



Nokia Service Router Linux

Release 24.3

Segment Routing Guide

3HE 20244 AAAA TQZZA
Edition: 01
March 2024

Nokia is committed to diversity and inclusion. We are continuously reviewing our customer documentation and consulting with standards bodies to ensure that terminology is inclusive and aligned with the industry. Our future customer documentation will be updated accordingly.

This document includes Nokia proprietary and confidential information, which may not be distributed or disclosed to any third parties without the prior written consent of Nokia.

This document is intended for use by Nokia's customers ("You"/"Your") in connection with a product purchased or licensed from any company within Nokia Group of Companies. Use this document as agreed. You agree to notify Nokia of any errors you may find in this document; however, should you elect to use this document for any purpose(s) for which it is not intended, You understand and warrant that any determinations You may make or actions You may take will be based upon Your independent judgment and analysis of the content of this document.

Nokia reserves the right to make changes to this document without notice. At all times, the controlling version is the one available on Nokia's site.

No part of this document may be modified.

NO WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY OF AVAILABILITY, ACCURACY, RELIABILITY, TITLE, NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, IS MADE IN RELATION TO THE CONTENT OF THIS DOCUMENT. IN NO EVENT WILL NOKIA BE LIABLE FOR ANY DAMAGES, INCLUDING BUT NOT LIMITED TO SPECIAL, DIRECT, INDIRECT, INCIDENTAL OR CONSEQUENTIAL OR ANY LOSSES, SUCH AS BUT NOT LIMITED TO LOSS OF PROFIT, REVENUE, BUSINESS INTERRUPTION, BUSINESS OPPORTUNITY OR DATA THAT MAY ARISE FROM THE USE OF THIS DOCUMENT OR THE INFORMATION IN IT, EVEN IN THE CASE OF ERRORS IN OR OMISSIONS FROM THIS DOCUMENT OR ITS CONTENT.

Copyright and trademark: Nokia is a registered trademark of Nokia Corporation. Other product names mentioned in this document may be trademarks of their respective owners.

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

© 2024 Nokia.

Table of contents

1	About this guide.....	5
1.1	Precautionary and information messages.....	5
1.2	Conventions.....	5
2	What's new.....	7
3	About segment routing.....	8
3.1	IS-IS extensions for segment routing.....	9
3.1.1	Router Capability TLV advertisement.....	10
3.1.2	Prefix SID processing.....	10
3.1.3	Adjacency SID processing.....	13
3.1.4	Datapath programming by SID type.....	14
3.2	Supported functionality.....	15
4	SR-MPLS configuration on the default network instance.....	17
4.1	Defining the SRGB and enabling SR-MPLS.....	17
4.2	Defining protocol-independent prefix SIDs.....	18
5	SR-MPLS configuration on the IS-IS instance.....	20
5.1	Configuring an IS-IS node SID.....	20
5.2	Dynamic adjacency SIDs configuration.....	21
5.2.1	Defining the dynamic SRLB.....	21
5.2.2	Enabling dynamic adjacency SID assignment for the IS-IS instance.....	22
5.3	Static adjacency SIDs configuration.....	23
5.3.1	Defining the static SRLB.....	23
5.3.2	Configuring static adjacency SID assignment for the IS-IS instance.....	24
5.3.3	Configuring static adjacency SIDs for an interface.....	25
5.4	Overriding adjacency SID assignment mode on an interface.....	25
5.5	Enabling SR-MPLS on the IS-IS instance.....	27
6	BGP shortcuts configuration over segment routing tunnels.....	28
6.1	Configuring BGP shortcuts over segment routing.....	28
7	Segment routing display commands.....	30

7.1	Displaying the SID database.....	30
7.2	Displaying label block information.....	32
7.3	Displaying tunnel table entries.....	32
8	LSP ping and trace for segment routing tunnels.....	36
8.1	Performing an LSP ping to a segment routing prefix.....	36
8.2	Performing an LSP trace to a segment routing prefix.....	37

1 About this guide

This document describes configuration details for the Segment Routing feature set used with the Nokia Service Router Linux (SR Linux).

This document is intended for network technicians, administrators, operators, service providers, and others who need to understand how the router is configured.

**Note:**

This manual covers the current release and may also contain some content that will be released in later maintenance loads. See the *SR Linux Release Notes* for information on features supported in each load.

Configuration and command outputs shown in this guide are examples only; actual displays may differ depending on supported functionality and user configuration.

1.1 Precautionary and information messages

The following are information symbols used in the documentation.



DANGER: Danger warns that the described activity or situation may result in serious personal injury or death. An electric shock hazard could exist. Before you begin work on this equipment, be aware of hazards involving electrical circuitry, be familiar with networking environments, and implement accident prevention procedures.



WARNING: Warning indicates that the described activity or situation may, or will, cause equipment damage, serious performance problems, or loss of data.



Caution: Caution indicates that the described activity or situation may reduce your component or system performance.



Note: Note provides additional operational information.



Tip: Tip provides suggestions for use or best practices.

1.2 Conventions

Nokia SR Linux documentation uses the following command conventions.

- **Bold** type indicates a command that the user must enter.
- Input and output examples are displayed in Courier text.
- An open right-angle bracket indicates a progression of menu choices or simple command sequence (often selected from a user interface). Example: **start** > **connect to**.
- A vertical bar (|) indicates a mutually exclusive argument.

- Square brackets ([]) indicate optional elements.
- Braces ({ }) indicate a required choice. When braces are contained within square brackets, they indicate a required choice within an optional element.
- *Italic* type indicates a variable.

Generic IP addresses are used in examples. Replace these with the appropriate IP addresses used in the system.

2 What's new

Topic	Location
LSP ping and trace for segment routing tunnels	LSP ping and trace for segment routing tunnels

3 About segment routing

Segment routing performs shortest path routing or source routing using the concept of abstract segment. A segment can represent the local prefix of a node, a specific adjacency of the node (interface or next hop), a service context, or a specific explicit path over the network. Each segment is identified by a segment ID (SID).

With segment routing, the source router can define the end-to-end path for packets to reach a destination using an ordered list of segments represented by their SIDs. Each SID represents actions that subsequent nodes in the network execute, such as forwarding a packet to a specific destination or interface.

Unlike a typical Multiprotocol Label Switching (MPLS) architecture, a segment routing network does not require LDP or RSVP-TE to set up tunnels. Instead, segment routing extensions to the IGP (such as IS-IS) provide support for allocation and signaling of SID information.

SR-MPLS

When segment routing is instantiated over the MPLS data plane (referred to as SR-MPLS), a SID is represented by a standard MPLS label or an index that maps to an MPLS label. To route a packet to a destination, the source router pushes one or more MPLS labels, which represent the required SIDs onto the packet. Each subsequent node forwards the packet based on the outer label, removes the outer label, and forwards the packet to the next segment on the path.

The MPLS data plane requires no modifications to support SR-MPLS.

SR Linux supports SR-MPLS over both IPv4 (SR-MPLS IPv4) and IPv6 (SR-MPLS over IPv6).

Node SID

A prefix SID is a segment type that represents the ECMP-aware shortest path to reach a particular IP prefix from any IGP topology location. A node SID is a type of prefix SID that identifies a specific node in the IGP topology using a loopback address as the prefix. When a node SID is included in an MPLS label stack, it instructs the receiving node to forward the packet to the destination along the ECMP-aware shortest path determined by the IGP.

You allocate the node SIDs to use in the network in a similar fashion as the loopback IP addresses. Like all prefix SIDs, node SIDs must be unique within the domain, and are allocated in the form of an MPLS label (or index). The range of labels available for prefix and node SIDs is defined in the SRGB.

Protocol-independent prefix SIDs

You can set a prefix SID at the network instance level that is shared by multiple IGPs (currently only IS-IS).

SR Linux supports up to four protocol-independent prefix SIDs to be associated with the default network instance. By default, all prefix SIDs are set as node SIDs. However, you can disable the node SID flag as required.



Note: As an alternative, you can also define a node SID under the IS-IS configuration context. In that case, the node SID is not protocol-independent but is specific to that IGP. However, in that case, the node SID flag cannot be disabled.

You can configure an IS-IS node SID and a protocol-independent prefix SID on a single interface. In this case, the IS-IS IGP overrides the protocol-independent prefix SID configuration, and only the IGP node SID is advertised.

A protocol-independent prefix SID provides more flexibility over an IS-IS node SID. If required, you can override the protocol-independent prefix SID with an IGP node SID.

SRGB

The Segment Routing Global Block (SRGB) defines the range of MPLS labels available for global segments, such as prefix and node SIDs. The SRGB is defined as one contiguous block of MPLS labels. Nokia strongly recommends using identical SRGBs on all nodes within the SR domain. The use of identical SRGBs simplifies troubleshooting because the same MPLS label represents the same prefix or node SID on each node.

The SRGB configuration is supported on the default network instance only.

SRLB

While the SRGB defines a range of MPLS labels for global segments, the Segment Routing Label Block (SRLB) defines a range of MPLS labels for local segments, such as adjacency SIDs. An adjacency SID is a segment type that represents an instruction to forward packets over a specific (one-hop) adjacency between two nodes in the IGP.

The SRLB is configurable per IGP protocol instance (on the default network instance only). It is defined as one contiguous block of MPLS labels per IGP instance.

SR Linux can support a dynamic SRLB, a static SRLB, or a combination of both.

With a dynamic SRLB, SR Linux dynamically assigns adjacency SIDs as required. With a static SRLB, you must define the static adjacency SID values manually.



Note: Static adjacency SID configuration is not supported for broadcast interfaces.

See the *SR Linux MPLS Guide* for more information about MPLS and MPLS labels.

3.1 IS-IS extensions for segment routing

Segment routing with IS-IS (also known as SR-ISIS) refers to the segment routing extensions of the IS-IS IGP protocol and the forwarding entries created by those extensions. The SR-ISIS extensions advertise TLVs specific to segment routing, which propagate SIDs across the domain.

With segment routing enabled on the IS-IS instance, IS-IS extensions can support TLVs to advertise IPv4/IPv6 prefix SIDs. The following sub-TLVs are defined in RFC 8667 and are supported in the implementation of SR-ISIS:

- Prefix SID sub-TLV
- Adjacency SID sub-TLV
- SR-Capabilities sub-TLV
- SR-Algorithm sub-TLV

By default, SR Linux advertises each configured prefix SID as a node SID. You can disable the node SID flag for protocol-independent prefix SIDs but not for IS-IS-specific node SIDs.

The following sections describes the behaviors and limitations of the SR-ISIS TLV and sub-TLVs.

3.1.1 Router Capability TLV advertisement

In SR Linux, advertisement of the Router Capability TLV (TLV 242) with the sub-TLVs related to segment routing is automatically enabled when segment routing is configured on the IS-IS instance. The originated TLV 242 encodes the router ID and always indicates domain-wide scope. There is no explicit configuration to enable or disable router capability advertisement.

The following table describes the TLV 242 sub-TLVs that SR Linux supports.

Table 1: TLV 242 sub-TLVs

Sub-TLVs	Description
SR Capabilities sub-TLV	Advertises one SRGB block and data plane capability: <ul style="list-style-type: none"> If the SRGB block is not configured, the entire SR Capabilities sub-TLV is suppressed SR Linux does not support multiple, discontinuous SRGB blocks I = 1 indicates support for IPv4 V = 1 indicates support for IPv6
SR Algorithm sub-TLV	The algorithm field is set to 0, indicating shortest path first (SPF) algorithm based on link metric. (The value is not checked on receive.)
SR Local Block sub-TLV	(Optional) Advertises the SRLB. SR Linux does not support multiple, discontinuous SRLB blocks.
SRMS Preference sub-TLV	Not supported.

Receiving Router Capability TLVs

SR Linux decodes the received Router Capability TLVs and maintains the details in the LSDB YANG state model.

3.1.2 Prefix SID processing

SR Linux uses Prefix SID sub-TLVs to advertise IPv4/IPv6 prefix SIDs.

Origination of IPv4 and IPv6 Prefix SID sub-TLVs

SR Linux originates Prefix SID sub-TLVs with the following processing rules and flag encoding:

- originates a single Prefix SID sub-TLV per IS-IS IP Reachability TLV
- encodes the 32-bit index in the Prefix SID sub-TLV while the 24-bit label is not supported
- both IPv4 and IPv6 Prefix SID sub-TLVs originate within MT = 0

The following table describes the default flag encoding for Prefix SID sub-TLVs.

Table 2: Default flag encoding for Prefix SID sub-TLV

Flag	Default encoding	Description
R-flag	R = 1 (for node SID)	Re-advertisement flag When R = 1, the Prefix SID sub-TLV and its corresponding IP reachability TLV are propagated between levels.
	R = 0 (for prefix SID)	Re-advertisement flag R = 0 initially, but this setting can change to R = 1 during propagation by other routers.
N-flag	N = 1	Node SID flag Set by default, meaning that SR Linux advertises each configured prefix SID as a node SID. You can disable the node SID flag for protocol-independent prefix SIDs but not for IGP-specific prefix SIDs. If the referenced interface is <code>system0.0</code> , the node SID flag must be set.
P-flag	P = 1	No-PHP Flag Always set, meaning the label for the prefix SID is pushed by the PHP router when forwarding to this router. When the SR Linux PHP router (with P-flag set to 1) processes a received prefix SID with the P-flag set to zero, it uses implicit-null for the outgoing label toward the router that advertised it.
E-flag	E = 0	Explicit null flag Always set to 0, meaning that the router is requesting no explicit null termination. However, the SR Linux PHP router processes a received prefix SID with the E-flag set to 1 as long as the P-flag is also set to 1. In this case, the PHP router pushes explicit-null for the outgoing label toward the router that advertised it. The system ignores the value of the E-flag if the P-flag is not set.
V-flag	V = 0	Value flag Always set to 0 to indicate an index for the SID, instead of a specific value.
L-flag	L = 0	Local flag Always set to 0 to indicate that the SID index value is not locally significant.
Algorithm field	algorithm 0	Always set to 0 to indicate shortest path first (SPF) algorithm based on link metric. This field is not checked on a received Prefix SID sub-TLV.

Receiving IPv4 and IPv6 Prefix SID sub-TLVs

SR Linux processes a Prefix SID sub-TLV using the following rules:

- decodes Prefix SID sub-TLVs that are received in TLVs 135, 235, 236, and 237 for representation in the YANG state model of the LSDB
- ignores Prefix SID sub-TLVs received in TLVs 235 and 237 for further processing; these sub-TLVs do not appear in the state representation of the local SID database (only in the LSDB)
- if the local flag is set or the algorithm is not 0, deems a Prefix SID sub-TLV in a received TLV 135 or TLV-236 invalid and makes it inactive
 - If the MPLS label value implied by the received Prefix SID sub-TLV falls outside the range of the SRGB block, deems the node SID invalid and makes it inactive (reason = sid-index-out-of-range). No forwarding state is created for invalid entries.
- does not resolve a Prefix SID sub-TLV received with the N-flag set and a prefix length different than 32 (for IPv4) or 128 (for IPv6)
- does not resolve a Prefix SID sub-TLV received in TLV 135 or TLV 236 if the L-flag is set
- resolves a Prefix SID received within an IP reachability TLV based on the following route preference:
 1. SID received via L1 in a Prefix SID sub-TLV part of the IP reachability TLV
 2. SID received via L2 in a Prefix SID sub-TLV part of the IP reachability TLV
- resolves a Prefix SID sub-TLV received without the N-flag set as long as the prefix length is 32 (for IPv4) or 128 (for IPv6)
- processes only the first Prefix SID sub-TLV if multiple are received within the same IS-IS IP reachability TLV

Inter-level propagation of prefix SIDs

SR Linux performs inter-level propagation of prefix SIDs as follows:

- An L1L2 router propagates a prefix (and Prefix SID sub-TLV) received in an IP reachability TLV from L1 to L2. A router in L2 sets up an SR-ISIS tunnel to the L1 router via the L1L2 router, which acts as an LSR.
- If an ABR summarizes a prefix, the Prefix SID sub-TLV is not propagated with the summarized route between levels. To propagate the node SID for a /32 prefix, you must disable route summarization.
- By default, an L1L2 router does not propagate a prefix (or Prefix SID sub-TLV) received in an IP reachability TLV from L2 to L1.
- Using an export policy, an L1L2 router propagates a prefix (and Prefix SID sub-TLV) received in an IP reachability TLV from L2 to L1. A router in L1 sets up an SR-ISIS tunnel to the L2 router via the L1L2 router, which acts as an LSR.
 - L2 to L1 route leaking occurs when a level 2 IS-IS route is matched by the IS-IS export policy.

IS-IS prefix SID database

In the YANG state model, the IS-IS prefix SID database captures information about:

- advertised IS-IS node SIDs
- remotely originated node and prefix SIDs learned and marked as active by IS-IS

Global SID Database

In the YANG state model, the global SID database captures information about:

- configured protocol-independent prefix SIDs

- advertised IS-IS node SIDs
- remotely originated node and prefix SIDs learned and marked as active by IS-IS

Prefix SID conflicts

Remotely signaled prefix SID entries can cause conflicts with local prefix SIDs in the following cases:

- **prefix conflict**

If you configure a SID index value for a prefix/node SID that creates an overlap with an existing ILM entry in the SRGB block, a prefix conflict occurs.

In SR Linux, all local prefix/node SID entries have a lower numerical preference than remote prefix SID entries learned via IS-IS. Therefore the local prefix/node SID remains active and advertised, but in the SID database the entry appears with a status of: `prefix-conflict = true`. All other conflicting entries show as inactive in the IS-IS prefix SID database and may no longer be visible in the global SID database.

- **SID conflict**

After SR Linux removes the inactive prefix SID entries, a SID conflict can still occur if the same SID is assigned to multiple IS-IS prefixes. In this case, the prefix SID with the lowest SID index remains active, but in the SID database the entry appears with a status of: `sid-conflict = true`. All other conflicting entries show as inactive in the IS-IS prefix SID database and may no longer be visible in the global SID database.

- **SID out of range**

When a prefix SID advertised from another router has a SID index or label value that is not within the locally defined SRGB range of the network instance, SR Linux determines that it is out of range. All other conflicting entries show as inactive in the IS-IS prefix SID database and may no longer be visible in the global SID database.

3.1.3 Adjacency SID processing

SR Linux uses Adjacency SID sub-TLVs to advertise IPv4/IPv6 adjacency SIDs.

Origination of IPv4 and IPv6 Adjacency SID sub-TLVs

By default, SR Linux does not automatically assign adjacency SIDs, but you can enable dynamic adjacency SID assignment on the IS-IS instance. SR Linux supports origination of IPv4 and IPv6 Adjacency SID sub-TLVs in TLVs 22 and 222. Dynamic adjacency SIDs are allocated as follows:

- SR-ISIS assigns dynamic adjacency SIDs on all P2P and LAN IS-IS interfaces in all levels (except for interfaces individually configured with an adjacency SID assignment of **none** or **static**).
- If an interface is enabled for IPv4, SR-ISIS allocates IPv4 adjacency SIDs.
- If the interface is enabled for IPv6, SR-ISIS allocates IPv6 adjacency SIDs.
- For each IS-IS interface, the YANG state indicates all IPv4 and IPv6 adjacency SIDs that are currently active and programmed in the ILM table.

The following table describes the flag encoding for the Adjacency SID sub-TLVs.

Table 3: Flag encoding for Adjacency SID sub-TLV

Flag	Encoding	Description
F-flag	F = 0 (for IPv4)	Address-family flag for IPv4 adjacency encapsulation.
	F = 1 (for IPv6)	Address-family flag for IPv6 adjacency encapsulation.
B-flag	B = 0	Backup flag. Set to zero and is not processed on receipt.
V-flag	V = 1	Value flag. Always set to 1, indicating that the adjacency SID carries a value.
L-flag	L = 1	Local flag. Always set to 1, indicating that the adjacency SID has local significance.
S-flag	S = 0	Set flag. Always set to zero because assigning adjacency SIDs to parallel links between neighbors is not supported. A received adjacency SID with the S-flag set is not processed.
P-Flag	P = 1 (static)	Set to 1 when static adjacency-sid is configured on the interface, indicating that the adjacency SID allocation is persistent.
	P = 0 (dynamic)	Set to 0 when dynamic adjacency-sid is configured on the interface.
weight octet	N/A	Not supported and is set to all zeros.

Receiving IPv4 and IPv6 Adjacency SID sub-TLVs

- LAN Adjacency SID sub-TLVs received in TLVs 22 and 222 are decoded for representation in the YANG state model of the LSDB. If there are multiple Adjacency SID sub-TLVs in a particular TLV, only the first Adjacency SID sub-TLV is processed.
- LAN Adjacency SID sub-TLVs received in TLVs 23 and 223 are ignored for further processing.
- A LAN Adjacency SID sub-TLV in a received TLV 22/222 is invalid and not used if either of the following are true:
 - The V-flag is not set.
 - The L-flag is not set.
- A received adjacency SID with the S-flag set is not processed.

3.1.4 Datapath programming by SID type

The following table describes the datapath programming by SID type.

Table 4: Datapath programming by SID type

SID type	Datapath programming
IS-IS node SID	<p>When advertised, the IS-IS node SID creates an ILM entry that matches the associated MPLS label value (SRGB start-label + index) and performs a POP operation.</p> <p>The POP operation indicates that the default network instance is the context for the IP header lookup that follows.</p>
Protocol-independent prefix SID	<p>When configured, the protocol-independent prefix SID (whether advertised or not), creates an ILM entry that matches the associated MPLS label value (SRGB start-label + index) and performs a POP operation.</p> <p>The POP operation indicates that the default network instance is the context for the IP header lookup that follows.</p>
IS-IS adjacency SID (P2P and LAN)	<p>When advertised, the adjacency SID:</p> <ul style="list-style-type: none"> • Creates an ILM entry that matches the associated MPLS label value and forwards matching packets to the adjacent neighbor using a SWAP to implicit null. <ul style="list-style-type: none"> – If the adjacency SID is IPv4, forwarding is via an IPv4 ARP entry. – If the adjacency SID is IPv6, forwarding is via an IPv6 ND entry. • Creates a TTM entry with SR-ISIS as the tunnel type. The destination of the tunnel is the IPv4/IPv6 interface address of the neighbor.
Remote prefix SID	<p>When received and valid:</p> <ul style="list-style-type: none"> • The remote prefix SID: <ul style="list-style-type: none"> – Creates an ILM entry that matches on the localized MPLS label value (SRGB start-label + index) and performs a SWAP operation. – Creates an IPv4/IPv6 tunnel entry in TTM with SR-ISIS as the tunnel type. • If IS-IS has a non-ECMP route toward the IPv4/IPv6 prefix, then the FEC linked to the ILM (or representing the tunnel) has one next-hop label forwarding entry (NHLFE). The NHLFE pushes label N, where $N = (\text{SRGB start label of the next-hop IS-IS neighbor}) + \text{SID index}$. • If IS-IS has an ECMP route toward the IPv4/IPv6 prefix, then the FEC linked to the ILM (or representing the tunnel) has multiple ECMP members, each corresponding to one NHLFE. Each NHLFE pushes label N_i, where $N_i = (\text{SRGB start label of next-hop IS-IS neighbor}) + \text{SID index}$.

3.2 Supported functionality

SR-MPLS over IPv4 and SR-MPLS over the IPv6 dataplane on 7250 IXR-6/10

- maximum of one label pushed per packet (corresponding to a remote node SID)
- maximum of one label popped per packet (corresponding to a local node SID)

- transit and egress LSR ECMP hashing
- ingress LSR ECMP hashing
- prefix/node SID (IPv4/IPv6)
- adjacency SID (IPv4/IPv6)

MPLS label management

- Configurable SRGB for the default network instance:
 - SRGB is a reference to a static shared label range.
 - SRGB must be one contiguous block.
- Configurable SRLB for the IS-IS instance of the default network instance:
 - SR Linux supports a static SRLB (dedicated or shared), a dynamic SRLB (dedicated only), or a combination of both.
 - Each SRLB must be one contiguous block.

IS-IS segment routing extensions

- IPv4/IPv6 prefix/node SID
- SID collision handling (no event generated when collision detected)

BGP shortcuts over segment routing tunnels

- SR-ISIS-IPv4 tunnels in TTM can resolve BGP IPv4 routes.
- SR-ISIS-IPv6 tunnels in TTM can resolve:
 - BGP IPv4 routes
 - BGP IPv6 routes
- SR-ISIS tunnels are programmed into TTM with a non-configurable preference value of 11.
- SR-ISIS tunnels that correspond to an adjacency-SID have a metric of 0.
- SR-ISIS tunnels that correspond to a remote prefix-SID have a metric that reflects the IGP cost to reach the advertising router.

SR-MPLS OAM

- ICMP tunneling
- RFC 4950 extensions

ICMP tunneling

SR Linux supports ICMP extensions for MPLS to support debugging and tracing in MPLS and SR-MPLS networks. With ICMP tunneling enabled, ICMP messages can be tunneled to the endpoint of the tunnel and then returned through IP routing. For more information, see the *SR Linux MPLS Guide*.

4 SR-MPLS configuration on the default network instance

Segment routing on the MPLS data plane is supported on the default network instance only. To configure available SR-MPLS options for the default network instance, perform the following tasks:

1. [Defining the SRGB and enabling SR-MPLS](#)
2. [Defining protocol-independent prefix SIDs](#) (optional)

4.1 Defining the SRGB and enabling SR-MPLS

About this task

The SRGB references a static shared MPLS label range that must be one contiguous block of labels. To configure the SRGB, you must enter the **segment-routing mpls** context on the default network instance, which enables SR-MPLS.



Note: Nokia strongly recommends to using identical SRGBs on all nodes within the SR domain.

Procedure

Step 1. Define a static MPLS label range in the system context.

Step 2. Assign that MPLS label range to be available for use by the SRGB.

Example: Define the MPLS label range for the SRGB

The following example defines static shared MPLS label range `srgb-range-1` for the SRGB.

```
--{ * candidate shared default }--[ ]--
# info system mpls label-ranges static srgb-range-1
system {
  mpls {
    label-ranges {
      static srgb-range-1 {
        shared true
        start-label 16001
        end-label 16999
      }
    }
  }
}
```

Example: Assign the MPLS label range to the SRGB

The following example assigns static label range `srgb-range-1` to the SRGB.

```
--{ * candidate shared default }--[ ]--
# info network-instance default segment-routing mpls global-block
network-instance default {
  segment-routing {
```

```

mpls {
  global-block {
    label-range srgb-range-1
  }
}

```

4.2 Defining protocol-independent prefix SIDs

About this task

You can configure a prefix SID that is shared by multiple IGPs (currently only IS-IS) in a network instance. SR Linux supports up to four protocol-independent prefix SIDs to be associated with the default network instance. By default, all prefix SIDs are set as node SIDs. However, you can disable the node SID flag as required.



Note: As an alternative, you can define a node SID under the IS-IS configuration context. In that case, the node SID is not protocol-independent but is specific to that IGP, and the node SID flag cannot be disabled.

A single interface can be configured with both an IS-IS node SID and a protocol-independent prefix SID. In this case, the IS-IS IGP overrides the protocol independent prefix SID configuration, and only the IGP node SID is advertised.

A protocol-independent prefix SID is typically preferred over an IS-IS node SID because the protocol-independent prefix SID provides more flexibility. If required, you can override the protocol-independent prefix SID with an IGP node SID.

Prerequisites

- Configure the primary address of a loopback (loN.n) or system0.0 subinterface with the required node SID prefix.

Procedure

Step 1. To configure the local prefix SID, you must specify the following:

- local-prefix SID index (1 to 4)
- loopback subinterface that owns the advertised prefixes
- IPv4 or IPv6 (or both) label index for the SID, referencing the SRGB base, using one of the following options:
ipv4-label-index | ipv6-label-index

Step 2. Optionally, you can disable the node SID flag on the defined prefix SID, using the following option:

node-sid {true | false}

If the referenced interface is system0.0, the node SID flag cannot be set to false.

Example: Define local prefix SID

The following example defines local prefix SID 2 on loopback interface lo0.2. Prefix SID 2 is associated with IPv4 label index 102 and IPv6 label index 202, and the node SID flag is set to false.

```
--{ * candidate shared default }--[ ]--  
# info network-instance default segment-routing mpls local-prefix-sid 2  
  network-instance default {  
    segment-routing {  
      mpls {  
        local-prefix-sid 2 {  
          interface lo0.2  
          ipv4-label-index 102  
          ipv6-label-index 202  
          node-sid false  
        }  
      }  
    }  
  }  
}
```

5 SR-MPLS configuration on the IS-IS instance

SR-MPLS is configurable on one IS-IS instance only (on the default network instance). Before you can enable SR-MPLS you must configure IS-IS as the IGP on the segment routing nodes in your network. See the *SR Linux Routing Protocols Guide*.

SR Linux supports a static or dynamic SRLB, or a combination of both. With dynamic SRLB, SR Linux automatically assigns the dynamic adjacency SIDs for the associated interfaces. With static SRLB, you must manually define static adjacency SIDs for the associated interfaces.

To configure SR-MPLS on an IS-IS instance:

1. Specify the node SID (a protocol-independent prefix SID or IS-IS node SID). See:
 - [Defining protocol-independent prefix SIDs](#)
 - [Configuring an IS-IS node SID](#)
2. Define the Segment Routing Label Block (SRLB) and configure the adjacency SID mode using one of the following methods:
 - [Dynamic adjacency SIDs configuration](#)
 - [Static adjacency SIDs configuration](#)
3. (Optional) Set interface-specific adjacency SID modes. See:
 - [Overriding adjacency SID assignment mode on an interface](#)
4. Enable SR-MPLS. See:
 - [Enabling SR-MPLS on the IS-IS instance](#)

5.1 Configuring an IS-IS node SID

Prerequisites

- Configure the primary address of a loopback (loN.n) or system0.0 subinterface with the prefix to associate with the node SID, and set the subinterface as an IS-IS interface.

About this task

Unlike protocol-independent prefix SIDs, IS-IS node SIDs are specific to the IS-IS IGP. Also, you cannot disable the node SID flag for IS-IS node SIDs.



Note: The IS-IS node SID configuration is optional if you have already defined a protocol-independent prefix SID. If both are applied to the same interface, the IS-IS IGP node SID configuration overrides the protocol-independent prefix SID configuration.

To define the node SID in the IS-IS interface configuration, you must specify an index value that is relative to the SRGB. In this case, the MPLS label value is calculated as follows:

Local Label (Prefix/Node SID) = SRGB start-label + {SID index}

For example, if the SRGB range starts at 16000, and the index value is 1, the resulting MPLS label value for the SID is 16001.

Procedure

To configure the node SID corresponding to an IPv4 or IPv6 prefix, assign an index value to the loopback (or system) subinterface in the IS-IS instance.

Example: Configure node SID

The following example adds IPv4 node SID index 1 to loopback subinterface lo0.1.

```
--{ * candidate shared default }--[ ]--
# info network-instance default protocols isis instance sr-isis-1 interface lo0.1
network-instance default {
  protocols {
    isis {
      instance sr-isis-1 {
        interface lo0.1 {
          segment-routing {
            mpls {
              ipv4-node-sid {
                index 1
              }
            }
          }
        }
      }
    }
  }
}
```

5.2 Dynamic adjacency SIDs configuration

To enable allocation of dynamic adjacency SIDs for all interfaces in the IS-IS instance, perform the following tasks:

1. [Defining the dynamic SRLB](#)
2. [Enabling dynamic adjacency SID assignment for the IS-IS instance](#)



Note: You can optionally override the dynamic adjacency SID setting for the IS-IS instance by configuring individual interfaces to use static adjacency SIDs or to use no adjacency SID assignment at all. See section [Overriding adjacency SID assignment mode on an interface](#) for more information.

5.2.1 Defining the dynamic SRLB

About this task

To assign dynamic adjacency SID labels, you must configure a dynamic SRLB that references a dynamic range of MPLS labels. The dynamic SRLB must reference a dedicated (non-shared) label range consisting of one contiguous block of labels.

Procedure

Step 1. Define a dynamic MPLS label range.

Step 2. Assign that MPLS label range to the SRLB.

Example: Define the dynamic MPLS label range for the SRLB

The following example defines a dynamic MPLS label range (`srlb-dynamic-1`) for use by the SRLB.

```
--{ * candidate shared default }--[ ]--
# info system mpls label-ranges dynamic srlb-dynamic-1
system {
  mpls {
    label-ranges {
      dynamic srlb-dynamic-1 {
        start-label 15001
        end-label 15999
      }
    }
  }
}
```

Example: Assign the dynamic MPLS label range to the SRLB

The following example assigns the defined dynamic MPLS label range (`srlb-dynamic-1`) to the SRLB.

```
--{ * candidate shared default }--[ ]--
# info network-instance default protocols isis dynamic-label-block
network-instance default {
  protocols {
    isis {
      dynamic-label-block srlb-dynamic-1
    }
  }
}
```

5.2.2 Enabling dynamic adjacency SID assignment for the IS-IS instance

About this task

You can enable dynamic adjacency SID allocation for the network instance. The default setting is disabled (**false**). When enabled, IS-IS assigns a dynamic adjacency SID to all IS-IS interfaces in all levels, except for interfaces configured at the interface level with an adjacency SID assignment of **none** or **static**. You can also set the hold time that is applied to the dynamically allocated adjacency SIDs.

Procedure

To enable dynamic adjacency SID assignment, in the IS-IS instance configuration, set the segment routing **dynamic-adjacency-sids** option to **all-interfaces true**.

Example: Enable dynamic adjacency SIDs for all interfaces

The following example enables dynamic adjacency SIDs for IS-IS instance `sr-isis-1` and sets the hold-time to 20 seconds.

```
--{ * candidate shared default }--[ ]--
# info network-instance default protocols isis instance sr-isis-1 segment-routing mpls
dynamic-adjacency-sids
network-instance default {
```

```

protocols {
  isis {
    instance sr-isis-1 {
      segment-routing {
        mpls {
          dynamic-adjacency-sids {
            all-interfaces true
            hold-time 20
          }
        }
      }
    }
  }
}

```

5.3 Static adjacency SIDs configuration

To enable assignment of static adjacency SIDs for all interfaces in the IS-IS instance, perform the following tasks:

1. [Defining the static SRLB](#)
2. [Configuring static adjacency SID assignment for the IS-IS instance](#)
3. [Configuring static adjacency SIDs for an interface](#)



Note:

- Static adjacency SID configuration is not supported for broadcast interfaces.
- You can optionally override the static adjacency SID setting by configuring individual interfaces to use dynamic adjacency SIDs. See [Overriding adjacency SID assignment mode on an interface](#).

5.3.1 Defining the static SRLB

About this task

To assign static adjacency SID labels, you must configure a static SRLB that references a static range of MPLS labels. The static SRLB must reference a dedicated or shared label range consisting of one contiguous block of labels.

Procedure

- Step 1.** Define a static MPLS label range.
- Step 2.** Assign that MPLS label range to the static SRLB.

Example: Define the static MPLS label range for the SRLB

The following example defines a static MPLS label range (`srlb-static-10`) for use by the SRLB.

```

--{ * candidate shared default }--[ ]--
# info system mpls label-ranges static srlb-static-10
system {
  mpls {
    label-ranges {

```

```

        static srlb-static-10 {
            shared true
            start-label 14001
            end-label 14999
        }
    }
}

```

Example: Assign the static MPLS label range to the SRLB

The following example assigns the defined static MPLS label range (`srlb-static-10`) to the static SRLB label block.

```

--{ * candidate shared default }--[ ]--
# info network-instance default protocols isis instance sr-isis-1 segment-routing mpls
static-label-block
network-instance default {
    protocols {
        isis {
            instance sr-isis-1 {
                segment-routing {
                    mpls {
                        static-label-block srlb-static-10
                    }
                }
            }
        }
    }
}

```

5.3.2 Configuring static adjacency SID assignment for the IS-IS instance

About this task

You can configure the IS-IS instance to use static adjacency SID assignment. In this case, IS-IS does not assign any dynamic adjacency SIDs to any interfaces.

Procedure

To specify static adjacency SID assignment for all interfaces in the IS-IS instance, set the segment routing `dynamic-adjacency-sids` option to **all-interfaces false**.

Example: Enable static adjacency SIDs for all interfaces

```

--{ * candidate shared default }--[ ]--
# info network-instance default protocols isis instance sr-isis-1 segment-routing mpls
dynamic-adjacency-sids
network-instance default {
    protocols {
        isis {
            instance sr-isis-1 {
                segment-routing {
                    mpls {
                        dynamic-adjacency-sids {
                            all-interfaces false
                        }
                    }
                }
            }
        }
    }
}

```



```

    }
  }
}

```

5.3.3 Configuring static adjacency SIDs for an interface

About this task

If you set the adjacency SID mode to static for the IS-IS instance or for specific interfaces, you must manually set the value for the adjacency SIDs for the applicable interfaces.

Procedure

To define the static adjacency SID value for an interface, use the following options:

- `ipv4-adjacency-sid static <adjacency-SID>`
- `ipv6-adjacency-sid static <adjacency-SID>`

Example: Set static IPv4 and IPv6 adjacency SID values

```

--{ * candidate shared default }--[ ]--
# info network-instance default protocols isis instance sr-isis-1 interface ethernet-2/1.1
segment-routing mpls
  network-instance default {
    protocols {
      isis {
        instance sr-isis-1 {
          interface ethernet-2/1.1 {
            segment-routing {
              mpls {
                ipv4-adjacency-sid {
                  static 13001
                }
                ipv6-adjacency-sid {
                  static 14001
                }
              }
            }
          }
        }
      }
    }
  }
}

```

5.4 Overriding adjacency SID assignment mode on an interface

About this task

You can override the adjacency SID configuration set on the IS-IS instance by configuring individual interfaces with a different adjacency SID assignment mode. The SID adjacency options available for an interface are:

- **dynamic**
IS-IS dynamically allocates one or more dynamic adjacency SIDs for this interface for each enabled address family. On a broadcast interface, the TLV 22/222 corresponding to the adjacency with the

Designated IS (DIS) includes a LAN adjacency SID sub-TLV that reports all the adjacent systems on the LAN. In this case, you must also configure a dynamic SRLB to automatically assign the adjacency SIDs to the dynamic interfaces.

- **static**
IS-IS does not assign dynamic adjacency SIDs. Instead, you must statically configure an adjacency SID for the interface. In this case, you must also configure a static SRLB. This option is not available if the interface type is broadcast (LAN).
- **none**
No adjacency SIDs are allocated. If no SR-MPLS traffic is flowing on a particular interface, set the adjacency SID assignment to **none** to save resources that the dynamic adjacency SIDs would otherwise consume.

Procedure

Under the IS-IS interface, set the following segment routing options:

- **ipv4-adjacency-sid assignment [static | none | dynamic]**
- **ipv6-adjacency-sid assignment [static | none | dynamic]**

Example: Set IPv4 and IPv6 adjacency SID assignments to static

```
--{ * candidate shared default }--[ ]--
# info network-instance default protocols isis instance sr-isis-1 interface ethernet-1/2.1
network-instance default {
  protocols {
    isis {
      instance sr-isis-1 {
        interface ethernet-1/2.1 {
          segment-routing {
            mpls {
              ipv4-adjacency-sid {
                assignment static
              }
              ipv6-adjacency-sid {
                assignment static
              }
            }
          }
        }
      }
    }
  }
}
```

Example: Set IPv4 and IPv6 adjacency SID assignment to dynamic

```
--{ * candidate shared default }--[ ]--
# info network-instance default protocols isis instance sr-isis-1 interface ethernet-1/2.2
network-instance default {
  protocols {
    isis {
      instance sr-isis-1 {
        interface ethernet-1/2.2 { {
          segment-routing {
            mpls {
              ipv4-adjacency-sid {
                assignment dynamic
              }
              ipv6-adjacency-sid {
```


6 BGP shortcuts configuration over segment routing tunnels

When you have segment routing enabled in your network, you can specify a segment routing tunnel as the next hop for a BGP route. This capability is referred to as BGP shortcuts over segment routing.

SR-ISIS-IPv4 tunnels in TTM can resolve BGP IPv4 routes.

SR-ISIS-IPv6 tunnels in TTM can resolve BGP IPv4 and BGP IPv6 routes.

To configure BGP shortcuts, you must configure the BGP protocol with the allowed tunnel types and the required tunnel resolution mode.

6.1 Configuring BGP shortcuts over segment routing

Procedure

Step 1. In the default network instance, define the tunnel-resolution mode for the BGP protocol. This setting determines the order of preference and the fallback when using tunnels in the tunnel table instead of routes in the FIB. Available options are as follows:

- **require**
requires tunnel table lookup only
- **prefer**
prefers tunnel table lookup over FIB lookup
- **disabled (default)**
performs FIB lookup only

Step 2. Set the allowed tunnel types for next-hop resolution.

Example: Configure IPv4 BGP shortcuts

The following example shows the BGP next-hop resolution configuration to allow IPv4 SR-ISIS tunnels, with the tunnel mode set to prefer.

```
--{ * candidate shared default }--[ ]--
# info network-instance default protocols bgp afi-safi ipv4-unicast ipv4-unicast next-hop-
resolution ipv4-next-hops tunnel-resolution
network-instance default {
  protocols {
    bgp {
      afi-safi ipv4-unicast {
        ipv4-unicast {
          next-hop-resolution {
            ipv4-next-hops {
              tunnel-resolution {
                mode prefer
                allowed-tunnel-types [
                  sr-isis
                ]
              }
            }
          }
        }
      }
    }
  }
}
```


7 Segment routing display commands

SR Linux supports display commands to provide operational information about segment routing, including the following:

- SID database
- label block information
- tunnel-table entries

7.1 Displaying the SID database

Procedure

Use the **info from state** command to display information about the SID database.

Example: Display the global SID database

```
--{ candidate shared default }--[ ]--
# info from state network-instance default segment-routing mpls sid-database
network-instance default {
  segment-routing {
    mpls {
      sid-database {
        prefix-sid 10.20.1.1/32 sid-label-value 3010 protocol isis protocol-instance 0
protocol-multi-topology 0 algorithm 0 {
  active true
      }
        prefix-sid 10.20.1.2/32 sid-label-value 3020 protocol isis protocol-instance 0
protocol-multi-topology 0 algorithm 0 {
  active true
      }
        prefix-sid 10.20.1.3/32 sid-label-value 3030 protocol isis protocol-instance 0
protocol-multi-topology 0 algorithm 0 {
  active true
      }
        prefix-sid 10.20.1.4/32 sid-label-value 3040 protocol isis protocol-instance 0
protocol-multi-topology 0 algorithm 0 {
  active true
      }
        prefix-sid 10.20.1.5/32 sid-label-value 3050 protocol isis protocol-instance 0
protocol-multi-topology 0 algorithm 0 {
  active true
      }
        prefix-sid 10.20.1.6/32 sid-label-value 3060 protocol isis protocol-instance 0
protocol-multi-topology 0 algorithm 0 {
  active true
      }
        prefix-sid 2001:db8:20:1::1/128 sid-label-value 3610 protocol isis protocol-
instance 0 protocol-multi-topology 0 algorithm 0 {
  active true
      }
        prefix-sid 2001:db8:20:1::2/128 sid-label-value 3620 protocol isis protocol-
instance 0 protocol-multi-topology 0 algorithm 0 {
  active true
      }
      }
    }
  }
}
```



```
IPv4 tunnel table of network-instance "default"
```

IPv4 Prefix	Encaps Type	Tunnel Type	Tunnel ID	FIB	Metric	Preference	Last Update	Next-hop (Type)	Next-hop
10.20.1.1/32	mpls	sr-isis	13010	Y	10	11	2021-11-10T 14:17:32.202Z	10.10.1.1 (mpls)	ethernet-1/33.1
10.20.1.3/32	mpls	sr-isis	13030	Y	10	11	2021-11-10T 14:18:02.299Z	10.10.3.3 (mpls)	ethernet-1/34.1
10.20.1.4/32	mpls	sr-isis	13040	Y	10	11	2021-11-10T 14:18:26.729Z	10.10.4.4 (mpls)	ethernet-1/8.1
10.20.1.5/32	mpls	sr-isis	13050	Y	10	11	2021-11-10T 14:19:02.343Z	10.10.23.5 (mpls)	ethernet-1/9.1
10.20.1.6/32	mpls	sr-isis	13060	Y	10	11	2021-11-10T 14:19:38.140Z	10.10.14.6 (mpls)	ethernet-1/11.1

```
5 SR-ISIS tunnels, 5 active, 0 inactive
```

```
IPv6 tunnel table of network-instance "default"
```

IPv6 Prefix	Encaps Type	Tunnel Type	Tunnel ID	FIB	Metric	Preference	Last Update	Next-hop (Type)	Next-hop
2001:db8:20:1::1/128	mpls	sr-isis	13610	Y	10	11	2021-11-10T 14:17:32.204Z	fe80::201:1ff:feff:20(mpls)	ethernet-1/33.1
2001:db8:20:1::3/128	mpls	sr-isis	13630	Y	10	11	2021-11-10T 14:18:03.316Z	fe80::201:3ff:feff:21(mpls)	ethernet-1/34.1
2001:db8:20:1::4/128	mpls	sr-isis	13640	Y	10	11	2021-11-10T 14:18:31.638Z	fe80::201:4ff:feff:20(mpls)	ethernet-1/8.1
2001:db8:20:1::5/128	mpls	sr-isis	13650	Y	10	11	2021-11-10T 14:19:04.882Z	fe80::201:5ff:feff:20(mpls)	ethernet-1/9.1
2001:db8:20:1::6/128	mpls	sr-isis	13660	Y	10	11	2021-11-10T 14:19:43.545Z	fe80::201:6ff:feff:4(mpls)	ethernet-1/11.1

```
5 SR-ISIS tunnels, 5 active, 0 inactive
```

```
--{ running }--[ ]--
```

Example: Show IPv4 tunnel-table entries

```
--{ running }--[ ]--
```

```
# show network-instance default tunnel-table ipv4
```

```
IPv4 tunnel table of network-instance "default"
```

IPv4 Prefix	Encaps Type	Tunnel Type	Tunnel ID	FIB	Metric	Preference	Last Update	Next-hop (Type)	Next-hop
10.20.1.1/32	mpls	sr-isis	13010	Y	10	11	2021-11-10T 14:17:32.202Z	10.10.1.1 (mpls)	ethernet-1/33.1
10.20.1.3/32	mpls	sr-isis	13030	Y	10	11	2021-11-10T 14:18:02.299Z	10.10.3.3 (mpls)	ethernet-1/34.1
10.20.1.4/32	mpls	sr-isis	13040	Y	10	11	2021-11-10T 14:18:26.729Z	10.10.4.4 (mpls)	ethernet-1/8.1
10.20.1.5/32	mpls	sr-isis	13050	Y	10	11	2021-11-10T 14:19:02.343Z	10.10.23.5 (mpls)	ethernet-1/9.1
10.20.1.6/32	mpls	sr-isis	13060	Y	10	11	2021-11-10T 14:19:38.140Z	10.10.14.6 (mpls)	ethernet-1/11.1

```

|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
|         |         |         |         |         |         |         |         |         |         |
|-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
5 SR-ISIS tunnels, 5 active, 0 inactive
|-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

Example: Show IPv6 tunnel-table entries

```

--{ running }--[ ]--
# show network-instance default tunnel-table ipv6
-----
IPv6 tunnel table of network-instance "default"
-----
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|IPv6 Prefix|Encaps|Tunnel|Tunnel|FIB|Metric|Prefe|Last Update|Next-hop|Next-hop|
|            |Type  |Type  |ID     |   |      |rence|            |(Type)  |          |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|2001:db8:20:|mpls  |sr-isis|13610  |Y  |10   |11   |2021-11-10T|fe80::201:|ethernet-1/33.1|
|1::1/128     |      |      |       |   |     |     |14:17:32.202Z|1ff:feff:20(mpls)|
|2001:db8:20:|mpls  |sr-isis|13630  |Y  |10   |11   |2021-11-10T|fe80::201:|ethernet-1/34.1|
|1::3/128     |      |      |       |   |     |     |14:18:02.299Z|3ff:feff:21(mpls)|
|2001:db8:20:|mpls  |sr-isis|13640  |Y  |10   |11   |2021-11-10T|fe80::201:|ethernet-1/8.1  |
|1::4/128     |      |      |       |   |     |     |              |4ff:feff:20(mpls)|
|2001:db8:20:|mpls  |sr-isis|13650  |Y  |10   |11   |2021-11-10T|fe80::201:|ethernet-1/9.1  |
|1::5/128     |      |      |       |   |     |     |              |5ff:feff:20(mpls)|
|2001:db8:20:|mpls  |sr-isis|13660  |Y  |10   |11   |2021-11-10T|fe80::201:|ethernet-1/11.1|
|1::6/128     |      |      |       |   |     |     |              |6ff:feff:4(mpls)|
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
5 SR-ISIS tunnels, 5 active, 0 inactive
-----

```

Example: Show IPv4 tunnel-table entries by destination prefix

```

--{ running }--[ ]--
# show network-instance default tunnel-table ipv4 10.20.1.6/32 type sr-isis detail
-----
Show report for network instance "default" tunnel table
-----
=====
Destination      : 10.20.1.6/32
Encapsulation    : mpls
Tunnel Type      : sr-isis
Metric           : 10
Preference       : 11
Last Update      : 2021-11-10T14:19:38.140Z
FIB Status       : active
Next-hops
  10.10.14.6 (mpls) via [ethernet-1/11.1]
  pushed MPLS labels : [53060]
=====

```

Example: Show IPv6 tunnel-table entries by destination prefix

```

# show network-instance default tunnel-table ipv6 2001:db8:10:20:1::6/128 detail
-----
Show report for network instance "default" tunnel table
-----
=====
Destination      : 2001:db8:10:20:1::6/128
=====

```

```
Encapsulation      : mpls
Tunnel Type       : sr-isis
Metric            : 10
Preference        : 11
Last Update       : 2021-11-10T14:19:43.545Z
FIB Status        : active
Next-hops
  2001:db8:201:6ff:feff:4 (mpls) via [ethernet-1/11.1]
    pushed MPLS labels : [53660]
=====
```

8 LSP ping and trace for segment routing tunnels

To check connectivity and trace the path to any midpoint or endpoint of an SR-ISIS shortest path tunnel, SR Linux supports the following OAM commands:

- **tools oam lsp-ping sr-isis prefix-sid <prefix>**
- **tools oam lsp-trace sr-isis prefix-sid <prefix>**

Supported parameters include **destination-ip**, **source-ip**, **timeout**, and **ecmp-next-hop-select**. However, the only mandatory parameter is the **prefix-sid**.

Results from the **lsp-ping** and **lsp-trace** operations are displayed using **info from state** commands.

8.1 Performing an LSP ping to a segment routing prefix

Procedure

To check connectivity to a segment routing prefix, use the **tools oam lsp-ping sr-isis** command. To display the results, use the **info from state oam lsp-ping sr-isis** command, specifying the session ID output from the **lsp-ping**.

Example: Perform an LSP ping to a destination segment routing prefix

```
# tools oam lsp-ping sr-isis prefix-sid 10.20.1.6/32
/oam/lsp-ping/sr-isis/prefix-sid[prefix=10.20.1.6/32]:
  Initiated LSP Ping to prefix 10.20.1.6/32 with session id 49152
```

Example: Display results of the LSP ping

```
--{ + running }--[ ]--
# info from state oam lsp-ping sr-isis prefix-sid 10.20.1.6/32 session-id 49152
  oam {
    lsp-ping {
      sr-isis {
        prefix-sid 10.20.1.6/32 {
          session-id 49152 {
            test-active false
            statistics {
              round-trip-time {
                minimum 4292
                maximum 4292
                average 4292
                standard-deviation 0
              }
            }
            path-destination {
              ip-address 127.0.0.1
            }
            sequence 1 {
              probe-size 48
              request-sent true
              out-interface ethernet-1/33.1
              reply {
```

```
received true
reply-sender 10.20.1.6
udp-data-length 40
mpls-ttl 255
round-trip-time 4292
return-code replying-router-is-egress-for-fec-at-stack-depth-n
return-subcode 1
}
}
}
}
}
}
}
}
```

8.2 Performing an LSP trace to a segment routing prefix

Procedure

To trace the path to any midpoint or endpoint of a segment routing tunnel, use the **tools oam lsp-trace** command. To display the results, use the **info from state oam lsp-trace sr-isis** command, specifying the session ID output from the **lsp-trace**.

Example: Perform an LSP trace to a destination segment routing prefix

```
# tools oam lsp-trace sr-isis prefix-sid 10.20.1.6/32
/oam/lsp-trace/sr-isis/prefix-sid[prefix=10.20.1.6/32]:
  Initiated LSP Trace to prefix 10.20.1.6/32 with session id 49153
```

Example: Display results of the LSP trace

```
--{ + running }--[ ]--
# info from state oam lsp-trace sr-isis prefix-sid 10.20.1.6/32 session-id 49153
oam {
  lsp-trace {
    sr-isis {
      prefix-sid 10.20.1.6/32 {
        session-id 49153 {
          test-active false
          path-destination {
            ip-address 127.0.0.1
          }
          hop 1 {
            probe 1 {
              probe-size 76
              probes-sent 1
              reply {
                received true
                reply-sender 10.20.1.2
                udp-data-length 60
                mpls-ttl 1
                round-trip-time 2768
                return-code label-switched-at-stack-depth-n
                return-subcode 1
              }
            }
          }
          downstream-detailed-mapping 1 {
            mtu 1500
            address-type ipv4-numbered
            downstream-router-address 10.10.4.4
          }
        }
      }
    }
  }
}
```

```
        downstream-interface-address 10.10.4.4
        mpls-label 1 {
            label 27000
            protocol isis
        }
    }
}
hop 2 {
    probe 1 {
        probe-size 76
        probes-sent 1
        reply {
            received true
            reply-sender 10.20.1.4
            udp-data-length 60
            mpls-ttl 2
            round-trip-time 3414
            return-code label-switched-at-stack-depth-n
            return-subcode 1
        }
        downstream-detailed-mapping 1 {
            mtu 1500
            address-type ipv4-numbered
            downstream-router-address 10.10.9.6
            downstream-interface-address 10.10.9.6
            mpls-label 1 {
                label 27000
                protocol isis
            }
        }
    }
}
hop 3 {
    probe 1 {
        probe-size 76
        probes-sent 1
        reply {
            received true
            reply-sender 10.20.1.6
            udp-data-length 32
            mpls-ttl 3
            round-trip-time 4429
            return-code replying-router-is-egress-for-fec-at-stack-depth-n
            return-subcode 1
        }
    }
}
}
}
}
}
}
```


Customer document and product support



Customer documentation

[Customer documentation welcome page](#)



Technical support

[Product support portal](#)



Documentation feedback

[Customer documentation feedback](#)