



SR-TE Per-Flow (Class) ODN and Automated Steering (PCE Delegated)

This chapter describes how Segment Routing Traffic Engineering (SR-TE) works with the Per-flow policy (PFP) On-Demand Next-hop (ODN) and auto steering (Per flow ODN/AS) mechanism. This chapter contains the following sections:

- [Feature Information for SR-TE Per-Flow \(Class\) ODN and Automated Steering \(PCE Delegated\)](#), on page 1
- [Information About SR-TE Per-Flow \(Class\) ODN and Automated Steering \(PCE Delegated\)](#), on page 2
- [BGP Color Extended Community and VRF Prefix Coloring](#), on page 3
- [Support for PFP with RIB Path](#), on page 6
- [Configuring SR-TE Per-Flow Class \(ODN\) and Automated Steering \(PCE Delegated\)](#), on page 7
- [Verifying SR-TE Per-Flow Class \(ODN\) and Automated Steering \(PCE Delegated\)](#), on page 9

Feature Information for SR-TE Per-Flow (Class) ODN and Automated Steering (PCE Delegated)

Table 1: Feature History

Feature Name	Release	Description
Support for PFP with RIB Path	Cisco IOS XE 17.9.1a	This feature enables you to configure forwarding class in a per-flow policy using the Routing Information Base (RIB) path option. Instead of configuring a per-destination policy, the RIB option uses the IGP shortest path to the policy destination.

Feature Name	Release	Description
Attaching Extended Color Communities to BGP VRF	Cisco IOS XE 17.7.1a Cisco IOS XE 17.11.1a	<p>This feature introduces new methods of attaching extended color communities to a prefix. A color community is an indicator of the bandwidth or latency level of the traffic being sent to the prefix. The following are the new ways of attaching them to the prefix:</p> <ul style="list-style-type: none"> • VRF export coloring • VRF import coloring • Route redistribution coloring in BGP • Neighbor inbound coloring <p>From Cisco IOS XE 17.11.1a, this feature is extended to the following platforms:</p> <ul style="list-style-type: none"> • Cisco Catalyst 8300 Series Edge Platforms • Cisco Catalyst 8500 Series Edge Platforms • Cisco Catalyst 8000V Edge Software
SR-TE Per-Flow (Class) ODN and Automated Steering (PCE Delegated)	Cisco IOS XE Amsterdam 17.4	This feature lets you steer traffic with SR-TE PFP based on the QoS markings on the packets. The traffic is then switched onto the appropriate path based on the forward classes of the packet.

Information About SR-TE Per-Flow (Class) ODN and Automated Steering (PCE Delegated)

The Segment Routing-Traffic Engineering (SR-TE) Per-Flow policy (PFP) On-Demand Next-Hop (ODN) with auto-steering (Per-Flow ODN/AS) is a mechanism that allows the steering of traffic on a segment routing policy based on the attributes of the packets. SR-TE PFP ODN with auto steering (Per flow ODN/AS) is a mechanism that allows the steering of traffic on an SR policy based on the attributes of the packets. Packets are classified using Cisco's Modular QoS CLI (MQC) framework and then marked using internal tags known as Forward Classes (FCs). A PFP is then used to route the marked packets based on the mappings between an FC and its corresponding path. This means that the traffic is steered based on its QoS markings and switched to the appropriate path based on the FC of the packet.

A PFP is identified by $\langle color, endpoint \rangle$. It is configured with a per-flow forwarding class table with up to eight entries, with each entry indexed by an FC and pointing to a Per Destination Policy (PDP).



Note The following features are supported:

- 250 PFP+PDP (Combination)
- 6 PE and 6 VPE
- 10k VPNV4 prefix limit
- L3VPN Inter AS Option B for SR PFP
- IPv6 over PFP

Restrictions for SR-TE Per-Flow (Class) ODN and Automated Steering (PCE Delegated)

- Dynamic change in the Quality of Service (QoS) policy is not supported.
- PIC core over SR-TE tunnel PIC edge is not supported.
- VPLS over SR-TE is not supported.
- Configure the set forward class to 0 to take the default path for nonforward class.
- BGP Labeled Unicast (BGP-LU) (RFC 3107) is not supported for SR ODN PFP Auto Steering.
- L2VPN over PFP tunnels is not supported.
- Performance measurement over PFP is not supported.
- MPLS ping or trace route over PFP is not supported.
- Auto route announcement over PFP or PDP is not supported.
- PIC is not supported over PFP.

BGP Color Extended Community and VRF Prefix Coloring

In the Segment Routing Traffic Engineering mechanism, the prefix that needs an SR-TE routing path is associated with a color-extended community (an attribute that assigns color to the prefixes). Currently, BGP has the capability to attach the color-extended community based only on the neighbor command routemap outbound configuration. To color the prefixes based on attributes such as Source-VRF, Destination-VRF, CE-neighbor, and Source protocol, the following ways of attaching color are introduced:

- VRF Export Coloring
- VRF Import Coloring
- Route Redistribution Coloring into BGP

- Neighbor In-bound Coloring

Additionally, in Cisco IOS XE releases prior to 17.7.1a, any new color-extended community attached to the prefix replaces the existing color-extended community available in the prefix. To be able to add the new color-extended community to the existing list of color-extended communities instead of replacing, the keyword **additive** is added to the **route-map** command as part of Cisco IOS XE 17.7.1a:

```
route-map SRTE-color-map permit
set extcommunity color < 1-4294967295> [additive]
```



Note When a BGP update is received with multiple color-extended communities, the highest color value in the list is used for SR policy creation, and the binding SID corresponding to the SR policy is used as the routing path for the received BGP path. If the SR policy corresponding to the highest color is not available, BGP uses the interface as the routing path for the update.

Supported Platforms

From Cisco IOS XE 17.7.1a, this feature is supported on:

- Cisco ASR 1000 Series platforms

From Cisco IOS XE 17.11.1a, this feature is supported on:

- Cisco Catalyst 8300 Series Edge platforms
- Cisco Catalyst 8500 Series Edge platforms
- Cisco Catalyst 8000V Edge software

Attaching a Color-Extended Community

The following ways of attaching color-extended communities are available:

- **VRF Export Coloring:** The following configuration attaches a color extended community to the VPN prefix as per the export route map color-extended community associated with the VRF. This enables the association of the color-extended community based on the source VRF of the VPN prefix:

```
route-map SRTE-color-map permit
set extcommunity color < 1-4294967295> [additive]
vrf def SRTE-VRF
rd 1:1
!
address-family ipv4
export map SRTE-color-map
exit-address-family
!
address-family ipv6
export map SRTE-color-map
exit-address-family
```

- **VRF Import Coloring:** The following configuration attaches a color-extended community to an imported VRF prefix as per the import route map color-extended community associated with the VRF. This enables the attachment of the color-extended community to a prefix based on the VRF the prefix is imported to:

```

route-map SRTE-color-map permit
set extcommunity color < 1-4294967295> [additive]
vrf def SRTE-VRF
rd 1:1
!
address-family ipv4
import map SRTE-color-map
exit-address-family
!
address-family ipv6
import map SRTE-color-map
exit-address-family

```

- **Route Redistribution Coloring into BGP:** The following configuration attaches a color-extended community as part of the redistribution routes to BGP. This associates the color-extended community to a prefix based on the source protocol owning the prefix:

```

route-map SRTE-color-map permit
set extcommunity color < 1-4294967295> [additive]
router bgp <ASnum>
address-family ipv4
redistribute <source-protocol> route-map SRTE-color-map
or
network <address> mask <network-mask> route-map SRTE-color-map
exit-address-family
!
address-family ipv4 vrf <vrf-name>
redistribute <source-protocol> route-map SRTE-color-map
or
network <address> mask <network-mask> route-map SRTE-color-map
exit-address-family
!
address-family ipv6
redistribute <source-protocol> route-map SRTE-color-map
or
network <address>/masklen route-map SRTE-color-map
exit-address-family
!
address-family ipv6 vrf <vrf-name>
redistribute <source-protocol=> route-map SRTE-color-map
or
network <address>/masklen route-map SRTE-color-map
exit-address-family

```

- **Neighbor Inbound Coloring:** The following configuration attaches a color-extended community as part of the inbound route map processing attached to the neighbor. This attaches a color-extended community based on the neighbor advertising the prefix:

```

route-map SRTE-color-map permit
set extcommunity color < 1-4294967295> [additive]
router bgp <ASnum>
address-family ipv4
neighbor <address> route-map SRTE-color-map in
exit-address-family
!
address-family vpv4
neighbor <address> route-map SRTE-color-map in
exit-address-family
!
address-family ipv4 vrf <vrf-name>
neighbor <address> route-map SRTE-color-map in
exit-address-family
!
address-family ipv6

```

```

neighbor <address> route-map SRTE-color-map in
exit-address-family
!
address-family vpnv6
neighbor <address> route-map SRTE-color-map in
exit-address-family
!
address-family ipv6 vrf <vrf-name>
neighbor <address> route-map SRTE-color-map in
exit-address-family

```

Support for PFP with RIB Path

PFP consists of a bundle output chain element (OCE), and each hash of the bundle OCE consists of a PDP policy (PDP tunnel). In this scenario, a PDP policy is created for the default IGP or RIB learned path. This means that a separate PDP policy is created for every default IGP or RIB learned path. Therefore, this implementation will eventually increase the number of policies and will not scale.

From Cisco IOS XE 17.9.1a, you can configure forwarding class in a PFP using the RIB path option. Instead of configuring a PDP, the RIB option uses the IGP shortest path to the policy destination.

PFP has a binding SID, similar to the PDP. The traffic-steering mechanism is also the same as PDP, either through BSID or through RIB.

A PFP is in the operational UP state based on the following conditions:

- The default FC is configured with a PDP, and it is in the operational UP state.
- The default FC is configured with the RIB path and is resolved.



Note The state of the nondefault FC does not affect the PFP state.

After a packet is steered on the PFP, according to the FC marked by Modular QoS CLI (MQC) at ingress, the following scenarios show the path of the packet:

- If PFP is in the Down state, the packet is dropped.
- If no FC is attached to a packet, the packet is forwarded with a default FC in PFP.
- If an FC is attached to a packet that points to a resolved RIB path or an operational PDP, the packet is forwarded to it.
- If an FC attached on a packet points to a nonexisting unresolved RIB path or a nonoperational PDP, the packet is forwarded to the default FC.

Example: Configuring PFP with RIB Path

The following example shows how to configure PFP using both the RIB path and color:

```

segment-routing traffic-eng
policy PERFLOW
color 10 end-point 1.1.1.1
binding-sid mpls 15001

```

```

candidate-path
  preference 1
  per-flow
    forward-class 0 rib
    forward-class 1 color 20
    forward-class 2 color 30

```

The following example shows how to configure the ODN PFP using both the RIB path and color:


```

segment-routing traffic-eng
on-demand color 10
candidate-path
  preference 1
  per-flow
    forward-class 0 rib
    forward-class 1 color 20
    forward-class 2 color 30

```

Configuring SR-TE Per-Flow Class (ODN) and Automated Steering (PCE Delegated)

Consider the following topology:



Perform the following steps to configure ODN for PFP:

1. Configure QOS on PE1:

```

class-map DSCP
  match DSCP AF41

```

- Set the forward class on the class map:

```

policy-map per-flow
  class DSCP
    set forward-class 1

```

- Attach the policy map on the corresponding interface:

```

interface GigabitEthernet0/0/3
  service-policy type eppr input PFP

```

2. Configure SR-TE PFP on PE1:

- Set the forward class on PFP:

```
on-demand color 4500
authorized
candidate-paths
  preference 2
  per-flow
    forward-class 0 color 100
    forward-class 0 rib
    forward-class 2 color 102
```

- Attach the segment list to PDP:

```
policy perflow_pdp
color 100 end-point 10.5.5.5
candidate-paths
  preference 2
  explicit segment-list srtel weight 10
  !
  constraints
    segments
      dataplane mpls
```

- Set the segment list to SR-TE:

```
segment-routing traffic-eng
  segment-list name srtel
    index 1 mpls label 16002
    index 2 mpls label 16005
```

3. Configure SR-TE PFP on PE2:

```
ip prefix-list pfp seq 5 permit 10.35.0.0/16 le 32
```

- Attach the route map to PFP:

```
route-map pfp permit 10
  match ip address prefix-list pfp
  set extcommunity color 4500
```

- Activate the BGP routes:

```
router bgp 100
!
address-family vpnv4
  neighbor 10.1.1.1 activate
  neighbor 10.1.1.1 send-community extended
  neighbor 10.1.1.1 route-map pfp out
```

4. View the PFP output:

```
show segment-routing traffic-eng policy name *6.6.6.6|4090 detail

Name: *6.6.6.6|4090 (Color: 4090 End-point: 6.6.6.6)
Owners : BGP
Status:
Admin: up, Operational: up for 01:29:41 (since 06-21 14:09:05.510)
Candidate-paths:
Preference 1 (BGP):
Per-flow Information (active):
Forward PDP PDP BSID RW
Class Color Status Status
-----
```



```

0 rib n/a n/a
1 129 up Done
2 130 up Done
3 131 up Done
4 132 up Done
Default Forward Class: 0
Attributes:
Binding SID: 39
Allocation mode: dynamic
State: Programmed
IPv6 caps enabled
Tunnel ID: 65568 (Interface Handle: 0x26)
Per owner configs:
BGP
Binding SID: dynamic
Stats:
5 minute output rate 0 bits/sec, 0 packets/sec
Packets: 500524 Bytes: 88056352

Event history:
Timestamp Client Event type Context: Value
-----:-----:-----:-----
06-21 14:09:05.489 BGP Policy created Name: BGP
06-21 14:09:05.490 BGP Set colour Colour: 4090
06-21 14:09:05.490 BGP Set end point End-point: 6.6.6.6
06-21 14:09:05.490 BGP Set dynamic pce Path option: per flow
06-21 14:09:05.510 BGP BSID allocated FWD: label 39
06-21 14:09:05.510 FH Resolution Policy state UP Status: PFP RESOLVED CP: 1
06-21 14:09:05.551 FH Resolution REOPT triggered Status: REOPTIMIZED CP: 1
06-21 14:09:05.576 FH Resolution REOPT triggered Status: REOPTIMIZED CP: 1
06-21 14:09:05.602 FH Resolution REOPT triggered Status: REOPTIMIZED CP: 1
06-21 14:09:05.626 FH Resolution REOPT triggered Status: REOPTIMIZED CP: 1

```

Verifying SR-TE Per-Flow Class (ODN) and Automated Steering (PCE Delegated)

Use the following command to verify SR-TE Per-Flow Class (ODN) and Automated Steering (PCE Delegated):

```

show segment-routing traffic-eng policy name *10.5.5.5|4500

Name: *10.5.5.5|4500 (Color: 4500 End-point: 10.5.5.5)
Owners : BGP
Status:
Admin: up, Operational: up for 00:03:50 (since 09-07 16:07:02.938)
Candidate-paths:
Preference 2 (BGP):
Per-flow Information (active):
Forward PDP PDP BSID RW
Class Color Status Status
-----:-----:-----:-----
0 100 up Done
1 101 up unknown Pending
2 102 up unknown Pending
Default Forward Class: 0
Attributes:
Binding SID: 72
Allocation mode: dynamic
State: Programmed
IPv6 caps enabled
Tunnel ID: 65675 (Interface Handle: 0x2D)

```

```
Per owner configs:  
BGP  
Binding SID: dynamic  
Stats:  
5 minute output rate 0 bits/sec, 0 packets/sec  
Packets: 9 Bytes: 584
```