



# Configure Segment Routing Path Computation Element

The Segment Routing Path Computation Element (SR-PCE) provides stateful PCE functionality by extending the existing IOS-XR PCEP functionality with additional capabilities. SR-PCE is supported on the MPLS data plane and IPv4 control plane.



**Note** The Cisco IOS XRv 9000 is the recommended platform to act as the SR-PCE. Refer to the [Cisco IOS XRv 9000 Router Installation and Configuration Guide](#) for more information.

**Table 1: Feature History Table**

Feature Name	Release Information	Feature Description
SR-PCE: Single PCE scale enhancement	Release 7.5.1	With this feature, support for a single PCE is enhanced to 50000 nodes, 100000 LSPs, 500000 links, and 2000 PCEP sessions.

- [About SR-PCE, on page 2](#)
- [Usage Guidelines and Limitations, on page 3](#)
- [Configure SR-PCE, on page 3](#)
- [PCE Override of PCC Initiated Policies, on page 8](#)
- [PCE-Initiated SR Policies, on page 11](#)
- [SR-PCE Flexible Algorithm Multi-Domain Path Computation, on page 13](#)
- [Exclude Network Resources during Path Computation over SR-TE Policies, on page 18](#)
- [Configure the Shortest Path for Disjoint Candidate Paths, on page 23](#)
- [Enable Strict Disjointness for SR-TE policies in the PCE, on page 25](#)
- [ACL Support for PCEP Connection, on page 27](#)
- [Anycast SID-Aware Path Computation, on page 27](#)
- [SR-PCE IPv4 Unnumbered Interface Support, on page 32](#)
- [Inter-Domain Path Computation Using Redistributed SID, on page 34](#)
- [Configuring the North-Bound API on SR-PCE, on page 38](#)

## About SR-PCE

*Table 2: Feature History Table*

Feature Name	Release Information	Feature Description
TCP Authentication Option	Release 7.3.1	This feature introduces support for TCP Authentication Option (TCP-AO), which replaces the TCP Message Digest 5 (MD5) option, which was used for authenticating PCEP (TCP) sessions by using a clear text or encrypted password.

The path computation element protocol (PCEP) describes a set of procedures by which a path computation client (PCC) can report and delegate control of head-end label switched paths (LSPs) sourced from the PCC to a PCE peer. The PCE can request the PCC to update and modify parameters of LSPs it controls. The stateful model also enables a PCC to allow the PCE to initiate computations allowing the PCE to perform network-wide orchestration.



**Note** For more information on PCE, PCC, and PCEP, refer to the [Path Computation Element](#) section in the *MPLS Configuration Guide for Cisco ASR 9000 Series Routers*.

SR-PCE learns topology information by way of IGP (OSPF or IS-IS) or through BGP Link-State (BGP-LS).

SR-PCE is capable of computing paths using the following methods:

- TE metric—SR-PCE uses the TE metric in its path calculations to optimize cumulative TE metric.
- IGP metric—SR-PCE uses the IGP metric in its path calculations to optimize reachability.
- LSP Disjointness—SR-PCE uses the path computation algorithms to compute a pair of disjoint LSPs. The disjoint paths can originate from the same head-end or different head-ends. Disjoint level refers to the type of resources that should not be shared by the two computed paths. SR-PCE supports the following disjoint path computations:
  - Link – Specifies that links are not shared on the computed paths.
  - Node – Specifies that nodes are not shared on the computed paths.
  - SRLG – Specifies that links with the same SRLG value are not shared on the computed paths.
  - SRLG-node – Specifies that SRLG and nodes are not shared on the computed paths.

When the first request is received with a given disjoint-group ID, the first LSP is computed, encoding the shortest path from the first source to the first destination. When the second LSP request is received with the same disjoint-group ID, information received in both requests is used to compute two disjoint paths: one path from the first source to the first destination, and another path from the second source to the second destination. Both paths are computed at the same time.

### TCP Authentication Option

TCP Message Digest 5 (MD5) authentication has been used for authenticating PCEP (TCP) sessions by using a clear text or encrypted password. This feature introduces support for TCP Authentication Option (TCP-AO), which replaces the TCP MD5 option.

TCP-AO uses Message Authentication Codes (MACs), which provides the following:

- Protection against replays for long-lived TCP connections
- More details on the security association with TCP connections than TCP MD5
- A larger set of MACs with minimal system and operational changes

TCP-AO is compatible with Master Key Tuple (MKT) configuration. TCP-AO also protects connections when using the same MKT across repeated instances of a connection. TCP-AO protects the connections by using traffic key that are derived from the MKT, and then coordinates changes between the endpoints.



---

**Note** TCP-AO and TCP MD5 are never permitted to be used simultaneously. TCP-AO supports IPv6, and is fully compatible with the proposed requirements for the replacement of TCP MD5.

---

## Usage Guidelines and Limitations

To ensure PCEP compatibility, we recommend that the Cisco IOS XR version on the SR-PCE be the same or later than the Cisco IOS XR version on the PCC or head-end.

## Configure SR-PCE

This task explains how to configure SR-PCE.

### Before you begin

The Cisco IOS XRv 9000 is the recommended platform to act as the SR-PCE.

### SUMMARY STEPS

1. **configure**
2. **pce**
3. **address ipv4** *address*
4. **state-sync ipv4** *address*
5. **tcp-buffer size** *size*
6. **password** { **clear** | **encrypted** } *password*
7. **tcp-ao** *key-chain* [**include-tcp-options**] [**accept-ao-mismatch-connection**]
8. **segment-routing** { **strict-sid-only** | **te-latency** }
9. **timers**
10. **keepalive** *time*
11. **minimum-peer-keepalive** *time*

12. `reoptimization time`
13. `exit`

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b> <b>Example:</b> RP/0/RSP0/CPU0:router# <code>configure</code>	Enters global configuration mode.
<b>Step 2</b>	<b>pce</b> <b>Example:</b> RP/0/RSP0/CPU0:router(config)# <code>pce</code>	Enables PCE and enters PCE configuration mode.
<b>Step 3</b>	<b>address ipv4 address</b> <b>Example:</b> RP/0/RSP0/CPU0:router(config-pce)# <code>address ipv4 192.168.0.1</code>	Configures a PCE IPv4 address.
<b>Step 4</b>	<b>state-sync ipv4 address</b> <b>Example:</b> RP/0/RSP0/CPU0:router(config-pce)# <code>state-sync ipv4 192.168.0.3</code>	Configures the remote peer for state synchronization.
<b>Step 5</b>	<b>tcp-buffer size size</b> <b>Example:</b> RP/0/RSP0/CPU0:router(config-pce)# <code>tcp-buffer size 1024000</code>	Configures the transmit and receive TCP buffer size for each PCEP session, in bytes. The default buffer size is 256000. The valid range is from 204800 to 1024000.
<b>Step 6</b>	<b>password {clear   encrypted} password</b> <b>Example:</b> RP/0/RSP0/CPU0:router(config-pce)# <code>password encrypted pwd1</code>	Enables TCP MD5 authentication for all PCEP peers. Any TCP segment coming from the PCC that does not contain a MAC matching the configured password will be rejected. Specify if the password is encrypted or clear text. <b>Note</b> TCP-AO and TCP MD5 are never permitted to be used simultaneously.
<b>Step 7</b>	<b>tcp-ao key-chain [include-tcp-options] [accept-ao-mismatch-connection]</b> <b>Example:</b> RP/0/RSP0/CPU0:router(config-pce)# <code>tcp-ao</code>	Enables TCP Authentication Option (TCP-AO) authentication for all PCEP peers. Any TCP segment coming from the PCC that does not contain a MAC matching the configured key chain will be rejected.

	Command or Action	Purpose
	<code>pce_tcp_ao include-tcp-options</code>	<ul style="list-style-type: none"> <li>• <b>include-tcp-options</b>—Includes other TCP options in the header for MAC calculation.</li> <li>• <b>accept-ao-mismatch-connection</b>—Accepts connection even if there is a mismatch of AO options between peers.</li> </ul> <p><b>Note</b> TCP-AO and TCP MD5 are never permitted to be used simultaneously.</p>
<b>Step 8</b>	<b>segment-routing</b> { <b>strict-sid-only</b>   <b>te-latency</b> } <b>Example:</b> RP/0/RSP0/CPU0:router(config-pce) # <b>segment-routing strict-sid-only</b>	Configures the segment routing algorithm to use strict SID or TE latency. <p><b>Note</b> This setting is global and applies to all LSPs that request a path from this controller.</p>
<b>Step 9</b>	<b>timers</b> <b>Example:</b> RP/0/RSP0/CPU0:router(config-pce) # <b>timers</b>	Enters timer configuration mode.
<b>Step 10</b>	<b>keepalive</b> <i>time</i> <b>Example:</b> RP/0/RSP0/CPU0:router(config-pce-timers) # <b>keepalive 60</b>	Configures the timer value for locally generated keep-alive messages. The default time is 30 seconds.
<b>Step 11</b>	<b>minimum-peer-keepalive</b> <i>time</i> <b>Example:</b> RP/0/RSP0/CPU0:router(config-pce-timers) # <b>minimum-peer-keepalive 30</b>	Configures the minimum acceptable keep-alive timer that the remote peer may propose in the PCEP OPEN message during session establishment. The default time is 20 seconds.
<b>Step 12</b>	<b>reoptimization</b> <i>time</i> <b>Example:</b> RP/0/RSP0/CPU0:router(config-pce-timers) # <b>reoptimization 600</b>	Configures the re-optimization timer. The default timer is 1800 seconds.
<b>Step 13</b>	<b>exit</b> <b>Example:</b> RP/0/RSP0/CPU0:router(config-pce-timers) # <b>exit</b>	Exits timer configuration mode and returns to PCE configuration mode.

## Configure the Disjoint Policy (Optional)

This task explains how to configure the SR-PCE to compute disjointness for a pair of LSPs signaled by PCCs that do not include the PCEP association group-ID object in their PCEP request. This can be beneficial for deployments where PCCs do not support this PCEP object or when the network operator prefers to manage the LSP disjoint configuration centrally.

### SUMMARY STEPS

1. **disjoint-path**
2. **group-id** *value* **type** {**link** | **node** | **srlg** | **srlg-node**} [**sub-id** *value*]
3. **strict**
4. **lsp** {**1** | **2**} **pcc** **ipv4** *address* **lsp-name** *lsp\_name* [**shortest-path**]

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>disjoint-path</b> <b>Example:</b> <pre>RP/0/RSP0/CPU0:router(config-pce)# disjoint-path</pre>	Enters disjoint configuration mode.
Step 2	<b>group-id</b> <i>value</i> <b>type</b> { <b>link</b>   <b>node</b>   <b>srlg</b>   <b>srlg-node</b> } [ <b>sub-id</b> <i>value</i> ] <b>Example:</b> <pre>RP/0/RSP0/CPU0:router(config-pce-disjoint)# group-id 1 type node sub-id 1</pre>	<p>Configures the disjoint group ID and defines the preferred level of disjointness (the type of resources that should not be shared by the two paths):</p> <ul style="list-style-type: none"> <li>• <b>link</b>—Specifies that links are not shared on the computed paths.</li> <li>• <b>node</b>—Specifies that nodes are not shared on the computed paths.</li> <li>• <b>srlg</b>—Specifies that links with the same SRLG value are not shared on the computed paths.</li> <li>• <b>srlg-node</b>—Specifies that SRLG and nodes are not shared on the computed paths.</li> </ul> <p>If a pair of paths that meet the requested disjointness level cannot be found, then the paths will automatically fallback to a lower level:</p> <ul style="list-style-type: none"> <li>• If the requested disjointness level is SRLG or node, then link-disjoint paths will be computed.</li> <li>• If the requested disjointness level was link, or if the first fallback from SRLG or node disjointness failed, then the lists of segments encoding two shortest paths, without any disjointness constraint, will be computed.</li> </ul>

	Command or Action	Purpose
Step 3	<b>strict</b> <b>Example:</b> RP/0/RSP0/CPU0:router(config-pce-disjoint)# <b>strict</b>	(Optional) Prevents the automatic fallback behavior of the preferred level of disjointness. If a pair of paths that meet the requested disjointness level cannot be found, the disjoint calculation terminates and no new path is provided. The existing path is not modified.
Step 4	<b>lsp {1   2} pcc ipv4 address lsp-name lsp_name [shortest-path]</b> <b>Example:</b> RP/0/RSP0/CPU0:router(config-pce-disjoint)# <b>lsp 1 pcc ipv4 192.168.0.1 lsp-name rtrA_t1 shortest-path</b> RP/0/RSP0/CPU0:router(config-pce-disjoint)# <b>lsp 2 pcc ipv4 192.168.0.5 lsp-name rtrE_t2</b>	Adds LSPs to the disjoint group.  The <b>shortest-path</b> keyword forces one of the disjoint paths to follow the shortest path from the source to the destination. This option can only be applied to the the first LSP specified.

## Global Maximum-delay Constraint

This feature allows a PCE to compare the cumulative latency of a computed path against a global maximum-delay constraint value. If the latency of the computed path exceeds this global constraint, the path is not considered valid. This ensures that all latency-based paths computed by the PCE and signaled to the PCCs in the network do not exceed this maximum-delay constraint.

```
pce
constraints
  bounds
    cumulative
    type
      latency <1-4294967295> Bound metric value in microseconds
```

### Configuration

To configure a PCE for specifying maximum cumulative latency metric, you must complete the following configurations:

```
RP/0/RSP0/CPU0:ios(config)# pce
RP/0/RSP0/CPU0:ios(config-pce)# constraints
RP/0/RSP0/CPU0:ios(config-pce-constr)# bounds
RP/0/RSP0/CPU0:ios(config-pce-constr-bounds)# cumulative
RP/0/RSP0/CPU0:ios(config-pce-constr-bounds-type)# type latency 1000000
RP/0/RSP0/CPU0:ios(config-pce-constr-bounds-type)#
```

### Verification

Verify using the **show** command:

```
RP/0/RSP0/CPU0:ios(config-pce-constr-bounds-type)# show
Wed Oct 12 22:18:22.962 UTC
pce
constraints
  bounds
```

```

cumulative
  type latency 1000000
!
!
!
!
!

```

## PCE Override of PCC Initiated Policies

**Table 3: Feature History Table**

Feature Name	Release Information	Feature Description
PCE Override of PCC-Initiated Policies	Release 7.7.1	<p>You can now override the Label Switched Paths (LSP) attributes in the Path Computation Elements (PCEs) to improve the path computation of the Path Computation Client (PCC)-Initiated policies. The overriding is based on the matching criteria of the Path Computation Element Protocol (PCEP) Peer and LSP attributes.</p> <p>You can configure the PCE override rule for one, many, or all PCCs.</p> <p>This feature introduces the following commands:</p> <ul style="list-style-type: none"> <li>• <a href="#">override-rules (PCE)</a></li> <li>• <a href="#">pce try-regex</a></li> <li>• <a href="#">show pce override-rules detail</a></li> </ul>

There are fewer PCEs in the network compared to the PCCs. However, PCE versions may be newer than the PCC version because it is easy to update the lesser PCEs than the higher number of PCCs. This ensures PCE supports more features with the latest Internet Assigned Numbers Authority (IANA) code points than PCCs.

IOS-XR PCE implementation is based on the Path Computation Element Protocol (PCEP) standards and the implementations of the PCEP vary leading to interoperability challenges. To overcome these challenges, you can configure PCE with the override-rules (OVR). These OVRs have a filter part and an overriding part. The matching criteria filter matches all the (Segment Routing) SR policies against these OVRs, suppose they match, then the overriding part applies to these SR policies attributes.

Matching is done when the LSP is created in the PCE LSP database and the LSP is not re-evaluated after receiving the next PCEP report message for the same LSP. Based on these matching criteria, you can override the values of the LSP attributes.

For example, suppose that a PCC is outdated and does not support reporting the latency metric type. In that case, PCC reports to the LSP with the metric-type that it supports and the PCE applies the OVRs that you configure to update the metric-type upon reception.

### Limitations and guidelines

- Applicable only for PCC initiated policies
- This feature does not support SRv6 policies.
- Supports only Segment Routing (SR) setup type



- Supports only ten override rules.
- Lower sequence numbers have higher precedence. LSP can be matched to one override sequence at a time. It will be matched to the highest preference sequence and further matching stops.
- Show commands like “show pce lsp” display values received from PCC, so applied changes are visible only in values, which are based on path-computation.
- In a PCE High Availability case, where there are multiple PCEs for a particular PCC, configure this feature likewise on all PCEs, to allow seamless transition of delegation between the PCEs.
- PCE (North Bound) NB-API clients only see overridden attributes in their communication with PCEs.
- The PCE OVR feature is seamless for PCCs. This means that PCCs will not know if the tunnel’s attributes are overridden. PCEs uses overridden attributes internally and sends only the original attributes to the PCCs.
- Only one color or color range per override rule is supported.
- The filtration option "all" under peer and LSP cannot be configured with another filtration option under the same config submodule. For example, all and regex cannot be provided together under the LSP matching criteria.
- Only IPv4 Access Control Lists (ACL)s are supported for matching PCEP peers.
  - If the ACL is empty or there are no entries within the ACL, then it is implicit permit.
  - If the ACL is not configured but the name is specified, then it is implicit permit.
  - In all other cases, if no entries are matched in the ACL, then it is implicit deny.
  - The behavior aligns with ACL behavior in other areas on an XR router.

### Configuration steps: PCE override of PCC-Initiated policies

To configure a PCE override of PCC-Initiated policies, you must complete the following configurations:



**Note** Before you apply the configuration you can use the **try-regex** utility to test the regexes. This utility is an optional step and does not alter the system. This example shows how to verify if the override-rule for the matching-criteria *lsp name <regex>* works:

```
Router# pce try-regex ^cp_c_[0-9]+$ cp_c_5000
Regex verification utility
Regex:      ^cp_c_[0-9]+$
Test string: cp_c_5000
Result:     Matched
```

1. Enter the PCE configuration mode.
2. Create sequence and matching criteria.
3. Create override-rules.

```
Rrouter(config)# pce
Rrouter(config-pce)# override-rules
```

```

Router(config-pce-ovr-rule)# Sequence 100
Router(config-pce-ovr-rule-seq)# matching-criteria
Router(config-pce-ovr-rule-crit)# peer
Router(config-pce-ovr-rule-peer)# all
Router(config-pce-ovr-rule-peer)# exit
Router(config-pce-ovr-rule-crit)# lsp colors 0-50
Router(config-pce-ovr-rule-crit)# exit
Router(config-pce-ovr-rule-seq)# override metric type igp
Router(config-pce-ovr-rule-seq)# override constraints bandwidth 1000
Router(config-pce-ovr-rule-seq)# commit

```

### Running configuration

- Match all LSPs from all peers and modify metric type to latency:

```

pce
  override-rules
    sequence 1
      matching-criteria
        peer
        all
        lsp
        all
      override
        metric
        type latency
  !

```

- Match LSP name using regex and peer based on provided ACL and change metric type to IGP:

```

pce
  override-rules
    sequence 1
      matching-criteria
        peer
        access-list ipv4 PCC1
      !
      lsp
        name ^cfg_test2_.*$
      !
      override
        metric
        type igp
  !

```

- Match policy using specified Segment Routing (SR) policy color from all peers and change bandwidth to 1000 kbps:

```

pce
  override-rules
    sequence 1
      matching-criteria
        peer
        all
      !
      lsp
        colors 10
      !
      override

```

```
constraints
  bandwidth 1000
```

```
!
```

### Verification

Show command displays operational values, which are applied during the override. For example, a bandwidth value is converted into IEEE float format. The value that is displayed in the show command may not be equal to the value configured, because of the limited precision of IEEE format used in PCEP.

Verify using the show command **show pce override-rules [sequence <sequence>]**

Show command output with two override rules that are configured and each matched one LSP are:

```
Router# show pce override-rules details

PCE's Override Rule database:
-----
Sequence number: 1
Matching criteria:
  Peer:
    IPv4 ACL name: PCC1
    LSP
    Regex: ^cfg_test1_.*$
Override:
  Metric type: Latency
  Constraints:
Matching LSPs:
  Peer: 192.168.0.1, Tunnel name: cfg_test1_xxx_discr_5

Sequence number: 2
Matching criteria:
  Peer:
    IPv4 ACL name: PCC1
    LSP
    Regex: ^cfg_test2_.*$
Override:
  Metric type: IGP
  Constraints:
Matching LSPs:
  Peer: 192.168.0.1, Tunnel name: cfg_test2_xxx_discr_5
```

## PCE-Initiated SR Policies

Use cases based on centralized optimization, such as congestion mitigation solutions, rely on the ability of the PCE to signal and instantiate SR-TE policies in the network. We refer to this as PCE-initiated SR-TE policies.

PCE-initiated SR-TE policies can be triggered via Crossworks Network Controller (recommended approach) or via CLI at the PCE.

For more information on configuring SR-TE policies, see the [SR-TE Policy Overview](#).

The PCE deploys the SR-TE policy using PCC-PCE communication protocol (PCEP).

1. PCE sends a PCInitiate message to the PCC.
2. If the PCInitiate message is valid, the PCC sends a PCRpt message; otherwise, it sends PCErr message.
3. If the PCInitiate message is accepted, the PCE updates the SR-TE policy by sending PCUpd message.

You can achieve high-availability by configuring multiple PCEs with SR-TE policies. If the head-end (PCC) loses connectivity with one PCE, another PCE can assume control of the SR-TE policy.

### Configuration Example: PCE-Initiated SR Policy with Explicit SID List

To configure a PCE-initiated SR-TE policy, you must complete the following configurations:

1. Enter PCE configuration mode.
2. Create the segment list.




---

**Note** When configuring an explicit path using IP addresses of intermediate links, the SR-TE process prefers the protected Adj-SID of the link, if one is available.

---

3. Create the policy.

```
/* Enter PCE configuration mode and create the SR-TE segment lists */
Router# configure
Router(config)# pce

/* Create the SR-TE segment lists */
Router(config-pce)# segment-routing
Router(config-pce-sr)# traffic-eng
Router(config-pce-sr-te)# segment-list name addr2a
Router(config-pce-sr-te-sl)# index 10 address ipv4 10.1.1.2
Router(config-pce-sr-te-sl)# index 20 address ipv4 10.2.3.2
Router(config-pce-sr-te-sl)# index 30 address ipv4 10.1.1.4
Router(config-pce-sr-te-sl)# exit

/* Create the SR-TE policy */
Router(config-pce-sr-te)# peer ipv4 10.1.1.1
Router(config-pce-sr-te)# policy P1
Router(config-pce-sr-te-policy)# color 2 end-point ipv4 2.2.2.2
Router(config-pce-sr-te-policy)# candidate-paths
Router(config-pce-sr-te-policy-path)# preference 50
Router(config-pce-sr-te-policy-path-preference)# explicit segment-list addr2a
Router(config-pce-sr-te-pp-info)# commit
Router(config-pce-sr-te-pp-info)# end
Router(config)#
```

### Running Config

```
pce
segment-routing
traffic-eng
segment-list name addr2a
index 10 address ipv4 10.1.1.2
index 20 address ipv4 10.2.3.2
index 30 address ipv4 10.1.1.4
!
peer ipv4 10.1.1.1
policy P1
color 2 end-point ipv4 2.2.2.2
candidate-paths
preference 50
```

```

explicit segment-list addr2a
!
!

```

## SR-PCE Flexible Algorithm Multi-Domain Path Computation

Table 4: Feature History Table

Feature Name	Release Information	Feature Description
SR-PCE Flexible Algorithm Multi-Domain Path Computation	Release 7.3.1	With this feature, SR-PCE can use Flexible Algorithms to compute multi-domain paths.

Flexible Algorithm provides a traffic engineered path automatically computed by the IGP to any destination reachable by the IGP. With the SR-PCE Flexible Algorithm Multi-Domain Path Computation feature, SR-PCE can use Flexible Algorithms to compute multi-domain paths. See the [Enabling Segment Routing Flexible Algorithm](#) chapter for information about Segment Routing Flexible Algorithm.

The SR-PCE Flexible Algorithm Multi-Domain Path Computation feature incorporates the following functionality:

- BGP-LS has been augmented to allow selected nodes to advertise the Flexible Algorithm definition (FAD) to the SR-PCE
- PCEP has been augmented (vendor-specific object) to allow a PCC to indicate SR policy constraint based on the Flexible Algorithm instance number
- SR-PCE algorithms have been augmented to compute paths based on a Flexible Algorithm constraint

The SR-PCE Flexible Algorithm multi-domain path computation requires the following:

- The same Flexible Algorithm instance ID is used across domains.
- The metric for those Flexible Algorithm instances must be the same across domains.
- The affinity constraints for those Flexible Algorithm instances may be different across domains.
- Multiple Flexible Algorithms can exist in a domain.

For example, considering a multi-domain topology (Domain 1 and Domain 2), the following scenarios meet the requirements listed above:

Scenario	Domain 1	Domain 2
Scenario 1	Flexible Algorithm 128, metric delay	Flexible Algorithm 128, metric delay
Scenario 2	Flexible Algorithm 128, metric delay	Flexible Algorithm 128, metric delay, exclude affinity blue
Scenario 3	Flexible Algorithm 128, metric delay, exclude affinity yellow	Flexible Algorithm 128, metric delay, exclude affinity blue



```

address-family ipv4 unicast
  metric-style wide
  router-id 10.1.1.8
  segment-routing mpls
!
interface Loopback0
  passive
  address-family ipv4 unicast
  prefix-sid absolute 16008
  prefix-sid algorithm 128 absolute 16808
!

```

### Configuration on Node 4 (ABR/ASBR)

#### IS-IS and Flexible Algorithm Configuration

```

router isis 1
  is-type level-2-only
  net 49.0001.0000.0000.0004.00
  distribute link-state instance-id 100
  flex-algo 128
  metric-type delay
  advertise-definition

  address-family ipv4 unicast
    metric-style wide
    router-id 10.1.1.4
    segment-routing mpls
  !
  interface Loopback0
    passive
    address-family ipv4 unicast
    prefix-sid absolute 16004
    prefix-sid algorithm 128 absolute 16804
  !
router isis 2
  is-type level-2-only
  net 49.0002.0000.0000.0004.00
  distribute link-state instance-id 200
  flex-algo 128
  metric-type delay
  advertise-definition

  address-family ipv4 unicast
    metric-style wide
    router-id 10.1.1.4
    segment-routing mpls
  !
  interface Loopback0
    passive
    address-family ipv4 unicast
    prefix-sid absolute 16004
    prefix-sid algorithm 128 absolute 16804
  !

```

#### BGP-LS Configuration

```

router bgp 65000
  bgp router-id 10.1.1.4
  address-family link-state link-state
  !
  neighbor-group AS65000-LS-group

```





```
pcc
 source-address ipv4 10.1.1.1
 pce address ipv4 10.1.1.10
 precedence 10
 !
 report-all
 !
 !
 !
```

### Configuration on PCE

```
pce
 address ipv4 10.1.1.10
 rest
 !
 !
router bgp 65000
 bgp router-id 10.1.1.10
 address-family link-state link-state
 !
 neighbor-group AS65000-LS-group
 remote-as 65000
 update-source Loopback0
 address-family link-state link-state
 !
 !
 neighbor 10.1.1.4
 use neighbor-group AS65000-LS-group
 description *** To Node-4 ***
 !
 !
 neighbor 10.1.1.5
 use neighbor-group AS65000-LS-group
 description *** To Node-5 ***
 !
 !
 !
```

# Exclude Network Resources during Path Computation over SR-TE Policies

*Table 5: Feature History Table*

Feature Name	Release Information	Feature Description
Exclude Network Resources during Path Computation over SR-TE Policies	Release 24.1.1	

Feature Name	Release Information	Feature Description
		<p>You can reduce the risk of attacks due to less secure network resources and IP addresses, improve network performance affected by overloaded or congested paths, and ensure higher levels of network stability and availability that could be impacted by resources experiencing issues and requiring maintenance. These improvements are possible because you can now exclude specific network resources using their IP addresses during path computation for traffic over SR-TE policies.</p> <p>Previously, you could not exclude network resources during path computation.</p> <p>The feature introduces these changes:</p> <p><b>CLI:</b></p> <ul style="list-style-type: none"> <li>• The <b>resources</b> option is introduced in the <b>segment-routing traffic-eng policy</b> and <b>segment-routing traffic-eng on-demand color</b> commands.</li> <li>• The <b>resource-list</b> keyword is introduced in the <b>segment-routing traffic-eng</b> command.</li> </ul> <p><b>YANG Data Models:</b></p> <ul style="list-style-type: none"> <li>• <code>Cisco-IOS-XR-infra-xtc-oper.yang</code></li> <li>• <code>Cisco-IOS-XR-infra-xtc-agent-oper.yang</code></li> <li>• <code>Cisco-IOS-XR-infra-xtc-agent-cfg.yang</code></li> </ul> <p>See (<a href="#">GitHub</a>, <a href="#">Yang Data Models Navigator</a>)</p>

You can configure the SR-TE policies as either single or disjoint candidate paths that do not include certain IP addresses or subnets when traffic is routed through the network.

The key benefits of the feature are:

- **Security** - Excluding certain IP addresses or network resources can help protect sensitive resources from unauthorized access or attacks.
- **Performance** - By excluding certain resources, traffic can be directed away from congested parts of the network.
- **Isolation** - Excluding resources can help isolate different parts of the network, preventing issues in one part of the network from affecting others.

## Configuration to exclude network resources during Path Computation

To exclude network resources during path computation for traffic, you must complete these high-level tasks in order:

1. Configure the network resources or IP addresses that you want to exclude from the network list.
2. Associate the excluded network resources to candidate paths for SR-TE or ODN SR-TE policies.

### Configure the network resources or IP addresses to exclude from the network list

Perform the following task to configure a list of IPv4 addresses that you want to exclude from the network resource list:

```
Router(config)#segment-routing traffic-eng
Router(config-sr-te)#resource-list node_resc_list
Router(config-sr-te-rl)#index 1 ipv4 10.10.10.1
Router(config-sr-te-rl)#index 2 ipv4 10.10.10.8
```

#### Running Configuration

```
!
segment-routing
 traffic-eng
  resource-list node_resc_list
  index 1 ipv4 10.10.10.1
  index 2 ipv4 10.10.10.8
!
```

### Associate the excluded network resources to candidate paths for SR-TE policies

Perform the following task to associate the excluded IPv4 addresses to one or more candidate paths for SR-TE policies:

```
Router(config)#segment-routing traffic-eng
Router(config-sr-te)#policy dynamic pcep_policy
Router(config-sr-te-policy)#candidate-paths
Router(config-sr-te-policy-path)#preference 100
Router(config-sr-te-policy-path-pref)#constraints resources exclude resource-list
node_resc_list
```

#### Running Configuration

```
!
segment-routing
 traffic-eng
  policy dynamic pcep_policy
  candidate-paths
```

```

    preference 100
    constraints
    resources
    exclude resource-list node_resc_list
    !
    !
    !
    !
    !
    !
    !

```

## Verification

```
Router#show segment-routing traffic-eng policy endpoint ipv4 100.2.1.1 color 8001
```

```

Wed Aug 30 22:21:50.014 UTC
SR-TE policy database
-----
Color: 8001, End-point: 100.2.1.1
Name: srte_c_8001_ep_100.2.1.1
Status:
Admin: up Operational: up for 00:00:50 (since Aug 30 22:20:59.341)
Candidate-paths:
Preference: 100 (configuration) (active)
Name: dynamic_pcep_policy
Requested BSID: 8001
PCC info:
Symbolic name: cfg_dynamic_pcep_policy_discr_100
PLSP-ID: 14042
Constraints:
Protection Type: protected-preferred
Maximum SID Depth: 12
Exclude Resources: node_resc_list
10.10.10.1
10.10.10.8
Dynamic (pce 100.7.1.1) (valid)
Metric Type: TE, Path Accumulated Metric: 440
SID[0]: 41111 [Adjacency-SID, 101.1.5.1 - 101.1.5.2]
SID[1]: 21600 [Prefix-SID, 100.6.1.1]
SID[2]: 41600 [Adjacency-SID, 101.2.6.2 - 101.2.6.1]
Attributes:
Binding SID: 8001
Forward Class: Not Configured
Steering labeled-services disabled: yes
Steering BGP disabled: no
IPv6 caps enable: yes
Invalidation drop enabled: no
Max Install Standby Candidate Paths: 0

```

## Associate the excluded network resources to candidate paths for ODN SR-TE policies

Perform the following task to associate the excluded IPv4 addresses for ODN SR-TE policies:

```

Router(config)#segment-routing
Router(config-sr)#traffic-eng
Router(config-sr-te)#on-demand color 7001
Router(config-sr-te-color)#constraints resources exclude resource-list node_resc_list

```

## Running Configuration

```

!
segment-routing
 traffic-eng
  on-demand color 7001

```

```

constraints
  resources
    exclude resource-list node_resc_list
  !
!
!
!
!

```

## Verification

```
Router#show segment-routing traffic-eng policy endpoint ipv4 100.2.1.1 color 7001
```

```
Wed Aug 30 22:53:01.079 UTC
```

```
SR-TE policy database
```

```

-----
Color: 7001, End-point: 100.2.1.1
Name: srte_c_7001_ep_100.2.1.1
Status:
  Admin: up Operational: up for 00:31:56 (since Aug 30 22:21:04.869)
Candidate-paths:
  Preference: 200 (BGP ODN) (inactive) (shutdown)
  Requested BSID: dynamic
  Constraints:
    Protection Type: protected-preferred
    Maximum SID Depth: 12
    Exclude Resources: node_resc_list
    10.10.10.1
    10.10.10.8
  Dynamic (inactive)
    Metric Type: IGP, Path Accumulated Metric: 0
  Preference: 100 (BGP ODN) (active)
  Requested BSID: dynamic
  PCC info:
    Symbolic name: bgp_c_7001_ep_100.2.1.1_discr_100
    PLSP-ID: 14044
  Constraints:
    Protection Type: protected-preferred
    Maximum SID Depth: 12
    Exclude Resources: node_resc_list
    10.10.10.1
    10.10.10.8
  Dynamic (pce 100.7.1.1) (valid)
    Metric Type: IGP, Path Accumulated Metric: 440
    SID[0]: 41111 [Adjacency-SID, 101.1.5.1 - 101.1.5.2]
    SID[1]: 21600 [Prefix-SID, 100.6.1.1]
    SID[2]: 41600 [Adjacency-SID, 101.2.6.2 - 101.2.6.1]
Attributes:
  Binding SID: 51679
  Forward Class: Not Configured
  Steering labeled-services disabled: yes
  Steering BGP disabled: no
  IPv6 caps enable: yes
  Invalidation drop enabled: no
  Max Install Standby Candidate Paths: 0

```

# Configure the Shortest Path for Disjoint Candidate Paths

Table 6: Feature History Table

Feature Name	Release Information	Feature Description
Configure the Shortest Path for Disjoint Candidate Paths	Release 24.1.1	<p>You can now configure the available disjoint paths for Label Switched Paths (LSPs) to prefer the shortest path between two points in the network. This configuration ensures that traffic is routed along the most efficient route in the network.</p> <p>Previously, you could not configure the shortest path preference for disjoint LSPs.</p> <p>The feature introduces these changes:</p> <p><b>CLI:</b></p> <p>The <b>shortest-path</b> keyword is introduced in the <b>policy candidate-paths constraints disjoint-path</b> command.</p> <p><b>YANG Data Models:</b></p> <ul style="list-style-type: none"> <li>• Cisco-IOS-XR-infra-xtc-oper.yang</li> <li>• Cisco-IOS-XR-infra-xtc-agent-oper.yang</li> <li>• Cisco-IOS-XR-infra-xtc-agent-cfg.yang</li> </ul> <p>See (<a href="#">GitHub</a>, <a href="#">Yang Data Models Navigator</a>)</p>

The configuration enables the available disjoint paths to indicate preference for the shortest path between two points in the network. This way the traffic is contained within specific network planes as per the affinity constraints while still provisioning Label Switched Paths (LSPs) that are diverse from each other to meet the path disjointness requirement.

In earlier releases, if path disjointness was configured for LSPs, the Path Computation Element (PCE) prioritized finding disjoint paths for the LSPs over adhering to the affinity constraints.

## Configure the shortest path for disjoint candidate paths

Perform the following task to indicate the disjoint path preference for the shortest path in the network:

```
Router(config)#segment-routing traffic-eng
Router(config-sr-te)#policy dynamic_pcep_policy_disjoint
```

```

Router(config-sr-te-policy)#candidate-paths
Router(config-sr-te-policy-path)#preference 100
Router(config-sr-te-policy-path-pref)#constraints disjoint-path group-id 1 type link
shortest-path

```

### Running Configuration

```

!
segment-routing
 traffic-eng
  policy dynamic_pcep_policy
  candidate-paths
  preference 100
  constraints
  disjoint-path group-id 1 type link shortest-path
!
!
!
!
!
!

```

### Verification

```
Router#show pce lsp name cfg_dynamic_pcep_policy_disjoint_discr_100 detail
```

Output received:

Wed Aug 30 18:33:57.807 UTC

PCE's tunnel database:

-----

PCC 100.1.1.1:

Tunnel Name: cfg\_dynamic\_pcep\_policy\_disjoint\_discr\_100

Color: 30115

Interface Name: srte\_c\_30115\_ep\_100.2.1.1

LSPs:

LSP[0]:

source 100.1.1.1, destination 100.2.1.1, tunnel ID 9031, LSP ID 24

State: Admin up, Operation active

Setup type: Segment Routing

Binding SID: 30115

Maximum SID Depth: 12

Preference: 100

Bandwidth: requested 0 kbps, applied 0 kbps

Protection type: protected-preferred

Prefix-SID algorithm: 0 (set: 0)

PCEP information:

PLSP-ID 0x36e9, flags: D:1 S:0 R:0 A:1 O:2 C:0

LSP Role: Exclude LSP

State-sync PCE: None

PCC: 100.1.1.1

LSP is subdelegated to: None

Reported path:

Metric type: TE, Accumulated Metric 30

SID[0]: Node, Label 21200, Address 100.2.1.1

Computed path: (Local PCE)

Computed Time: Wed Aug 30 18:31:12 UTC 2023 (00:02:45 ago)

Metric type: TE, Accumulated Metric 30

SID[0]: Node, Label 21200, Address 100.2.1.1

Reverse path: (Local PCE)

None

Computed Time: Wed Aug 30 18:31:12 UTC 2023 (00:02:45 ago)

Recorded path:

None



```

Disjoint Group Information:
  Type Link-Disjoint, Group 1 (SP)
SR Policy Association Group:
  Color: 30115, Endpoint: 100.2.1.1
  Policy Name: srte_c_30115_ep_100.2.1.1
  Preference: 100
  CP Name: dynamic_pcep_policy_disjoint

```

## Enable Strict Disjointness for SR-TE policies in the PCE

Table 7: Feature History Table

Feature Name	Release Information	Feature Description
Enable Strict Disjointness for SR-TE Policies in the PCE	Release 24.1.1	<p>You can now enforce strict disjoint constraints for Label Switched Paths (LSPs) and minimize the risk of a single point of failure that affects multiple LSPs. With this feature, if disjoint paths cannot be found for LSPs, the Path Computation Element (PCE) does not return any path.</p> <p>Previously, enforcing strict disjoint constraints for disjoint paths for LSPs was not possible.</p> <p>The feature introduces these changes:</p> <p><b>CLI:</b></p> <p>The <b>fallback disable</b> keyword is introduced in the <a href="#">segment-routing traffic-eng policy</a> and <a href="#">segment-routing traffic-eng on-demand color</a> commands.</p> <p><b>YANG Data Models:</b></p> <ul style="list-style-type: none"> <li>• <a href="#">Cisco-IOS-XR-infra-xtc-oper.yang</a></li> <li>• <a href="#">Cisco-IOS-XR-infra-xtc-agent-oper.yang</a></li> <li>• <a href="#">Cisco-IOS-XR-infra-xtc-agent-cfg.yang</a></li> </ul> <p>See (<a href="#">GitHub</a>, <a href="#">Yang Data Models Navigator</a>)</p>

When you enable **fallback disable**, the Path Computation Element (PCE) does not return any path for LSPs that are not disjoint if there are no disjoint candidate paths. In other words, when the PCE does not find multiple LSPs that satisfy the disjointness requirement, it returns no path. This indicates that there might be a networking issue that prevents the establishment of the desired redundancy, and you can investigate and address the underlying problem.

Without the option, the PCE can relax disjointness by applying an objective function or use a local policy when no objective function is requested.

Enforcing strict disjointness ensures that LSPs do not share any common links or nodes along their paths, minimizing the risk of a single point of failure affecting multiple LSPs. The feature is critical to maintain network reliability and ensures that traffic is properly routed in the event of a network failure.

## Configure Strict Disjointness in the PCE

### Enable strict disjointness for ODN SR-TE policies

Perform the following steps to enable strict disjointness for ODN SR-TE policies:

```
Router(config)#segment-routing traffic-eng
Router(config-sr-te)#on-demand color 4
Router(config-sr-te-color)#dynamic
Router(config-sr-te-color-dyn)#disjoint-path group-id 1 type node fallback disable
Router(config-sr-te-color-dyn)#commit
```

### Running Configuration

```
segment-routing
 traffic-eng
  on-demand color 4
  dynamic
   disjoint-path group-id 1 type node fallback disable
  !
 !
 !
 !
```

### Enable strict disjointness for SR-TE policies

Perform the following steps to enable strict disjointness for SR-TE policies:

```
Router(config)#segment-routing traffic-eng
Router(config-sr-te)#policy foo
Router(config-sr-te-policy)#color 1 end-point ipv4 10.10.10.1
Router(config-sr-te-policy)#candidate-paths preference 100
Router(config-sr-te-policy-path-pref)#constraints disjoint-path group-id 1 type node fallback
disable
Router(config-sr-te-policy-path-pref)#commit
```

### Running Configuration

```
segment-routing
 traffic-eng
  policy foo
   color 1 end-point ipv4 10.10.10.1
   candidate-paths
    preference 100
    dynamic
     pcep
    !
    metric
     type latency
    !
   constraints
    disjoint-path group-id 1 type node fallback disable
  !
```

```

!
!
!
!
!

```

### Verification

```

Router#sh pce association
Wed Mar 13 14:38:27.173 PDT

PCE's association database:
-----
Association: Type Node-Disjoint, Group 1, Strict
Associated LSPs:
  LSP[0]:
    PCC 10.10.10.1, tunnel name cfg_foo_discr_100,  PLSP ID 4, tunnel ID 8, LSP ID 2,
    Configured on PCC
  LSP[1]:
    PCC 10.10.10.1, tunnel name bgp_c_4_ep_10.10.10.2_discr_100,  PLSP ID 1, tunnel ID 5,
    LSP ID 3, Configured on PCC
Status: Satisfied

```

## ACL Support for PCEP Connection

PCE protocol (PCEP) (RFC5440) is a client-server model running over TCP/IP, where the server (PCE) opens a port and the clients (PCC) initiate connections. After the peers establish a TCP connection, they create a PCE session on top of it.

The ACL Support for PCEP Connection feature provides a way to protect a PCE server using an Access Control List (ACL) to restrict IPv4 PCC peers at the time the TCP connection is created based on the source address of a client. When a client initiates the TCP connection, the ACL is referenced, and the client source address is compared. The ACL can either permit or deny the address and the TCP connection will proceed or not.

Refer to the Implementing Access Lists and Prefix Lists chapter in the *IP Addresses and Services Configuration Guide for Cisco ASR 9000 Series Routers* for detailed ACL configuration information.

To apply an ACL to the PCE, use the **pce peer-filter ipv4 access-list *acl\_name*** command.

The following example shows how to configure an ACL and apply it to the PCE:

```

pce
 address ipv4 10.1.1.5
 peer-filter ipv4 access-list sample-peer-filter
!
ipv4 access-list sample-peer-filter
 10 permit ipv4 host 10.1.1.6 any
 20 permit ipv4 host 10.1.1.7 any
 30 deny ipv4 any any
!

```

## Anycast SID-Aware Path Computation

This feature allows the SR-TE head-end or SR-PCE to compute a path that is encoded using Anycast prefix SIDs of nodes along the path.


An Anycast SID is a type of prefix SID that identifies a set of nodes and is configured with n-flag clear. The set of nodes (Anycast group) is configured to advertise a shared prefix address and prefix SID. Anycast routing enables the steering of traffic toward multiple advertising nodes, providing load-balancing and redundancy. Packets addressed to an Anycast address are forwarded to the topologically nearest nodes.



**Note** For information on configuring Anycast SID, see [Configuring a Prefix-SID on the IS-IS Enabled Loopback Interface](#) and [Configuring a Prefix-SID on the OSPF-Enabled Loopback Interface](#).


This example shows how Anycast SIDs are inserted into a computed SID list.

The following figure shows 3 isolated IGP domains without redistribution and without BGP 3107. Each Area Border Router (ABR) 1 through 4 is configured with a node SID. ABRs 1 and 2 share Anycast SID 16012 and ABRs 3 and 4 share Anycast SID 16034.



Consider the case where nodes A and Z are provider edge (PE) routers in the same VPN. Node A receives a VPN route with BGP next-hop to node Z. Node A resolves the SR path to node Z based on ODN behaviors with delegation of path computation to SR-PCE.

Before considering Anycast SIDs, the head-end router or SR-PCE computes the SID list.



Assume that the computed path from node A to node Z traverses node 2 and node 4. This translates to SID list {16002, 16004, 1600Z} when node SIDs are leveraged to encode the path.

When an Anycast SID-aware path is requested, the path computation algorithm performs the following:


- **Path Computation**—Computes the path according to optimization objectives and constraints
- **Path Encoding**—Encodes the path in a SID list leveraging node-SIDs and adj-SIDs as applicable
- **Anycast SID Replacement**—Reiterates the original SID list by replacing node SIDs with Anycast SIDs present on the nodes along the computed path.

If a node has multiple Anycast SIDs, the algorithm considers them according to their weights. See [Weighted Anycast SIDs, on page 31](#).


- **Optimality Validation**—The new paths are validated against the original optimization objectives and constraints (maintain same cumulative metric as original SID list and do not violate path constraints).
- **Anycast SID Promotion**—If the optimality validation is successful, then the Anycast-encoded SID list is signaled and instantiated in the forwarding.

The following figure depicts cumulative metrics between nodes in the network.

Under these conditions, the optimality check is met, and therefore, the Anycast-encoded SID list from node A to node Z is {16012,16034,1600Z}.



The Anycast SID aware path computation also provides resiliency. For example, if one of the ABRs (in this case, ABR 1) becomes unavailable or unreachable, the path from node A to node Z {16012,16034,1600Z} will still be valid and usable.



### Configuration Examples

1. Configure Prefix SIDs on the ABR nodes.
  - a. Configure each node with a node SID.
  - b. Configure each group of nodes with a shared Anycast SID.

See [Configuring a Prefix-SID on the IS-IS Enabled Loopback Interface](#) and [Configuring a Prefix-SID on the OSPF-Enabled Loopback Interface](#).

2. Configure SR policies to include Anycast SIDs for path computation using the **anycast-sid-inclusion** command.

This example shows how to configure a local SR policy to include Anycast SIDs for PCC-initiated path computation at the head-end router:

```
Router(config)# segment-routing traffic-eng
Router(config-sr-te)# policy FOO
Router(config-sr-te-policy)# color 10 end-point ipv4 10.1.1.10
Router(config-sr-te-policy)# candidate-paths
Router(config-sr-te-policy-path)# preference 100
Router(config-sr-te-policy-path-pref)# dynamic
Router(config-sr-te-pp-info)# anycast-sid-inclusion
```

### Running Configuration

Use the **anycast-sid-inclusion** command to request Anycast SID-aware path computation for the following SR policy types:

- Local SR policy with PCC-initiated path computation at the head-end router:

```
segment-routing
 traffic-eng
  policy FOO
    color 10 end-point ipv4 10.1.1.10
    candidate-paths
      preference 100
      dynamic
      anycast-sid-inclusion
```

- Local SR policy with PCC-initiated/PCE-delegated path computation at the SR-PCE:

```

segment-routing
traffic-eng
policy BAR
color 20 end-point ipv4 10.1.1.20
candidate-paths
preference 100
dynamic
  pcep
  anycast-sid-inclusion

```

- On-demand SR policies with a locally computed dynamic path at the head-end, or centrally computed dynamic path at the SR-PCE:

```

segment-routing
traffic-eng
on-demand color 10
dynamic
  anycast-sid-inclusion

```

- On-demand SR policies with centrally computed dynamic path at the SR-PCE:

```

segment-routing
traffic-eng
on-demand color 20
dynamic
  pcep
  anycast-sid-inclusion

```

## Weighted Anycast SIDs

**Table 8: Feature History Table**

Feature Name	Release Information	Feature Description
Weighted Anycast SID-Aware Path Computation	Release 7.3.1	<p>This feature extends Anycast SIDs with weighted nodes.</p> <p>Weighted Anycast nodes advertise a cost (weight) along with the Anycast SID. Traffic is then distributed according to the weights.</p> <p>Weighted Anycast SIDs allow for highly available paths with node redundancy and path optimality that provide Fast Re-Route (FRR) for node failure of service provider edge (PE) routers and ABR/ASBRs nodes in multi-domain networks.</p>

Weighted Anycast nodes advertise a cost along with the Anycast SID. This cost serves as a weight. The native SR path computation algorithms are augmented to compute optimum paths relying on Weighted Anycast SIDs during path encoding. Traffic to the SID is then distributed according to the weights.

The following example shows how node SID, Anycast SID, and Weighted Anycast SID are applied on node 1:

```

router isis 1
 interface Loopback0
  address-family ipv4 unicast
  prefix-sid absolute 16001 // Node SID
  !
  !
 interface Loopback1
  prefix-attributes anycast
  address-family ipv4 unicast
  prefix-sid absolute 16012 // Anycast SID - (prefer node 1 or 2)
  !
  !
 interface Loopback2
  prefix-attributes anycast
  address-family ipv4 unicast
  weight 1
  prefix-sid absolute 17012 // Weighted Anycast SID (prefer node 1)
  !
  !
 interface Loopback3
  prefix-attributes anycast
  address-family ipv4 unicast
  weight 100000
  prefix-sid absolute 18012 // Weighted Anycast SID (prefer node 2)
  !
  !
  !

```

## SR-PCE IPv4 Unnumbered Interface Support

This feature allows IPv4 unnumbered interfaces to be part of an SR-PCE topology database.

An unnumbered IPv4 interface is not identified by its own unique IPv4 address. Instead, it is identified by the router ID of the node where this interfaces resides and the local SNMP index assigned for this interface.

This feature provides enhancements to the following components:

- IGPs (IS-IS and OSPF):
  - Support the IPv4 unnumbered interfaces in the SR-TE context by flooding the necessary interface information in the topology
- SR-PCE:




---

**Note** SR-PCE and path computation clients (PCCs) need to be running Cisco IOS XR 7.0.2 or later.

---

- Compute and return paths from a topology containing IPv4 unnumbered interfaces.
- Process reported SR policies from a head-end router that contain hops with IPv4 unnumbered adjacencies.

PCEP extensions for IPv4 unnumbered interfaces adhere to IETF RFC8664 “PCEP Extensions for Segment Routing” (<https://datatracker.ietf.org/doc/rfc8664/>). The unnumbered hops use a Node or



Adjacency Identifier (NAI) of type 5. This indicates that the segment in the explicit routing object (ERO) is an unnumbered adjacency with an IPv4 ID and an interface index.

- SR-TE process at the head-end router:
  - Compute its own local path over a topology, including unnumbered interfaces.
  - Process PCE-computed paths that contain hops with IPv4 unnumbered interfaces.
  - Report a path that contains hops with IPv4 unnumbered interfaces to the PCE.

### Configuration Example

The following example shows how to configure an IPv4 unnumbered interface:

```
RP/0/0/CPU0:rtrA(config)# interface GigabitEthernet0/0/0/0
RP/0/0/CPU0:rtrA(config-if)# ipv4 point-to-point
RP/0/0/CPU0:rtrA(config-if)# ipv4 unnumbered Loopback0
```

To bring up the IPv4 unnumbered adjacency under the IGP, configure the link as point-to-point under the IGP configuration. The following example shows how to configure the link as point-to-point under the IGP configuration:

```
RP/0/0/CPU0:rtrA(config)# router ospf one
RP/0/0/CPU0:rtrA(config-ospf)# area 0
RP/0/0/CPU0:rtrA(config-ospf-ar)# interface GigabitEthernet0/0/0/0
RP/0/0/CPU0:rtrA(config-ospf-ar-if)# network point-to-point
```

### Verification

Use the **show ipv4 interface** command to display information about the interface:

```
RP/0/0/CPU0:rtrA# show ipv4 interface GigabitEthernet0/0/0/0 brief
Tue Apr  2 12:59:53.140 EDT
Interface                               IP-Address      Status          Protocol
GigabitEthernet0/0/0/0                 192.168.0.1    Up              Up
```

This interface shows the IPv4 address of Loopback0.

Use the **show snmp interface** command to find the SNMP index for this interface:

```
RP/0/0/CPU0:rtrA# show snmp interface
Tue Apr  2 13:02:49.190 EDT
ifName : Null0                ifIndex : 3
ifName : Loopback0            ifIndex : 10
ifName : GigabitEthernet0/0/0/0 ifIndex : 6
```

The interface is identified with the pair (IPv4:192.168.0.1, index:6).

Use the **show ospf neighbor** command to display the adjacency:

```
RP/0/0/CPU0:rtrA# show ospf neighbor gigabitEthernet 0/0/0/0 detail
...
Neighbor 192.168.0.4, interface address 192.168.0.4
  In the area 0 via interface GigabitEthernet0/0/0/0
  Neighbor priority is 1, State is FULL, 6 state changes
...
Adjacency SIDs:
  Label: 24001,      Dynamic, Unprotected
```

**Neighbor Interface ID: 4**

The output of the **show pce ipv4 topology** command is enhanced to display the interface index instead of the IP address for unnumbered interfaces:

```
RP/0/0/CPU0:sr-pce# show pce ipv4 topology
...
Link[2]: unnumbered local index 6, remote index 4
  Local node:
    OSPF router ID: 192.168.0.1 area ID: 0 ASN: 0
  Remote node:
    TE router ID: 192.168.0.4
    OSPF router ID: 192.168.0.4 area ID: 0 ASN: 0
  Metric: IGP 1, TE 1, Latency 1 microseconds
  Bandwidth: Total 125000000 Bps, Reservable 0 Bps
  Admin-groups: 0x00000000
  Adj SID: 24001 (unprotected)
```

The output of **show pce lsp detail** command includes unnumbered hops:

```
RP/0/0/CPU0:sr-pce# show pce lsp detail
...
Reported path:
  Metric type: TE, Accumulated Metric 3
  SID[0]: Adj unnumbered, Label 24001, local 192.168.0.1(6), remote 192.168.0.4(4)
  SID[1]: Adj unnumbered, Label 24002, local 192.168.0.4(7), remote 192.168.0.3(7)
  SID[2]: Adj unnumbered, Label 24000, local 192.168.0.3(5), remote 192.168.0.2(5)
Computed path: (Local PCE)
  Computed Time: Wed Apr 03 11:01:46 EDT 2019 (00:01:06 ago)
  Metric type: TE, Accumulated Metric 3
  SID[0]: Adj unnumbered, Label 24001, local 192.168.0.1(6), remote 192.168.0.4(4)
  SID[1]: Adj unnumbered, Label 24002, local 192.168.0.4(7), remote 192.168.0.3(7)
  SID[2]: Adj unnumbered, Label 24000, local 192.168.0.3(5), remote 192.168.0.2(5)
```

## Inter-Domain Path Computation Using Redistributed SID

**Table 9: Feature History Table**


Feature Name	Release Information	Feature Description
SR-PCE: Inter-Domain Computation using Redistributed SID (SRTE - BGP-LU Domains Interworking)	Release 7.3.1	This feature adds new functionality to the SR-PCE that enables it to compute a path for remote non-SR end-point device distributed by BGP-LU.

A Path Computation Element (PCE) computes SR-TE paths based on SR topology database that stores connectivity, state, and TE attributes of SR network nodes and links. BGP Labeled Unicast (BGP-LU) provides MPLS transport across IGP boundaries by advertising loopbacks and label binding of impact edge and border routers across IGP boundaries.

This feature adds new functionality to the SR-PCE that enables it to compute a path for remote non-SR end-point device distributed by BGP-LU.

The remote end-point device in the BGP-LU domain is unknown to the SR-PCE. For the SR-PCE to know about the end-point device, the gateway ABR/ASBR learns the end-point prefix via BGP-LU. The prefix is then redistributed to SR-PCE topology database from the gateway ABR/ASBR. SR-PCE then can compute the best path from the head-end device to the selected gateway router.


The following topology shows an SR domain and a BGP-LU domain, with a gateway ABR/ASBR between the two domains.



1. The gateway ABR/ASBR is configured with BGP/IGP helper to learn the remote prefix through BGP-LU and redistribute the remote prefix to the IGP helper, then to SR-PCE.
2. The SR-PCE selects the best gateway node to BGP-LU domain and computes the path to reach the remote prefix through the gateway node.
3. The head-end device in the SR domain requests a path to the remote destination and signals the SR profile interworking with the BGP-LU domain.

The BGP-LU prefix advertisement to SR-PCE Traffic Engineer Database (TED) is done by creating an IGP helper on the ABR/ASBR to redistribute BGP-LU prefix information to IGP. IGP then sends the prefix information to the SR-PCE via BGP-LS.

If there are multiple ABR/ASBRs advertising the same remote BGP-LU prefix, the SR-PCE selects the best gateway node to the BGP-LU domain using the accumulative metric from the head-end device to the gateway and the advertised metric from the gateway to the destination.



## Example: Inter-Domain Path Computation Using Redistributed SID

The following examples show the configurations for the IGP helper, BGP-LU, and proxy BGP-SR:

### Configuration on the End-Point Device

Configure the end-point device to allocate a label for the BGP-LU prefix on the end-point device:

```
router bgp 3107
  bgp router-id 1.0.0.8
  address-family ipv4 unicast
    network 1.0.0.8/32 route-policy bgplu-com
    allocate-label all

route-policy bgplu-com
  set community (65002:999)
end-policy
```

### Configuration on the Gateway ABR/ASBR

1. Configure the remote prefix set and create the route policy for the BGP-LU domain:

```
prefix-set bgplu
  1.0.0.7/32,
  1.0.0.8/32,
  1.0.0.101/32,
  1.0.0.102/32
end-set
!

route-policy bgp2isis
  if destination in bgplu then
    pass
  else
    drop
  endif
```

```
end-policy
!  
end
```

## 2. Configure the helper IGP instance on the Loopback interface:

```
router isis 101
 is-type level-2-only
 net 49.0001.0000.1010.1010.00
 distribute link-state instance-id 9999
 nsf cisco
 nsf lifetime 120
 address-family ipv4 unicast
 metric-style wide
 maximum-paths 64
 router-id Loopback10
 redistribute bgp 3107 metric 200 route-policy bgp2isis
 segment-routing mpls sr-prefer
!  
interface Loopback10 >>> this loopback is for gateway SR-TE node-id
 passive
 address-family ipv4 unicast
 prefix-sid index 2001 explicit-null
```

## 3. Configure the gateway proxy BGP-SR and SR Mapping Server to allocate SR labels:

```
router bgp 3107
 address-family ipv4 unicast
 segment-routing prefix-sid-map
 allocate-label all

segment-routing
 global-block 16000 23999
 mapping-server
 prefix-sid-map
 address-family ipv4
 1.0.0.7/32 2007
 1.0.0.8/32 2008
 1.0.0.101/32 2101
 1.0.0.102/32 2102
```

## Configuring the North-Bound API on SR-PCE

Table 10: Feature History Table

Feature Name	Release Information	Feature Description
SR-Multicast: Tree-SID Label Range Validation in North-Bound API	Release 7.10.1	<p>The SR-PCE provides a north-bound HTTP-based API to establish communication between the SR-PCE and the Cisco Crosswork Optimization Engine.</p> <p>This release adds support for the following:</p> <ul style="list-style-type: none"> <li>• This feature allows the router to allocate the SR-Multicast Tree Segment Identifier (Tree-SID) Label automatically from the label range the operator configures. The operators can hence provision without passing the Label via North Bound API. Note that the label range configuration is mandatory for this feature to work.</li> <li>• If the label is not valid, an error message is returned to the application.</li> <li>• If the label is valid, the PCE computes, then program the Tree-SID multicast on the network nodes.</li> </ul> <p>For more information, refer to the <a href="#">Cisco Crosswork Optimization Engine User Guides</a>.</p>

Table 11: Feature History Table

Feature Name	Release Information	Feature Description
SR-PCE: Stateful North-Bound API for Policies Containing IPv4 Unnumbered Interfaces	Release 7.9.1	<p>The SR-PCE provides a north-bound HTTP-based API to allow communication between the SR-PCE and the Cisco Crosswork Optimization Engine.</p> <p>This release adds support for the following:</p> <ul style="list-style-type: none"> <li>• Reporting, provisioning, and updating of policies with IPv4 unnumbered interfaces in segment lists</li> <li>• Computing policy paths with IPv4 unnumbered interfaces in segment lists</li> <li>• Displaying IPv4 unnumbered hops in initiated tunnel database on PCE</li> </ul> <p>For more information, refer to the <a href="#">Cisco Crosswork Optimization Engine User Guides</a>.</p>

Table 12: Feature History Table

Feature Name	Release Information	Feature Description
SR-PCE: Stateful North-Bound API for Tree-SID	Release 7.5.1	<p>The SR-PCE provides a north-bound HTTP-based API to allow communication between the SR-PCE and the Cisco Crosswork Optimization Engine.</p> <p>This release adds stateful north-bound APIs to support real-time monitoring of Tree-SID states on the SR-PCE using a subscription model.</p> <p>For more information, refer to the <a href="#">Cisco Crosswork Optimization Engine User Guides</a>.</p>

Table 13: Feature History Table

Feature Name	Release Information	Feature Description
SR-PCE: North-Bound API for SRv6 and Flexible Algorithm in Cisco Optimization Engine (COE) v3.0 release	Release 7.3.2	<p>The SR-PCE provides a north-bound HTTP-based API to allow communication between the SR-PCE and the Cisco Crosswork Optimization Engine.</p> <p>This release adds support for the following:</p> <ul style="list-style-type: none"> <li>• Reporting of Flexible Algorithm participation and definitions</li> <li>• SRv6 topology information (nodes, links, Node uSIDs and Adj uSIDs)</li> <li>• SRv6 uSID list and uB6 SIDs allocated for a policy</li> </ul> <p>For more information, refer to the <a href="#">Cisco Crosswork Optimization Engine User Guides</a>.</p>

The SR-PCE provides a north-bound HTTP-based API to allow communication between SR-PCE and external clients and applications.

Over this API, an external application can leverage the SR-PCE for topology discovery, SR policy discovery, and SR policy instantiation.

The Cisco Crosswork Optimization Engine is an application that leverages the SR-PCE. For more information, refer to the [Cisco Crosswork Optimization Engine User Guides](#).

Use the following commands under PCE configuration mode to configure the API to allow communication between SR-PCE and external clients or applications.

Command	Description
<b>rest authentication basic</b>	(Optional) Specify basic (plaintext) authentication. By default, authentication is disabled.
<b>rest username</b> <i>password</i> { <b>clear</b>   <b>encrypted</b> } <i>password</i>	Add credentials when connecting to API.  <b>Note</b> This command is used only if authentication is configured.



Command	Description
<code>rest sibling ipv4 address</code>	<p>Opens a synchronization channel to another PCE in the same high availability (HA) pair.</p> <p><b>Note</b> For more information regarding SR-PCE HA pairs, refer to the <a href="#">Multiple Cisco SR-PCE HA Pairs</a> chapter of the <a href="#">Cisco Crosswork Optimization Engine 1.2.1 User Guide</a>.</p>
Command	Description
<code>api authentication {basic   digest}</code>	<p>Specify the type of authentication:</p> <ul style="list-style-type: none"> <li>• <b>basic</b> – Use HTTP Basic authentication (plaintext)</li> <li>• <b>digest</b> – Use HTTP Digest authentication (MD5)</li> </ul>
<code>api username password {clear   encrypted} password</code>	Add credentials when connecting to API.
<code>api sibling ipv4 address</code>	<p>Opens a synchronization channel to another PCE in the same high availability (HA) pair.</p> <p><b>Note</b> For more information regarding SR-PCE HA pairs, refer to the <a href="#">Multiple Cisco SR-PCE HA Pairs</a> chapter of the <a href="#">Cisco Crosswork Optimization Engine 1.2.1 User Guide</a>.</p>

### Example: Configuring API on SR-PCE

```
pce
address ipv4 10.1.1.100
rest
user admin
password encrypted 1304131F0202
!
authentication basic
sibling ipv4 10.1.1.200
!
!
end
```

```
pce
address ipv4 10.1.1.100
api
user admin
password encrypted 1304131F0202
!
authentication digest
sibling ipv4 10.1.1.200
!
```

```
!
end
```

The following example shows the current active connections:

```
RP/0/0/CPU0:pce1# show tcp brief | i 8080
Thu Aug  6 00:40:15.408 PDT
0xe9806fb8 0x60000000      0      0  :::8080          :::0          LISTEN
0xe94023b8 0x60000000      0      0  10.1.1.100:50487 10.1.1.200:8080 ESTAB
0xeb20bb40 0x60000000      0      0  10.1.1.100:8080  10.1.1.200:44401 ESTAB
0xe98031a0 0x60000000      0      0  0.0.0.0:8080    0.0.0.0:0    LISTEN
```

The first and fourth entries show the API server listening for IPv4 and IPv6 connections.

The second and third entries show the established sibling connection between PCE1 (10.1.1.100) and PCE2 (10.1.1.200).