID : 6

Record Size : 52

Template layout

Field	ID	Ent.ID	Offset	Size
TLOC TABLE OVERLAY SESSION ID	12435	9	0	4
tloc local color	12437	9	4	16
tloc remote color	12439	9	20	16
tloc tunnel protocol	12440	9	36	8
tloc local system ip address	12436	9	44	4
tloc remote system ip address	12438	9	48	4

Client: Flow Monitor CISCO-MONITOR-art_ipv4

Exporter Format: IPFIX (Version 10)

Template ID : 0

Verify Performance Monitoring Configuration

Source ID : 0
 Record Size : 208
 Template layout

Field	ID	Ent.ID	Offset	Size
interface input snmp	10		0	4
connection client ipv4 address	12236	9	4	4
connection server ipv4 address	12237	9	8	4
ip dscp	195		12	1
ip protocol	4		13	1
ip ttl	192		14	1
connection server transport port	12241	9	15	2
connection initiator	239		17	1
timestamp absolute monitoring-interval	359		18	8
flow observation point	138		26	8
overlay session id input	12432	9	34	4
routing vrf service	12434	9	38	4
application id	95		42	4
interface output snmp	14		46	4
flow direction	61		50	1
flow sampler	48		51	1
overlay session id output	12433	9	52	4
timestamp absolute first	152		56	8
timestamp absolute last	153		64	8
connection new-connections	278		72	4
connection sum-duration	279		76	8
connection server counter bytes long	232		84	8
connection server counter packets long	299		92	8
connection client counter bytes long	231		100	8
connection client counter packets long	298		108	8
connection server counter bytes network	8337	9	116	8

connection client counter bytes network	8338	9	124	8	
connection delay response to-server sum	9303	9	132	4	
connection server counter responses	9292	9	136	4	
connection delay response to-server his	9300	9	140	4	
connection client counter packets retra	9268	9	144	4	
connection delay application sum	9306	9	148	4	
connection delay response client-to-ser	9309	9	152	4	
connection transaction duration sum	9273	9	156	4	
connection transaction duration min	9275	9	160	4	
connection transaction duration max	9274	9	164	4	
connection transaction counter complete	9272	9	168	4	
connection client counter bytes retrans	9267	9	172	4	
connection server counter bytes retrans	9269	9	176	4	
connection server counter packets retra	9270	9	180	4	
connection delay network long-lived to-	9255	9	184	4	
connection delay network to-client num-	9259	9	188	4	
connection delay network long-lived to-	9254	9	192	4	
connection delay network to-server num-	9258	9	196	4	
connection delay network long-lived cli	9256	9	200	4	
connection delay network client-to-serv	9257	9	204	4	

Client: Flow Monitor CISCO-MONITOR-media_ipv4

Exporter Format: IPFIX (Version 10)

Template ID : 0

Source ID : 0

Record Size : 180

Template layout

Field	ID	Ent.ID	Offset	Size
ipv4 source address	8		0	4
ipv4 destination address	12		4	4

interface input snmp	10	8	4
ip dscp	195	12	1
ip protocol	4	13	1
ip ttl	192	14	1
ipv6 source address	27	15	16
ipv6 destination address	28	31	16
transport source-port	7	47	2
transport destination-port	11	49	2
connection initiator	239	51	1
timestamp absolute monitoring-interval	359	52	8
flow observation point	138	60	8
overlay session id input	12432	9	68
routing vrf service	12434	9	72
application id	95	76	4
routing forwarding-status	89	80	1
interface output snmp	14	81	4
flow direction	61	85	1
flow sampler	48	86	1
overlay session id output	12433	9	87
transport rtp ssrc	4254	9	91
transport rtp payload-type	4273	9	95
counter bytes long	1	96	8
counter packets	2	104	4
timestamp absolute first	152	108	8
timestamp absolute last	153	116	8
connection new-connections	278	124	4
transport packets expected counter	4246	9	128
transport packets lost counter	4251	9	132
transport packets lost rate	4253	9	136
transport rtp jitter mean	4255	9	140
transport rtp jitter minimum	4256	9	144
transport rtp jitter maximum	4257	9	148

```
| counter bytes rate          | 4235 | 9 | 152 | 4 |
| application media bytes counter | 4236 | 9 | 156 | 4 |
| application media bytes rate   | 4238 | 9 | 160 | 4 |
| application media packets counter | 4239 | 9 | 164 | 4 |
| application media packets rate  | 4241 | 9 | 168 | 4 |
| transport rtp jitter mean sum   | 4325 | 9 | 172 | 8 |
```

For more information, see the [show flow exporter](#) command page.



CHAPTER 14

Enhanced Policy Based Routing

Table 39: Feature History

Feature Name	Release Information	Description
Enhanced Policy Based Routing for Cisco Catalyst SD-WAN	Cisco IOS XE Catalyst SD-WAN Release 17.4.1a Cisco vManage Release 20.4.1	This release extends Enhanced Policy Based Routing (ePBR) to Cisco Catalyst SD-WAN. ePBR is a protocol-independent traffic-steering mechanism that routes traffic based on flexible policies for traffic flows. You can create ePBR policies using CLI add-on templates in Cisco SD-WAN Manager.

- [Overview of ePBR, on page 237](#)
- [Configure ePBR, on page 239](#)
- [Monitor ePBR, on page 242](#)

Overview of ePBR

Enhanced Policy Based Routing (ePBR) is an advanced version of Policy Based Routing (PBR). With this feature, traffic forwarding is based on policies rather than routing tables, and gives you more control over routing. ePBR extends and complements the existing mechanisms provided by routing protocols. ePBR is an advanced local data policy that routes traffic based on flexible match criteria such as IPv4 and IPv6 addresses, port numbers, protocols, or packet size.

ePBR matches traffic using flexible Cisco Common Classification Policy Language (C3PL language). It supports matching prefixes, applications, Differentiated Services Code Point (DSCP), Security Group Tags (SGT), and so on. With ePBR, based on match conditions, you can configure a single or multiple next hops for traffic forwarding. You also have the option to configure Internet Protocol Service Level Agreement (IP SLA) tracking. If a configured next hop is unavailable, traffic is routed to the next available hop through dynamic probing enabled by the IP SLA tracker.

Features and Benefits

- Supports both IPv4 and IPv6.
- Supports multiple next hops; and if the next hop isn't reachable, ePBR automatically switches to the next available hop.


- You have the option to configure IP SLA tracking. If this is configured, the next hop is selected only when the IP SLA probe is successful.
SLA probes can be configured in the same or a different VRF.
- If the current hop isn't reachable, syslog messages are generated and the user is notified of the same.

How ePBR Works

- ePBR is applicable to unicast routing only and is based on traffic matching using C3PL.
- All packets received on an ePBR-enabled interface are passed through policy maps. The policy maps used by ePBR dictate the policy, determining where to forward packets.
- ePBR policies are based on a classification criteria (match) and an action criteria (set) that are applied to traffic flow.
- To enable ePBR, you must create a policy map that specifies the packet match criteria and desired policy-route action. Then you associate the policy map on the required interface.
- The match criteria is specified in a class. The policy map then calls the class and takes action based on the set statement.
- The set statements in ePBR policies define the route in terms of next hops, DSCP, VRFs, and so on.

Usage Example

Figure 16: Traffic Redirection with ePBR



This example shows that traffic is coming into VPN 1 interface. Based on the classification configured on VPN 1, the traffic overrides the regular route forwarding and is redirected to a next-hop in VPN 100, where additional network services are applied to the incoming traffic. Network services, such as WAN optimization,

are then applied on the redirected traffic before it is forwarded to the Cisco Catalyst SD-WAN overlay network through VPN 0.

Configure ePBR

To configure ePBR using Cisco SD-WAN Manager, [create a CLI add-on feature template and attach it to the device template](#).

This section provides examples of ePBR configurations that you can add to the CLI add-on template.

Configure ePBR for IPv4

In the following example:

- The extended ACLs define the network or the host.
- Class maps match the parameters in the ACLs.
- Policy maps with ePBR then take detailed actions based on the set statements configured.
- Multiple next-hops are configured. ePBR chooses the first available next-hop.

```
ip access-list extended test300
 100 permit ip any 192.0.2.1 0.0.0.255
ip access-list extended test100
 100 permit ip any 192.0.2.20 0.0.0.255
!
class-map match-any test300
 match access-group name test300
class-map match-any test100
 match access-group name test1
!
policy-map type eubr test300
 class test300
  set ipv4 vrf 300 next-hop 10.0.0.2 10.0.40.1 10.0.50.1 ...
policy-map type eubr test100
 class test100
  set ipv4 vrf 100 next-hop 10.10.0.2 10.20.20.2 10.30.30.2 ...
!
interface GigabitEthernet0/0/1
 service-policy type eubr input test300
interface GigabitEthernet0/0/2
 service-policy type eubr input test100
```

Configure IPv4 Tracking

This example shows how to configure ePBR along with tracking. In the example:

- IP SLA operations of type ICMP Echo are configured and ACLs are defined.
- Class maps are then used to match parameters in the ACLs and the policy map takes action based on the set statements configured.
- The number 10 in `set ipv4 vrf 300 next-hop verify-availability 10.10.0.2 10 track 2` represents the sequence number.

```
ip sla 1
```

```

    icmp-echo 10.0.0.2
    vrf 100
ip sla schedule 1 life forever start-time now
track 1 ip sla 1 state
ip sla 2
    icmp-echo 10.10.0.2
    vrf 300
ip sla schedule 2 life forever start-time now
track 2 ip sla 2 state
ip access-list extended test300
 100 permit ip any 10.10.0.2 0.0.0.255
ip access-list extended test100
 100 permit ip any 10.10.0.3 0.0.0.255
class-map match-any test300
 match access-group name test300
class-map match-any test100
 match access-group name test100
policy-map type epbr test300
 class test300
   set ipv4 vrf 300 next-hop verify-availability 10.10.0.2 10 track 2
policy-map type epbr test100
 class test100
   set ipv4 vrf 100 next-hop verify-availability 10.0.0.2 10 track 1
!
interface GigabitEthernet0/0/1
 service-policy type epbr input test300
interface GigabitEthernet0/0/2
 service-policy type epbr input test100

```

Configure ePBR for IPv6

In the following example:

- The extended ACLs define the network or the host.
- Class maps are used to match the parameters in the ACLs.
- Policy maps with ePBR then take detailed actions based on the set statements configured. .
- Single or multiple next-hop addresses can be configured. ePBR selects the first available next-hop address

```

ipv6 access-list test300_v6
 sequence 100 permit ipv6 any 2001:DB81::/32
ipv6 access-list test100_v6
 sequence 100 permit ipv6 any 2001:DB82::/32
!
class-map match-any test300_v6
 match access-group name test300_v6
class-map match-any test100_v6
 match access-group name test100_v6
policy-map type epbr test300_v6
 class test300_v6
   set ipv6 vrf 300 next-hop 2001:DB8::1
policy-map type epbr test100_v6
 class test100_v6
   set ipv6 vrf 100 next-hop 2001:DB8::2 2001:DB8:FFFF:2 ...
!
interface GigabitEthernet0/0/1
 service-policy type epbr input test300_v6
interface GigabitEthernet0/0/2
 service-policy type epbr input test100_v6

```

Configure IPv6 Tracking

This example shows how to configure ePBR for IPv6 along with tracking enabled. In this example:

- IP SLA operations of type ICMP Echo are configured and ACLs are defined.
- Class maps are then used to match parameters in the ACLs and the policy map takes action based on the set statements configured.
- Tracking is configured such that if the result of the IP SLA is unavailable, the packets aren't sent to the next-hop configured on the class.

```
ip sla 3
  icmp-echo 2001:DB8::1
  vrf 100
ip sla schedule 3 life forever start-time now
track 3 ip sla 3 state
ip sla 4
  icmp-echo 2001:DB8::2
  vrf 300
ip sla schedule 4 life forever start-time now
track 4 ip sla 4 state
ipv6 access-list test300_v6
  sequence 100 permit ipv6 any 2001:DB8::/32
ipv6 access-list test100_v6
  sequence 100 permit ipv6 any 2001:DB8::1/32
class-map match-any test300_v6
  match access-group name test300_v6
class-map match-any test100_v6
  match access-group name test100_v6
policy-map type epbr test300_v6
  class test300_v6
    set ipv6 vrf 300 next-hop verify-availability 2001:DB8::2 10 track 4
policy-map type epbr test100_v6
  class test100_v6
    set ipv6 vrf 100 next-hop verify-availability 2001:DB8::1 10 track 3
interface GigabitEthernet0/0/1
  service-policy type epbr input test300_v6
interface GigabitEthernet0/0/2
  service-policy type epbr input test100_v6
```

Configure ePBR for IPv4 with Multiple Next Hops and SLA Tracking

In the following example:

- IP SLA operations of type ICMP Echo are configured and ACLs are defined.
- Class maps are then used to match parameters in the ACLs and the policy map takes action based on the set statements configured.
- Tracking is configured for next hops such that if the previous IP address isn't reachable, and the IP SLA confirms the next hop as reachable, packets flow to the next hop address.

```
ip sla 1
  icmp-echo 10.0.0.2
  vrf 100
ip sla schedule 1 life forever start-time now
track 1 ip sla 1 state
ip sla 2
  icmp-echo 10.10.0.2
  vrf 300
ip sla schedule 2 life forever start-time now
```

```

track 2 ip sla 2 state
ip sla 3
  icmp-echo 10.20.0.2
  vrf 400
ip sla schedule 3 life forever start-time now
track 3 ip sla 3 state
ip access-list extended test300
  100 permit ip any 192.0.2.1 255.255.255.0
ip access-list extended test100
  100 permit ip any 192.0.2.10 255.255.255.0
!
class-map match-any test300
  match access-group name test300
class-map match-any test100
  match access-group name test100
!
policy-map type epbr test300
  class test300
    set ipv4 vrf 300 next-hop verify-availability 10.10.0.2 10 track 2
    set ipv4 vrf 400 next-hop verify-availability 10.20.0.2 11 track 3
policy-map type epbr test100
  class test100
    set ipv4 vrf 100 next-hop verify-availability 10.0.0.2 10 track 1
!
interface GigabitEthernet0/0/1
  service-policy type epbr input test300
interface GigabitEthernet0/0/2
  service-policy type epbr input test100
!

```



Note When next hops are configured along with the tracker, if the next hop is unreachable or if the IP SLA fails, the next available hop is selected. This means that when the tracker is configured, both next hop availability and IP SLA results are checked.

Monitor ePBR

ePBR can't be monitored through Cisco SD-WAN Manager. To verify your configuration or monitor ePBR statistics, use the show commands described below.

Verify Availability of Next Hop

The following is sample output from the **show platform software epbr track** command.

```

Device# show platform software epbr track
Track Object:
obj num:2:
  track:0x7F94B4376760
  seq:10, nhop:123.0.0.2, nhop_reachable:1, track_handle:0x7F94AFDAE240,
  global:0, vrf_name:300, track_reachable:1
  parent:0x7F94B4383778, oce:0x7F94B81193A8
obj num:1:
  track:0x7F94B8187810
  seq:10, nhop:100.0.0.2, nhop_reachable:1, track_handle:0x7F94AFDAE1D0,
  global:0, vrf_name:100, track_reachable:1
  parent:0x7F94B8187778, oce:0x7F94B81188B8

```

In this example, `nhop_reachable` has the value 1, which indicates that the next hop is reachable. `track_reachable` represents the result of SLA probe and has the value 1, which indicates that the next hop is reachable. If the next hop isn't reachable, the value would be 0 for these parameters.

View Next Hop Configuration

Use the `show platform software epbr R0 feature-object redirect` to view the next hop configuration.



Note To be able to view this output, you must have tracker configured.

```
Device# show platform software epbr r0 feature-object redirect
FMAN EPBR Redirect Feature Objectep

Feature Object ID: 9876543211
  Flags: 0x3
  Table ID: 0x4
  Next-hop: 10.10.10.2
  P2P ADJ-ID: 0

Feature Object ID: 1234567890
  Flags: 0x3
  Table ID: 0x2
  Next-hop: 172.16.0.0
  P2P ADJ-ID: 0
```




CHAPTER 15

Forward Error Correction

Table 40: Feature History

Feature Name	Release Information	Description
Forward Error Correction	Cisco IOS XE SD-WAN Release 16.11.x	Feature introduced. FEC is a mechanism to recover lost packets on a link by sending extra “parity” packets for every group of 4 packets.

Forward Error Correction (FEC) is a mechanism to recover lost packets on a link by sending extra “parity” packets for every group of 4 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 Catalyst SD-WAN devices.



Note We recommend Cisco IOS XE Release 17.6.3 as the minimum release when using FEC.

- [Supported Devices for Forward Error Correction, on page 245](#)
- [Configure Forward Error Correction for a Policy, on page 245](#)
- [Monitor Forward Error Correction Tunnel Information, on page 246](#)
- [Monitor Forward Error Application Family Information, on page 247](#)
- [Monitor Forward Error Correction Status Using the CLI, on page 247](#)

Supported Devices for Forward Error Correction

The forward error correction is supported on all the Cisco IOS XE Catalyst SD-WAN devices.

Configure Forward Error Correction for a Policy

- Step 1** From the Cisco SD-WAN Manager menu, choose **Configuration > Policies**.
- Step 2** Click **Centralized Policy** and then click **Add Policy**.
- Step 3** Click **Next**.

- Step 4** Click **Next** again and then click **Configure Traffic Rules**.
- Step 5** Click **Traffic Data**, and from the **Add Policy** drop-down list, choose **Create New**.
- Step 6** Click **Sequence Type**.
- Step 7** From the **Add Data Policy** pop-up menu, choose **QoS**.
- Step 8** Click **Sequence Rule**.
- Step 9** In the **Applications/Application Family List**, choose one or more applications or lists.
- Step 10** Click **Accept**.
- Step 11** Click **Actions** and click **Loss Correction**.
- Step 12** In the **Actions** area, choose one of the following:
- **FEC Adaptive**: Only send FEC information when the loss detected by the system exceeds the packet loss threshold.
 - **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 13** Click **Save Match and Actions**.
- Step 14** Click **Save Data Policy**.
- Step 15** Click **Next** and take these actions to create a centralized policy:
- a) Enter a **Name** and a **Description**.
 - b) Select **Traffic Data Policy**.
 - c) Choose VPNs and a site list for the policy.
 - d) Save the policy.

Monitor Forward Error Correction Tunnel Information

- Step 1** From the Cisco SD-WAN Manager menu, choose **Monitor > Devices**.
Cisco vManage Release 20.6.x and earlier: From the Cisco SD-WAN Manager menu, choose **Monitor > Network**.
- Step 2** Choose 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 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
 - 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 From the Cisco SD-WAN Manager menu, choose **Monitor > Devices**.

Cisco Catalyst SD-WAN Control Components Release 20.6.1 and earlier: From the Cisco SD-WAN Manager menu, choose **Monitor > Network**.

Step 2 Choose a device group.

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

Note In Cisco Catalyst SD-WAN Control Components Release 20.7.1 and earlier releases, **SAIE Applications** is called **DPI Applications**.

The FEC Recovery Rate application information includes the following:

- A graph for which you can choose the following perspective:
 - Application Usage—Usage of various types of traffic for the selected application families, in KB.
- 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 SD-WAN Manager tooltip for the **Packet Delivery Performance** column.

- Traffic usage, in KB, MB, or GB for the selected application family.
-

Monitor Forward Error Correction Status Using the CLI

Use the **show sdwan tunnel statistics fec** command to verify the FEC status on a Cisco IOS XE Catalyst SD-WAN device:

```

Device# show sdwan tunnel statistics fec
tunnel stats ipsec 80.80.10.19 80.80.10.25 12346 12366
fec-rx-data-pkts      0
fec-rx-parity-pkts   0
fec-tx-data-pkts     0
fec-tx-parity-pkts   0
fec-reconstruct-pkts 0
fec-capable          true
fec-dynamic          false
tunnel stats ipsec 80.80.10.19 80.80.10.50 12346 12346
fec-rx-data-pkts     122314
fec-rx-parity-pkts   30578
fec-tx-data-pkts     125868
fec-tx-parity-pkts   31467
fec-reconstruct-pkts 3
fec-capable          true
fec-dynamic          false

```

The following table describes the FEC counters related to the output shown in the **show sdwan tunnel statistics fec** command:

Name of Counter	Description
fec-rx-data-pkts	Displays the number of data packets received by the device.
fec-rx-parity-pkts	Displays the number of parity packets received by the device.
fec-tx-data-pkts	Displays the number of data packets sent by the device.
fec-tx-parity-pkts	Displays the number of parity packets sent by the device.
fec-reconstruct-pkts	Displays the number of received packets reconstructed by the device.



CHAPTER 16

Packet Duplication

Table 41: Feature History

Feature Name	Release Information	Description
Packet Duplication for Noisy Channels	Cisco IOS XE Catalyst SD-WAN Release 16.12.1b	This feature helps mitigate packet loss over noisy channels, thereby maintaining high application QoE for voice and video.
Packet Duplication for Large Packets Using Underlay Fragmentation	Cisco IOS XE Catalyst SD-WAN Release 17.15.1a Cisco Catalyst SD-WAN Control Components Release 20.15.1	This feature enables packet duplication even when the packet size is greater than the path maximum transmission unit (PMTU) discovered on the duplicate tunnel. With the help of underlay fragmentation, this feature uses adjacency MTU instead of tunnel PMTU to provide this capability. Cisco SD-WAN Manager provides a chart for viewing packet duplication information for tunnels.

- [Information about Packet Duplication, on page 249](#)
- [Configure Packet Duplication Using Centralized Policy, on page 250](#)
- [Configure Packet Duplication Using Policy Groups, on page 251](#)
- [Configure Underlay Fragmentation Using Cisco SD-WAN Manager, on page 252](#)
- [Restrictions for Packet Duplication, on page 252](#)
- [Monitor Packet Duplication Statistics for a Device, on page 252](#)
- [Monitor Tunnel Information for a Device , on page 252](#)

Information about Packet Duplication

Cisco IOS XE Catalyst 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 Catalyst SD-WAN devices. If one of the packets is lost, a copy of the packet is forwarded to the server. Receiving Cisco IOS XE Catalyst 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.

Packet Duplication for Large Packets Using Underlay Fragmentation

When packets are intercepted for duplication, the system queries the IP database using the incoming tunnel ID. It then fetches the duplicate tunnel object. The system compares the packet length with the path maximum transmission unit (PMTU) of the duplicate tunnel. If the packet length is smaller than the duplicate tunnel's PMTU, the packets are duplicated.

Starting from Cisco IOS XE Catalyst SD-WAN Release 17.15.1a, packet duplication with underlay fragmentation ensures that packets get duplicated even when the packet size is more than the PMTU of duplicate tunnel.

For more information on how to enable underlay fragmentation, see [VFR and Underlay Fragmentation](#).

To monitor packet duplication statistics, see [View Loss Percentage, Latency, Jitter, Octet, and Packet Duplication Information for Tunnels](#).

Supported Traffic

Cisco IOS XE Catalyst SD-WAN Devices support packet duplication for the following traffic types:

From Cisco IOS XE Catalyst SD-WAN Release 16.12.1b:

IPv4 traffic over IPv4 tunnel

From Cisco IOS XE Catalyst SD-WAN Release 17.15.1a:

- IPv4 traffic over IPv6 tunnel
- IPv6 traffic over IPv4 tunnel
- IPv6 traffic over IPv6 tunnel

Configure Packet Duplication Using Centralized Policy

1. From the Cisco SD-WAN Manager menu, choose **Configuration > Policies**.
2. Select **Centralized Policy** at the top of the page and then click **Add Policy**.
3. Click **Next** twice to select **Configure Traffic Rules**.
4. Select **Traffic Data**, and from the **Add Policy** drop-down list, click **Create New**.
5. Click **Sequence Type** in the left pane.
6. From the **Add Data Policy** pop-up window, select **QoS**.
7. Click **Sequence Rule**.
8. In the **Applications/Application Family List/Data Prefix**, select one or more applications or lists.
9. Click **Actions** and choose **Loss Correction**.
10. In the Actions area, select the **Packet 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.
11. Click **Save Match and Actions**.
 12. Click **Save Data Policy**.
 13. Click **Next** and take these actions to create a Centralized Policy:
 - a. Enter a Name and a Description.
 - b. Select **Traffic Data Policy**.
 - c. Choose **VPNs/site list** for the policy.
 - d. Save the policy.

Configure Packet Duplication Using Policy Groups

Minimum supported release: Cisco Catalyst SD-WAN Control Components Release 20.14.1

1. Select **Configuration > Policy Groups**.
2. Click **Application Priority & SLA**.
3. Click **Add Application Priority & SLA Policy**. Provide a policy name and description.
4. Enable **Advanced Layout** in the top right pane.
5. Click **Add Traffic Policy**.
6. Enter a name for the policy and specify VPNs.
7. In the **Direction** drop-down list, select **All**.
8. In the **Default Action**, click **Accept**.
9. Click **Add**.
10. Click **Add Rules**.
11. Click **Match**. Select appropriate match condition.
12. Click **Action > Loss Correction**.
13. In the **Type** drop-down list, choose **Packet Duplication**.
14. Click **Save Match and Actions**.
15. Click **Save Policy**.

Configure Underlay Fragmentation Using Cisco SD-WAN Manager

Minimum supported releases: Cisco IOS XE Catalyst SD-WAN Release 17.15.1a

1. From the Cisco SD-WAN menu, choose **Configuration > Configuration Groups**.
2. Click **Transport & Management Profile**.
3. Select the desired transport profile and click **Edit**.
4. Click **Edit Ethernet Interface > Tunnel**.
5. Enable **Allow Fragmentation** and **MTU To Max**.
6. Click **Save**.

Restrictions for Packet Duplication

- Packet duplication interop, forward error correction (FEC), and TCP optimization on Cisco IOS XE Catalyst SD-WAN devices is not supported between Cisco IOS XE Release 16.x and Cisco IOS XE Catalyst SD-WAN Release 17.x versions.
- Packet duplication cannot work in conjunction with local or remote TLOC in the policy. Data policy or AAR is not configured when specifying the packet duplicated tunnel.

Monitor Packet Duplication Statistics for a Device

1. From the Cisco SD-WAN Manager menu, choose **Monitor > Devices**.
2. Choose a device.
3. For a device, in the **Action** column, click ... and choose **Real Time**.
4. In the **Device Options** drop-down menu, click **Tunnel Packet Duplication Statistics**.

Monitor Tunnel Information for a Device

1. From the Cisco SD-WAN Manager menu, choose **Monitor > Devices**.
2. Click a device name.
3. In the left pane, click **Tunnel** in the WAN area.

The right pane displays information about tunnel connection information, including loss percentage, latency, jitter, octets, and packet duplication.
4. In the right pane, click **Chart Options** to choose the format in which you want to view the information.



CHAPTER 17

Policy Configuration Tagging

Table 42: Feature History

Feature Name	Release Information	Description
Support for Cisco Catalyst SD-WAN Policy Configuration Tagging Using the Cisco Catalyst SD-WAN Controller CLI Template	Cisco IOS XE Catalyst SD-WAN Release 17.9.1a Cisco vManage Release 20.9.1	<p>This feature allows you to group multiple policy objects under a tag. The tag mechanism, when used in Cisco Catalyst SD-WAN centralized or localized policies, provides the following functionalities:</p> <ul style="list-style-type: none"> • Controls the download speed of a policy between the Cisco Catalyst SD-WAN Controller and the Cisco IOS XE Catalyst SD-WAN devices. • Improves management of defined lists in the Cisco Catalyst SD-WAN Controller. • Better organizes the configurations of the intent-based network.

- [Supported Devices for Policy Configuration Tagging, on page 254](#)
- [Restrictions for Policy Configuration Tagging, on page 254](#)
- [Information About Policy Configuration Tagging, on page 254](#)
- [Benefits of Policy Configuration Tagging, on page 257](#)
- [Configure Policy Configuration Tagging Using a CLI Template, on page 258](#)
- [Verify Tag-Instances Configuration Using the CLI, on page 260](#)

Supported Devices for Policy Configuration Tagging

Table 43: Supported Devices and Releases

Release	Supported Devices
Cisco IOS XE Catalyst SD-WAN Release 17.9.1a and later	<ul style="list-style-type: none"> • Cisco Catalyst 8500 Series Edge Platforms • Cisco Catalyst 8300 Series Edge Platforms • Cisco Catalyst 8200 Series Edge Platforms • Cisco Catalyst 8200 uCPE Series Edge Platforms • Cisco ASR 1000 Series Aggregation Services Routers • Cisco ISR 1000 and ISR 4000 Series Integrated Services Routers (ISRs) • Cisco ISR 1100 and ISR 1100X Series Integrated Services Routers (ISRs) • Cisco IR1101 Integrated Services Router Rugged • Cisco CSR 1000v Series Cloud Services Routers (CSR 1000V) • Cisco Catalyst 8000V Edge Software (Catalyst 8000V)

For details on supported models for each of these device families, refer to [Cisco SD-WAN Device Compatibility](#) page.

Restrictions for Policy Configuration Tagging

- Only data-prefix-lists, data-ipv6-prefix-lists, and app-lists tag members are supported.
- Configuration of both direction and direction-less tags within the same TAG is not supported.
- Configuration of tags using only Cisco SD-WAN Controller CLI templates is supported.
- Multi-tenancy is not supported.
- Configuration of number of tags is limited to maximum of 255.
- Configuration of objects per tag is limited to 64.

Information About Policy Configuration Tagging

The policy configuration tagging feature allows you to group policy objects and to assign tag values to various traffic flows by defining a policy. You can name the tags based on the functionality of the policy objects used

to achieve the intent-based network configurations. These tags that are provisioned through the Cisco SD-WAN Controller are used in the policy rules for traffic classification.

You can assign unique tag IDs while creating each of the tags.

You can define members under a tag name, which are referenced directly under tag objects. The members can be directional or directionless. Supported tag member types are:

- Data-prefix-list
- Data-ipv6-prefix-list
- App-list


Data-prefix-list and data-ipv6-prefix-list are directional attributes, which are matched as source or destination keywords in the data-policy match statements. App-list is a directionless attribute. You can use directionless keyword such as application id in the app-list policy match statements. Directional and directionless attributes cannot be grouped under the same tag.

You can apply the configured tags in a match criterion under localized and centralized policies. Devices process the tag configurations and apply the configurations to the data plane when the tag is referenced in the policy.

You can use the configuration type feature to tag objects in a configuration. The configuration tags are used in Cisco Catalyst SD-WAN centralized policy such as data policy, and app-aware routing policy and localized access-list policy. The following tag attributes are used in a policy match sequence statement:

- Source-tag-instance
- Destination-tag-instance
- Tag-instance

Figure 17: Policy Configuration Tagging in a Cisco Catalyst SD-WAN Network



As shown in the figure, at the Cisco SD-WAN Controller you can configure the tags using the policy objects with unique tag IDs. Once the tag IDs are assigned these tags are pushed to the Cisco IOS XE Catalyst SD-WAN devices in the network, which reference these tags. The devices then extract the policy list objects from the tags, which are used in the policy rules.

Features of Policy Configuration Tagging

- Supports only configuration type tag.
- Supports tagging a group of objects configuration.
- Supported tag members are data-prefix-lists, data-ipv6-prefix-lists, and app-lists.
- Supports defining configuration tags through a tag-centric model called **Defined Tag**.
- Supports adding configuration only through Cisco SD-WAN Controller CLI templates from Cisco SD-WAN Manager.

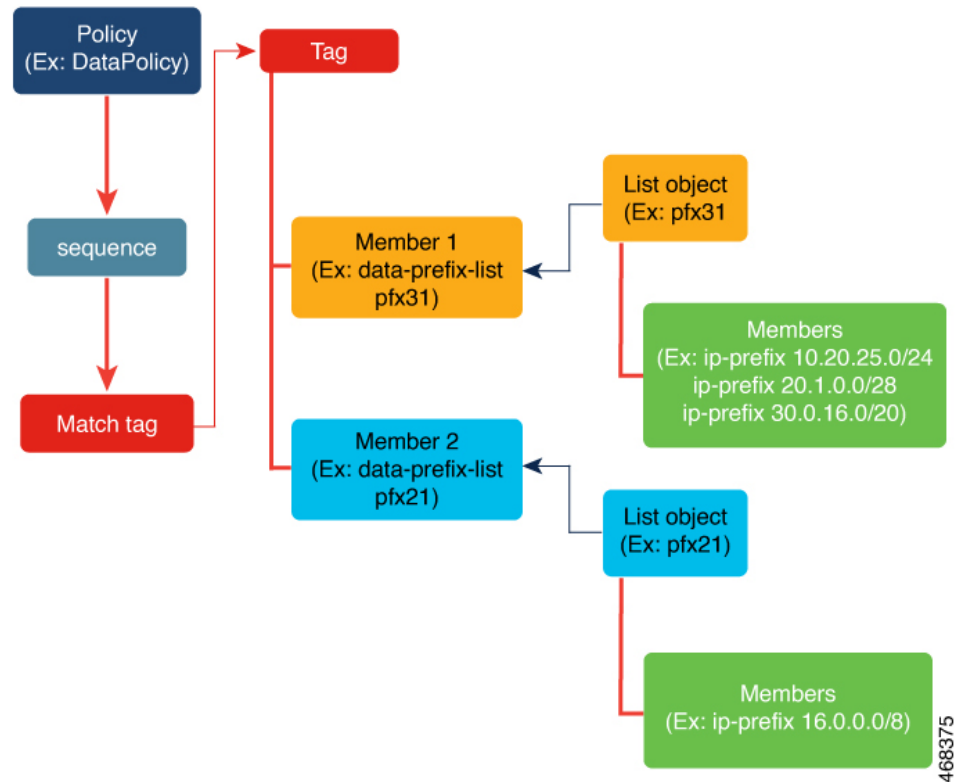
Tag Workflow

1. In Cisco SD-WAN Controller, create a tag that is based on the network intent.
2. Add the following policy list object members:
 - Data-prefix for each location
 - App-lists for applications

The policy list objects can be defined anytime in the workflow, even after adding them in the tag instances.

3. Push these tags to the Cisco IOS XE Catalyst SD-WAN devices in the network.
4. Create a policy with multiple match sequences and include the tag objects in the Cisco Catalyst SD-WAN data-policy, app-aware-routing policy, and access-list policy.
5. If you add or remove a tag, the status is automatically reflected in the policy.
6. Update the policy to include new tag objects.

Figure 18: Tagging Workflow with Examples



Benefits of Policy Configuration Tagging

The benefits of using policy configuration tagging are:

- Enables reusability of policy objects.
- Enables faster policy download on a device with reduced configuration size and sequences.
- Tag sharing across different policies is supported.
- Enables visibility or correlation across the network in a user-defined intent.
- Controls the policy configuration download speed between the Cisco SD-WAN Controller and the Cisco IOS XE Catalyst SD-WAN devices.
- Improves management of the defined lists in the controller.

- Better organization of the configurations for the intent-based network.

Configure Policy Configuration Tagging Using a CLI Template

Before You Begin

Ensure that the controllers and the edge devices are all updated to the latest versions—Cisco Catalyst SD-WAN Control Components Release 20.9.x, Cisco vManage Release 20.9.1, and Cisco IOS XE Catalyst SD-WAN Release 17.9.1a.

Configure Policy Configuration Tagging Using a CLI Template

For more information about using CLI templates, see [CLI Add-On Feature Templates](#) and [CLI Templates](#).



Note By default, CLI templates execute commands in global config mode.

This section provides example CLI configurations to configure tag-instances and centralized policy using Cisco SD-WAN Controller CLI templates.

Creating Policy Configuration Tagging

1. Configure a new object tag-instance on Cisco SD-WAN Controller:

```
tag-instances [tag-instance] [lists]
```

2. Create tag-instance with member attributes such as app-lists, data-ipv6-prefix-list, and data-prefix-list. Configure tag instances with a global unique ID for each of the tag names. The tag configuration is pushed to only those devices which reference these TAGs:

```
tag-instance tag-instance-name [id global-unique-id] [app-list app-list-name]
[data-prefix-list prefix-list-name] [data-ipv6-prefix-list ipv6-prefix-list-name]
```

3. Configure tag-instance lists:

```
lists[app-list app-list-name] [data-prefix-list prefix-list-name]
[data-ipv6-prefix-list ipv6-prefix-list-name]
```

Adding Tag-Instances in a policy match criteria

1. Configure localized access-list policy (ACLs and IPv6 ACLs) to include destination or source tag instances in matching attributes:

```
match [destination-tag-instance dest-tag-name | source-tag-instance
src-tag-name]
```

2. Configure centralized data policy to include destination-tag-instance, source-tag-instance, or tag-instance in matching attributes:

```
match [destination-tag-instance dest-tag-name | source-tag-instance
src-tag-name | tag-instance tag-name]
```

3. Configure centralized Application Aware Route (AAR) policy to include destination-tag-instance, source-tag-instance, or tag-instance in matching attributes:

```

match[destination-tag-instance dest-tag-name | source-tag-instance
src-tag-name | tag-instance tag-name]

```

Here's the complete configuration example for creating tag-instances, including the tag instances as matching attribute in localized and centralized policies:

```

****Tag Configuration****
tag-instances
tag-instance blue
  id 2000
  data-ipv6-prefix-list v6_pfx1 v6_pfx2
!
tag-instance orange
  id 3000
  app-list appl1 appl2
!
lists
data-prefix-list pfx1
  ip-prefix 10.0.0.1/32
!
data-ipv6-prefix-list v6_pfx1
  ipv6-prefix 2001::1/128
!
app-list appl1
  app amazon
!
!
!
****Localized Policy****
policy
lists
data-prefix-list pfx1
  ip-prefix 10.20.24.0/24
!
!
access-list acl
sequence 10
  match
    source-tag-instance blue
  !
  action accept
  count acl_input_wc
!
!
default-action drop
!
!
****Centralized Policy ****
policy
data-policy DP1
vpn-list vpn1
sequence 100
  match
    tag-instance orange
  !
  action accept
!
!
sequence 200
  match
    source-tag-instance blue
!
!

```

```

    action drop
      count count1
    !
  !
sequence 300
match
  destination-tag-instance blue
!
action accept
!

```

Verify Tag-Instances Configuration Using the CLI

The following is a sample output from the **show sdwan tag-instances from-vsmart** command displaying the downloaded tags from Cisco SD-WAN Controller on Cisco IOS XE Catalyst SD-WAN device:

```

Device# show sdwan tag-instances from-vsmart
tag-instances-from-vsmart
tag-instance APP_facebook_TAG9
  id 60000
  app-list apps_facebook
tag-instance APP_office_TAG10
  id 70000
  app-list apps_ms apps_zoom
tag-instance APP_webex_TAG8
  id 50000
  app-list apps_webex
lists data-prefix-list multicast_pfx
  ip-prefix 10.10.20.30/8
lists data-prefix-list pfx1
  ip-prefix 10.20.24.0/24
lists data-prefix-list pfx21
  ip-prefix 172.16.10.10/8
lists data-prefix-list pfx22
  ip-prefix 172.16.20.20/16
  ip-prefix 192.168.10.20/8
lists data-ipv6-prefix-list v6_pfx1
  ipv6-prefix 2001::/64
lists data-ipv6-prefix-list v6_pfx21
  ipv6-prefix 2001::1/128
  ipv6-prefix 2001::/64
lists app-list apps_facebook
  app dns
  app facebook
lists app-list apps_ms
  app ms-office-365
  app ms-office-web-apps
  app ms-services
  app ms-teams
  app pop3
lists app-list apps_webex
  app sip
  app webex-audio
  app webex-control
  app webex-media
  app webex-meeting
  app webex-video
lists app-list apps_zoom
  app zoom-meetings

```


The following is a sample output from the **show sdwan policy from-vsmart** command displaying the policy that is downloaded from the Cisco SD-WAN Controller on Cisco IOS XE Catalyst SD-WAN device:

```
Device# show sdwan policy from-vsmart
from-vsmart sla-class SLA1
  latency 100
from-vsmart data-policy DATA_POLICY
  direction from-service
  vpn-list vpn_1
  sequence 11
    match
      destination-port      5060
      protocol              17
      source-tag-instance   DP_V4_TAG1
      destination-tag-instance DP_V4_TAG3
    action accept
      count src_dst_legacy_v4
  sequence 21
    match
      source-tag-instance DP_V4_TAG1
    action drop
      count src_v4
  sequence 31
    match
      source-tag-instance   DP_V4_TAG2
      destination-tag-instance DP_V4_TAG3
      tag-instance          APP_webex_TAG8
    action drop
      count src_dst_app_v4
  sequence 41
    match
      source-tag-instance   DP_V4_TAG1
      destination-tag-instance DP_V4_TAG3
      tag-instance          APP_facebook_TAG9
    action accept
      count src_dst_app2_v4
```

The following is a sample output from the **show platform software common-classification** command displaying the tag information from a forwarding manager on a forwarding plane (FMAN-FP):

```
Device# show platform software common-classification F0 tag all
Total Number of TAGs: 9
tag id      tag name          tag type      num clients  num sets    num member types
total members
-----
900         special_TAG7      Per Type OR   0             2           1
  2
10000      DP_V4_TAG1        Per Type OR   1             1           1
  1
11000      DP_V4_TAG2        Per Type OR   1             2           1
  2
12000      DP_V4_TAG3        Per Type OR   1             6           1
  6
20000      DP_V6_TAG4        Per Type OR   1             1           1
  1
21000      DP_V6_TAG5        Per Type OR   1             2           1
  2
50000      APP_webex_TAG8    Per Type OR   1             1           1
  1
60000      APP_facebook_TAG9 Per Type OR   1             1           1
  1
70000      APP_office_TAG10 Per Type OR   1             2           1
  2
```

```

Device# show platform software common-classification f0 tag 1 summary
TAG ID: 1
TAG TYPE: Per Type OR
TAG Name: net1
Is Dummy: F

client data:
  client id      client name
  -----
  166            SDWAN

member data:
  Prefix List    6
  App List       3

Device# show platform software common-classification f0 tag 1 prefixList
member details:
member detail type      member id      member data
-----
IPv4 Prefix List        65537          100
IPv6 Prefix List        65538          101
IPv4 Prefix List        65540          103
IPv6 Prefix List        65541          104
IPv6 Prefix List        65544          107
IPv4 Prefix List        65546          109

Device# show platform software common-classification f0 tag 1 applist
member details:
member detail type      member id      member data
-----
App List                65539          102
App List                65542          105
App List                65545          108

Device# show platform software common-classification f0 tag 1 set
Total Number of SETs: 18
Set ID      member detail type      member id      member data
-----
1           IPv4 Prefix List        65537          100
1           App List                65539          102
2           IPv4 Prefix List        65537          100
2           App List                65542          105
3           IPv4 Prefix List        65537          100
3           App List                65545          108
4           IPv6 Prefix List        65538          101
4           App List                65539          102
5           IPv6 Prefix List        65538          101

```



CHAPTER 18

Integrate Cisco IOS XE Catalyst SD-WAN Device with Cisco ACI

Table 44: Feature History

Feature Name	Release Information	Description
Integration with Cisco ACI	Cisco IOS XE Catalyst SD-WAN Release 16.12.1b	The Cisco IOS XE Catalyst 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 Catalyst SD-WAN Controller. The Cisco Catalyst SD-WAN Controller, which is configured in Cisco ACI as an external device manager that provides Cisco IOS XE Catalyst 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 using 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 264](#)
- [Verify Cisco ACI Registration, on page 264](#)
- [SLA Classes, on page 264](#)
- [Data Prefixes, on page 265](#)
- [VPNs, on page 265](#)
- [Map Data Prefix and VPN to SLA, on page 265](#)
- [Create an App-Route-Policy, on page 265](#)
- [Map ACI Sites, on page 266](#)
- [Unmap ACI Sites, on page 267](#)
- [Delete a Controller, on page 267](#)

Guidelines to Integrate with Cisco ACI

The general steps that you perform in Cisco SD-WAN Manager to configure the integration are:

1. Verify that Cisco ACI has registered the desired controller as a partner with a Cisco Catalyst SD-WAN Controller, as described in the procedure, [Verify Cisco ACI Registration, on page 264](#).
2. Attach devices to the Cisco Catalyst SD-WAN Controller, as described in the Map ACI Sites section.

The following guidelines apply when integrating Cisco SD-WAN Manager with Cisco ACI:

- Only new Cisco IOS XE Catalyst 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 SD-WAN Manager policy builder.
- Before you apply WAN SLA policies, establish a connection between the Cisco Catalyst SD-WAN Controller and the Cisco APIC. For instructions, see [Cisco ACI and Cisco IOS XE Catalyst 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 SD-WAN Manager, perform the following steps in the Cisco SD-WAN Manager to verify that Cisco ACI has registered the desired controller as a Cisco SD-WAN Manager partner:

1. In Cisco SD-WAN Manager, select **Administration > Integration Management**.
The Integration Management page displays.
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 SD-WAN Manager 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 SD-WAN Manager, 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 SD-WAN Manager.

To view these data prefixes, follow these steps:

1. In Cisco SD-WAN Manager, 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 SD-WAN Manager. These VPNs become available in Cisco SD-WAN Manager automatically. You do not need to configure the VPNs in Cisco SD-WAN Manager.

To view these VPNs, follow these steps:

1. In Cisco SD-WAN Manager, select **Configuration > Policies**.
2. From the Custom Options drop-down menu, select **Lists**.
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 SD-WAN Manager. You can view these mappings in Cisco SD-WAN Manager 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:

1. In Cisco SD-WAN Manager, select **Configuration > Policies**.
2. Click the **More Actions** icon at the right of a row that contains a centralized policy, and then click **Edit**.
3. Select **Traffic Rules**.
4. Select **Add Policy > Create New**.
5. Click **ACI Sequence Rules**.
6. From the VPN drop-down, choose a VPN ID. Cisco SD-WAN Manager displays a list of data prefixes and SLA classes that are mapped to this VPN. (These mappings were sent by Cisco ACI.)
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**.
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 SD-WAN Manager creates the policy.
9. To apply a site list and a VPN list to the policy, select **Policy Application**, then select **Application-Aware Routing**, and click **New Site Lists and VPN List**.
10. Select a site list and a VPN list for the policy.
11. Add sequence rules to the policy as needed.
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:

1. In Cisco SD-WAN Manager, select **Administration > Integration Management**.
2. Click the **More Actions** icon to the right of the row for the applicable site and select **Attach Devices**.
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**.
4. Click the arrow pointing right to move the device to the Selected Devices column on the right.



Note 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.

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:

1. In Cisco SD-WAN Manager, select **Administration > Integration Management**.
The Integration Management page displays.
2. Click the **More Actions** icon to the right of the row for the applicable site and select **Detach Devices**.
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**.
4. Click the arrow pointing right to move the device to the Selected Devices column on the right.



Note 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.

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 SD-WAN Manager. Deleting an ACI partner from Cisco SD-WAN Manager 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.

1. In Cisco SD-WAN Manager, select **Administration > Integration Management**.
2. Detach all devices that are attached to the controller.
For instructions, see the Detach Devices from a Controller section.
3. Click the **More Actions** icon to the right of the row for the applicable site and select **Delete Controller**.



CHAPTER 19

Custom Applications

Table 45: Feature History

Feature Name	Release Information	Description
Support for Defining Custom Applications	Cisco IOS XE Catalyst SD-WAN Release 17.3.1a Cisco vManage Release 20.3.1	This feature adds support for defining custom applications.

- [Information About Custom Applications](#), on page 269
- [Configure Custom Applications Using Cisco SD-WAN Manager](#), on page 272
- [Verify Custom Applications](#), on page 273

Information About Custom Applications

Cisco Network-Based Application Recognition (NBAR) is a Cisco technology that performs the SD-WAN Application Intelligence Engine (SAIE) flow on network traffic to identify network applications according to their traffic characteristics.



Note In Cisco vManage Release 20.7.1 and earlier releases, the SAIE flow is called the deep packet inspection (DPI) flow.

The specific traffic characteristics of a network application are called an application signatures. Cisco packages the signature for an application, together with other information, as a protocol. Cisco packages a large set of protocols, covering numerous commonly occurring network applications, as a Protocol Pack. Cisco updates and distributes Protocol Packs regularly. They provide a database of network application signatures for NBAR to use to identify network application traffic.

The term network applications is defined broadly, and may include all of the following, and more:

- Social media websites
- Voice over IP (VoIP) applications
- Streaming audio and video, such as Cisco Webex
- Cloud applications, such as for cloud storage

- SaaS applications
- Custom network applications specific to an organization

Identifying applications is useful for monitoring network traffic, configuring application-aware traffic policy, and more.

To summarize network application signatures, protocols, and Protocol Packs, and how NBAR uses them:

- The traffic of a network application has unique characteristics that can be used to identify the traffic as belonging to that specific application. These characteristics are called application signatures.
- Cisco packages the signature for a specific network application as a protocol.
- Cisco packages a large set of protocols, covering commonly occurring internet applications, as Protocol Packs.
- Cisco NBAR performs the SAIE flow on traffic to gather the information required to identify the sources of the traffic, and uses protocols, such as those provided in Protocol Packs, to match that information to specific network applications. The result is that NBAR identifies the network applications producing traffic in the network.

Cisco Software-Defined Application Visibility and Control (SD-AVC) uses Cisco NBAR application identification to provide information about application usage within a network.

Custom Applications

In addition to the standard protocols provided in a Protocol Pack, you can define protocols, called custom applications, to identify internet traffic, often for uncommon network applications that are of specific interest to their organization. Custom applications augment the protocols provided in a Protocol Pack.

You can use custom applications in the same way as any other protocol when configuring:

- Cisco Catalyst SD-WAN policies
- Application Quality of Experience (AppQoE) policies, such as application-aware routing, TCP acceleration, and Quality of Service (QoS)



Note The following terms are used in the documentation of related technologies, and are equivalent: custom applications, custom protocols, user-defined applications

Custom Applications in Cisco Catalyst SD-WAN

Cisco Software-Defined AVC (SD-AVC) is a component of Cisco Application Visibility and Control (AVC). It functions as a centralized network service, operating with specific participating devices in a network. One function of Cisco SD-AVC, which is included as a component of Cisco Catalyst SD-WAN, is to create and manage custom applications. Cisco Catalyst SD-WAN uses this Cisco SD-AVC functionality, through SD-AVC REST APIs, to enable you to define custom applications within Cisco Catalyst SD-WAN.

As a Cisco Catalyst SD-WAN user, you can use Cisco SD-WAN Manager to define custom applications. Cisco SD-AVC then pushes the custom applications to devices in the network. The devices in the network use the custom applications and other application protocols to analyze traffic traversing the devices.

The process of defining a custom protocol includes choosing criteria to identify network traffic as coming from a specific network application. The criteria can include characteristics of hosts originating the traffic, such as server names, IP addresses, and so on.

Priority of Protocols and Custom Applications

It is possible to define custom applications that match some of the same traffic as a protocol included in the Protocol Pack operating with Cisco NBAR. When matching traffic, custom applications have priority over Protocol Pack protocols. Deploying SD-AVC within an existing network does not require any changes to the network topology.

Restrictions for Custom Applications

- Maximum number of custom applications: 1100
- Maximum number of L3/L4 rules: 20000
- Maximum number of server names: 50000
- For server names, maximum instances of wildcard followed by a period (.): 50000
Example: *.cisco.com matches www.cisco.com, developer.cisco.com
- For server names, maximum instances of prefix wildcard as part of server name: 256
Example: *ample.com matches www.example.com
- Mapping the same domain to two different custom applications is not supported.
- DNS traffic and application traffic need to be in the same VRF for SD-AVC to perform first packet classification.
- Creating custom applications through CLI is not supported in Cisco Catalyst SD-WAN policy.
- Activation of custom applications:
 - When using Cisco vManage Release 20.5.1 releases and earlier: For devices using releases earlier than Cisco IOS XE Catalyst SD-WAN Release 17.5.1a, the activation of custom applications is as follows:
 - A custom application created in Cisco SD-WAN Manager is not activated for visibility functionality (monitoring traffic) or control functionality (traffic policy) until a policy that makes use of the custom application is applied.
 - When using Cisco vManage Release 20.5.1 or later: For devices using Cisco IOS XE Catalyst SD-WAN Release 17.5.1a or later, the activation of custom applications is as follows:
 - A custom application created in Cisco SD-WAN Manager is activated immediately for application visibility functionality only (monitoring traffic), such as for protocol-discovery counters and Flexible NetFlow (FNF). When activated for visibility functionality only, custom applications do not affect traffic policy.
 - When the custom application is used by a policy, it becomes activated for control functionality (traffic policy) also.

Configure Custom Applications Using Cisco SD-WAN Manager

Prerequisites

Install Cisco SD-AVC as a component of Cisco Catalyst SD-WAN. For information on how to enable SD-AVC on Cisco SD-WAN Manager, see [Information on how to enable SD-AVC for Cisco SD-WAN devices](#).

Perform the following steps to configure custom applications:

1. In Cisco SD-WAN Manager, select **Configuration** > **Policies**.
2. Select **Centralized Policy**.
3. Click **Custom Options** and select **Centralized Policy** > **Lists**.
4. Click **Custom Applications**, and then click **New Custom Application**.
5. To define the application, provide an application name and enter match criteria. The match criteria can include one or more of the attributes provided: server names, IP addresses, and so on. You do not need to enter match criteria for all fields.

The match logic follows these rules:

- Between all L3/L4 attributes, there is a logical AND. Traffic must match all conditions.
- Between L3/L4 and Server Names, there is a logical OR. Traffic must match either the server name or the L3/L4 attributes.

Field	Description
Application Name	(mandatory) Enter a name for the custom application. Maximum length: 32 characters
Server Names	One or more server names, separated by commas. You can include an asterisk wildcard match character (*) only at the beginning of the server name. Examples: *cisco.com, *.cisco.com (match www.cisco.com, developer.cisco.com, ...)
L3/L4 Attributes	
IP Address	Enter one or more IPv4 addresses, separated by commas. Example: 10.0.1.1, 10.0.1.2 Note The subnet prefix range is 24 to 32.

Field	Description
Ports	Enter one or more ports or port ranges, separated by commas. Example: 30, 45-47
L4 Protocol	Select one of the following: TCP, UDP, TCP-UDP

- Click **Add**. The new custom application appears in the table of custom applications.



Note To check the progress of creating the new custom application, click **Tasks** (clipboard icon). A panel opens, showing active and completed processes.

Example Custom Application Criteria

Criteria	How to configure fields
Domain name	Server Names: cisco.com
Set of IP addresses, set of ports, and L4 protocol	IP Address: 10.0.1.1, 10.0.1.2 Ports: 20, 25-37 L4 Protocol: TCP-UDP
Set of ports and L4 protocol	Ports: 30, 45-47 L4 Protocol: TCP

Verify Custom Applications

Verify Custom Applications in Cisco SD-WAN Manager

After you define a custom application, it appears in the **Custom Application List**, which shows all available protocols and custom applications. The **Custom Application List** is available here:

Configuration > Policies > Centralized Policy > Add Policy > Custom Applications.

Verify Protocols and Custom Applications on a Device

Use the **show ip nbar protocol-id** command to display all protocols and custom applications that are loaded on the router. It is helpful to filter the results. For example, to display all protocols and custom applications with "custom" in the name, use this:

```
vm5#show ip nbar protocol-id | include custom
custom_amazon          3899          PPKD LOCAL
custom_facebook        3284          PPKD LOCAL
```

See [show ip nbar protocol-id](#).



CHAPTER 20

Service Insertion

Table 46: Feature History

Feature Name	Release Information	Description
Service Insertion Using Workflows	Cisco IOS XE Catalyst SD-WAN Release 17.13.1a Cisco Catalyst SD-WAN Manager Release 20.13.1	With this feature, you can create a service chain from the Workflow Library and configure a service chain action for a policy. A service chain inserts a set of services in the flow of traffic and can be designed to affect traffic according to your needs.
Trusted and Untrusted Postures	Cisco IOS XE Catalyst SD-WAN Release 17.14.1a Cisco Catalyst SD-WAN Manager Release 20.14.1	With this feature, you can configure trusted traffic to flow to a trusted high availability pair in a service chain.

- [Information About Service Insertion](#), on page 276
- [Restrictions for Service Insertion](#) , on page 280
- [Use Cases for Service Insertion](#) , on page 280
- [Configure Service Insertion](#), on page 281
- [Configure Service Chain Actions in a Data Policy](#), on page 282
- [Traffic Steering to a Service Chain](#), on page 283
- [Path Preference](#), on page 286
- [Share Service Chains Across User VPN](#), on page 286
- [Separate Interfaces for Transmitted and Received Traffic](#), on page 287
- [Service Chaining Trusted and Untrusted Traffic](#), on page 287
- [Service Chain Between Two Routers](#), on page 288
- [Configure Fall Back and Restrict Behavior for Traffic Through a Service Chain](#), on page 288
- [Interfaces for Attaching Services in a Service Chain to a Router](#), on page 289
- [Service Chaining with Software Defined Cloud Interconnect Bring Your Own Service](#) , on page 289
- [Configure Service Insertion Using a CLI Template](#), on page 290

Information About Service Insertion

Service insertion, also known as *service chaining*, refers to placing one or more network or security services into the path of specific data traffic within the Cisco Catalyst SD-WAN overlay fabric. These services are defined in a service chain, which is a set of services that traffic routes through. The traffic is routed according to service chain actions that you configure for a data policy.

A service chain can be in any device, and can be used in any topology, including full mesh, hub-spoke, and Cisco Catalyst SD-WAN Multi-Region Fabric (MRF).

Cisco Catalyst SD-WAN service chaining is flexible, fully automated, and can be deployed on a per VPN basis. Service chaining includes the following key feature:

- Service chaining can be used for overlay, local ingress and egress, inter- and intra-VPN, transit, branch-to-branch, branch-to-internet, branch-to-cloud, and cloud-to-cloud traffic
- Automatic forwarding of traffic through all services in a chain
- Services attachment methods of IPv4, IPv6, dual stack, and tunneled
- Configurable high availability across instances of a single service
- Built-in load balancing across instances of a single service, which supports equal cost multipath routing (ECMP) across high availability pairs
- Advanced service tracking
- Service chain sharing across multiple user VPNs, which can be different or the same as user traffic VPNS
- Traffic steering methods using control policy, data policy, interface ACL, and supported match conditions
- Fall back and restrict behavior
- Path preference and symmetric routing
- Security services to and from service transports
- Trusted and untrusted high availability pairs and traffic marking (from Cisco Catalyst SD-WAN Manager Release 20.14.1)
- Periodic on demand state notifications for serviceability
- Cisco Catalyst SD-WAN Manager orchestration: Workflow based service chaining and traffic policy configuration

Service Insertion Capabilities

The following table provides information about the capabilities of the service chaining feature in releases before and after Cisco Catalyst SD-WAN Manager Release 20.13.1.


Capability	Releases Before Cisco Catalyst SD-WAN Manager Release 20.13.1	Releases From Cisco Catalyst SD-WAN Manager Release 20.13.1
Multiple services in a chain	No native support	Native support

Capability	Releases Before Cisco Catalyst SD-WAN Manager Release 20.13.1	Releases From Cisco Catalyst SD-WAN Manager Release 20.13.1
Traffic steering	Control policy	Control policy, data policy, interface ACL
Policy binding	Remote	Remote and local
Traffic type	IPv4	IPv4, IPv6, dual stack, tunnel
Load balancing	Across 4 IP addresses that serve as service endpoints	Across 4 instances of active-backup pairs for every traffic type
High availability	As provided by load balancing	Active and backup pairs
Tracking	To one connection per service instance	To every connection toward an abstract service
Configurable tracker probes	Not supported	All trackers are individually configurable
Behind-the-service tracking	Not supported	Supported
Affinity (service routes and data policy)	Not supported	Supported
TLOC preference	Supported	Supported
Fall back, restrict	Not supported	Supported
Tunnel connected services	Not supported	Supported
Shared service VPN	Not supported	Supported
To and from service transports	Not supported	Supported
Trusted and untrusted postures	Not supported	Supported from Cisco Catalyst SD-WAN Manager Release 20.14.1
Periodic and on-demand serviceability	Not supported	Supported
Cisco Catalyst SD-WAN Manager orchestration	Uses feature templates	Uses the Workflow Library and configuration groups (feature templates are not supported)
Deployment	On premises	On premises, cloud, middle mile colocated
Service instance type	Physical	Physical or virtual

Service Insertion Key Concepts and Implementation

The following figure illustrates the basic concepts of service chaining and the general steps involved in service chain creation and execution.

Figure 19: Service Insertion Concepts and Steps



1	<p>Services bring up:</p> <ul style="list-style-type: none"> • Bring up and connect services to the Cisco Catalyst SD-WAN routers. • Use Cisco Catalyst SD-WAN Manager to bring up desired services.
2	<p>Service chain configuration and advertising:</p> <ul style="list-style-type: none"> • Use the Workflow Library or CLI commands to configure a service chain for routers, shown as SC1 • Use a configuration group to configure SC-HUB1. The Workflow Library configuration adds auto-generated service chain configuration to the Service VPN part of the configuration group base on your inputs. • SC-HUB advertises the service chain to the Cisco SD-WAN Controller


3	<p>Service chain policy:</p> <ul style="list-style-type: none"> • Match traffic or routes and perform service chain actions • Apply a service chain policy to sites where traffic originates • The Cisco SD-WAN Controller resolves and advertises the service chain to target sites
4	<p>Traffic steering:</p> <ul style="list-style-type: none"> • Traffic is steered from the source (B1) to SC-HUB • The first service in the service chain is executed • Traffic returns to SC-HUB from the first service • The second service in the service chain is executed • Traffic returns to SC-HUB from the second service • Traffic is forwarded to the destination (B2)

The following figure illustrates the key elements in service insertion. In this figure, SC-HUB1 is the router to which the service chain is attached.

In the hairpin model, traffic is sent by SC-HUB1 to a service in the service chain, and the service returns the traffic to SC-HUB1. SC-HUB1 then either forwards the traffic to the next service in the service chain or to the destination if the traffic is returning from the last service in the service chain.

In the exit Model, traffic is sent by SC-HUB1 to a service in a service chain, and the service forwards the traffic to the destination. Traffic may return from the destination to the service, which returns it to SC-HUB1.

Figure 20: Service Insertion Key Elements



Restrictions for Service Insertion

- A service chain can include up to four service types. Each service type can have multiple instances of the service, either as high availability pairs that are load-balanced by the feature, or behind a third-party load balancer.
- The services in a service chain must be in a single VPN.
- If you are using a dual stack service in a service chain, every service in that service chain must have a dual stack high-availability pair.
- A specific device interface should not be used for more than one service in a particular service chain.
- A specific interface can be used in different service chains only if it is used for the same service type in each of the service chains.
- All the interfaces and tunnels for the services in a service chain should be part of the VPN in which the service chain is defined.
- More than one tracker should not be associated with a given interface. For example, if endpoint-tracker tracker1 is associated with GigabitEthernet1, a different tracker cannot be associated with GigabitEthernet1.

Use Cases for Service Insertion

- Service chaining can be used when traffic from a less secure region of a network should pass through a firewall to ensure that it has not been tampered with.
- Service chaining can be used in a network that consists of multiple VPNs, where each represents a different function or organization, to ensure that traffic between VPNs flows through a firewall. For example, in a campus, interdepartmental traffic might go through a firewall, while intradepartmental traffic might be routed directly.

- Service chaining can be used to ensure regulatory compliance, such as Payment Card Industry Data Security Standard (PCI DSS), where PCI traffic should flow through firewalls in a centralized data center or regional hub.

Configure Service Insertion

Beginning with Cisco Catalyst SD-WAN Manager Release 20.13.1, you can configure service insertion by using the **Workflow Library**. From the **Workflow Library**, you can create a new service chain or modify an existing one. A service chain can contain up to four service types.

The workflow guides you through configuring several steps, including:

- Configuring the name and description of the service chain
- Specifying the services in the service chain and the order of the services in the chain
- Provide attachment parameters for the services in the chain, which are used when you attach the service chain to routers
- For each service type, specify the VPN and configure options such as load balancing, high availability and tracking

To create or modify a service chain:

1. From the Cisco SD-WAN Manager menu, choose **Workflows > Workflow Library**.
2. Click **Define and Configure Service Chain**.
3. Follow the prompts in the workflow.

Ensure that you define a tracker. Tracker configuration is critical to avoid blackhauling. Defining a tracker ensures that the service chain is determined to be in the UP state and is used. If the IP address of a service chain firewall is used with an ICMP-based tracker, ensure that the firewall allows ICMP on the appropriate interface.

Ensure that the service chain can route returning traffic back into the Cisco Catalyst SD-WAN fabric. To do so, use dynamic routing protocols between the service chain and Cisco Catalyst SD-WAN router (service chain hub) or use static routes.

Attach the service chain to the appropriate Cisco Catalyst SD-WAN SC-Hub router. The service chain does not need to be attached to branch routers.

After you configure service insertion, perform the following actions as needed:

- Configure service chain actions for a data policy to route traffic through a service chain. See [Configure Service Chain Actions in a Data Policy](#).
- Use a control policy, data policy, or interface access control list to direct traffic to a service chain. See [Traffic Steering to a Service Chain](#).
- Configure TLOC preference or affinity preference to choose the preferred path for traffic to a service chain. See [Path Preference](#).
- Configure separate interfaces for transmitted and received traffic. See [Separate Interfaces for Transmitted and Received Traffic](#).

- Configure trusted traffic to flow to a trusted high availability pair. See [Service Chaining Trusted and Untrusted Traffic](#).
- Configure fall back or restrict behavior for traffic that travels through a service chain. See [Configure Fall Back and Restrict Behavior for Traffic Through a Service Chain](#).

Configure Service Chain Actions in a Data Policy

Beginning with Cisco Catalyst SD-WAN Manager Release 20.13.1, you can route traffic through a service chain by configuring service chain actions for a data policy.

1. From the Cisco SD-WAN Manager menu, choose **Configuration > Policies**.
2. Click **Custom Options** then click **Traffic Policy** under **Centralized Policy**.
3. Click the **Traffic Data** tab.
4. Click **Add Policy** and click **Create New**.
5. Click **Sequence Type** and choose **Service Chaining** from the **Add Data Policy** dialog box.
6. Click the **Actions** tab.
7. Click **Service**.
8. Configure the fields that the following table describes.

Table 47: Service Chain Action Fields

Field	Description
Service: Type	Choose a type of service for the service chain.
Service: VPN	VPN in which the service chain is hosted. Range: 0 through 65530
Service: TLOC IP	Enter the IP address of the Transport Locator (TLOC) for applying the services in the service chain.
Color	Choose a color for the TLOC.
Encapsulation	Choose the encapsulation type for the TLOC.
Service: TLOC List	Choose a predefined TLOC list to use for applying services to branch traffic.
Local	Check the Local check box if the service chain is hosted locally. If you do not check this check box, the service chain is hosted remotely.

Field	Description
Restrict	<p>Check this option to cause packets to be dropped if the service chain goes down. If you configure this policy with the Local option, packets are dropped locally. If you configure this policy with the Remote option, packets are dropped on the remote host.</p> <p>This option is unchecked by default (traffic falls back to routing).</p>

Traffic Steering to a Service Chain

You can direct traffic to a service chain by using a control policy, data policy, or interface access control list.

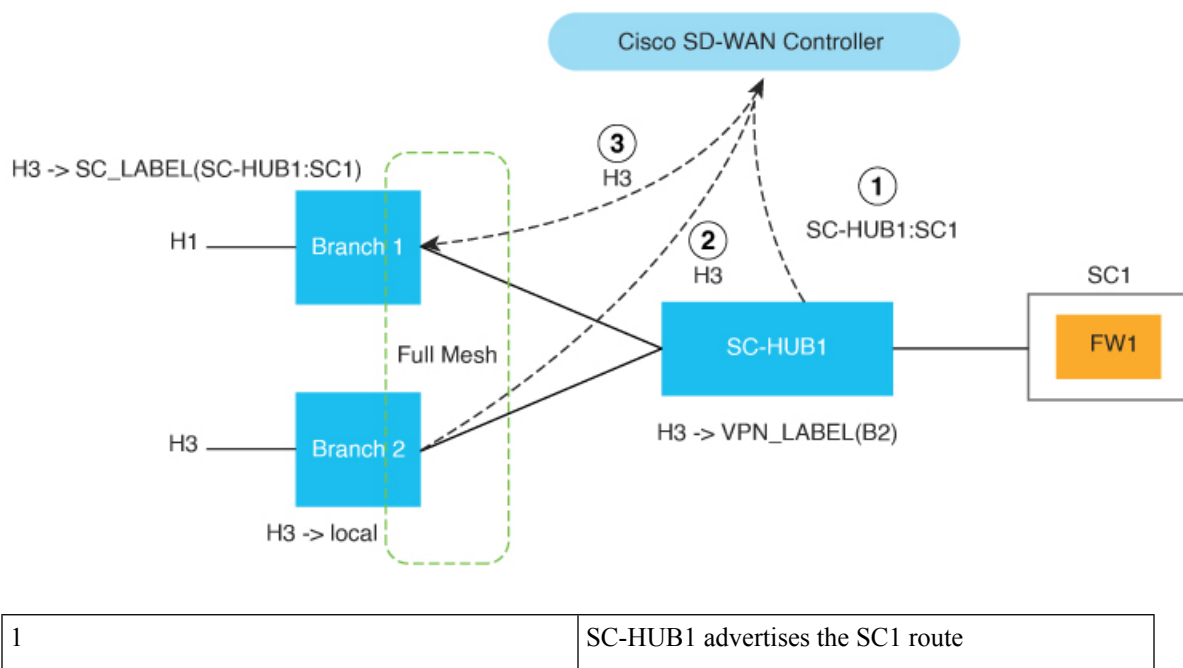
Traffic Steering Using a Control Policy

You can use a control policy to modify Cisco Overlay Management Routes, also referred to as vRoutes, to direct traffic to a service chain instead of the original destination.

The following figure shows an example of the use of a control policy to direct traffic to a service chain.

In this example, the policy causes service chain 1 (SC1) to be applied to traffic that flows between H1 (host 1) and H3 (host 3). The policy sets SC1 as the next hop for H1 and H3 traffic routes. Before the policy is in effect, traffic flows from B2 (branch 2) to B1 (branch 1). After the policy is in effect, traffic flows from B2 to SC-HUB1:SC1 to B1.

Figure 21: Traffic Steering with a Control Policy



2	B2 advertises the H3 route to the Cisco SD-WAN Controller
3	The control policy results in overriding the H3 route next hop to SC1 and the Cisco SD-WAN Controller advertises the H3 route to B1

Example configuration:


```
Control-policy name
  sequence number
  match route
  action accept
  set service-chain sc_name [tloc|tloc-list name] [vpn vpn]
apply-policy site-list site_list control-policy name out
```

Traffic Steering Using a Data Policy

You can use a data policy to match traffic and operate in the context of source VPNs during forwarding.

The following figure shows an example of the use of a data policy to specify service chaining intent in a remote branch.

Figure 22: Traffic Steering with Traffic Service Chaining Intent Specified at a Remote Branch



The following example shows configuration for traffic steering with a data policy when traffic intent is specified on a remote device. In this example:

- **match criteria** specifies applications to be matched to source and destination IP address combinations
- **restrict|fallback** configures restrict or fall back
- **tloc|tloc-list list** specifies the traffic path preference using TLOC ranking



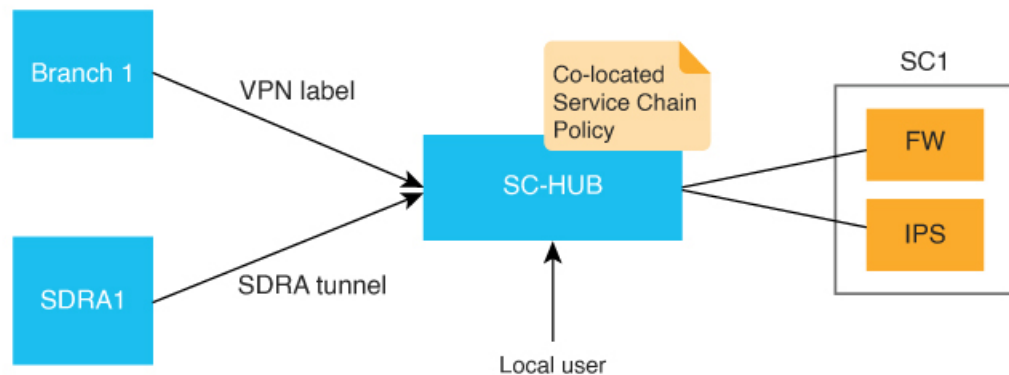
Note `set attribute trust-posture` is available from Cisco Catalyst SD-WAN Manager Release 20.14.1.

```

policy
  data-policy name
  vpn-list name
  sequence 100
  match criteria
  action accept
    set service-chain sc_name vpn vpn {restrict|fallback} [tloc|tloc-list list]
  set attribute trust-posture {trusted | untrusted}
  apply-policy site-list remote-sites data-policy name from-service
  
```

The following figure shows an example of the use of a data policy to specify the service chaining intent locally in the device to which the service chain is attached.

Figure 23: Traffic Steering with Traffic Service Chaining Intent Specified on a Local Device



The following example shows configuration for service chaining intent on a local device. In this example, **local** indicates that traffic needs to be directed to a service chain locally.

```

set service-chain SC1 [vpn vpn] local [restrict|fallback]
apply-policy site-list SC-HUB-sites data-policy policy {from-service|from tunnel}|from-tunnel}
  
```


Traffic Steering Using an Interface Access Control List

You can use an interface access control list (ACL) to service chain traffic that is incoming or outgoing on a specified interface. In some situations, the traffic forwarding decision may need to come from a prior routing lookup or data policy.

This approach is useful when all traffic from an interface should be directed through a service chain.

The following figure shows an example of the use of an ACL to direct traffic through a service chain.

Figure 24: Traffic Steering with an ACL



The following example shows configuration for traffic steering using an ACL.

```
access-list list
  sequence number
  match criteria
  action accept
    set service-chain SC1 [vpn vpn] {restrict|fallback}
interface interface
  access-list list {in|out}
```

Path Preference

You can use TLOC preference or affinity preference to choose the preferred path for traffic to a service chain.

To do so, configure a TLOC list to direct traffic only over certain TLOCs or to prefer certain TLOCs over others. The TLOC list can be specified with **tloc-list** as part of a service chain action in a data policy or a control policy.

To configure affinity preference, use **affinity-group preference** in branch sites to set the affinities of branches, and use **affinity-group** in service chain hubs to set the affinities of VPNs. The data policy **set service chain** action is compliant with affinity by default.

You can configure the following command to disable consideration of affinity in a data policy:

data-policy-ignore-affinity-metric

If both TLOC preference and affinity preference are configured, the affinity preference is evaluated first, then the TLOC preference is evaluated.


Share Service Chains Across User VPN

A service chain VPN can be shared across multiple user VPNs, and traffic between VPNs can be serviced chained in any VPN. Sharing a service chain does not require additional configuration. If source and destination VPNs are different, route leaking is required between the source and destination VPN.

The following figure illustrates the sharing of service chains across user VPNs. In this figure:

- SC1 (service chain 1) is attached to VPN100 can be automatically shared by traffic in VPN1 (H1) and VPN2 (H4)
- Traffic between VPN1 (H1) and VPN2 (H4) can be service chained in VPN1 or VPN2 or in a shared service chain (VPN100)

Figure 25: Service Chain Sharing Across VPNs




Separate Interfaces for Transmitted and Received Traffic

You can use the **service** command to configure separate interfaces for transmitted and received traffic through a service chain. In this situation, transmitted and received traffic are tracked independently. For more information, see [service](#).

The following figure illustrates this approach.

Figure 26: Separate Interfaces for Transmitted and Received Traffic



Service Chaining Trusted and Untrusted Traffic


Minimum supported releases: Cisco IOS XE Catalyst SD-WAN Release 17.14.1a, Cisco Catalyst SD-WAN Manager Release 20.14.1

You can configure trusted traffic to flow to a trusted high availability pair. In this situation, untrusted traffic flows to an untrusted high availability pair.

Use the **set attribute trust-posture untrusted action** in data policy to mark a packet as trusted or untrusted. The default trust-posture of a packet is trusted.

The following figure illustrates the flow of trusted and untrusted traffic.

Figure 27: Trusted and Untrusted Traffic



Example configuration:


```
service-chain SC1
  service netsvc1
    sequence 10
    service-transport-ha-pair 1
    attribute trust-posture {trusted|untrusted}
```

Service Chain Between Two Routers

If the router that is transmitting traffic to a service chain is different from the router that is receiving traffic from the service chain, configure the same service chain in each device. The service chain can have only one service and is for intra-VPN traffic only.

The following figure illustrates this approach.

Figure 28: Service Chain Between Two Router



Configure Fall Back and Restrict Behavior for Traffic Through a Service Chain

You can configure fall back or restrict behavior for traffic that travels through a service chain.

When **fallback** is configured in the **set service-chain** action, traffic falls back to routing if a service chain goes down or if the TLOCs that are specified in a policy are not available.

When **restrict** is configured in the **set service-chain** action, packets are dropped if a service chain goes down or if the TLOCs that are specified in a policy are not available. The restrict behavior is suitable for security services such as a firewall.

Fall back and restrict can be specified in a centralized data policy (remote or collocated) and an interface ACL.



Note If an egress ACL is used to direct traffic to a service chain, all packets continue to the destination even if the restrict behavior is configured because the forwarding decision is made before the state of the service chain is detected.

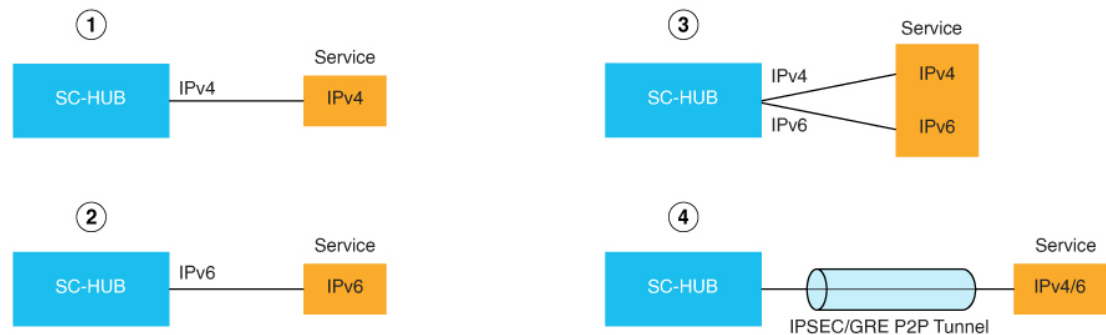
Interfaces for Attaching Services in a Service Chain to a Router

The services in a service chain must be in a single VPN, called a *service chain VPN*, or *SC-VPN*.

The services in a service chain can be attached to a Cisco Catalyst SD-WAN router through any combination of an IPv4, IPv6, dual stack, or tunnel interface.

The following figure illustrates the interfaces for attaching services in a service chain to a router.

Figure 29: Attaching Services to a Router



1	IPv4 attachment
2	IPv6 attachment
3	Dual stack attachment
4	Tunnel attachment

Service Chaining with Software Defined Cloud Interconnect Bring Your Own Service

The Software Defined Cloud Interconnect (SDCI) establishes connections between branch sites and the cloud through network service providers, including Megaport and Equinix. The SDCI bring your own service

(BYOS) functionality establishes a centralized location for service inspection by connecting a service chain to the Cisco Catalyst 8000v Edge Software (Catalyst 8000v) SDCI gateways that are deployed in the middle mile network. BYOS enables the seamless integration of external services with the SDCI infrastructure. Colocated data policies, also known as centralized data policies, are enforced on these gateways within the middle mile network for selective data traffic inspection.

In this context, a branch site represents the first mile, a service provider acts as the middle mile, and the cloud serves as the last mile.

The BYOS service inspection for SDCI allows service chaining in the following situations:

- Connecting branch sites to cloud workloads through middle mile providers using the C8000v SDCI gateway.
- Interconnecting branch site through the middle mile provider using the Catalyst 8000v SDCI gateway.
- Facilitating intercloud traffic connectivity by the middle mile provider through the Catalyst 8000v SDCI gateway.

Configure Service Insertion Using a CLI Template

For more information about using CLI templates, see [CLI Add-On Feature Templates](#) and [CLI Templates](#).



Note By default, CLI templates execute commands in global config mode.

The section provides a sample CLI configuration for service insertion.

1. Create a service chain:


```
service-chain chain-number
```
2. Configure a description for the service chain:


```
service-chain-description description
```
3. Specify the services that are in the service chain and configure related options:


```
service service-type service-parameters
```
4. (Optional, from Cisco Catalyst SD-WAN Manager Release 20.14.1) Configure the trust posture for the services that are in the service chain:


```
service service-type service-transport-ha-pair value attribute trust-posture {trusted | untrusted}
```
5. (Optional) Configure all Cisco Catalyst SD-WAN bidirectional forwarding (BFD) sessions to be brought down:


```
service-chain-affect-bfd
```
6. Specify the name of the VPN that hosts all services in the service chain:


```
service-chain-vrf vrf
```
7. (Optional, enabled by default) Enable endpoint tracking for services in the service chain:


```
track-enable
```
8. (Optional, enabled by default) Enable the service chain, which makes it active on devices:


```
service-chain-enable
```



CHAPTER 21

Service Chaining



Note Beginning with Cisco IOS XE Catalyst SD-WAN Release 17.13.1a and Cisco Catalyst SD-WAN Manager Release 20.13.1, see [Service Insertion](#) for information about service chaining.

Services in the Network

Services such as firewall, load balancer, and intrusion detection and prevention (IDP) are often run within a virtualized environment, and they may physically be centralized in one location or in several locations for redundancy. Services may be internal, cloud based, or external subscriptions. Networks must be able to reroute traffic from any location in the network through such services.

Customers want the ability to internally spawn or externally subscribe to new services on demand—for capacity, redundancy, or simply to select best-of-breed technologies. For example, if a firewall site exceeds its capacity, a customer can spawn a new firewall service at a new location. Supporting this new firewall would require the configuration of policy-based, weighted load distribution to multiple firewalls.

Following are some of the reasons to reroute a traffic flow through a service or chain of services:

- Traffic flow from a less secure region of a network must pass through a service, such as a firewall, or through a chain of services to ensure that it has not been tampered with.
- For a network that consists of multiple VPNs, each representing a function or an organization, traffic between VPNs must traverse through a service, such as a firewall, or through a chain of services. For example, in a campus, interdepartmental traffic might go through a firewall, while intradepartmental traffic might be routed directly.
- Certain traffic flows must traverse a service, such as a load balancer.

Today, the only way to reroute traffic flow is by provisioning every routing node—from the source to the service node to the systems beyond the service node—with a policy route. This is done either by having an operator manually configure each node or by using a provisioning tool that performs the configuration for each node on behalf of the operator. Either way, the process is operationally complex to provision, maintain, and troubleshoot.


Provisioning Services in the Cisco Catalyst SD-WAN Overlay Network

In the Cisco Catalyst SD-WAN solution, the network operator can enable and orchestrate all service chaining from a central controller, that is, from the Cisco SD-WAN Controller. No configuration or provisioning is required on any of the devices.

The general flow of service chaining in a Cisco Catalyst SD-WAN network is as follows:

- Devices advertise the services available in their branch or campus—such as firewall, IDS, and IDP—to the Cisco SD-WAN Controllers in their domain. Multiple devices can advertise the same services.
- Devices also advertise their OMP routes and TLOCs to the Cisco SD-WAN Controllers.
- For traffic that requires services, the policy on the Cisco SD-WAN Controller changes the next hop for the OMP routes to the service landing point. In this way, the traffic is first processed by the service before being routed to its final destination.

The following figure illustrates how service chaining works in the Cisco Catalyst SD-WAN solution. The network shown has a centralized hub router that is connected to two branches, each with a device. The standard network design implements a control policy such that all traffic from branch site 1 to branch site 2 travels through the hub router. Sitting behind the hub router is a firewall device. So now, assume we want all traffic from site 1 to site 2 to first be processed by the firewall. Traffic from the device at site 1 still flows to the hub router, but instead of sending it directly to site 2, the hub router redirects the traffic to the firewall device. When the firewall completes its processing, it returns all cleared traffic to the hub, which then passes it along to the device at site 2.



Service Route SAFI

The hub and local branch devices advertise the services available in their networks to the Cisco SD-WAN Controllers in its domain using service routes, which are sent by way of OMP using the service route Subsequent Address Family Identifier (SAFI) bits of the OMP NLRI. The Cisco SD-WAN Controllers maintain the service routes in their RIB, and they do not propagate these routes to the devices.

Each service route SAFI has the following attributes:

- VPN ID (vpn-id)—Identifies the VPN that the service belongs to.
- Service ID (svc-id)—Identifies the service being advertised by the service node. The Cisco Catalyst SD-WAN software has the following predefined services:
 - FW, for firewall (maps to svc-id 1)
 - IDS, for Intrusion Detection Systems (maps to svc-id 2)
 - IDP, for Identity Providers (maps to svc-id 3)
 - netsvc1, netsvc2, netsvc3, and netsvc4, which are reserved for custom services (they map to svc-id 4, 5, 6, and 7, respectively)
- Label—For traffic that must traverse a service, the Cisco SD-WAN Controller replaces the label in the OMP route with the service label in order to direct the traffic to that service.
- Originator ID (originator-id)—The IP address of the service node that is advertising the service.
- TLOC—The transport location address of the device that is “hosting” the service.
- Path ID (path-id)—An identifier of the OMP path.

Service Chaining Policy

To route traffic through a service, you provision either a control policy or a data policy on the Cisco SD-WAN Controller. You use a control policy if the match criteria are based on a destination prefix or any of its attributes. You use a data policy if the match criteria include the source address, source port, DSCP value, or destination port of the packet or traffic flow. You can provision the policy directly using the CLI, or it can be pushed from Cisco SD-WAN Manager.

The Cisco SD-WAN Controller maintains OMP routes, TLOC routes, and service routes in its route table. A given OMP route carries a TLOC and the label associated with it. On a Cisco SD-WAN Controller, a policy can be applied that changes the TLOC and its associated label to be that of a service.

Tracking the Health of the Service Chain

Beginning with Cisco IOS XE Catalyst SD-WAN Release 17.3.1a, Cisco Catalyst SD-WAN periodically probes devices providing network services to test whether they are operational. Tracking the availability of devices in the service chain helps to prevent a null route, which can occur if a policy routes traffic to a service device which is not available. By default, Cisco Catalyst SD-WAN writes the tracking results to a service log, but this can be disabled.

Limitations

- Service insertion over tunnel interface is not supported on Cisco IOS XE Catalyst SD-WAN devices.
- Control policy based service-chain action on locally hosted service-chain is not supported.
- Configuring service-chain and AppQoE on the same device is not supported irrespective of the data-policy or control-policy based actions.
- [Configure Service Chaining, on page 294](#)
- [Service Chaining Configuration Examples, on page 295](#)
- [Monitor Service Chaining, on page 303](#)

Configure Service Chaining

Here is the workflow for configuring service chaining for a device managed by Cisco Catalyst SD-WAN:

1. Service devices are accessed through a specific VRF. In the VPN template that corresponds to the VRF for a service device, configure service chaining, specifying the service type and device addresses. By default, the tracking feature adds each service device status update to the service log. You can disable this in the VPN template.
2. Attach the VPN template to the device template for the device managed by Cisco Catalyst SD-WAN.
3. Apply the device template to the device.

Configure Service Chaining Using Cisco SD-WAN Manager

To configure service chaining for a device.

1. In Cisco SD-WAN Manager, create a VPN template.
2. Click **Service**.
3. In the **Service** section, click **New Service** and configure the following:
 - **Service Type:** Select the type of service that the service device is providing.
 - **IP Address:** IP Address is the only working option.
 - **IPv4 Address:** Enter between one and four addresses for the device.
 - **Tracking:** Determines whether the periodic health updates of the service device are recorded in the system log. Default: On



Note Maximum number of services: 8

4. Click **Add**. The service appears in the table of configured services.

CLI Equivalent for Cisco IOS XE Catalyst SD-WAN Devices

The following table shows how configuration of service chaining by CLI corresponds to configuration in Cisco SD-WAN Manager. CLI configuration differs between Cisco IOS XE Catalyst SD-WAN devices and Cisco vEdge devices. The CLI example below is for a Cisco IOS XE Catalyst SD-WAN device.

CLI (Cisco IOS XE Catalyst SD-WAN device)	Cisco SD-WAN Manager
<pre>service firewall vrf 10</pre>	<p>In Cisco SD-WAN Manager, configure service insertion in the VPN template for a specific VRF—VRF 10 in this example.</p> <p>Select the service type from the drop-down —firewall in this example.</p>

CLI (Cisco IOS XE Catalyst SD-WAN device)	Cisco SD-WAN Manager
<pre>no track-enable</pre> <p>Note Default: enabled</p>	When adding a service in the VPN template Service , select On or Off for Tracking .
<pre>ipv4 address 10.0.2.1 10.0.2.2</pre>	In the VRF template Service , enter one or more IP addresses for the service device providing a specific service.

CLI Example

```
sdwan
  service firewall vrf 10
  ipv4 address 10.0.2.1 10.0.2.2
commit
```

CLI Equivalent for Cisco vEdge Devices

The following table shows how configuration of service chaining by CLI corresponds to configuration in Cisco SD-WAN Manager. CLI configuration differs between Cisco IOS XE Catalyst SD-WAN devices and Cisco vEdge devices. The CLI example below is for a Cisco vEdge device.

CLI (Cisco vEdge device)	Cisco SD-WAN Manager
<pre>vpn 10</pre>	In Cisco SD-WAN Manager, configure service insertion in the VPN template—VPN 10 in this example. Select the service type from the drop-down—firewall in this example.
<pre>service FW address 10.0.2.1</pre>	Select the service type from the drop-down—firewall in this example. Provide one or more addresses for the service device.
<pre>no track-enable</pre> <p>Note Default: enabled</p>	When adding a service in the VPN template Service , select On or Off for Tracking .

CLI Example


```
vpn 10
  service FW address 10.0.2.1
commit
```

Service Chaining Configuration Examples

Service chaining control policies direct data traffic to service devices that can be located in various places in the network before the traffic is delivered to its destination. For service chaining to work, you configure a centralized control policy on the Cisco SD-WAN Controller, and you configure the service devices themselves on the device collocated in the same site as the device. To ensure that the services are advertised to the Cisco SD-WAN Controller, the IP address of the service device must resolve locally.

This topic provides examples of configuring service chaining.

Route Intersite Traffic through a Service



A simple example is to route data traffic traveling from one site to another through a service. In this example, we route all traffic traveling from the device at Site 1 to the device at Site 2 through a firewall service that sits behind a hub (whose system IP address is 100.1.1.1). To keep things simple, all devices are in the same VPN.

For this scenario, you configure the following:

- On the hub router, you configure the IP address of the firewall device.
- On the Cisco SD-WAN Controller, you configure a control policy that redirects traffic destined from Site 1 to Site 2 through the firewall service.
- On the Cisco SD-WAN Controller, you apply the control policy to Site 1.

Here is the configuration procedure:

1. On the hub router, provision the firewall service, specifying the IP address of the firewall device. With this configuration, OMP on the hub router advertises one service route to the Cisco SD-WAN Controller. The service route contains a number of properties that identify the location of the firewall, including the TLOC of the hub router and a service label of `svc-id-1`, which identifies the service type as a firewall. (As mentioned above, before advertising the route, the device ensures that the firewall's IP address can be resolved locally.)

```
sdwan
service firewall vrf 10
  ipv4 address 10.1.1.1
```

2. On the Cisco SD-WAN Controller, configure a control policy that redirects data traffic traveling from Site 1 to Site 2 through the firewall. Then, also on the Cisco SD-WAN Controller, apply this policy to Site 1.


```
policy
  lists
    site-list firewall-sites
      site-id 1
  control-policy firewall-service
    sequence 10
      match route
        site-id 2
      action accept
      set service FW vpn 10
    default-action accept
  apply-policy
    site-list firewall-sites control-policy firewall-service out
```

This policy configuration does the following:

- Create a site list called **firewall-sites** that is referenced in the **apply-policy** command and that enumerates all the sites that this policy applies to. If you later want to scale this policy so that all traffic destined to Site 2 from other sites should also first pass through the firewall, all you need to do is add the additional site IDs to the **firewall-sites** site list. You do not need to change anything in the **control-policy firewall-service** portion of the configuration.
- Define a control policy named **firewall-service**. This policy has one sequence element and the following conditions:
 - Match routes destined for Site 2.
 - If a match occurs, accept the route and redirect it to the firewall service provided by the Hub router, which is located in VPN 10.
 - Accept all nonmatching traffic. That is, accept all traffic not destined for Site 2.
- Apply the policy to the sites listed in **firewall-list**, that is, to Site 1. The Cisco SD-WAN Validator applies the policy in the outbound direction, that is, on routes that it redistributes to Site 1. In these routes:
 - The TLOC is changed from Site 2's TLOC to the hub router's TLOC. This is the TLOC that the Cisco SD-WAN Controller learned from the service route received from the hub router. It is because of the change of TLOC that traffic destined for Site 2 is directed to the hub router
 - The label is changed to **svc-id-1**, which identifies the firewall service. This label causes the hub router to direct the traffic to the firewall device.

When the hub router receives the traffic, it forwards it to the address 10.1.1.1, which is the system IP address of the firewall. After the firewall has finished processing the traffic, the firewall returns the traffic to the hub router, and this router then forwards it to its final destination, which is Site 2.

Route Inter-VPN Traffic through a Service Chain with One Service per Node



A service chain allows traffic to pass through two or more services before reaching its destination. The example here routes traffic from one VPN to another through services located in a third VPN. The services are located behind different hub routers. Specifically, we want all traffic from device-1 in VPN 20 and that is destined for prefix x.x.0.0/16 in VPN 30 on device-2 to go first through the firewall behind Hub-1 and then through the custom service netvc1 behind Hub-2 before being sent to its final destination.

For this policy to work:

- VPN 10, VPN 20, and VPN 30 must be connected by an extranet, such as the Internet
- VPN 10 must import routes from VPN 20 and VPN 30. Routes can be selectively imported if necessary.

- VPN 20 must import routes from VPN 30. Routes can be selectively imported if necessary.
- VPN 30 must import routes from VPN 20. Routes can be selectively imported if necessary.

For this scenario, you configure four things:

- You configure the IP address of the firewall device on the Hub-1 router.
- You configure the IP address of the custom service device on the Hub-2 router.
- On the Cisco SD-WAN Controller, you configure a control policy that redirects traffic destined from Site 1 to Site 2 through the firewall device.
- On the Cisco SD-WAN Controller, you configure a second control policy that redirects traffic to the custom service device.

Here is the configuration procedure:

1. Configure the firewall service on Hub-1. With this configuration, OMP on the Hub-1 router advertises a service route to the Cisco SD-WAN Controller. The service route contains a number of properties that identify the location of the firewall, including the TLOC of the hub router and a service label of `svc-id-1`, which identifies the service type as a firewall.

```
sdwan
service firewall vrf 10
  ipv4 address 10.1.1.1
```

2. Configure the custom service `netvc1` on Hub-2. With this configuration, OMP on the Hub-2 router advertises a service route to the Cisco SD-WAN Controller. The service route contains the TLOC of the Hub-2 and a service label of `svc-id-4`, which identifies the custom service.

```
sdwan
service netvc1 vrf 10
  ipv4 address 2.2.2.2
```

3. Create a control policy on the Cisco SD-WAN Controller for first service in the chain—the firewall—and apply it to Site 1, which is the location of the device-1 router:

```
policy
  lists
    site-list firewall-custom-service-sites
      site-id 1
  control-policy firewall-service
    sequence 10
      match route
        vpn 30
        site-id 2
      action accept
      set service FW
    default-action accept
  apply-policy
    site-list firewall-custom-service-sites control-policy firewall-service out
```

This policy configuration does the following:

- Create a site list called **firewall-custom-service-sites** that is referenced in the **apply-policy** command and that enumerates all the sites that this policy applies to.
- Define a control policy named **firewall-service** that has one sequence element and the following conditions:
 - Match routes destined for both VPN 30 and Site 2.

- If a match occurs, accept the route and redirect it to a firewall service.
 - If a match does not occur, accept the traffic.
- Apply the policy to the sites in the **firewall-custom-service-sites** site list, that is, to Site 1. The Cisco SD-WAN Controller applies this policy in the outbound direction, that is, on routes that it redistributes to Site 1. In these routes:
 - The TLOC is changed from Site 2's TLOC to the Hub-1 router's TLOC. This is the TLOC that the Cisco SD-WAN Controller learned from the service route received from the hub. It is because of the change of TLOC that traffic destined for Site 2 is directed to the Hub-1 router.
 - The label is changed to svc-id-1, which identifies the firewall service. This label causes the Hub-1 router to direct the traffic to the firewall device.

When the Hub-1 router receives the traffic, it forwards it to the address 10.1.1.1, which is the system IP address of the firewall. After the firewall completes processing the traffic, it returns the traffic to the Hub-1 router, which, because of the policy defined in the next step, forwards it to the Hub-2 router.

4. Create a control policy on the Cisco SD-WAN Controller for the second service in the chain, which is the custom service, and apply it to the site of the Hub-1 router:

```

policy
  site-list custom-service
    site-id 3
  control-policy netsvc1-service
    sequence 10
    match route
      vpn 30
      site-id 2
    action accept
      set service netsvc1
    default-action accept
  apply-policy
    site-list custom-service control-policy netsvc1-service out
  
```


This policy configuration does the following:

- Create a site list called **custom-service** that is referenced in the **apply-policy** command and that enumerates all the sites that this policy applies to.
- Define a control policy named **netsvc1-service** that has one sequence element and the following conditions:
 - Match routes destined for both VPN 30 and Site 2.
 - If a match occurs, accept the route and redirect it to the custom service.
 - If a match does not occur, accept the traffic.
- Apply the policy to the sites in the **custom-service** list, that is, to Site 3. The Cisco SD-WAN Controller applies this policy in the outbound direction, that is, on routes that it redistributes to Site 3. In these routes:
 - The TLOC is changed from Site 2's TLOC to the Hub-2 router's TLOC. This is the TLOC that the Cisco SD-WAN Controller learned from the service route received from the Hub-2 router. It is because of the change of TLOC that traffic destined for Site 2 is directed to the Hub-2 router.

- The label is changed to svc-id-4, which identifies the custom service. This label causes the Hub-2 to direct the traffic to the device that is hosting the custom service

When the Hub-2 routers receives the traffic, it forwards it to the address 2.2.2.2, which is the system IP address of the device hosting the custom service. After the traffic has been processed, it is returned to the Hub-2 router, which then forwards it to its final destination, Site 2.

Route Inter-VPN Traffic through a Service Chain with Multiple Services per Node



If a service chain has more than one service that is connected to the same node, that is, both services are behind the same device, you use a combination of control policy and data policy to create the desired service chain. The example here is similar to the one in the previous section, but instead has a firewall and a custom service (netsec-1) behind a single hub router. Here, we want all data traffic from device-1 in VPN 20 destined for prefix x.x.0.0/16 on device-2 in VPN 30 to first go through the firewall at Hub-1, then through the custom service netsec1, also at Hub-1, and then to its final destination.

For this policy to work:

- VPN 10, VPN 20, and VPN 30 must be connected by an extranet, such as the Internet.
- VPN 10 must import routes from VPN 20 and VPN 30. Routes can be selectively imported if necessary.
- VPN 20 must import routes from VPN 30. Routes can be selectively imported if necessary.
- VPN 30 must import routes from VPN 20. Routes can be selectively imported if necessary.

For this scenario, you configure the following:

- On the hub router, you configure the firewall and custom services.
- On the Cisco SD-WAN Controller, you configure a control policy that redirects data traffic from Site 1 that is destined to Site 2 through the firewall.
- On the Cisco SD-WAN Controller, you configure a data policy that redirects data traffic to the custom service.

Here is the configuration procedure:

1. On the hub router, configure the firewall and custom services:

```
sdwan
service firewall vrf 10
  ipv4 address 10.1.1.1
service netsec1 vrf 10
  ipv4 address 2.2.2.2
```


With this configuration, OMP on the hub router advertises two service routes to the Cisco SD-WAN Controller, one for the firewall and the second for the custom service netsvc1. Both service routes contain the TLOC of the Hub-1 router and a service label that identifies the type of service. For the firewall service, the label is svc-id-1, and for the custom service, the label is svc-id-4.

2. On the Cisco SD-WAN Controller, configure a control policy controller to reroute traffic destined for VPN 30 (at Site 2) to firewall service that is connected to Hub-1 (at Site 3), and apply this policy to Site 1:

```
policy
  lists
    site-list device-1
      site-id 1
    control-policy firewall-service
      sequence 10
      match route
        vpn 30
      action accept
      set service FW
  apply-policy
    site-list device-1 control-policy firewall-service out
```

3. On the Cisco SD-WAN Controller, configure a data policy that redirects, or chains, the data traffic received from the firewall device to the custom service netsvc1. Then apply this policy to Hub-1. This data policy routes packets headed for destinations in the network x.x.0.0/16 to the IP address 2.2.2.2, which is the system IP address of the device hosting the custom service.

```
policy
  lists
    site-list device-2
      site-id 2
    site-list Hub-1
      site-id 3
    prefix-list svc-chain
      ip-prefix x.x.0.0/16
    vpn-list vpn-10
      vpn 10
  data-policy netsvc1-policy
    vpn-list vpn-10
    sequence 1
    match
      ip-destination x.x.0.0/16
    action accept
    set next-hop 2.2.2.2
  apply-policy
    site-list Hub-1 data-policy netsvc1-policy from-service
```

Active or Backup Scenario with Service Chaining

When using **set service** action to configure active or backup control policy with **set service** action for service chaining, if total number of available paths (summary of active and standby paths) is more than configured **send-path-limit**, do not set preference directly to routes. Ensure to use **set tloc-list** action to set preferences together with **set service** action. Otherwise, you may see cases where either only active or only backup paths are advertised to a particular spoke router.

For example, in the Cisco SD-WAN Controller OMP table, there are eight active and backup paths. Based on the best-path calculation, the paths are sorted in the following order:

backup1, backup2, backup3, backup4, active1, active2, active3, active4

When **send-path-limit 4** is configured, if you apply the first policy, only the four backup paths are sent. If you apply the second policy, two active and two backup paths are sent.

Example of policy susceptible for failures if **send-path-limit** is lower than total number of active and backup paths:

```
control-policy SET_SERVICE_ACTIVE-BACKUP
sequence 10
match route
prefix-list _AnyIpv4PrefixList
site-list HUBS_PRIMARY
tloc-list INTERNET_TLOCS
!
action accept
set
preference 200
service FW vpn 10
!
!
sequence 20
match route
prefix-list _AnyIpv4PrefixList
site-list HUBS_SECONDARY
tloc-list INTERNET_TLOCS
!
action accept
set
preference 100
service FW vpn 10
!
!
!
default-action accept
!
!
```

Example of the same policy but fixed according to recommendations:

```
policy
lists
tloc-list HUBS_PRIMARY_INTERNET_TLOCS
tloc 10.0.0.0 color biz-internet encap ipsec preference 200
tloc 10.0.0.1 color biz-internet encap ipsec preference 200
!
tloc-list HUBS_SECONDARY_INTERNET_TLOCS
tloc 10.255.255.254 color biz-internet encap ipsec preference 100
tloc 10.255.255.255 color biz-internet encap ipsec preference 100
!
!
control-policy SET_SERVICE_ACTIVE-BACKUP_FIXED
sequence 10
match route
prefix-list _AnyIpv4PrefixList
site-list HUBS_PRIMARY
tloc-list INTERNET_TLOCS
!
action accept
set
service FW vpn 10 tloc-list HUBS_PRIMARY_INTERNET_TLOCS
!
!
sequence 20
match route
```

```

prefix-list _AnyIpv4PrefixList
site-list HUBS_SECONDARY
tloc-list INTERNET_TLOCS
!
action accept
set
  service FW vpn 10 tloc-list HUBS_SECONDARY_INTERNET_TLOCS
!
!
!
default-action accept
!
!
```

Monitor Service Chaining

You can monitor different aspects of service chaining on hub and spoke devices.



Note Configuring a service device to operate as part of the service chain is called service insertion.

- On a hub device, view the configured services.
 - From the Cisco SD-WAN Manager menu:

View the configured services on the **Real Time** monitoring page (**Monitor** > **Devices** > *hub-device* > **Real Time**). For **Device Options**, select **OMP Services**.

Cisco vManage Release 20.6.x and earlier: View the configured services on the **Real Time** monitoring page (**Monitor** > **Network** > *hub-device* > **Real Time**). For **Device Options**, select **OMP Services**.
- On a spoke device, view the details of the service chain path.
 - **Using Cisco SD-WAN Manager:**

View the service chain path on the **Traceroute** page (**Monitor** > **Devices** > *spoke-device* > **Troubleshooting** > **Connectivity** > **Trace Route**). Enter the destination IP, VPN, and source interface for the desired path.

Cisco vManage Release 20.6.x and earlier: View the service chain path on the **Traceroute** page (**Monitor** > **Network** > *spoke-device* > **Troubleshooting** > **Connectivity** > **Trace Route**). Enter the destination IP, VPN, and source interface for the desired path.
 - **Using the CLI:**

Use the **traceroute** command. For information, see the [Cisco Catalyst SD-WAN Command Reference](#).

Example: View a Service Chain Path Between Two Spoke Devices

The following example shows how to view the path between two spokes before and after adding a service chain between them, using Cisco SD-WAN Manager or the CLI.

For clarity, the example presents a scenario of two spoke devices, a hub device, and a service device providing a firewall service, and shows how to configure the firewall service chain.

Here are the details for each device in the scenario:

Device	Address
Hub, through interface ge0/4	10.20.24.15
Spoke 1	10.0.3.1
Spoke 2	10.0.4.1
Service device (firewall service)	10.20.24.17

Configuration of the three devices:

```

Hub
====
vm5# show running-config vpn 1
vpn 1
 name ospf_and_bgp_configs
 service FW
  address 10.20.24.17
 exit
router
 ospf
  router-id 10.100.0.1
  timers spf 200 1000 10000
  redistribute static
  redistribute omp
  area 0
   interface ge0/4
   exit
  exit
 !
 !
interface ge0/4
 ip address 10.20.24.15/24
 no shutdown
 !
interface ge0/5
 ip address 10.30.24.15/24
 no shutdown
 !
 !

```

```

Spoke 1
=====
vpn 1
 name ospf_and_bgp_configs
 interface ge0/1
  ip address 10.0.3.1/24
  no shutdown
 !
 !

```

```

Spoke2
=====
vpn 1
 interface ge0/1
  ip address 10.0.4.1/24
  no shutdown

```

!
!

1. Without Service Insertion:

At this point, no service insertion policy has been configured, so executing **traceroute** on Spoke 1 to display the path details to Spoke 2 (10.0.4.1) shows a simple path to Spoke 2:

→ **Spoke 2 (10.0.4.1)**

```
vm4# traceroute vpn 1 10.0.4.1
Traceroute 10.0.4.1 in VPN 1
traceroute to 10.0.4.1 (10.0.4.1), 30 hops max, 60 byte packets
 1 10.0.4.1 (10.0.4.1) 7.447 ms 8.097 ms 8.127 ms
```

Similarly, viewing the Traceroute page in Cisco SD-WAN Manager shows a simple path from Spoke 1 to Spoke 2.

2. With Service Insertion:

The following Cisco SD-WAN Controller policy configures service insertion for a firewall service, using the firewall service device described above.

```
vm9# show running-config policy
policy
  lists
    site-list firewall-sites
      site-id 400
  !
  !
  control-policy firewall-services
  sequence 10
  match route
    site-id 600
  !
  action accept
  set
    service FW vpn 1
  !
  !
  !
  default-action accept
  !
  !
vm9# show running-config apply-policy
apply-policy
  site-list firewall-sites
  control-policy firewall-services out
  !
  !
```

After configuring the service insertion, executing **traceroute** on Spoke 1 (10.0.3.1) to display the path details to Spoke 2 (10.0.4.1) shows this path:

→ **Hub (10.20.24.15) → Firewall service device (10.20.24.17) → Hub (10.20.24.15) → Spoke 2 (10.0.4.1)**

```
Traceroute -m 15 -w 1 -s 10.0.3.1 10.0.4.1 in VPN 1
traceroute to 10.0.4.1 (10.0.4.1), 15 hops max, 60 byte packets
 1 10.20.24.15 (10.20.24.15) 2.187 ms 2.175 ms 2.240 ms
 2 10.20.24.17 (10.20.24.17) 2.244 ms 2.868 ms 2.873 ms
 3 10.20.24.15 (10.20.24.15) 2.959 ms 4.910 ms 4.996 ms
 4 10.0.4.1 (10.0.4.1) 5.045 ms 5.213 ms 5.247 ms
```

Similarly, viewing the **Traceroute** page in Cisco SD-WAN Manager shows each step of the path from Spoke 1 to Spoke 2, through the hub and firewall service device.



CHAPTER 22

Lawful Intercept

The Lawful Intercept 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 48: Feature History

Feature Name	Release Information	Description
Encryption of Lawful Intercept Messages	Cisco IOS XE Catalyst SD-WAN Release 16.12.1b	This feature encrypts lawful intercept messages between a Cisco IOS XE Catalyst SD-WAN device and a media device using static tunnel information.

- [Information About Lawful Intercept, on page 307](#)
- [Prerequisites for Lawful Intercept, on page 310](#)
- [Install Lawful Intercept using Cisco Catalyst SD-WAN Manager, on page 311](#)
- [Lawful Intercept MIBs, on page 312](#)
- [Restrict Access to Trusted Hosts \(Without Encryption\), on page 312](#)
- [Restrict Trusted Mediation Device, on page 313](#)
- [Configure Lawful Intercept, on page 313](#)
- [Configure Lawful Intercept Using CLI, on page 313](#)
- [Encrypt Lawful Intercept Traffic , on page 314](#)
- [Verify Static Tunnel with Media Device Gateway, on page 316](#)

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. Cisco SD-WAN Manager ensures that Cisco SD-WAN Manager 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 30: Cisco Catalyst SD-WAN Lawful Intercept Workflow

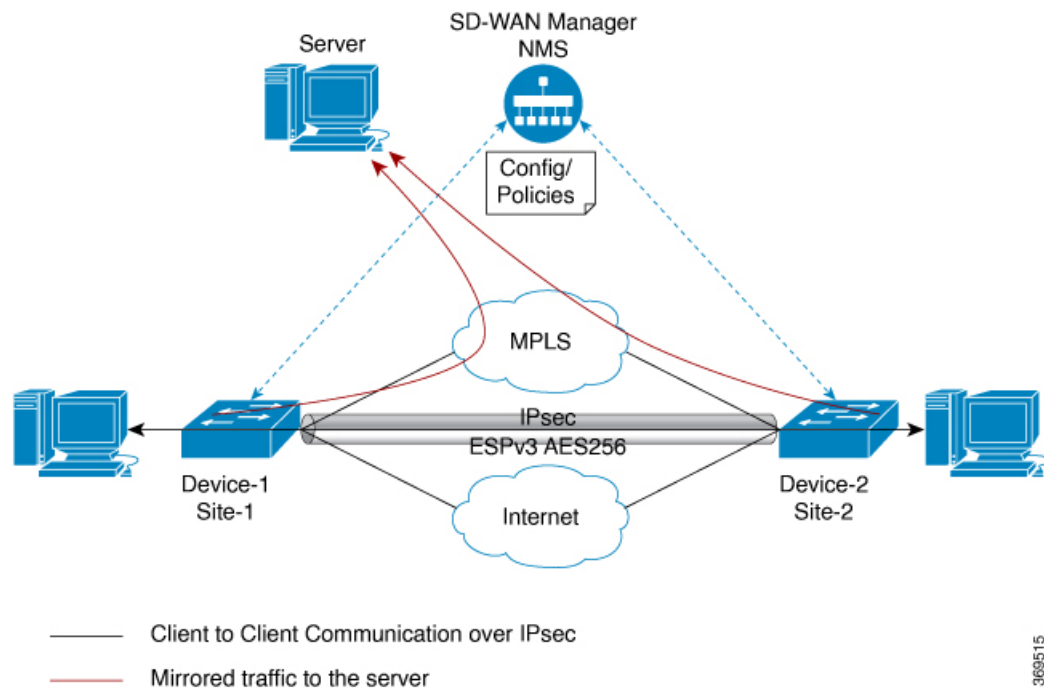



Figure 31: Cisco Catalyst SD-WAN Lawful Intercept Process



369514

Licence-based Lawful Intercept

Cisco Catalyst SD-WAN solution is a term-based licensed feature. This feature license enables the Cisco SD-WAN Manager component of the Cisco Catalyst SD-WAN solution and allows the customer to access the Lawful Intercept function. Once the Lawful Intercept license is enabled on the solution, Cisco SD-WAN Manager provides a new privilege in the Manage Users menu of the Cisco SD-WAN Manager 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 SD-WAN Manager.
2. Create an lawful intercept admin (liadmin) user on Cisco SD-WAN Manager. The **liadmin** user must be associated with the user group, Basic.
3. Login to Cisco SD-WAN Manager as **liadmin** user and configure Lawful Intercept specific templates.

4. Cisco SD-WAN Manager automatically pushes templates to all Cisco IOS XE Catalyst SD-WAN devices with Lawful Intercept compatible images.
5. Configuration is pushed to device from Cisco SD-WAN Manager 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 Catalyst 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 Cisco SD-WAN Manager using the VPN Interface Ethernet page of Feature Template. See VPN Interface Ethernet section in Templates topic.

Install Lawful Intercept using Cisco Catalyst SD-WAN Manager



Note The following process must be repeated for every Cisco SD-WAN Manager node.

1. Connect to a Cisco SD-WAN Manager 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 SD-WAN Manager.
7. Login to Cisco SD-WAN Manager 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 Catalyst 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 Catalyst 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 Cisco SD-WAN Manager 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 Catalyst 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 Catalyst 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 Catalyst 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
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
!

```

□ pre-shared key, should be same on cEdge

□ Tunnel address

□ MD GW phy interface

□ cEdge wan interface

Verify Static Tunnel with Media Device Gateway

The IPsec tunnel between the Cisco IOS XE Catalyst 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 23

Lawful Intercept 2.0

Table 49: Feature History

Feature Name	Release Information	Description
Lawful Intercept 2.0	Cisco vManage Release 20.9.1	This feature introduces Lawful Intercept Version 2.0. In the Lawful Intercept 2.0 feature, key information is provided to a law enforcement agency (LEA) by the Cisco Catalyst SD-WAN routers and control components so that they can decrypt the Cisco Catalyst SD-WAN IPsec traffic captured by the Managed Service Provider (MSP). This helps the LEA decrypt the encrypted network traffic information. For information on Lawful Intercept 1.0, see the chapter Lawful Intercept in the Cisco Catalyst SD-WAN Policies Configuration Guide.
Lawful Intercept 2.0 Enhancements	Cisco vManage Release 20.10.1	This feature enhances the Cisco SD-WAN Manager GUI and the troubleshooting options available for the Lawful Intercept feature in Cisco Catalyst SD-WAN. <ul style="list-style-type: none"> • Cisco SD-WAN Manager GUI enhancements: <ul style="list-style-type: none"> • A Sync to vSmart button to synchronize a newly created intercept configuration with the Cisco SD-WAN Controller. • A toggle button to enable or disable an intercept. • A progress page to display the status of synchronization and activation. • A red dot on the task list icon in the Cisco SD-WAN Manager toolbar to indicate any new lawful intercept tasks. • A task list pane to view a list of active and completed lawful intercept tasks. • An intercept retrieve option Get IRI to retrieve key information or Intercept Related Information (IRI) from the Cisco SD-WAN Controller. • Ability to troubleshoot Cisco SD-WAN Controller and Cisco SD-WAN Manager using the debug logs and admin tech files.

Feature Name	Release Information	Description
Lawful Intercept 2.0 Enhancements	Cisco Catalyst SD-WAN Manager Release 20.12.1	This feature extends Lawful Intercept to multitenancy mode, and provides support for Cisco SD-WAN Manager clusters. For more information on Cisco SD-WAN Manager clusters, see Cluster Management .


- [Information About Lawful Intercept 2.0, on page 318](#)
- [Prerequisites for Cisco Catalyst SD-WAN Lawful Intercept 2.0, on page 320](#)
- [Benefits of Cisco Catalyst SD-WAN Lawful Intercept 2.0, on page 320](#)
- [Configure Lawful Intercept 2.0 Workflow, on page 320](#)
- [Create a Lawful Intercept Administrator, on page 320](#)
- [Create a Lawful Intercept API User, on page 321](#)
- [Create an Intercept, on page 321](#)
- [Retrieve an Intercept, on page 323](#)
- [Troubleshooting Cisco SD-WAN Controller for Lawful Intercept from Cisco SD-WAN Manager, on page 323](#)

Information About Lawful Intercept 2.0

Cisco Catalyst SD-WAN's Lawful Intercept feature allows an LEA to get a copy of network traffic for analysis or evidence. This is also referred as traffic mirroring. See the chapter [Lawful Intercept](#) in the Cisco Catalyst SD-WAN Policies Configuration Guide.

From Cisco vManage Release 20.9.1, Cisco Catalyst SD-WAN implements a new architecture for Lawful Intercept, as shown in the following figure.

Figure 32: Lawful Intercept 2.0 Architecture



The following are the characteristics of the new architecture:

- Traffic mirroring is outside the scope of Cisco Catalyst SD-WAN. The LEA works with the corresponding service provider to capture network traffic for mirroring.



Note In the illustration above, the service provider is an underlay connection and the IPsec tunnel is an overlay connection.

- Because the captured network traffic is encrypted, Cisco SD-WAN Manager and Cisco SD-WAN Controller provide key information to the LEA.
- The LEA retrieves the keys from Cisco SD-WAN Manager to decrypt Cisco Catalyst SD-WAN IPsec traffic. The LEA ensures that they retrieve key information is retrieved during each rekey period. The rekey period is provided by the service provider. For more information about retrieving keys, see [Retrieve an Intercept, on page 323](#). For information on rekey period, see [Configure Data Plane Security Parameters](#).

A Lawful Intercept administrator is solely responsible for configuring intercepts and creating Lawful Intercept API users who perform Lawful Intercepts. A Cisco SD-WAN Manager administrator can create an account for the Lawful Intercept administrator; the administrator must be a member of the **li-admin** group. For more information about creating an account for a Lawful Intercept administrator, see [Create Lawful Intercept Administrator](#).

Prerequisites for Cisco Catalyst SD-WAN Lawful Intercept 2.0

- A Cisco SD-WAN Controller must be set to **Manager mode**.
- For more information about decrypting the IPsec traffic in Cisco Catalyst SD-WAN, contact Cisco Support or Cisco Sales team.

Benefits of Cisco Catalyst SD-WAN Lawful Intercept 2.0

- It is not necessary to configure edge devices for Lawful Intercepts.



Note To configure an intercept, an administrator must select the edge devices that have to be included in the intercept. This is necessary because the key information that is retrieved from Cisco SD-WAN Manager also includes the keys for the selected devices.

- The service provider captures the data traffic for interception. Traffic is not intercepted from the edge devices.

Configure Lawful Intercept 2.0 Workflow



Note The Lawful Intercept feature can be configured only through Cisco SD-WAN Manager, and not through the CLI.

To configure Lawful Intercept in Cisco SD-WAN Manager, perform the following steps:

1. [Create Lawful Intercept Administrator](#)
2. [Create Lawful Intercept API User](#)
3. [Create an Intercept](#)

Create a Lawful Intercept Administrator

Using the Admin account in Cisco SD-WAN Manager, create an account for the Lawful Intercept administrator.

1. From the Cisco SD-WAN Manager menu, choose **Administration > Lawful Intercept**.
2. Click **Add User** to create a Lawful Intercept administrator user account.
3. In the **Full Name** field, enter a full name for the Lawful Intercept administrator.

4. In the **User Name** field, enter a user name for the Lawful Intercept administrator. The user name must be prefixed with **li-**.
5. In the **Password** field, enter a password for the Lawful Intercept administrator.
6. Confirm the password in the **Confirm Password** field.
7. From the **User Group** drop-down list, choose **li-admin**, and then click **Add**.

The newly created Lawful Intercept administrator user account is displayed in the **Users** window.

Create a Lawful Intercept API User

The Lawful Intercept API User account is for those users of LEA who log in and retrieve key information using Cisco SD-WAN Manager's REST API. These are the users who perform a lawful intercept of the Cisco Catalyst SD-WAN IPsec traffic.

The LEA use

`https://{vmanage_ip}/dataservice/li/intercept/retrieve/<intercept_id>` to retrieve the key information.

To create a Lawful Intercept API user, perform the following steps:

1. Log in to Cisco SD-WAN Manager as a Lawful Intercept administrator.



Note When a Lawful Intercept administrator logs in to Cisco SD-WAN Manager, only the **Monitor** and **Administration** options are available in the Cisco SD-WAN Manager menu.

2. From Cisco SD-WAN Manager menu, choose **Administration** > **Lawful Intercept**.
3. Click **Add User** to create an Lawful Intercept API user account.
4. In the **Full Name** field, enter a full name for the Lawful Intercept API user.
5. In the **User Name** field, enter a user name for the Lawful Intercept API user. The user name must be prefixed with **li-**.
6. In the **Password** field, enter a password for the Lawful Intercept API user.
7. Confirm the password in the **Confirm Password** field.
8. From the **User Group** drop-down list, choose **li-api**, and click **Add**.

The newly created Lawful Intercept API user account is displayed in the **Users** window. The LEA can log in to Cisco SD-WAN Manager using the Lawful Intercept API user account to retrieve key information.

Create an Intercept

Minimum supported release: Cisco vManage Release 20.9.1 and Cisco Catalyst SD-WAN Control Components Release 20.9.1

Configure an intercept to collect intercept data. To configure an intercept, do the following:

1. From the Cisco SD-WAN Manager menu, choose **Administration > Lawful Intercept**.
2. Click the **Intercepts** tab, and then click **Add Intercepts**.
3. Cisco Catalyst SD-WAN Manager Release 20.12.1 and later releases:
From the **Tenant** drop-down list, choose a tenant. For more information about adding a tenant, see [Add a New Tenant](#).
4. In the **Intercept ID** field, enter a number. Enter a minimum of two digits and a maximum of 25 digits.
5. In the **Description** field, enter a description for the intercept.
6. By default the **Enable** toggle button is enabled. However, the intercept remains in an inactive state after it is created.
7. Click **Next**.
In single-tenant mode, the **Add Edge Devices** pop-up window displays all the edge devices in the Cisco Catalyst SD-WAN network.
Cisco Catalyst SD-WAN Manager Release 20.12.1 and later releases:
In multi-tenant mode, the **Add Edge Devices** pop-up window displays all the single-tenant edge devices associated with the selected tenant.
8. Click one or more edge device names to add to the intercept and click **Next**.
Cisco SD-WAN Manager provides the keys for the edge devices selected here.



Note Specify an intercept warrant for all the edge devices that are added to the intercept.

When an edge device is added for interception, all its peer devices, which are connected in the same network, are also available for Lawful Interception.

9. The **Add LI API users** pages displays all the LI-API users created by the Lawful Intercept administrator.
10. Click one or more user names to add to the intercept. The users selected here can retrieve key information that is required for interception from Cisco SD-WAN Manager. For information on how keys are retrieved for an intercept, see [Retrieve an Intercept](#).
11. Click **Summary**
The summary of the intercept is displayed.
12. Click **Submit**. The **Intercepts** page displays the configured intercept.
13. Click **Sync to vSmart** to synchronize the configured intercept configuration in Cisco SD-WAN Manager with Cisco SD-WAN Controller.

A progress page displays the status of the synchronization and activation. After successful synchronization, the **Activate State** field displays a green check mark.



Note The **Activate State** field displays a green check mark status only if Cisco SD-WAN Controller is set to **Manager** mode.

If there are any additional Lawful Intercept tasks, a red dot is displayed on the task list icon in the Cisco SD-WAN Manager toolbar. Click the tasks list icon to view a list of all the active and completed Lawful Intercept tasks. You can view up to 500 latest Lawful Intercept tasks.

If an intercept is modified, the **Sync to vSmart** button is enabled. Click **Sync to vSmart** to synchronize the intercept configuration in Cisco SD-WAN Manager with Cisco SD-WAN Controller.



Note The **Sync to vSmart** button is enabled only when a new intercept is created, or when an intercept is edited or deleted.

To retrieve key information that is required for interception, click **...**, and then click **Get IRI**. The IRI is retrieved from Cisco SD-WAN Controller and displayed in Cisco SD-WAN Manager.

Retrieve an Intercept

An LEA is responsible to periodically retrieve key information because this information is required to decrypt the traffic captured by the MSP.

An LEA can retrieve key information by using [Cisco Catalyst SD-WAN Manager REST APIs](#).

1. An LEA logs in to Cisco SD-WAN Manager as a Lawful Intercept API user.
2. After a Lawful Intercept API user is authenticated, the LEA sends a request using the Cisco SD-WAN Manager REST APIs specifying the intercept ID that it wants to get the key information for.
3. When a request from the LEA is received by Cisco SD-WAN Manager, Cisco SD-WAN Manager forwards the request to the Cisco SD-WAN Controller on which intercepts are configured.
4. Cisco SD-WAN Controller then retrieves the key information for the specified intercept ID and returns the key information to Cisco SD-WAN Manager in JSON format.

Troubleshooting Cisco SD-WAN Controller for Lawful Intercept from Cisco SD-WAN Manager

Minimum supported release: Cisco vManage Release 20.10.1 and Cisco Catalyst SD-WAN Control Components Release 20.10.1

Cisco SD-WAN Manager offers debug logs and admin tech files to troubleshoot any issues in Cisco SD-WAN Controller and Cisco SD-WAN Manager.

Debug Logs

Use debug logs to troubleshoot Cisco SD-WAN Controller from Cisco SD-WAN Manager.

To view the debug logs in Cisco SD-WAN Manager:

1. From Cisco SD-WAN Manager menu, choose **Administration > Lawful Intercept**.
2. Click the **Devices** tab.

3. Click ... adjacent to the device that you want to view the debug logs, and choose **Debug Log**.
4. In the **Log Files** drop-down list, choose the name of the log file.

The lower part of the window displays the log information.

Admin Tech Files

Use debug logs and admin tech files to troubleshoot Cisco SD-WAN Manager and Cisco SD-WAN Controller from Cisco SD-WAN Manager. For more information about generating an admin tech file, see [Generate Admin-Tech Files](#).



CHAPTER 24

Troubleshoot Cisco Catalyst SD-WAN Policies

- [Overview](#), on page 325
- [Support Articles](#), on page 325
- [Feedback Request](#), on page 326
- [Disclaimer and Caution](#), on page 326

Overview

This chapter provides links to documents authored by Cisco subject matter experts (SMEs). They aim to help you resolve technical issues without requiring a support ticket. If these documents are unable to resolve your issue, we recommend visiting the applicable [Cisco Community](#). There is a wealth of information and advice available from fellow Cisco customers who may have experienced this issue already and provided a solution. If you are not able to find a resolution on the Community, it may be best that you raise a support ticket at [Cisco Support](#). In cases where a support ticket has to be raised, these documents provide guidance about the data that should be collected and added to the support ticket. Specify the support document you referred, and TAC can create an improvement request with the document owner.

Support Articles

The documents in this section were created using specific software and hardware listed in the Components Used section of each article. However, this does not mean that they are limited to what is listed in Components Used, and generally remain relevant for later versions of software and hardware. Note that there could be some changes in the software or hardware that can cause commands to stop working, the syntax to change, or GUIs and CLIs to look different from one release to another.

The following are the support articles associated with this technology:

Document	Description
Cisco Catalyst SD-WAN - Configure Route Leaking	This video shows how to configure Route Leaking in Cisco Catalyst SD-WAN.
Collect an Admin-Tech in Cisco Catalyst SD-WAN Environment and Upload to TAC Case	This document describes how to initiate an <code>admin-tech</code> in a Cisco Catalyst SD-WAN environment.

Document	Description
Configure AAR Policy on Cisco Catalyst SD-WAN	This video shows how to configure Application Aware Routing Policy on Cisco Catalyst SD-WAN.
Configure Cisco Catalyst SD-WAN Router to Restrict SSH Access	This document describes the process to restrict SSH connection to a Cisco Catalyst SD-WAN router.
Configure a Control Policy for Region Topology	This video shows how to configure a control policy for regional topology so the sites on different regions can reach the internet through the closest DC.
Configure Active/Standby Hub and Spoke Topology on Cisco Catalyst SD-WAN	This document describes the steps to configure and validate an Active Standby Hub and Spoke Topology on Cisco Catalyst SD-WAN.
Configure a Data Policy to Overwrite a Control Policy	This video shows how to configure a data policy to complete the task: Users from Sites in Region 1 must access AWS networks through DC in Region 2. Everything else must flow via DC on Region 1.
Determine Policy Drops on cEdge with FIA Trace	This video shows how to determine policy traffic drops on cEdge with FIA Trace.
Troubleshoot Cisco Catalyst Controller Policy Push Activation Errors	This document describes some common errors seen during the activation of a Cisco SD-WAN Controller policy from Cisco SD-WAN Manager in an Cisco Catalyst SD-WAN overlay network.
Understand BFD Protocol Relationship with App-Aware Routing	This document describes the relationship that exists between the BFD Hello packets and the App-Aware Routing Tunnel statistics.

Feedback Request

Your input helps. A key aspect to improving these support documents is customer feedback. Note that these documents are owned and maintained by multiple teams within Cisco. If you find an issue specific to the document (unclear, confusing, information missing, etc):

- Provide feedback using the **Feedback** button located at the right panel of the corresponding article. The document owner will be notified, and will either update the article, or flag it for removal.
- Include information regarding the section, area, or issue you had with the document and what could be improved. Provide as much detail as possible.

Disclaimer and Caution

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.