

COE-PCE Initiated SR Policy with OSPF and IS-IS SR-TE Autoroute Announce

As part of a tactical TE solution, the Path Computation Element (PCE) can provision a Segment Routing Traffic Engineering (SR-TE) policy to mitigate link congestion.

Autoroute announcement is a steering mechanism in which IGPs automatically use the policy for destination's downstream of the policy end point. Autoroute announcement is performed using Cisco Crossworks Optimization Engine (COE). COE provides real-time network optimization allowing operators to maximize network utilization effectively and increase service velocity.

You can configure COE-PCE initiated SR policy in the following ways:

- PCE Initiated SR Policy with OSPF SR-TE Autoroute Announce — It enables a steering mechanism in which IGPs automatically use the SR-TE policy for destination's downstream of the policy end point.
- PCE-Initiated SR Policy with IS-IS SR-TE Autoroute Announce — It enables System-to-Intermediate System (IS-IS) interaction with traffic engineering to receive the SR-TE policies via autoroute announcement notifications.

- COE-PCE Initiated SR Policy with OSPF Autoroute Announce, on page 1
- SR-PCE: Support for PCE-Initiated SR Policy and ISIS Autoroute, on page 8
- LSR Support for Autoroute Announce SR Policies, on page 10
- Support of BGP PIC for Short LCM Policies, on page 11

COE-PCE Initiated SR Policy with OSPF Autoroute Announce

Table 1: Feature History

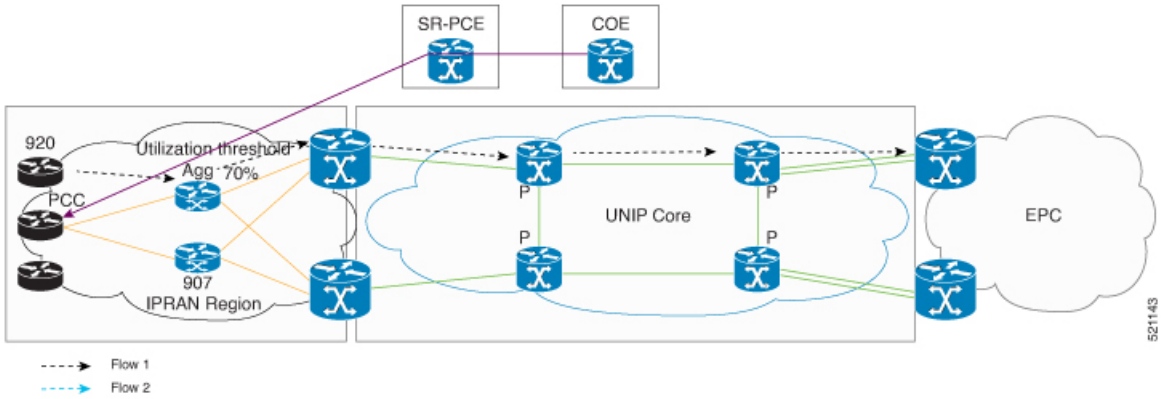
Feature Name	Release Information	Feature Description
PCE Initiated SR Policy with OSPF Autoroute Announce	Cisco IOS XE Bengaluru 17.4.1	This feature enables a steering mechanism in which IGPs automatically use the policy for destination's downstream of the policy end point.

A PCE collects various pieces of network information to determine traffic flows causing link congestion. The PCE computes a suitable path to divert those flows and to alleviate the congestion. The PCE then deploys the SR-TE policy to divert the traffic leading to the congestion using the Stateful Path Computation Element Protocol (PCEP) to provision the policy. When the congestion is alleviated, the SR-TE policy is removed.

The PCEP message contains SID list to be deployed by the head-end. Path Computation Client (PCC) profiles allow activation of autoroute announce for the policy provisioned by PCEP, using the profile IDs. The profile ID on the PCE and PCC should match, otherwise the policy is not provisioned. For example, if the PCE provisions a policy with profile ID 1 and the head-end where the policy is being provisioned also has the PCC profile ID 1 configured with autoroute announce, COE-PCE initiated SR policy is activated for that policy.

COE-PCE Initiated SR Policy

Figure 1: COE-PCE Initiated SR Policy



The preceding topology shows how an SR-PCE policy is initiated from COE:

- SR policy is configured on the COE with profile ID.
- COE pushes the SR policy to PCE and PCE forwards the SR policy to PCC.
- Profile ID on PCC is matched with the profile ID on COE-PCE.
- OSPF autoroute announce is configured on the PCC.
- The policy gets provisioned.
- The data traffic now adheres to the SR policy that is pushed from the COE.
- Complete SR Policy manipulation occurs only on COE.

Restrictions for PCE Initiated SR Policy

- A maximum of 500 SR policies are supported.
- Only native COE is supported.
- Effective Cisco IOS XE Bengaluru 17.5.1, Bandwidth optimization based on SR tactical policy is supported on RSP3.
- Bandwidth optimization by using COE is not supported.

- PIC core and PIC edge are not supported over SR-TE tunnel till Cisco IOS XE Cupertino Release 17.8.1. Starting with Cisco IOS XE Release 17.9.1, PIC core is supported for short LCM policies with 0, 1, or 2 SR labels.
- Effective Cisco IOS XE Bengaluru 17.5.1, ECMP over SR-TE is supported on RSP3.
- 6PE and 6VPE are not supported with three and four transport labels.
- IPv6 is not supported.
- A maximum of 10,000 VPNv4 prefix limits are supported.
- BGP LU (RFC 3107) is not supported for intra-AS and inter-AS.

ECMP Over SR-TE

Table 2: Feature History

Feature Name	Release Information	Feature Description
ECMP over SR-TE Policy	Cisco IOS XE Bengaluru 17.5.1	This feature allows you to configure ECMP over SR-TE policies. In case of multiple paths, this feature enables mitigation of local congestion through load balancing. This feature is supported only on Cisco ASR 900 RSP3 module.

The following sections explain how local congestion can be mitigated and how ECMP can be deployed over SR-TE policies to attain load balancing.

Note The traffic that is load balanced over multiple paths is HW-load balanced.

Restrictions for ECMP over SR-TE Policies

Cisco ASR 900 RSP3 module supports **sr_5_label_push_enable** and **sr_pfp_enable** templates. Following restrictions apply for different template combinations.

With **sr_5_label_push_enable** template:

- Only one service label is supported with LB over SR-TE tunnels with three or four TE labels. This service label includes L3VPN, L2VPN, 6PE, 6VPE, and RFC 3107 BGP-LU label.
- 6PE and 6VPE are not supported with three and four SR-TE tunnel labels.
- Segment routing is not supported in **enable_portchannel_qos_multiple_active** template.
- HW load balancing for L2VPN/EVPN services is not supported if the L2VPN/EVPN destination has a static route configured over SR-TE tunnel.

With **sr_pfp_enable** template:

- SR PM HW time stamping is not supported.
- VLAN COS marking is not supported.
- HW load balancing is not supported.
- Policer based hierarchical QOS on the ingress is not supported.
- Short-Pipe tunneling mode is not supported.

Other Restrictions:

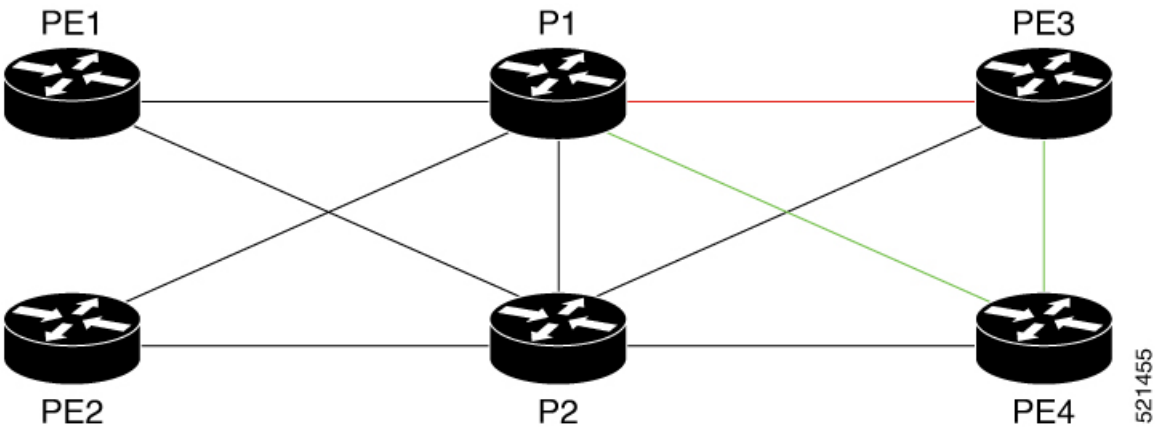
- PIC core over SR-TE tunnels are not supported till Cisco IOS XE Cupertino Release 17.8.1.
- PIC edge over SR TE tunnels are not supported till Cisco IOS XE Cupertino Release 17.8.1.
- PIC edge multipath over SR TE tunnels are not supported till Cisco IOS XE Cupertino Release 17.9.1.
- W-ECMP is not supported.
- Next hop ECMP is not supported within an SR policy.
- Local congestion mitigation (LCM) is applicable only for best effort traffic. All other delay sensitive traffic uses safe SIDs (Flex Algo 128). Delay sensitive traffic is not redirected using the LCM tunnels.

Local Congestion Mitigation

In today's network deployments it is important for every router in the network to have the capability to provision the traffic in such a way that it avoids the congestion based on the amount of traffic ingressing and egressing out of it. In order to provision this congestion mitigation, it is essential for the routers to support Equal Cost Multi-Path (ECMP) load balancing, that is, distributing the traffic based on the number of paths available to reach the destination.

Congestion mitigation helps the routers to move certain traffic to a different path than the current path, using the tactical SR policies. When the link congestion threshold is crossed, the COE (Cisco Optimization Engine) that monitors the link congestion based on the interface counters, pushes these tactical policies using PCE. These PCE initiated tactical policies that are used for local congestion mitigation (LCM) are deployed when necessary and only best effort traffic is load balanced over these tactical SR-TE policies.

Figure 2: Illustration of Local Congestion Mitigation



In the above topology, let us assume that the best effort traffic is coming in to P1 from PE1 and PE2 for the destination PE3 and the link between P1 and PE3 is congested. To mitigate the congestion between P1 and PE3, ECMP paths from P1 and PE3 are required. With segment routing this is achieved by deploying multiple tactical SR policies from P1 to PE3, one through directly connected link P1-PE3 and the other through the path P1-PE4-PE3. These policies are called tactical policies and are used to avoid local congestion mitigation by load balancing the best effort traffic over these tactical policies. The LCM is applicable only for best effort traffic. All other delay sensitive traffic would use safe SIDs (Flex Algo 128). Delay sensitive traffic is not redirected using the LCM tunnels. Originating traffic is directed on non-LCM tunnels and transiting traffic with safe-SIDs is treated as normal label entry traffic and forwarded accordingly.

In the above topology, any node may deploy LCM tactical tunnels to mitigate congestion over a particular link. These nodes transit or sometimes originate the traffic to the LCM tunnel end points or even beyond the tunnel end points.

Let us assume that PE nodes originate the traffic and P nodes are transit node for the traffic originated somewhere else. Based on these combinations following are the different types of traffic that have to be considered:

As a PE Node,

- L3VPN best effort traffic
- L2VPN best effort traffic
- Global traffic

As a P node,

- Any traffic that comes in for a non-flexible algorithm 0 label is treated as an entry swap on the Label lookup.
- Any traffic that comes in for flexible algorithm 0 label is treated as a swap case or it may be translated to pop and push stack of labels, if there is an LCM created for that outgoing link based on congestion.

Based on the number of TE labels that the LCM tunnels have to push, the number of labels outside of TE labels can be either one or two (service labels).

Load Balancing

At the head end, following are the different types of traffic that is subjected to load balancing. The traffic type here includes both best effort and delay sensitive.

As a PE Node,

- L3VPN traffic
- L2VPN traffic
- Global traffic

As a P node,

- Any traffic that comes in is treated based on the Label lookup.

Autoroute Announcement

Autoroute announcement or bandwidth optimization is used to steer traffic away from congested links and better utilize the network.

The PCEP message contains SID list to be deployed by the head-end. Path Computation Client (PCC) profiles allow autoroute announce to be activated for the policy instantiated by PCEP, using the profile IDs. For example, if the PCE instantiates a policy with profile ID 1 and the head-end where the policy is being instantiated has the PCC profile ID 1 configured with autoroute announce, PCE initiated SR policy is activated for that policy.

Autoroute announce can be configured under both policies created with strict SID and policies created with non-strict SID. The main difference between configuring autoroute under policies created with strict SID (assume A) and non-strict SID (assume B) is that with A, the lookup entry will be programmed only in RIB whereas with B, the lookup entry will be programmed in RIB and LFIB for flexible algorithm label 0.

Static Route Configuration

By adding a static route to the same destination but with different tunnels having the same endpoint, a load balancing is formed for the route over the tunnels configured. This is applicable for all types of traffic.

Next Hop ECMP within a SR Policy

If there is a SR policy created to a destination with a set of SIDs and the SR policy headend have multiple equal paths to reach the next hop, no ECMP is formed to reach the next hop within the SR policy.

Configuring ECMP over SR-TE Policy with OSPF Autoroute Announce

The below configuration shows how to configure ECMP over SR-TE policy with OSPF autoroute announce and PCE initiated segement routing policy with profile ID as 100.

```
pce
    address ipv4 13.13.13.13
segment-routing traffic-eng
   peer ipv4 10.0.0.1
    segment-list name ss1
    index 1 mpls label 18021
    index 2 mpls label 18023
    policy 100
     binding-sid mpls 15999
     color 100 end-point ipv4 12.12.12.12
     candidate-paths
      preference 10
       explicit segment-list ss1
    !
       constraints
       segments
       dataplane mpls
 profile-id 100
```

Now, to push the PCE initiated OSPF autoroute announce from PCE to PCC, the profile IDs on PCE and PCC must match. The below configuration shows the PCC configuration and that the profile ID is matching with PCE and thus the autoroute announce is enabled. Based on the number of autoroute policies configured on the ECMP link, packets are load balanced on the ECMP links for the non-strict SIDs.

```
segment-routing traffic-eng
  pcc
  pce address 13.13.13.13 source-address 10.0.0.1
  profile 100
```

```
   autoroute
    include all
```

Verifying SR Policy with Autoroute Announce

```
ASR903-R1#show segment-routing traffic-eng policy all

Name: *12.12.12.12|100 (Color: 100 End-point: 12.12.12.12)
Owners : PCEP
Status:
Admin: up, Operational: up for 66:41:16 (since 09-18 16:56:50.444)
Candidate-paths:
Preference 10 (PCEP):
PCC profile: 100
Dynamic (pce 13.13.13.13) (active)
Metric Type: TE, Path Accumulated Metric: 5
16003 [Prefix-SID, 3.3.3.3]
16012 [Prefix-SID, 12.12.12.12]
Attributes:
Binding SID: 15999
Allocation mode: explicit
State: Programmed
Autoroute:
Include all
```

Verifying OSPF Autoroute for IGP

Use the following two commands to verify the OSPF Autoroute for IGP:

```
ASR903-R1#show ip cef 12.12.12.12 -------------IGP ROUTE
12.12.12.12/32
nexthop 12.12.12.12 Tunnel65536 ------------Tunnel pushed for IGP ROUTE
```

```
ASR903-R1# show ip cef 12.12.12.12 internal
12.12.12.12/32, epoch 3, RIB[I], refcnt 6, per-destination sharing
  sources: RIB
  feature space:
    IPRM: 0x00028000
    Broker: linked, distributed at 1st priority
    LFD: 12.12.12.12/32 0 local labels
        contains path extension list
  ifnums:
    Tunnel65536(64)
  path list 3C97B678, 3 locks, per-destination, flags 0x49 [shble, rif, hwcn]
    path 3E393010, share 1/1, type attached nexthop, for IPv4
     MPLS short path extensions: [rib | lblmrg | srlbl] MOI flags = 0x1 label implicit-null

      nexthop 12.12.12.12 Tunnel65536, IP midchain out of Tunnel65536 2FFE3D00
  output chain:
    IP midchain out of Tunnel65536 2FFE3D00
    label [16012|16012]
    FRR Primary (0x3D9D4CE0)
      <primary: TAG adj out of Port-channel1, addr 100.0.0.2 3C9559C0>
      <repair:  TAG adj out of BDI1110, addr 111.0.0.2 3C954FC0>
```

Verify the Tunnel ID on the SR Policy

```
ASR903-R1# show segment-routing traffic-eng policy name margin detail
Name: Margin (Color: 1000 End-point: 12.12.12.12)
  Owners : CLI
  Status:
    Admin: up, Operational: up for 00:50:52 (since 09-16 11:00:06.697)
  Candidate-paths:
    Preference 10 (CLI):
```

```
      Dynamic (pce 13.13.13.13) (active)
        Metric Type: TE, Path Accumulated Metric: 5
          16012 [Prefix-SID, 12.12.12.12]
  Attributes:
    Binding SID: 15900
      Allocation mode: explicit
      State: Programmed
  IPv6 caps enabled
  Tunnel ID: 65536 (Interface Handle: 0x15B)
  Per owner configs:
    CLI
      Binding SID: 15900
  Stats:
    Packets: 535473  Bytes: 805338440
Event history:
    Timestamp           Client          Event type          Context: Value
    ---------           ------          ----------          -------: -----
    09-16 11:00:06.377  CLI             Policy created      Name: CLI
    09-16 11:00:06.418  CLI             Set colour          Colour: 1000
    09-16 11:00:06.418  CLI             Set end point       End-point: 12.12.12.12
    09-16 11:00:06.446  CLI             Set binding SID     BSID: Binding SID set
    09-16 11:00:06.577  CLI             Set dynamic         Path option: dynamic
    09-16 11:00:06.620  CLI             BSID allocated      FWD: label 15900
    09-16 11:00:06.637  FH Resolution   Policy state UP     Status: PATH RESOLVED
    09-16 11:00:06.697  FH Resolution   Policy state DOWN   Status: PATH NOT RESOLVED
    09-16 11:00:06.706  CLI             Set dynamic pce     Path option: dynamic pce
    09-16 11:00:07.240  FH Resolution   Policy state UP     Status: PATH RESOLVED
    09-16 11:00:09.520  FH Resolution   REOPT triggered     Status: REOPTIMIZED
```

SR-PCE: Support for PCE-Initiated SR Policy and ISIS Autoroute

Table 3: Feature History

Feature Name	Release Information	Description
SR-PCE: Support for PCE-Initiated SR Policy and ISIS Autoroute	Cisco IOS XE Cupertino 17.7.1	This feature enables System-to-Intermediate System (IS-IS) interaction with traffic engineering to receive the SR-TE policies via autoroute announcement notification. These notifications are used as IGP shortcuts during SPT computation and route calculation and are installed as nexthops for applicable routes in Routing Information Base (RIB) or MPLS Forwarding Infrastructure (MFI).

Prior to Cisco IOS XE Cupertino Release 17.7.1, PCE-initiated SR policy was only supported on OSPF protocol. Starting with Cisco IOS XE Cupertino Release 17.7.1, PCE-initiated SR policy is supported also on System-to-Intermediate System (IS-IS) that interacts with traffic engineering to receive the SR-TE policies via autoroute announcement notification. These notifications are used as IGP shortcuts during SPT computation and route calculation and are installed as nexthops for applicable routes in Routing Information Base (RIB) or MPLS Forwarding Infrastructure (MFI).

<need information to add network diagram>

Configure PCE-Initiated SR Policy and ISIS Autoroute

To configure PCE-initiated SR policy and ISIS autoroute:

```
policy Margin
  color 1000 end-point 12.12.12.12
  binding-sid mpls 15900
  candidate-paths
   preference 10
    constraints
     segments
      dataplane mpls
     !
dynamic
     pcep
     metric
      margin
       absolute 5
segment-routing traffic-eng
                              pcc
  pce address 13.13.13.13 source-address 10.0.0.1
  profile 100
   autoroute
    include all
```

Verification of IS-IS Autoroute Configuration

Use the **show isis segment-routing policy** command to verify the configuration of IS-IS segment routing policy.

```
Router2#show isis segment-routing policy

            ISIS Router with ID (2.2.2.2) (Process ID 1)

    Codes: SS - Strict SPF, SP - Default SPF
           r - relative, a - absolute, c - constant, n - none

Endpoint IP  Level/System-ID  Interface     Metric   Attributes       Last Updated
15.0.0.15    L1/R5.00         Tunnel65538     0      SS:n 00:01:14      00:01:14
16.0.0.16    L1/R6.00         Tunnel65539    -3      SS:r 00:00:03      00:01:21
```

LSR Support for Autoroute Announce SR Policies

Table 4: Feature History

Feature Name	Release Information	Description
LSR Support for Autoroute Announce SR Policies	Cisco IOS XE Cupertino 17.9.1	This feature enables Label Switch Routing (LSR) and thus helps to forward labeled (EOS0, EOS1) traffic over three or four labeled segment routing autoroute static tunnels.

Label Switch Routing (LSR) is a high-performance packet forwarding technology that integrates the performance and traffic management capabilities of data link layer (layer 2) switching with the scalability, flexibility, and performance of network layer (layer 3) routing. Each LSR in the network makes an independent, local decision as to which label value to use to represent a forwarding equivalence class. Each LSR informs its neighbors of the label bindings it has made.

Starting with Cisco IOS XE Cupertino Release 17.9.1, the introduction of LSR helps to forward labeled (EOS0, EOS1) traffic over three or four labeled segment routing autoroute static tunnels. Traffic can be destined to tunnel end-point or beyond end-point. Prior to this release, any MPLS traffic that was forwarded to segment routing tunnel with three or more labels would get dropped.

You must use metro aggregation services license to enable the feature.

Scenario: LSR Support for Autoroute Announce SR Policies

The figure below shows a network topology where routers R1 and R2 are configured with MPLS labels and routers R4, R5, and R6 are configured with the segment routing labels. The router R1 sends traffic to router R7 that is placed beyond the SR network. The router R3 is designated to translate the MPLS labels to SR labels and forward the traffic from router R1 to Router R7.

Prior to this release, router R3 was unable to translate MPLS label to SR label and hence traffic used to be dropped. Starting with this release, when you enable the feature and configure the router R3 with a static policy (autoroute announce with three or more labels), traffic flows uninterruptedly from router R1 to router R7.

Figure 3: LSR Support for Autoroute Announce SR Policies



Configure LSR Support for Autoroute Announce SR Policies

To configure LSR support for autoroute announce SR Policies:

1. **Enable the Feature**:

```
platform segment-routing traffic-eng lsr-over-extended-te-enable
```

2. **Shut the Static Policy**:

```
policy PE11-PE13
  shutdown
  color 50 end-point 13.13.13.13
  autoroute
   include all
  !
  candidate-paths
   preference 100
    explicit segment-list Prefix
    !
    constraints
     segments
      dataplane mpls
```

3. **Unshut the Static Policy**:

```
policy PE11-PE13
  no shutdown
  color 50 end-point 13.13.13.13
  autoroute
   include all
  !
  candidate-paths
   preference 100
    explicit segment-list Prefix
    !
    constraints
     segments
      dataplane mpls
```

Verification of LSR Support for Autoroute Announce SR Policies

Undisrupted traffic flow verifies the feature configuration.

Support of BGP PIC for Short LCM Policies

Table 5: Feature History

Feature Name	Release Information	Feature Description
Support of BGP PIC for Short LCM Policies	Cisco IOS XE Cupertino 17.9.1	This feature introduces the support of BGP Prefix Independent Convergence (PIC) and helps you to enable BGP PIC core and BGP PIC edge for short local congestion mitigation (LCM) policies. This feature helps to minimise the convergence time after a network failure. You should only configure LCM policies or the SR policies with 0, 1, and 2 SR labels.

The BGP PIC (Prefix Independent Convergence) Edge for IP and MPLS-VPN feature improves BGP convergence after a network failure. This convergence is applicable to both core and edge failures and can be used in both IP and MPLS networks. BGP PIC core and PIC edge for short LCM policies are not supported with BGP Labeled Unicast (LU).

Starting with Cisco IOS XE Cupertino Release 17.9.1, when you configure BGP Prefix Independent Convergence (PIC) core and BGP PIC edge for short local congestion mitigation (LCM) policies or SR policies, the convergence time is minimised. As a result, traffic can quickly and easily flow from one path to the other after a network failure. You should only configure LCM policies or the SR policies with 0, 1, and 2 SR labels to avoid any traffic loss or delay. During a traffic congestion, the traffic can move from one path to the other in less than 50 milliseconds for BGP PIC core and in less than a second for BGP PIC edge. Prior to this release, BGP PIC core and BGP PIC edge over any SR policy were not supported because of which convergence time could range between a second to a minute.

For more information on BGP PIC core and PIC edge, see IP Routing: BGP Configuration Guide.

Configure BGP PIC Short LCM Policies

Configure BGP PIC Core:

```
cef table output-chain build favor convergence-speed

 segment-list name Secondarytunnel
  index 1 mpls adjacency 138.0.0.2
  index 2 mpls label 18333
 !
 segment-list name primarytunnel
  index 1 mpls label 18333
policy Margin3
 color 3 end-point 3.3.3.3
 autoroute
 include all
  !
  candidate-paths
  preference 1
  explicit segment-list primarytunnel
    !
  constraints
  segments
  dataplane mpls
policy Margin4
  color 4 end-point 3.3.3.3
  autoroute
   include all
  !
  candidate-paths
   preference 1
    explicit segment-list Secondarytunnel
    !
    constraints
     segments
      dataplane mpls
     !
    !
```

Verification of BGP PIC Short LCM Policies Configuration

Use the **show segment-routing traffic-engineering policy** *name* command to verify the segment routing policy name configuration.

```
Router#show segment-routing traffic-engineering policy Margin3

Name: Margin3 (Color: 3 End-point: 3.3.3.3)
  Owners : CLI
  Status:
    Admin: up, Operational: up for 00:06:55 (since 05-19 08:28:37.189)
  Candidate-paths:
    Preference 1 (CLI):
      Explicit: segment-list primetunnel (active)
        Weight: 1, Metric Type: TE
          18333 [Prefix-SID, 3.3.3.3]
  Attributes:
    Binding SID: 19
      Allocation mode: dynamic
      State: Programmed
    Autoroute:
      Include all (Strict)
Router#show segment-routing traffic-engineering policy  Margin4

Name: Margin4 (Color: 4 End-point: 3.3.3.3)
  Owners : CLI
  Status:
    Admin: up, Operational: up for 00:06:58 (since 05-19 08:28:37.206)
  Candidate-paths:
    Preference 1 (CLI):
      Explicit: segment-list Secondarytunnel (active)
        Weight: 1, Metric Type: TE
          44 [Adjacency-SID, 138.0.0.1 - 138.0.0.2]
          18333 [Prefix-SID, 3.3.3.3]
  Attributes:
    Binding SID: 20
      Allocation mode: dynamic
      State: Programmed
    Autoroute:
      Include all (Strict)
```

Use the show ip cef vrf command to display Cisco Express Forwarding (CEF)-related Virtual Routing and Forwarding (VRF) information.

```
Router#show ip cef vrf L3VPN 150.0.0.0 internal | sec output   output chain:
    loadinfo 39E0E004, per-session, 2 choices, flags 0403, 4 locks
      flags [Per-session, for-rx-IPv4, recursive]
      translation map 39DBE764 owned by path list 38188B28, 9610 locks
        2 choices, 16 buckets, flags 0x1
        Path index      [  0  1  0  1  0  1  0  1  0  1  0  1  0  1  0  1 ]
        Repair path     [  -  X  -  X  -  X  -  X  -  X  -  X  -  X  -  X ]

        Path available  [  X  X  X  X  X  X  X  X  X  X  X  X  X  X  X  X ]
        Translation map [  0  0  2  2  4  4  6  6  8  8 10 10 12 12 14 14 ]
      16 hash buckets
        < 0 > label 24
              loadinfo 308BE6E4, per-session, 5 choices, flags 0111, 9612 locks
                flags [ Per-session, for-mpls-not-at-eos, indirection]
< 1 > label 17005-(local:17005)
                      TAG midchain out of Tunnel65537 388777E0
                      label [explicit-null|explicit-null](ptr:0x39887F90)-(local:18)
                      FRR Primary (0x39C6E860)
                        <primary: TAG adj out of Port-channel1, addr 100.0.0.2 38873420>
                        <repair:  TAG midchain out of MPLS-SR-Tunnel1 39BFD300
```

```
                    label 20
                    TAG adj out of BDI138, addr 138.0.0.2 39BFDB00>
< 2 > label 17005-(local:17005)
      TAG midchain out of Tunnel65538 388773E0
      label [explicit-null|explicit-null](ptr:0x39888090)-(local:19)
```