Cisco Catalyst SD-WAN Policies Configuration Guide, Releases 26.x and Later

Cisco SD-WAN
Updated October 21, 2026

© 2026 Cisco Systems, Inc. All rights reserved.

Full Cisco Trademarks with Software License

Full Cisco Trademarks with Software License

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 NONINFRINGEMENT 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)

Topics included

1 Read Me First....11

2 Policy Overview....15

Policy architecture....17
Centralized control policy architecture....17
Policy validation for Cisco Catalyst SD-WAN....24
Cisco Catalyst SD-WAN Controller Policy components....25
TLOC attributes used in policies....27
Cisco Catalyst SD-WAN Route attributes used in policies....27
Best practice for Cisco Catalyst SD-WAN Controller policy processing and application....28
Cisco Cisco Catalyst SD-WAN Controller policy operation....29
Control policy....30
Data policies....31
VPN membership policy operation....33
Configure and execute Cisco SD-WAN Controller policies....34
Policy status and health monitoring....35

3 Centralized Policy....37

Feature history for centralized policy....38
Overview of centralized policies....39
Types of centralized policies....39
Configure centralized policies....40
Configure centralized policies using the CLI....40
Centralized policies using classic policies....46
Centralized policies configuration examples....84

4 Localized policy....95

Localized policies....96
QoS for router-generated Cisco SD-WAN Manager traffic....96
Localized policies....96
Types of localized policies....96
Configure policies using a configuration group....98
Configure Forwarding Classes and QoS using a policy group....99
Configure objects and policies using a configuration group....100
Configure ACLs using a configuration group....107
Match parameters....111
Action parameters....114
Configure localized policy for IPv4 using the CLI....115

Configure localized policy for IPv6 using the CLI....118
Localized data policy configuration examples....120
Configure localized policy using classic policies....121
Start the policy configuration wizard....121
Configure groups of interest for localized policy....122
Configure forwarding classes/QoS....124
Configure ACLs....125
Configure route policies....128
Match parameters....129
Action parameters....132
Configure policy settings....133
Apply localized policy in a device template....134
Activate a localized policy....135

5 Default AAR and QoS policies....137
Default AAR and QoS policies....138
Default AAR and QoS policies....138
Prerequisites for default AAR and QoS policies....139
Restrictions for default AAR and QoS policies....139
Supported devices for default AAR and QoS policies....139
Default AAR and QoS policy use cases....139
Configure default AAR and QoS policies using Cisco SD-WAN Manager....140
Monitor default AAR and QoS policies....144

6 Redirect DNS in Cisco Catalyst SD-WAN Service-Side VPNs....147
Feature history for redirect DNS in a service-side VPN....148
Redirect DNS in a service-side VPN....148
Benefits of redirect DNS....149
Restrictions for redirect DNS in a service-side VPN....149
Use cases for redirect DNS in a service-side VPN....150
Configure redirect DNS in a service-side VPN....151
Configure redirect DNS in service-side VPN using policy groups....151
Configure redirect DNS in a service-side VPN using CLI commands....153
Configure redirect DNS in a service-side VPN using classic policies....155
Configuration examples for redirect DNS....156
Verify redirect DNS in a service-side VPN....157

7 Device Access Policy....159
Feature history for device access policy....160
Device access policies in Cisco SD-WAN Manager....160
Configure device access policy using the CLI....161
Configure IPv4 device access policy using configuration groups....162
Configure IPv6 device access policy using configuration groups....163

Configure device access policy using localized policy....164
ACL statistics and counters examples....166
ACL policy verification on an SNMP server....167
ACL policy verification on SSH....170

8 Application Intelligence Engine Flow....173
SAIE flow overview....174
SAIE flow configuration using Cisco SD-WAN Manager....174
Apply centralized policy for SAIE flow....175
Monitor running applications....175
View SAIE applications....176
Action parameters for configuring SAIE flow....177
Configure SAIE using the CLI....178

9 Application-Aware Routing....183
Feature history for application-aware routing....184
Application-aware routing....185
Components of application-aware routing....186
Fundamentals of application-aware routing policy....187
Default action of application-aware routing policy....188
Application-aware routing for multicast protocols....189
Restrictions for multicast protocols....189
SLA....189
SLA-preferred colors....190
SLA classes....190
Classification of tunnels into SLA classes....191
Measure loss, latency, and jitter....192
Calculate average loss, latency, and jitter....192
Configure SLA....193
Configure traffic rules....194
Best tunnel path....205
Variance configuration for best tunnel path....205
Verify variance configuration for best tunnel path....206
Per-class application-aware routing....207
Application probe classes....208
Configure an application probe class....209
Add app-probe-class to an SLA class....210
Default DSCP values....211
Configure default DSCP on a Cisco BFD template....211
Configure BFD template using configuration groups 212
Apply policies to sites and VPNs....212
Restrictions for policy application....213
Apply policies to sites and VPNs using classic policies....213
Apply policies to sites and VPNs using CLI commands....214

Verify applied policy....214
Activate an application-aware routing policy....214
Configure application-aware routing policies....215
Configure application-aware routing using CLI commands....215
Configure application-aware routing policies using classic policies....217
Configure an application-aware routing policy using CLI....218
Data plane tunnel performance monitoring....224
Parameters for sliding window size....225
Monitor tunnel statistics and routing paths....226
Enable application visibility on Cisco IOS XE Catalyst SD-WAN devices....226

10 Enhanced Application-Aware Routing....227
Enhanced Application-Aware Routing in Cisco SD-WAN....228
Pfr measurements....228
Application-aware routing design and measurement intervals....229
Benefits of enhanced application-aware routing....230
Guidelines of enhanced application-aware routing....230
Compatibility with Cisco IOS XE catalyst SD-WAN devices without enhanced application-aware routing....230
Restrictions for enhanced application-aware routing....231
Prerequisites for enhanced application-aware routing....231
Supported devices for enhanced application-aware routing....231
Configure enhanced application-aware routing....231
Configure enhanced application-aware routing using a configuration group....232
Configure enhanced application-aware routing using CLI commands....232
Configure enhanced application-aware routing using templates....233
Verify enhanced application-aware routing configuration....234
Monitor enhanced application-aware routing using Cisco Catalyst SD-WAN Manager....236
Enhanced application-aware routing troubleshooting....236

11 Traffic Flow Monitoring....239
Feature history for traffic flow monitoring....240
Traffic flow monitoring....242
Traffic flow monitoring with cflowd....242
IPFIX information elements for Cisco IOS XE Catalyst SD-WAN devices....243
Flexible NetFlow for VPN0 interface....244
Flexible NetFlow export spreading....246
Flexible NetFlow export of BFD metrics....247
Cflowd traffic flow monitoring with SAIE flows....248
Maximum FNF record rate for aggregated data....249
Restrictions for traffic flow monitoring....250
Restrictions for enabling collect loopback in flow telemetry when using loopbacks as TLOCs....250
Configure traffic flow monitoring....250
Configure traffic flow monitoring on Cisco IOS XE Catalyst SD-WAN devices using policy groups....250
Configure traffic flow monitoring on Cisco IOS XE Catalyst SD-WAN devices using classic policies....252

Configure cflowd traffic flow monitoring using the CLI....259
Configure flexible netflow on VPN0 interface....261
Configure flexible NetFlow with the export of BFD metrics using the CLI....262
Configuration examples for flexible NetFlow export of BFD metrics....263
Apply and enable cflowd policy....265
Configure cflowd traffic flow monitoring....266
Configure the maximum FNF record rate for aggregated data using the CLI....272
Verify traffic flow monitoring....273
Verify collect loopback....273
Verify interface binding on the device....276
Verify the flexible NetFlow configuration on the VPN0 interface....276
Verify the flexible NetFlow configuration with the export of BFD metrics....280

12 Application Performance Monitor....283
Application performance monitor....284
Restrictions for performance monitoring implementation....285
Configure application performance monitor....285
Performance monitoring configuration verification....287

13 Enhanced Policy Based Routing....293
Enhanced policy based routing....294
EPBR configuration....295
EPBR monitoring commands....298

14 Forward Error Correction....301
Supported devices for forward error correction....303
Configure forward error correction for a policy....303
Forward error correction status monitoring using CLI....304

15 Packet Duplication....305
Feature history for packet duplication....306
Packet duplication....306
Configure packet duplication....307
Configure packet duplication using policy groups....307
Configure packet duplication on additional tunnels using the CLI....308
Configure packet duplication using classic policies....309
Configure underlay fragmentation using configuration groups....309
Restrictions for packet duplication....310
Monitor packet duplication statistics for a device using the CLI....310
Monitor packet duplication statistics for a device using Cisco SD-WAN Manager....311
Monitor tunnel information for a device....312

16 Policy Configuration Tagging.....315
Feature history for policy configuration tagging.....316
Supported devices for policy configuration tagging.....316
Restrictions for policy configuration tagging.....317
Policy configuration tagging.....317
Benefits of policy configuration tagging.....320
Configure policy configuration tagging using CLI commands.....321
Tag-instances verification.....323

17 Custom Applications.....327
Custom applications.....328
Restrictions for custom applications.....330
Configure custom applications using Cisco SD-WAN Manager.....331
Verify custom applications.....333

18 Service Insertion.....335
Feature history for service insertion.....336
Service insertion.....336
Restrictions for service insertion.....338
Use cases for service insertion.....338
Service insertion configuration methods.....339
Configure service insertion using a CLI template.....339
Configure service insertion using the workflow library.....340
Configure service chain actions in a data policy.....341
Traffic steering to a service chain.....342
How traffic steering using a control policy works.....342
How data policies steer traffic.....344
Traffic steering using an interface access control list.....345
Configure path preference.....347
How service chains share across user VPNs.....347
How separate interfaces for transmitted and received traffic work.....348
Service chaining trusted and untrusted traffic.....349
How service chains work between two routers.....350
Fall back and restrict behavior configuration for service chain traffic.....350
How service chain interfaces work.....351
Service chaining with software defined cloud interconnect bring your own service.....352

19 Service Insertion for Equinix.....353
Feature history for service insertion for Equinix.....354
Service insertion for Equinix.....354
Requirements for service insertion for Equinix.....354
Service insertion configuration workflow for Equinix.....355

Define and configure service chain....355
Instantiate service chain....356
Attach a service chain to the Cisco SD-WAN device....356
Detach a service chain instance....357
Configuration verification for Equinix service insertion....357

20 Service Chaining....359
Service chaining in Cisco SD-WAN....360
Service provisioning and chaining in Cisco Catalyst SD-WAN....361
Restrictions for service chaining....362
Configure service chaining....362
Configure service chaining for Cisco IOS XE Catalyst SD-WAN devices using CLI commands....363
Configure service chaining using templates....364
Service chain configuration examples....364
Monitor service chaining 371

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
- User Documentation for Cisco SD-WAN Release 20

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.

2 Policy Overview

Topics:

- Policy architecture
- Policy validation for Cisco Catalyst SD-WAN
- Cisco Catalyst SD-WAN Controller Policy components
- Best practice for Cisco Catalyst SD-WAN Controller policy processing and application
- Cisco Cisco Catalyst SD-WAN Controller policy operation
- Configure and execute Cisco SD-WAN Controller policies
- Policy status and health monitoring

Describes how policies influence data traffic and routing information within the Cisco IOS XE Catalyst SD-WAN overlay network.

A policy is a network control mechanism that

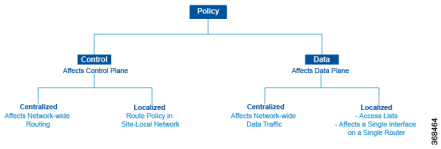
- influences the flow of data traffic and routing information among Cisco IOS XE Catalyst SD-WAN devices in the overlay network,
- comprises routing policy for control plane traffic flow and data policy for data plane traffic flow, and
- implements enterprise-specific traffic control requirements through basic policies and advanced features.

Table 1: Feature history

Feature	Release Information	Description
Policy enforcement failure	Cisco IOS XE Catalyst SD-WAN Release 17.18.1a Cisco Catalyst SD-WAN Manager Release 20.18.1	This feature aims to detect when a policy download fails and raises an alarm (policy-enforcement-status). Additionally, this feature introduces new **service-path** show commands.
Policy validation in Cisco SD-WAN	Cisco IOS XE Catalyst SD-WAN Release 26.1.1 Cisco Catalyst SD-WAN Manager Release 26.1.1.1	This feature ensures network reliability and operational efficiency by automatically validating Cisco Catalyst SD-WAN policies for accuracy, platform compliance, and alignment with network requirements before deployment.

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 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.

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 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 in the network.

Centralized data policy applies to data traffic by directing traffic flows to take specific paths through the network, or by adding Quality of Service (QoS) on packets.

Localized Policy

Localized policy refers to policy provisioned locally on Cisco IOS XE Catalyst SD-WAN devices. Localized policy also comprises two components:

- Control policy, which affects the routing information that Cisco IOS XE Catalyst SD-WAN devices advertise to their local attached networks
- Data policy, which applies access control lists (ACLs) and QoS to packets being transmitted from and received on Cisco IOS XE Catalyst SD-WAN device router interfaces

Localized control policy affects the routing information that Cisco IOS XE Catalyst SD-WAN devices propagate within their local site network.

When data plane traffic is unpolicied, 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

Describes the structural division between control and data planes, as well as centralized versus localized functions.

Policy architecture is a framework that

- implements overlay network-wide policies called Cisco SD-WAN Validator policy or centralized policy
- affects the flow of both control plane traffic and data plane traffic across the overlay network, and
- supports traditional routing policies on Cisco IOS XE Catalyst SD-WAN devices for local routing protocol control.

Policy implementation details

The Cisco Catalyst SD-WAN policy architecture provides centralized policy configuration on a Cisco SD-WAN Controller. Cisco SD-WAN Controller policy affects control plane traffic through 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. It also affects data plane traffic that travels between the Cisco IOS XE Catalyst SD-WAN devices across the overlay network.

You can also create routing policies on the Cisco IOS XE Catalyst SD-WAN devices. These policies are 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

Explains centralized control policy architecture that distinguishes between routing protocol policies (Control) and access control lists or firewall filters (Data).

A centralized control policy architecture is a network management framework that

- enables centralized orchestration of network-wide routing decisions by a central authority instead of hop-by-hop implementation
- allows influence over network routes advertised by Cisco SD-WAN Controllers through policies provisioned centrally
- ensures routing policy results are distributed to devices using the Overlay Management Protocol (OMP) while maintaining centralized policy configuration.

Architecture components and implementation

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.

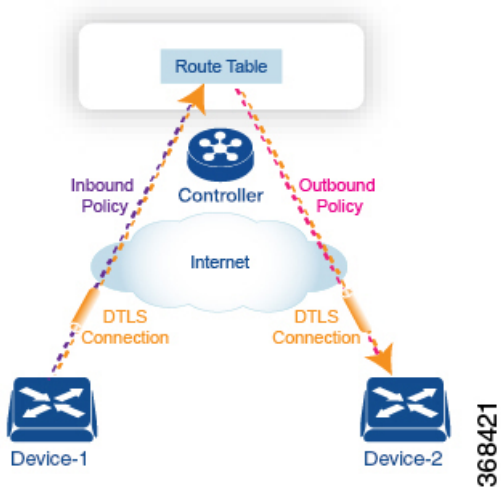
Centralized control policy allows you to influence the network routes advertised by the Cisco SD-WAN Controllers. This type of policy, which is provisioned centrally on the Cisco SD-WAN Controller, affects both the route information that the Cisco SD-WAN Controller stores in its primary route table and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

Describes the categorization of OMP routes into three types: OMP routes (prefixes), TLOC routes (transport locations), and Service routes (network services).

Route types are OMP route categorizations that

- carry prefix information that the devices learn from routing protocols
- identify transport locations and their properties, and
- identify available network services on local-site networks.

Route type details

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

Describes the standard state where all route information is shared unmodified across all controllers and devices.

Default behavior without centralized control policy is a network state that

- operates when no centralized control policy is provisioned on the Cisco SD-WAN Controller
- enables unrestricted sharing of all route information across controllers and devices within a domain, and
- results in automatic redistribution of OMP, TLOC, and service routes throughout the network.

Route advertisement and redistribution behavior

The default behavior produces the following route advertisement and redistribution patterns 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

Explains how provisioning a policy allows for filtering, modifying, or rejecting routes before they are distributed.

A centralized control policy is a network management mechanism that

- filters, modifies, or rejects route information before distribution to Cisco IOS XE Catalyst SD-WAN devices in a domain
- controls route information stored in the Cisco Catalyst SD-WAN Controller's route table or advertised by the Cisco Catalyst SD-WAN Controller, and
- applies to specific sites in the overlay network in either inbound or outbound direction with respect to the Cisco Catalyst SD-WAN Controller.

Control policy direction behaviors

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

Control policy behavior depends on the direction of application:

- **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. Routes are either accepted or rejected as the first step. 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.
- **Outbound direction**: Filters or modifies the routes that the Cisco Catalyst SD-WAN Controller redistributes to the Cisco IOS XE Catalyst SD-WAN devices. Routes are either accepted or rejected as the first step. 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 is a second type of centralized data policy that 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. 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

Describes practical scenarios for using policies to customize how data moves through the overlay.

Examples of modifying traffic flow with centralized control policy are practical scenarios that demonstrate how you can use centralized control policies to modify the flow of data traffic through the overlay network.

Basic examples of traffic flow modification

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.

Arbitrary topologies

Describes how to use control policies to build specific network structures, such as hub-and-spoke models, to save resources.

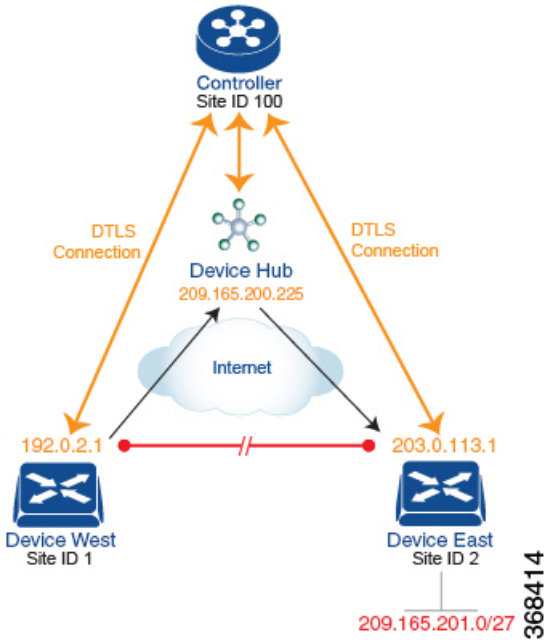
An arbitrary topology is a network design approach that

- uses control policies to create specific network structures like hub-and-spoke models,
- minimizes CPU resource overhead by reducing the number of direct IPsec tunnel connections, and
- enables scalable network architectures for large deployments with hundreds or thousands of branches.

Hub-and-spoke topology implementation

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.

West-East branch hub-and-spoke configuration

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.

Traffic engineering

Explains using TLOC preferences and path tracking to favor specific hub devices and monitor end-to-end path availability.

Traffic engineering is a control policy capability that

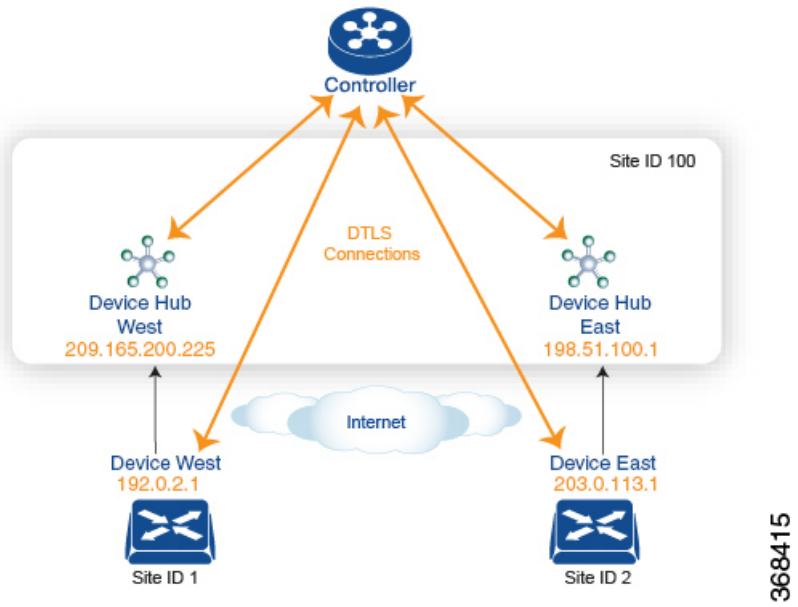
- allows you to design and provision traffic flow through the network using TLOC preferences

- enables path tracking to monitor end-to-end path availability between source and destination devices, and
- provides service chaining functionality by redirecting packets through network services such as firewalls.

Traffic engineering implementation methods

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

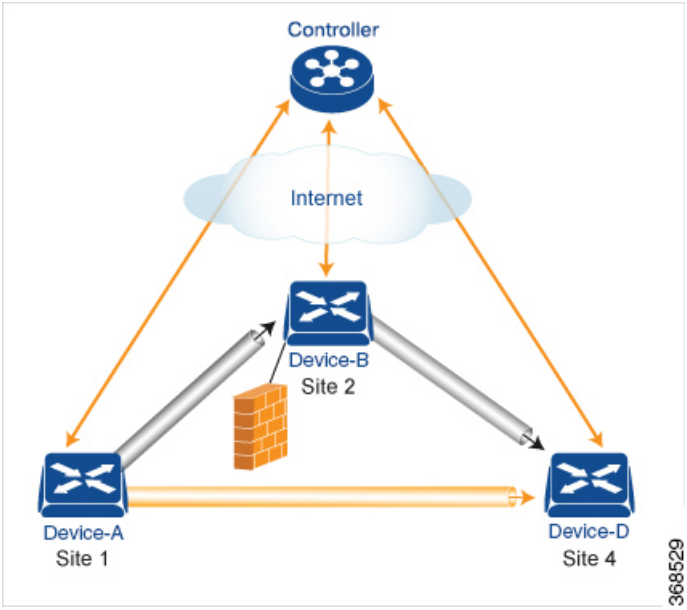


The figure shows that Site ID 100 has two hub devices, one that serves the West side of the network and a second that serves the East side. Data traffic from the Device West must be handled by the Device West hub, and similarly, data traffic from the Device East branch must go through the Device East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



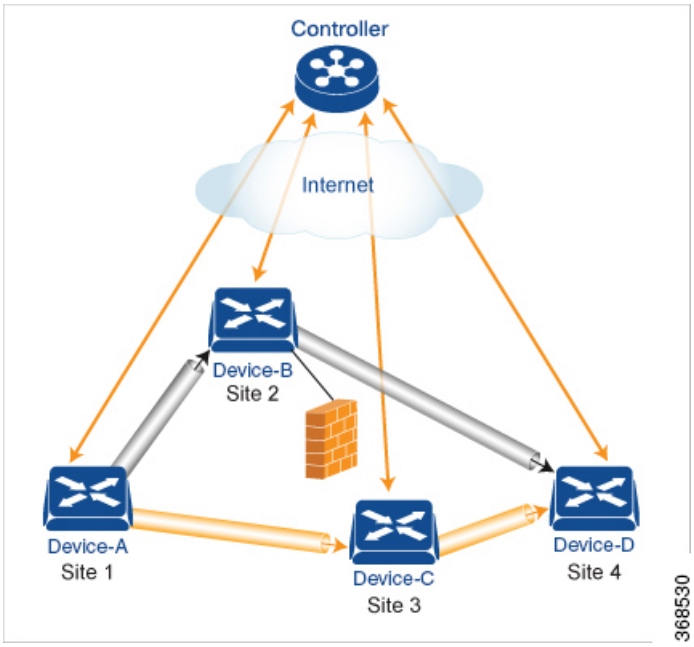
Note

Ensure that the Traffic Engineering (TE) service is enabled on the intermediate hub to utilize the traffic engineering functionalities.

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

Figure 6: Traffic 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.

Centralized policy-based configuration on prefixes and IP headers

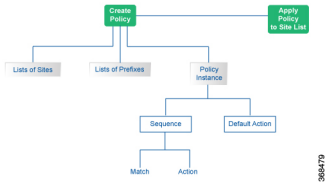
Describes how data policies use ordered match-action sequences to evaluate and handle IP packets based on headers.

A centralized data policy based on source and destination prefixes and on headers in IP packets is a series of numbered sequences that

- consists of ordered match-action pairs that are evaluated from lowest sequence number to highest sequence number
- takes the associated action when a packet matches one of the match conditions and stops policy evaluation on that packet, and
- drops and discards packets by default if no parameters in any of the sequences match.

Configuration components

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



Policy validation for Cisco Catalyst SD-WAN

Explains the policy validation for Cisco Catalyst SD-WAN devices.

Policy validation in Cisco Catalyst SD-WAN is a process that

- automatically checks if your policies are accurate,

- ensures that policies comply with platform capabilities, and
- confirms that policies align with your network requirements.

Application list: An application list is a collection of applications grouped together for policy matching, allowing you to apply policies to multiple applications at once.

Application family: An application family refers to a category of related applications such as Cisco Webex, Microsoft Teams, and Zoom. Application categories simplify policy management and enforcement.

Key features of policy validation

This process helps verify that all configurations operate within supported limits before deployment.

- Centralized policy checks: From Cisco IOS XE Catalyst SD-WAN Release 26.1.1 policy validation is centrally managed, enabling quicker error detection and ensuring configurations remain within supported limits.
- Enhanced device alerts: Devices send detailed alerts to Cisco SD-WAN Manager for proactive monitoring.
- Filters in a sequence: Each policy sequence supports up to 64 filters, giving you flexibility to define granular traffic matching criteria.
- Entries in a list: You can create and modify lists with a combined total of up to 8192 entries including existing and new entries to ensure scalability for complex network requirements.

Cisco Catalyst SD-WAN Controller Policy components

Describes the three essential building blocks of policies: lists, policy definitions, and policy applications.

Cisco Catalyst SD-WAN Controller policy components are a configuration framework that

- consists of three essential building blocks: lists, policy definitions, and policy applications
- enables centralized management and enforcement of overlay network-wide policies on Cisco Catalyst SD-WAN Control Components, and
- ensures consistency in policy enforcement across the overlay network through centralized configuration.

Policy component architecture

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 2: The three building blocks of Cisco SD-WAN Controller policies

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

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 site1  --------→ Apply the defined policy towards the sites in
site-list
  control-policy prefer_local out
  !
policy
 lists
 site-list site1
  site-id 100
 tloc-list prefer_site1 ---→ Define the lists required for apply-policy and
for use within the policy
  tloc 192.0.2.1 color mols encap ipsec preference 400
 control-policy prefer_local
  sequence 10
   match route
    site-list sitele ------->Lists previously defined used within policy
  !
   action accept
    set
     tloc-list prefer_site
    !
   !
  !
```

TLOC attributes used in policies

Lists attributes like IP address, color, and encapsulation that can be matched or modified within TLOC routes.

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.

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

Lists attributes like origin, preference, and service used to identify and modify OMP routes.

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 3:

Route Attribute	Function	Application Point Set By	Application Point Modify By
Origin	Source of the route, either BGP, OSPF, connected, static. Match origin command is available only for Cisco vEdge devices but not on Cisco IOS XE Catalyst SD-WAN devices.	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

Best practice for Cisco Catalyst SD-WAN Controller policy processing and application

Outlines how policies are sequentially processed and the importance of consistent configuration across controllers.

Policy processing and application guidelines

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.

- 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.

Site-list policy application requirements

Cisco SD-WAN Controller policy is applied on a per-site-list basis.

- 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.

Policy distribution and routing considerations

Policy distribution and routing decisions require understanding of how information flows through the overlay network.

- 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.

Requirement: Consistent configuration across controllers

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.

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

Explains at a high level how control, data, and VPN membership policies function to manage the network.

Cisco Cisco Catalyst SD-WAN Controller policy operation is a network management mechanism that

- operates on routing information carried in OMP updates through control policy
- affects data traffic through data policy, and
- controls the distribution of VPN routing tables through VPN membership.

Basic policy types

The basic Cisco SD-WAN Controller policies are:

- Control Policy

- Data Policy
- VPN Membership

Control policy

Explains how control policy customizes routing decisions in the control plane to influence paths through the overlay or local branch networks.

A control policy is a routing policy mechanism that

- operates on routes and routing information in the control plane of the overlay network
- customizes network-wide routing decisions that determine or influence routing paths through the overlay network, and
- allows modification of OMP route attributes before they are placed in the route table or advertised to devices.

Control policy types and operation

Two types of control policy provide routing customization:

- Centralized control policy: Provisioned on the Cisco SD-WAN Controller and customizes network-wide routing decisions through the overlay network
- Local control policy: Provisioned on a Cisco IOS XE Catalyst SD-WAN device and 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: 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.
- Service routes: Advertise the network services, such as firewalls, available to VPN members at the local site.

By default, no centralized control policy is provisioned. In this bare, unpolicied 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, edge devices do not perform any policy actions on routes learned from the Cisco SD-WAN Controllers.

The figure below shows the OMP route flow in the overlay network and illustrates the effects of centralized control policy. Centralized control policy affects the exchange of routing information between the Cisco IOS XE Catalyst SD-WAN devices and the Cisco SD-WAN Controllers. As shown in the figure, you can apply a policy in two directions:

- Inbound–A policy applied to routes being received by the Cisco SD-WAN Controller from the Cisco IOS XE Catalyst SD-WAN devices in the network
- Outbound–A policy applied to routes being sent by the Cisco SD-WAN Controller to the Cisco IOS XE Catalyst SD-WAN devices in the network

For inbound policy, the routes received from the Cisco IOS XE Catalyst SD-WAN devices are modified before they are placed in the Cisco SD-WAN Controller's route table, and the modified routes are also sent to all other Cisco IOS XE Catalyst SD-WAN devices. Inbound policies affect all the devices in the network domain because the Cisco SD-WAN Controller advertises the modified routes to all the Cisco IOS XE Catalyst SD-WAN devices in the VPN.

For an outbound policy, the routes in the Cisco SD-WAN Controller's route table are not modified. Instead, the Cisco SD-WAN Controller reads a route from its route table, creates a copy of the route, modifies the copy as directed by the policy, and sends this modified copy to the destination devices. The original route in the Cisco SD-WAN Controller's route table remains unchanged after the update. Again, note that the direction is outbound from the perpspective 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 policies

Describes policies that influence data traffic flow based on IP header fields or specific router interfaces using centralized or localized rules.

A data policy is a network traffic control mechanism that

- 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
- controls data traffic traveling over IPsec connections between Cisco IOS XE Catalyst SD-WAN devices, and
- examines fields in data packet headers, looking at source and destination addresses and ports, protocol and DSCP values.

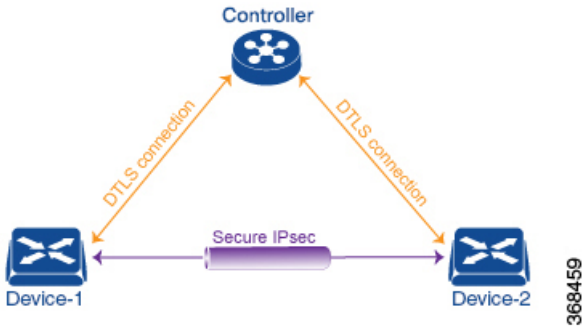
Data policy types

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

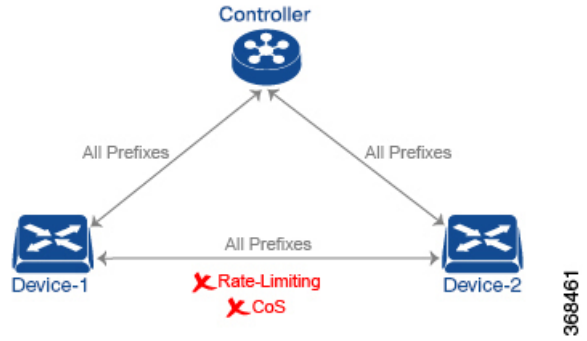
- **Centralized data policy**: Controls the flow of data traffic based on the source and destination addresses and ports and DSCP fields in the packet's IP header (referred to as a 5-tuple), and based on network segmentation and VPN membership. These types of data policy are provisioned centrally, on 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 polices 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.



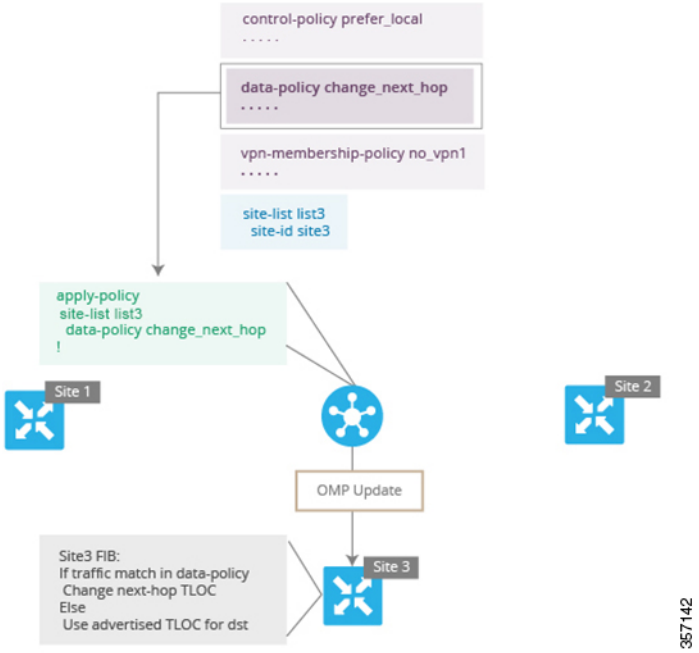
With no access lists provisioned on a Cisco IOS XE Catalyst SD-WAN device, all data traffic is transmitted at line rate and with equal importance, using one of the interface's queues. Using access lists, you can provision class of service, which allows you to classify data traffic by importance, spread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

Data policy 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.

Data policy topology example

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 7: 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.

Restrictions for data policy

Outlines the limitations and supported configurations for applying data policies per site on Cisco SD-WAN Controller.

You can apply only one data policy per site in each direction on Cisco SD-WAN Controller regardless of the number of VPNs configured.

- **Supported directions:**
  - From-service
  - From-tunnel
  - All- this direction represents both from-service and from-tunnel traffic.
- **Supported configurations:** You can use any one of the following:
  - One data policy for from-service direction and One data policy for from-tunnel direction, or
  - One data policy for all direction.
- **Unsupported configurations:**
  - One data policy for from-service and another for all at the same site.
  - One data policy for from-tunnel and another for all at the same site.

  Applying multiple data policies in the same direction at a site is not supported and can result in unintended behavior.

VPN membership policy operation

Describes how VPN membership policy restricts the distribution of VPN routing tables to specific devices to enforce business-specific participation limits.

VPN membership policy operation is a network routing mechanism that

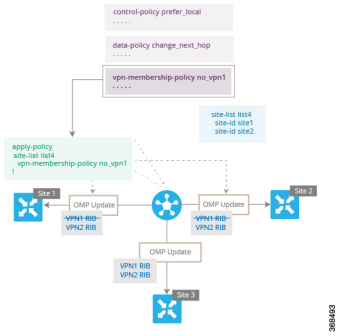
- affects the VPN route tables that are distributed to particular Cisco IOS XE Catalyst SD-WAN devices
- filters route updates sent to devices, removing routes for VPNs they do not service, and
- enforces business usage model restrictions on device participation in particular VPNs.

VPN membership policy behavior

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 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. Direction is not set when applying VPN membership policy.

Figure 8: VPN membership topology



VPN membership policy filtering scenario

The 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.

Configure and execute Cisco SD-WAN Controller policies

Describes where policies are configured (centralized) and whether they execute on the controller or the local device.

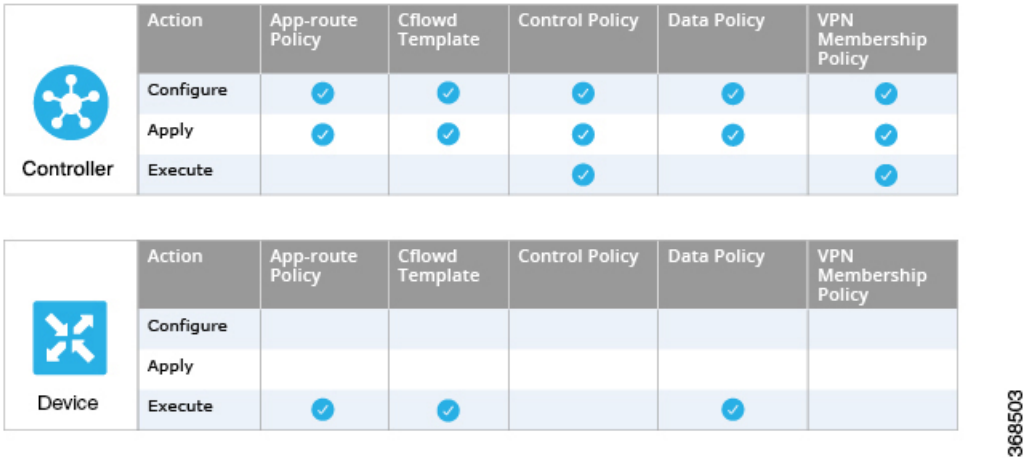
Cisco SD-WAN Controller policy configuration and execution is a centralized policy management system that

- configures all policies centrally on the Cisco IOS XE Catalyst SD-WAN devices using policy definitions and lists
- applies policies on the Cisco IOS XE Catalyst SD-WAN devices with apply-policy and lists, and
- executes policies on different locations depending on the policy type.

Policy execution locations

The actual Cisco SD-WAN Controller policy execution location depends on the type of policy, as shown in this figure:

Figure 9: Cisco SD-WAN Controller Policy



Policy execution locations are categorized as follows:

- 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.
- 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.

Policy status and health monitoring

Lists tools and commands, such as service-path show commands, used to verify policy enforcement and pathing.

Tools and commands are available to verify policy enforcement and pathing in SD-WAN environments.

Service-path show commands

```
vm5#show sdwan policy service-path vpn 101 interface GigabitEthernet5.101
source-ip ... dest-ip ... protocol 1

Next Hop: Blackhole
AAR Policy:
  Policy name: AAR1
  VPN list name: VPN101
  Sequence number: 1
Data Policy:
  Policy name: DP1
  VPN list name: VPN101
  Sequence number: 1
Drop reason: Tunnels matching local color preference not found, dropping due
to strict criteria


vm5#show sdwan policy tunnel-path vpn 101 interface GigabitEthernet5.101
source-ip ... dest-ip ... protocol 1

Next Hop: IPsec
  Source: ... 12346 Destination: ... 12366 Local Color: lte Remote Color: lte
```

```
 Remote System IP: ...
AAR Policy:
  Policy name: AAR1
  VPN list name: VPN101
  Sequence number: 1
Data Policy:
  Policy name: DP1
  VPN list name: VPN101
  Sequence number: 1
Exceptions:
  Tunnels matching SLA not found, SLA criteria ignored
```

3 Centralized Policy

Topics:

- Feature history for centralized policy
- Overview of centralized policies
- Configure centralized policies
- Centralized policies configuration examples

Introduces centralized policy management capabilities and guides users through configuration, verification, and implementation examples for network-wide policy control.

Feature history for centralized policy

Lists the features and enhancements that have been introduced for centralized policy functionality across different software releases.

This reference provides information about the features and enhancements that have been introduced for centralized policy functionality across different software releases.

Table 4: 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.
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.
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).
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.
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.
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
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.
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**.
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 toCisco SD-WAN Manager AAR policy and view the applied recommendations on Cisco SD-WAN Manager.
Remote Preferred Color in Data Policy	Cisco IOS XE Catalyst SD-WAN Release 17.15.1aCisco 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 using classic policies* section for more information.

Overview of centralized policies

Provides an overview of centralized policies in the network overlay.

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

Describes the different types of centralized policies used in SD-WAN overlay networks to control traffic routing and data flow.

Centralized policies are network management mechanisms that

- control traffic routing and data flow throughout the overlay network from a central location
- operate on different planes to affect either routing decisions or data traffic flow, and
- are provisioned on the Cisco Catalyst SD-WAN Controller controller and applied network-wide.

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 vEdge devices. The effects of centralized control policy are seen in how Cisco vEdge devices direct the overlay network's data traffic to its destination.

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 vEdge 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.

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. These 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

Configure centralized policies

Use one of these methods to configure centralized policies:

- Configure centralized policies using the CLI
- Configure centralized policies using classic policies

Configure centralized policies using the CLI

This section helps you configure centralized policies using the command line interface.

To configure a centralized control policy using the CLI:

Procedure

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

   Example

   ```
   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:

   Example

   ```
   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:

   Example

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

4. Create a series of match-action pair sequences:

   Example

```
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:

   Example

```
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:

   Example

```
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
```

> **Note**
>
> Match origin command is available only for Cisco vEdge devices but not on Cisco IOS XE Catalyst SD-WAN devices.

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:

   Example

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

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

   Example

   ```
   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 device so that the Cisco Catalyst SD-WAN Controller knows how to reach the services:

    Example

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

    Specify the VPN is which the service is located and one to four IP addresses to reach the service device or devices. If multiple devices provide the same service, the device load-balances the traffic among them. Note that the Cisco IOS XE 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.

Configure a VPN membership data policy

The high-level steps for configuring a VPN membership data policy are shown below:

Procedure

---

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

Example

```
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:

Example

```
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.

Example

```
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}
```

4. Define policing parameters, as needed:

   Example

   ```
   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:

   Example

   ```
   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:

   Example

   ```
   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:

   Example

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

8. Define actions to take when a match occurs:

   Example

   ```
   vSmart(config-sequence-number)# action (accept | drop) [count counter-name]
   [log] [tcp-optimization]
   vSmart(config-sequence-number)# action acccept 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:

    Example

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

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

Example

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

Centralized policies using classic policies

Explains how to configure centralized policies using the Cisco SD-WAN Manager policy configuration wizard.

Centralized policies using classic policies are network management configurations that

- use the Cisco SD-WAN Manager policy configuration wizard to guide policy creation and editing
- consist of multiple operations including creating groups, configuring topology, and defining traffic rules, and
- require activation for the centralized policy to take effect in the overlay network.

Configuration operations

The wizard consists of these 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

Configure the policy configuration wizard to begin creating centralized policies.

Start the policy configuration wizard to create and configure centralized policies for your SD-WAN deployment.

The policy configuration wizard provides a guided interface for creating centralized policies. Use this procedure when you need to create new policies for traffic management, security, or service chaining.

Procedure

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.

The policy configuration wizard opens with the **Create Groups of Interest** window displayed, ready for policy configuration.

Configure groups of interest for centralized policy

This topic helps you configure groups of interest for a 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:

Procedure

1. Configure application:
   a) In the groups of interest list, click **Application** list type.
   b) Click **New Application List**.
   c) Enter a name for the list.
   d) Choose either **Application** or **Application Family**.

      Application can be the names of one or more applications, such as **Third Party Control**, **ABC News**, **Mircosoft 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**.

      > **Note**
      >
      > A single app-list can hold either app names or application families. It cannot hold more than 10 application families. When you try to put an application in atmost one app-list. When you have to reuse the application, keep the number of shared applications less than or equal to 10. Keep the number of lists in an applictaion less than or equal to five. Do not keep more than one shared app in the same list.

   e) In the **Select** drop-down, in the 'Search' filter, select the required applications or application families.

      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.

2. Configure color:
   a) In the groups of interest list, click **Color**.
   b) Click **New Color List**.
   c) Enter a name for the list.
   d) 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.

   e) Click **Add**.

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

3. Configure community:
   a) In the group of interest list, click **Community**.
   b) Click **New Community List**.
   c) Enter a name for the community list.
   d) 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.

   e) 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.

   f) Click **Add**.
4. Configure data prefix:
   a) In the **Groups of Interest** list, click **Data Prefix**.
   b) Click **New Data Prefix List**.
   c) Enter a name for the list.
   d) Choose either **IPv4** or **IPv6**.
   e) In the **Add Data Prefix** field, enter one or more data prefixes separated by commas.
   f) Click **Add**.
5. Configure policer:
   a) In the groups of interest list, click **Policer**.
   b) Click **New Policer List**.
   c) Enter a name for the list.
   d) Define the policing parameters:
      1. In the **Burst** field, enter the maximum traffic burst size, a value from 15,000 to 10,000,000 bytes.
      2. In the **Exceed** field, select the action to take when the burst size or traffic rate is exceeded. It can be **drop**, which sets the packet loss priority (PLP) to **low**.

         You can use the **remark** action to set the packet loss priority (PLP) to **high**.
      3. In the **Rate** field, enter the maximum traffic rate, a value from 0 through 2^{64} – 1 bits per second (bps).

   e) Click **Add**.

> **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.

6. Configure prefix:
   a) In the groups of interest list, click **Prefix**.
   b) Click **New Prefix List**.
   c) Enter a name for the list.
   d) In the **Add Prefix** field, enter one or more data prefixes separated by commas.
   e) Click **Add**.
7. Configure site:
   a) In the groups of interest list, click **Site**.
   b) Click **New Site List**.
   c) Enter a name for the list.
   d) In the **Add Site** field, enter one or more site IDs separated by commas.

      For example, 100 or 200 separated by commas or in the range, 0 - 4294967295.

   e) Click **Add**.
8. Configure app probe class:
   a) In the groups of interest list, click **App Probe Class**.
   b) Click **New App Probe Class**.
   c) Enter the probe class name in the **Probe Class Name** field.
   d) Select the required forwarding class from the **Forwarding Class** drop-down list.
   e) 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.

   f) Click **Save**.
9. Configure SLA class:
   a) In the groups of interest list, click **SLA Class**.
   b) Click **New SLA Class List**.
   c) Enter a name for the list.
   d) Define the SLA class parameters:
      1. In the **Loss** field, enter the maximum packet loss on the connection, a value from 0 through 100 percent.
      2. In the **Latency** field, enter the maximum packet latency on the connection, a value from 0 through 1,000 milliseconds.
      3. In the **Jitter** field, enter the maximum jitter on the connection, a value from 1 through 1,000 milliseconds.
      4. Select the required app probe class from the **App Probe Class** drop-down list.

   e) (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.

f) 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

g) Enter the **Loss Variance (%)**, **Latency Variance (ms)**, and the **Jitter Variance (ms)** for the selected criteria.

h) Click **Add**.

10. Configure TLOC:

a) In the groups of interest list, click **TLOC**.

b) Click **New TLOC List**. The **TLOC List** popup displays.

c) Enter a name for the list.

d) In the **TLOC IP** field, enter the system IP address for the TLOC.

e) In the **Color** field, select the TLOC's color.

f) In the **Encap** field, select the encapsulation type.

g) In the **Preference** field, optionally select a preference to associate with the TLOC.

The range is 0 to 4294967295.

h) Click **Add TLOC** to add another TLOC to the list.

i) 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.

11. Configure VPN:

a) In the groups of interest list, click **VPN**.
b) Click **New VPN List**.
c) Enter a name for the list.
d) In the **Add VPN** field, enter one or more VPN IDs separated by commas.

For example, 100 and 200 separated by commas. Range: 1 to 65525, excluding 512. For details see the VRF range behavior change described here.

Click **Add**.

12. Configure region:

Minimum release: Cisco SD-WAN 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**.

a) In the groups of interest list, click **Region**.
b) Click **New Region List**.
c) In the **Region List Name** field, enter a name for the region list.
d) 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.

e) Click **Add**.

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

13. Configure preferred color group:

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.

a) In the groups of interest list, click **Preferred Color Group**.
b) Click **New Preferred Color Group**.
c) In the **Preferred Color Group Name** field, enter a name for the preferred color group.
d) In the **Primary Colors** pane, do the following:

1. Choose the color preference from the **Color Preference** drop-down list.
2. 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	Choose the color preference from the drop-down list. The options are: • default • 3g • biz-internet • blue • bronze • custom1 • custom2, and so on You can select multiple colors.
Path Preference	Choose the path preference from the drop-down list. The options are: • 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. **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.

e) In the **Secondary Colors** pane, do the following:

1. Choose the color preference in the **Color Preference** drop-down list.
2. Choose the path preference from the **Path Preference** drop-down list.

f) In the **Tertiary Colors** pane, do the following:

1. Choose the color preference from the **Color Preference** drop-down list.
2. Choose the path preference from the **Path Preference** drop-down list.

g) 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

Describes the integration of WAN insight (WANI) into Cisco SD-WAN Manager for predictive path recommendations.

Integrating WAN insight (WANI) into Cisco SD-WAN Manager is a cloud-based analytics service integration that

- provides comprehensive insights into application and network performance for Cisco Catalyst SD-WAN
- enables Predictive Path Recommendations to be applied to actionable centralized AAR policy to influence forwarding decisions, and
- applies recommendations as TLOC preferences in AAR policies pushed to Cisco SD-WAN Controller.

Implementation details

Cisco SD-WAN Analytics collects and stores metadata about traffic flows in its cloud storage and produces analytics based on this collected data. The analytics service is available with Cisco DNA Advantage and Cisco DNA Premier software subscriptions. 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.

For more information about using Predictive Path Recommendations, see Predictive Path Recommendations.

To apply predictive path recommendations when there are recommendations in Cisco SD-WAN Analytics:

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 when integrating WAN insight:

- 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

Provides predictive network optimization capabilities that use statistical models to examine historical data and generate path recommendations for optimal application traffic routing.

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

Configure network topology and VPN membership using hub-and-spoke, mesh, or custom control policies to define how sites communicate within the SD-WAN overlay network.

Configure network topology to define the communication patterns between sites in your SD-WAN deployment and establish VPN membership for routing policies.

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

You can configure three types of topologies: Hub-and-Spoke for centralized communication, Mesh for full connectivity between sites, and Custom Control for advanced route and TLOC control policies.

Procedure

1. Configure a hub-and-spoke topology.

   Hub-and-Spoke

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

   b. Enter a name for the hub-and-spoke policy.

   c. Enter a description for the policy.

   d. In the **VPN List** field, select the VPN list for the policy.

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

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

g. In the right pane, add hub sites to the network topology:

   1. Click **Add Hub Sites**.
   2. In the **Site List** field, select a site list for the policy component.
   3. Click **Add**.
   4. Repeat these steps to add more hub sites to the policy component.

h. In the right pane, add spoke sites to the network topology:

   1. Click **Add Spoke Sites**.
   2. In the **Site List Field**, select a site list for the policy component.
   3. Click **Add**.
   4. Repeat these steps to add more spoke sites to the policy component.

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

j. Click **Save Hub-and-Spoke Policy**.

2. Configure a mesh topology.

   Mesh

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

   b. Enter a name for the mesh region policy component.

   c. Enter a description for the mesh region policy component.

   d. In the **VPN List** field, select the VPN list for the policy.

   e. Click **New Mesh Region**.

   f. In the **Mesh Region Name** field, enter a name for the individual mesh region.

   g. In the **Site List** field, select one or more sites to include in the mesh region.

   h. Click **Add**.

   i. Repeat these steps to add more mesh regions to the policy.

   j. Click **Save Mesh Topology**.

3. Configure a custom control policy for route control.

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

   a. In the **Add Topology** drop-down, select **Custom Control (Route & TLOC)**.

   b. Enter a name for the control policy.

   c. Enter a description for the policy.

   d. In the left pane, click **Sequence Type**. The **Add Custom Control Policy** popup displays.

   e. Select **Route**. A policy component containing the text string **Route** is added in the left pane.

   f. Double-click the **Route** text string, and enter a name for the policy component.

   g. In the right pane, click **Sequence Rule**. The **Match/Actions** box opens, and **Match** is selected by default.

h. 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.

i. 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.

j. Click **Save Match and Actions**.

k. Click **Sequence Rule** to configure more sequence rules, as desired. Drag and drop to re-order them.

l. Click **Sequence Type** to configure more sequences, as desired. Drag and drop to re-order them.

m. Click **Save Control Policy**.

4. Configure a custom control policy for TLOC control.

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

   a. In the **Add Topology** drop-down, select **Custom Control (Route & TLOC)**.

   b. Enter a name for the control policy.

   c. Enter a description for the policy.

   d. In the left pane, click **Sequence Type**. The **Add Custom Control Policy** popup displays.

   e. Select **TLOC**. A policy component containing the text string **TLOC** is added in the left pane.

   f. Double-click the **TLOC** text string, and enter a name for the policy component.

   g. In the right pane, click **Sequence Rule**. The **Match/Actions** box opens, and **Match** is selected by default.

   h. 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.

   i. 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.

   j. Click **Save Match and Actions**.

   k. Click **Sequence Rule** to configure more sequence rules, as desired. Drag and drop to re-order them.

   l. Click **Sequence Type** to configure more sequences, as desired. Drag and drop to re-order them.

   m. Click **Save Control Policy**.

The topology and VPN membership configuration is saved. 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.

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

> **Note**
>
> The Cisco Catalyst SD-WAN policy supports maximum up to 1024 sequences, including default sequence.

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

Configure the system to import a previously created topology configuration.

Import an existing topology to reuse previously configured network structures and policies.

Use this procedure when you need to incorporate a topology that was previously created and configured in your environment.

Before you begin

Follow these steps to import existing topology:

Procedure

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

The existing topology is imported and you can proceed to configure traffic rules in the wizard.

Create a VPN membership policy

Configure a VPN membership policy to define network topology relationships between sites and VPNs.

Create a VPN membership policy to establish which sites can access specific VPNs in your network topology.

VPN membership policies define the relationship between site lists and VPN lists in your network configuration. Only one VPN membership policy can be added at a time, so all required site lists and VPN lists must be included in a single policy.

Procedure

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.

The VPN membership policy is created and you can proceed to configure traffic rules for your network topology.

Configure traffic rules

This topic helps you configure traffic rules for a centralized data policy.

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 SD-WAN Manager Application Intelligence Engine (SAIE) flow, see the SD-WAN Application Intelligence Engine Flow section in the *Cisco Catalyst SD-WAN Policies Configuration Guide, Cisco IOS XE Catalyst SD-WAN Release 26.x*.

Before you begin

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

Procedure

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.

> **Note**
>
> Adding multiple application matches, subnet matches, or combining several match conditions in a single sequence may increase the total filter count beyond the allowed maximum.
>
> If this total filter count is reached, an error message `Filter count exceeded MAX (64)` is displayed when applying a data policy configuration.

Match Condition	Procedure
None (match all packets)	Do not specify any match conditions.
Applications /Application Family List	**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: 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**. This match condition is available for IPv6 traffic from Cisco IOS XE Catalyst SD-WAN Release 17.9.1a and Cisco SD-WAN Release 20.9.1.

Match Condition	Procedure
Destination Data Prefix	**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.
Destination Port	**a.** In the **Match** conditions, click **Destination Port**. **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 [-]).
DNS Application List	Add an application list to enable split DNS. **a.** In the **Match** conditions, click **DNS Application List**. **b.** In the drop-down, select the application family. This match condition is available for IPv6 traffic from Cisco IOS XE Catalyst SD-WAN Release 17.9.1a and Cisco SD-WAN Release 20.9.1.
DNS	Add an application list to process split DNS. **a.** In the **Match** conditions, click **DNS**. **b.** In the drop-down, select **Request** to process DNS requests for the DNS applications, and select **Response** to process DNS responses for the applications.
DSCP	**a.** In the **Match** conditions, click **DSCP**. **b.** In the **DSCP** field, type the DSCP value, a number from 0 through 63.
Packet Length	**a.** In the **Match** conditions, click **Packet Length**. **b.** In the **Packet Length** field, type the length, a value from 0 through 65535.
PLP	**a.** In the **Match** conditions, click **PLP** to set the **Packet Loss Priority**. **b.** In the **PLP** drop-down, select **Low** or **High**. To set the PLP to **High**, apply a policer that includes the **exceed remark** option.
Protocol	**a.** In the **Match** conditions, click **Protocol**. **b.** In the **Protocol** field, type the Internet Protocol number, a number from 0 through 255.

Match Condition	Procedure
ICMP Message	To match ICMP messages, in the Protocol field, set the Internet Protocol Number to 1, or 58, or both. **Note** This field is available from Cisco IOS XE Catalyst SD-WAN Release 17.4.1a, Cisco SD-WAN Release 20.4.1.
Source Data Prefix	a. In the **Match** conditions, click **Source Data Prefix**. b. To match a list of source prefixes, select the list from the drop-down. c. To match an individual source prefix, enter the prefix in the **Source** field.
Source Port	a. In the **Match** conditions, click **Source Port**. b. In the **Source** field, enter the port number. Specify a single port number, a list of port numbers (with numbers separated by a space), or a range of port numbers (with the two numbers separated with a hyphen [-]).
TCP	a. In the **Match** conditions, click **TCP**. b. In the **TCP** field, **syn** is the only option available.

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.	**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.	**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.	**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.

Action Condition	Description	Procedure
Log	Minimum release: Cisco IOS XE Catalyst SD-WAN Release 17.11.1a and Cisco Catalyst SD-WAN Control Components Release 20.11.1 Click **Log** to enable logging. 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.	**a.** In the **Action** conditions, click **Log** to enable logging.
Policer	Apply a policer to matching data packets.	**a.** In the **Match** conditions, click **Policer**. **b.** In the **Policer** drop-down field, select the name of a policer.
Loss Correction	Apply loss correction to matching data packets. 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. FEC is supported only for IPSEC tunnels, it is not supported for GRE tunnels. • **FEC Adaptive** – Corresponding packets are subjected to FEC only if the tunnels that they go through have been deemed unreliable based on measured loss. If you choose **FEC Adaptive**, an additional field, **Loss Threshold**, displays that allows you to specify the packet loss threshold for automatically enabling FEC. Adaptive FEC starts to work at 2% packet loss; this value is configurable. You can specify a loss threshold of 1 to 5%. The default packet loss threshold is 2%. • **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.	**a.** In the **Match** conditions, click **Loss Correction**. **b.** In the **Loss Correction** field, select **FEC Adaptive**, **FEC Always**, or **Packet Duplication**.

Click **Save Match and Actions**.

16. Create additional sequence rules as desired. Drag and drop to re-arrange them.
17. Click **Save Data Policy**.

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

Match parameters - control policy

Lists the available match parameters for OMP AND TLOC routes in control policies.

For OMP AND TLOC routes, you can match attributes to control routing behavior AND traffic flow in your SD-WAN network.

Table 5: Control policy match parameters

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.
Community List	List of one OR more BGP communities. In the **Community List** field, you can specify: · *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	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.
Criteria OR	Compares each regex string in the community list against the community string of the route. The OR condition is applicable across multiple community lists AND is valid for all devices. Starting from Cisco IOS XE Catalyst SD-WAN Release 17.5.1a, the Community **Types** AND **Criteria** fields are available.

Match Condition	Description
OMP Tag	Tag value associated with the route OR prefix in the routing database on the device. The range is 0 through 4294967295.
Origin	Protocol from which the route was learned.
Originator	IP address from which the route was learned.
Path Type	In a Hierarchical SD-WAN architecture, match a route by its path type, which can be one of the following: • **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. **Note** This option is available beginning with Cisco vManage Release 20.8.1.
Preference	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.
Prefix List	One OR more prefixes. Specifies the name of a prefix list.
Not available in Cisco SD-WAN Manager.	One OR more overlay network site identifiers.
Site	Individual site identifier. The range is 0 through 4294967295.

Match Condition	Description
Region	Region defined for Hierarchical SD-WAN. The range is 1 to 63. **Note** This option is available beginning with Cisco vManage Release 20.7.1.
Role	In a Hierarchical SD-WAN architecture, match by the device type, which can be **Border Router** OR **Edge Router**. **Note** This option is available beginning with Cisco vManage Release 20.8.1.
TLOC	Individual TLOC address. **Note** To use the **set TLOC** AND **set TLOC-list** commands, you must use the set-VPN command.
VPN	Individual VPN identifier. Range: 1 to 65525, excluding 512. For details see the VRF range behavior change described here. When a policy is pushed with duplicate VPN IDs, the latest VPN ID overwrites the first one. However, starting from Cisco IOS XE Catalyst SD-WAN Release 17.9.5, the first one is not overridden.
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.

Match Condition	Description
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 6:

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 Catalyst SD-WAN Release 17.9.1a and Cisco SD-WAN 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	Choose one of the following: • **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. **Note** Minimum releases: Cisco IOS XE Catalyst SD-WAN Release 17.9.1a, Cisco SD-WAN Release 20.9.1
DNS Application List	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. This match condition is available for IPv6 traffic from Cisco IOS XE Catalyst SD-WAN Release 17.9.1a and Cisco SD-WAN Release 20.9.1.
DNS	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**.
DSCP	Specifies the DSCP value.
Packet length	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 [-]).
Packet Loss Priority (PLP)	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.
Protocol	Specifies Internet protocol number. The range is 0 through 255.

Match Condition	Description
ICMP Message	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. When you select Protocol as Both, the **ICMP Message or ICMPv6 Message** field displays. **Note** This field is available from Cisco IOS XE Catalyst SD-WAN Release 17.4.1a and Cisco SD-WAN Release 20.4.1 .
Source Data Prefix	Specifies the group of source prefixes or an individual source prefix.
Source Port	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 [-]).
TCP Flag	Specifies the TCP flag, syn.

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 SD-WAN 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 7: 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
	1	reassembly-timeout
12		parameter-problem
	0	general-parameter-problem
	1	option-missing
	2	no-room-for-option
13	0	timestamp-request
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 8: 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

Lists the action parameters available for control policy configuration and their functions.

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:

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 these actions:

Action Condition	Description
Export To	Export the route the 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 SD-WANCisco 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. *address* is the system IP address. *color* can be one of **3g**, **biz-internet**, **blue**, **bronze**, **custom1**, **custom2**, **custom3**, **default**, **gold**, **green**, **lte**, **metro-ethernet**, **mpls**, **private1** through **private6**, **public-internet**, **red**, and **silver**. *encapsulation* can be **gre** or **ipsec**. Optionally, set a preference value (from 0 to 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.
TLOC Action	Direct matching routes or TLOCs using the mechanism specified by *action*, and enable end-to-end tracking of whether the ultimate destination is reachable. Setting 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.

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

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	Minimum release: Cisco IOS XE Catalyst SD-WAN Release 17.11.1a and Cisco vManage Release 20.11.1. Click **Log** to enable logging. 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. For information on **policy log-rate-limit** CLI, see **policy log-rate-limit** command in the Cisco Catalyst SD-WAN Qualified Command Reference Guide.

Action Condition	Description
Redirect DNS	

Action Condition	Description
	Redirects DNS requests to a particular DNS server or Umbrella. Redirecting requests is optional, but if you do so, you must specify both actions. For an inbound policy, **redirect-dns host** allows the DNS response to be correctly forwarded back to the requesting service VPN. For an outbound policy, specify the IP address of the DNS server. **redirect-dns umbrella** is only supported in Direct Internet Interface (DIA) use cases. It is not supported in SIG/SSE or overlay scenarios. When using **redirect-dns umbrella**, you do not need to explicitly configure **nat use-vpn 0**. **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). **Note** You can set only local TLOC preferences with redirect-dns as actions on the same sequence, but not remote TLOC. **Note** You cannot configure Redirect DNS and SIG at the same time. NAT DIA fallback and redirect-dns IP actions are supported at the same time in data policy beginning with Cisco IOS XE Catalyst SD-WAN Release 26.1.1. **Note** 1. Starting from Cisco IOS XE Release 17.15.5, data policy action **redirect-dns IP** supports SIG/SSE. You can configure **sig-action sse** action along with **redirect-dns IP** action. Data policy action **redirect-dns umbrella** supports Direct Internet Interface (DIA). You can configure **nat use-vpn 0** action along

Action Condition	Description
	with **redirect-dns umbrella** action. 2. Starting from Cisco IOS XE Catalyst SD-WAN Release 26.1.1, data policy action **redirect-dns IP** supports direct internet interface fallback. You can configure **nat use-van 0** and **nat fallback** action along with **redirect-dns IP** action.
TCP Optimization	Fine-tune TCP to decrease round-trip latency and improve throughout for matching TCP traffic.
Secure Internet Gateway	Redirect application traffic to a SIG. **Note** Before you apply a data policy for redirecting application traffic to a SIG, you must have configured the SIG tunnels. For more information on configuring Automatic SIG tunnels, see Automatic Tunnels . For more information on configuring Manual SIG tunnels, see Manual Tunnels. 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 SD-WAN Release 20.8.1 .

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.
DSCP	DSCP value. The range is 0 through 63.
Forwarding Class	Name of the forwarding class.

Action Condition	Description
Local TLOC	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. The encapsulation options are: **ipsec** and **gre**. 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. By default, encapsulation is **ipsec**.
Next Hop	Sets the next hop IP address to which the packet should be forwarded. **Note** Starting from Cisco IOS XE Catalyst SD-WAN Release 17.5.1a and Cisco SD-WAN 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.
Policer	Applies a policer. Specifies the name of policer configured with the **policy policer** command.
Service	Specifies a service to redirect traffic to before delivering the traffic to its destination. The TLOC address or list of TLOCs identifies the remote TLOCs to which the traffic should be redirected to reach the service. In the case of multiple TLOCs, the traffic is load-balanced among them. The VPN identifier is where the service is located. Standard services: **FW**, **IDS**, **IDP** Custom services: **netsvc1**, **netsvc2**,**netsvc3**, **netsvc4** TLOC list is configured with a **policy lists tloc-list** list. 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.
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. Range: 1 to 65525, excluding 512. For details see the VRF range behavior change described here.

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 9:

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

Configure policy application to sites and VPNs through the Apply Policies interface.

Apply policies to specific sites and VPNs to control network behavior and traffic management across your SD-WAN infrastructure.

Use the Apply Policies to Sites and VPNs page to associate configured policies with specific network locations and VPN segments. Different policy blocks require different site and VPN list associations based on their functionality.

Procedure

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:

   Table 10: Policy block associations

If policy block type is...	Then...
Topology	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**.
Application-Aware Routing	Click **New Site List** and **VPN list**. Choose one or more site lists, and choose one or more VPN lists. Click **Add**.
Traffic Data	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**.
cflowd	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.

The policy is successfully applied to the selected sites and VPNs and is available in the policies table for management and deployment.

Configure NAT fallback on Cisco IOS XE Catalyst SD-WAN Devices

Configure NAT fallback using Cisco SD-WAN Manager by creating and configuring a data policy.

This task enables NAT fallback functionality on Cisco IOS XE Catalyst SD-WAN Devices, allowing traffic to fall back to direct internet access (DIA) when tunnel connectivity is unavailable.

NAT DIA fallback provides redundancy for internet-bound traffic by automatically switching to direct internet access when the primary tunnel path fails. This configuration ensures continuous connectivity and improved network resilience.

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

Before you begin

Follow these steps to configure NAT fallback on Cisco IOS XE Catalyst SD-WAN Devices:

Procedure

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**.
11. Click **Centralized Policy** and for the required centralized policy, click **...** and select **Edit**.
12. Click **Traffic Rules** and select **Traffic Data**.
13. From the **Add Policy** drop-down, select **Import Existing**.
14. Select the NAT policy that you created from the **Policy** drop-down.
15. Click **Policy Application** and select **Traffic Data**.
16. Click **+ New Site List and VPN List**.
17. Select the direction, VPN, and site as required.
18. Click **Add**.
19. Click **Save Policy Changes**.
20. Click to select **VPN**, and **Site** from the drop-down.

NAT fallback is configured and enabled. Traffic will automatically fall back to direct internet access when tunnel connectivity is unavailable.

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

The following NAT fallback actions/commands are now supported:

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

Activate a centralized policy

Configure centralized policies to manage network traffic and routing across connected devices in your SD-WAN environment.

Activating a centralized policy sends that policy to all connected Cisco SD-WAN Controllers, enabling consistent policy enforcement across your SD-WAN network.

Centralized policies provide a unified approach to managing network behavior across your SD-WAN infrastructure. You can also view, copy, edit, and delete existing policies as needed to maintain your network configuration.

Procedure

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**.
4. To view centralized policies, select a policy from the **Centralized Policy** list.
5. 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.

6. For a policy created using Cisco SD-WAN Manager policy configuration wizard, click **...** and select **Preview**.

   This policy is displayed in text format.

7. To copy a policy, select a policy from the **Centralized Policy** list and click **...** and select **Copy**.
8. In the Policy Copy popup window, enter the policy name and a description of the policy.

> **Note**
>
> If you are upgrading to 18.4.4 version, data policy names need to be under 26 characters.

> **Note**
>
> Starting with the Cisco SD-WAN release 19.3Cisco IOS XE Release 17.2, 127 characters are supported for policy names for these policy types:
>
> - Central route policy
> - Local route policy
> - Local Access Control lOst (ACL)
> - Local IPv6 ACL
> - Central data policy
> - Central app route policy
> - QoS map
> - Rewrite rule
>
> All other policy names support 32 characters.

9. Click **Copy**.
10. To edit policies created using the Cisco SD-WAN Manager policy configuration wizard, click **...** and select **Edit**.
11. Edit the policy as needed and click **Save Policy Changes**.
12. To edit polices created using the CLI method, in the **Custom Options** drop-down, click **CLI Policy**.
13. For the desired policy, click **...** and select **Edit**.
14. Edit the policy as needed and click **Update**.
15. To delete policies, select a policy from the **Centralized Policy** list and click **...** and select **Delete**.
16. Click **OK** to confirm deletion of the policy.

The centralized policy is activated and distributed to all connected devices, or the selected policy management action is completed successfully.

Centralized policies configuration examples

This topic provides configuration examples involving centralized policies.

General centralized policy example

Control access

Restrict traffic

Traffic engineering

Creating arbitrary topologies

Community example

SIG data policy fallback

Ranking color preference example

Data policy for IPv6 applications example

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

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 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
  !
 !
!
```

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
```

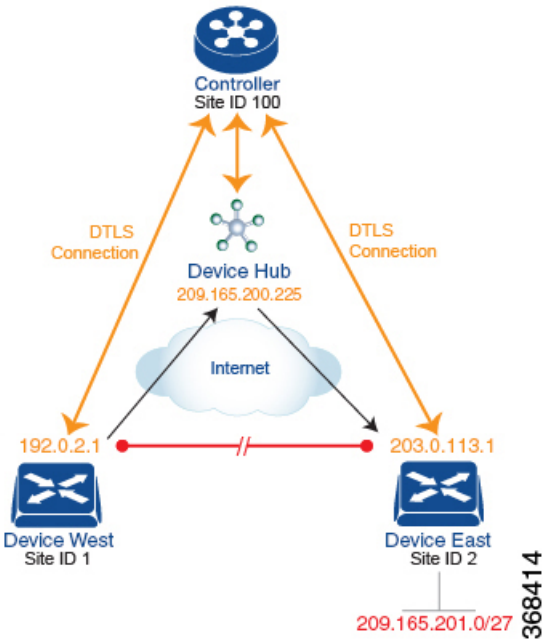
This examples illustrates how to disallow certain types of data traffic from being sent from between VPNs. This policy drops data traffic on port 25, which carries SMTP mail traffic, that originates in 209.165.201.0/27. However, the policy accepts all other data traffic, including non-SMTP traffic from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

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 encap ipsec
apply-policy
  site west-sites control-policy change-tloc out
```

A rough English translation of this policy is:

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

This control policy is configured on the Cisco 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.

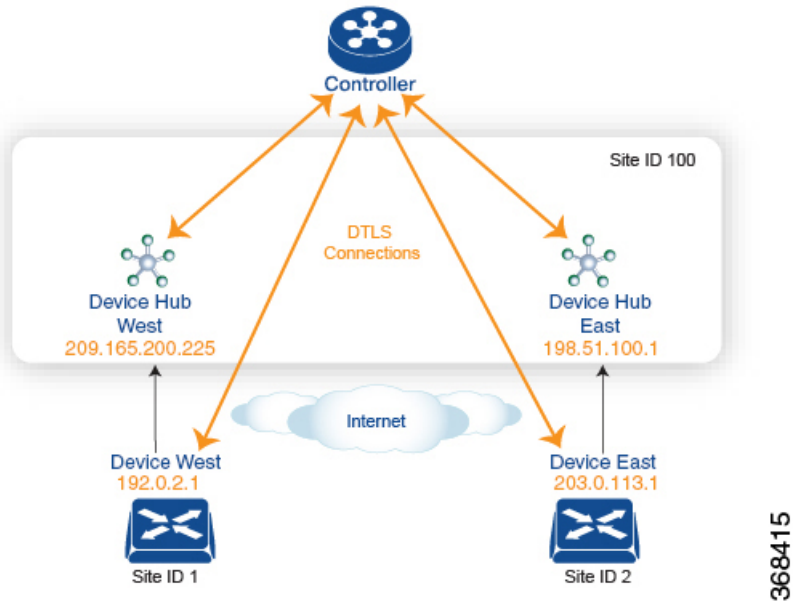
If the West side of the network had many sites instead of just one and each site had its own device, it would be straightforward to apply this same policy to all the sites. To do this, you simply add the site IDs of all the sites in the site-list west-sites list. This is the only change you need to make 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 encap ipsec
apply-policy
  site west-sites control-policy change-tloc out
```

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.

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*
```

From Cisco IOS XE Catalyst SD-WAN Release 17.8.1a and Cisco SD-WAN 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
```

```
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
   !
```

```
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
 !
!
```

4 Localized policy

Topics:

- Localized policies
- QoS for router-generated Cisco SD-WAN Manager traffic
- Localized policies
- Configure policies using a configuration group
- Configure localized policy for IPv4 using the CLI
- Configure localized policy for IPv6 using the CLI
- Localized data policy configuration examples
- Configure localized policy using classic policies

Localized policies

Introduces localized policy concepts and provides configuration guidance for IPv4 and IPv6 implementations, including Quality of Service for router-generated traffic.

A localized policy is a network configuration mechanism that enables administrators to define and apply network policies directly on network devices at specific locations.

Localized policy types and components

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.

QoS for router-generated Cisco SD-WAN Manager traffic

Provides information about prioritizing router-generated Cisco SD-WAN Manager traffic using QoS policies and class maps.

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.

Table 11: 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.

Localized policies

Describes policies that are provisioned locally through the CLI or through a device template.

A localized policy is 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.

A localized policy is a policy that is provisioned locally through the CLI on the Cisco vEdge devices, or through a Cisco SD-WAN Manager device template.

Types of localized policies

Describes the different categories of localized policies that operate on control plane and data plane traffic in SD-WAN networks.

Localized policies are network management configurations that

- operate directly on individual Cisco vEdge deviceCisco IOS XE Catalyst SD-WAN device devices rather than being centrally distributed,
- control either control plane routing behavior or data plane traffic flow at the local site level, and

- affect network operations within the specific branch or site where the device is deployed.

Localized control policy characteristics

Control policy operates on the control plane traffic in the Cisco SD-WANCisco 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 vEdge deviceCisco 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 vEdge deviceCisco 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 vEdge deviceCisco 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 characteristics

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

Localized data policy, so called because it is provisioned on the local Cisco vEdge deviceCisco 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 vEdge deviceCisco 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 vEdge deviceCisco IOS XE Catalyst SD-WAN devices, services enabled by default include:

- DHCP (for DHCPv4 and DHCPv6)
- DNS

- ICMP

You can also enable services for BGP, Netconf, NTP, OSPF, SSHD, and STUN.

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.

Data packet mirroring

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 vEdge deviceCisco IOS XE Catalyst SD-WAN devices support 1:1 mirroring; that is, a copy of every packet is sent to the alternate destination.

Configure policies using a configuration group

Provisions localized policies, including Access Control Lists, Quality of Service, and Route Policies, on Cisco IOS XE Catalyst SD-WAN devices using a centralized configuration group. It outlines the procedures for configuring objects, forwarding classes, ACLs, and route policies, and deploying the final configuration to target devices.

Use this task to provision localized policies, such as Access Control Lists (ACLs), Quality of Service (QoS), and Route Policies, on Cisco IOS XE Catalyst SD-WAN devices using a centralized configuration group approach.

Before you begin

- A configuration group must be created.
- A Transport & Management profile must be configured.

Follow these steps to configure policies using a configuration group:

Procedure

1. Configure **Objects and Policies** using a configuration group. For more information, refer to Configure objects and policies using a configuration group on page 100.
2. Configure **Forwarding Classes and QoS** using a configuration group. For more information, refer to Configure Forwarding Classes and QoS using a policy group on page 99.
3. Configure ACLs using a configuration group. For more information, refer to Configure ACLs using a configuration group on page 107.
4. Configure route policy using a configuration group. For more information, refer to Create Route Policy Using Configuration Group.

The localized policy is saved to the configuration group and is ready to be applied to the target devices.

What to do next

Deploy the configuration group to the target devices to push the localized policy settings to the network. For more information, Deploy a Configuration Group.

Configure Forwarding Classes and QoS using a policy group

Defines forwarding classes and Quality of Service maps within a policy group to prioritize network traffic and optimize bandwidth on Cisco IOS XE Catalyst SD-WAN devices. This configuration ensures consistent traffic management across your network infrastructure.

Use this task to define forwarding classes and Quality of Service (QoS) maps within a policy group to prioritize network traffic and ensure efficient bandwidth utilization on your Cisco IOS XE Catalyst SD-WAN devices.

Before you begin

A policy group must be created.

Follow these steps to configure forwarding classes and QoS using a policy group:

Procedure

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Policy Groups**.
2. Select the desired policy group and click **Edit**.
3. Click **Application Priority and SLA**.
4. Click **Add Traffic Policy** or select an existing policy to edit.
5. Click **Advanced Layout** to switch to the advanced view.
6. Click **QoS Queue**.
7. Click **Add QoS Policy** to define the QoS queue parameters.

   Table 12: QoS Queue

Field	Description
Queuing Model	Choose a value from the drop-down list for the queuing model.
Policy Name	Provide a name for the policy.
Interface	Specify a value for the interface.
Forwarding class	Choose a value for the forwarding class from the drop-down list.
Bandwidth %	Specify the maximum bandwidth. The range is 1-99.
Drops	Choose a value for the drop type from the following options: • **Random Early** • **Tail**
Scheduling type	Specify how to prioritize data packets for transmission to the destination by configuring the schedule type. The default is Weighted Round Robin (WRR).

8. Click **Save**.

The forwarding classes and QoS maps are saved to the policy group and are ready to be applied to the target devices.

What to do next

Deploy the policy group to the target devices to push the QoS settings to the network. For more information, Deploy Policy Group Workflow.

Configure objects and policies using a configuration group

Defines reusable policy objects, such as AS Path lists, community lists, and prefix lists, within a configuration group to serve as essential building blocks for localized policies.

Use this task to define reusable policy objects, such as AS Path lists, community lists, and prefix lists, within a configuration group. These objects serve as the building blocks for your localized policies.

Before you begin

Follow these steps to configure policy objects within a configuration group:

Procedure

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Configuration Groups**.
2. Select the desired configuration group and click **Edit**
3. Click **Objects and Profiles**.
4. Perform one of the following tasks based on your requirements:
   - Configure an AS Path list using a configuration group on page 100.
   - Configure Standard Community using a configuration group on page 101.
   - Configure Data Prefix using a configuration group on page 102.
   - Configure Extended Community using a configuration group on page 103.
   - Configure Class Map using a configuration group on page 104.
   - Configure Mirror using a configuration group on page 105.
   - Configure Policer using a configuration group on page 105.
   - Configure Prefix using a configuration group on page 106.

The policy objects are saved to the configuration group and are available to be referenced in your policy sequences.

What to do next

After defining your policy objects, you can reference them while configuring policies such as Access Control Lists (ACLs) or Route Policies. For more information, refer to Configure policies using a configuration group on page 98.

Configure an AS Path list using a configuration group

Creates a list of BGP AS paths that you can reference in route policies to control route acceptance, preference, or advertisement. This configuration enables precise management of route flow and traffic path selection across your network devices.

Use this task to create a list of BGP AS paths that you can reference in route policies to control route acceptance, preference, or advertisement.

Before you begin

- A configuration group must be created.
- A Transport and Management profile must be configured.

Follow these steps to configure an AS Path list:

Procedure

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Configuration Groups**.
2. Click **Objects and Profiles**
3. Select the **AS Path** policy object.
4. Click **Add AS Path**.
5. Enter a **Name** and **Description** for the AS Path list.
6. Enter the AS path number in the **AS path list ID** field. The range is 1 to 65535.
7. Enter the AS path in the **AS path list** field, separating AS numbers with a comma. You can write each AS as a single number or as a regular expression.
8. Click **Save**.

The AS Path list is saved to the configuration group and is ready to be referenced in route policy configurations.

What to do next

Deploy the configuration group to the target devices to push the localized policy settings to the network. For more information, Deploy a Configuration Group.

Configure Standard Community using a configuration group

Creates a list of standard or expanded BGP communities that you can reference in route policies to control route acceptance or advertisement. This configuration enables you to manage route propagation and influence traffic flow across your network devices.

Use this task to create a list of standard or expanded BGP communities that you can reference in route policies to control route acceptance or advertisement.

Before you begin

- A configuration group must be created.
- A Transport and Management profile must be configured.

Follow these steps to configure a community list:

Procedure

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Configuration Groups**.
2. Click **Objects and Profiles**.
3. Select the **Standard Community** policy object.
4. Click **Add Standard Community**.
5. Enter a **Name** and **Description** for the Standard Community.
6. In the **Standard Community** field, enter one or more communities separated by commas. You can use formats such as aa:nn (AS number and network number), internet, local-as, no-advertise, or no-export.

   The format example is given in the field.

Field	Description
Standard Community	• *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.

7. Click **Save**.

The community list is saved to the configuration group and is ready to be referenced in route policy configurations.

What to do next

Deploy the configuration group to the target devices to push the localized policy settings to the network. For more information, Deploy a Configuration Group.

Configure Data Prefix using a configuration group

Creates a list of IP prefixes that you can reference in data policies to permit or restrict traffic based on source or destination addresses. This configuration enables granular traffic control across your network.

Use this task to create a list of IP prefixes that you can reference in data policies to permit or restrict traffic based on source or destination addresses.

Before you begin

- A configuration group must be created.
- A Transport and Management profile must be configured.

Follow these steps to configure a data prefix list:

Procedure

---

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Configuration Groups**.
2. Click **Objects and Profiles**.
3. Select the **Data Prefix** policy object.
4. Click **Add Data Prefix**.
5. Enter a **Name** and **Description** for Data Prefix.

6. In the **Data Prefix** field, enter one or more IP prefixes separated by commas. You can enter IP prefixes for IPv4 or IPv6.
7. Click **Save**.

The data prefix list is saved to the configuration group and is ready to be referenced in data policy configurations.

What to do next

Deploy the configuration group to the target devices to push the localized policy settings to the network. For more information, Deploy a Configuration Group.

Configure Extended Community using a configuration group

Creates a list of BGP extended communities that you can reference in route policies for advanced route control. This configuration enables precise route filtering and traffic path manipulation across your network.

Use this task to create a list of BGP extended communities that you can reference in route policies for advanced route control.

Before you begin

- A configuration group must be created.
- A Transport and Management profile must be configured.

Follow these steps to configure an extended community list:

Procedure

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Configuration Groups**.
2. Click **Objects and Profiles**.
3. Select the **Extended Community** policy object.
4. Click **Add Extended Community**.
5. Enter a **Name** and **Description** for Extended Community.
6. In the **Extended Community** field, enter the community details.

   The format example is given in the field.

Field	Description
Extended Community	• **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.

The extended community list is saved to the configuration group and is ready to be referenced in route policy configurations.

What to do next

Deploy the configuration group to the target devices to push the localized policy settings to the network. For more information, Deploy a Configuration Group.

Configure Class Map using a configuration group

Defines a class map within a configuration group to categorize forwarding classes and map them to specific queues for Quality of Service management. This configuration enables efficient traffic prioritization and bandwidth allocation across your Cisco IOS XE Catalyst SD-WAN devices.

Use this task to create a class map within a configuration group to define forwarding classes and map them to specific queues for Quality of Service (QoS) management.

Before you begin

- A configuration group must be created.
- A Transport and Management profile must be configured.

Follow these steps to configure a class map:

Procedure

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Configuration Groups**.
2. Click **Objects and Profiles**.
3. Select the **Class Map** policy object.
4. Click **Add Class Map**.
5. Enter a **Name** and **Description** for Class Map.
6. Select the required **Queue** from the drop-down list..
7. Click **Save**.

The class map is saved to the configuration group and is ready to be referenced in QoS policy configurations.

What to do next

Deploy the configuration group to the target devices to push the localized policy settings to the network. For more information, Deploy a Configuration Group.

Configure Mirror using a configuration group

Specifies source and destination IP addresses for traffic mirroring to facilitate network troubleshooting and traffic analysis. This configuration enables you to capture and monitor data packets across your network.

Use this task to specify source and destination IP addresses for traffic mirroring to facilitate network troubleshooting and traffic analysis.

Before you begin

- A configuration group must be created.
- A Transport and Management profile must be configured.

Follow these steps to configure a mirror list:

Procedure

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Configuration Groups**.
2. Click **Objects and Profiles**.
3. Select the **Mirror** policy object.
4. Click **Add Mirror**.
5. Enter a **Name** and **Description** for Mirror.
6. In the **Remote Destination IP** field, enter the IP address of the destination for which to mirror the packets.
7. In the **Source IP** field, enter the IP address of the source of the packets to mirror.
8. Click **Save**.

   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.

The mirror list is saved to the configuration group and is ready to be referenced in localized data policy configurations.

What to do next

Deploy the configuration group to the target devices to push the localized policy settings to the network. For more information, Deploy a Configuration Group.

Configure Policer using a configuration group

Defines policing parameters, such as rate and burst size, to control traffic flow and prevent network congestion within a configuration group. This configuration allows you to enforce bandwidth limits and maintain network performance across your Cisco IOS XE Catalyst SD-WAN devices.

Use this task to define policing parameters, such as rate and burst size, to control traffic flow and prevent network congestion.

Before you begin

- A configuration group must be created.

- A Transport and Management profile must be configured

Follow these steps to configure a policer:

Procedure

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Configuration Groups**.
2. Click **Objects and Profiles**.
3. Select the **Policer** policy object.
4. Click **Add Policer**.
5. Enter a **Name** and **Description** for Policer.
6. In the **Burst (bytes)** field, enter the maximum traffic burst size in bytes.
7. In the **Exceed** drop-down list, choose the action **Drop** or **Remark**.
8. In the **Rate (bps)**, enter the maximum traffic rate in bits per second (bps).

   The following table describe the options for configuring the policer.

Field	Description
Burst (bytes)	Specifies the maximum traffic burst size. Range is from 15000 to 10000000.
Exceed	Specifies an action to take when the burst size or traffic rate is exceeded. The options are: **Drop**–Sets the packet loss priority (PLP) to low. **Remark**–Sets the PLP to high. The default option is **Drop**.
Rate	Specifies the maximum traffic rate. It can be a value from 8 through 2^{64} bps (8 through 100000000000).

9. Click **Save**.

The policer is saved to the configuration group and is ready to be referenced in data policy or ACL configurations.

What to do next

Deploy the configuration group to the target devices to push the localized policy settings to the network. For more information, Deploy a Configuration Group.

Configure Prefix using a configuration group

Creates a list of IP prefixes that you can reference in policies to match specific network traffic or routes. This configuration enables precise traffic and route control across your network.

Use this task to create a list of IP prefixes that you can reference in policies to match specific network traffic or routes.

Before you begin

- A configuration group must be created.
- A Transport and Management profile must be configured.

Follow these steps to configure a prefix list:

Procedure

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Configuration Groups**.
2. Click **Objects and Profiles**.
3. Select the **Prefix** policy object.
4. Click **Add Prefix**.
5. Enter a **Name** and **Description** for Prefix.
6. In the **Prefix** field, enter one or more IP prefixes separated by commas. You can enter IP prefixes for IPv4 or IPv6.
7. Under **Add Prefix**, enter the prefix for the list. Optionally, click the **Choose a file** link to import a prefix list.
8. Click **Save**.

The prefix list is saved to the configuration group and is ready to be referenced in control or data policy configurations.

What to do next

Deploy the configuration group to the target devices to push the localized policy settings to the network. For more information, Deploy a Configuration Group.

Configure ACLs using a configuration group

Configure ACL IPv4 or ACL IPv6 features within configuration groups to control network traffic access.

This task allows you to configure access control lists (ACLs) for both IPv4 and IPv6 traffic using configuration groups, providing granular control over network traffic filtering and management.

ACLs in configuration groups enable centralized management of traffic filtering policies across multiple devices in your SD-WAN deployment.

Before you begin

On the **Configuration** > **Configuration Groups** page, choose **SD-WAN** as the solution type.

Follow these steps to configure ACLs using a configuration group:

Procedure

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Configuration Groups**.
2. Create and configure an Ethernet Interface feature in a Transport & Management profile.
3. Click the plus (+) adjacent to the Ethernet Interface feature to display the Add Feature pane.
4. Add a ACL IPv4 or ACL IPv6 feature to the Ethernet Interface feature.

   The following table describe the options for configuring the ACL IPv4 feature.

   Table 13: ACL IPv4

Field	Description
ACL Sequence Name	Specifies the name of the ACL sequence.

Field	Description
Condition	Specifies the ACL condition. The options are: • DSCP • Packet Length • Protocol • Source Data Prefix • Source Port • Destination Data Prefix • Destination Port • TCP • ICMP Message (Only in Cisco Catalyst SD-WAN Manager Release 20.15.1)
Action Type	Specifies the action type. The options are: Accept or Reject.
Accept Condition	Specifies the accept condition type. The options are: • Counter • DSCP • Log • Next Hop • Service Chain • Policer

The following table describe the options for configuring the ACL IPv6 feature.

Table 14: ACL IPv6

Field	Description
ACL Sequence Name	Specifies the name of the ACL sequence.

Field	Description
Condition	Specifies the ACL condition. The options are: • Next Header • Packet Length • Source Data Prefix • Source Port • Destination Data Prefix • Destination Port • TCP • Traffic Class • ICMP Message (Only in Cisco Catalyst SD-WAN Manager Release 20.15.1)
Action Type	Specifies the action type. The options are: Accept or Reject.
Accept Condition	Specifies the accept condition type. The options are: • Counter • Log • Next Hop • Traffic Class • Service Chain • Policer

Table 15: Match 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.

Match Condition	Description
ICMP Message	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. 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. **Note** This field is available from Cisco IOS XE Release 17.4.1, Cisco SD-WAN Release 20.4.1 Cisco vManage Release 20.4.1. For `ICMP-msg` and `icmp6-msg` message types, refer to the ICMP Message Types/Codes and Corresponding Enumeration Values table in the Centralized chapter.
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.
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**

Table 16: Action parameters

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 vEdge deviceCisco IOS XE Catalyst SD-WAN device.
Drop	Discards the packet. This is the default action.

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.

Description	Value or Range
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 17.5.1aCisco SD-WAN Release 20.5.1, **Use Default Route when Next Hop is not available** field is available next to Next Hop action parameter.

The ACL IPv4 or ACL IPv6 feature is successfully added to the Ethernet Interface feature within the configuration group, enabling traffic filtering based on the configured match and action parameters.

What to do next

Also see Deploy a configuration group.

Match parameters

Lists the parameters that access lists and route policies can match for IP prefixes and header fields.

Access lists can match IP prefixes and fields in the IP headers. Route policies can match various routing parameters including BGP attributes and OSPF tags.

Access list parameters

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.

Match Condition	Description
ICMP Message	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. 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. **Note** This field is available from Cisco IOS XE Release 17.4.1, Cisco SD-WAN Release 20.4.1 Cisco vManage Release 20.4.1. For **ICMP-msg** and **icmp6-msg** message types, refer to the ICMP Message Types/Codes and Corresponding Enumeration Values table in the Centralized chapter.
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.
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.

Match Condition	Description
Community List	List of one of more BGP communities. In **Community List**, you can specify: · *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: · **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.
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 optionss are: EGP (default), IGP, Incomplete. **Note** Match origin is available only for Cisco vEdge devices not on Cisco IOS XE Catalyst SD-WAN devices.
OSPF Tag	Specifies the OSPF tag number. The range is 0 through 4294967295.
Peer	Specifies the peer IP address.

Action parameters

Lists action parameters for access lists and route policies.

Action parameters define how packets are processed when they match conditions in access lists or route policies. For access lists, packets can be accepted or dropped, counted, classified, mirrored, or policed. For route policies, routes can be accepted or rejected with various attribute modifications.

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 vEdge deviceCisco IOS XE Catalyst SD-WAN device.
Drop	Discards the packet. This is the default action.

For a packet that is accepted, these 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 17.5.1aCisco SD-WAN Release 20.5.1, **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, these actions can be configured:

Description	Value or Range
Aggregator	Sets 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 17.5.1aCisco SD-WAN Release 20.5.1, **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 localized policy for IPv4 using the CLI

Configure localized IPv4 policy using CLI commands on Cisco IOS XE Catalyst SD-WAN devices.

Configure localized policy for IPv4 traffic management using CLI commands to control access, QoS, and policing on router interfaces.

Use this procedure to implement IPv4 policy configurations that control traffic flow, quality of service, and access control on Cisco IOS XE Catalyst SD-WAN devices through command-line interface commands.

Procedure

1. Create lists of IP prefixes as needed:

   Example

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

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

   Example

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

   Note

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

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

   Example

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

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

   Example

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

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

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

   Example

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

6. Configure a WAN interface without a shaping configuration:

   Example

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

7. Configure a WAN interface with a shaping configuration:

   Example

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

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

   Example

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

9. Define policing parameters:

   Example

```
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
```

The localized IPv4 policy is configured with QoS classes, rewrite rules, policing parameters, and access lists applied to the appropriate interfaces.

Configure localized policy for IPv6 using the CLI

Configure localized policy for IPv6 using the CLI by defining mirroring and policing parameters, creating access lists with match-action sequences, and applying them to interfaces.

This task allows you to implement traffic control and security policies for IPv6 traffic by creating access lists with specific match criteria and corresponding actions such as dropping, counting, mirroring, or policing packets.

Use this configuration when you need to control IPv6 traffic flow through your network interfaces by applying specific filtering, mirroring, and policing policies. This is particularly useful for traffic management, security enforcement, and network monitoring.

Before you begin

Follow these steps to configure localized policy for IPv6 using the CLI:

Procedure

1. Define mirroring parameters (for unicast traffic only):

   Example

   ```
   vEdge(config)# policy mirror mirror-name
   vEdge(config-mirror)# remote-dest ip-address source ip-address
   ```

2. Define policing parameters:

   Example

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

   ```
   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
   ```

3. Create an access list instance:

   Example

   ```
   vEdge(config)# policy ipv6 access-list list-name
   ```

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

4. Create a series of match-action pair sequences:

Example

```
vEdge(config-ipv6-access-list)# sequence number
vEdge(config-sequence)#
```

```
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).

5. Define match parameters for packets:

Example

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

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

6. Define actions to take when a match occurs:

Example

```
vEdge(config-sequence)# action drop
vEdge(config-sequence)# action count counter-name
vEdge(config-sequence)# action log
vEdge(config-sequence)# action accept class class-name
vEdge(config-sequence)# action accept mirror mirror-name
vEdge(config-sequence)# action accept policer policer-name
```

```
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
```

7. Create additional numbered sequences of match-action pairs within the access list, as needed.
8. If you want nonmatching packets to be accepted, configure the default action for the access list:

Example

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

If a packet does not match any of the conditions in one of the sequences, it is rejected by default.

9. Apply the access list to an interface:

Example

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

```
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.

The IPv6 access list is configured and applied to the specified interface, enabling traffic filtering based on the defined match criteria and actions. The policy will now process IPv6 packets according to the configured sequences and take the appropriate actions such as dropping, counting, mirroring, or policing traffic as specified.

Localized data policy configuration examples

Provides straightforward examples of configuring localized data policy to influence traffic flow across the Cisco Catalyst SD-WAN domain.

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.

Mirroring example

This example illustrates how to configure a mirror instance to automatically send a copy of certain types of data packet to a specified destination for analysis. After you configure the mirror instance, include it in an access list. Here, "mirror-m1" is configured with the host at source address 10.20.23.16 AND destination host at 10.2.2.11. The mirror instance is then included in the access list "acl2," which is configured so that data packets originating from the host at source address 10.20.24.17 AND going to the destination host at 10.20.25.18 are mirrored to the destination host at 10.2.2.11 with the source address of the originating host as 10.20.23.16.

```
policy
 mirror m1
  remote-dest 10.2.2.11 source 10.20.23.16
 !
!

vm5# show running-config policy access-list acl2
policy
 access-list acl2
  sequence 1
   match
    source-ip      10.20.24.17/32
    destination-ip 10.20.25.18/32
   !
   action accept
    mirror m1
   !
  !
  default-action drop
 !
!
```

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
  !
  !
```

Configure localized policy using classic policies

Describes how to configure localized policies using the Cisco SD-WAN Manager policy configuration wizard.

A localized policy configuration is a UI policy builder that

- uses the Cisco SD-WAN Manager policy configuration wizard
- consists of five windows to configure and modify localized policy components, and
- allows configuration of some or all components depending on the specific policy being created.

Localized policy components

The wizard configures and modifies these localized policy components:

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

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

Configure access to the policy configuration wizard to begin creating localized policies.

Start the policy configuration wizard to create and configure localized policies for your network.

The policy configuration wizard provides a guided interface for creating localized policies. Use this wizard when you need to establish network policies for specific devices or groups of devices in your SD-WAN deployment.

Procedure

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, where you can begin configuring your policy.

Configure groups of interest for localized policy

Configure groups of interest lists to define collections of network parameters that can be referenced in localized policies for BGP routing control.

Groups of interest allow you to create reusable lists of network parameters such AS paths, communities, prefixes, and other BGP attributes that can be referenced in localized policies to control routing behavior.

In Create Groups of Interest, create lists of groups to use in a localized policy. These lists help organize and manage network policies by grouping related parameters together for easier reference and maintenance.

Before you begin

Follow these steps to configure groups of interest for localized policy:

Procedure

1. Configure AS Path lists to specify BGP AS paths.
   a. In the group of interest list, click **AS Path**.
   b. Click **New AS Path List**.
   c. Enter a name for the list.
   d. Enter the AS path, separating AS numbers with a comma.
   e. 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.
2. Configure Community lists to create groups of BGP communities.

   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.
   a. In the group of interest list, click **Community**.
   b. Click **New Community List**.
   c. Enter a name for the community list.
   d. In the **Add Community** field, enter one or more data prefixes separated by commas in any of these 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.

e. Click **Add**.

3. Configure Data Prefix lists to specify IP prefixes.
   a. In the **Group of Interest** list, click **Data Prefix**.
   b. Click **New Data Prefix List**.
   c. Enter a name for the list.
   d. Enter one or more IP prefixes.
   e. 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.

4. Configure Extended Community lists for BGP extended communities.
   a. In the group of interest list, click **Extended Community**.
   b. Click **New Extended Community List**.
   c. Enter a name for the list.
   d. Enter the BGP extended community in these 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.

   e. Click **Add**.

5. Configure Class Map lists to define traffic classes with associated queues.
   a. In the group of interest list, click **Class Map**.
   b. Click **New Class List**.
   c. Enter a name for the class.
   d. Select a required queue from the **Queue** drop-down list.
   e. Click **Save**.

6. Configure Mirror lists to define packet mirroring parameters.
   a. In the group of interest list, click **Mirror**.
   b. Click **New Mirror List**. The Mirror List popup displays.
   c. Enter a name for the list.
   d. In the **Remote Destination IP** field, enter the IP address of the destination for which to mirror the packets.

e. In the **Source IP** field, enter the IP address of the source of the packets to mirror.

f. 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.

7. Configure Policer lists to define traffic policing parameters.

   a. In the group of interest list, click **Policer**.

   b. Click **New Policer List**.

   c. Enter a name for the list.

   d. In the **Burst (bps)** field, enter maximum traffic burst size. It can be a value from 15000 to 10000000 bytes.

   e. 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.

   f. In the **Rate (bps)** field, enter the maximum traffic rate. It can be value from 8 through 2^{64} bps (8 through 100000000000).

   g. Click **Add**.

   > **Note**
   >
   > When you define policers within localized policy lists, the device applies all defined policers to all localized policies. This happens because localized policies do not recognize which policers are referenced in templates. The device sends all policers through the Localized Policy CLI to ensure proper functionality.

8. Configure Prefix lists to specify IP prefixes with protocol selection.

   a. In the group of interest list, click **Prefix**.

   b. Click **New Prefix List**.

   c. Enter a name for the list.

   d. In the **Internet Protocol** field, click either **IPv4** or **IPv6**.

   e. Under **Add Prefix**, enter the prefix for the list. (An example is displayed.) Optionally, click the green **Import** link on the right-hand side to import a prefix list.

   f. Click **Add**.

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

You have successfully configured groups of interest lists for use in localized policies. These lists can now be referenced in policy configurations to control BGP routing behavior, traffic classification, and network security policies.

Configure forwarding classes/QoS

Configure QoS mappings and policy rewrite rules to manage traffic forwarding classes and quality of service parameters.

This task allows you to set up quality of service (QoS) mappings and policy rewrite rules to control traffic prioritization and forwarding behavior on network devices.

When you first open the **Forwarding Classes/QoS** page, **QoS Map** is selected by default. 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 vEdge deviceCisco 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 vEdge deviceCisco IOS XE Catalyst SD-WAN device, you must enable QoS scheduling and shaping. To enable QoS parameters for traffic that the Cisco vEdge deviceCisco IOS XE Catalyst SD-WAN device receives from transport-side interfaces and service-side interfaces.

Procedure

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**.
9. To import an existing QoS mapping, in **QoS**, click the **Add QoS Map** drop-down and select **Import Existing**.

   The **Import Existing Application QoS Map Policy** popup displays.

You have successfully configured QoS mappings and policy rewrite rules for traffic forwarding classes. The QoS configuration will control traffic prioritization and quality of service parameters on your network devices.

Configure ACLs

Configure access control lists (ACLs) to control network traffic by defining match conditions and corresponding actions.

Configure access control lists to define policies that control network traffic flow and security by specifying match conditions and corresponding actions for IPv4 or IPv6 traffic.

Access control lists provide a method to filter network traffic based on various criteria. You can create IPv4 or IPv6 ACL policies, import existing policies, and configure default actions for unmatched traffic.

Procedure

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 and select one of 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.

   If you click **Add IPv4 ACL Policy**, the **Add IPv4 ACL Policy** page appears.

If you click **Add IPv6 ACL Policy**, the **Add IPv6 ACL Policy** page appears.

3. Enter a name and description for the ACL in the **ACL Policy** page.
4. In the left pane, click **Add ACL Sequence**.

   An **Access Control List** box is displayed in the left pane.

5. Double-click the **Access Control List** box, and type a name for the ACL.
6. In the right pane, click **Add Sequence Rule** to create a single sequence in the ACL.

   Match is selected by default.

7. Click a match condition.
8. On the left, enter the values for the match condition.

   On the right enter the action or actions to take if the policy matches.

9. Repeat Steps 6 through 8 to add match-action pairs to the ACL.

The ACL policy is configured and saved. You can now proceed to configure additional policies or continue with route policy configuration.

What to do next

To configure **Device Access Policy**, see Device Access Policy.

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

Explicit and implicit access lists

Describes the two types of access lists available for network traffic control and their interaction when both are applied to the same traffic.

Explicit and implicit access lists are network security mechanisms that

- control traffic flow through different configuration methods
- operate at different interface levels and VPN scopes, and
- interact with specific precedence rules when both types apply to the same traffic.

Access list types and configuration

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 vEdge deviceCisco IOS XE Catalyst SD-WAN devices, these 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 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.

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

Table 17: Traffic handling with combined explicit and implicit ACLs

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)

Note

During software upgrades, you can connect to a remote repository for downloading images and configure tunnel interface using **allow-service all** to implicitly permit all the services, or set up an explicit ACL to allow SCP packets. This is a sample explicit ACL configuration:

```
policy
 access-list allow-remote-scp
  sequence 1
   match
    source-ip      10.0.0.0/32
    destination-ip 10.0.0.1/32
    source-port    22
    protocol       6
   !
   action accept
   !
  !
  default-action accept
 !
!
sdwan
 interface GigabitEthernet1
  tunnel-interface
   encapsulation ipsec
   color gold
   no allow-service dhcp
   no allow-service dns
   allow-service icmp
   no allow-service sshd
   no allow-service netconf
   no allow-service ntp
   no allow-service ospf
   no allow-service stun
   no allow-service https
   no allow-service snmp
   no allow-service bfd
  exit
  access-list allow-remote-scp in
 exit
```

Configure route policies

Configure routing policies to control routing behavior and traffic flow in your network.

Configure routing policies to define how routes are selected, filtered, and modified based on specific match conditions and actions.

Route policies allow you to control routing behavior by creating rules that match specific conditions and apply corresponding actions to those routes.

Before you begin

Follow these steps to configure route policies:

Procedure

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.

The route policy is configured and ready to control routing behavior based on the defined match conditions and actions.

Match parameters

Lists the parameters that access lists and route policies can match for IP prefixes and header fields.

Access lists can match IP prefixes and fields in the IP headers. Route policies can match various routing parameters including BGP attributes and OSPF tags.

Access list parameters

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.

Match Condition	Description
ICMP Message	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. 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. **Note** This field is available from Cisco IOS XE Release 17.4.1, Cisco SD-WAN Release 20.4.1 Cisco vManage Release 20.4.1. For **ICMP-msg** and **icmp6-msg** message types, refer to the ICMP Message Types/Codes and Corresponding Enumeration Values table in the Centralized chapter.
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.
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.

Match Condition	Description
Community List	List of one of more BGP communities. In **Community List**, you can specify: · *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: · **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.
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 optionss are: EGP (default), IGP, Incomplete. **Note** Match origin is available only for Cisco vEdge devices not on Cisco IOS XE Catalyst SD-WAN devices.
OSPF Tag	Specifies the OSPF tag number. The range is 0 through 4294967295.
Peer	Specifies the peer IP address.

Action parameters

Lists action parameters for access lists and route policies.

Action parameters define how packets are processed when they match conditions in access lists or route policies. For access lists, packets can be accepted or dropped, counted, classified, mirrored, or policed. For route policies, routes can be accepted or rejected with various attribute modifications.

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 vEdge deviceCisco IOS XE Catalyst SD-WAN device.
Drop	Discards the packet. This is the default action.

For a packet that is accepted, these 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 17.5.1aCisco SD-WAN Release 20.5.1, **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, these actions can be configured:

Description	Value or Range
Aggregator	Sets 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 17.5.1aCisco SD-WAN Release 20.5.1, **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

Configure policy settings in Policy Overview to define localized master policy parameters and application settings.

Configure policy settings to establish network traffic monitoring, application tracking, QoS scheduling, and ACL logging parameters for your localized master policy.

Policy settings define how network traffic is monitored and managed across your network infrastructure. These settings control various aspects including traffic flow monitoring, application tracking, and quality of service scheduling.

Before you begin

Follow these steps to configure policy settings:

Procedure

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**.

The policy settings are configured and saved. The localized master policy is now ready for deployment with the specified monitoring, tracking, and scheduling parameters.

Apply localized policy in a device template

Configure a localized policy within a device template to apply specific policy settings to individual devices.

Apply a localized policy to a device template to customize policy settings for specific devices or groups of devices within your SD-WAN deployment.

Device templates allow you to standardize configurations across multiple devices while incorporating localized policies for specific requirements. You can either create a new device template with a localized policy or add a localized policy to an existing template.

Before you begin

Ensure you have already configured the localized policy that you want to apply to the device template.

Follow these steps to apply a localized policy in a device template:

Procedure

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Templates**.
2. Create a new device template or select an existing template to edit.
   - 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 vEdge deviceCisco 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.

- 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.

3. Click **Additional Templates** located directly beneath the **Description** field.

   The screen scrolls to the **Additional Templates** section.

4. From the **Policy** drop-down, select the name of the policy you configured in the above procedure.
5. Click **Create** (for a new template) or **Update** (for an existing template).

The localized policy is successfully applied to the device template and will be deployed to devices that use this template.

Activate a localized policy

Configure a localized policy by activating, viewing, copying, editing, or deleting it as needed.

Activating a localized policy pushes the policy to all reachable Cisco SD-WAN Controllers in the network, enabling centralized policy management and enforcement across your SD-WAN infrastructure.

Localized policies provide granular control over network traffic and routing decisions. You can manage these policies through the Cisco SD-WAN Manager interface, where you can activate, view, copy, edit, and delete policies as needed for your network requirements.

Before you begin

Follow these steps to activate a localized policy:

Procedure

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.
8. To view localized policies, click **Localized Policy**, and select a policy.
9. 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.

The localized policy is successfully activated and pushed to all reachable Cisco SD-WAN Controllers in the network. You can view, copy, edit, or delete the policy as needed for ongoing network management.

5 Default AAR and QoS policies

Topics:

- Default AAR and QoS policies
- Prerequisites for default AAR and QoS policies
- Restrictions for default AAR and QoS policies
- Supported devices for default AAR and QoS policies
- Default AAR and QoS policy use cases
- Configure default AAR and QoS policies using Cisco SD-WAN Manager
- Monitor default AAR and QoS policies

Lists the development milestones and release information for default AAR and QoS policy feature, including feature description.

This reference provides information about the default application-aware routing (AAR) and quality of service (QoS) policies, including their development milestones, release information, and feature description.

Table 18: 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 the edge 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. You can configure Default AAR and QoS policies.

Default AAR and QoS policies

Explains how the system uses Network-Based Application Recognition (NBAR) to categorize over 1,000 applications into three business-relevance categories (business-relevant, business-irrelevant, and default) to prioritize traffic.

A default AAR and QoS policy is a network traffic management policy that

- uses Network-Based Application Recognition (NBAR) to identify and categorize more than 1,000 applications
- groups applications into three business-relevance categories (business-relevant, business-irrelevant, and default), and
- prioritizes traffic based on the business relevance of applications.

Business-relevance categories for applications

The workflow presents a set of more than 1,000 applications that can be identified by NBAR, an application recognition technology built into edge 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 edge devices in the network.

NBAR application recognition example

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 method that NBAR uses to recognize traffic is DNS Snooping. For NBAR to correctly categorize certain types of traffic, the unencrypted DNS traffic must pass through the router. 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.

Default AAR and QoS policies

Explains the primary advantages of default AAR and QoS policies, including the ability to customize bandwidth allocations and prioritize applications based on their importance to business operations.

A default AAR and QoS policy is a network management feature that

- enables management and customization of bandwidth allocations
- allows prioritization of applications based on their relevance to business operations, and
- supports flexible configuration to meet organizational requirements.

Prerequisites for default AAR and QoS policies

Outlines the requirements for using default AAR and QoS policies, including knowledge of relevant applications and familiarity with SLAs and QoS markings to prioritize traffic.

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

Restrictions for default AAR and QoS policies

Outlines the limitations of default AAR and QoS policies, including the lack of IPv6 support and the requirement that business-relevant application groups must contain at least one application.

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

Lists the Cisco hardware and software platforms compatible with default AAR and QoS policies.

This reference identifies the Cisco hardware and software platforms that support default AAR and QoS policies.

The following Cisco devices are compatible with 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

Default AAR and QoS policy use cases

Explains the ideal scenario for this feature: quickly creating and deploying consistent AAR and QoS policies across all devices when setting up a Cisco Catalyst SD-WAN network.

A default AAR and QoS policy use case is a deployment scenario that

- enables quick creation and deployment of consistent AAR and QoS policies
- applies these policies across all devices in a network, and
- supports efficient setup of a Cisco Catalyst SD-WAN network.

Configure default AAR and QoS policies using Cisco SD-WAN Manager

Configure default AAR and QoS policies using Cisco SD-WAN Manager to define application groups, path preferences, SLA classes, and bandwidth mappings.

This task enables you to configure default Application-Aware Routing (AAR) and Quality of Service (QoS) policies using Cisco SD-WAN Manager. It guides you through defining application groups, path preferences, SLA classes, and bandwidth mappings to optimize network performance and traffic management.

Use this task when you need to implement or modify default AAR and QoS policies in your Cisco SD-WAN environment. This workflow is relevant for administrators who want to ensure optimal application performance and bandwidth allocation across the network using Cisco SD-WAN Manager.

Before you begin

Ensure you have access to Cisco SD-WAN Manager and the necessary permissions to configure policies.

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

Procedure

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Policies**. Click **Add Default AAR & QoS**.

   The **Process Overview** page is displayed.

2. Click **Next**. Based on the requirements of your network, move the applications between the **Business Relevant**, **Default**, and **Business Irrelevant** groups.

   The **Recommended Settings based on your selection** page is displayed.

   > **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.

3. Click **Next**.

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

4. 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.

5. Click **Next**.

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

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**. If necessary, customize the bandwidth percentage values for each queue.

6. Click **Next**.

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

   For each policy, enter a prefix name and description.

7. 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.

   Table 19: Workflow steps and effects

Workflow Step	Affects the Following
Recommended Settings based on your selection	AAR and data policies
Path Preferences (optional)	AAR policies
App Route Policy Service Level Agreement (SLA) Class: • 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

8. Click **Configure**.

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

9. 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 edge devices, attach it to a localized policy through device templates.

After completing this workflow, default AAR, data, and QoS policies are created and applied as configured. You can view and edit the policies as needed, and the network will use the defined application groups, path preferences, SLA classes, and bandwidth mappings for optimized traffic management.

What to do next

For details on mapping application lists to queues, refer to the following information.

- 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

Configure navigation in Cisco SD-WAN Manager to verify the status of default Application Aware Routing (AAR), traffic data, and QoS policies. This task enables you to view policy lists and confirm their deployment using the platform's policy menus.

This task allows you to monitor and verify the status of default Application Aware Routing (AAR) and Quality of Service (QoS) policies in Cisco SD-WAN Manager.

- Ensures visibility into created AAR, traffic data, and QoS policies.
- Supports validation of policy deployment and effectiveness.

Use this task when you need to check the operational status or configuration of default AAR and QoS policies within Cisco SD-WAN Manager.

Monitoring these policies is essential for maintaining optimal network performance and ensuring that traffic management rules are correctly applied.

- Applicable after policy creation or modification.
- Useful for troubleshooting or routine verification of policy status.

Before you begin

Access to Cisco SD-WAN Manager is required.

- Ensure you have appropriate permissions to view configuration and policy menus.

Follow these steps to monitor default AAR and QoS policies:

Procedure

1. Monitor default AAR policies in Cisco SD-WAN Manager.

- From the Cisco SD-WAN Manager menu, choose **Configuration** > **Policies**.
- Click **Custom Options**.
- Choose **Traffic Policy** from **Centralized Policy**.
- Click **Application Aware Routing**.

  A list of AAR policies is displayed.
- Click **Traffic Data**.

  A list of traffic data policies is displayed.

2. Monitor QoS policies in Cisco SD-WAN Manager.

- From the Cisco SD-WAN Manager menu, choose **Configuration** > **Policies**.
- Click **Custom Options**.
- Choose **Forwarding Class/QoS** from **Localized Policy**.
- Click **QoS Map**.

  A list of QoS policies is displayed.

After completing these steps, you will be able to view lists of default AAR, traffic data, and QoS policies in Cisco SD-WAN Manager. You can confirm their status and verify that the policies are correctly deployed and operational.

What to do next

For further verification or troubleshooting, consult the referenced documentation or review policy logs as needed.

- Refer to external resources for advanced policy validation.

6 Redirect DNS in Cisco Catalyst SD-WAN Service-Side VPNs

Topics:

- Feature history for redirect DNS in a service-side VPN
- Redirect DNS in a service-side VPN
- Configure redirect DNS in a service-side VPN
- Verify redirect DNS in a service-side VPN

Describes how to configure DNS redirection within Cisco Catalyst SD-WAN service-side VPNs, enabling optimized DNS resolution for branch users and improving network performance and security through centralized DNS management.

Feature history for redirect DNS in a service-side VPN

Lists the development milestones and release information for redirect DNS in a service-side VPN, including feature description.

This table describes the developments of this feature, by release.

Table 20: Feature history

Feature name	Release information	Description
Redirect DNS in a Service-Side VPNs	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 servce-side VPN hosts and DNS redirects inside the service VPNs. You can configure redirect DNS using Cisco SD-WAN Manager.

Redirect DNS in a service-side VPN

Explains Redirect DNS, outlines policy options, and describes differences between dns-app-list and app-list for traffic classification in SD-WAN.

A redirect DNS policy is a network configuration that

- enables devices to respond to DNS queries using a specific configuration and associated host table cache
- allows multiple DNS databases to be configured on the device, and
- lets the software select a DNS name server configuration for each query, either 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.

Key differences between dns-app-list and app-list

Considerations for using dns-app-list and app-list

Redirect DNS in service-side VPN

The dns-app-list and app-list constructs serve different purposes in SD-WAN traffic classification:

- **dns-app-list:** This construct is specifically designed to capture DNS traffic related to applications. It allows processing DNS requests and responses on an application-by-application basis, which is useful for scenarios like split DNS configurations. This approach is beneficial when DNS traffic must be handled differently for each application.
- **app-list:** This is used to match actual application traffic. It is suitable for identifying and processing traffic that is directly related to applications, rather than DNS queries associated with those applications.
- Use dns-app-list to differentiate DNS traffic for specific applications, especially in environments where DNS traffic needs to be managed separately from the application traffic itself.
- Use app-list for general application traffic classification where there is no need to separately identify or process DNS traffic.

Redirect DNS 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.

Benefits of redirect DNS

Describes the benefits of enabling DNS redirects within a service-side VPN to selectively forward DNS.

Describes the

- Supports DNS redirects within a service-side VPN.
- Allows you to redirect DNS requests based on a match condition.

Restrictions for redirect DNS in a service-side VPN

Describes key limitations and restrictions for Redirect DNS, including protocol support, NAT requirements, and unsupported traffic scenarios

Restrictions for redirect DNS in a service-side VPN

- You cannot send a redirect DNS request without NAT configuration when the request comes from the same VPN using the same port but from a different host.
- If you configure the DNS server IP address using NAT, you cannot change it 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 supports only IPv4 traffic and does not support 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

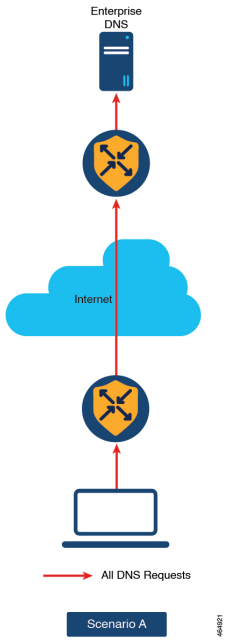
Lists service-side VPN use cases for redirect DNS, including unconditional forwarding of all DNS traffic and conditional redirection based on specific applications or domains.

This reference describes use cases for redirect DNS in a service-side VPN, including unconditional forwarding of all DNS traffic to an enterprise DNS server and conditional redirection based on specific applications or domains.

Unconditional redirect DNS

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

Figure 10: Unconditional redirect DNS

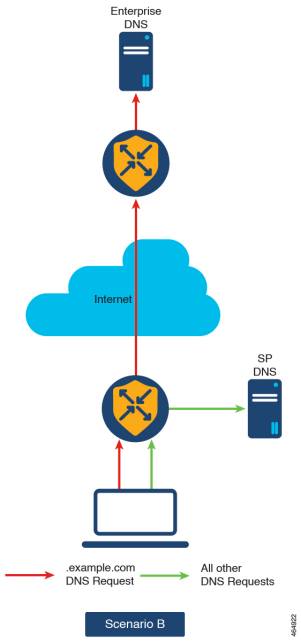


Conditional redirect DNS

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 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 11: Conditional redirect DNS



Configure redirect DNS in a service-side VPN

Describes how to configure Redirect DNS in Cisco Catalyst SD-WAN service VPNs using available methods such as policy groups, classic policies, or CLI-based configuration.

Use one of these methods to configure redirect DNS in a service-side VPN:

- Policy groups
- CLI commands
- Classic policies

Configure redirect DNS in service-side VPN using policy groups

Use these steps to configure redirect DNS in service-side VPN using policy groups.

Procedure

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Policy Groups** > **Application Priority and SLA**.
2. Create or edit an existing application priority and SLA policy, and select **Advanced Mode**.
3. Click **Add Traffic Policy**.

   Enter the name of the VPN and set the direction to Service.
4. Click **Add Rules**.
5. In the rule configuration page, set Protocol to **IPv4**.
6. To choose one or more **Match** conditions, click the fields and set the values as described.

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

Match Condition	Procedure
None (match all the packets)	Do not specify any match conditions.
Applications / Application Family List / Custom Applications	**a.** In the **Match** conditions menu, click **Applications/Application Family List**. **b.** From the drop-down list, choose the application family. **c.** To create an application list: **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.** From the **Select Application** drop-down list, choose the corresponding applications or application families. **5.** Click **Save**.
DNS Application List	Add an application list to enable split DNS: **a.** In the **Match** conditions menu, click **DNS Application List**. **b.** From the drop-down list, choose the application family.
DNS	Add an application list to process split DNS: **a.** In the **Match** conditions menu, click **DNS**. **b.** From the drop-down list, choose **Request** to process DNS requests for the DNS applications.
Destination Data Prefix	**a.** In the **Match** conditions menu, click **Destination Data Prefix**. **b.** To match a list of destination prefixes, from the **Data Prefix** drop-down list, choose a list. **c.** To match an individual destination prefix, enter the prefix in the **Destination: IP Prefix** field.
Destination Port	**a.** In the **Match** conditions menu, click **Destination Port**. **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 two numbers separated with a hyphen [-]).
DSCP	**a.** In the **Match** conditions menu, click **DSCP**. **b.** In the **DSCP** field, enter the DSCP value—a number from 0 through 63.
Packet Length	**a.** In the **Match** conditions menu, click **Packet Length**. **b.** In the **Packet Length** field, enter the length—a value from 0 through 65535.

Match Condition	Procedure
PLP	**a.** In the **Match** conditions menu, click **PLP** to set the **Packet Loss Priority**. **b.** From the **PLP** drop-down list, choose **Low** or **High**.
Protocol	**a.** In the **Match** conditions menu, click **Protocol**. **b.** In the **Protocol** field, enter the Internet Protocol number–a number from 0 through 255.
Source Data Prefix	**a.** In the **Match** conditions menu, click **Source Data Prefix**. **b.** To match a list of source prefixes, from the **Source Data Prefix List** drop-down list. choose a data prefix list. **c.** To match an individual source prefix, enter the prefix in the **Source** field.
Source Port	**a.** In the **Match** conditions menu, click **Source Port**. **b.** In the **Source** field, enter the port number. Specify a single port number, a list of port numbers (with numbers separated by a space), or a range of port numbers (with the two numbers separated with a hyphen [-]).

7. Under Actions, select **Accept**, enable **Redirect DNS**, and enter the DNS server IP address.
8. Click **Save**.

Configure redirect DNS in a service-side VPN using CLI commands

Describes how to configure Redirect DNS in a service-side VPN using CLI, including required policy lists, match-action sequences, and policy application on vSmart controllers.

This section provides the minimum policy components required to enable redirect DNS with a centralized data policy.

Use these steps to configure redirect DNS in a service-side VPN using CLI commands:

Procedure

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

   Example

   ```
   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.

Example

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

Refer to these lists in the **match** section of the data policy.

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

   Example

   ```
   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.

   Example

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

5. Create a series of match-action pair sequences.

   Example

   ```
   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. 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.

   Example

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

   For the *list-name* argument, specify the application list name.

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

   Example

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

8. Direct DNS requests to a particular DNS server or configure DNS response forwarding for inbound traffic. By default, the DNS servers configured in the VPN in which the policy is applied are used to process DNS lookups for the applications.

Example

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

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

For outbound traffic, configure the IP address of the DNS server. For inbound traffic, include the action to correctly forward the DNS response back to the service VPN.

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

Example

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

Configure redirect DNS in a service-side VPN using classic policies

Use these steps to configure Redirect DNS in a service-side VPN using classic policies.

Procedure

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**.

   In the **Name** and **Description**, enter a name and a description for the data policy.
5. Click **Sequence Type**.

   The **Add Data Policy** dialog box is displayed.

   a) 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.

      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.

   b) 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.
6. From the **Protocol** drop-down list, choose **IPv4** to apply the policy only to IPv4 address families.
7. To choose one or more **Match** conditions, click the fields and set the values as described.

   Not all match conditions are available for all policy sequence types.
8. To select the actions to take on matching data traffic, click the **Actions** menu.
9. To drop matching traffic, click **Drop**.

   The available policy actions are listed on the right side.

Configuration examples for redirect DNS

Provides CLI configuration examples for unconditional and conditional redirect DNS using centralized data policies in a service-side VPN.

This section provides examples of CLI configuration for redirect DNS, including unconditional and conditional scenarios, using centralized data policies in a service-side VPN.

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
```

Verify redirect DNS in a service-side VPN

escribes how to verify Redirect DNS configuration in a service-side VPN using the show sdwan policy from-vsmart command to confirm applied data policies and redirect actions.

This section explains how to verify the Redirect DNS configuration in a service-side VPN by using the **show sdwan policy from-vsmart** command and reviewing the output for applied data policies and redirect actions.

This 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
```

7 Device Access Policy

Topics:

- Feature history for device access policy
- Device access policies in Cisco SD-WAN Manager
- Configure device access policy using the CLI
- ACL statistics and counters examples
- ACL policy verification on an SNMP server
- ACL policy verification on SSH

Provides configuration guidelines for controlling device access based on security requirements.

Feature history for device access policy

This reference provides the feature history for device access policy on SNMP and SSH, including release information and feature descriptions.

This table describes the developments of this feature, by release.

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 policies in Cisco SD-WAN Manager

Explains how device access policies control traffic to local services on Cisco SD-WAN devices by filtering access from specific sources.

A device access policy is a security mechanism that

- controls IP traffic access to local services like SSH and SNMP on Cisco IOS XE Catalyst SD-WAN devices,
- protects the CPU from malicious device access traffic by applying filters to incoming connections, and
- defines rules that traffic must meet to pass through an interface before any other policies are applied.

Device access policy behavior

Starting from Cisco IOS XE Catalyst 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.

When both Access Control List (ACL) and Firewall (FW) policies are configured on Cisco IOS XE Catalyst SD-WAN devices, packet processing follows a defined order: ACLs are evaluated first at the interface level, and only packets permitted by the ACL proceed to the FW policy evaluation at the zone or global level. Traffic must be allowed by both ACL and FW

policies to be forwarded. This ensures ACLs have higher precedence in initial filtering, providing clear and predictable traffic control.

This behavior applies to device access policies configured via configuration groups and should be considered when defining ACL sequences and actions to ensure proper traffic flow and security enforcement.

Note

Device access policy is not applicable for VPN 512.

Configure device access policy using the CLI

Configure device access policies to control network access and enhance security by restricting connections based on specified criteria.

Configure device access policies to control and restrict access to network devices based on specified criteria such as IP addresses, ports, and protocols.

Procedure

1. Configure the SNMP access control list and community settings.

   Example

   ```
   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
   ```

2. Configure the SSH access control list and apply it to VTY lines.

   Example

   ```
   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.

The device access policies are now configured and will control access to the network devices based on the specified rules and criteria.

Configure IPv4 device access policy using configuration groups

Configure an IPv4 device access policy within a System profile using configuration groups in Cisco SD-WAN Manager.

Configure an IPv4 device access policy to control network access based on IPv4 addresses and sequences within a configuration group.

Use configuration groups to create reusable policy configurations that can be applied across multiple devices in your SD-WAN network.

Before you begin

On the **Configuration** > **Configuration Groups** page, choose **SD-WAN** as the solution type.

Follow these steps to configure an IPv4 device access policy using configuration groups:

Procedure

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Configuration Groups**.
2. Create and configure IPv4 Device Access Policy in a System profile.
   a) Configure the basic settings.

   Table 22: Basic Settings

Field	Description
Type	Choose a feature from the drop-down list.
Feature Name*	Enter a name for the feature. The name can be up to 128 characters and can contain only alphanumeric characters.
Description	Enter a description of the feature. The description can be up to 2048 characters and can contain only alphanumeric characters.

   b) Configure the ACL sequence settings.

   Table 23: ACL Sequence

Field	Description
ACL Sequence Name	Enter a name for the ACL Sequence.
Action Type	Choose one of the following actions for the ACL policy: • **Accept** • **Drop**

Field	Description
Default Action	The **Default Action** in the left pane is to drop the packets. Change the default action by clicking the ellipsis (**...**) icon.
Condition	• **Device Access Protocol (required)**: Choose a carrier from the drop-down list. For example, SNMP, SSH. • **Source Data Prefix**: Select an existing source data prefix or provide a source IP address. For example, 10.0.0.0/12. • **Source Port**: Enter the list of source ports when you have chosen SSH as the device access protocol. The range is 0 through 65535. • **Destination Data Prefix**: Select an existing destination data prefix or provide a destination IP address when you have chosen SSH as the device access protocol. For example, 10.0.0.0/12.

The IPv4 device access policy is configured within the System profile and can be deployed to devices through configuration groups.

What to do next

Also see Deploy a Configuration Group.

Configure IPv6 device access policy using configuration groups

Configure IPv6 device access policies through configuration groups in the SD-WAN Manager interface.

This task enables you to configure IPv6 device access policies using configuration groups, providing centralized management and deployment of security policies across SD-WAN devices.

IPv6 device access policies control which devices can access network resources over IPv6 connections. Configuration groups provide a template-based approach to deploy these policies consistently across multiple devices in your SD-WAN network.

Before you begin

On the **Configuration** > **Configuration Groups** page, choose **SD-WAN** as the solution type.

Follow these steps to configure IPv6 device access policy using configuration groups:

Procedure

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Configuration Groups**.
2. Create and configure IPv6 Device Access Policy in a System profile.

Table 24: IPv6 Device Access Policy

Field	Description
Add ACL Sequence	
ACL Sequence Name	Enter a name for the ACL Sequence.
Action Type	Choose one of the following actions for the ACL policy: • **Accept** • **Drop**
Default Action	The **Default Action** in the left pane is to drop the packets. Change the default action by clicking the ellipsis (**...**) icon.
Condition	• **Device Access Protocol (required)**: Choose a carrier from the drop-down list. For example, SNMP, SSH. • **Source Data Prefix**: Select an existing source data prefix or provide a source IP address. For example, 10.0.0.0/12. • **Source Port**: Enter the list of source ports when you have chosen SSH as the device access protocol. The range is 0 through 65535. • **Destination Data Prefix**: Select an existing destination data prefix or provide a destination IP address when you have chosen SSH as the device access protocol. For example, 10.0.0.0/12.

The IPv6 device access policy is configured within the configuration group and ready for deployment to SD-WAN devices.

What to do next

Also see Deploy a Configuration Group.

Configure device access policy using localized policy

Configure device access policy using localized policy to handle SNMP and SSH traffic directed towards the control plane.

Configure device access policy to handle SNMP and SSH traffic directed towards the control plane on s using destination ports based on the device access policy.

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.

Before you begin

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.

Follow these steps to configure device access policy using localized policy:

Procedure

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**
   >
   > From Cisco IOS XE Catalyst SD-WAN Release 17.9.1a and Cisco SD-WAN 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 SD-WAN 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	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.

Match Condition	Description
Destination Data Prefix	Enter the destination IP address. For example, 10.0.0.0/12.
VPN	Enter the VPN ID. Range: 1 to 65525, excluding 512. For details see the VRF range behavior change described here.

9. Click **Actions** to configure the action conditions.
10. Click **Accept** or **Drop** to accept or drop packets, respectively.
11. Click **Counter** if you want to count the packets that match the ACL policy criteria.

    By default, the counter is incremented when a packet matches all the conditions of the match criteria. From the **Counter Name** drop-down list, select the name of the counter, or click **Global** to choose from a list of already configured counters.
12. Click **Log** to enable logging of packets that match the ACL policy criteria.
13. Click **Save Match And Actions** to save all the conditions for the ACL policy.
14. Click **Save Device Access Control List Policy** to apply the selected match conditions to an action.
15. 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.

ACL statistics and counters examples

Provides examples for configuring and displaying ACL statistics and counters using YANG and CLI commands.

ACL statistics and counters help you monitor access control list usage and rule matches. You can configure and view this data through YANG models or CLI commands.

YANG configuration examples

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#
```

CLI display and clear commands

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}
```

ACL policy verification on an SNMP server

Provides information about verifying ACL policy on SNMP servers.

From Cisco IOS XE Catalyst SD-WAN Release 17.2.1r, 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;
            }
          }
        }
```

See the sample test log for SNMP-server ACL settings below:

```
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: R0/0: nesd: 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: R0/0: nesd: 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#
```

ACL policy verification on SSH

Provides verification methods for ACL policy implementation on SSH servers using Virtual Teletype lines.

From 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 acccess-list {
        tailf:cli-drop-node-name;
        tailf:cli-compact-syntax;
        tailf:cli-reset-container;
        tailf:cli-suppress-mode;
        tailf:cli-delete-when-empty;
        key "direction";
        leaf direction {
          type enumeration {
            enum "in";
            enum "out";
          }
        }
        leaf access-list {
          tailf:cli-drop-node-name;
          tailf:cli-prefix-key;
          type ios-types:exp-acl-type;
          mandatory true;
        }
        leaf vrf-also {
          description
            "Same access list is applied for all VRFs";
          type empty;
        }
      }
    }
```

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
```

8 Application Intelligence Engine Flow

Topics:

- SAIE flow overview
- SAIE flow configuration using Cisco SD-WAN Manager
- Configure SAIE using the CLI

Describes the Application Intelligence Engine (SAIE) flow and provides overview information about configuring the flow using the GUI or CLI interface.

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

SAIE flow overview

Describes the SD-WAN Application Intelligence Engine (SAIE) flow capabilities and configuration methods.

The SD-WAN Application Intelligence Engine (SAIE) flow is a packet inspection capability that

- looks into the packet past the basic header information to determine packet contents
- records packet information for statistical purposes or performs actions on the packet, and
- provides increased visibility into network traffic for understanding usage patterns and correlating network performance information.

SAIE flow benefits and configuration

> **Note**
>
> In Cisco SD-WAN 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 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)

SAIE flow configuration using Cisco SD-WAN Manager

Describes the sequential wizard screens that guide you through creating and editing SAIE flow policy components.

A SAIE flow configuration using Cisco SD-WAN Manager is a policy configuration process that

- guides you through creating and editing SD-WAN Application Intelligence Engine (SAIE) flow policy components,
- uses the Cisco SD-WAN Manager policy configuration wizard with sequential screens, and
- applies policies to sites and VPNs in the overlay network.

Configuration wizard screens

The wizard consists of the following sequential screens:

- 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 SAIE flow

Configure a centralized data policy for the SAIE flow by applying it to a list of sites in the overlay network.

Apply a centralized data policy for the SAIE flow to ensure it takes effect across specified sites in the overlay network.

A centralized data policy for the SAIE flow must be applied to a list of sites in the overlay network to ensure proper enforcement. The policy can be applied using either Cisco SD-WAN Manager or the CLI interface.

Procedure

1. Choose your preferred method to apply the centralized policy.
   - To apply a centralized policy in Cisco SD-WAN Manager, see *Configure Centralized Policy Using Cisco SD-WAN Manager*.
   - To apply a centralized policy using the CLI, continue with the next step.
2. Apply the data policy using the CLI command.

   Example

   ```
   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.

The centralized data policy for the SAIE flow is successfully applied to the specified sites and is actively enforcing policy rules on data traffic passing through the network.

Monitor running applications

Configure application visibility and monitor running applications on Cisco vEdge devices.

Enable the SAIE infrastructure to gain visibility into application traffic and monitor running applications on your network devices.

Application visibility is required to enable the Service Assurance and Insights Engine (SAIE) infrastructure on Cisco vEdge devices for monitoring network applications.

> **Note**
>
> In Cisco Cisco SD-WAN Release 20.7.x and earlier releases, the SAIE flow is called the deep packet inspection (DPI) flow.

Before you begin

Follow these steps to monitor running applications:

Procedure

1. Enable application visibility on the device.

   Example

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

2. Display information about the running applications using the appropriate show commands.

   Use the **show app DPI supported-applications**, **show app DPI applications**, and **show app DPI flows** commands on the device.

Application visibility is enabled on the device, and you can view information about supported applications, active applications, and application flows using the show commands.

View SAIE applications

View the list of all application-aware applications supported by the Cisco SD-WAN Manager software on the router.

This task allows you to view the list of all application-aware applications supported by the Cisco SD-WAN Manager software on the router.

Use this task to monitor and review SAIE applications running on your SD-WAN devices through the Cisco SD-WAN Manager interface.

Procedure

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.

You can now view the SAIE applications running on the selected device and the list of supported applications.

Action parameters for configuring SAIE flow

Lists the action parameters that you can configure when data traffic matches the conditions in the match portion of a centralized data policy.

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.

Configuration access

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.

Basic action parameters

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

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** *counter-name*	Name of a counter. Use the **show policy access-lists counters** command on the Cisco device.
Discard the packet. This is the default action.	Click **Drop**	**drop**	–

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

Parameters for accepted packets

Default action

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

Description	Cisco SD-WAN Manager	CLI Command	Value or Range
DSCP value.	Click **Accept**, then action **DSCP**.	**set DSCP** *value*	0 through 63

Description	Cisco SD-WAN Manager	CLI Command	Value or Range
Forwarding class.	Click **Accept**, then action **Forwarding Class**.	**set forwarding-class** *name*	Text
Next hop. Specify the IP address of the next-hop router for traffic that matches the match criteria. The next-hop router must be directly connected to the local router.	Click **Accept**, then action **Next Hop**.	**set next-hop** *IP-address*	IP address
Policer. Limit the bandwidth used by the traffic flow.	Click **Accept**, then action **Policer**.	**set policer** *name*	Name of a policer configured using the **policy policer** command
Traffic-engineering path. Specify a path for the traffic to follow. The traffic is placed in a traffic-engineering tunnel.	Click **Accept**, then action **Traffic Engineering**.	**set traffic-engineering path** *name*	Name of a traffic-engineering path configured with the **policy traffic-engineering path** command
VPN. Redirect matching traffic to a specific VPN. The redirected traffic is routed to the VPN based on the VPN's route table.	Click **Accept**, then action **VPN**.	**set VPN** *VPN-id*	0 through 65530

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 SAIE using the CLI

Configure a centralized data policy for the SD-WAN Application Intelligence Engine (SAIE) flow using CLI commands.

This task enables you to configure a centralized data policy for the SD-WAN Application Intelligence Engine (SAIE) flow to control application traffic behavior across overlay network sites.

The SAIE flow provides application intelligence and traffic control capabilities in SD-WAN deployments. Use this configuration when you need to implement centralized policies for application traffic management across multiple sites.

Note

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

Follow these steps to configure a centralized data policy for the SD-WAN Application Intelligence Engine (SAIE) flow:

Procedure

1. Create a list of overlay network sites to which the data policy is to be applied using the **apply-policy** command:

   Example

   ```
   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:

   Example

   ```
   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
   ```

   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.

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

   Example

   ```
   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:

   Example

   ```
   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:

Example

```
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:

   Example

   ```
   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:

   Example

   ```
   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:

   Example

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

9. Define additional match parameters for data packets:

   Example

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

10. Define actions to take when a match occurs:

    Example

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

11. For packets that are accepted, define the actions to take:

    Example

    ```
    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. Apply the data policy to one or more sites:

Example

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

What to do next

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`

9 Application-Aware Routing

Topics:

- Feature history for application-aware routing
- Application-aware routing
- SLA
- Best tunnel path
- Per-class application-aware routing
- Application probe classes
- Default DSCP values
- Configure BFD template using configuration groups
- Apply policies to sites and VPNs
- Activate an application-aware routing policy
- Configure application-aware routing policies
- Data plane tunnel performance monitoring

Describes Application-Aware Routing, a feature that dynamically selects the best network path for application traffic based on real-time performance metrics, ensuring optimal user experience and efficient use of network resources.

Feature history for application-aware routing

Lists the development milestones and release information for AAR, including feature description.

This table describes the developments of this feature, by release.

Table 25: Feature History

Feature	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. You can configure AAR for IPv6 traffic in Cisco vManage.
Application-Aware Routing for Hub and Spoke	Cisco vManage Release 20.9.1	This feature enables you to configure an application-aware routing policy on a hub device and apply the policy to packets that the hub device receives from a spoke device and sends to another spoke device. Applying the application-aware routing policy helps you realize the required SLA for spoke-hub-spoke traffic.
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.
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. You can configure best tunnel path to pick the best path while configuring SLA class.

Feature	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).
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.
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.
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.

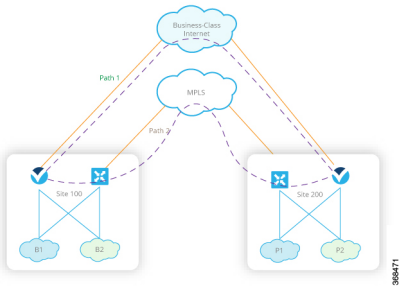
Application-aware routing

Describes how application-aware routing optimizes traffic paths and performance on IOS XE Catalyst SD-WAN devices using SLA-based path selection.

An application-aware routing (AAR) is a network optimization feature that

- tracks network and path characteristics of the data plane tunnels between Cisco IOS XE Catalyst SD-WAN devices
- uses collected information such as packet loss, latency, and jitter to compute optimal paths for data traffic, and
- considers factors in path selection beyond standard routing protocols, offering advantages such as optimized traffic flow and reduced network costs.

This image shows how SD-WAN builds multiple tunnels over Internet and MPLS and dynamically steers application traffic over the best-performing path.



Benefits of AAR

- Direct application traffic over WAN links that meet the required SLA for packet loss, latency, and jitter during normal network operation.
- Detect brownouts or soft failures in real time and automatically redirect traffic to the best available path to minimize performance degradation.
- Re-adjust traffic paths automatically once network conditions recover.
- Reduce network costs by load-balancing traffic more efficiently across available links.
- Improve application performance without upgrading the WAN.

TLOCs support on Cisco IOS XE Catalyst SD-WAN devices

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.

Components of application-aware routing

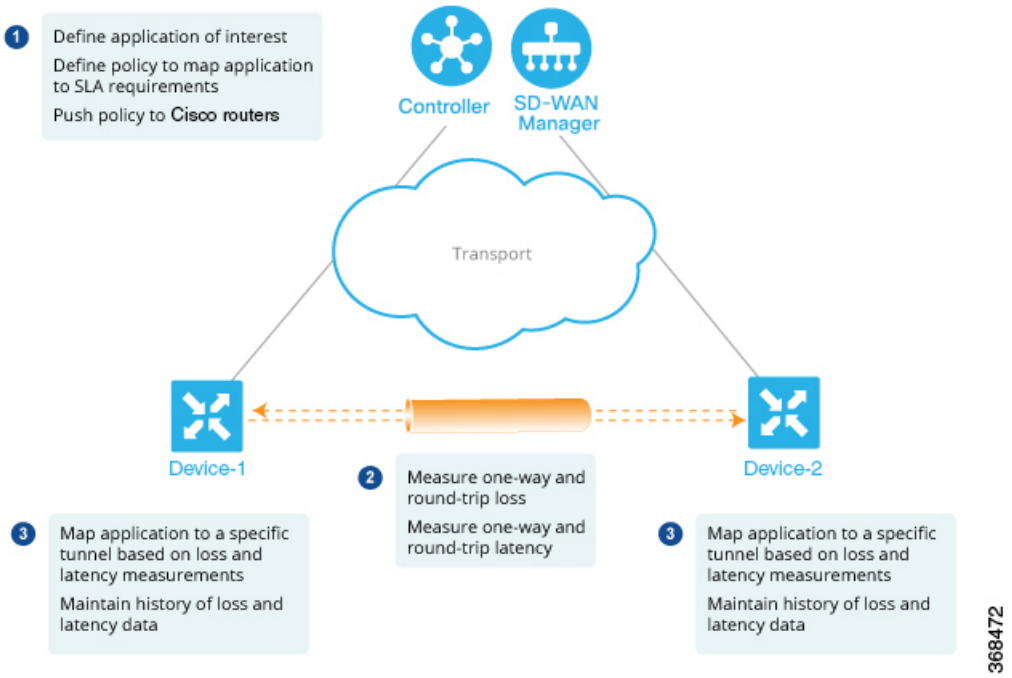
Lists the elements that comprise the application-aware routing solution and describes their roles in identifying, monitoring, and mapping application traffic.

The Cisco IOS XE Catalyst SD-WAN Application-Aware Routing solution consists of three elements: identification, monitoring and measuring, and mapping application traffic to a specific transport tunnel.

The application-aware routing solution includes the following 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.

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.

Fundamentals of application-aware routing policy

Provides general procedures for application-aware routing configuration.

Scope and purpose of 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.

It maps applications to the data plane tunnel performance characteristics required to transmit their 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.

Policy type and deployment model

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.

Like any policy, an application-aware routing policy contains a series of numbered (ordered) match-action sequences that the system evaluates from the lowest sequence number to the highest. When a data packet matches a condition, the system applies an SLA action to determine the data plane tunnel used to transmit the packet.

Default behavior and policy nature

If a packet does not match any parameters in 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. This is because application-aware routing policy accepts nonmatching traffic by default, it is classified as a positive policy. Other types of policies in the Cisco IOS XE Catalyst SD-WAN software are negative policies, because they drop non-matching traffic by default.

IPv6 support enhancement

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.

Default action of application-aware routing policy

Provides the default behavior and configuration options for application-aware routing policy when no match conditions are met.

This section describes how the default action in an application-aware routing policy determines packet handling when no match conditions are met.

The policy's default action defines how it handles packets that match none of the match conditions. For application-aware routing policy, if you do not configure a default action, the policy accepts all data packets and transmits them based on normal routing decisions, without considering 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.

> **Note**
>
> When a loopback interface with a public IP address is configured as a TLOC and bound to a NAT-enabled physical WAN interface, DIA forwarding works. However, strict color-based exit selection using centralized data policy local-tloc or local-tloc-list is not supported if the bound physical WAN interface is not itself configured as a TLOC.

Application-aware routing for multicast protocols

Provides information about application-aware routing support for multicast protocols, including feature history and multicast traffic classification on IOS SD-WAN devices.

This reference describes application-aware routing support for multicast protocols, and multicast traffic classification on Cisco IOS XE Catalyst SD-WAN devices.

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.

Multicast traffic classification

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

Outlines the restrictions for multicast protocols. Review these restrictions to ensure proper protocol configuration.

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

SLA

Explains SLA classes, their parameters, configuration limits, and updated default threshold behavior in Cisco Catalyst SD-WAN.

A service-level agreement (SLA) class is a policy construct that

- Defines maximum jitter, latency, packet loss, or their combination for data plane tunnels,
- Drives actions in application-aware routing, and
- Applies to tunnels formed by local and remote TLOC pairs.

You can define only four unique SLA classes in an application-aware route policy and configure up to eight SLA classes on devices in later releases. Earlier releases supports only four SLA configuration.

SLA components

SLA comprises of the following parameters:

Table 26: SLA Components

Description	Command	Value or range
Maximum acceptable packet jitter on the data plane tunnel	**jitter** *milliseconds*	1–1000 milliseconds
Maximum acceptable packet latency on the data plane tunnel.	**latency** *milliseconds*	1–1000 milliseconds
Maximum acceptable packet loss on the data plane tunnel.	**loss** *percentage*	1–100 percent

Threshold values for SLA class lists

Starting with Cisco IOS XE Catalyst SD-WAN Release 17.15.1a, the threshold values for SLA class lists are adjusted, as detailed in the table below. These changes may affect your network traffic behavior upon upgrading.

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

From Cisco Catalyst SD-WAN Manager Release 20.12.x, the threshold values for default SLA classes are not configurable.

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 the data flow matches the app-route and data policy sequences, the following expected behaviors occur:

- If the preferred colors in application-aware routing meet the SLA requirements and overlap with the 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 some colors overlap with the data policy and meet the SLA in application-aware routing, those 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 for forwarding. If the data policy specifies preferred colors, those colors are selected; otherwise, load balancing occurs across all colors in the data policy.

SLA classes

Lists the supported SLA classes and platform-specific limits for Cisco SD-WAN devices.

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

The SLA enhancement supports multitenancy by allowing different SLA classes to be pushed to 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.

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.

Maximum SLA classes supported on Cisco IOS XE Catalyst SD-WAN devices

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.

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	6	15
ASR 1002 X - 16GB	6	15
ASR 1002 HX - 16GB	6	15
ASR 1001 X - 16GB	6	15
ISR 4451 X	6	7
ISR 4431	6	7
Catalyst 8300 Edge Platforms	NA	7
Catalyst 8500 Edge platforms - 16GB	NA	15
Any other Cisco IOS XE Catalyst SD-WAN devices	6	6

Classification of tunnels into SLA classes

Lists the three main parts of the process for classifying tunnels into SLA classes for application-aware routing.

Classifying tunnels into one or more SLA classes for application-aware routing consists of three main 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

Provides information about how application-aware routing uses BFD Hello packets to measure loss, latency, and jitter on overlay network tunnels.

Application-aware routing uses BFD Hello packets to measure loss, latency, and jitter on overlay network tunnels, ensuring tunnel liveness.

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 uses a color to identify 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. You can configure this interval with the **bfd color interval** command. You can configure the BFD Hello packet interval per tunnel.

> **Note**
>
> We recommend using the BFD hello interval timer as 1 second or higher and BFD multiplier as 7 seconds.

Calculate average loss, latency, and jitter

Provides information about how application-aware routing calculates average loss, latency, and jitter for each tunnel using BFD polling intervals.

Application-aware routing uses BFD polling to collect packet latency, loss, jitter, and other statistics for each tunnel.

Tunnel polling and metrics collection

BFD periodically polls all tunnels on Cisco IOS XE Catalyst SD-WAN devices to collect packet latency, loss, jitter, and other statistics for 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."

Poll interval behavior and configuration

By default, the poll interval is 10 minutes. With the default BFD Hello packet interval set to 1 second, application-aware routing uses information from about 600 BFD Hello packets in one poll interval to calculate tunnel loss, latency, and jitter. The poll interval is user-configurable using the **bfd app-route poll-interval** command. Application-aware routing configures the poll interval per Cisco IOS XE Catalyst SD-WAN device, and it applies to all tunnels originating on the device.

Reducing the poll interval without reducing the BFD Hello packet interval may affect the quality of the loss, latency, and jitter calculations. For example, if you set the poll interval to 10 seconds and keep the BFD Hello packet interval at 1 second, the system uses only 10 Hello packets to calculate tunnel loss, latency, and jitter.

Sliding window of poll intervals

Application-aware routing preserves the loss, latency, and jitter information from each poll interval for six poll intervals. At the seventh poll interval, the system discards the earliest polling interval and stores 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.

Configure SLA

Lists the available methods to configure SLA in Cisco SD-WAN for AAR.

Use one of these methods to configure SLA in Cisco SD-WAN for AAR.

- Policy groups
- Classic policies

Configure SLA class using policy groups

Use these steps to configure SLA Class using policy groups.

Procedure

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Policy Groups** > **Application Priority & SLA**.
2. Create or edit an existing application priority and SLA policy and select **Advanced Mode**.
3. From the left pane, click **SLA Class**, and then click **Add SLA Class**.
4. Define the SLA class parameters.
   - In the **Loss** field, enter the maximum packet loss on the connection, a value from 0 through 100 percent.
   - In the **Latency** field, enter the maximum packet latency on the connection, a value from 1 through 1,000 milliseconds.
   - In the **Jitter** field, enter the maximum jitter on the connection, a value from 1 through 1,000 milliseconds.
5. 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.

Configure SLA class using classic policies

Configure an SLA class using classic policies to define loss, latency, and jitter parameters for application-aware routing. This enables precise traffic management and ensures optimal network performance.

Use these steps to configure an SLA class using classic policies to set thresholds for loss, latency, and jitter, enabling application-aware routing and optimal path selection in your SD-WAN deployment.

Procedure

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

   Centralized Policy is 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 the 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

Lists the available methods to configure traffic rules in Cisco SD-WAN for AAR.

Use one of these methods to configure traffic rules in Cisco SD-WAN for AAR.

- Policy groups
- Classic policies

Configure traffic rules using classic policies

Configure application-aware routing policies to define traffic rules using classic policies.

Use this task to define or modify traffic rules for application-aware routing in Cisco SD-WAN environments.

Procedure

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.

   a) 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.

   b) In the right pane, click **Sequence Rule**.

      The **Match/Actions** dialog box opens, and **Match** is selected by default.

      The available policy match conditions appear below the dialog box.

4. In the **Protocol** drop-down list, choose one of the following options:
   - **IPv4**
   - **IPv6**
   - **Both**

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

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

Table 29: 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 *Cisco Catalyst SD-WAN Cloud OnRamp Configuration Guide, Cisco IOS XE Release 17.x*.

> **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

You can use **Destination Region** in a Cisco Catalyst SD-WAN network using Cisco Catalyst SD-WAN Multi-Region Fabric.

Choose one of the following options from the drop-down list:

- **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.

For more information on how to configure Multi-Region Fabric, see the *Cisco Catalyst SD-WAN Multi-Region Fabric (also Hierarchical SD-WAN) Configuration Guide.*

Destination 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 [-]).
Traffic To	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.
DNS (to enable split DNS)	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.
DSCP	Type the DSCP value, a number from 0 through 63.
PLP	Choose **Low** or **High**. To set the PLP to **High**, apply a policer that includes the **exceed remark** option.
Protocol	Type the internet protocol number, a number from 0 through 255.
ICMP Message	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. 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. **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. When **Protocol** is selected as **Both**, the **ICMP Message or ICMPv6 Message** field displays.
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 [-]).

6. Cliick **Actions** to select actions for the matched data traffic. Set the values as described in the following table:

Table 30: 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	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. Click **SLA Class List**. In the **SLA Class** drop-down list, choose one or more SLA classes. Optionally, when the **Preferred Color** is not selected, you can choose the prefererred 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. 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. 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. 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. 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. Click **Fallback to best path** to select the best available tunnel to avoid a packet drop. **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. 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: `SLA Class selected, does not have Fallback Best Tunnel enabled. Please change the SLA class or change to Strict/Drop.` Click **Load Balance** to load balance traffic across all the tunnels.
Cloud SLA	**Cloud SLA** enables traffic to use the best path selection with Cisco Catalyst SD-WAN Cloud OnRamp for SaaS. Click **Cloud SLA**.

7. Click **Save Match and Actions**.

   Create additional sequence rules as desired. Drag and drop to re-arrange them.

8. Click **Save Application Aware Routing Policy**.
9. Click **Next** to move to **Apply Policies to Sites and VPNs** in the wizard.

Configure traffic class using policy groups

Use these steps to configure traffic Class using policy groups.

Procedure

---

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Policy Groups** > **Application Priority & SLA**.
2. Create or edit an existing application priority and SLA policy and select **Advanced Mode**.
3. Click **Add Traffic Policy**.
4. Enter the name of the VPN and set the direction to **Service**.
5. Click **Add Rules**.
6. In the **Protocol** drop-down list, choose one of the following options:
   - IPv4
   - IPv6
   - Both

   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 31: 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 *Cisco Catalyst SD-WAN Cloud OnRamp Configuration Guide, Cisco IOS XE Release 17.x*.

> **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

You can use **Destination Region** in a Cisco Catalyst SD-WAN network using Cisco Catalyst SD-WAN Multi-Region Fabric.

Choose one of the following options from the drop-down list:

- **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.

For more information on how to configure Multi-Region Fabric, see the *Cisco Catalyst SD-WAN Multi-Region Fabric (also Hierarchical SD-WAN) Configuration Guide.*

Destination 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 [-]).
Traffic To	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.
DNS (to enable split DNS)	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.
DSCP	Type the DSCP value, a number from 0 through 63.
PLP	Choose **Low** or **High**. To set the PLP to **High**, apply a policer that includes the **exceed remark** option.
Protocol	Type the internet protocol number, a number from 0 through 255.
ICMP Message	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. 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. **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. When **Protocol** is selected as **Both**, the **ICMP Message or ICMPv6 Message** field displays.
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 [-]).

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

Table 32: 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	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. Click **SLA Class List**. In the **SLA Class** drop-down list, choose one or more SLA classes. Optionally, when the **Preferred Color** is not selected, you can choose the prefererred 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. 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. 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. 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. 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. Click **Fallback to best path** to select the best available tunnel to avoid a packet drop. **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. 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: `SLA Class selected, does not have Fallback Best Tunnel enabled. Please change the SLA class or change to Strict/Drop.` Click **Load Balance** to load balance traffic across all the tunnels.
Cloud SLA	**Cloud SLA** enables traffic to use the best path selection with Cisco Catalyst SD-WAN Cloud OnRamp for SaaS. Click **Cloud SLA**.

9. Click **Save Match and Actions**.

   Create additional sequence rules as desired. Drag and drop to re-arrange them.

10. Click **Save Application Aware Routing Policy**.
11. Click **Next** to move to **Apply Policies to Sites and VPNs** in the wizard.

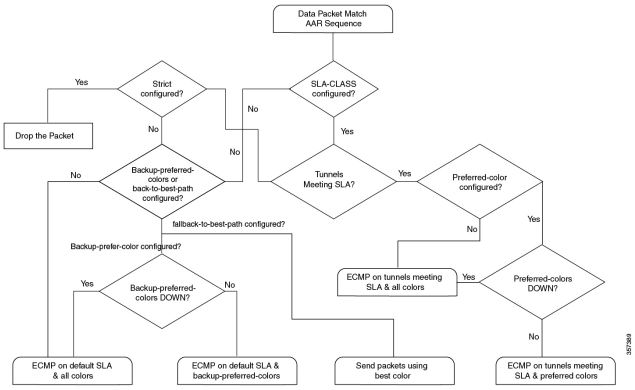
Best tunnel path

Lists policy actions for configuring the best application-aware routing tunnel selection when a SLA is not met.

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

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

Figure 12: Flow Chart for Application-Aware Routing Tunnel Selection



Variance configuration for best tunnel path

Provides information about variance configuration for best tunnel path selection in Cisco SD-WAN Manager and its effect on SLA class requirements.

This section explains how variance configuration affects best tunnel path selection in Cisco SD-WAN Manager when using best of worst (BOW) logic for SLA class requirements.

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

> **Note**
>
> Since processing cost scales with network complexity and other factors, BOW is not recommended for high-scale implementation unless the potential need for this forwarding has been validated against the performance cost.

BOW employs these scenarios based on the availability of SLA-compliant paths:

- When an SLA-compliant path exists, that tunnel is used for forwarding. BOW does not choose the traffic path, but its configured state is still evaluated during the polling cycle.
- If one tunnel fails but another remains compliant, AAR continues to use the compliant tunnel. BOW is not needed for forwarding, but measurement-state processing for the failed tunnel can still cause TLOC/state-reconciliation work.

- If all tunnels fail, BOW ranks the available tunnels and selects the least-degraded candidate when enabled. If BOW is disabled, behavior follows the strict/non-strict policy semantics documented for that policy.

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

Based on 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 during 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 variance configuration for best tunnel path

Provides examples of how configuration variance affects best tunnel path selection for latency, jitter, and loss in SD-WAN policies.

This section explains how configuring variance for latency, jitter, or loss impacts the selection of the best tunnel path in SD-WAN policies.

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
```

- When latency variance is not configured, tunnel T1 is selected as the best tunnel.
- When a latency variance of 10 ms is configured, tunnels T1, T2, and T3 all qualify as best tunnels.

- The rangeif from 110 ms to 120 ms.
- This range is calculated as best latency (110 ms) + latency variance (10 ms).
- Use the following formula to determine eligible tunnels: (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
```

- When jitter variance is not configured, tunnel T3 is selected as the best tunnel.
- When a jitter variance of 10 ms is configured, tunnels T1 and T3 qualify as the best tunnels.
- The range is from 152 ms to 162 ms.
- This range is calculated as best jitter (152 ms) + jitter variance (10 ms).
- Use the following formula to determine eligible tunnels: (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
```

1. When loss variance is not configured, tunnel T1 is selected as the best tunnel.
2. When a loss variance of 1% is configured, tunnels T1 and T2 qualify as the best tunnels.
3. The range is from 2% to 3%.
4. This range is calculated as best loss (2%) + loss variance (1%).
5. Use the following formula to determine eligible tunnels: (best_loss, best_loss + loss_variance)

Per-class application-aware routing

Explains how per-class application-aware routing uses SLA metrics and policies to direct application traffic based on loss, latency, and jitter measurements for accurate network performance.

A per-class application-aware routing policy is a network routing mechanism that

- uses SLA metrics such as loss, latency, and jitter measured between TLOCs
- applies policies that constrain paths for forwarding applications based on SLA classes, and

- requires active probing or passive monitoring to measure metrics on all paths to the traffic destination.

SLA metrics overview

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

However, SLA metrics derived from these high-priority packets do not reflect the priority that actual data traffic receives as it flows through the edge device. Depending on the application class, data traffic can carry different DSCP values in the network. Therefore, the network requires a more accurate representation of loss, latency, and jitter for real traffic profiles so it can direct different traffic types to the appropriate tunnels.

Application-aware routing and SLA policies

Application-aware routing uses policies to constrain the paths available for forwarding applications. These constraints are typically defined as SLA classes containing loss, latency, and jitter requirements that the network must meet.

To enforce these policies, the system measures SLA metrics across all paths to the traffic destination. It performs these measurements using either active probing or passive monitoring.

Active probing methods

Active probing generates synthetic traffic and injects it alongside real traffic. The network expects both probe traffic and real traffic to follow the same forwarding behavior.

Examples of active probing mechanisms include:

- BFD probing
- ICMP probes
- Periodic HTTP requests
- IP SLA measurements

The Cisco Catalyst SD-WAN solution uses BFD-based probes for active measurements.

Passive monitoring methods

Passive monitoring relies on the Cisco Catalyst SD-WAN Application Intelligence Engine (SAIE) to observe real traffic flows and measure loss, latency, and jitter directly from live traffic. For example, the system monitors RTP and TCP traffic to gather these metrics.

In Cisco vManage Release 20.7.1 earlier releases, the SAIE flow was referred to as the deep packet inspection (DPI) flow.

Application probe classes

Explains how application probe classes define color, and DSCP mapping for applications in a Cisco SD-WAN network site.

A application probe class is a policy component that

- defines forwarding class, color, and DSCP markings for application-aware routing probes,
- determines the QoS queue for BFD echo requests at the egress tunnel port, and
- controls DSCP usage and probe behavior when combined with SLA classes.

App-Probe class behavior and scope

Color and DSCP mapping are local to each Cisco SD-WAN network site, although some colors remain consistent across sites. BFD packet-loss priority is fixed to low.

When BFD packets use an SLA class, they use the same DSCP value. When BFD packets use both an SLA class and an app-probe-class, probes are sent separately for each SLA app-probe-class in a round-robin manner.

When the application route policy is applied at a site, only the colors relevant to the site are used. The maximum number of SLA class supported is platform dependent. For example, in ASR1000 series, upto 15 SLA classes are configurable. The maximum number of app-probe-class is always six across all the platforms.

Configure an application probe class

Lists all available methods to configure an application probe class.

Use one of these methods to configure an application probe class

- Policy groups
- CLI commands
- Classic policies

Configure app probe class using policy groups

Use these steps to configure app probe class using policy groups.

Procedure

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Policy Groups** > **Objects and Profiles**.
2. From the left pane, click **App Probe Class,** and then click **Add app probe class**.
3. Enter the probe class name in the **Probe Class Name** field.
4. Choose the required class map from the **Class Map** drop-down list.
5. Choose the appropriate color from the **Color** drop-down list and enter the **DSCP** value.
6. Click **+** icon, to add more entries as required.
7. Click **Save**.

Application probe class configuration using CLI

Provides steps to configure application probe class and map it with SLA class using CLI commands.

Configure the application probe class and map it with the SLA class using CLI commands.

Use this example to configure an application probe class for real-time video and map it to an SLA class.

```
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 DSCP value using the BFD template:

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

Configure an application probe class using classic policies

Configure an application probe class using classic policies. This task guides you through creating and assigning probe classes and forwarding classes in the centralized policy workflow.

Use these steps to configure an application probe class, define its probe classes, and associate them with forwarding classes for application-aware routing.

Procedure

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, 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

Sescribes how to create an SLA class list in Cisco SD-WAN Manager by defining loss, latency, jitter, and optional app probe settings.

Use these steps to create a new SLA class and associate it with an app probe class to enable performance monitoring based on loss, latency, and jitter criteria.

Before you begin

Procedure

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 the SLA class list.
4. Enter the required values for **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.

Default DSCP values

Lists the default DSCP value used in DSCP control traffic and provides guidance on configuration options and considerations for network service providers.

The DSCP control traffic uses the default DSCP value of 48. However, you can change the default value and configure it on the edge devices. All 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 default DSCP on a Cisco BFD template

Configure the default DSCP value on a Cisco BFD template to ensure proper marking of BFD packets.

Use these steps to configure the default DSCP value on a Cisco BFD template, ensuring that BFD packets are marked appropriately for traffic prioritization and quality of service.

Procedure

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

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

2. Click **Add Template** and select a device from the device list in the left pane.
3. In the right pane, select the BFD template listed under Basic Information.

   Enter **Template Name** and **Description** in the respective fields.

4. In the **Basic Configuration** pane, enter **Multiplier** and **Poll Interval (milliseconds)**.
5. In the **Default DSCP value for BFD Packets** field, enter the required device specific value or choose the default value for DSCP.
6. (Optional) In the **Color** pane, choose the required color from the drop-down list.
7. Enter the required **Hello Interval (milliseconds)** and **Multiplier**.
8. Choose the **Path MTU Discovery** value.
9. Enter the **BFD Default DSCP value for tloc color**.

10. Click **Add**.

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

Configure BFD template using configuration groups

Configure BFD template using configuration groups

Use these steps to configure BFD template using configuration groups.

Procedure

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Configuration Groups**.
2. Create and configure a BFD feature in a System profile.
3. Configure **basic settings**.

Table 33: Basic Configuration

Field	Description
Poll Interval(In Millisecond)	Specify how often BFD polls all data plane tunnels on a router to collect packet latency, loss, and other statistics used by application-aware routing. Range: 1 through 4,294,967,296 (2^{32} – 1) milliseconds Default: 600,000 milliseconds (10 minutes)
Multiplier	Specify the value by which to multiply the poll interval, to set how often application-aware routing acts on the data plane tunnel statistics to figure out the loss and latency and to calculate new tunnels if the loss and latency times do not meet the configured SLAs. Range: 1 through 6 Default: 6
DSCP Values for BFD Packets(decimal)	Specify the Differentiated Services Code Point (DSCP) value of the BFD packets that is used in the DSCP control traffic. Range: 0-63 Default: 48

Apply policies to sites and VPNs

Describes the available methods to apply policies to sites and VPNs.

Use one of these methods to apply policies to sites and VPNs.

- Policy groups: Add the application priority & SLA policy to a policy group, associate the required devices, and activate the policy group. For more information, refer to Deploy Policy Group Workflow in *Policy Groups Configuration Guide*.
- CLI commands

- Classic policies

Restrictions for policy application

Site ID requirement

- For all app-route-policy policies applied using **apply-policy** commands, site IDs across all site lists must be unique.
- Site lists must not contain overlapping site IDs.

  Example of overlap:

  - site-list 1: site-id 1-100
  - site-list 2: site-id 70-130
  - Sites 70-100 exist in both lists. If these two site lists are applied to different app-route-policy policies, committing the configuration on the target will fail.

- The same non-overlapping site ID requirement applies to:
  - 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 an **apply-policy** containing a cflowd-template command)

Allowed overlaps between different policy types

- Site lists can contain overlapping site IDs only when they are applied to different policy types.
- Example:
  - site-list 1 (1-100) applied to a control policy
  - site-list 2 (70-130) applied to a data policy.
- This configuration is allowed.

Apply policies to sites and VPNs using classic policies

Descibes how to configure policies to sites and VPNs in the overlay network using the policy configuration wizard and to associate policy blocks with VPNs and sites.

Use these steps to apply policy blocks to sites and VPNs in the overlay network to enforce application-aware routing and centralized policies.

Before you begin

Ensure you have created the required policy blocks in the previous steps of the policy configuration wizard.

Have lists of sites and VPNs ready for association with policy blocks.

Procedure

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.

    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 SD-WANCisco IOS XE Catalyst SD-WAN devices at the specified sites.

Apply policies to sites and VPNs using CLI commands

To apply an application-aware route policy to a list of sites in the overlay network, use this CLI command:

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

Verify applied policy

Describes the commands used to verify and monitor policies configured on the controller and pushed to routers.

Use these commands to verify and monitor policies configured on the controller and pushed to routers.

- 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.

Activate an application-aware routing policy

Describes how to activate an application-aware routing policy in Cisco SD-WAN Manager.

Use this task to activate an application-aware routing policy to control how application traffic is routed across the network based on defined criteria.

Procedure

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

   By default, the **Centralized Policy** page is displayed.

2. For the desired policy, click **...** and select **Activate**.

   The Activate Policy window opens and displays the IP addresses of the reachable Cisco SD-WAN Controllers where the policy will 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.

Configure application-aware routing policies

Describes the available methods to configure application-aware routing policies.

Use one of these steps to configure application-aware routing policies:

- CLI commands
- Classic policies

Configure application-aware routing using CLI commands

Configure application-aware routing using CLI commands to define policies, match criteria, and actions for application traffic across overlay networks.

Use these steps to control how application traffic is routed across overlay networks, ensuring that traffic meets specific SLA criteria such as latency, jitter, and loss.

Procedure

1. Create lists of overlay network sites to which the policy will be applied.

   Example

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

   a. Include one **site-id** command for each site.

   b. For contiguous site IDs, specify a range using a dash (-).

   c. Create additional site lists as needed.

2. Define SLA classes with traffic characteristics to apply to matching application data traffic.

   Example

   ```
   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-tunnelcriterialatencylossjitter
   ```

3. Create Lists of Applications, IP Prefixes, and VPNs.

Example

```
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.

   Example

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

5. Create match-action sequences: Define one or more numbered sequences that match data traffic and apply SLA classes.
   a) Create a sequence.

      Example

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

   b) Define match parameters for data packets.

      Example

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

   c) Define the action to take if a match occurs.

      Example

      ```
      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>
      ```

6. Configure action options for SLA classes:
   a) Configure basic SLA Class

      ```
      sla-class sla-class-name
      ```

      Forwards matching traffic as long as one tunnel meets the SLA.

      - If one tunnel matches, traffic goes through it.
      - If multiple tunnels match, traffic is distributed among them.
      - If no tunnel matches, traffic uses any available tunnel.

b) Configure preferred tunnel color.

```
sla-class sla-class-name
          preferred-color color
```

Traffic prefers a tunnel of the specified color.

- 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** *color*–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.

c) Count matching packets/bytes:

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

d) Place a sampled set of packets into syslog:

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

Match-action pairs are evaluated in numerical order. If no match occurs, the default action applies. For application-aware routing, the default is to accept nonmatching traffic. To apply SLA parameters to nonmatching packets:

7. If a packet does not match any of the conditions in one of the sequences, configure the default action.

Example

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

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.

8. Apply the policy to a site list.

Example

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

Configure application-aware routing policies using classic policies

Provides the steps to configure application-aware routing policies using Cisco SD-WAN Manager and describes the policy configuration wizard workflow.

This section describes how to configure application-aware routing policies using classic policies, including the sequence of steps and requirements for policy activation.

To configure application-aware routing policy, use the Cisco SD-WAN Manager policy configuration wizard. For Centralized Policy configuration details, refer to Configure Centralized Policies.

The wizard consists of four sequential windows that guide you through the process of creating and editing policy components:

Procedure

1. Create applications or groups of interest.

   Create lists that group related items and reference them in the match or action components of a policy. For configuration details, see Configure Groups of Interest.

2. Configure traffic rules.

   Define the match and action conditions of the policy.

3. Apply policies to sites and VPN.

   Associate the created policy blocks with sites and VPNs in the overlay network.

   In the first three policy configuration wizard windows, you create the policy components or blocks. In the final window, you apply these policy blocks to sites and VPNs

4. Activate the policy, to make an application-aware routing policy effective.

Configure an application-aware routing policy using CLI

Configure application-aware routing policy to direct ICMP traffic to links with latency of 50 milliseconds or less, and apply the policy to specific sites and VPNs. Verify and activate the configuration to ensure correct operation.

This task enables you to configure application-aware routing policy on a Cisco Catalyst SD-WAN Controller. The policy directs ICMP traffic to links with latency of 50 milliseconds or less.

You configure application-aware routing policy on a Cisco Catalyst SD-WAN Controller.

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.

Procedure

1. Define the SLA parameters to apply to matching ICMP traffic.

   Example

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

   This example directs ICMP traffic to links with latency of 50 milliseconds or less.

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

Example

```
vSmart(config-sla-class-test_sla_class)# exit
vSmart(config-sla-class-test_sla_class)# lists vpn-list vpn_1_list vpn 1
vSmart(config-vpn-list-vpn_1_list)# exit
vSmart(config-lists)# site-list site_500 site-id 500
vSmart(config-site-list-site_500)#
```

3. Configure the application-aware routing policy and specify the protocol numbers for ICMP, TCP, and UDP.

   Example

   ```
   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)#
   ```

   In this example, protocol 1 is ICMP, protocol 6 is TCP, and protocol 17 is UDP.

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

   Example

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

5. Display the configuration changes.

   Example

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

6. Validate that the configuration contains no errors.

   Example

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

7. Activate the configuration.

Example

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

8. Exit from configuration mode.

Example

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

Complete configuration example:

```
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
 !
!
```

Multicast protocol example:

```
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
 !
!
```

Remote color preference example:

```
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
 !
!
```

Data plane tunnel performance monitoring

Provides details about how BFD monitors data plane tunnel performance and how application-aware routing uses BFD statistics to determine tunnel transmission quality.

Monitor data plane tunnel health and performance using BFD statistics for application-aware routing decisions.

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.

Sliding window operation and statistics

- For each sliding window time period, application-aware routing receives 600 BFD Hello packets (BFD Hello interval × polling interval: 1 second × 600 seconds = 600 Hello packets). These packets measure packet loss and latency on the data plane tunnels.
- Application-aware routing retains the statistics for 1 hour (polling interval × multiplier: 10 minutes × 6 = 60 minutes).
- Application-aware routing places the statistics in six separate buckets, indexed 0 through 5. Bucket 0 contains the latest statistics, and bucket 5 contains the oldest. Every 10 minutes, the system places the newest statistics in bucket 0, discards the statistics in bucket 5, and moves the remaining statistics to the next bucket.
- Every 60 minutes (every hour), application-aware routing acts on the loss and latency statistics. It calculates the mean loss and latency across all buckets in all sliding windows and compares the result with the specified SLAs for the tunnel. If the calculated value satisfies the SLA, application-aware routing takes no action. 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 because the multiplier is 6. Although application-aware routing always retains six buckets of data, the multiplier determines how many buckets it uses to calculate the loss and latency. For example, if the multiplier is 3, it uses buckets 0, 1, and 2.

Adjusting parameters for faster detection

These default values take action only every hour, so they work well for a stable network. To capture network failures more quickly and allow application-aware routing to calculate new tunnels more often, adjust the values of these three parameters.

For example, if you change the polling interval to 1 minute (60,000 milliseconds), application-aware routing reviews tunnel performance characteristics every minute, but it performs loss and latency calculations based on only 60 Hello packets. If application-aware routing determines that it needs a new tunnel, it may take more than 1 minute to reset the tunnel.

Parameters for sliding window size

Use these parameters to configure the sliding window size for BFD-based application-aware routing.

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

Monitor tunnel statistics and routing paths

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

```
Device# show sdwan app-route stats

                                                                                                                  TX    RX
                                        SRC    DST    MEAN  MEAN            TOTAL          AVERAGE  AVERAGE  DATA  DATA
SRC IP       DST IP          PROTO  PORT   PORT   LOSS  LATENCY  INDEX  PACKETS  LOSS  LATENCY  JITTER   PKTS  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
...
```

Enable application visibility on Cisco IOS XE Catalyst SD-WAN devices

Provides steps to enable application visibility on Cisco IOS XE Catalyst SD-WAN devices for monitoring applications in all VPNs in the LAN.

Enable application visibility on Cisco IOS XE Catalyst SD-WAN devices to monitor all applications running in all VPNs in the LAN without configuring application-aware routing policy.

To configure application visibility on the router, use:

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

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

10 Enhanced Application-Aware Routing

Topics:

- Enhanced Application-Aware Routing in Cisco SD-WAN
- Configure enhanced application-aware routing
- Verify enhanced application-aware routing configuration
- Monitor enhanced application-aware routing using Cisco Catalyst SD-WAN Manager
- Enhanced application-aware routing troubleshooting

Enhanced Application-Aware Routing in Cisco SD-WAN

Explains how BFD collects loss, latency, and jitter metrics, how poll intervals and multipliers influence SLA-based path decisions, and how Enhanced Application-Aware Routing improves detection speed and traffic switchover time.

A Enhanced Application-Aware Routing is a Cisco Catalyst SD-WAN capability that

- speeds detection of tunnel performance issues
- enables faster traffic redirection to meet SLA thresholds, and
- reduces switchover time compared to default behavior.

Role of BFD in Performance Monitoring

Bidirectional Forwarding Detection (BFD) detects link failures and collects performance routing (PfR) data for SD-WAN tunnels (IPsec and GRE).

Each BFD hello packet gathers:

- Latency: Round-trip time between BFD echo request and reply.
- Jitter: Variation in packet arrival delay (irregular timing of packets).
- Loss: Number of echo requests that do not receive replies.

With the default BFD hello timer of 1 second, the system collects one PfR sample every second.

Poll Interval and Data Averaging

The system collects PfR samples throughout the poll interval (default: 10 minutes).

During each poll interval, it calculates the average for latency, jitter, and loss.

To make dynamic path decisions based on application-aware routing SLAs, the system uses a default multiplier of 6.

This multiplier evaluates multiple poll-interval averages before making routing decisions.

A poll interval average represents the time between consecutive data-collection events and indicates how frequently the system samples network performance.

Convergence Time Considerations

Convergence time is the time required for the network to recover and resume normal operation after a disruption.

- Default detection of slowly degrading WAN circuits: 10 minutes to 1 hour
- With the lowest recommended settings (2-minute poll interval × 6 intervals): 2 to 12 minutes

Using very low poll intervals can cause false positives and traffic instability due to insufficient performance data.

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.

Pfr measurements

Lists PfR measurement metrics, their sources, and descriptions for application-aware routing. Provides details on loss, latency, and jitter as measured by BFD.

PfR measurements provide insight into network performance by tracking loss, latency, and jitter using BFD.

Table 34: Pfr measurements

Metric	Source	Description
Loss	BFD	Measured as loss of BFD packet at 1pps or one packet in n_app_probe_class (n-apc) sec If the application probe class (APC) configuration is not set, BFD packet loss occurs at a rate of 1 packet per second (1 pps). With the APC configuration, packet loss reduces to 1 packet every N seconds. For more information see, Application Probe Class.
Latency	BFD	RTT measurements 1 pps or one packet in n-apc sec Without the application probe class (APC) configuration, RTT packet loss occurs at a rate of 1 packet per second (1 pps). With the APC configuration, packet loss reduces to 1 packet every N seconds.
Jitter	BFD	Variation in RTT

Application-aware routing design and measurement intervals

Lists default intervals, calculation methods, and recommendations for application-aware routing design and measurement accuracy.

This section provides information about default intervals, calculation methods, and expert recommendations for application-aware routing design and measurement accuracy.

- The default BFD hello-interval is 1 second, and the app-route/SLA poll-interval is 10 minutes.
- The BFD hello-interval defines how often the BFD (Bidirectional Forwarding Detection) protocol sends hello packets to detect the liveliness of a network path. The app-route/SLA poll-interval defines how often the network monitoring system collects data or measures network metrics related to application routes or Service Level Agreements (SLAs).
- By default, the system calculates to 60 minutes using 1 pps x 600 sec x 6 buckets. This calculation multiplies 1 packet per second (pps) by 600 seconds (10 minutes) and then by 6 buckets, resulting in a default poll-interval of 60 minutes.
- Experts suggest using a poll-interval of 120 seconds (2 minutes) and a multiplier of 5, which produces a 10-minute interval. This recommendation helps achieve a specific monitoring frequency.
- Reducing the poll-interval or multiplier improves detection time but can create false positives because PfR metrics rely on a small number of samples. Decreasing these values speeds up detection of network performance issues but increases the likelihood of false positives due to fewer data samples. You must balance detection time with the accuracy of PfR (Performance Routing) metrics.

- To improve measurement accuracy at a faster rate, reduce the BFD hello-interval. Decreasing the interval at which hello packets are exchanged allows more frequent detection of network path liveliness and improves measurement accuracy.

Benefits of enhanced application-aware routing

Lists the benefits of enhanced application-aware routing, including improved metric measurements, faster SLA detection and enforcement, and smoother transitions for network performance.

This list provides the key benefits of enhanced application-aware routing:

- Improved the PfR metrics: Inline data now improves PfR metric measurements (loss, latency, and jitter) by providing more accurate and detailed analysis. Cisco IOS XE Catalyst SD-WAN devices process and inspect this data directly at the network edge, enabling real-time inspection and decision-making.
- Quick Enhanced-App-Route Detection and SLA Enforcement: Reduced the PfR poll interval to a minimum of 10 seconds to enhance App-Route detection and SLA enforcement. Cisco IOS XE Catalyst SD-WAN devices can now quickly detect circuit degradation and remove tunnels from SLA forwarding when thresholds are not met, ensuring efficient and reliable performance.
- Speed: Increased the speed of SLA switch-over.
- SLA Dampening: Introduced SLA dampening to enable smoother transitions back to SLA forwarding. Tunnels now undergo a dampening period before re-entering SLA forwarding to prevent disruptions and instability and to minimize negative effects on network performance.
- Measurements: Enhanced measurement of loss, latency, and jitter.

Guidelines of enhanced application-aware routing

Outlines key guidelines for enhanced application-aware routing, including supported tunnel types, interface compatibility, SLA behavior, and performance considerations.

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.
- Does not impact tunnel scale and maintains minimal impact on memory and performance.
- Provides support with and without app-probe class configuration in SLA classes. For more information on app-probe class, see Application Probe Class.
- SLA dampening is supported.

Compatibility with Cisco IOS XE catalyst SD-WAN devices without enhanced application-aware routing

Lists compatibility considerations for Cisco IOS XE Catalyst SD-WAN devices that do not run enhanced application-aware routing.

Explains compatibility for Cisco IOS XE Catalyst SD-WAN devices when enhanced application-aware routing (EAAR) is not enabled or not available.

1. In these scenarios:
   - On the local side: When 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: When 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.

   It falls back to using BFD-based measurements where support for older releases or disabled features is required.
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, it reverts to using BFD-based measurements.

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

Restrictions for enhanced application-aware routing

Outlines the restrictions that apply to enhanced application-aware routing, including limitations on loopback unbind mode and per-queue measurement for GRE tunnels. Review these restrictions before enabling the feature.

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.
- GRE tunnels do not support per-queue measurement. Per-queue measurement monitors and analyzes network traffic on a per-queue basis by measuring and collecting metrics and statistics for each individual queue in a network device or system. A queue acts as a buffer where packets are stored before transmission or processing.

Prerequisites for enhanced application-aware routing

Lists the 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.

Supported devices for enhanced application-aware routing

Provides details on the devices that support enhanced application-aware routing.

Cisco IOS XE Catalyst SD-WAN device

Configure enhanced application-aware routing

Describes how to deploy enhanced app-aware routing configurations from Cisco Catalyst SD-WAN Manager to Cisco IOS XE Catalyst SD-WAN devices.

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.

Use one of these methods to configure enhanced application-aware routing:

- Configuration groups
- CLI commands
- Templates

Configure enhanced application-aware routing using a configuration group

Configure enhanced application-aware routing using a configuration group in Cisco Catalyst SD-WAN Manager.

Use these steps to configure enhanced application-aware routing (EAAR) using a configuration group in Cisco SD-WAN Manager.

Procedure

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

   Table 35: Enhanced app-route modes and parameters

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

7. Click **Save**.

Configure enhanced application-aware routing using CLI commands

Configure enhanced application-aware routing using CLI commands.

Use these steps to configure enhanced application-aware routing using CLI commands to enforce SLAs and optimize routing decisions based on application performance metrics.

Procedure

1. Enable enhanced PfR measurements for SLA enforcement.

Example

```
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 and local routers.

- The remote Cisco IOS XE Catalyst SD-WAN device must provide loss statistics to the local Cisco IOS XE Catalyst SD-WAN device.
- The local Cisco IOS XE Catalyst SD-WAN device utilizes these metrics to enforce SLAs.

2. Modify poll interval and multiplier settings for SLA enforcement and switchover.

Example

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

The default poll-interval is 10 seconds and multiplier is 6 for SLA enforcement and switchover. The aggressive mode setting for app route pfr multiplier is 6 by default; it is 5 for moderate mode.

Use the **bfd enhanced-app-route ignore-local-loss** command to ignore local loss on the router caused by QoS queues.

By default, loss includes both local and WAN loss.

3. Configure SLA dampening time and multiplier.

Example

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

SLA dampening time is the waiting period before returning the tunnel to SLA buckets after meeting the SLA. The default time is 120 seconds. The aggressive mode setting for dampening multiplier is 120 by default.

Configure enhanced application-aware routing using templates

Configure enhanced application-aware routing using a feature template in Cisco SD-WAN Manager.

Procedure

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 modes:

Table 36: Enhanced app-aware routing modes and parameters

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 a CLI template.

6. Click **Save**.

Verify enhanced application-aware routing configuration

Provides methods to verify enhanced application-aware routing (EAAR) configuration and operation using show commands.

To verify the enhanced application-aware routing configuration and display the configured parameters 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

                                                        SOURCE TLOC      REMOTE
```

```
 TLOC                                    DST PUBLIC     DST PUBLIC

SYSTEM IP         SITE ID   STATE  COLOR                  COLOR
    SOURCE IP          IP                        PORT        ENCAP  BFD-LD
FLAGS       UPTIME
------------------------------------------------------------------------------------------------------------------------
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) statistics details for each tunnel, across different intervals of measurements, for every configured APC.

Note

The total-packets field in the **show sdwan app-route stats summary** output displays values in multiples of 64 or 0 and includes both application packets and BFD probes. For per-interval application packet counts, refer to the tx-data-pkts and rx-data-pkts fields.

```
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
              TOTAL                      AVERAGE  AVERAGE     TX DATA     RX
DATA     IPV6 TX         IPV6 RX
INDEX    PACKETS  LOSS     LATENCY  JITTER           PKTS           PKTS
     DATA PKTS  DATA PKTS
------------------------------------------------------------------------------------------------------------------------
0            664              0             0                 0
       0                  0                 0                       0
1            663              0             0                 0
       0                  0                 0                       0
2            666              0             0                 0
       0                  0                 0                       0
3            664              0             0                 0
       0                  0                 0                       0
4            662              0             0                 0
```

```
          0                    0                    0                    0
5               664                    0                    0                    0
          0                    0                    0                    0
```

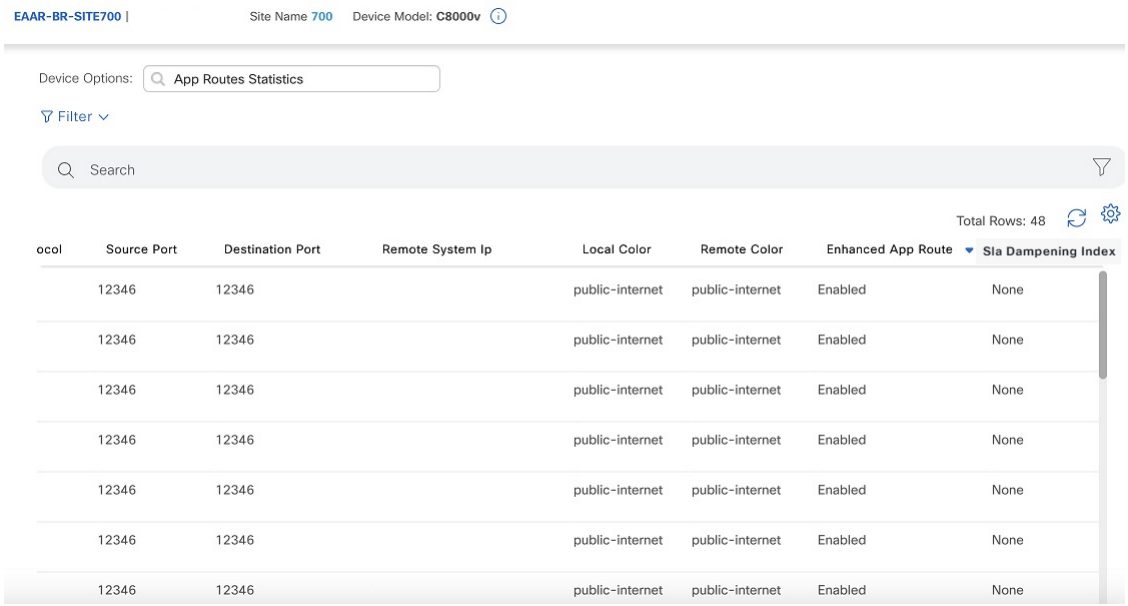
Monitor enhanced application-aware routing using Cisco Catalyst SD-WAN Manager

Describes how to monitor enhanced application-aware routing using Cisco SD-WAN Manager.

Procedure

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**.

   The following image shows the App Routes Statistics view:



Enhanced application-aware routing troubleshooting

Provides steps and example commands to troubleshoot enhanced application-aware routing on the device, including configuration and status verification.

This section provides troubleshooting steps and example command outputs for 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
                                        SOURCE TLOC      REMOTE TLOC
                            DST PUBLIC                       DST PUBLIC
SYSTEM IP        SITE ID   STATE       COLOR            COLOR           SOURCE
 IP                        IP                              PORT        ENCAP
BFD-LD    FLAGS             UPTIME
-------------------------------------------------------------------------------
172.16.0.0        100       down        private1         lte
10.0.0.0                      10.0.0.1                       12367       ipsec
  20011     EAAR                NA
172.16.0.1        500       down        private1         3g
10.0.0.0                      10.0.0.1                       12366       ipsec
  20013     EAAR                NA
172.16.0.0        600       down        private1         3g
10.0.0.0                      10.0.0.1                       12366       ipsec
  20007     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              private1
remote-color             3g
sla-class-index          0
```

```
fallback-sla-class-index None
enhanced-app-route       Enabled
sla-dampening-index      None
```

11 Traffic Flow Monitoring

Topics:

- Feature history for traffic flow monitoring
- Traffic flow monitoring
- Restrictions for traffic flow monitoring
- Configure traffic flow monitoring
- Verify traffic flow monitoring

Introduces traffic flow monitoring capabilities and guides through the configuration and verification process for monitoring network traffic patterns.

Feature history for traffic flow monitoring

Lists traffic flow monitoring feature history information.

Traffic flow monitoring feature history provides details about new features, enhancements, and capabilities introduced across different software releases.

Table 37: 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. Configure Cflowd traffic flow monitoring on Cisco SD-WAN Managers.
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. To enable logging of dropped packets, check the **Implicit ACL Logging** check box and to configure how often the packet flows are logged, enter the value in the **Log Frequency** field.
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. Configure Cflowd traffic flow monitoring to collect ToS, sampler ID, and remarked DSCP values in netflow 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. Configure this feature using the CLI template and also add-on CLI template.

Feature Name	Release Information	Description
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	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.
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.
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** > **Control WAN Interface Information**(choose the device option), in Cisco SD-WAN Manager to display the binding relationship between the loopback and physical interfaces.

Feature Name	Release Information	Description
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.
DTA Support for FNF Statistics	Cisco IOS XE Catalyst SD-WAN Release 17.18.1a Cisco Catalyst SD-WAN Manager Release 20.18.1	Cisco IOS XE Catalyst SD-WAN devices use DTA to handle all FNF statistics.

Traffic flow monitoring

Describes the process of monitoring and analyzing network traffic flows to track data movement, measure performance, and identify patterns across network infrastructure.

The following sections describe traffic flow monitoring.

Traffic flow monitoring with cflowd

Describes a flow analysis tool for analyzing Flexible NetFlow traffic data in overlay networks.

Cflowd is a flow analysis tool that

- analyzes Flexible NetFlow (FNF) traffic data
- monitors traffic flowing through Cisco vEdge deviceCisco 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.

Cflowd configuration and operation

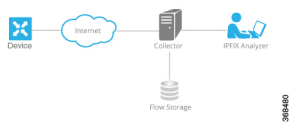
For a traffic flow, Cflowd periodically sends template reports to flow collector. These reports contain information about the flows and the data is extracted from the payload of these reports.

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 vEdge deviceCisco 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 vEdge deviceCisco IOS XE Catalyst SD-WAN devices do not cache any of the records that are exported to a collector.

> **Note**
>
> From Cisco IOS XE Catalyst SD-WAN Release 17.9.1a, netFlow on Secure Internet Gateway (SIG) tunnels is supported on Cisco vEdge deviceCisco IOS XE Catalyst SD-WAN devices. However, netflow is not supported on regular IPSec tunnel.

Cflowd and SNMP traffic monitoring comparison:

Table 38: Cflowd and SNMP monitoring differences

Cflowd	SNMP
Monitors service side traffic (LAN to WAN, WAN to LAN, LAN to LAN and DIA)	Monitors all traffic including non-service side traffic
Byte counting starts from L3 header	Byte counting starts from L2 header
Does not monitor BFD traffic	Monitors BFD traffic

IPFIX information elements for Cisco IOS XE Catalyst SD-WAN devices

Lists the IPFIX information elements that the Cisco IOS XE Catalyst SD-WAN device Cflowd software exports to collectors.

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	–
destinationIPv6Address	28	Yes	IPv6 destination address in the IP packet header.	ipv6Address (16 bytes)	default	–

Information Element	Element ID	Exported to External Collector	Description	Data Type	Data Type Semantics	Units or Range
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
application business-relevance	12244	Yes	Application business-relevance	varibale length	identifier	–

Flexible NetFlow for VPN0 interface

Describes Flexible NetFlow for VPN0 interface capabilities on Cisco IOS XE SD-WAN devices.

Flexible NetFlow for VPN0 interface is a network monitoring capability that

- enables bidirectional traffic visibility on a VPN0 interface of a Cisco IOS XE Catalyst SD-WAN device from Cisco IOS XE Catalyst SD-WAN Release 17.7.1a
- provides statistics on packets flowing through the device and helps to identify the tunnel or service VPNs, and
- provides visibility for all traffic (both ingress and egress) hitting VPN0 on Cisco IOS XE SD-WAN devices.

Flexible NetFlow components and configuration

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 39: Flexible netflow components

	Cisco Catalyst SD-WAN Flexible Netflow	Cisco SD-WAN Flexible Netflow VPN0 from Cisco vManage 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** and attach on VPN0 interface. Supported on CLI template and add-on CLI feature template in Cisco SD-WAN Manager.
Interface	Cisco Catalyst SD-WAN tunnel interface and service VPN interface	VPN0 interface except Cisco Catalyst SD-WAN tunnel and VPN 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

Consider these limitations when implementing Flexible Netflow on VPN0 interface configurations.

Flexible Netflow on VPN0 has specific limitations that affect interface support, record configuration, flow reporting, and protocol handling.

- 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, and 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.

- 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

Describes a feature that spreads the export of records in the monitor cache over a time interval to improve collector performance.

Flexible NetFlow export spreading is a feature that

- spreads out the export of records in the monitor cache over a time interval to improve collector performance
- prevents collector performance issues when multiple network devices export records simultaneously from synchronized caches, and
- distributes export bursts evenly across the cache timeout period instead of sending all records at once.

Export spreading operation

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. 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 these three spread intervals:

- **app-tables**: application-table, application-attributes option template
- **tloc-tables**: tunnel-tloc-table option template
- **other-tables**: other option templates

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.

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.

Spreading interval operation

This example shows 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 export of BFD metrics

Describes the Flexible NetFlow export of BFD metrics feature that enables exporting BFD telemetry data to external collectors for network monitoring and analysis.

Flexible NetFlow export of BFD metrics is a network monitoring feature that

- exports BFD telemetry data to external FNF collectors for analyzing average jitter, average latency, and loss per tunnel
- measures jitter and latency in microseconds and loss in units of 0.01%, and
- provides enhanced monitoring and faster collection of network state data.

Feature requirements and configuration

Minimum releases: Cisco IOS XE Catalyst SD-WAN Release 17.10.1a and Cisco Catalyst SD-WAN Control Components Release 20.10.1

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.

How the export of BFD metrics works

Describes how BFD metrics are exported through IPFIX packets to external collectors via the Forwarding TABLE Manager.

Minimum releases: Cisco IOS XE Catalyst SD-WAN Release 17.10.1a and Cisco Catalyst SD-WAN Control Components Release 20.10.1

Summary

The key components involved in the export of BFD metrics are:

- Cisco IOS XE Catalyst SD-WAN device: Responsible for sending IP Flow Information Export (IPFIX) packets to an external collector.
- Forwarding TABLE Manager (FTM): Generates the source metrics for BFD export.
- Cisco SD-WAN Controllers or Cisco SD-WAN Manager: Configuration points for BFD export interval settings.

Workflow

These stages describe how BFD metrics are exported:

1. 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.
   - 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**
   - 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.

FNF statistics management on Cisco IOS XE Catalyst SD-WAN devices

Explains how Cisco IOS XE Catalyst SD-WAN devices manage FNF statistics from Cisco IOS XE Catalyst SD-WAN Release 17.18.1a.

FNF statistics management on Cisco IOS XE Catalyst SD-WAN devices is a network monitoring capability that

- uses Data Tracking Agent (DTA) to export BFD metrics from FNF starting from Cisco IOS XE Catalyst SD-WAN Release 17.18.1a
- stabilizes the network by efficiently managing high IPFIX traffic, and
- ensures network performance remains unaffected.

Monitoring information

See Monitor BFD Metrics for information on monitoring.

Cflowd traffic flow monitoring with SAIE flows

Explains Cflowd traffic flow monitoring with SAIE flows feature for monitoring network traffic

Cflowd traffic flow monitoring with SAIE flows is a network monitoring feature that

- provides real-time device options for monitoring both Cflowd flows and SAIE flows
- allows filtering for displaying specific applications or application families running within a VPN, and
- enables traffic analysis on selected Cisco IOS XE Catalyst SD-WAN device devices.

Feature requirements and capabilities

Minimum releases: Cisco IOS XE Catalyst SD-WAN Release 17.10.1a and Cisco vManage Release 20.10.1

This feature enables you to choose two Cisco SD-WAN Manager real-time device options for monitoring both Cflowd flows and SAIE flows.

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 SAIE flows, see the SD-WAN Application Intelligence Engine Flow chapter.

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

Lists the benefits of Cflowd traffic flow monitoring with SAIE flows.

Cflowd traffic flow monitoring with SAIE flows provides several key benefits for network operations and performance management.

- 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

Requirement: prerequisites for cflowd traffic flow monitoring with SAIE flows

Outlines the minimum releases and configuration requirements for Cflowd traffic flow monitoring with SAIE flows.

You must meet the 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, refer to Configure global application visibility.
- For more information on configuring global flow visibility, refer to Configure global flow visibility.

Restrictions for cflowd traffic flow monitoring with SAIE flows

Outlines the limitations and constraints when using 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.

Maximum FNF record rate for aggregated data

Describes the maximum rate of aggregated traffic data FNF records that devices can send to reduce performance demands.

A maximum FNF record rate is a network configuration parameter that

- limits the rate (records per minute) of aggregated traffic data FNF records that a device can send
- reduces performance demands (CPU and memory) on the device, and
- helps manage network traffic when there is a large number of applications producing network traffic.

Raw and aggregated traffic flow data

Minimum supported releases: Cisco IOS XE Catalyst SD-WAN Release 17.14.1a, Cisco Catalyst SD-WAN Control Components Release 20.14.1

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.

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 the CLI.

Restrictions for traffic flow monitoring

Describes the notes, limitations, and restrictions that apply to traffic flow monitoring functionality.

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

Outlines restrictions for flow telemetry enhancement when using loopbacks as TLOCs.

Consider these restrictions when enabling collect loopback in flow telemetry with loopback 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 Acccess (DIA) scenario, is not supported.
- Multi-tenant scenario is not supported.
- All IP or IPv6 visibility features are sub-features to flow-visibility and app-visibility. You must enable flow-visibility or app-visibility before enabling the sub-features.

Configure traffic flow monitoring

Describes how to configure traffic flow monitoring on Cisco devices. Provides detailed information and procedures for setting up monitoring capabilities to track and analyze network traffic patterns.

The following sections provide information about configuring traffic flow monitoring.

Configure traffic flow monitoring on Cisco IOS XE Catalyst SD-WAN devices using policy groups

Use this task to enable flow visibility and application visibility for IPv4 and IPv6 traffic on Cisco IOS XE Catalyst SD-WAN devices within a configuration group to monitor network traffic patterns.

Before you begin

- A policy group must be created.
- Ensure that you have configured Cflowd collector details in the Cisco SD-WAN Manager menu from **Configuration** > **Network Hierarchy** > **Collectors** > **Cflowd**.

> **Note**
>
> The Cflowd configuration applies to the global level and not the site level.
>
> The additional settings that you configure are applied to the Cisco SD-WAN Controllers while deploying the application priority and SLA policy. For more information about configuring Cflowd, see the section Configure Cflowd in *Configure Collectors in a Network Hierarchy*.

- If you require manual configuration or are using a release that does not support the Policy Group method, use the CLI Add-on Profile method. For more information, refer to CLI Add-On Profile.

Follow these steps to traffic flow monitoring using a policy group:

Before you begin

Follow these steps to configure traffic flow monitoring on using policy groups:

Procedure

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Policy Groups**.
2. Select the desired policy group and click **Edit**.
3. Click **Application Priority and SLA**.
4. Click **Add Traffic Policy** or select an existing policy to edit.
5. Click **Additional Settings** to monitor traffic flow on incoming packets in the LAN for application and flow visibility over IPv4, IPv6, or both network addresses.
6. Enable traffic flow monitoring:

   From Cisco Catalyst SD-WAN Manager Release 20.13.x, the configuration for **Flow Visibility** and **Application Visibility**, and **Cflowd** policy is available under **Additional Settings** .

   Table 40: Additional Settings

Field	Description
Application Visibility	Monitor all the applications running in all VPNs over IPv4, IPv6, or both networks in the LAN.
Flow Visibility	Monitor traffic flow over IPv4, IPv6, or both network addresses in the LAN.

7. Click **Save**.

The flow visibility and application visibility configurations are saved to the policy group and are ready to be applied to the target devices.

What to do next

Deploy the policy group to the target devices to push the visibility settings to the network. For more information, Deploy Policy Group Workflow..

Configure traffic flow monitoring on Cisco IOS XE Catalyst SD-WAN devices using classic policies

Configure Cflowd traffic flow monitoring on devices using Flexible NetFlow (FNF) to export traffic data.

This task enables you to set up Cflowd traffic flow monitoring on Cisco IOS XE Catalyst SD-WAN devices using classic policies.

Cflowd traffic flow monitoring uses Flexible NetFlow (FNF) to export traffic data from your network devices. This configuration is performed using classic policies on Cisco IOS XE Catalyst SD-WAN devices.

Procedure

Perform the following steps to configure Cflowd monitoring.

- Configure global flow visibility on page 252
- Configure global flow visibility using configuration groups on page 255
- Configure global application visibility on page 255
- Configure a Cflowd monitoring policy on page 256
- View cflowd information on page 258

Traffic flow monitoring with Cflowd is configured on your Cisco IOS XE Catalyst SD-WAN device, and traffic data will be exported using Flexible NetFlow.

Configure global flow visibility

Configure Cflowd visibility globally on all Cisco IOS XE Catalyst SD-WAN devices to perform traffic flow monitoring on traffic from all VPNs in the LAN.

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.

Flow visibility enables monitoring of network traffic flows using Cflowd technology. This configuration allows you to monitor both IPv4 and IPv6 traffic flows across your SD-WAN network.

Procedure

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.

> **Note**
>
> Starting with Cisco SD-WAN Release 26.1.1, the FNF default cache size changes automatically based on the length of the defined FNF flow. For example, if you enable more optional features, the FNF default cache size decreases.

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 classification feature-manager EXMEM-usage** command to display the EXMEM memory usage for various clients.

```
Device# show platform hardware qfp 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 flow visibility using configuration groups

Configures global flow visibility on Cisco IOS XE Catalyst SD-WAN devices using configuration groups to monitor traffic flows across all service-side VPNs. This configuration enables network administrators to gain visibility into traffic patterns and troubleshoot network issues efficiently.

Enable Cflowd visibility globally on all Cisco IOS XE Catalyst SD-WAN devices to perform traffic flow monitoring for traffic originating from all VPNs in the LAN.

Before you begin

-

Follow these steps to configure global flow visibility:

Procedure

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

-

---

1. Click **Transport & Management Profile**.
2. Select the desired transport profile and click **Edit**.
3. Click **Edit Ethernet Interface** > **Tunnel**.
4. Enable **Allow Fragmentation** and **MTU To Max**.
5. Click **Save**.

What to do next

-

Configure global application visibility

Configure Cflowd visibility globally on all Cisco IOS XE Catalyst SD-WAN devices to perform traffic flowing monitoring on traffic coming to the router from all VPNs in the LAN.

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.

Procedure

---

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 a Cflowd monitoring policy

Configure a policy for Cflowd traffic flow monitoring using the Cisco SD-WAN Manager policy configuration wizard.

Configure a policy for Cflowd traffic flow monitoring to enable network traffic analysis and monitoring capabilities in your SD-WAN environment.

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.

Procedure

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** option is supported for cEdge devices only. Export spreading distributes the export of flow records over time to reduce the burst effect. This helps prevent situations where a large number of flows expire simultaneously and are exported at the same time, which can overwhelm the collector or cause network congestion.

    The default is **Disable**.

    g. To enable BFD metrics export, in the **BFD Metrics Export** field, click **Enable**.

Starting from Cisco IOS XE Catalyst SD-WAN Release 17.10.1a and Cisco vManage Release 20.10.1, BFD metrics export is supported for cEdge devices only. BFD metrics export provides additional information about the Bidirectional Forwarding Detection sessions and their associated metrics.

The default is **Disable**.

h. In the **Exporting Interval** field, 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

Lists the commands available to view Cflowd information on the Cisco IOS XE Catalyst SD-WAN device.

To view Cflowd information, use the commands on the Cisco IOS XE Catalyst SD-WAN device.

Available commands

- 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.

Sample output

This 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-proto 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

Configure Cflowd traffic flow monitoring using the CLI to enable flow visibility and monitoring on network devices.

This task configures Cflowd traffic flow monitoring to enable network visibility and flow monitoring capabilities on managed devices. The configuration includes setting up flow templates, collectors, and applying policies to monitor network traffic patterns.

Cflowd traffic flow monitoring provides visibility into network traffic patterns and flow characteristics. This configuration is performed from the Cisco SD-WAN Controller that controls the Cisco IOS XE Catalyst SD-WAN devices in your network.

Before you begin

Ensure you have CLI access to the Cisco SD-WAN Controller controller.

Follow these steps to configure Cflowd traffic flow monitoring using the CLI:

Procedure

1. Configure a Cflowd template to specify flow visibility and flow sampling parameters:

   Example

```
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.

2. To collect TOS, DSCP output and TLOC loopback in flow monitor:

   Example

   ```
   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
   ```

   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.

3. Configure a flow collector:

   Example

   ```
   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.

4. Configure a data policy that defines traffic match parameters and that includes the action **cflowd**:

Example

```
vSmart(config)# policy data-policy policy-name
vSmart(config-data-policy)# sequence number
vSmart(config-sequence)# match match-parameters
vSmart(config-sequence)# action cflowd
```

5. 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.

   Example

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

6. Apply the data policy to the sites in the overlay network that contain the Cisco IOS XE Catalyst SD-WAN devices:

   Example

   ```
   vSmart(config)# apply-policy site-list list-name
   vSmart(config-site-list)# data-policy policy-name
   vSmart(config-site-list)# cflowd-template template-name
   ```

Cflowd traffic flow monitoring is now configured and applied to the specified network sites. The system will collect flow data according to the configured template parameters and send it to the specified collectors for analysis and monitoring.

Configure flexible netflow on VPN0 interface

Configure FNF on a VPN0 interface using a CLI template or the CLI add-on template with the ezPM profile.

This task enables Flexible Netflow (FNF) on a VPN0 interface to monitor traffic metrics using the Easy Performance Monitor (ezPM) profile.

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.

Procedure

Configure the performance monitor context using the sdwan-FNF profile and specify the traffic monitor settings.

Example

```
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>
```

FNF is successfully configured on the VPN0 interface with the specified performance monitor context and traffic monitoring capabilities.

Configure flexible NetFlow with the export of BFD metrics using the CLI

Configure Flexible Netflow with BFD metrics export using the CLI.

Configure Flexible NetFlow to export BFD metrics data for network monitoring and analysis purposes.

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.

Procedure

1. Enable the export of BFD metrics.

   Example

   ```
   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.

   Example

   ```
   policy
    cflowd-template template-name
      collector vpn vpn-id address ip-address port port transport transport
      source-interface interface
      no bfd-metrics-export
   ```

BFD metrics export is configured according to your requirements. The system will export BFD metrics data at the specified intervals if enabled, or stop exporting if disabled.

Complete configuration example

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

Provides configuration examples for Flexible NetFlow export of BFD metrics.

This reference provides configuration examples for Flexible NetFlow export of Bidirectional Forwarding Detection (BFD) metrics. Minimum supported releases: Cisco IOS XE Catalyst SD-WAN Release 17.10.1a and Cisco Catalyst SD-WAN Control Components Release 20.10.1.

This 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
```

This 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'}
```

This example displays cflowd Forwarding Table Manager (FTM) statistics:

Minimum supported release: Cisco IOS XE Release 17.4.1a

```
device# show sdwan app-fwd cflowd statistics ftm
      ftm-flow-rate-limit      :      0
      ipfix-data-flow-rate     :      0
      ipfix-data-packet-rate   :      0
      flow-rate-limit-drop     :      0
      app-aggregation-db-cnt   :      0
      app-aggregation-aged-cnt :      0
      app-aggregation-drop-cnt :      0
      app-aggregation-high-watermark :0
======================================================================
      ftm-fw-zone-pair         :      0
      ftm-fw-zone              :      0
      ftm-utd-policy           :      0
      ftm-utd-policy-aged      :      0
      ftm-utd-urlf-url         :      0
      ftm-utd-urlf-url-aged    :      0
      ftm-utd-amp-filename     :      0
      ftm-utd-amp-filename-aged:      0
      ftm-utd-amp-malname      :      0
      ftm-utd-amp-malname-aged :      0
      ftm-c3pl-class           :      0
      ftm-c3pl-policy          :      0
```

This example displays cflowd Data Tracking Agent (DTA) statistics:

Minimum supported release: Cisco IOS XE Catalyst SD-WAN Release 17.18.1a

```
device# show sdwan app-fwd cflowd statistics dta
DTA Common Statistics Summary

dta_flow_rate_limit: 27000
dta_punt_packet_rate_limit: 1928
dta_flow_seq_drop: 0
dta_fnf_cb_cnt: 5405
dta_lowq_active_cnt: 0
ipfix_data_flow_rate: 0
ipfix_data_packet_rate: 0
```

Apply and enable cflowd policy

Configure and activate Cflowd policy to enable traffic flow monitoring on specified sites.

Apply and activate Cflowd policy to enable traffic flow monitoring and data collection on sites in the overlay network.

A centralized data policy must be applied to a list of sites and activated before it takes effect. The Cflowd template must be associated with the data policy for traffic flow monitoring capabilities.

Before you begin

Follow these steps to apply and enable Cflowd policy:

Procedure

1. Apply the centralized data policy to a site list.

   Example

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

2. Associate the Cflowd template with the data policy to activate it.

   Example

   ```
   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.

3. Commit the configuration to activate it.

   After you successfully activate the configuration by issuing a **commit** command, the Cisco Catalyst SD-WAN Controller pushes the data policy to the Cisco vEdge deviceCisco 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.

4. Display the centralized data policy configuration on the Cisco Catalyst SD-WAN Controller.

   Example

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

5. Display the centralized data policy that has been pushed to the Cisco IOS XE Catalyst SD-WAN device.

   Example

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

6. Enable Cflowd visibility directly on Cisco vEdge deviceCisco IOS XE Catalyst SD-WAN devices for traffic flow monitoring on all VPNs.

   Example

   ```
   Device(config)# policy flow-visibility
   ```

   You can enable Cflowd visibility directly on Cisco vEdge deviceCisco 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.

   > **Note**
   >
   > Do not attach the flow monitor to a certain interface after configuring the flow or app visibility. The **policy flow-visibility** command applies the global flow monitor. You need not attach the monitor to any interface again manually.

7. Monitor the applications using the show commands on the device.

   To monitor the applications, use the **show app cflowd flows** and **show app cflowd statistics** commands on the device.

The Cflowd policy is successfully applied and activated, enabling traffic flow monitoring across the specified sites. The policy configuration is pushed to the devices and flow visibility is enabled for monitoring applications.

Configure cflowd traffic flow monitoring

Configure Cflowd traffic flow monitoring with centralized data policy to monitor traffic and send flow data to collectors.

This task enables you to configure Cflowd traffic flow monitoring with a centralized data policy to monitor network traffic and send flow information to collectors for analysis.

Cflowd traffic monitoring is configured on a Cisco Catalyst SD-WAN Controller controller using centralized data policy. This approach allows you to monitor specific traffic types and send flow data to collectors for network analysis and troubleshooting.

Before you begin

Follow these steps to configure Cflowd traffic flow monitoring:

Procedure

1. Create a Cflowd template to define the location of the collector and to modify Cflowd timers.

   Example

```
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.

   Example

```
vsmart(config)# policy lists vpn-list vpn_1 vpn 1
```

3. Create a list of sites to apply the data policy to.

   Example

```
vsmart(config)# policy lists site-list cflowd-sites site-id 400,500,600
```

4. Configure the data policy.

   Example

```
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.

   Example

```
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.

Example

```
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
```

Cflowd traffic monitoring is configured and active. The system monitors TCP traffic from the specified VPNs and sites, sending flow data to the configured collector. You can verify the configuration and check flow statistics using the appropriate show commands.

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
```

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_ipv6-16679632136623 63:
  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
```

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
```

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
```

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 the CLI

Configure the maximum rate for a device to send aggregated traffic data to Cisco SD-WAN Manager.

This task configures the maximum FNF record rate to control the volume of aggregated traffic data that a device sends per minute.

Minimum supported releases: Cisco IOS XE Catalyst SD-WAN Release 17.14.1a, Cisco Catalyst SD-WAN Control Components Release 20.14.1

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.

Procedure

Configure the maximum rate (FNF records per minute) for a device to send aggregated traffic data to Cisco SD-WAN Manager.

Example

```
policy app-agg-node max-records-per-minute
```

The device is configured with the specified maximum FNF record rate for sending aggregated traffic data.

Set maximum FNF records to 1000 per minute

The following configures a device to send a maximum of 1000 FNF records per minute of aggregated traffic data.

```
policy app-agg-node 1000
```

Restore default maximum FNF records

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

Describes how to verify that traffic flow monitoring is functioning correctly and provides information about the verification process and related procedures.

The following sections provide information about verifying traffic flow monitoring.

Verify collect loopback

Provides commands and sample outputs to verify the ingress and egress interface output for collect loopback.

You can verify the ingress and egress interface output using these commands.

show sdwan app-fwd cflowd flows

This 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-proto 1
 tcp-cntrl-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 this command.

show sdwan app-fwd cflowd table

This 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
   APPQOE  APPQOE

     TCP

      SLA  COLOR                                        FEC   FEC   DUP D
DUP D  DUP   CXP                    CXP          SSL    SSL      EN     SSL
 EN   DE     SSL DE   SSL      SSL      SSL                      APPQOE  DRE
    DRE
                                                    SRC    DEST         IP
     CNTRL  ICMP    TOTAL  TOTAL                          EGRESS INTF
  INGRESS INTF                                             DROP    DROP
     NOT  NOT    QUEUE                DSCP    SAMPLER  D     R     PKTS
PKTS   R     D              SERVICE  PATH  REGION  READ   WRITTEN  READ
WRITTEN  READ   WRITTEN  SERVICE  TRAFFIC  POLICY  APPQOE  APPQOE   PASS
INPUT   INPUT    APPQOE
VPN  SRC IP                  DEST IP                  PORT   PORT   DSCP
PROTO  BITS   OPCODE  PKTS   BYTES   START TIME               NAME
    NAME             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    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      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        0
     0         0
0    10.0.5.5                 10.0.15.10                58048  22     4     6
      24     0       41     1752    Tue Jun 27 09:21:06 2023  internal0/0/rp:0
  GigabitEthernet9  unknown      network-service     No Drop       0        0
      0    0       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  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       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        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        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       2      1           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       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        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       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        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      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        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       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        0
     0         0
```

Verify interface binding on the device

Provides verification of interface binding on the device using a show command.

You can verify the interface binding on the device using this command.

show sdwan control local-properties WAN-interface-list

This 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

                          PUBLIC          PUBLIC PRIVATE         PRIVATE
                            PRIVATE
MAX   RESTRICT/            LAST         SPI TIME    NAT  VM          BIND
INTERFACE                 IPv4            PORT   IPv4            IPv6
                            PORT    VS/VM COLOR            STATE
CNTRL CONTROL/     LR/LB  CONNECTION   REMAINING   TYPE CON REG     INTERFACE


                                                                      STUN
                                          PRF IDs
--------------------------------------------------------------------------------
GigabitEthernet1              10.0.10.10      12346  10.0.10.10      ::
                                12346    2/1  lte              up     2
no/yes/no   No/No  0:20:20:27   0:01:14:20  N    5  Default N/A

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 the flexible NetFlow configuration on the VPN0 interface

Lists commands to verify Flexible NetFlow configuration on the VPN0 interface including flow records, monitors, cache details, and exporter information.

This reference provides verification commands for Flexible NetFlow (FNF) configuration on the VPN0 interface, including commands to view flow records, monitor configuration, cache details, and exporter settings.

View flexible netflow record configuration summary

You can verify FNF record configuration using this command.

```
Device# show flow record <monitor-context-name>
```

Note

The monitor name is used as temp0 in these examples.

This 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
```

This 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
```

This 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
```

This 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

This 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
……
```

This 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
```

This 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 the flexible NetFlow configuration with the export of BFD metrics

Provides verification commands and sample outputs for Flexible NetFlow configuration with BFD metrics export.

Minimum releases: Cisco IOS XE Catalyst SD-WAN Release 17.10.1a and Cisco Catalyst SD-WAN Control Components Release 20.10.1

Flow exporter configuration verification

The 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)
```

Flow exporter statistics verification

The 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
```

Flow exporter templates verification

The 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

12 Application Performance Monitor

Topics:

- Application performance monitor
- Restrictions for performance monitoring implementation
- Configure application performance monitor
- Performance monitoring configuration verification

Provides details on the Application Performance Monitor (APM) feature in Cisco Catalyst SD-WAN which is a simplified framework to configure intent-based performance monitors.

The Application Performance Monitor (APM) feature in Cisco Catalyst SD-WAN provides an express method for configuring an intent-based performance monitor with the help of predefined monitoring profiles.

Table 41: 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.

Application performance monitor

Describes a simplified framework to configure intent-based performance monitors, enabling real-time, end-to-end application performance visibility filtered by client, network, and server segments.

An application performance monitor is a simplified framework that

- enables you to configure intent-based performance monitors
- provides real-time, end-to-end application performance visibility filtered by client segments, network segments, and server segments, and
- collects performance metrics for specific traffic to help optimize application performance.

Key concepts in application performance monitoring

Application performance monitoring includes the following key concepts:

- 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.
- 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.
- Traffic Monitoring Specifications: You can choose to filter performance metrics using classification and sampler.

> **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 include the following filters:

- 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.

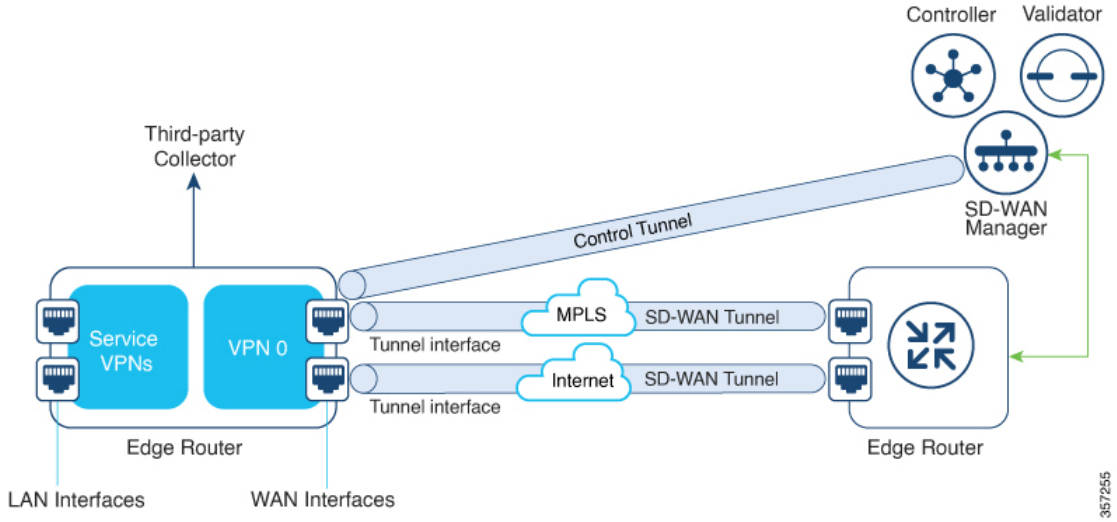
Application performance monitor provides the following features and benefits:

- 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.

The following image shows the performance monitoring workflow:

Figure 13: 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.

Restrictions for performance monitoring implementation

Outlines the limitations and usage constraints of performance monitoring.

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:

- 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.
- 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

Configure application performance monitor globally on all WAN tunnel interfaces or on specific WAN tunnel interfaces.

Enable monitoring of traffic metrics for application response time (ART), media monitors, or both to analyze network performance and troubleshoot issues.

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.

Before you begin

Procedure

1. Enable performance monitor globally on all Cisco Catalyst SD-WAN tunnel interfaces.

   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.

   Example

   ```
   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
   ```

2. Enable performance monitor on a specific tunnel interface.

   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.

   Example

   ```
   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
   ```

3. Specify additional monitoring filters and sampling rate.

   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-replace** . 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.

Example

```
class-map match-any match-audio
 match protocol rtp-audio
!
performance monitor context CISCO-APP-MONITOR profile sdwan-performance
 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
```

Application performance monitoring is configured and enabled according to your specified configuration. Traffic metrics for ART and media are now monitored on the configured interfaces.

Performance monitoring configuration verification

Lists details about the essential commands used to confirm the performance MONITOR configuration.

Performance monitoring configuration verification provides administrators with essential commands to confirm that performance monitors are properly configured and operational. This verification process ensures that traffic monitors are enabled, interfaces are correctly configured, and exporters are functioning as expected.

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
======================
 Tunnel1
```

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:                                0High Watermark:
                  0

Flows added:                                     0
Flows aged:                                      0
  - Active timeout      (  1800 secs)           0
  - Inactive timeout    (    15 secs)           0
  - Event aged                                   0
  - Watermark aged                               0
  - Emergency aged                               0
```

13 Enhanced Policy Based Routing

Topics:

- Enhanced policy based routing
- EPBR configuration
- EPBR monitoring commands

Lists the development milestones and release information for OMP, including feature description.

Enhanced Policy Based Routing (ePBR) is a protocol-independent traffic-steering mechanism that routes traffic based on flexible policies for traffic flows in Cisco Catalyst SD-WAN.

Table 42: 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.

Enhanced policy based routing

Explains that ePBR is an advanced version of Policy Based Routing that forwards traffic based on flexible policies (using C3PL language) rather than standard routing tables.

Enhanced Policy Based Routing (ePBR) is an advanced version of Policy Based Routing (PBR) that

- forwards traffic based on policies rather than routing tables
- provides more control over routing by extending and complementing existing mechanisms provided by routing protocols, and
- routes traffic based on flexible match criteria such as IPv4 and IPv6 addresses, port numbers, protocols, or packet size.

Features and capabilities

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.

ePBR provides these key features:

- 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.
- If the current hop isn't reachable, syslog messages are generated and the user is notified of the same.

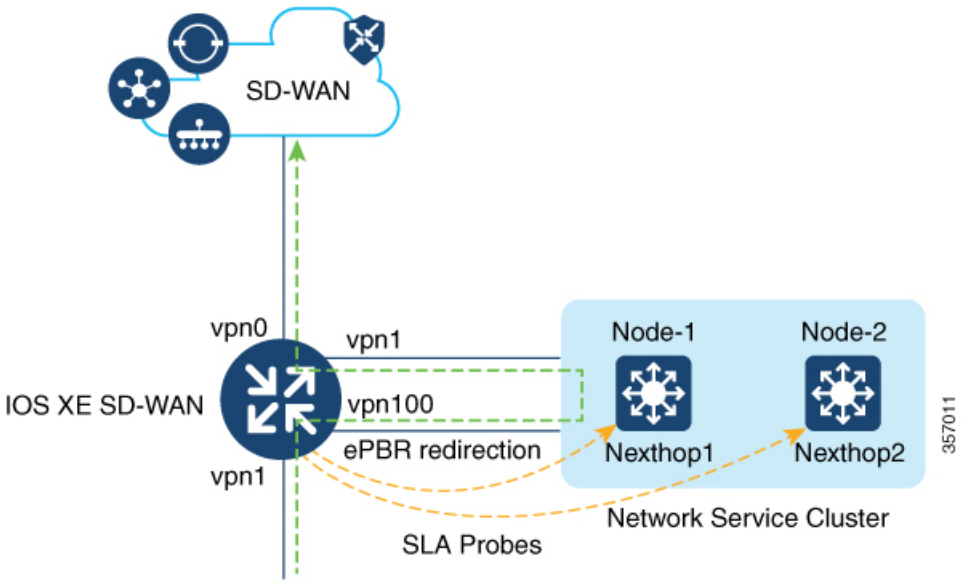
SLA probes can be configured in the same or a different VRF.

ePBR operates using the following mechanisms:

- 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.

Traffic redirection with ePBR

Figure 14: 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.

EPBR configuration

Provides the general method for implementation using CLI add-on feature templates within Cisco SD-WAN Manager and configuration examples using ACLs, class maps, and policy maps.

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 epbr test300
 class test300
  set ipv4 vrf 300 next-hop 10.0.0.2 10.0.40.1 10.0.50.1 ...
policy-map type epbr 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 epbr input test300
interface GigabitEthernet0/0/2
 service-policy type epbr 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

The following example shows how to configure IPv6 tracking with ePBR:

```
ip sla 1
  icmp-echo 2001:DB8::1
  vrf 100
ip sla schedule 1 life forever start-time now
track 1 ip sla 1 state
ip sla 2
  icmp-echo 2001:DB8::2
  vrf 300
ip sla schedule 2 life forever start-time now
track 2 ip sla 2 state
ip sla 3
  icmp-echo 2001:DB8::3
  vrf 400
ip sla schedule 3 life forever start-time now
track 3 ip sla 3 state
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 verify-availability 2001:DB8::2 10 track 2
  set ipv6 vrf 400 next-hop verify-availability 2001:DB8::3 11 track 3
```

```
policy-map type epbr test100_v6
 class test100_v6
  set ipv6 vrf 100 next-hop verify-availability 2001:DB8::1 10 track 1
!
interface GigabitEthernet0/0/1
 service-policy type epbr input test300_v6
interface GigabitEthernet0/0/2
 service-policy type epbr input test100_v6
!
```

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.

EPBR monitoring commands

Lists CLI show commands for ePBR monitoring rather than using the Cisco SD-WAN Manager interface.

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
```

14 Forward Error Correction

Topics:

- Supported devices for forward error correction
- Configure forward error correction for a policy
- Forward error correction status monitoring using CLI

Explains a packet loss recovery mechanism that sends extra parity packets to recover lost data on network links.

Forward Error Correction (FEC) is a packet loss recovery mechanism that

- recovers lost packets on a link by sending extra "parity" packets for every group of 4 packets
- allows up to a single lost packet in the group to be recovered as long as the receiver receives a subset of packets in the group (at-least N-1) and the parity packet, and
- is supported on Cisco IOS XE Catalyst SD-WAN devices.

Feature history and supported traffic

Table 43: 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.
Adaptive FEC Based on Packet Loss Threshold	Cisco IOS XE Release 17.6.2 Cisco vManage Release 20.6.1	This feature lets you configure the mean loss percentage threshold level to trigger adaptive FEC. Adaptive FEC is turned on only if the packet loss level exceeds the configured threshold and turns off once the packet loss level in the network has improved.
IPv6 support for Forward Error Correction	Cisco IOS XE Catalyst SD-WAN Release 17.15.1a	IPv6 support for Forward Error Correction is added.

> **Note**
>
> We recommend Cisco IOS XE Release 17.6.3 as the minimum release when using FEC.

> **Note**
>
> Monitoring FEC through Cisco SD-WAN Manager is not supported.

Cisco IOS XE Catalyst SD-WAN Devices support forward error correction 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

Supported devices for forward error correction

Lists the Cisco IOS XE Catalyst SD-WAN devices that support 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

Configure forward error correction (FEC) in a centralized policy to reduce packet loss and improve network reliability.

Configure forward error correction (FEC) to automatically send redundant information with data transmissions, helping to recover lost packets and maintain network reliability. FEC can be configured to operate adaptively based on loss thresholds or to always send correction data.

Forward error correction is configured as part of a centralized policy's QoS data policy. This feature helps maintain data integrity across network links by providing packet recovery mechanisms when transmission errors occur.

Before you begin

Follow these steps to configure forward error correction for a policy:

Procedure

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Policies** > **Centralized Policy** > **Add Policy**.
2. Click **Next**, **Next** again and then click **Configure Traffic Rules**.
3. Click **Traffic Data**, and from the **Add Policy** drop-down list, choose **Create New**.
4. Click **Sequence Type**. From the **Add Data Policy** pop-up menu, choose **QoS**.
5. Click **Sequence Rule**. In the **Applications/Application Family List**, choose one or more applications or lists.
6. Click **Accept**. Click **Actions** and click **Loss Correction**.
7. 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.

     Loss Threshold: Specify the packet loss threshold for automatically enabling FEC.

     > **Note**
     >
     > The **Loss Threshold** field is displayed only if you choose **FEC Adaptive**.

     The default packet loss threshold is 2%.

     You can specify a loss threshold of 1 to 5%.
   - **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.
8. Click **Save Match and Actions**.
9. Click **Save Data Policy**. 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.

The forward error correction configuration is now part of your centralized policy. The policy will automatically apply FEC based on your selected mode (adaptive or always) to help maintain data integrity and reduce packet loss for the specified applications.

Forward error correction status monitoring using CLI

Provides FEC status verification methods and counter descriptions for tunnel statistics monitoring.

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 192.0.2.1 203.0.113.1 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 192.0.2.1 203.0.113.1 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.

15 Packet Duplication

Topics:

- Feature history for packet duplication
- Packet duplication
- Configure packet duplication
- Configure underlay fragmentation using configuration groups
- Restrictions for packet duplication
- Monitor packet duplication statistics for a device using the CLI
- Monitor packet duplication statistics for a device using Cisco SD-WAN Manager
- Monitor tunnel information for a device

Feature history for packet duplication

Lists the feature history information for packet duplication functionality across different software releases.

This reference provides feature history information for packet duplication functionality, showing the evolution of capabilities across different software releases.

Table 44: 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. IPv6 support for Packet Duplication is also added.
Packet duplication on more than one additional tunnel	Cisco IOS XE Catalyst SD-WAN Release 26.1.1	This feature allows you to generate multiple copies for each original packet on more than one additional tunnels.

Packet duplication

Explains a network feature that duplicates packets on alternate available SD-WAN tunnels to mitigate packet loss.

Packet duplication is a network feature that

- duplicates packets on alternate available SD-WAN tunnels to mitigate packet loss
- enables receiving Cisco IOS XE Catalyst SD-WAN devices to forward one packet to the server and discard the duplicate packets, and
- is suitable for edges with multiple WAN links.

Packet duplication capabilities

You can view the tunnel packet duplication statistics using device CLI or the Cisco SD-WAN Manager device dashboard.

> **Note**
>
> From Cisco IOS XE Catalyst SD-WAN Release 17.18.1a, the tunnel selection algorithm has been optimized. While selecting a duplicate tunnel, the highest preference is given to the tunnel which has a different local color compared to that of the primary tunnel.

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.

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 these 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

Configure packet duplication

Use one of these methods to configure packet duplication:

- Configure packet duplication using policy groups
- Configure packet duplication on additional tunnels using the CLI
- Configure packet duplication using classic policies

Configure packet duplication using policy groups

Configure packet duplication using policy groups to implement loss correction for network traffic.

Configure packet duplication through policy groups to enhance network reliability by creating duplicate packets for critical traffic flows.

Packet duplication is a loss correction technique that sends duplicate copies of packets across different paths to ensure reliable delivery of critical network traffic. This feature is available starting from Cisco Catalyst SD-WAN Control Components Release 20.14.1.

Procedure

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**.

The packet duplication policy is configured and applied to the specified traffic flows, providing loss correction through duplicate packet transmission.

Configure packet duplication on additional tunnels using the CLI

When packet duplication is enabled, by default, packets are duplicated on one additional tunnel. Using the **packet-duplication max-preselect**<*2-4*> command, you can enable duplication on 2-4 additional tunnels.

Before you begin

If you do not configure this command, the duplication factor remains 1 by default.

This configuration only configures the duplication number. The actual duplication is enabled using the policy configuration method.

Procedure

Configure the duplication number.

Example

```
config-tr
    sdwan
     packet-duplication max-preselect <2-4>
```

Configure packet duplication using classic policies

Configure packet duplication through classic policies to reduce packet loss by duplicating packets and sending them on available links.

Configure packet duplication to reduce packet loss by duplicating packets and sending them on the next available links.

Use classic policies to configure packet duplication for applications that require enhanced reliability and reduced packet loss. This feature is configured through centralized policies in Cisco SD-WAN Manager.

Procedure

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.

The packet duplication policy is configured and applied to the selected applications and VPNs. Packets for the specified applications will be duplicated and transmitted over multiple links to reduce packet loss.

Configure underlay fragmentation using configuration groups

Configure underlay fragmentation settings using transport and management profiles in configuration groups.

Configure underlay fragmentation to optimize data transmission across network interfaces by enabling fragmentation and MTU settings through configuration groups.

Underlay fragmentation configuration is necessary when network interfaces need to handle data packets that exceed the maximum transmission unit (MTU) size. This task applies to environments where network optimization is required for efficient data flow.

Minimum supported releases: Cisco IOS XE Catalyst SD-WAN Release 17.15.1a

Before you begin

Follow these steps to configure underlay fragmentation using configuration groups:

Procedure

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**.

The underlay fragmentation configuration is now enabled for the selected transport profile, allowing the system to handle packet fragmentation and optimize MTU settings for improved network performance.

Restrictions for packet duplication

Outlines limitations and constraints for packet duplication functionality on network devices.

Consider these restrictions when implementing packet duplication:

- Packet duplication interop along with 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. Due to latency differences between tunnels, there may be a transient period where duplicate packets are forwarded to the application.
- 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.
- Packet duplication is supported only on Cisco IOS XE Catalyst SD-WAN devices but not on all the edge devices.
- At the hub, packet duplication cannot be performed for traffic received from spoke1 and forwarded to spoke2 and also vice versa. Packet duplication is not supported for traffic arriving from the WAN (from-tunnel direction) within policy. It is only supported for service-side traffic (from-service direction).

Monitor packet duplication statistics for a device using the CLI

Provides sample output from the show sdwan tunnel statistics pkt-dup command for monitoring packet duplication statistics.

This reference shows sample output from the **show sdwan tunnel statistics pkt-dup** command.

```
Device#show sdwan tunnel statistics pkt-dup
tunnel stats ipsec 192.0.2.1 203.0.113.1 12366 12406
pktdup-rx                           1313388
pktdup-rx-other                     526666
pktdup-rx-this                      1324958
```

```
pktdup-rx-fwd                        540190
pktdup-rx-fwd-dup-tun                789569
pktdup-tx                            0
pktup-tx-other                       0
pktdup-tx-dup-tun-selection-failed   0
pktdup-tx-dup-tun-sent-failed        0
pktdup-capable                       true
pktdup-tx-intercepted               0
pktdup-max-preselect                4
```

Monitor packet duplication statistics for a device using Cisco SD-WAN Manager

Configure monitoring of packet duplication statistics for network devices through the real-time monitoring interface.

This task enables you to view and analyze packet duplication counters for tunnel traffic on a specific device, helping you monitor network performance and troubleshoot connectivity issues.

Packet duplication statistics provide insights into tunnel packet handling, including original packets received, duplicate packets processed, and failed duplication attempts. These metrics are essential for monitoring network redundancy and tunnel performance.

Procedure

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**.

   The following table describes the packet duplication counters that appear in the **Tunnel Packet Duplication Statistics** pane:

Name of Counter	Description
pktdup-rx	Displays number of original packets received by primary tunnel.
pktdup-rx-other	Displays the number of duplicate packets received by this duplicate tunnel.
pktdup-rx-this	N/A
pktdup-rx-fwd	Displays the number of packets received and forwarded by this tunnel.
pktdup-rx-fwd-dup-tun	Displays the number of packets received and forwarded by the duplicate tunnel(s).
pktdup-tx	Displays the total number of duplicate packets generated with respect to this tunnel traffic which will be sent out on duplicate tunnel(s).
pktdup-tx-other	Displays the total number of duplicate packets sent out by this tunnel, but these packets belongs to some other original tunnel.

Name of Counter	Description
pktdup-tx-dup-tun-selection-failed	Displays the total number of failed duplication because duplicate tunnels are not available.
pktdup-tx-dup-tun-sent-failed	Displays the total number of duplicated packets failed to be sent out by duplicate tunnel(s).
pktdup-tx-intercepted	Displays the total number of packets which need to be duplicated with respect to this tunnel (equivalent to policy-match).
pktdup-capable	Indicates that this tunnel is capable of supporting the packet duplication functionality.

The tunnel packet duplication statistics are displayed with real-time counter values showing packet processing metrics for the selected device.

Monitor tunnel information for a device

Configure monitoring of tunnel connection information for a specific device, including loss percentage, latency, jitter, octets, and packet duplication metrics.

Monitor tunnel connection information for a device to view metrics such as loss percentage, latency, jitter, octets, and packet duplication counters.

Use tunnel monitoring to analyze the performance and health of tunnel connections on your devices. The monitoring interface provides detailed packet duplication counters and connection statistics to help troubleshoot network issues.

Procedure

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.

Table 45: Packet duplication counters

Packet Duplication Counters	Description	Mapping to CLI Counters
RX	Displays the number of original packets received on this primary tunnel.	pktdup-RX
RX Fwd	Displays the number of original packets forwarded by this primary tunnel.	pktdup-RX-fwd

Packet Duplication Counters	Description	Mapping to CLI Counters
RX Dup Fwd	Displays the number of duplicated packets forwarded by duplicate tunnels. RX Fwd + RX Dup Fwd = RX	pktdup-RX-fwd-dup-tun
TX Intercepted	Displays the number of packets intercepted for duplication.	pktdup-TX-intercepted
TX Dup	Displays the number of successful duplicate packets on the duplicate tunnel(s).	(pktdup-TX) - (pktdup-TX-dup-tun-sent-failed)
TX Dup Failed	Displays the number of failed duplications.	(pktdup-TX-dup-tun-selection-failed) + (pktdup-TX-dup-tun-sent-failed)

You can view tunnel connection information and packet duplication counters for the selected device in your chosen chart format.

16 Policy Configuration Tagging

Topics:

- Feature history for policy configuration tagging
- Supported devices for policy configuration tagging
- Restrictions for policy configuration tagging
- Policy configuration tagging
- Benefits of policy configuration tagging
- Configure policy configuration tagging using CLI commands
- Tag-instances verification

Introduces policy configuration tagging for network devices and guides through implementation using CLI commands to organize and manage network policies efficiently.

Feature history for policy configuration tagging

Describes a feature that allows grouping multiple policy objects under tags to control policy download speed, improve management of defined lists, and better organize intent-based network configurations.

Table 46: 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	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: • 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

Lists the devices and releases that support policy configuration tagging functionality in Cisco SD-WAN environments.

This reference provides information about the devices and software releases that support policy configuration tagging in Cisco SD-WAN deployments. For detailed information on supported models within each device family, refer to the Cisco SD-WAN Device Compatibility page.

Table 47: Supported Devices and Releases

Release	Supported Devices
Cisco IOS XE Catalyst SD-WAN Release 17.9.1a and later	• Cisco Catalyst 8500 Series Edge Platforms • Cisco Catalyst 8300 Series Edge Platforms • Cisco Catalyst 8200 Series Edge Platforms and 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)

Restrictions for policy configuration tagging

Outlines the limitations and constraints that apply when configuring policy tags in the system.

Consider the following restrictions when configuring policy tags:

- 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 Catalyst 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.

Policy configuration tagging

Explains a network feature that groups policy objects and assigns tag values to traffic flows for intent-based network configurations.

Policy configuration tagging is a network feature that

- allows you to group policy objects and assign tag values to various traffic flows by defining a policy
- enables you to name tags based on the functionality of policy objects used to achieve intent-based network configurations, and
- provisions tags through the Cisco Catalyst SD-WAN Controller for use in policy rules for traffic classification.

Tag configuration details

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.

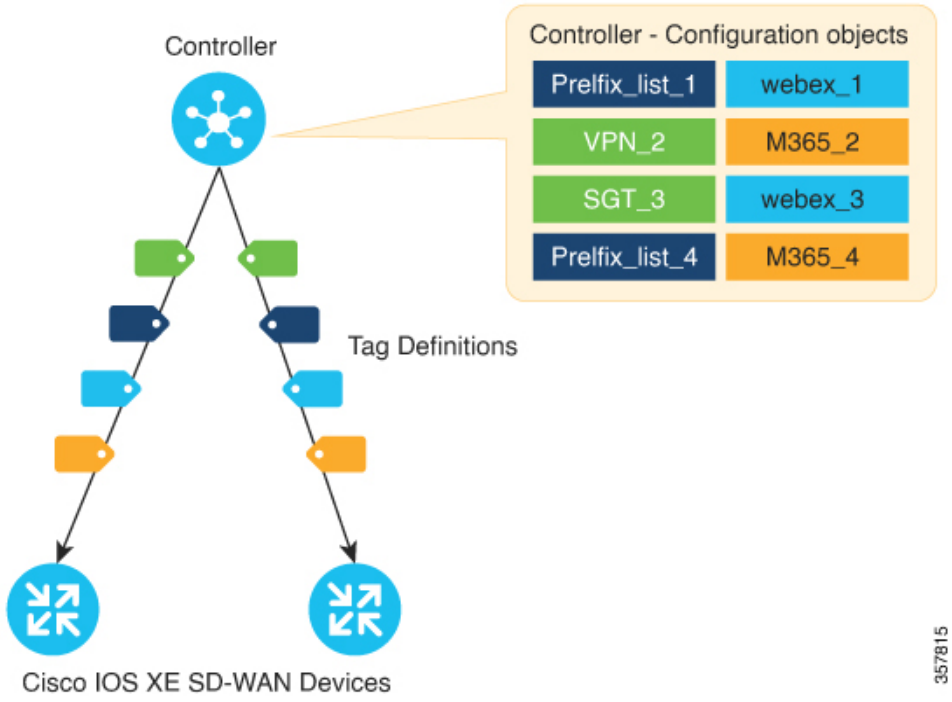
You can apply the configured tags in a match criterion under localized and centralized policies. Devices process the tag configurations and apply the configurations to the data plane when the tag is referenced in the policy.

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

The following figure shows policy configuration tagging in Cisco Catalyst SD-WAN:

Figure 15: Policy configuration tagging in Cisco Catalyst SD-WAN



As shown in the figure, at the Cisco Catalyst 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.

The policy configuration tagging feature provides the following capabilities:

- 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 Catalyst SD-WAN Controller CLI templates from Cisco SD-WAN Manager.

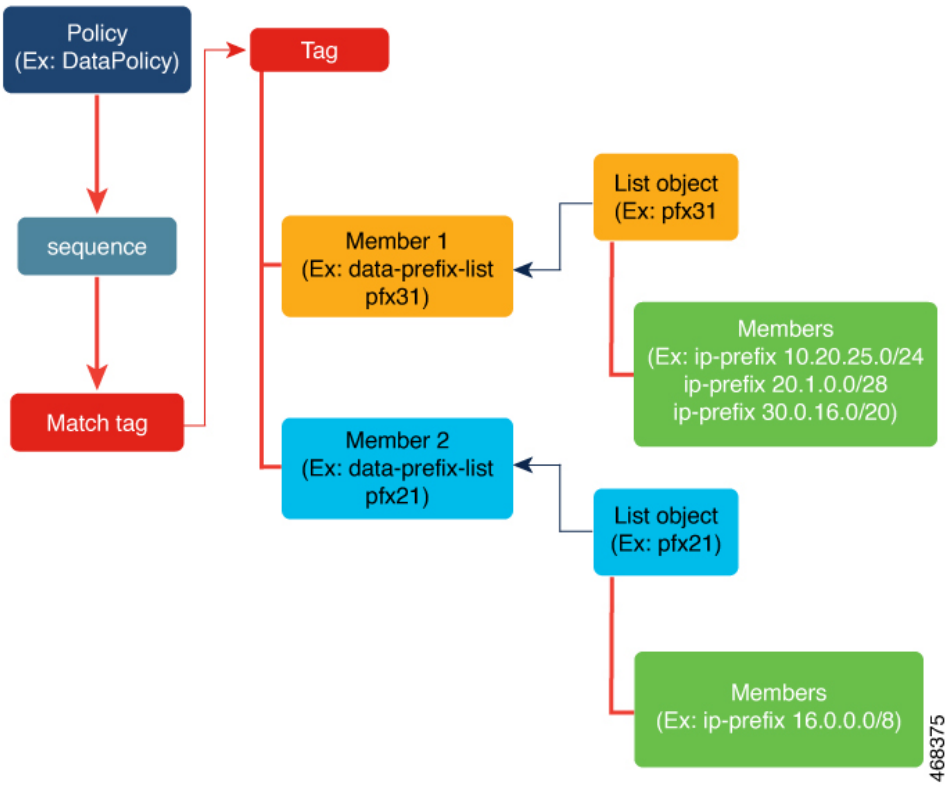
The tag workflow includes the following steps:

1. In Cisco Catalyst 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.

The following figure shows the tagging workflow with examples:

Figure 16: Tagging workflow with examples



Benefits of policy configuration tagging

Explains the advantages of implementing policy configuration tagging in network management systems.

A policy configuration tag is a network management feature that

- enables reusability of policy objects
- improves policy download performance by reducing configuration size and sequences, and
- provides better organization of configurations for intent-based networks.

Policy configuration tagging advantages

Policy configuration tagging enhances network operations in several ways:

- Enables reusability of policy objects.
- Enables faster policy downloads on devices due to reduced configuration size and sequence complexity.
- Supports tag sharing across different policies, allowing for greater flexibility and efficiency.
- Makes it possible to correlate policies across the network, aligning with user-defined intents.
- Controls policy configuration download speeds between the Cisco Catalyst SD-WAN Controller and the Cisco IOS XE Catalyst SD-WAN devices.
- Improves management of the defined lists in the network controller.
- Contributes to improved organization of configuration objects in intent-based networks.

Configure policy configuration tagging using CLI commands

Configure policy configuration tagging using CLI commands to create tag-instances and apply them to centralized policies.

This task allows you to configure policy configuration tagging using CLI commands to create tag-instances with member attributes and include them in localized and centralized policies for enhanced traffic matching and control.

Policy configuration tagging enables more granular traffic matching by allowing you to create tag-instances that can be referenced in access control lists, data policies, and Application Aware Route (AAR) policies. 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.

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.

Follow these steps to configure policy configuration tagging using CLI commands:

Procedure

1. Configure a new object tag-instance on Cisco Catalyst SD-WAN Controller:

   Example

   ```
   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:

   Example

   ```
   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:

   Example

   ```
   lists[app-list app-list-name] [data-prefix-list prefix-list-name] [data-ipv6-prefix-list
   ipv6-prefix-list-name]
   ```

4. Configure localized access-list policy (ACLs and IPv6 ACLs) to include destination or source tag instances in matching attributes:

Example

```
match [destination-tag-instance dest-tag-name | source-tag-instance src-tag-name]
```

5. Configure centralized data policy to include destination-tag-instance, source-tag-instance, or tag-instance in matching attributes:

   Example

   ```
   match [destination-tag-instance dest-tag-name | source-tag-instance src-tag-name |
   tag-instance tag-name]
   ```

6. Configure centralized Application Aware Route (AAR) policy to include destination-tag-instance, source-tag-instance, or tag-instance in matching attributes:

   Example

   ```
   match [destination-tag-instance dest-tag-name | source-tag-instance src-tag-name | tag-instance
     tag-name]
   ```

You have successfully configured policy configuration tagging using CLI commands. The tag-instances are created and can be referenced in localized and centralized policies for enhanced traffic matching and control.

Complete configuration example

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
    !
```

Tag-instances verification

Provides sample outputs and commands for verifying downloaded tag information from Cisco SD-WAN Controller on Cisco IOS XE Catalyst SD-WAN device.

Use the following commands and sample outputs to verify tag-instances, policy information, and tag classification data on the device.

Tag instances from Cisco SD-WAN Controller

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
```

Policy from Cisco SD-WAN Controller

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
```

Platform software common classification

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
```

17 Custom Applications

Topics:

- Custom applications
- Configure custom applications using Cisco SD-WAN Manager
- Verify custom applications

Provides feature history for custom applications.

This content describes the features related to defining custom applications in Cisco's SD-WAN solutions.

Table 48: 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.

Custom applications

Describes custom applications used in Cisco NBAR technology.

A custom application is a protocol that

- identifies internet traffic for uncommon network applications of specific interest to an organization
- augments the protocols provided in a protocol pack, and
- can be defined by users to enhance application visibility and control.

Custom applications reference information

Custom applications are defined protocols that help identify specific network applications.

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.

> **Note**
>
> The following terms are used in the documentation of related technologies, and are equivalent: custom applications, custom protocols, user-defined applications

From Cisco vManage Release 20.5.1, packet trace output includes these terms:

- Classification name: This is the application classification of traffic for policy purposes.
- Classification visibility name: This is the application classification of the traffic for traffic visibility purposes, such as for protocol-discovery counters and Flexible NetFlow (FNF).

The **NBAR** section of the Network-wide Path Insights (NWPI) **Insight advanced** views uses the same terms. In Cisco SD-WAN Manager, see this in **Tools** > **Network-wide Path Insights**.

In this example, the classification name and classification visibility name are different because custom applications are enabled only for visibility functionality (monitoring traffic).

```
Packet number in flow: 12
Classification state: Final
Classification name: ssl
Classification ID: 1312 [CANA-L7:453]
Candidate classification sources:
```

```
N/A
Early cls priority: 255
Permit apps list id: 0
Sdavc Early prioirty as app: 0
Classification visibility name: cisco_com
Classification visibility ID: 3714 [21:3714]
Number of matched sub-classifications: 0
Number of extracted fields: 0
Is PA (split) packet: False
Is FIF (first in flow) packet: False
TPH-MQC bitmask value: 0x4
Source MAC address: 00-00-5E-00-53-00
Destination MAC address: 00-00-5E-00-53-01
Traffic Categories: N/A
```

In this example, the Classification name and Classification visibility name match because custom applications are enabled for visibility functionality (monitoring traffic) and control functionality (traffic policy).

```
Packet number in flow: 12
Classification state: Final
Classification name: cisco_com
Classification ID: 1312 [CANA-L7:453]
Candidate classification sources:
N/A
Early cls priority: 255
Permit apps list id: 0
Sdavc Early prioirty as app: 0
Classification visibility name: cisco_com
Classification visibility ID: 3714 [21:3714]
Number of matched sub-classifications: 0
Number of extracted fields: 0
Is PA (split) packet: False
Is FIF (first in flow) packet: False
TPH-MQC bitmask value: 0x4
Source MAC address: 00-00-5E-00-53-00
Destination MAC address: 00-00-5E-00-53-01
Traffic Categories: N/A
```

Custom application functionality

- Releases earlier than Cisco vManage Release 20.5.1, creating a custom application has the following effect:

  A custom application does not affect visibility functionality (monitoring traffic) or control functionality (traffic policy) until a policy that makes use of the custom application is applied.

- Releases Cisco vManage Release 20.5.1, or later

  Creating a custom application has the following effect:

  - A custom application affects application visibility functionality only (monitoring traffic), such as for protocol-discovery counters and Flexible NetFlow (FNF), and does not affect traffic policy.
  - When the custom application is used by a policy, it affects control functionality (traffic policy) also.

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.

Custom application names

From Cisco vManage Release 20.5.1, when you create a new custom application and provide a name for the application, appends "-Custom" to the name. This ensures that the name does not conflict with other application types, such as those from Protocol Packs, or cloud-sourced applications.

Custom applications usage

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)

Restrictions for custom applications

Outlines restrictions for custom applications, including limits on custom application counts, server name wildcards, and supported features. Review these constraints before configuring custom applications.

Ensure that the following restrictions are adhered to when configuring custom applications:

- Maximum number of applications

  Maximum number of custom applications: 1100

- Maximum number of rules

  Maximum number of L3/L4 rules: 20000 and maximum number of server names: 50000

- Wildcards
  - Wildcard followed by a period

    For server names, maximum instances of wildcard followed by a period (.): 50000

    Example: ***.cisco.com** matches www.cisco.com, developer.cisco.com

    > **Note**
    >
    > ***.cisco.com**: Configuring the custom application server name as *.cisco.com (including the dot) results in an exact suffix match. The rule matches any domain ending in **.cisco.com**, regardless of the subdomains. Consequently, both **www.cisco.com** and **developer.cisco.com** match, whereas **www.cisco.com.test.org.us** does not match.

  - Prefix wildcard

    For server names, maximum instances of prefix wildcard as part of server name: 256

    Example: *ample.com matches www.example.com

> **Note**
>
> ***cisco.com**: Configuring the custom application server name as *cisco.com (omitting the dot) results in a substring match. Consequently, any domain string containing ***cisco.com**, such as **www.cisco.com** or **www.cisco.com.test.org.us**, will be matched.

- One domain, two applications

  Mapping the same domain to two different custom applications is not supported.

- First-packet classification

  To enable first packet classification with SD-AVC, unencrypted DNS traffic must be redirected through the router running on Cisco IOS XE Catalyst SD-WAN devices and must be visible within the same VRF as the application traffic.

- CLI

  Creating custom applications though CLI is not supported in Cisco Catalyst SD-WAN policy.

- Disaster recovery

  SD-AVC state replication in disaster recovery (DR) is not supported in Cisco Catalyst SD-WAN Manager Release 20.12.1 and earlier. DR switchover resets SD-AVC state for Cisco Catalyst SD-WAN Manager Release 20.12.1 and earlier releases. Post switchover, custom applications will not be present in SD-AVC.

- DNS traffic unencrypted

  For NBAR to correctly categorize traffic based on domain name, the DNS traffic needs to pass unencrypted through the router.

- Custom application name restriction

  Cisco Catalyst SD-WAN does not support creating a custom application with "test" as part of the name, such as test, test1, test2, and so on. SD-WAN Manager permits you to create the application, but it does not deploy correctly to the devices in the network.

Configure custom applications using Cisco SD-WAN Manager

Configure custom applications using Cisco SD-WAN Manager.

Configure custom applications in Cisco SD-WAN Manager to enable advanced traffic identification AND policy enforcement for Cisco SD-AVC as part of Cisco SD-WAN Manager.

- Install Cisco SD-AVC as a component of Cisco SD-WAN Manager.

This task is relevant when you need to define AND manage custom applications for traffic classification in Cisco SD-WAN deployments using Cisco SD-WAN Manager. Custom applications allow for granular control AND visibility of network traffic beyond predefined application lists.

Perform the following steps to configure custom applications:

- Access the application catalog in Cisco SD-WAN Manager.

Before you begin

Install Cisco SD-AVC as a component of Cisco SD-WAN Manager. 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.

Follow these steps to configure custom applications using Cisco SD-WAN Manager:

Procedure

1. From Cisco SD-WAN Manager, select **Configuration** > **Application Catalog**.
2. Click **Applications**.

   The application catalog is displayed.
3. Select **Custom Application**.

   For Cisco IOS XE Catalyst SD-WAN Release 17.14.1a AND earlier, do the following:

   a. In Cisco SD-WAN Manager, select **Configuration** > **Policies**.

   b. Select **Centralized Policy**.

   c. Click **Custom Options** AND select **Centralized Policy** > **Lists**.

   d. Click **Custom Applications**.

   The custom application options are displayed.
4. Click **New Custom Application**.

   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. From Cisco IOS XE Catalyst SD-WAN Release 17.16.1a, Cisco SD-WAN Manager appends "-Custom" to the name. This ensures that the name does not conflict with other application types, such as those from Protocol Packs, OR cloud-sourced applications.
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	

Field	Description
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.
Ports	Enter one OR more ports OR port ranges, separated by space. Example: 30, 45-47
L4 Protocol	Select one of the following: TCP, UDP, TCP-UDP

The new custom application appears in the table of custom applications.

5. Click **Add**.

   To check the progress of creating the new custom application, click **Tasks** (clipboard icon). A panel opens, showing active AND completed processes.

   The custom application is added successfully.

After completing these steps, the custom application is available for use in policies AND traffic classification. See the custom application functionality section in the Custom Application Overview section.

Example custom application criteria

Refer to the example custom application criteria section for guidance.

Criteria	How to configure fields
Domain name	**Server Names**: Custom
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

Provides information on verifying custom applications and protocols.

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           PPDK LOCAL
custom_facebook                  3284           PPDK LOCAL
```

See show ip nbar protocol-id.

18 Service Insertion

Topics:

- Feature history for service insertion
- Service insertion
- Restrictions for service insertion
- Use cases for service insertion
- Service insertion configuration methods
- Configure service chain actions in a data policy
- Traffic steering to a service chain
- Configure path preference
- How service chains share across user VPNs
- How separate interfaces for transmitted and received traffic work
- Service chaining trusted and untrusted traffic
- How service chains work between two routers
- Fall back and restrict behavior configuration for service chain traffic
- How service chain interfaces work
- Service chaining with software defined cloud interconnect bring your own service

Provides a method for directing network traffic through specific service functions or appliances in a network path.

Service insertion enables network administrators to intercept and redirect traffic flows through intermediate service functions such as firewalls, load balancers, intrusion detection systems, or other network appliances without requiring changes to the underlying network topology.

Feature history for service insertion

Lists the features and their release information for service insertion functionality.

This reference provides a chronological overview of service insertion features introduced in Cisco Catalyst SD-WAN releases, documenting when capabilities were added and their key functionality.

Table 49: 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.

Service insertion

Describes placing network or security services into the path of specific data traffic within the Cisco Catalyst SD-WAN overlay fabric.

Service insertion, also known as *service chaining*, is a network capability that

- places one or more network or security services into the path of specific data traffic within the Cisco Catalyst SD-WAN overlay fabric
- routes traffic through services defined in a service chain according to service chain actions configured for a data policy, and
- can be deployed in any device and topology, including full mesh, HUB-spoke, and Cisco Catalyst SD-WAN Multi-Region Fabric (MRF).

Service insertion capabilities

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 features:

- 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

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.

Table 50: Service insertion capabilities

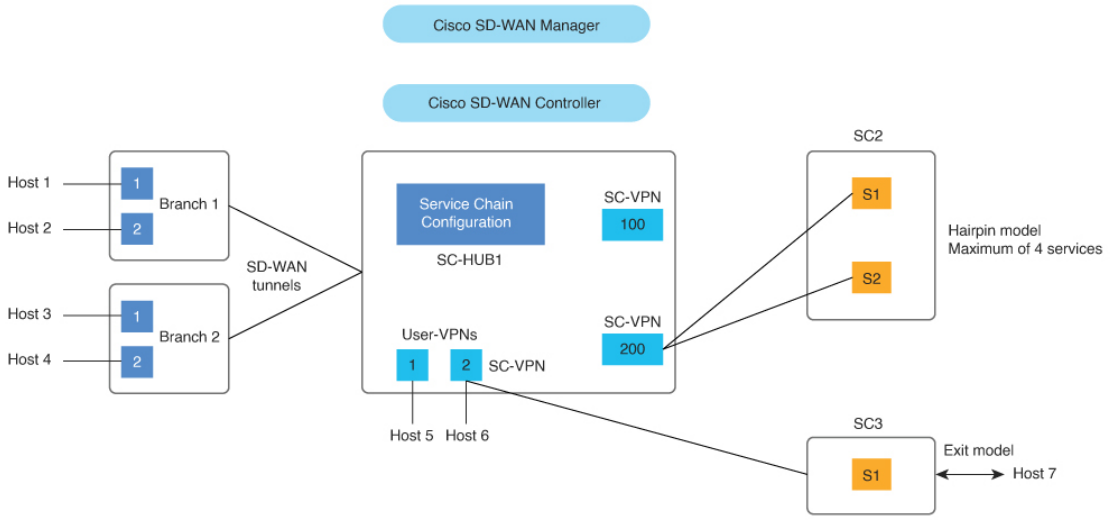
Capability	Releases before Cisco Cisco Catalyst SD-WAN Manager Release 20.13.1	Releases from Cisco Cisco Catalyst SD-WAN Manager Release 20.13.1
Multiple services in a chain	No native support	Native support
Traffic steering	Control policy	Control policy, data policy, interface ACL
Policy binding	Remote	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
Supported deployments	• Single-arm deployment • Single-arm deployment with multiple interfaces • Bump-in-the-wire deployment	• Single-arm deployment • Single-arm deployment with multiple interfaces • Bump-in-the-wire deployment

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 17: Service insertion key elements



Restrictions for service insertion

Outlines the configuration limitations and constraints that apply to service insertion implementation.

Adhere to the following limitations when configuring service insertion to ensure proper functionality and avoid configuration conflicts.

- 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

Lists common scenarios where service chaining provides network security and compliance benefits.

Service insertion use cases demonstrate how traffic can be directed through security services like firewalls to enhance network protection, ensure inter-VPN security, and meet regulatory compliance requirements.

- 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.

Service insertion configuration methods

Lists the available methods for configuring service insertion.

Use one of these methods to configure service insertion:

- Configure service insertion using a CLI template
- Configure service insertion using the workflow library

Configure service insertion using a CLI template

Configure service insertion using CLI templates to create and manage service chains with specified services and parameters.

This task configures service insertion using CLI templates to create service chains that define the sequence of services through which traffic is processed.

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.

Procedure

1. Create a service chain:

   Example

   ```
   service-chain chain-number
   ```

2. Configure a description for the service chain:

   Example

   ```
   service-chain-description description
   ```

3. Specify the services that are in the service chain and configure related options:

   Example

   ```
   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:

Example

```
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:

   Example

   ```
   service-chain-affect-bfd
   ```

6. Specify the name of the VPN that hosts all services in the service chain:

   Example

   ```
   service-chain-vrf vrf
   ```

7. (Optional, enabled by default) Enable endpoint tracking for services in the service chain:

   Example

   ```
   track-enable
   ```

8. (Optional, enabled by default) Enable the service chain, which makes it active on devices:

   Example

   ```
   service-chain-enable
   ```

The service insertion is configured using the CLI template with the specified service chain parameters and settings.

Configure service insertion using the workflow library

Configure service insertion by using the Workflow Library to create or modify service chains with UP to four service types.

Service insertion allows you to configure network services in a specific order to process traffic before it reaches its destination. Using the workflow library provides a guided approach to create and manage service chains effectively.

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

Procedure

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 tacker. 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.

The service chain is configured and ready to be used with data policies and traffic steering configurations.

What to do next

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

Configure service chain actions in a data policy to route traffic through a service chain.

This task enables you to route traffic through a service chain by configuring service chain actions for 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.

Procedure

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 51: 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: 1 to 65525, excluding 512. For details see the VRF range behavior change described here.
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.
Restrict	Check this option to cause packets to be dropped if the service chain goes down. f 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. This option is unchecked by default (traffic falls back to routing).

Service chain actions are configured in the data policy, enabling traffic to be routed through the specified service chain.

Traffic steering to a service chain

Describes how to direct network traffic to a service chain using three methods: control policy, data policy, or interface access control list for traffic steering implementation.

You can direct traffic to a service chain by using a control policy, data policy, or interface access control list.

How traffic steering using a control policy works

Describes how a control policy modifies Cisco Overlay Management Routes (vRoutes) to direct traffic to a service chain instead of the original destination.

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.

Summary

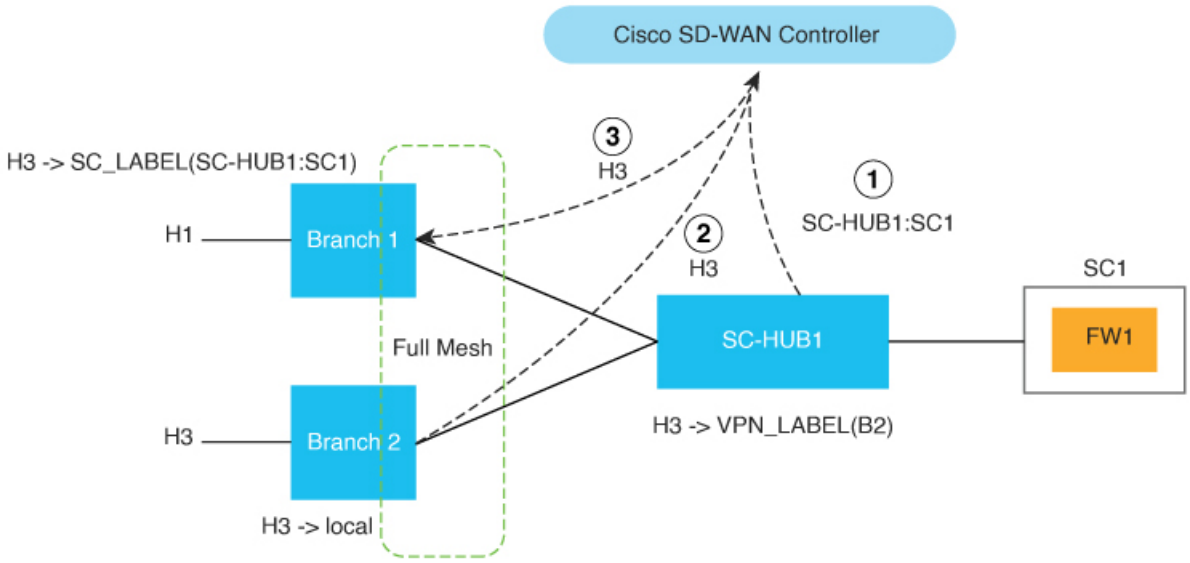
The key components involved in traffic steering using a control policy are:

- Control policy: Modifies vRoutes to redirect traffic to service chains
- Service chain hub (SC-HUB1): Advertises service chain routes
- Branch routers (B1, B2): Source and destination points for traffic flow
- Cisco SD-WAN Controller: Processes and advertises modified routes based on control policy
- Service chain (SC1): Traffic processing path applied to matching flows

The policy causes service chain 1 (SC1) to be applied to traffic that flows between H1 (host 1) and H3 (host 3) by setting SC1 as the next hop for H1 and H3 traffic routes.

Workflow

Figure 18: Traffic Steering with a Control Policy



These stages describe how traffic steering using a control policy works:

1. SC-HUB1 advertises the SC1 route to make the service chain available for traffic steering.
2. B2 advertises the H3 route to the Cisco SD-WAN Controller to establish the original destination path.
3. The control policy results in overriding the H3 route next hop to SC1 and the Cisco SD-WAN Controller advertises the modified H3 route to B1.

   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.

Note

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
```

How data policies steer traffic

Describes how data policies match traffic and operate in the context of source VPNs during forwarding to enable traffic steering and service chaining.

You can use a data policy to match traffic and operate in the context of source VPNs during forwarding.

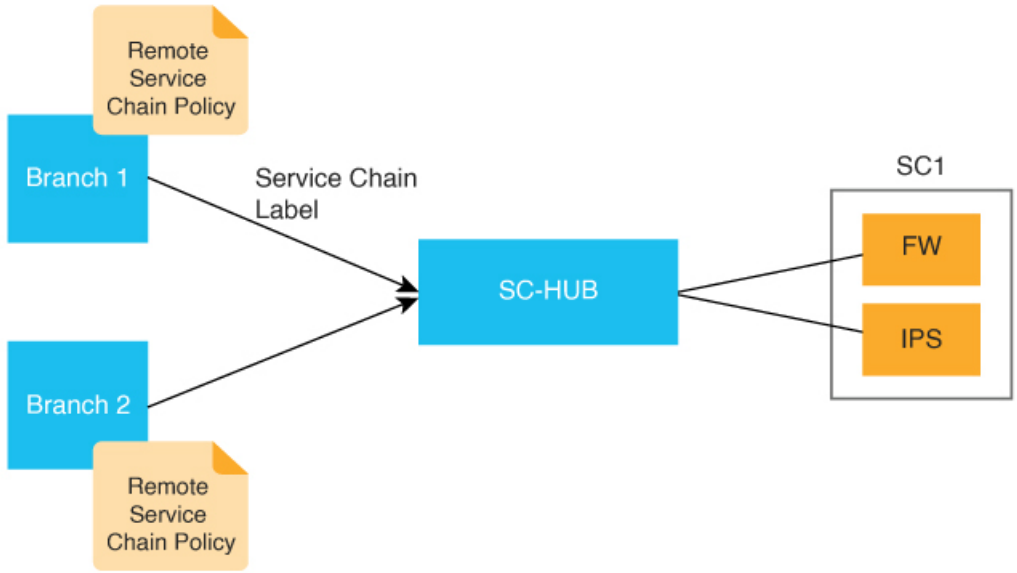
Summary

The key components involved in traffic steering using data policies are:

- Data policy: Matches traffic and defines forwarding actions based on specified criteria
- Match criteria: Specifies applications to be matched to source and destination IP address combinations
- Service chain: Defines the path for traffic processing through network services
- TLOC ranking: Specifies traffic path preference using TLOC ranking

Workflow

Figure 19: Traffic Steering with Traffic Service Chaining Intent Specified at a Remote Branch



These stages describe how data policies steer traffic:

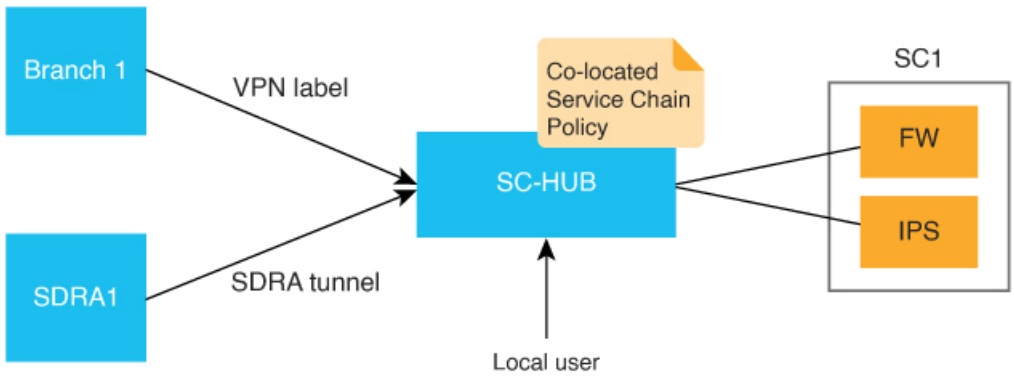
1. The system configures the data policy with match criteria and service chaining parameters for remote branch traffic steering.
   - **match criteria** specifies applications to be matched to source and destination IP address combinations
   - **restrict|fallback** configures restrict or fall back
   - **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
```

2. The system applies service chaining intent locally on the device to which the service chain is attached.

Figure 20: Traffic Steering with Traffic Service Chaining Intent Specified on a Local Device



In this configuration, **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

Describes how to use an interface access control list (ACL) to service chain traffic that is incoming or outgoing on a specified interface.

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.

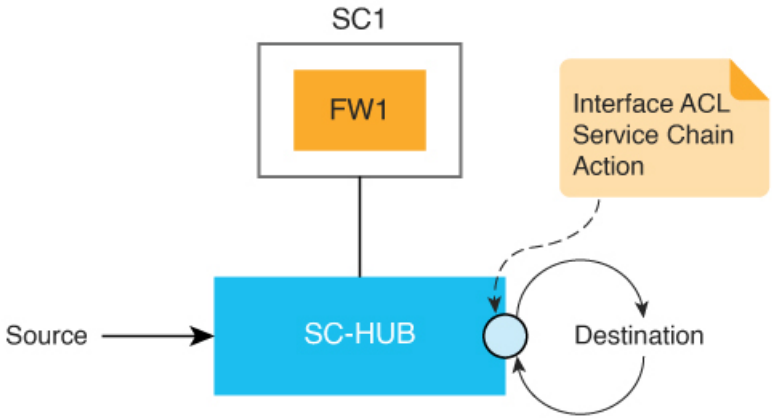
Summary

The key components involved in traffic steering using an interface ACL are:

- **Interface ACL**: Controls traffic flow on specified interfaces and directs traffic through service chains
- **Service chain**: Processes traffic according to configured policies and routing decisions
- **Traffic matching criteria**: Determines which traffic flows are subject to service chaining

Workflow

Figure 21: Traffic Steering with an ACL



Traffic steering using an interface ACL involves the following stages:

1. Configure the access control list with matching criteria and service chain action.

```
access-list list
  sequence number
    match criteria
    action accept
     set service-chain SC1 [vpn vpn] {restrict|fallback}
```

2. Apply the access list to the interface for incoming or outgoing traffic.

```
interface interface
 access-list list {in|out}
```

3. The system processes traffic according to the configured ACL rules and directs matching traffic through the specified service chain.

> **Note**
>
> Service chaining can be used in dual-site setups using Virtual Router Redundancy Protocol (VRRP) for redundancy. Ensure your service chaining policy does not block VRRP control traffic. Service chaining policies often have a default action to `drop` traffic that does not meet the match criteria. This can accidentally block VRRP control packets. To prevent this, set the default action in your service chaining policy to `accept` VRRP control traffic. This will ensure VRRP works correctly for redundancy.

Configure path preference

Configure TLOC preference or affinity preference to choose the preferred path for traffic to a service chain.

You configure path preference to control how traffic flows through service chains by directing traffic over specific TLOCs or by setting affinity preferences between branch sites and service chain hubs.

You can use TLOC preference or affinity preference to choose the preferred path for traffic to a service chain. If both TLOC preference and affinity preference are configured, the affinity preference is evaluated first, then the TLOC preference is evaluated.

Procedure

1. 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.

2. 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.

3. (Optional) Configure the following command to disable consideration of affinity in a data policy: **data-policy-ignore-affinity-metric**

Traffic to the service chain follows the configured path preferences, with affinity preference taking priority over TLOC preference when both are configured.

How service chains share across user VPNs

Describes how service chain VPNs can be shared across multiple user VPNs and how traffic between different VPNs can be service chained.

Service chain sharing allows a single service chain VPN to be utilized by traffic from multiple user VPNs without requiring additional configuration. When source and destination VPNs differ, route leaking is required between the source and destination VPN.

Summary

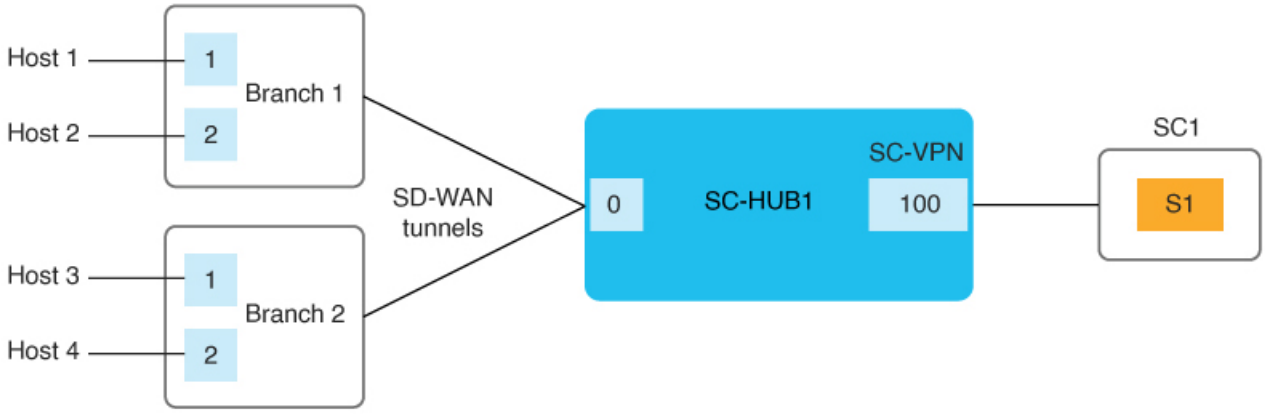
The key components involved in the service chain sharing process are:

- Service chain VPN: A dedicated VPN that provides service chaining functionality and can be shared across multiple user VPNs

- User VPNs: Individual VPNs that can utilize shared service chains for traffic processing
- Route leaking: Required mechanism when source and destination VPNs are different to enable traffic flow

Workflow

Figure 22: Service Chain Sharing Across VPNs



These stages describe how service chains share across user VPNs:

1. The service chain VPN is automatically made available to traffic from multiple user VPNs.
   - SC1 (service chain 1) is attached to VPN100 can be automatically shared by traffic in VPN1 (H1) and VPN2 (H4)
2. Traffic between different VPNs can be service chained through various VPN locations.
   - Traffic between VPN1 (H1) and VPN2 (H4) can be service chained in VPN1 or VPN2 or in a shared service chain (VPN100)

How separate interfaces for transmitted and received traffic work

Describes how separate interfaces can be configured for transmitted and received traffic through a service chain, where traffic flows are tracked independently.

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.

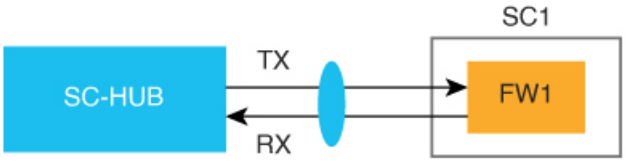
Summary

The key components involved in the separate interface configuration for transmitted and received traffic are:

- Service chain: Routes traffic through multiple network services with independent interfaces for each direction
- Transmitted traffic interface: Handles outgoing traffic flow through the service chain
- Received traffic interface: Handles incoming traffic flow through the service chain
- Traffic tracking system: Monitors transmitted and received traffic independently

Workflow

Figure 23: Separate Interfaces for Transmitted and Received Traffic



Service chaining trusted and untrusted traffic

Describes how the system routes trusted and untrusted traffic through different high availability pairs in service chaining configurations.

Minimum supported releases: Cisco IOS XE Catalyst SD-WAN Release 17.14.1a, Cisco Catalyst SD-WAN Manager Release 20.14.1

In service chaining configurations, the system can separate trusted and untrusted traffic flows to maintain security and performance isolation. The default trust-posture of a packet is trusted unless explicitly configured otherwise.

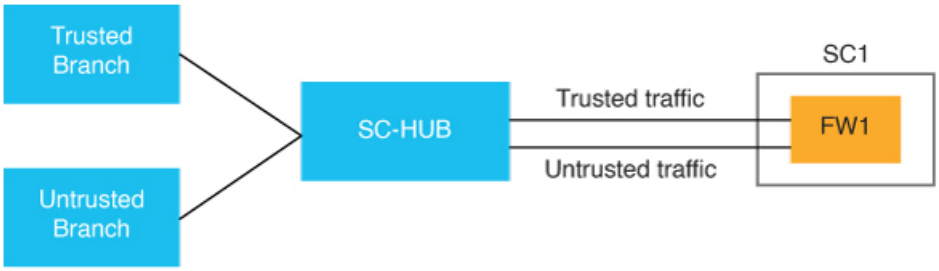
Summary

The key components involved in service chaining trusted and untrusted traffic are:

- Data policy: Uses the **set attribute trust-posture untrusted action** command to mark packets as trusted or untrusted
- Trusted high availability pair: Receives and processes trusted traffic flows
- Untrusted high availability pair: Receives and processes untrusted traffic flows
- Service chain configuration: Defines the service transport HA pairs and their trust-posture attributes

Workflow

Figure 24: Trusted and Untrusted Traffic



These stages describe how the system processes and routes trusted and untrusted traffic through service chains:

1. The system evaluates incoming packets and determines their trust-posture based on data policy configuration. By default, packets are marked as trusted unless explicitly configured otherwise.
2. The system routes trusted traffic to the configured trusted high availability pair within the service chain.
3. The system routes untrusted traffic to the configured untrusted high availability pair within the service chain.
4. The service chain configuration defines the trust-posture attribute for each service transport HA pair using the following example configuration:

```
service-chain SC1
    service netsvc1
        sequence 10
```

```
        service-transport-ha-pair 1
            attribute trust-posture {trusted|untrusted}
```

How service chains work between two routers

Describes how service chains operate when the transmitting router differs from the receiving router, requiring identical configuration on each device for intra-VPN traffic processing.

When the router that transmits traffic to a service chain differs from the router that receives traffic from the service chain, a specific configuration approach is required. This scenario applies to intra-VPN traffic only and involves service chains with a single service.

Summary

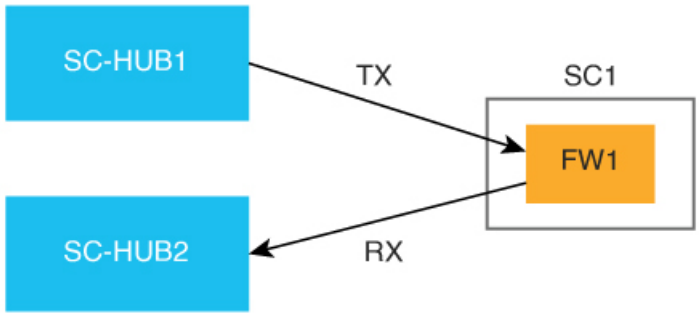
The key components involved in the service chain process between two routers are:

- Transmitting router: Sends traffic to the service chain for processing
- Receiving router: Receives processed traffic from the service chain
- Service chain: Contains only one service and processes intra-VPN traffic

The process requires identical service chain configuration on both routers to ensure proper traffic flow and processing.

Workflow

Figure 25: Service Chain Between Two Router



These stages describe how service chains operate between two routers:

1. The transmitting router sends intra-VPN traffic to the configured service chain for processing.
2. The service chain processes the traffic using its single configured service.
3. The receiving router receives the processed traffic from the service chain and continues routing operations.

Fall back and restrict behavior configuration for service chain traffic

Provides configuration options for fall back or restrict behavior when traffic travels 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.

How service chain interfaces work

Describes the interfaces used for attaching services in a service chain to a router and the attachment methods available.

The services in a service chain must be in a single VPN, called a *service chain VPN*, or *SC-VPN*.

Summary

The key components involved in attaching services in a service chain to a router are:

- Cisco Catalyst SD-WAN router: Acts as the attachment point for service chain services
- Service chain VPN (SC-VPN): Contains all services in the service chain
- Interface types: Provide various attachment methods including IPv4, IPv6, dual stack, or tunnel interfaces

Workflow

Figure 26: Attaching Services to a Router



These stages describe how services in a service chain attach to a router through different interface types:

1. The services can be attached through any combination of the following interface types:

1	IPv4 attachment
2	IPv6 attachment
3	Dual stack attachment
4	Tunnel attachment

Service chaining with software defined cloud interconnect bring your own service

Explains how BYOS functionality enables centralized service inspection by connecting service chains to SDCI gateways deployed in the middle mile network.

Service chaining with software defined cloud interconnect bring your own service is a functionality that

- 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
- enables the seamless integration of external services with the SDCI infrastructure, and
- enforces colocated data policies, also known as centralized data policies, on these gateways within the middle mile network for selective data traffic inspection.

Network architecture

The Software Defined Cloud Interconnect (SDCI) establishes connections between branch sites and the cloud through network service providers, including Megaport and Equinix. 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.

19 Service Insertion for Equinix

Topics:

- Feature history for service insertion for Equinix
- Service insertion for Equinix
- Requirements for service insertion for Equinix
- Service insertion configuration workflow for Equinix
- Configuration verification for Equinix service insertion

Outlines the implementation and configuration of service insertion capabilities specifically designed for Equinix network environments and infrastructure requirements.

Feature history for service insertion for Equinix

Lists the development milestones and release information for service insertion for Equinix, including feature description.

Service insertion for Equinix enables deployment of Palo Alto Networks firewall on Equinix infrastructure and attachment of service chains to Equinix interconnect gateway through the **Workflow Library** in Cisco SD-WAN Manager.

This table describes the developments of this feature, by release.

Table 52: Feature history

Feature Name	Release Information	Description
Service Insertion for Equinix	Cisco IOS XE Catalyst SD-WAN Release 17.15.1a Cisco Catalyst SD-WAN Manager Release 20.15.1	With this feature, you can deploy Palo Alto Networks firewall on Equinix and attach a service chain to Equinix interconnect gateway from the **Workflow Library** in Cisco SD-WAN Manager.

Service insertion for Equinix

Describes a feature that enables defining, creating, instantiating, and deploying service chains in Equinix through Cisco SD-WAN Manager.

Service insertion for Equinix is a feature that

- enables the use of Cisco SD-WAN Manager to define, create, instantiate, and deploy service chains in Equinix, and
- supports creating service chains that include a single standalone or stateful Palo Alto Networks Firewall service instance per service chain.

Service insertion workflow automation

The Cisco SD-WAN Manager service insertion workflow automates configurations on a Cisco IOS XE Catalyst SD-WAN device. Based on the selection of the service and the service type, the service instance creates a configuration group for service insertion which has the details of the service instances defined in the workflow.

Requirements for service insertion for Equinix

Outlines the prerequisites required before configuring service insertion for Equinix environments.

Ensure that you meet all prerequisites before configuring service insertion for Equinix to avoid deployment issues and ensure successful integration.

- Ensure that you have an active Equinix account. If you dont already have an account, you can create an account on the Equinix portal. Refer to the *New User Equinix Fabric Portal Access* documentation from Equinix.
- Ensure that you have an active Equinix billing account. If you dont already have one, you can create billing accounts for each region in which you would like to deploy an Interconnect Gateway using this account. Refer to the *Billing Account Management* documentation from Equinix.
- Ensure that the Equinix account is associated to the Cisco SD-WAN Manager. For more information, see Associate Equinix Account with Cisco SD-WAN Manager.

- Deploy the Equinix Interconnect Gateway to a configuration group. For more information, see Create Interconnect Gateway at an Equinix Location.
- Purchase the required Palo Alto Networks firewall licenses.

Service insertion configuration workflow for Equinix

Lists the workflow steps to configure service insertion for Equinix.

The workflow to configure service insertion for Equinix includes defining and configuring service chains, instantiating service chains, attaching service chains to Cisco SD-WAN devices, and configuring service chain actions for data policies.

1. Define and configure service chain: Creates a service chain definition. The service chain definition comprises of a service type and the order of the service. After you provide the details in the workflow, Cisco SD-WAN Manager creates a configuration group. You cannot edit the parameters in the configuration group after it's creation.
2. Instantiate service chain: Instantiates a service chain by deploying an instance of the service in Equinix.
3. Attach a service chain to the Cisco SD-WAN Device: Attach the service chain to Equinix interconnect gateway.

   Detach a service chain instance: Detach the service chain instance from Equinix interconnect gateway
4. Configure service chain actions for a data policy to route traffic through a service chain. See Configure Service Chain Actions in a Data Policy in *Cisco Catalyst SD-WAN Policies Configuration Guide, Cisco IOS XE Catalyst SD-WAN Release 26.x*.

Define and configure service chain

Configure a service chain definition workflow in the Equinix environment using firewall services.

Define and configure a service chain to establish network security services using Palo Alto Networks firewall in the Equinix environment.

Service chains enable you to deploy and manage network security services in cloud environments. This task focuses on configuring firewall services within the Equinix platform through the workflows interface.

Before you begin

Follow these steps to define and configure service chain:

Procedure

1. In the Cisco SD-WAN Manager menu, click **Workflows** > **Workflows Library** > **Define and Configure Service Chain**.
2. Follow the on-screen instructions to complete the service chain definition workflow.

   Following are some of the parameters in the workflow:
   - In **Select Environment** page, choose **Equinix** as the environment for the service chain definition.
   - Choose only firewall as the service type. For each service chain intance you can choose only one Palo Alto Networks firewall.
   - The **Size**, **Flavor**, and **SW-Version** in the workflow refers to the Palo Alto Networks firewall that you choose.
   - The **sshPublickey Value** is retrieved from the Equinix account that you choose.
   - Enter the management plane bandwidth (Mbps) between the router and the service in the **Management Plane** field.

- The name entered in the **User Name** field is the user name of the Palo Alto Networks firewall credentials.
- Choose **Stateful** or **Standalone** type of HA mode for the service. If you choose **Stateful** HA mode, ensure that you have the appropriate active licenses to proceed with the instantiation of the service chain.

The service chain definition workflow is completed and the firewall service is configured in the Equinix environment with the specified parameters.

Instantiate service chain

Configure a service chain instantiation using the Workflows Library to deploy network services.

This task enables you to instantiate a service chain through the Cisco SD-WAN Manager Workflows Library, allowing you to deploy and configure network services such as Palo Alto Networks firewall services.

Use this workflow when you need to deploy network service chains in your environment. The workflow provides guided steps to configure service parameters and deploy services through the Equinix platform integration.

Procedure

---

1. In the Cisco SD-WAN Manager menu, click **Workflows** > **Workflows Library** > **Instantiate Service Chain**.
2. Follow the on-screen instructions to instantiate the service chain.

   Following are some of the parameters of the workflow:

   - Ensure that you choose an active Equinix billing account for the service insertion workflow.
   - Ensure that you provide Palo Alto Networks firewall parameters based on your requirements. This workflow only instantiates the Palo Alto Networks firewall service. For any Palo Alto Networks firewall security policies and lifecycle management see the Palo Alto Networks firewall documentation.
   - The **Size**, **Flavor**, and **SW-Version** are retrieved from the Equinix portal based on the selected Equinix account and the Palo Alto Networks firewall you choose in the Instantiate Service Chain workflow.

The service chain is successfully instantiated and deployed in your network environment with the configured parameters.

Attach a service chain to the Cisco SD-WAN device

Configure a service chain attachment to a Cisco SD-WAN router using the workflows library or service chain instances.

This task attaches a service chain to a Cisco SD-WAN device to enable service insertion functionality and ensure proper traffic routing through network services.

Service chains enable traffic to be routed through network services before reaching its destination. When attaching service chains to Cisco SD-WAN devices, you can select only one interconnect gateway at a time and should attach the service chain to the appropriate router without needing to attach it to branch routers.

Before you begin

- 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.

Follow these steps to attach a service chain to the Cisco SD-WAN device:

Procedure

1. In the Cisco SD-WAN Manager menu, click **Workflows** > **Workflows Library** > **Attach Service Chain to Cisco SD-WAN Router**.
2. Follow the on-screen instructions to attach the service chain to a Cisco IOS XE Catalyst SD-WAN device.
   - You can select only one interconnect gateway at a time to attach a service chain.
   - Attach the service chain to the appropriate Cisco IOS XE Catalyst SD-WAN device. You dont need to attach the service chain to the branch routers.

> **Note**
>
> Alternatively, attach the Cisco IOS XE Catalyst SD-WAN device to a service chain using **Configuration** > **Service Insertion** > **Service Chain Instances**. Click **...** next to the instance name you wish to attach a device and click **Attach**. This takes you to the service chain attachment workflow under **Workflows** > **Workflows Library** > **Attach Service Chain to Cisco SD-WAN Router**.

The service chain is successfully attached to the Cisco SD-WAN device, enabling service insertion functionality for traffic routing through network services.

Detach a service chain instance

Configure detachment of a device from a service chain instance to remove the service chain configuration from the device.

Detach a device from a service chain instance to remove the service chain configuration and stop applying the configured services to traffic passing through that device.

Use this task when you need to remove a device from a service chain instance, effectively stopping the application of the service chain's configured services on that specific device.

Procedure

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Service Insertion**.
2. Click **Service Chain Instances** to view the list of service chain instances and the devices attached to each instance.
3. Click **...** adjacent to the instance that you wish to detach the device and choose **Detach**.
4. In the confirmation dialog box, click **Detach**.

The device is detached from the service chain instance, and the service chain configuration is removed from the device.

Configuration verification for Equinix service insertion

Provides methods to verify service chain configurations after instantiation in Cisco SD-WAN.

- After you instantiate the service chain, you can view the service chain instances in **Configuration** > **Service Insertion**. This page displays a list of service chain instances along with the Cisco IOS XE Catalyst SD-WAN device attached to each service chain instance.

- To view the details of the device that is attached to the service chain instance, click the device in the **Attached to** column. You can use the information about service TX and RX interfaces and the IP addresses on the Cisco SD-WAN Manager for the configurations on the Palo Alto Networks Firewall. Configure the service TX IP address on the service TX interface, and the service RX IP address on the service RX interface on the Palo Alto Networks Firewall.
- To view details such as service type, password, IP address, and so on, of the service used in the service chain instance, click the **Number of Services** option.

20 Service Chaining

Topics:

- Service chaining in Cisco SD-WAN
- Configure service chaining
- Monitor service chaining

Service chaining in Cisco SD-WAN

Explains service chaining in Cisco Catalyst SD-WAN, including service advertisement, policy-based traffic steering, routing behavior, health tracking, and limitations.

Service chaining in Cisco Catalyst SD-WAN is a centralized orchestration capability that

- advertises available services using OMP service routes with defined service identifiers.
- steers traffic through services by modifying OMP next-hop, TLOC, and labels via policy.
- tracks service availability to prevent routing traffic to unavailable service devices.

Services in the network

Services such as firewall, load balancer, and intrusion detection and prevention (IDP) often run within a virtualized environment and may physically be centralized in one location or in several locations for redundancy. Services may be internal, cloud based, or external subscriptions. Networks must 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 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 requires the configuration of policy-based, weighted load distribution to multiple firewalls.

Reasons to reroute traffic through 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.

In a network that consists of multiple VPNs, each representing a function or an organization, traffic between VPNs must traverse a service, such as a firewall, or a chain of services.

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, operators reroute traffic flow by provisioning every routing node from the source to the service node to the systems beyond the service node with a policy route. Operators manually configure each node or use a provisioning tool that performs the configuration for each node on behalf of the operator. This process is operationally complex to provision, maintain, and troubleshoot.

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.

Service provisioning and chaining in Cisco Catalyst SD-WAN

Explains how Cisco Catalyst SD-WAN provisions services and implements service chaining to direct traffic through network services before reaching its destination.

This section explains how Cisco Catalyst SD-WAN provisions services and implements service chaining to direct traffic through network services before reaching its destination.

Provisioning services in the Cisco Catalyst SD-WAN overlay network

In the Cisco Catalyst SD-WAN solution, the network operator enables and orchestrates all service chaining from a central controller, the Cisco SD-WAN Controller. The solution does not require configuration or provisioning on the devices.

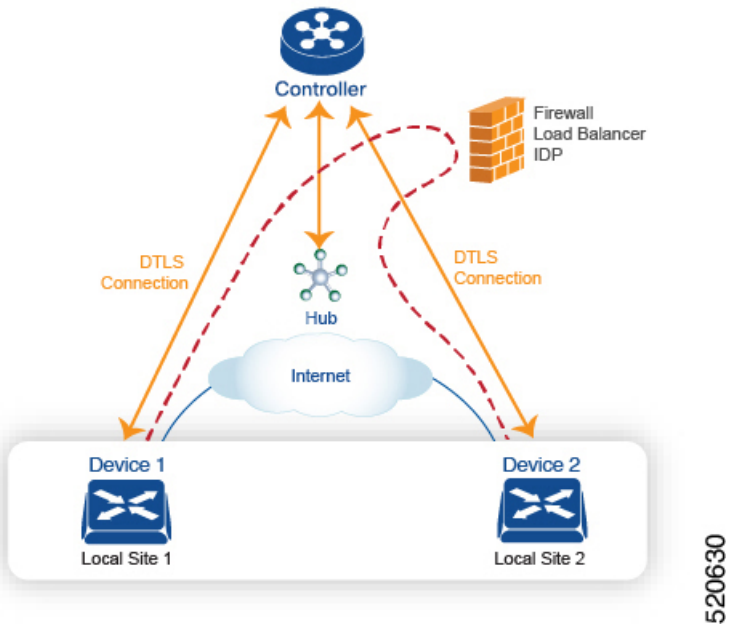
The general flow of service chaining in a Cisco Catalyst SD-WAN network is as follows:

- 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 example illustrates how service chaining works in the Cisco Catalyst SD-WAN solution.

The network has a centralized hub router connected to two branches, each with a device. The standard network design implements a control policy so that all traffic from branch site 1 to branch site 2 travels through the hub router. A firewall device sits behind the hub router.

When the design requires all traffic from site 1 to site 2 to be processed by the firewall, traffic from the device at site 1 still flows to the hub router. Instead of sending the traffic directly to site 2, the hub router redirects it to the firewall device. After the firewall completes its processing, it returns the cleared traffic to the hub, which then forwards it to the device at site 2.



Service route SAFI

Hub and local branch devices advertise the services available in their networks to the Cisco SD-WAN Controllers in the domain. They use service routes to do this. OMP sends these service routes 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. They do not propagate these routes back to the devices.

Each service route SAFI has the following attributes:

- VPN ID (vpn-id): Identifies the VPN to which the service belongs.
- Service ID (svc-id): Identifies the service that the service node advertises. Cisco Catalyst SD-WAN software defines these predefined services: FW for firewall (svc-id 1), IDS for Intrusion Detection Systems (svc-id 2), IDP for Identity Providers (svc-id 3), and netsvc1 through netsvc4 for custom services (svc-id 4-7).
- Label: For traffic that must traverse a service, the Cisco SD-WAN Controller replaces the label in the OMP route with the service label to direct traffic to that service.
- Originator ID (originator-id): The IP address of the service node that advertises the service.
- TLOC: The transport location address of the device that hosts the service.
- Path ID (path-id): An identifier of the OMP path.

- 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.

Restrictions for service chaining

Ensure to read the key restrictions for service chaining to ensure proper operation and performance on Cisco IOS XE Catalyst SD-WAN devices.

- Cisco IOS XE Catalyst SD-WAN devices do not support service insertion over a tunnel interface.
- You cannot use control policy-based service-chain actions on a locally hosted service chain.
- You cannot configure service-chain and AppQoE on the same device, regardless of whether you use data-policy or control-policy actions.

Configure service chaining

Provides the workflow and available methods for configuring service chaining using templates and CLI commands.

This section describes the workflow and available methods for configuring service chaining for a device managed by Cisco Catalyst SD-WAN.

Workflow for service chaining configuration

1. Access service devices through a specific VRF. In the VPN template that corresponds to the VRF for a service device, configure service chaining and specify 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.

Use one of these methods to configure service chaining:

1. CLI commands
2. Templates

Configure service chaining for Cisco IOS XE Catalyst SD-WAN devices using CLI commands

This section explains configuring service chaining on Cisco IOS XE Catalyst SD-WAN devices using CLI.

This section explains how to configure service chaining on Cisco IOS XE Catalyst SD-WAN devices using CLI commands and compares the process with Cisco SD-WAN Manager configuration.

Service chaining configuration comparison for Cisco IOS XE Catalyst SD-WAN devices using CLI commands

This 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.

Table 53: CLI and Cisco SD-WAN Manager service chaining configuration comparison

CLI (Cisco IOS XE Catalyst SD-WAN device)	Cisco SD-WAN Manager
`service firewall vrf 10`	In Cisco SD-WAN Manager, configure service insertion in the VPN template for a specific VRF—VRF 10 in this example. Select the service type from the drop-down list; firewall in this example.
`no track-enable`	When adding a service in the VPN template **Service**, select **On** or **Off** for **Tracking**.
Note Default: enabled	
`ipv4 address 10.0.2.1 10.0.2.2`	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
```

Configure service chaining using templates

Configure service chaining for a device using templates. This task enables you to define and add service devices to a VPN template, specifying service type, IP address, and tracking options.

Configure service chaining for a device by creating and modifying a VPN template. This allows you to define service devices and their properties for integration into the network.

Procedure

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

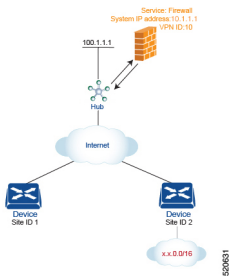
   Maximum number of services: 8
4. Click **Add**. The service appears in the table of configured services.

Service chain configuration examples

Provides configuration examples for service chaining, including routing intersite and inter-VPN traffic through service devices using control and data policies.

Service chaining control policies direct data traffic to service devices located in various places in the network before the traffic is delivered to its destination. This section provides configuration examples for implementing service chaining using centralized control policies on the Cisco SD-WAN Controller, and configuring service devices on the device collocated in the same site. To ensure services are advertised to the Cisco SD-WAN Controller, the IP address of the service device must resolve locally.

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, all traffic traveling from the device at Site 1 to the device at Site 2 is routed through a firewall service that sits behind a hub (system IP address 100.1.1.1). All devices are in the same VPN.

For this scenario, configure the following:

- On the hub router, configure the IP address of the firewall device.
- On the Cisco SD-WAN Controller, configure a control policy that redirects traffic destined from Site 1 to Site 2 through the firewall service.
- On the Cisco SD-WAN Controller, 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 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. 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
```

3. This policy configuration performs the following actions:

   It creates a site list called **firewall-sites** that the apply-policy command references and that enumerates all the sites to which this policy applies. If you later want to scale this policy so that all traffic destined for Site 2 from other sites also passes through the firewall first, you only need to 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.

   It defines a control policy named **firewall-service** with one sequence element and the following conditions:

   - It matches routes destined for Site 2.
   - If a match occurs, it accepts the route and redirects it to the firewall service provided by the hub router, which is located in VPN 10.
   - It accepts all nonmatching traffic, that is, all traffic not destined for Site 2.

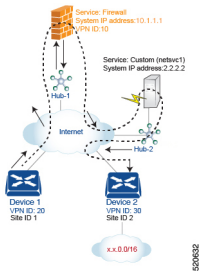
   It applies the policy to the sites listed in firewall-list, that is, to Site 1. The Cisco SD-WAN Validator applies the policy in the outbound direction, that is, on routes that it redistributes to Site 1. In these routes:

   - It changes the TLOC from Site 2's TLOC to the hub router's TLOC. The Cisco SD-WAN Controller learned this TLOC from the service route received from the hub router. Because of this TLOC change, traffic destined for Site 2 is directed to the hub router.
   - It changes the label 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 the traffic to the address 10.1.1.1, which is the system IP address of the firewall. After the firewall finishes processing the traffic, it returns the traffic to the hub router, and the hub router then forwards it to its final destination, 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 it reaches its destination. In this example, the configuration routes traffic from one VPN to another through services located in a third VPN. The services sit behind different hub routers. Specifically, the configuration routes all traffic from device-1 in VPN 20 that is destined for prefix x.x.0.0/16 in VPN 30 on device-2 first through the firewall behind Hub-1 and then through the custom service netsvc1 behind Hub-2 before sending it 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 netsvc1 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 netsvc1 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
```

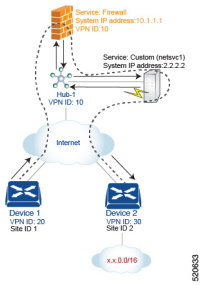
```
    sequence 10
      match route
        vpn 30
        site-id 2
      action accept
        set service netsvc1
    default-action accept
apply-policy
  site-list custom-service control-policy netsvc1-service out
```

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 connected to the same node, meaning both services sit behind the same device, you use a combination of control policy and data policy to create the desired service chain. This example is similar to the previous one, but it places a firewall and a custom service (netsvc-1) behind a single hub router. In this scenario, the configuration routes all data traffic from device-1 in VPN 20 that is destined for prefix x.x.0.0/16 on device-2 in VPN 30 first through the firewall at Hub-1, then through the custom service netsvc1 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 netsvc1 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 one for the custom service netsvc1. Both service routes include the TLOC of the Hub-1 router and a service label that identifies the service type. The firewall service uses label svc-id-1, and the custom service uses label 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, and 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.

This example shows a policy that can fail if the send-path-limit is lower than the 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
!
!
```

This example shows the same policy after applying the recommended fix:

```
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

Provides information about monitoring service chaining paths and configured services using Cisco SD-WAN Manager and CLI on hub and spoke devices.

You can monitor different aspects of service chaining on hub and spoke devices.

View configured services on a hub device

From the Cisco SD-WAN Manager menu, go to **Monitor** > **Devices** > **Hub-device** > **Real Time** and select **OMP Services** under Device Options to view the configured services.

In Cisco vManage Release 20.6.x and earlier, go to **Monitor** > **Network** > **Hub-device** > **Real Time**,and select **OMP Services** under Device Options to view the configured services.

View service chain path on a spoke device

From Cisco SD-WAN Manager, go to **Monitor** > **Devices** > **Spoke-device** > **Troubleshooting** > **Connectivity** > **Trace Route**, then enter the destination IP, VPN, and source interface to view the service chain path.

In Cisco vManage Release 20.6.x and earlier, go to **Monitor** > **Network** > **Spoke-device** > **Troubleshooting** > **Connectivity** > **Trace Route**, then enter the destination IP, VPN, and source interface to view the service chain path.

Using the CLI:

Use the **traceroute** command. For information, see the .

> **Note**
>
> Configuring a service device to operate as part of the service chain is called service insertion.

Viewing the service chain path using the CLI

Use the **traceroute** command. For more information, refer to Cisco Catalyst SD-WAN Command Reference.

Example: view a service chain path between two spoke devices

This example shows how to view the path between two spokes before and after adding a service chain between them by using Cisco SD-WAN Manager or the CLI. It presents a scenario with two spoke devices, a hub device, and a service device that provides firewall service. It also 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
 !
!
```

Without service insertion:

At this point, no service insertion policy is configured. When you run a traceroute from Spoke 1 to Spoke 2 (10.0.4.1), the output shows a direct 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, the Traceroute page in Cisco SD-WAN Manager shows a direct path from Spoke 1 to Spoke 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 you configure service insertion, running a traceroute from Spoke 1 (10.0.3.1) to Spoke 2 (10.0.4.1) shows the updated 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, the Traceroute page in Cisco SD-WAN Manager shows each step of the path from Spoke 1 to Spoke 2 through the hub and the firewall service device.