Segment Routing Traffic Engineering Commands

This chapter describes the commands used to configure and use Segment Routing Traffic Enginering.

Note All commands applicable to the Cisco NCS 5500 Series Router are also supported on the Cisco NCS 540 Series Router that is introduced from Cisco IOS XR Release 6.3.2. References to earlier releases in Command History tables apply to only the Cisco NCS 5500 Series Router.

Note

- Starting with Cisco IOS XR Release 6.6.25, all commands applicable for the Cisco NCS 5500 Series Router are also supported on the Cisco NCS 560 Series Routers.
- Starting with Cisco IOS XR Release 6.3.2, all commands applicable for the Cisco NCS 5500 Series Router are also supported on the Cisco NCS 540 Series Router.
- References to releases before Cisco IOS XR Release 6.3.2 apply to only the Cisco NCS 5500 Series Router.
- Cisco IOS XR Software Release 7.0.1 specific updates are not applicable for the following variants of Cisco NCS 540 Series Routers:
  - N540-28Z4C-SYS-A
  - N540-28Z4C-SYS-D
  - N540X-16Z4G8Q2C-A
  - N540X-16Z4G8Q2C-D
  - N540X-16Z8Q2C-D
  - N540-12Z20G-SYS-A
  - N540-12Z20G-SYS-D
  - N540X-12Z16G-SYS-A
  - N540X-12Z16G-SYS-D

accounting prefixes ipv6 mode

To enable SRv6 traffic accounting, use the **accounting prefixes ipv6 mode** command in XR Config mode.

accounting prefixes ipv6 mode per-prefix per-nexthop srv6-locator

Syntax Description

per-prefix	Enables accounting for every prefix.
per-nexthop	Enables accounting for every prefix and nexthop.
srv6-locator	Enables accounting only for Segment-routing SRv6 locator.

Command Default

None

Command Modes

XR Config

Command History

Release	Modification
Release 7.10.1	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

The following example shows how to enable SRv6 traffic accounting:

```
Router(config)#accounting prefixes ipv6 mode per-prefix per-nexthop srv6-locators
```

affinity (SR-TE)

To configure a named interface link admin group by assigning affinity to an interface, use the **affinity name** *NAME* command in SR-TE interface submode.

affinity name *name*

Syntax Description

name	Affinity color name

Command Default

None

Command Modes

SR-TE interface

Command History

Release	Modification
Release 6.3.1	This command was introduced.

Usage Guidelines

Named Interface Link Admin Groups let you assign, or map, up to 32 color names for affinity and attribute-flag attributes instead of 32-bit hexadecimal numbers. After mappings are defined, the attributes can be referred to by the corresponding color name in the CLI.

Example

The following example shows how to assign affinity to interfaces:

```
Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# interface TenGigE0/0/1/2
Router(config-sr-if)# affinity
Router(config-sr-if-affinity)# name RED
```

affinity-map (SR-TE)

To define an affinity map, use the **affinity-map name** *name* **bit-position** *bit-position* command in SR-TE sub-mode.

affinity-map name *name* **bit-position** *bit-position*

Syntax Description

name *name*	Specify the name of the affinity-map.
bit-position *bit-position*	Specify the bit position in the Extended Admin Group bitmask. The *bit-position* range is from 0 to 255.

Command Default

None

Command Modes

SR-TE configuration

Command History

Release	Modification
Release 6.3.1	This command was introduced.

Usage Guidelines

Configure affinity maps on the following routers:

- Routers with interfaces that have an associated admin group attribute.
- Routers that act as SR-TE head-ends for SR policies that include affinity constraints.

Example

The following example shows how to define affinity maps.

```
Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# affinity-map
Router(config-sr-te-affinity-map)# name RED bit-position 23
```

autoroute include ipv6 all

To enable IPv6 autoroute support for SR-TE policies with IPv4 endpoints, use the **autoroute include ipv6 all** command in the SR-TE policy and PCC profile modes. To disable this feature, use the **no** form of this command.

autoroute include ipv6 all
no autoroute include ipv6 all

Syntax Description

This command has no keywords or arguments.

Command Default

IPv6 autoroute support is disabled.

Command Modes

SR-TE policy

PCC profile

Command History

Release	Modification
Release 7.3.4	This command was introduced.

Usage Guidelines

The **include ipv6 all** command form enables autoroute support for IPv6 prefixes, for a specified SR-TE policy. This command can be used in the SR-TE policy and PCC profile modes.

Example

The following example shows how to configure the IPv6 autoroute function for an SR-TE policy with an IPv4 endpoint:

```
Router# configure
Router(config)# segment-routing traffic-eng policy pol12
Router(config-sr-te-policy)# autoroute include ipv6 all
Router(config-sr-te-policy)# commit
```

The following example shows how to configure the IPv6 autoroute function for a PCE-instantiated SR-TE policy with an IPv4 endpoint:

```
Router# configure
Router(config)# segment-routing traffic-eng pcc profile 10
Router(config-pcc-prof)# autoroute include ipv6 all
Router(config-pcc-prof)# commit
```

bfd timers

To specify how long to wait for new BFD session to come up, use the **bfd timers** command in SR-TE sub-mode.

bfd timers session-bringup *seconds*

Syntax Description

seconds	Specify how long to wait for new BFD session to come up, in seconds. The range is from 10 to 3600.

Command Default

The default BFD session bring-up timer is 60 seconds.

Command Modes

SR-TE configuration

Command History

Release	Modification
Release 6.3.1	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Example

The following example shows how to configure the BFD session timer.

```
Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# bfd timers session-bringup 90
```

bgp bestpath igp-metric sr-policy

To configure BGP best path selection based on SR policy metrics in an SR-TE domain, use the **bgp bestpath igp-metric sr-policy** command in BGP configuration mode on the headend router. To remove the configuration, use the **no** form of the command.

bgp bestpath igp-metric sr-policy

Syntax Description

This command has no keywords or arguments.

Command Default

BGP best path selection based on SR policy metrics is disabled.

Command Modes

BGP configuration

Command History

Release	Modification
Release 7.3.2	This command was introduced.

Example

The following example shows how to configure BGP best path selection based on SR policy metrics (over IGP metric) in an SR-TE domain:

```
RR # configure
RR (config) # router bgp 100
RR (config-bgp)# bgp bestpath igp-metric sr-policy
RR (config-bgp)# commit
RR (config-bgp)# end
```

bgp prefix-path-label ignore

To indicate BGP to ignore the programming of the service route's prefix label when recursing onto the BSID of an SR-TE policy, use the **bgp prefix-path-label ignore** command in SR-TE policy steering config mode.

bgp prefix-path-label ignore

Syntax Description

This command has no keywords or arguments.

Command Default

None

Command Modes

SR-TE policy steering

Command History

Release	Modification
Release 7.9.1	This command was introduced.

Usage Guidelines

This command can be configured for manual SR policies.

Example

The following example shows how to configure BGP to ignore the programming of the service route's prefix label when recursing onto the BSID of an SR-TE policy:

```
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# policy POLICY1
Router(config-sr-te-policy)# steering
Router(config-sr-te-policy-steering)# bgp prefix-path-label ignore
```

binding-sid (SR-TE)

To specify the binding SID (BSID) allocation behavior, use the **binding-sid** command in SR-TE sub-mode.

binding-sid { **dynamic disable** | **explicit** { **enforce-srlb** | **fallback-dynamic** } }

Syntax Description

dynamic disable	Disables dynamic binding SID allocation. Candidate paths without an explicit BSID will be considered invalid.
explicit enforce-srlb	Specifies strict SRLB enforcement. If the BSID is not within the SRLB, the policy stays down.
explicitfallback-dynamic	Specifies that, if the BSID is not available, the BSID is allocated dynamically and the policy comes up.

Command Default

None

Command Modes

SR-TE configuration

Command History

Release	Modification
Release 6.3.2	This command was introduced.

Usage Guidelines

Explicit BSIDs are allocated from the segment routing local block (SRLB) or the dynamic range of labels. A best-effort is made to request and obtain this BSID for the SR-TE policy. If requested BSID is not available (if it does not fall within the available SRLB or is already used by another application or SR-TE policy), the policy stays down.

Use this command to specify how the BSID allocation behaves if the BSID value is not available.

Example

The following example shows how to specify how the BSID allocation behaves if the BSID value is not available.

Fallback to dynamic allocation:

```
Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# binding-sid explicit fallback-dynamic
```

Strict SRLB enforcement:

```
Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# binding-sid explicit enforce-srlb
```

distribute link-state

To enable reporting of SRTE policies, use the **distribute link-state** command in the SR-TE configuration mode.

distribute link-state [**report-candidate-path-inactive**]

Table 1: Syntax Description:

Syntax	Description
report-candidate-path-inactive	Enables reporting of SRTE policies using BGP-LS.

Command Default

The reporting of policies to BGP-LS is disabled by default.

Command Modes

SR-TE configuration (config-sr-te)

Command History

Release	Modification
Release 24.1.1	Supports reporting of SR-TE policies using BGP- Link State for SRv6.
Release 7.10.1	This command was introduced and supports reporting of SR-TE policies using BGP- Link State for SR-MPLS.

Task ID

Task ID	Operation
distribute link-state	write/read

Example

This example shows how to enable BGP-LS reporting and syncing of SRTE Policies:

```
Router# config
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# distribute link-state
Router(config-sr-te-distribute-ls)# report-candidate-path-inactive
Router(config-sr-te-distribute-ls)# exit
```

effective-metric

effective-metric admin-distance metric-type { **igp** | **te** | **latency** | **hopcount** | **unknown** } **admin-distance** *distance*

Syntax Description

admin-distance metric-type	Specify the metric type.
admin-distance *distance*	Specify the admin distance for the specified metric type.

Command Default

None

Command Modes

SR-TE configuration

Command History

Release	Modification
Release 6.3.1	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Example

```
Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# effective-metric admin-distance metric-type te admin-distance 15
```

interface

To to assign affinity and configure the TE metric for an interface, use the **interface** command in SR-TE submode.

interface *type interface-path-id* { **affinity name** *name* | **metric** *value* }

Syntax Description

type	Interface type. For more information, use the question mark (?) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all possible interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.
affinity name *name*	Specifies the affinity color name. Configure this on routers with interfaces that have an associated admin group attribute.
metric *value*	Specifies the traffic engineering (TE) metric. The range is from 0 to 2,147,483,647.

Command Default

None

Command Modes

SR-TE configuration

Command History

Release	Modification
Release 6.3.1	This command was introduced.

Usage Guidelines

Configure this on routers with interfaces that have an associated admin group attribute.

Example

The following example show how to assign affinity to an interface.

```
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# interface TenGigE0/0/1/2
Router(config-sr-if)# affinity
Router(config-sr-if-affinity)# name RED
```

The following example show how to configure the TE metric for an interface.

```
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# interface TenGigE0/0/1/2
Router(config-sr-te-if)# metric 50
```

kshortest-paths

To set the maximum number of attempts for SR-TE to compute paths that satisfy cumulative metric bounds criteria, use the **kshortest-paths** command in SR-TE configuration mode. To revert to the default number of attempts (100), use the **no** form of the command.

kshortest-paths *max-attempts*

no kshortest-paths

Syntax Description

max-attempts	Maximum number of attempts. Choose a value between 1 and 200.

Command Default

100 attempts are made to compute paths that satisfy the cumulative metric bounds criteria.

Command Modes

SR-TE configuration (config-sr-te)

Command History

Release	Modification
Release 7.3.1	This command was introduced.

Usage Guidelines

By default, a maximum of 100 attempts are made. To update the value, you can use this command.

You can use the **show segment-routing traffic-eng policy color** command (**Number of K-shortest-paths** field) to see the K-shortest path algorithm computation result. For example, if the **Number of K-shortest-paths** field displays 4, it means that the K-shortest path algorithm took 4 computations to find the right path. The 4 shortest paths that are computed using K-shortest path algorithm did not respect the cumulative bounds, and the fifth shortest path was valid against the bounds.

Example

This example shows how to set the maximum number of attempts for computing paths that satisfy the cumulative metric bounds criteria:

```
Router# configure terminal
Router(config)# segment-routing traffic-eng
Router(config-sr-te)# kshortest-paths 120
Router(config-sr-te)# commit
```

logging

To enable SYSLOG alarms related to PCEP peer-status and SR-TE policies, use the **logging** command in SR-TE submode.

logging { **pcep peer-status** | **policy status** }

Syntax Description

pcep peer-status	Enables PCEP peer status SYSLOG alarms.
policy status	Enables SR-TE related SYSLOG alarms.

Command Default

None

Command Modes

SR-TE configuration

Command History

Release	Modification
Release 6.3.1	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Example

The following example shows how to enable logging for SR-TE policies.

```
Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# logging policy status
```

maximum-sid-depth

To customize the maximum number of SIDs advertised by the router or signaled by the PCC during PCEP session establishment, use the **maximum-sid-depth** command in SR-TE sub-mode or SR-TE ODN sub-mode.

maximum-sid-depth *value*

Syntax Description

value	Specifies the maximum number of SIDs advertised by the router or signaled by the PCC during PCEP session establishment. The range is from 1 to 255.

Command Default

The default MSD *value* is equal to the maximum MSD supported by the platform (555).

Command Modes

SR-TE configuration

SR-TE On-Demand Next-Hop (SR-ODN) configuration

Command History

Release	Modification
Release 6.3.2	This command was introduced.

Usage Guidelines

The default MSD *value* is equal to the maximum MSD supported by the platform (555).

Note The platform's SR-TE label imposition capabilities are as follows:

- Up to 5 transport labels when no service labels are imposed
- Up to 3 transport labels when service labels are imposed
- Up to 5 transport labels when no service labels are imposed
- Up to 3 transport labels when service labels are imposed
- Up to 5 transport labels when no service labels are imposed
- Up to 3 transport labels when service labels are imposed

For cases with path computation at PCE, a PCC can signal its MSD to the PCE in the following ways:

- During PCEP session establishment – The signaled MSD is treated as a node-wide property.
  - MSD is configured under **segment-routing traffic-eng maximum-sid-depth** *value* command.
- During PCEP LSP path request – The signaled MSD is treated as an LSP property.
  - On-demand (ODN) SR Policy: MSD is configured using the **segment-routing traffic-eng on-demand color** *color* **maximum-sid-depth** *value* command.

Note If the configured MSD values are different, the per-LSP MSD takes precedence over the per-node MSD.

After path computation, the resulting label stack size is verified against the MSD requirement.

- If the label stack size is larger than the MSD and path computation is performed by PCE, then the PCE returns a "no path" response to the PCC.
- If the label stack size is larger than the MSD and path computation is performed by PCC, then the PCC will not install the path.

Note A sub-optimal path (if one exists) that satisfies the MSD constraint could be computed in the following cases:

- For a dynamic path with TE metric, when the PCE is configured with the **pce segment-routing te-latency** command or the PCC is configured with the **segment-routing traffic-eng te-latency** command.
- For a dynamic path with LATENCY metric
- For a dynamic path with affinity constraints

For example, if the PCC MSD is 4 and the optimal path (with an accumulated metric of 100) requires 5 labels, but a sub-optimal path exists (with accumulated metric of 110) requiring 4 labels, then the sub-optimal path is installed.

Example

The following example shows how to configure the MSD during PCEP session establishment. The signaled MSD is treated as a node-wide property:

```
RP/0/RSP0/CPU0:ios(config)# segment-routing
RP/0/RSP0/CPU0:ios(config-sr)# traffic-eng
RP/0/RSP0/CPU0:ios(config-sr-te)# maximum-sid-depth 4
```

The following example shows how to configure the MSD during PCEP LSP path request for the On-demand (ODN) SR Policy. The signaled MSD is treated as an LSP property:

```
RP/0/RSP0/CPU0:ios(config)# segment-routing
RP/0/RSP0/CPU0:ios(config-sr)# traffic-eng
RP/0/RSP0/CPU0:ios(config-sr-te)# on-demand color 250
RP/0/RSP0/CPU0:ios(config-sr-te-color)# maximum-sid-depth 4
```

max-install-standby-cpaths

To configure standby candidate paths for all SR policies, for a specific policy, or for an ODN template, use the **max-install-standby-cpaths** command.

To disable the configuration for global SR policies, use the **no** form of this command.

max-install-standby-cpaths *value*

Syntax Description

value	Specifies the number of non-active CPs to program in forwarding. The range for *value* is from 1 to 3 for global SR policies, and from 0 (disable) to 3 for local and ODN policies.

Command Default

None

Command Modes

SR-TE configuration

SR-TE Policy configuration

SR-TE On-Demand Next-Hop (SR-ODN) configuration

Command History

Release	Modification
Release 7.6.1	This command was introduced.

Usage Guidelines

- Up to three non-active CPs can be programmed in the forwarding plane.
- Manually configured CPs are supported. This includes CPs with explicit paths or dynamic (head-end computed or PCE-delegated) paths.
- On-Demand instantiated CPs (ODN) are supported.
- BGP-initiated CPs are supported.
- PCE-initiated CPs via PCEP are not supported.
- Programming of non-active CPs is not supported with SRv6-TE policies, Per-Flow Policies (PFP), or point-to-multipoint SR policies (Tree-SID)
- PCEP reporting of additional CPs is supported, but the PCEP reporting does not distinguish between active and non-active CPs.
- Programming of non-active CPs can be enabled for all SR policies (global), for a specific policy (local), or ODN template.

  If enabled globally and locally or on ODN template, the local or ODN configuration takes precedence over the global configuration.

  - Programming of non-active CPs under global SR-TE and configuring policy path protection of an SR policy is supported. In this case, policy path protection takes precedence.
  - Programming of non-active CPs for a specific SR policy and configuring policy path protection of an SR policy is not supported.

- The number of policies supported could be impacted by the number of non-active CPs per policy. Programming non-active CPs in the forwarding plane consumes hardware resources (such as local label and ECMP FEC) when more candidate paths are pre-programmed in forwarding than are actually carrying traffic.
- The active CP will be in programmed state. The remaining CPs will be in standby programmed state.
- We recommend that you create separate PM sessions for active and standby candidate paths to monitor the health of the paths end-to-end.
- The protected paths for each CP is programmed in the respective LSPs. The protected paths of active CPs are programmed in the active LSP, and the protected paths of standby CPs are programmed in the standby LSP.
- If a candidate path with higher preference becomes available, the traffic will switch to it in Make-Before-Break (MBB) behavior.

Example

The following example shows how to configure standby candidate paths globally:

```
Router(config)# segment-routing traffic-eng
Router(config-sr-te)# max-install-standby-cpaths 2
Router(config-sr-te)#
```

The following example shows how to configure standby candidate paths for a specific SR policy:

```
Router(config)# segment-routing traffic-eng
Router(config-sr-te)# policy MyBackupPolicy
Router(config-sr-te-policy)# max-install-standby-cpaths 2
Router(config-sr-te-policy)#
```

The following example shows how to configure standby candidate paths for an SR ODN template:

```
Router(config)# segment-routing traffic-eng
Router(config-sr-te)# on-demand color 10
Router(config-sr-te-color)# max-install-standby-cpaths 1
Router(config-sr-te-color)#
```

The following example shows how to enable three standby CPs globally and disable standby CPs on SR policy and ODN template:

```
Router(config)# segment-routing traffic-eng
Router(config-sr-te)# max-install-standby-cpaths 3
Router(config-sr-te)# policy MyBackupPolicy
Router(config-sr-te-policy)# max-install-standby-cpaths 0
Router(config-sr-te-policy)# exit
Router(config-sr-te)# on-demand color 10
Router(config-sr-te-color)# max-install-standby-cpaths 0
Router(config-sr-te-color)#
```

max-metric

Use the **max-metric** command in the SR-TE sub-mode to set the protocol advertising maximum metric. This will render the router as a less preferable intermediate hop for other routers.

maximum-metric *default-route delay external interlevel level on-startup srv6-locator te*

Syntax Description

default-route	Override the default route metric with maximum metric.
delay	Apply max metric to delay metric.
external	Override metric of prefixes learned from another protocol with maximum metric.
interlevel	Override metric of prefixes learned from another ISIS level with maximum metric.
level	Set maximum metric for one level only.
on-startup	Set maximum metric temporarily after reboot.
srv6-locator	Override segment routing ipv6 locator metric with maximum metric.
te	Apply max-metric to TE metric.

Command Modes

SR-TE configuration

Command History

Release	Modification
Release 7.6.1	This command was introduced.
Release 7.8.1	This command was modified.

Example

The following example shows how to set the maximum metric for the SR-TE:

```
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# max-metric delay te
Router(config-sr-te)# commit

Router(config-sr-te)# #sh isis da de r100

IS-IS 1 (Level-2) Link State Database
LSPID    LSP Seq Num    LSP Checksum    LSP Holdtime/Rcvd    ATT/P/OL
F100.00.00  * 0x000000a    0x79ab        1190 /*              0/0/0
    Area Address:    49.0001
    LSP MTU:         1350
    NLPID:           0xcc
    NLPID:           0x8e
    MT:              Standard (IPv4 Unicast)
    MT:              IPv6 Unicast
```

```
IP Address:      2020:1000::100
Hostname:        100
Router Cap:      20.1.0.100 D:0 S:0
Metric: 16777214        IS-Extended r101.00
Metric: 16777214        IS-Extended r101.00
Metric: 16777214        MT (IPv6 Unicast) IS-Extended r101.00
Metric: 16777214        MT (IPv6 Unicast) IS-Extended r103.00
Metric: 16777214        IP-Extended 6.6.6.100/32
Metric: 16777214        IP-Extended 10.1.1.0/24
Metric: 16777214        IP-Extended 10.4.1.0/24
Metric: 16777214        IP-Extended 20.1.0.100/32
Metric: 16777214        MT (IPv6 Unicast) IPv6 2001:1000::/64
Metric: 16777214        MT (IPv6 Unicast) IPv6 2004:1000::/64
Metric: 16777214        MT (IPv6 Unicast) IPv6 2020:1000::100/128
Metric: 16777214        MT (IPv6 Unicast) IPv6 6060:1000::100/128
```

nexthop validation color-extcomm disable

To disable BGP Next-Hop validation on the route reflector in an SR-TE domain, use the **nexthop validation color-extcomm disable** command in BGP configuration mode. To remove the configuration, use the **no** form of the command.

nexthop validation color-extcomm disable

Syntax Description

This command has no keywords or arguments.

Command Default

BGP NH validation is not disabled in an SR-TE domain.

Command Modes

BGP configuration

Command History

Release	Modification
Release 7.3.2	This command was introduced.

Usage Guidelines

To fully enable Next-Hop soft validation for SR policy-installed routes, do the following:

- On the headend router, enable **nexthop validation color-extcomm sr-policy**
- On the route reflector, enable **nexthop validation color-extcomm disable**

Note BGP NH soft validation is enabled on the headend router while the usual BGP NH validation is disabled on the RR.

Example

The following example shows how to disable BGP Next-Hop validation on a RR in an SR-TE domain:

```
Headend # configure
Headend (config) # router bgp 100
Headend (config-bgp)# nexthop validation color-extcomm disable
Headend (config-bgp)# commit
Headend (config-bgp)# end
```

nexthop validation color-extcomm sr-policy

To enable BGP Next-Hop soft validation in an SR-TE domain, use the **nexthop validation color- extcomm sr-policy** command in BGP configuration mode.

nexthop validation color-extcomm sr-policy

Syntax Description

This command has no keywords or arguments.

Command Default

BGP NH validation is disabled.

Command Modes

BGP configuration

Command History

Release	Modification
Release 7.3.2	This command was introduced.

Usage Guidelines

To fully enable Next-Hop soft validation for SR policy-installed routes, do the following:

- On the headend router, enable nexthop validation color-extcomm sr-policy
- On the route reflector, enable nexthop validation color-extcomm disable

Note BGP NH soft validation is enabled on the headend router while the usual BGP NH validation is disabled on the RR.

Example

The following example shows how to configure BGP Next-Hop soft validation on the headend router in an SR-TE domain:

```
Headend # configure
Headend (config) # router bgp 100
Headend (config-bgp)# nexthop validation color-extcomm sr-policy
Headend (config-bgp)# commit
Headend (config-bgp)# end
```

Use this command to view BGP Soft Next-Hop Validation details.

```
Headend # show bgp process detail | i Nexthop

Use SR-Policy admin/metric of color-extcomm Nexthop during path comparison: enabled
ExtComm Color Nexthop validation: SR-Policy then RIB.
```

on-demand constraints

To configure the SR Flexible Algorithm constraints, use the **constraints segments sid-algorithm** command in SR-TE sub-mode.

To specify resource constraints for path computation for ODN SR-TE policies, use the **constraints resources** command in SR-TE configuration mode.

on-demand color *color* **constraints** { **segments sid-algorithm** *algo* | **resources** { **exclude resource-list** *name* | **exclude-group** *group_name* | **apply-group** *group_name* } }

Syntax Description

segments	Specify constraints for segments of a path in a network.
sid-algorithm *algo*	Specify the SR Flexible Algorithm value. The *algo* range is from 128 to 255.
resources	Specify resource constraints for path computation.
exclude	Exclude resources from path computation.
resource-list *name*	Specify the name of the resource-list to exclude from the path computation.

Command Default

None

Command Modes

SR-TE configuration

Command History

Release	Modification
Release 24.1.1	The resources option was introduced.
Release 7.9.1	For Cisco IOS XR Release 7.9.1, you must reconfigure all SR-ODN configurations with Flexible Algorithm constraints that use the on-demand dynamic sid-algorithm with the on-demand constraints command.
Release 7.4.1	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Example

The following example shows how to add an SR Flexible Algorithm constraint:

```
Router(config-sr-te-color)#constraints segments sid-algorithm 128
```

The following example shows how 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
Routerconfig-sr-te-color)#constraints resources exclude resource-list node_resc_list
```

on-demand dynamic affinity

To configure the affinity constraints for dynamic ODN paths, use the **on-demand dynamic affinity** command in SR-TE sub-mode.

on-demand color *color* **dynamic affinity** { **include-all** | **include-any** | **exclude-any** } [**name** *name*]

Syntax Description

affinity {**include-all** | **include-any** | **exclude-any**}	Specify the affinity type.
name *name*	Name of the affinity.

Command Default

None

Command Modes

SR-TE configuration

Command History

Release	Modification
Release 6.3.1	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Example

The following example shows how to configure the affinity contraints .

```
Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# on-demand color 10 dynamic
Router(config-sr-te-color-dyn)# affinity include-all name CROSS
Router(config-sr-te-color-dyn)#
```

on-demand dynamic bounds

To configure SR-TE ODN to calculate a shortest path with cumulative metric bounds, use the **on-demand dynamic bounds** command in SR-TE sub-mode.

on-demand color *color* **bounds cumulative type** { **hopcount** | **igp** | **latency** | **te** } *metric*

Syntax Description

type {**hopcount** | **igp** | **latency** | **te**}	Specify the metric type.
metric	Specify the bound metric value. Valid values are from 1 to 4294967295.

Command Default

None

Command Modes

SR-TE configuration

Command History

Release	Modification
Release 7.3.1	This command was introduced.

Usage Guidelines

When an SR policy is configured on a head-end node with these metric bounds, a path is finalized towards the specified destination only if it meets each of these criteria.

PCE-based cumulative metric bounds computations are not supported. You must use non-PCE (SR-TE topology) based configuration for path calculation, for cumulative bounds.

If you use PCE dynamic computation configuration with cumulative bounds, the PCE computes a path and validates against cumulative bounds. If it is valid, then the policy is created with this path on PCC. If the initial path doesn't respect the bounds, then the path is not considered, and no further K-shortest path algorithm is executed to find the path.

Example

The following example shows how to configure IGP, TE, hop count, and latency metric bounds for the SR-ODN color template:

```
Router(config-sr-te)# on-demand color 1000 dynamic
Router(config-sr-te-color-dyn) bounds cumulative
Router(config-sr-te-odc-bounds-type)# type igp 100
Router(config-sr-te-odc-bounds-type)# type te 60
Router(config-sr-te-odc-bounds-type)# type hopcount 6
Router(config-sr-te-odc-bounds-type)# type latency 1000
```

on-demand dynamic disjoint-path

To configure the disjoint-path constraints, use the **on-demand dynamic disjoint-path** command in SR-TE sub-mode.

on-demand color *color* **dynamic disjoint-path group-id** *id* **type** { **link** | **node** | **srlg** | **srlg-node** } [{ **sub-id** *sub_id* | **fallback disable** }]

Syntax Description

group-id *id*	Specify the group ID of the disjoint path. Valid values are from 1 to 65535.
type {**link** | **node** | **srlg** | **srlg-node** }	Specify the type of disjointness.
sub-id *id*	Specify the sub-group ID of the disjoint path. Valid values are from 1 to 65535.
fallback disable	Disable all fallback behavior in case the requested disjointness cannot be achieved.

Command Default

None

Command Modes

SR-TE configuration

Command History

Release	Modification
Release 24.1.1	The **fallback disable** keyword was introduced.
Release 6.3.1	This command was introduced.

Usage Guidelines

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):

- 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.

If a pair of paths that meet the requested disjointness level cannot be found, then the paths will automatically fallback to a lower level:

- If the requested disjointness level is SRLG or node, then link-disjoint paths will be computed.
- 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.

Example

```
Router(config-sr-te-color-dyn)# disjoint-path group-id 775 type link
```

The following example indicates how to configure strict disjointness for an ODN SR-TE policy:

```
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
```

on-demand dynamic metric

To configure the On-Demand dynamic path metric, use the **on-demand dynamic metric** command in SR-TE sub-mode.

on-demand color *color* **dynamic metric** { **margin** { **absolute** *value* | **relative** *percent* } *margin* | **type** { **hopcount** | **igp** | **latency** | **te** } }

Syntax Description

metric {**absolute** *value* | **relative** *percent*}	Specify the On-Demand dynamic path metric margin. The range for *margin* and *percent* is from 0 to 2147483647.
type { **hopcount** | **igp** | **latency** | **te** }	Specify the metric type for use in path computation.

Command Default

None

Command Modes

SR-TE configuration

Command History

Release	Modification
Release 6.3.1	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Example

```
Router(config-sr-te-color-dyn)# metric type te
Router(config-sr-te-color-dyn)# metric margin absolute 5
```

on-demand dynamic pcep

To indicate that only the path computed by SR-PCE should be associated with the on-demand SR policy, use the **on-demand dynamic pcep** command in SR-TE sub-mode.

on-demand color *color* **dynamic pcep**

Syntax Description

This command has no keywords or arguments.

Command Default

None

Command Modes

SR-TE configuration

Command History

Release	Modification
Release 6.3.1	This command was introduced.

Usage Guidelines

With this configuration, local path computation is not attempted; instead the head-end router will only instantiate the path computed by the SR-PCE.

Example

```
Router(config-sr-te)# on-demand color 10 dynamic pcep
```

on-demand dynamic sid-algorithm

Note For Cisco IOS XR Release 7.9.1, you must reconfigure all SR-ODN configurations with Flexible Algorithm constraints that use the on-demand dynamic sid-algorithm with the on-demand constraints command.

To configure the SR Flexible Algorithm constraints, use the **on-demand dynamic sid-algorithm** command in SR-TE sub-mode.

on-demand color *color* **dynamic sid-algorithm** *algo*

Syntax Description

sid-algorithm *algo*	Specify the SR Flexible Algorithm value . The *algo* range is from 128 to 255.

Command Default

None

Command Modes

SR-TE configuration

Command History

Release	Modification
Release 6.3.1	This command was introduced.
Release 7.4.1	This command was replaced by the **on-demand constraints** command.
Release 7.9.1	For Cisco IOS XR Release 7.9.1, you must reconfigure all SR-ODN configurations with Flexible Algorithm constraints that use the **on-demand dynamic sid-algorithm** with the **on-demand constraints** command.

Usage Guidelines

This command was replaced by the on-demand constraints command.

Example

```
Router(config-sr-te-color-dyn)# sid-algorithm 128
```

on-demand maximum-sid-depth

To customize the maximum SID depth (MSD) constraints advertised by the router, use the **on-demand maximum-sid-depth** command in SR-TE sub-mode.

on-demand color *color* **maximum-sid-depth** *value*

Syntax Description

maximum-sid-depth *value*	Specify the maximum SID depth. The range of *value* is 1 to 255.

Command Default

The default MSD *value* is equal to the maximum MSD supported by the platform (555).

Command Modes

SR-TE configuration

Command History

Release	Modification
Release 7.0.1	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Example

```
Router(config-sr-te-color)# maximum-sid-depth 5
```

on-demand steering

on-demand color *color* **steering** { **labeled-services disable** | **path-invalidation drop** }

Syntax Description

labeled-services disable	Disable steering of labeled-services for on-demand color policies. This configuration applies for a specific ODN color.
path-invalidation drop	Drop traffic but keep the SR policy up in the control plane.

Command Default

None

Command Modes

SR-TE configuration

Command History

Release	Modification
Release 7.0.1	This command was introduced.
Release 7.4.1	The **path-invalidation drop** keywords are introduced.

Usage Guidelines

- **labeled-services disable**: The SR-TE MPLS Label Imposition Enhancement feature increases the maximum label imposition capabilities of the platform.

  In previous releases, the platform supported:

  - Up to 5 MPLS transport labels when no MPLS service labels are imposed
  - Up to 3 MPLS transport labels when MPLS service labels are imposed

  With the SR-TE MPLS Label Imposition Enhancement feature, the platform supports the following:

  - Up to 12 MPLS transport labels when no MPLS service labels are imposed
  - Up to 9 MPLS transport labels when MPLS service labels are imposed

  This enhancement is enabled and disabled dynamically, as the label count changes. For example, if a path requires only 3 MPLS transport labels, the MPLS Label Imposition Enhancement feature is not enabled.

  You can disable labeled services for SR-TE policies. The label switching database (LSD) needs to know if labeled services are disabled on top of an SR-TE policy to perform proper label stack splitting.

- **path-invalidation drop**:

  By default, if an SR Policy becomes invalid, traffic would fall back to the native SR forwarding path. In some scenarios, a network operator may require that certain traffic be only carried over the path associated with an SR policy and never allow the native SR LSP to be used. This command is introduced to meet this requirement.

  With **path-invalidation drop** enabled, an SR policy that would become invalid (for example, no valid candidate path available) is programmed to drop traffic. At the same time, the SR policy stays up in the control plane to prevent prefixes mapped to the SR policy from falling back to the native SR LSP.

When the SR policy becomes valid again, forwarding over the SR policy resumes.

Example

The following example shows how enable the dropping of traffic when an On-Demand SR Policy becomes invalid.

```
Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# on-demand color 10
Router(config-sr-te-color)# steering
Router(config-sr-te-on-demand-color-steering)# path-invalidation drop
```

The following example shows how to disable steering of labeled-services for on-demand color policies:

```
Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# on-demand color 10
Router(config-sr-te-color)# steering
Router(config-sr-te-on-demand-color-steering)# labeled-services disable
```

path-invalidation drop

To enable the dropping of traffic when an SR Policy becomes invalid, use the **path-invalidation drop** command.

policy *policy* **steering path-invalidation drop**

on-demand color *color* **steering path-invalidation drop**

pcc profile *profile* **steering path-invalidation drop**

Syntax Description

This command has no keywords or arguments.

Command Default

Disabled

Command Modes

SR-TE Policy

SR-TE ODN

SR-TE PCC

Command History

Release	Modification
Release 7.4.1	This command was introduced.

Usage Guidelines

By default, if an SR Policy becomes invalid, traffic would fall back to the native SR forwarding path. In some scenarios, a network operator may require that certain traffic be only carried over the path associated with an SR policy and never allow the native SR LSP to be used. This command is introduced to meet this requirement.

With **path-invalidation drop** enabled, an SR policy that would become invalid (for example, no valid candidate path available) is programmed to drop traffic. At the same time, the SR policy stays up in the control plane to prevent prefixes mapped to the SR policy from falling back to the native SR LSP.

When the SR policy becomes valid, forwarding over the SR policy resumes.

Example

The following example shows how enable the dropping of traffic when an SR Policy becomes invalid.

```
Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# on-demand color 10
Router(config-sr-te-color)# steering
Router(config-sr-te-on-demand-color-steering)# path-invalidation drop
```

The following example shows how enable the dropping of traffic when an On-Demand SR Policy becomes invalid.

```
Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
```

```
Router(config-sr-te)# policy FOO
Router(config-sr-te-policy)# steering
Router(config-sr-te-policy-steering)# path-invalidation drop
```

The following example shows how enable the dropping of traffic when a PCE-Initiated SR Policy becomes invalid.

```
Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# pcc profile 7
Router(config-pcc-prof)# steering
Router(config-pcc-prof-steering)# path-invalidation drop
```

pcc pce address

To configure the SR-PCE address and options, use the **pcc pce address** command in SR-TE configuration mode.

pcc pce address ipv4 *address* [**keychain** *word* | **password** { **clear** | **encrypted** } *password* | **precedence** *0-255* | **tcp-ao** *word* [**include-tcp-options**]]

Syntax Description

keychain *keychain-name*	Configures keychain based authentication for PCC
password {**clear**	**encrypted**} *password*	Configures password for MD5 authentication
precedence *precedence*	Specifies the precedence for the PCC peer. The value range is from 0 to 255.
tcp-ao *tcp-ao-keychain-name*	Configures AO keychain based authentication
include-tcp-options	Includes other TCP options in the header.

Command Default

None

Command Modes

SR-TE configuration

Command History

Release	Modification
Release 6.3.1	This command was introduced.

Usage Guidelines

A PCE can be given an optional precedence. If a PCC is connected to multiple PCEs, the PCC selects a PCE with the lowest precedence value. If there is a tie, a PCE with the highest IP address is chosen for computing path. The precedence value range is from 0 to 255.

Example

The following shows how to configure the SR-PCE address.

```
Router(config)# segment-routing traffic-engineering
Router(config-sr-te)# pcc pce address ipv4 1.1.1.2 precedence 250
```

pcc report-all

To enable the PCC to report all SR policies in its database to the PCE, use the **pcc report-all** command in SR-TE configuration mode.

pcc report-all

Syntax Description This command has no keywords or arguments.

Command Default None

Command Modes SR-TE configuration

Command History

Release	Modification
Release 6.3.1	This command was introduced.

Usage Guidelines No specific guidelines impact the use of this command.

Example

The following example shows how to enable the PCC to report all SR policies in its database to the PCE:

```
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# pcc report-all
```

pcc source-address

To configure the PCC source address, use the **pcc source-address** command in SR-TE configuration mode.

pcc source-address ipv4 *address*

Syntax Description

address	Specifies the local IPv4 address of the PCC.

Command Default

None

Command Modes

SR-TE configuration

Command History

Release	Modification
Release 6.3.1	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Example

The following example shows how to configure the PCC source address:

```
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# pcc source-address ipv4 1.1.1.4
```

pcc timers

To configure PCEP-related timers, use the **pcc timers** command in SR-TE configuration mode.

pcc timers { **deadtimer** *seconds* | **delegation-timeout** *seconds* | **initiated** { **orphan** *seconds* | **state** *seconds* } | **keepalive** *seconds* }

Syntax Description

deadtimer *seconds*	Specifies how long the remote peers wait before bringing down the PCEP session if no PCEP messages are received from this PCC. The range is from 1 to 255 seconds.
delegation-timeout *seconds*	Specifies how long a delegated SR policy can remain up without an active connection to a PCE. The range is from 0 to 3600 seconds.
initiated orphan *seconds*	Specifies the amount of time that a PCE-initiated SR policy will remain delegated to a PCE peer that is no longer reachable by the PCC. The range is from 10 to 180 seconds.
initiated state *seconds*	Specifies the amount of time that a PCE-initiated SR policy will remain programmed while not being delegated to any PCE. The range is from 15 to 14440 seconds (24 hours).
keepalive *seconds*	Specifies how often keepalive messages are sent from PCC to its peers. The range is from 0 to 255 seconds.

Command Default

Deadtimer: 120 seconds

Delegation timeout: 60 seconds

Initiated orphan: 180 seconds

Initiated state: 600 seconds

Keepalive: 30 seconds

Command Modes

SR-TE configuration

Command History

Release	Modification
Release 6.3.1	This command was introduced.

Usage Guidelines

To better understand how the PCE-initiated SR policy timers operate, consider the following example:

1. PCE A instantiates SR policy P at head-end N.
2. Head-end N delegates SR policy P to PCE A and programs it in forwarding.
3. If head-end N detects that PCE A is no longer reachable, then head-end N starts the PCE-initiated orphan and state timers for SR policy P.
4. If PCE A reconnects before the orphan timer expires, then SR policy P is automatically delegated back to its original PCE (PCE A).

5. After the orphan timer expires, SR policy P will be eligible for delegation to any other surviving PCE(s).
6. If SR policy P is not delegated to another PCE before the state timer expires, then head-end N will remove SR policy P from its forwarding

Example

```
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# pcc
Router(config-sr-te-pcc)# timers keepalive 20
Router(config-sr-te-pcc)# timers deadtimer 60
Router(config-sr-te-pcc)# timers delegation-timeout 30
Router(config-sr-te-pcc)# timers initiated orphan 60
Router(config-sr-te-pcc)# timers initiated state 1200
```

policy bfd

To enable SBFD on an SR-TE policy or an SR on-demand (SR-ODN) color template and enter BFD configuration mode, use the **policy bfd** command in SR-TE configuration mode

policy *policy* **bfd** { **disable** | **invalidation-action** { **down** | **none** } | **logging session-state-change** | **minimum-interval** *interval* | **multiplier** *multiplier* | **reverse-path binding-label** *label* }

Syntax Description

disable	Disables BFD session.
invalidation-action {**down** | **none**}	Specifies the action to be taken when BFD session is invalidated. • **down**: LSP can only be operationally up if the BFD session is up. • **none**: BFD session state does not affect LSP state, use for diagnostic purposes
loggingsession-state-change	Displays a syslog when the state of the session changes.
minimum-interval *interval*	Specifies the interval between sending BFD hello packets to the neighbor. The range is from 50 to 30000 milliseconds.
multiplier *multiplier*	Specifies the number of times a packet is missed before BFD declares the neighbor down. The range is from 2 to 10.
reverse-path binding-label *label*	(SR-TE policy only) Spcifies BFD packets return to head-end by using a binding label.

Command Default

minimum-interval = 150

multiplier = 3

Command Modes

SR-TE policy

SR-TE ODN

Command History

Release	Modification
Release 7.0.1	This command was introduced.

Usage Guidelines

Do not use BFD with disjoint paths. The reverse path might not be disjoint, causing a single link failure to bring down BFD sessions on both the disjoint paths.

reverse-path binding-label: (**SR-TE policy only**) Use the **reverse-path binding-label** label command to specify BFD packets return to head-end by using a binding label.

By default, the S-BFD return path (from tail-end to head-end) is via IPv4. You can use a reverse binding label so that the packet arrives at the tail-end with the reverse binding label as the top label. This label is meant to point to a policy that will take the BFD packets back to the head-end. The reverse binding label is configured per-policy.

Note that when MPLS return path is used, BFD uses echo mode packets, which means the tail-end's BFD reflector does not process BFD packets at all.

The MPLS label value at the tail-end and the head-end must be synchronized by the operator or controller. Because the tail-end binding label should remain constant, configure it as an explicit BSID, rather than dynamically allocated.

Example

The following example shows how to enable SBFD on an SR-TE policy:

```
Router(config)# segment-routing traffic-eng
Router(config-sr-te)# policy POLICY1
Router(config-sr-te-policy)# bfd
Router(config-sr-te-policy-bfd)# invalidation-action down
Router(config-sr-te-policy-bfd)# minimum-interval 250
Router(config-sr-te-policy-bfd)# multiplier 5
Router(config-sr-te-policy-bfd)# reverse-path binding-label 24036
Router(config-sr-te-policy-bfd)# logging session-state-change
```

The following example shows how to enable SBFD on an SR-ODN color:

```
Router(config)# segment-routing traffic-eng
Router(config-sr-te)# on-demand color 10
Router(config-sr-te-color)# bfd
Router(config-sr-te-color-bfd)# minimum-interval 250
Router(config-sr-te-color-bfd)# multiplier 5
Router(config-sr-te-color-bfd)# logging session-state-change
Router(config-sr-te-color-bfd)# invalidation-action down
```

policy binding-sid mpls

To specify the explicit BSID, use the **policy binding-sid mpls** command in SR-TE policy mode.

binding-sid mpls *label*

Syntax Description

label	Explicit binding SID label

Command Default

None

Command Modes

SR-TE policy

Command History

Release	Modification
Release 6.3.1	This command was introduced.

Usage Guidelines

Explicit BSIDs are allocated from the segment routing local block (SRLB) or the dynamic range of labels. A best-effort is made to request and obtain the BSID for the SR-TE policy. If requested BSID is not available (if it does not fall within the available SRLB or is already used by another application or SR-TE policy), the policy stays down.

Example

The following example shows how to configure an SR policy to use an explicit BSID of 1000:

```
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# policy FOO
Router(config-sr-te-policy)# binding-sid mpls 1000
```

policy candidate-paths constraints affinity

To configure affiity constraints on an SR-TE policy, use the **policy candidate-paths constraints affinity** command in SR-TE configuration mode.

policy *policy* **candidate-paths preference** *preference* **constraints affinity** { **include-all** | **include-any** | **exclude-any** } **name** *name*

Syntax Description

policy *policy*	Specifies the name of the policy.
candidate-paths preference *preference*	Configures the candidate path preference. The range is from 1 to 65535.
constraints affinity {**include-allinclude-anyexclude-any**}	Configures the affinity constraints.
name *name*	Specifies the affinity name.

Command Default

None

Command Modes

SR-TE policy

Command History

Release	Modification
Release 6.3.1	This command was introduced.

Usage Guidelines

The candidate path with the highest preference is the active candidate path (highlighted below) and is installed in forwarding.

You can apply a color or name to links or interfaces by assigning affinity bit-maps to them. You can then specify an affinity (or relationship) between an SR policy path and link colors. SR-TE computes a path that includes or excludes links that have specific colors,or combinations of colors

Example

The following example shows how to associate affinity constraints for an SR-TE policy:

```
Router(config-sr-te)# policy POLICY1
Router(config-sr-te-policy)# color 20 end-point ipv4 1.1.1.4
Router(config-sr-te-policy)# candidate-paths
Router(config-sr-te-policy-path)# preference 200
Router(config-sr-te-policy-path-pref)# constraints affinity exclude-any red
```

policy candidate-paths constraints disjoint-path

To configure the disjoint-path constraints, use the **on-demand dynamic disjoint-path** command in SR-TE sub-mode.

policy *policy* **candidate-paths preference** *preference* **constraints disjoint-path group-id** *id* **type** { **link** | **node** | **srlg** | **srlg-node** } [{ **sub-id** *sub_id* | **shortest-path** | **fallback disable** }]

Syntax Description

group-id *id*	Specify the group ID of the disjoint path. Valid values are from 1 to 65535.
type {**link** | **node** | **srlg** | **srlg-node** }	Specify the type of disjointness.
sub-id *id*	Specify the sub-group ID of the disjoint path. Valid values are from 1 to 65535.
shortest-path	Enable shortest path computation for the selected candidate path.
fallback disable	Disable all fallback behavior in case the requested disjointness cannot be achieved.

Command Default

None

Command Modes

SR-TE configuration

Command History

Release	Modification
Release 24.1.1	The **shortest-path** and **fallback disable** keywords were introduced.
Release 6.3.1	This command was introduced.

Usage Guidelines

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):

- 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.

If a pair of paths that meet the requested disjointness level cannot be found, then the paths will automatically fallback to a lower level:

- If the requested disjointness level is SRLG or node, then link-disjoint paths will be computed.
- 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.

Example

```
Router(config-sr-te)# policy FOO
Router(config-sr-te-policy)# candidate-paths preference 100
Router(config-sr-te-poliilojkl,.cy-path-pref)# constraints disjoint-path group-id 775 type
 link
```

The following example indicates how to configure the shortest path preference for a disjoint path:

```
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
```

The following example indicates how to configure strict disjointness for a SR-TE policy:

```
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
```

policy candidate-paths constraints resources

To exclude IP addresses from the path computation for SR-TE policies, use the **policy candidate-paths constraints resources** command in the SR-TE configuration mode.

policy *policy* **candidate-paths preference** *preference* **constraints resources** { **exclude resource-list** *name* | **exclude-group** *group_name* | **apply-group** *group_name* }

Syntax Description

resources {**exclude-group** | **exclude** | **apply-group**}	Specify the resource constraints for path computation: • exclude. Excludes resources from the path computation. • exclude-group. Excludes the apply-group configuration from the group. • apply-group. Applies configuration from a group.
resource-list *name*	Specify the name of the resource-list to exclude from the path computation.

Command Default

None

Command Modes

SR-TE configuration

Command History

Release	Modification
Release 24.1.1	This command was introduced.

Usage Guidelines

None.

Example

The following example shows how to exclude a list of IPv4 addresses 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
```

The following example shows how 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
```

policy candidate-paths dynamic

To configure the SR-TE head-end or SR-PCE to compute a path that is encoded using Anycast prefix SIDs of nodes along the path, use the **policy candidate-paths dynamic** command.

policy *policy* **candidate-paths preference** *preference* **dynamic** { **anycast-sid-inclusion** | **pcep** }

Syntax Description

anycast-sid-inclusion	Specifies a PCC-initiated path computation at the head-end router, encoded using Anycast prefix SIDs of nodes along the path.
pcep	Specifies that the path computation is at the SR-PCE.

Command Default

None

Command Modes

SR-TE

Command History

Release	Modification
Release 6.3.1	This command was introduced.

Usage Guidelines

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.

Example

The following example shows how to request a PCC-initiated Anycast SID-aware 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 1.1.1.10
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
```

policy candidate-paths dynamic metric

policy *policy* **candidate-paths preference** *preference* **dynamic metric** { **margin** { **absolute** | **relative** } *margin* | **sid-limit** *value* | **type** { **hopcount** | **igp** | **latency** | **te** } }

Syntax Description

metric {**absolute**	**relative** } *margin*	Specify the On-Demand dynamic path metric margin. The range for *margin* is from 0 to 2147483647.
sid-limit *value*	Specify the maximun SID depth (MSD).
type { **hopcount**	**igp**	**latency**	**te** }	Specify the metric type for use in path computation.

Command Default

None

Command Modes

SR-TE configuration

Command History

Release	Modification
Release 6.3.1	This command was introduced.

Usage Guidelines

If the configured MSD values are different, the per-LSP MSD takes precedence over the per-node MSD.

Example

```
Router(config-sr-te-policy-path-pref)# dynamic metric type te
Router(config-sr-te-policy-path-pref)# dynamic metric margin absolute 5
```

policy candidate-paths explicit

policy *policy* **candidate-paths preference** *preference* **explicit segment-list** *sid_list* [**weight** *weight*]

Syntax Description

segment-list	*sid_list*	Specify the explicit segment list.
weight	*weight*	Path option weight. Range is from 1 to 4294967295.

Command Default

None

Command Modes

ST-TE policy

Command History

Release	Modification
Release 6.3.1	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Example

```
Router(config-sr-te)# policy POLICY1
Router(config-sr-te-policy)# color 10 end-point ipv4 1.1.1.4
Router(config-sr-te-policy)# candidate-paths
Router(config-sr-te-policy-path)# preference 100
Router(config-sr-te-policy-path-pref)# explicit segment-list SIDLIST1
```

policy candidate-paths per-flow

To map a forward class to a per-flow policy, use the **policy candidate-paths per-flow** command.

policy *policy* **candidate-paths preference** *preference* **per-flow forward-class** { *value* **color** *color* | **default** *value* }

Syntax Description

forward-class *value*	Specify the forward class (FC). Values are from 0 to 7.
color *color*	Specify the color of the policy.
default *value*	Explicitly specify a default FC.

Command Default

When not explicitly configured, FC 0 is the default FC.

Command Modes

SR-TE policy

Command History

Release	Modification
Release 7.2.1	This command was introduced.

Usage Guidelines

When not explicitly configured, FC 0 is the default FC.

A Per-Flow Policy (PFP) defines an array of FC-to-PDP mappings. A PFP can then be used to steer traffic into a given PDP based on the FC assigned to a packet.

A Per-Flow Policy (PFP) is considered valid as long as its default FC has a valid Per-Destination Policy (PDP).

A color associated with a PFP SR policy cannot be used by a non-PFP SR policy. For example, if a per-flow ODN template for color 100 is configured, then the system will reject the configuration of any non-PFP SR policy using the same color. You must assign different color value ranges for PFP and non-PFP SR policies.

Example

```
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# policy FOO
Router(config-sr-te-policy)# candidate-paths
Router(config-sr-te-policy-path)# preference 100
Router(config-sr-te-policy-path-pref)# per-flow
Router(config-sr-te-pol-cp-pfp)# forward-class 0 color 10
Router(config-sr-te-pol-cp-pfp)# forward-class 1 color 20
```

policy candidate-paths preference lock duration

To enable a new lock duration for the Protect candidate path, use the **policy candidate-paths preference lock duration** command in the SR-TE configuration mode. To remove the lock function for a Protect path, use the **no** form of the command.

policy *name* [**candidate-paths** [**preference** *preference* [**lock** [**duration** *seconds*]]]]

Syntax Description

candidate-paths [**preference** *preference*]	(Optional) Configures the candidate path preference. The range is from 1 to 65535.
lock [**duration** *seconds*]	(Optional) Enables the specified lock duration for the Protect candidate path. The default lock duration is 300 seconds.

Command Default

The default Protect path lock duration is 300 seconds.

Command Modes

SR-TE configuration (config-sr-te)

Command History

Release	Modification
Release 7.4.2	This command was introduced.

Usage Guidelines

When the Working path is invalid, the Protect path becomes active. After the Working path has recovered, the Protect path remains active until the default lock duration (300 seconds) expires. You can configure a different lock duration using this command.

The duration range is 0 (disabled) to 3000 seconds. If the lock duration is 0 (disabled), then the Working path becomes active as soon as it recovers. If duration is not specified, the Protect path remains active.

Example

This example shows how to enable a new lock duration of 600 seconds for the Protect candidate path:

```
RP/0/RSP0/CPU0:ios# configure
RP/0/RSP0/CPU0:ios(config)# segment-routing traffic-eng
RP/0/RSP0/CPU0:ios(config-sr-te)# policy foo candidate-paths preference 50 lock duration
600
RP/0/RSP0/CPU0:ios(config-sr-te)# commit
```

policy color end-point

To configure the SR-TE color and end-point address, use the **policy color end-point** command.

policy *policy* **color** *color* **end-point** { **ipv4** | **ipv6** } *ip_addr*

Syntax Description

color *color*	Specify the color of the SR policy.
end-point {**ipv4**|**ipv6**} *ip_addr*	Specify the IPv4 or IPv6 address of the end-point.

Command Default

None

Command Modes

SR-TE policy

Command History

Release	Modification
Release 6.3.1	This command was introduced.

Usage Guidelines

An SR-TE policy is identified as an ordered list (head-end, color, end-point):

- Head-end – Where the SR-TE policy is instantiated
- Color – A numerical value that distinguishes between two or more policies to the same node pairs (Head-end – End point)
- End-point – The destination of the SR-TE policy

  Every SR-TE policy has a color value. Every policy between the same node pairs requires a unique color value.

Example

```
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# policy POLICY1
Router(config-sr-te-policy)# color 10 end-point ipv4 1.1.1.4
```

policy ipv6 disable

To disable IPv6 encapsulation (IPv6 caps) for a particular color and IPv4 NULL end-point, use the **ipv6 disable** command is SR-TE configuration mode.

policy ipv6 disable

Syntax Description

This command has no keywords or arguments.

Command Default

None

Command Modes

SR-TE configuration mode

Command History

Release	Modification
Release 6.5.1	This command was introduced.

Usage Guidelines

IPv6 caps for IPv4 NULL end-point is enabled automatically when the policy is created in Segment Routing Path Computation Element (SR-PCE). The binding SID (BSID) state notification for each policy contains an "ipv6_caps" flag that notifies SR-PCE clients (PCC) of the status of IPv6 caps (enabled or disabled).

An SR-TE policy with a given color and IPv4 NULL end-point could have more than one candidate path. If any of the candidate paths has IPv6 caps enabled, then all of the remaining candidate paths need IPv6 caps enabled. If IPv6 caps is not enabled on all candidate paths of same color and end-point, traffic drops can occur.

You can disable IPv6 caps for a particular color and IPv4 NULL end-point using the **ipv6 disable** command on the local policy. This command disables IPv6 caps on all candidate paths that share the same color and IPv4 NULL end-point.

Example

This example shows how to disable IPv6 caps for a particular color and IPv4 NULL end-point:

```
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# policy P1
Router(config-sr-te-policy)# color 1 end-point ipv4 0.0.0.0
Router(config-sr-te-policy)# ipv6 disable
```

policy path-protection

To enable path-protection for an SR-TE policy's candidate paths, use the **policy path-protection** command in the SR-TE configuration mode. To disable SR-TE policy path-protection, use the **no** form of the command.

policy *name* [**path-protection**]

Syntax Description

path-protection	(Optional) Specifies that path-protection should be enabled for the specified policy.

Command Default

Path-protection is not enabled for an SR-TE policy's candidate paths.

Command Modes

SR-TE configuration (config-sr-te)

Command History

Release	Modification
Release 7.4.2	This command was introduced.

Example

This example shows how to enable SR-TE policy path-protection for the policy **foo**:

```
RP/0/RSP0/CPU0:ios# configure
RP/0/RSP0/CPU0:ios(config)# segment-routing traffic-eng
RP/0/RSP0/CPU0:ios(config-sr-te)# policy foo path-protection
RP/0/RSP0/CPU0:ios(config-sr-te-path-pref-protection)#commit
```

policy performance-measurement

To apply a performance measurement profile to an SR policy, use the **performance-measurement** command in SR-TE configuration mode.

{ **policy performance-measurement** [**delay-measurement delay-profile name** *name* [**logging delay-exceeded**]] | [**liveness-detection liveness-profile name** *name* [**invalidation-action** { **down** | **none** }] | **logging session-state-change**] | [**reverse-path label** *label*] }

Syntax Description

policy *policy*	Specifies the SR policy name.
liveness-detection	Enables end-to-end SR Policy Liveness Detection
invalidation-action {**none** | **down**}	Specifies the action to take when the PM liveness session goes down: • **down (default)**: The candidate path is immediately operationally brought down. • **none**: No action is taken. If logging is enabled, the failure is logged but the SR Policy operational state is not modified.
logging session-state-change	Enables Syslog messages when the session state changes.
logging delay-exceeded	Enables Syslog messages when the delay exceeds the threshold.
delay-profile name *profile*	Specifies the SR Policy delay profile name.
reverse-path label {*BSID-value* | *NODE-SID-value*}	Specifies the MPLS label to be used for the reverse path for the reply. If you configured liveness detection with ECMP hashing, you must specify the reverse path. The default reverse path uses IP Reply. • *BSID-value*: The Binding SID (BSID) label for the reverse SR Policy. (This is practical for manual SR policies with a manual BSID.) • *NODE-SID-value*: The absolute SID label of the (local) Sender Node to be used for the reverse path for the reply.

Command Default

None

Command Modes

SR-TE configuration

Command History

Release	Modification
Release 6.5.2	This command was introduced.
Release 7.3.1	The **liveness-detection** options were introduced.

Example

```
Router(config)# segment-routing traffic-eng
Router(config-sr-te)# policy TEST
Router(config-sr-te-policy)# color 4 end-point ipv4 10.10.10.10
Router(config-sr-te-policy)# performance-measurement
Router(config-sr-te-policy-perf-meas)# delay-measurement delay-profile name profile2
```

policy shutdown

To shutdown an SR policy, use the **policy** *name* **shutdown** command in SR-TE configuration mode.

policy *name* **shutdown**

Syntax Description

policy*name*	Specifies the SR policy name.

Command Default

None

Command Modes

SR-TE configuration mode

Command History

Release	Modification
Release 6.3.1	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Example

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

resource-list

To configure a list of IPv4 addresses that you want to exclude from the network resource list for a candidate path, use the **resource-list** command in SR-TE configuration mode.

resource-list *name* **index** *"1-65535"* **ipv4** *ipv4-addr*

Syntax Description

resource-list *name*	Specify the resource-list name to exclude from the path computation.
index *1-65535*	Specify the index entry. Ranges from 1–65535.
ipv4 *ipv4-addr*	Specify the IPv4 address that you want to exclude from the network resource list.

Command Default

None

Command Modes

SR-TE configuration mode

Command History

Release	Modification
Release 24.1.1	This command was introduced.

Usage Guidelines

None.

Example

The following example shows how 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
```

segment-list

To create a segment list for explicit policy path, use the **segment-list** command in SR-TE configuration mode.

segment-list [**name**] *name* **index** *index* **mpls** { **label** *label* | **adjacency** { *ipv4-addr* *ipv6-addr* } }

Syntax Description

index *index*	Specifies the index entry.
mpls	Enters MPLS configure mode.
label *label*	Specify the MPLS label value.
adjacency {*ipv4-addr* | *ipv6-addr*}	Specify the IP address.

Command Default

None

Command Modes

SR-TE configuration mode

Command History

Release	Modification
Release 6.3.1	This command was introduced.

Usage Guidelines

A segment list can use IPv4/IPv6 addresses (adjacency) or MPLS labels, or a combination of both.

- The IP address can be link or a Loopback address.
- Once you enter an MPLS label, you cannot enter an IP address.

Example

The following example shows how to create a segment list with IP addresses:

```
Router(config-sr-te)# segment-list name SIDLIST1
Router(config-sr-te-sl)# index 10 mpls adjacency 1.1.1.2
Router(config-sr-te-sl)# index 20 mpls adjacency ipv4 1.1.1.3
Router(config-sr-te-sl)# index 30 mpls adjacency ipv4 1.1.1.4
```

The following example shows how to create a segment list with MPLS labels:

```
Router(config-sr-te)# segment-list name SIDLIST2
Router(config-sr-te-sl)# index 10 mpls label 16002
Router(config-sr-te-sl)# index 20 mpls label 16003
Router(config-sr-te-sl)# index 30 mpls label 16004
```

The following example shows how to create a segment list with IP addresses and MPLS labels:

```
Router(config-sr-te)# segment-list name SIDLIST3
Router(config-sr-te-sl)# index 10 mpls adjacency ipv4 1.1.1.2
Router(config-sr-te-sl)# index 20 mpls label 16003
```

```
Router(config-sr-te-sl)# index 30 mpls label 16004
```

te-latency

To enable ECMP-aware path computation for TE metric, use the **te-latency** command in SR-TE configuration mode.

te-latency

Syntax Description This command has no keywords or arguments.

Command Default None

Command Modes SR-TE configuration mode

Command History

Release	Modification
Release 6.3.1	This command was introduced.

Usage Guidelines ECMP-aware path computation is enabled by default for IGP and LATENCY metrics

Example

This example shows how to enable ECMP-aware path computation for TE metric:

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

timers

To configure SR-TE reoptimization timers, use the **timers** command in SR-TE configuration mode.

timers { **candidate-path cleanup-delay** *seconds* | **cleanup-delay** *seconds* | **init-verify-restart** *seconds* | **init-verify-switchover** *seconds* | **init-verify-startup** *seconds* | **periodic-reoptimization** *seconds* | **install-delay** *seconds* }

Syntax Description

candidate-path cleanup-delay *seconds*	Specifies the delay before cleaning up candidate paths. Range of *seconds* is from 0 (immediate cleanup) to 86400.
cleanup-delay *seconds*	Specifies the delay before cleaning up previous path. Range of *seconds* is from 0 (immediate cleanup) to 300.
init-verify-restart *seconds*	Specifies the delay before topology convergence after topology starts populating for restart case. Range of *seconds* is from 10 to 10000.
init-verify-switchover *seconds*	Specifies the delay before topology convergence after topology starts populating for switchover case. Range of *seconds* is from 10 to 10000.
init-verify-startup *seconds*	Specifies the delay before topology convergence after topology starts populating for startup case. Range of *seconds* is from 10 to 10000.
install-delay *seconds*	Specifies the delay before switching to a reoptimized path. Range of *seconds* is from 0 (immediate cleanup) to 300.
periodic-reoptimization *seconds*	Specifies how often to perform periodic reoptimization of policies. Range of *seconds* is from 0 (disables reoptimization) to 86400.

Command Default

- **candidate-path cleanup-delay**: 120 seconds
- **cleanup-delay**: 10 seconds
- **init-verify-restart**: 40 seconds
- **init-verify-switchover**: 60 seconds
- **init-verify-startup**: 120 seconds
- **install-delay**: 10 seconds
- **periodic-reoptimization**: 600 seconds

Command Modes

SR-TE configuration mode

Command History

Release	Modification
Release 6.3.1	This command was introduced.

Usage Guidelines

No specific guidelines impact the use of this command.

Example

```
Router(config)# segment-routing traffic-eng
Router(config-sr-te)# timers
Router(config-sr-te-timers)# candidate-path cleanup-delay 600
Router(config-sr-te-timers)# cleanup-delay 60
Router(config-sr-te-timers)# init-verify-restart 120
Router(config-sr-te-timers)# init-verify-startup 600
Router(config-sr-te-timers)# init-verify-switchover 30
Router(config-sr-te-timers)# install-delay 60
Router(config-sr-te-timers)# periodic-reoptimization 3000
```