



NSP

Network Services Platform

Release 24.4

Path Control and Simulation Guide

3HE-20024-AAAA-TQZZA
Issue 1
April 2024

Legal notice

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.

© 2024 Nokia.

Contents

| | |
|---|-----------|
| About this document | 5 |
| 1 Path Control overview | 7 |
| 1.1 How does NSP enable path control? | 7 |
| 1.2 How do I navigate the Network Info and Statistics view?..... | 9 |
| 1.3 How can I view system activity logs? | 11 |
| 1.4 How do I modify the path control network map? | 12 |
| 1.5 How does NSP function as a PCE? | 14 |
| 1.6 What path control algorithms does NSP support? | 15 |
| 1.7 What are flex algo definitions? | 17 |
| 1.8 How is inventory synchronized?..... | 19 |
| 1.9 How are IRO objects used for path calculation? | 22 |
| 1.10 What SROS path computation methods does NSP support? | 23 |
| 1.11 How do I access path control APIs?..... | 24 |
| 1.12 How does path control implement user access control? | 25 |
| 1.13 How do I associate a workflow with a node, link, or LSP? | 27 |
| 1.14 How do I enable automatic VSR-NRC site switchover?..... | 29 |
| 1.15 How do I visualize detailed metrics? | 31 |
| 2 Path simulation overview | 33 |
| 2.1 How does NSP enable path simulation? | 33 |
| 2.2 How do I modify the path simulation network map?..... | 34 |
| 2.3 How does path simulation implement user access control? | 36 |
| 2.4 How do I access path simulation APIs? | 37 |
| 2.5 How do I import a network? | 38 |
| 2.6 How do I add an admin domain? | 39 |
| 2.7 How do I run a worst case failure scenario? | 40 |
| 3 LSPs | 41 |
| 3.1 What types of LSPs does NSP support? | 41 |
| 3.2 How does latency-based LSP rerouting work? | 48 |
| 3.3 How do I view LSP path history?..... | 52 |
| 3.4 What are association groups?..... | 53 |
| 3.5 How are BSIDs used to reduce MSD?..... | 54 |
| 3.6 How do I create PCE-initiated LSPs? | 55 |
| 3.7 How do I create an LSP? | 58 |

| | | |
|----------|--|-----------|
| 3.8 | How do I collect statistics? | 61 |
| 3.9 | How do I view an LSP's computed path?..... | 63 |
| 3.10 | How do I view the LSPs associated with a path profile policy?..... | 64 |
| 3.11 | How do I modify the demand bandwidth of an LSP? | 65 |
| 3.12 | How do I highlight a potential path on the network map? | 67 |
| 3.13 | How do I collect bit rate statistics using 7250 IXR? | 70 |
| 4 | NEs and links | 73 |
| 4.1 | What are BGP EPE links?..... | 73 |
| 4.2 | How do I place a link set into maintenance mode? | 75 |
| 4.3 | How do I turn down a link set? | 77 |
| 4.4 | How do I create an IGP link? | 79 |
| 4.5 | How do I create a BGP link? | 81 |
| 4.6 | How do I create a broadcast link? | 82 |
| 4.7 | How do I add a node? | 84 |
| 4.8 | How do I group NEs by region? | 86 |
| 5 | Policies | 87 |
| 5.1 | What is a path profile policy? | 87 |
| 5.2 | How do I create a path profile policy? | 92 |
| 5.3 | How do I apply a path profile override?..... | 95 |
| 5.4 | How do I create a router ID mapping policy? | 96 |
| 5.5 | How do I modify the system IP MPLS configuration policy? | 98 |
| 5.6 | What are SR policies?..... | 99 |
| 5.7 | How do I create an SR policy? | 100 |

About this document

Purpose

This document provides important contextual information and procedures that will enable readers to use NSP's path control and simulation functions.

Document support

Customer documentation and product support URLs:

- [Documentation Center](#)
- [Technical support](#)

How to comment

Please send your feedback to documentation.feedback@nokia.com.

1 Path Control overview

1.1 How does NSP enable path control?

1.1.1 Path control for NSP

NSP provides a view of the IGP topology and PCE LSPs. It also displays the status of the IGP network and provides functionality to optimize the network resources.

NSP leverages centralized, intelligent network control capabilities so that operators can rapidly adapt to changing demand and traffic patterns and run their networks more efficiently. It accepts path connection requests from service management operations, from OSS and orchestration systems, and from physical/virtual network elements, then calculates optimal paths through the network for a given set of business and technical constraints by leveraging centralized views of all available assets/topologies and their current state.

NSP's path control functions are based on a Path Computation Element (PCE) architecture that integrates standard protocols such as PCEP to open up path computation to external control. This allows PCCs to be enhanced with various path optimization algorithms that ensure optimal path placement across the network. NSP's path control implementation is stateful in nature and will maintain an up-to-date Traffic Engineering Database (TED), as well as the current RSVP-based label switched paths (LSP) and the segment routing path (SRP) state. It tracks RSVP BW and manages BW for the Segment-Routed TE paths as a unified state.

1.1.2 Multi-domain path computation

NSP supports path computation across multiple IGP instances. These instances are discovered as admin domains with stitching points on the common ASBR routers. The path traversal algorithm uses a flat graph and computes the shortest path based on the required metric. NSP PCE precomputes inter-domain path reachability, pruning domains in order to avoid visiting domains which may form loops or which don't reach the destination. All existing constraints apply, such as diversity and Max SID Depth. Both Segment-Routed TE paths and RSVP TE paths are supported and deployed.

1.1.3 Northbound interface for topology retrieval

NSP supports both an IETF-based NBI and a proprietary NBI model with extended attributes for topology retrieval. This is in addition to the existing IETF-based NBI. Using this NBI, northbound applications can obtain the IP and TE topology from NSP, including additional information (such as area number/level instance and node/link/prefix SIDs). Northbound applications and controllers, such as NSP's own IP/Optical coordination, can also modify the following TE attributes: SRLG, TE metric, IGP metric, Latency, and admin group.

1.1.4 Unnumbered interface support

NSP allows an LSP to be computed over a set of links where the interfaces are unnumbered. After configuring unnumbered interfaces and OSPF/ISIS link interfaces with TE enabled, the link appears

in the path control network map. Only TE-enabled OSPF/ISIS unnumbered interfaces are supported. If TE is later disabled, the links become operationally down.

If the BGP-LS message contains "local if index" and "remote if index", paring between incoming and outgoing links will be allowed. As a result, the two links will be shown as one on the network map. If the BGP-LS message only contains "local if index", paring is not possible. The LSP pathfinder algorithm is able to find and use unnumbered interfaces in the same manner as regular interfaces.

1.2 How do I navigate the Network Info and Statistics view?


1.2.1 The Network Info and Statistics view

The Network Info and Statistics view provides an at-a-glance summary of the state of the network and provides links to points of interest.

1.2.2 Network Summary

The Network Summary section of the Network Info and Statistics view is comprised of the following subsections:

What is the status of your network?

This subsection displays the overall status of the network, based on the health of the VSR-NRCs and PCEP sessions, which can be viewed individually using the provided **buttons** .

What is in your network?

This subsection displays the number of routers, IGP/BGP prefixes, and domains with areas that are in the network. Clicking on **Go to router list** will take you to the Routers view.

LSPs count

This subsection displays the number PCC-initiated LSPs and PCE-initiated LSPs that are in the network. Clicking on **Go to Path List** will take you to the LSPs view.

1.2.3 System status

The System status section of the Network Info and Statistics view is comprised of the following subsections:

Link status

This subsection displays the number of links in the network that are operationally Up, Down, or in Maintenance. Clicking on **Go to link list** will take you to the Links view.

LSP status

This subsection displays the number of LSPs in the network that are operationally Up or Down. Clicking on **Go to Path List** will take you to the LSPs view.

Link utilization distribution

This subsection displays the percentage of links in the network that are in Critical, High, Target, or Low utilization states. Clicking on **Go to link list** will take you to the Links view.

Path profiles override

This subsection displays the number of path profile overrides in the network that are Active or Failed. Clicking on **View in Path List** will take you to the LSPs view, where only the LSPs with path profile overrides configured are displayed.

TE updates

This subsection displays the number of TE updates that have occurred in the network within the selected time frame. Clicking on the drop-down menu will allow you to change the selected time frame.

1.2.4 System automation activities

The System automation activities section of the Network Info and Statistics view is comprised of the following subsections:



Note: The subsections will only display information based on the selected time frame: Last 5 minutes, Last 10 minutes, Last hour, or Total number.

Clicking on **system activity logs** will take you to that section of the dashboard.

PCEP activities

This subsection displays the number of PCEP requests, replies, reports, updates, and initiations that have occurred in the network within the selected time frame.

Path computation

This subsection displays the number of path computations, re-signal requests, and computation failures that have occurred in the network within the selected time frame.


Top 5 control plane summary

This subsection displays the 5 PCCs that have experienced the most PCEP requests, replies, reports, updates, and initiations within the selected time frame. Clicking on **View full list** will display a list of all PCCs and the amount of activity that has occurred on each.

1.3 How can I view system activity logs?

1.3.1 System activity logging

The System Activity Logging view is visible when **Path Control, System Activity Logging** is selected from the NSP menu.


This view displays a list of system activity logs, with the option to **Refresh**  the list, or click on any log to populate the Log details panel.

1.4 How do I modify the path control network map?


1.4.1 Available options

The path control network map is a persisted and stateful presentation of the IGP topology layer which is associated to managed data layers for location positioning. The following options can be used to configure the path control network map:

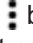
Refresh

The **Refresh**  button is located on the path control network map view. When clicked, the latest version of the map is fetched.


Search In Map

The **Search In Map**  button is located on the path control network map view. When clicked, a window opens, in which you can provide device-specific information in order to locate a Link, Router, Prefix, or Region (when nodes are clustered by region) on the network map.


Export Network

Export Network can be selected from the contextual menu when the **More**  button is clicked on the path control network map view. This option is used to export the current network as a .zip file so that it can be imported into the Simulation tool for testing purposes.


Save Map Changes

Save Map changes can be selected from the contextual menu when the **More**  button is clicked on the path control network map view. When the NSP's Map Layouts and Groups views are part of the deployment, The path control function fetches and utilizes the node positions from that data, but users can reposition nodes on the network map as required. When Save Map changes is selected, the new node positions are persisted only in the path control network map.

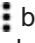
Restore Common Map Layout

Restore Common Map Layout can be selected from the contextual menu when the **More**  button is clicked on the path control network map view. When the NSP's Map Layouts and Groups views are part of the deployment, the path control function fetches and utilizes the node positions from that data, but users can reposition nodes on the network map as required. Users can then revert to the map configuration reflected in the NSP's Map Layouts and Groups views by selecting Restore Common Map Layout.


Rebuild

Rebuild can be selected from the contextual menu when the **More**  button is clicked on the path control network map view and Map Actions is chosen. This options deletes, then rebuilds the network map to ensure synchronization with the IGP topology layer.


Auto Layout

Auto Layout can be selected from the contextual menu when the **More**  button is clicked on the path control network map view and Map Actions is chosen. Runs layout algorithm on the server for the current map state and attempts to distribute the coordinates of the nodes' X and Y positions in an organic manner.

Clean-up References

Clean-up References can be selected from the contextual menu when the **More**  button is clicked on the path control network map view and Map Actions is chosen. When network elements go down, NSP does not automatically delete topology data. This allows NBI configurations to be preserved. However, when Clean-up References is selected, any such objects found in the IGP topology layer will be deleted. This prevents the accumulation of stale data. The network map is simultaneously rebuilt to reflect these changes.

Toggle Hidden Devices

Toggle Hidden Devices can be selected from the contextual menu when the **More**  button is clicked on the path control network map view and Map Actions is chosen. Depending on supported feature sets, some network elements may not appear on the network map. This option enables visibility of those entities.

1.5 How does NSP function as a PCE?

1.5.1 NSP as a PCE

NSP can function as a PCE, and contains the logic to calculate paths. The VSR-NRC is a component of NSP, but does not calculate any paths. The VSR-NRC terminates PCEP connections and conveys path request messages from PCCs to NSP. NSP computes the requested path and responds to the VSR-NRC, which conveys the response to the PCCs. The communication between VSR-NRC and PCCs is accomplished using the PCEP protocol.

In order for NSP to compute paths, it must discover the IGP topology. Topology discovery can be performed by peering the VSR-NRC directly in the IGP or using BGP-LS. If using IGP, the VSR-NRC must have full visibility of the topology. For multi-area topologies, this means that the VSR-NRC must be connected to every area, or to the ABRs/(L1/L2s) via IGP (OSPF or ISIS) adjacencies. If using BGP-LS, the VSR-NRC must be peered with a BGP speaker, ABRs/(L1/L2s) that are BGP speakers, or a Router Reflector that is peered to a BGP speaker in each IGP area. In order for BGP-LS discovery to be successful, each BGP speaker must support BGP-LS.

i **Note:** Only the VSR-NRC supports topology discovery for NSP. Do not use any other devices, such as the vCPAA, because they are not supported.

For more information about configuring the VSR-NRC for use with NSP, see the *NSP Installation and Upgrade Guide*.

1.6 What path control algorithms does NSP support?

1.6.1 Star algorithm

NSP provides a load-balancing and optimal-path-placement algorithm, known as the STAR algorithm. This algorithm uses an internal metric, calculated from the current value of the TE bandwidth reservation, to route the CSPF paths. Every path that is allocated on a TE link changes the internal metric for both the link and the overall path. Initially, all links have the same star weight, or metric, so the first path requests for CSPF traversal will choose the shortest path that satisfies all constraints. If there are multiple paths that satisfy the user constraints, then a path will be chosen randomly. This behavior is the same for normal CSPF.

Subsequent requests will choose paths that possess the least star weight, thereby ignoring the path that the normal CSPF algorithm would have chosen. The calculation of the star weight is based on a formula that uses the current link reservation. The user constraints are still satisfied. This balances the overall network utilization.

The STAR algorithm is invoked per LSP by associating that LSP to a path profile policy. The path profile template is defined using NSP and requires setting the objective to use STAR WEIGHT. The path profile policy is specified with the LSP definition and is conveyed to NSP via a PCE request message.

1.6.2 Disjoint optimal path computation algorithm

NSP provides support for disjoint path computation between a source destination pair and between two pairs of sources and destinations. Applications can use this algorithm to provide no-impact redundancy for a service offering. The algorithm provides node/link and SRLG types of disjoint path computations. The algorithm can also re-optimize an existing path if a second path request asks to be disjoint from the existing path. The ability to treat a pair of paths as mutually disjoint requires associating a path profile ID to the path request. In addition, a path group ID specification is also essential to implicitly identify the path pair from other path pairs. The disjoint optimal path calculation algorithm can also compute paths that are bidirectionally symmetric, to ensure that forward and reverse traffic use the hops while being disjoint.

i **Note:** NSP can only compute bidirectionally symmetric forward or reverse paths. For an RSVP LSP with primary and secondary path specification, the profile is applied to both paths. For example, if there are two RSVP LSPs between the respective distinct sources and destinations, the primary path of LSP 1 will be mutually disjoint from the primary path of LSP2, and vice versa for secondary paths. The algorithm cannot be applied to ensure the primary and secondary paths between the same source and destination pair are mutually disjoint.

1.6.3 Global concurrent optimization algorithm

NSP also supports optimizing the paths of existing LSPs by applying an optimization algorithm. This algorithm extracts the current resource availability on the current topology and reroutes the selected LSP paths such that the overall network consumption is minimized. The result is to utilize more network links, but also reduce the consumption on the links. LSPs must be delegated to NSP and must be preselected. Profiles do not have to be associated to the paths in order to use this algorithm. The LSPs to be optimized are selected manually using NSP's service management views.



Note: The LSPs that have a profile with the disjoint option enabled are excluded.

1.7 What are flex algo definitions?

1.7.1 Flex algo definitions

Normally, routers construct a single shortest path tree based on the IGP metric, and any deviations from this shortest path tree - due to traffic engineering requirements, for example - would require per-path CSPF signaling and control protocols such as RSVP or SR-TE. However, when flexible algorithm - or Flex Algo - is used, routers in the network can each compute unique trees that respect a combination of TE constraints and objectives beyond the base IGP Metric, without the need for per-path signaling and control.

For example, a Flex Algo instance can be configured to use link delay instead of IGP metric, and admin group rules to restrict paths to a specific subset of the network. Each router is configured with the same definition, ensuring they all construct the same view of the shortest path while considering these rules. When a path/packet which follows these rules needs to be sent, traffic can be steered along the network using a SR Prefix SID associated with the Flex Algo definition. Each router in the path would have pre-computed a shortest path next-hop which respects the defined rules. Unlike SR-TE, Flex Algo can achieve forwarding that respects the rules with a single segment in the packet. Note, however, that Flex Algo definition is a reduced form of TE compared to what SR-TE can achieve, and while multiple Flex Algos can exist in the network, the current maximum is 5.

Path control implementation

Flex Algo is inherently a router-based traffic engineering solution for various use cases and can operate without the use of a central path computation. With that in mind, path control can still provide various functions with Flex Algo, including:

- IGP Flex Algo Definition state visibility and assurance
- Flex Algo Path and Tree visualization and assurance
- Compute end-to-end paths across multiple domains with Flex Algo in each domain (local router unable to)
- Compute end-to-end paths where some domains use Flex Algo, and others use SR-TE (local router unable to)
- Leverage Flex Algo as an MSD reduction tool and combine with additional TE (PCE-computed diversity, bandwidth, latency constraints, etc. combined with a Flex Algo configured for Delay)

 **Note:** Path control does not support Flex Algo-based computations with PCEP LSPs.

Flex Algo Definition (FAD)

The Flex Algo Definition (FAD) is a core component of Flex Algo. It is a configuration definition applied on each router in the network that is partaking in a Flex Algo instance, and it is flooded to all routers in the local area to ensure consistency and prevent loops. The FAD is also exported to BGP-LS and can be discovered by a PCE/Controller.

The FAD is what makes up the traffic engineering rules to impose on the algorithm, and includes the options below:

| Parameter | Description |
|-------------------------|--|
| Algo | Specifies the assigned value of the Flex Algo instance. Range is 128-255. Algo 0 is also known as the native IGP Metric algorithm. |
| Priority | Specifies the priority of the Flex Algo instance. |
| Metric Type | Specifies the metric type of the Flex Algo instance. Options are IGP, TE-Metric, or Delay. |
| Admin Group Include-All | Specifies whether the Flex Algo instance will include all admin groups. |
| Admin Group Include-Any | Specifies whether the Flex Algo instance will include any admin groups. |
| Admin Group Exclude | Specifies whether the Flex Algo instance will exclude admin groups. |
| Exclude SRLG | Specifies whether the Flex Algo instance will exclude SRLGs. Not supported by SROS. |

Winning FAD

Since a FAD is configured individually on each router, there must be a distributed consensus to ensure that all routers agree on the same definition and converge to the same topological value. This is known as the winning FAD. The winning FAD occurs for each local link state flooding area. The winning FAD is selected based on the following criteria:

- From the advertisements of the FAD in the area (including both locally generated advertisements and received advertisements), the one(s) with the numerically-greatest priority value is selected.
- If there are multiple advertisements of the FAD with the same numerically-greatest priority, the one that is originated from the router with the numerically-greatest System-ID (in the case of IS-IS) or Router ID (in the case of OSPFv2 and OSPFv3) is selected.

1.8 How is inventory synchronized?

1.8.1 Synchronizing inventory

NSP requires inventory information for certain features. On startup, an initial synchronization is triggered that pages through the RESTCONF inventory find API and synchronizes inventory information. From there, ongoing synchronization consumes kafka messages from the nsp-yang-model.change-notif kafka topic. This handles the management, unmanagement, and modification of new NEs after NSP startup.

i **Note:** The openconfig logical inventory adaptors - such as sros-oc-logical-inventory - must be installed in MDM to enable it to synchronize openconfig interface information from the managed NE.

1.8.2 Xpaths

For inventory synchronization to work, the following xpaths must be populated in NSP by either MDM or NFM-P for each node being managed. You can confirm that this data is being populated using the NSP RESTCONF inventory find API. For more information, see the [Network Developer Portal](#).

Xpath: `/nsp-equipment:network/network-element`

Xpath: `openconfig-interfaces:interfaces/interface/subinterfaces/subinterface`

The following is an example of a POST call using the NSP RESTCONF inventory find API:

```
POST /restconf/operations/nsp-inventory:find
{
  "input": {
    "xpath-filter": "/nsp-equipment:network/network-element",
    "fields": "ne-id;ne-name;type;version;mac-address;ip-address;
admin-state;oper-state",
    "limit": 10
  }
}
```

1.8.3 Manual inventory synchronization

If synchronization is lost on either the NSP side or on the MDM side, the below APIs may be used to return to a synchronized state.

If NSP is out of sync, use the following API to trigger an initial synchronization again:

GET `/sdn/api/v4/system/resync-nms`

i **Note:** This also triggers re-synchronization for the Original Service Fulfillment application. If your network is large, or contains many services, this could potentially be a time-consuming operation.

If MDM is out of sync, use the following API to force MDM to re-synchronize logical inventory data from the node:

```
POST /restconf/operations/nsp-admin-resync:trigger-resync

{
  "nsp-admin-resync:input": {
    "plugin-id": "mdm",
    "network-element": [
      {
        "ne-id": "11.50.150.51",
        "sbi-classes": [
          { "class-id": "openconfig-interfaces:
/interfaces/interface/subinterfaces/subinterface/openconfig-if-ip:
ipv4/addresses/address"
        }
      ]
    ]
  }
}
```

Lastly, you can use the following API to set up a polling policy for a specific path (also known as re-synchronization policy):

```
PATCH /restconf/data/nsp-admin-resync:resync-policies/mdm/resync-policy=
default%20policy/node=SR-7750/version=22.7.R1/entity=
openconfig-interfaces%3A%2Finterfaces%2Finterface%2Fsubinterfaces%
2Fsubinterface%2Fopenconfig-if-ip:ipv4%2Faddresses%2Faddress

{
  "entity": [
    {
      "entity-type": "openconfig-interfaces:
/interfaces/interface/subinterfaces/subinterface/openconfig-if-ip:
ipv4/addresses/address",
      "period": 1,
      "admin-state": "enabled",
    }
  ]
}
```

```
"description": "Network Interface Polling modify"  
}  
]  
}
```

1.9 How are IRO objects used for path calculation?

1.9.1 IRO object for path calculation

NSP supports the IRO object specification within a PCC request. NSP computes a CSPF path from the source to the IRO object, and another CSFP path from the IRO object to the destination. If the second CSPF path visits any of the nodes in first CSPF path, the path computation fails.

When used with a path profile policy that contains the bidirectional disjoint specification, a forward LSP and its matching reverse LSP must share the same IRO configuration. This means that the list of addresses in the IRO path must be the same, but their order reversed. This is because the disjoint algorithm is natively bidirectional strict. If the reverse LSP contained IROs that did not exist in the forward path, no path would be found, because it would no longer be bidirectional strict.

1.10 What SROS path computation methods does NSP support?

1.10.1 Supported SROS path computation methods

NSP supports both Stateful and Stateless Bringup as described in draft-koldychev-pce-operational.

Nokia recommends using Stateless Bringup for interworking with SROS. This method uses the SROS CLI command “path-computation-method pce”.

The use of PCReq, followed by PCRep, enables the NSP PCE to be first to calculate paths. The NSP PCE can also consider all Path Profile and Association Group parameters - such as Diversity - when doing so.

Note that these PCReq and PCRep interactions add additional overhead. When the PCE is experiencing heavy load, PCReqs can timeout, and the subsequent retries from the PCC can exacerbate this as NSP attempts to manage the timeouts.

The use of the SROS CLI command “path-computation-method local-cspf” combined with “pce-control” is also supported, but not recommended for most deployments. In this scenario, the SROS router first calculates a path before delegating control to the PCE. Upon reception, the PCE will accept the path, if it is valid - otherwise, it calculates a new path. If the path is valid, it will be used, but may not immediately meet the Path Profile or Association Group parameters, such as Diversity. Only after the standard re-signal timer (which defaults to between 20 and 50 minutes) will all Path Profile and Association Group parameters be considered.

1.11 How do I access path control APIs?

1.11.1 Path control APIs

NSP's path control functions are available for OSS using programmable APIs. For general information about developer support, visit the [Network Developer Portal](#).

For information about the specific REST APIs used for path control, append **/sdn/doc** to the server URL. For example: `https://<NSP_cluster>:443/sdn/doc`.

Where *NSP_cluster* is the IP address of the NSP cluster.

1.12 How does path control implement user access control?

1.12.1 User access control for path control

NSP's path control function grants or restricts access to specific capabilities based on a user's permissions, as defined in NSP's Users and Security views. Customers upgrading to NSP Release 23.11 or later from any earlier NSP release must manually edit their existing roles in order to align with the new permissions defined below. Previous permissions are no longer enforced by the path control UI and APIs. See the *NSP System Administrator Guide* for more information.

In NSP Release 23.11, the following new permissions are introduced to further refine a user's access to path control functions:

- **Read** - View all available information. Cannot make any changes, except those allowed through Map Palette actions as those changes are applied only for the user that performed the action. The Read permission is a prerequisite for granting additional permissions.
- **Engineering** - Available for selection when Read is enabled. Allows the user to perform the following:
 - Configure global system configuration settings such as TWAMP test modes, Enable/disable historical data collection, SR policy configuration, TCA configuration policy, traffic data collection parameters, maintenance mode policy
 - Create, edit, and delete Path Profile policies
 - Create, edit, and delete Router ID Mapping policies
 - Update configuration parameters of links, such as srlg-value, latency, te-metric, igp-metric, administrative-group, admin-status, measuredIpBw, and measuredMplsBw
 - Clean up topology references
 - Trigger a data synchronization with the connected NMS
 - Force NSP Plugin to form a specific connection
 - Create, edit, and delete network map (Rebuild, Auto layout, Toggle hidden devices, Restore common map layout, Save map changes, and Export network)
 - Map Palette actions
- **Operations** - Available for selection when Read is enabled. Allows the user to perform the following:
 - Create, edit, delete, shutdown, and no shutdown on PCE-initiated LSP paths
 - Create, edit, delete, admin up, and admin down on SR Policies and Candidate Paths
 - Configure Path Profile Override for both PCE and PCC initiated LSP paths
 - Reset Failed Override Profile and Reset Controlled Reroute Path on LSPs
 - Activate and deactivate maintenance mode on nodes and links
 - Run workflows associated with nodes, links, and LSPs
 - Create, edit, and delete network map (Rebuild, Auto layout, Toggle hidden devices, Restore common map layout, Save map changes, and Export network)
 - Map Palette actions
- **Troubleshooting** - Available for selection when Read is enabled. Allows the user to perform the following:
 - Resignal, optimize, shutdown, no shutdown, on LSP paths
 - Resignal, admin up, and admin down on SR Policies and Candidate Paths

-
- Create, edit, and delete network map (Rebuild, Auto layout, Toggle hidden devices, Restore common map layout, Save map changes, and Export network)
 - Map Palette actions
 - Path Finder
 - Run Path Diagnostics

The above permissions can be combined to create the following roles from NSP's Users and Security views:



Note: These roles must be created by the customer and are not provided as pre-configured roles.

| Role | Permissions |
|----------------|--|
| Engineer | Engineering Operations Troubleshooting |
| LSP Operator | Operations Troubleshooting |
| Troubleshooter | Troubleshooting |
| Read-Only | Read |

1.13 How do I associate a workflow with a node, link, or LSP?

1.13.1 Associating workflows

This procedure is used to associate workflows with nodes, links, or LSPs in the NSP's path control views.

i **Note:** Workflow definitions must include the "IP/MPLS Optimization" tag in order to be eligible for association with a node, link, or LSP in the NSP's path control views.

Workflows must contain the following inputs:

- **payload** - includes maintenance mode, which is defined using NSP's path control function. NSP's path control function uses system configuration to set this field to either 'AUTOMATIC' or 'MANUAL'.
- **token_auth**
- **rest_gateway_host** - hard-coded in workflow definition. IP address of NSP REST server.
- **status** - hard-coded in workflow definition. Valid options: 'MAINTENANCE', 'UP'. Applicable only to node workflows.

The following is an example of required workflow tags and inputs - in this case, for placing a node into maintenance mode:

nodeIntoMaint:

```
type: direct
tags:
- IP/MPLS Optimization
input:
- token_auth
- payload
- rest_gateway_host: 'xxx.xxx.xxx.xxx'
- status: 'MAINTENANCE'
```


1.13.2 Steps

1

Using NSP's workflows function, import or create a workflow to associate with a node, link, or LSP. See the *NSP Network Automation Guide* for more information.

2

Perform one of the following:

- a. From the **Path Control, Network Map** view, select a node or link from the map, then click **More** , **Show workflows**.

-
- b. From the **Path Control, Routers** view, click on the Show workflows button in-line with any node.
 - c. From the **Path Control, Links** view, click on the Show workflows button in-line with any link.
 - d. From the **Path control, LSPs** view, click on the Show workflows button in-line with any LSP.

3

A dialog box appears with a list of existing workflows. Select the workflow created in [Step 1](#) and click **RUN WORKFLOW**.

END OF STEPS

1.14 How do I enable automatic VSR-NRC site switchover?

1.14.1 Automatic VSR-NRC site switchover

This procedure can be used to enable automatic switchover from the primary VSR-NRC site to the inactive VSR-NRC site.

If the primary VSR-NRC loses Cproto channel connectivity to the NSP host server, the NSP host server raises a cprotoChannelDown alarm. This alarm triggers a workflow. After a configurable down timer (with a default value of 5 minutes) expires, the workflow checks whether the cprotoChannelDown alarm has been cleared. If it has, cproto connectivity is up, and the workflow doesn't continue. However, if the alarm still exists, the workflow will ping the VSR-NRC at the inactive site to ensure it is in better shape before initiating site switchover.

1.14.2 Steps

Create a Kafka Trigger

1

From the **Workflows, Kafka Triggers** view, click **CREATE KAFKA TRIGGER**. The Create Kafka Trigger form opens.

2

Configure the parameters as follows:

- **Workflow (PUBLISHED):** switchover
- **Kafka Topic:** nsp-db-fm
- **Trigger Rule:** `[$?(@.alarmName == 'CprotoChannelDown' && @.severity == 'major')]`
- **Kafka Event:** CREATE
- Enable the **ENABLED** check box

3

Click **CREATE**. The Kafka Trigger is created.

Create an Environment

4


From the **Workflows, Environments** view, click **CREATE ENVIRONMENT**. The Create Environment form opens.

5

Add the following variables:

- **username:** The username used to log in to the NSP host server
- **password:** The password used with the above username
- **ping_duration:** 3

-
- **delay:** 300
 - **site1_vsr:** The IP address of the primary VSR-NRC
 - **site1_wfm:** The private IP address of the primary Workflow Manager
 - **site1_vip_advertised:** The advertised IP address of the primary NSP host server
 - **site2_vsr:** The IP address of the standby VSR-NRC
 - **site2_wfm:** The private IP address of the standby Workflow Manager
 - **site2_vip_advertised:** The advertised IP address of the standby NSP host server
 - **site2_sdn1:** The private IP address of the first member of the standby NSP host server cluster
 - **site2_sdn2:** The private IP address of the second member of the standby NSP host server cluster
 - **site2_sdn3:** The private IP address of the third member of the standby NSP host server cluster

 **Note:** If the NSP host server was deployed with 1+1 redundancy, the values of site2_sdn1, site2_sdn2, and site2_sdn3 should be set to 'null'.

Import or create workflow

6

Perform one of the following:

- a. From the **Workflows, All Workflows** view, choose **IMPORT > File System**.
- b. From the **Workflows, All Workflows** view, choose **IMPORT > GitHub**.
- c. From the **Workflows, All Workflows** view, click **⊕WORKFLOW**.

7

See the *NSP Network Automation Guide* for more information about creating workflows.

END OF STEPS

1.15 How do I visualize detailed metrics?

1.15.1 Visualizing detailed metrics

Users can access detailed metrics from the Grafana dashboard, which plots various internal system metrics over time. Some of the metrics that can be found on the Grafana dashboard include path updates per second, BGP-LS NLRI events per second, internal worker queue sizes, and LSP latency recompute rates.

1.15.2 Steps

1

From the **Path Control, Network Info and Statistics** view, click **Go to detailed Grafana metrics**. The Path Control metrics Grafana dashboard opens.

2

Browse the visualizations for resource utilization, worker queues, and internal message rates.

END OF STEPS

2 Path simulation overview

2.1 How does NSP enable path simulation?

2.1.1 Path simulation for NSP

NSP's path simulation function provides the ability to simulate changes in the IP topology that was discovered by using NSP's path control function.


This function is run in a separate VM and imports the IP topology and LSPs from NSP path control. The specific simulation functions supported are:

1. Modifying link attributes
2. Modifying the status links
3. Creating or deleting LSPs
4. Modifying the profile of imported LSPs
5. Optimizing LSPs via the GCO algorithm

For each change, the Simulate button is activated to determine the visual impact of that change.

The general functionality supported in addition to the above functionality are:

1. Import IP/TE topology only
2. Import LSPs
3. Import profiles
4. Import SR policies
5. Delete an imported topology


 **Note:** SR policies that are operationally down cannot be imported into path simulation.

2.2 How do I modify the path simulation network map?


2.2.1 Available options

The path simulation network map is a representation of an imported network. The following options can be used to configure the path simulation network map:


Refresh

The **Refresh**  button is located on the path simulation network map view. When clicked, the latest version of the map is fetched.


Search In Map

The **Search In Map**  button is located on the path simulation network map view. When clicked, a window opens, in which you can provide device-specific information in order to locate a Link, Router, Prefix, or Region (when nodes are clustered by region) on the network map.


Export Network

Export Network can be selected from the contextual menu when the **More**  button is clicked on the path simulation network map view. This option is used to export the current network as a .zip file so that it can be imported into the Simulation tool for testing purposes.


Save Map Changes

Save Map changes can be selected from the contextual menu when the **More**  is clicked on the path simulation network map view. When the NSP's Map Layout and Groups views are part of the deployment, its path simulation function fetches and utilizes the node positions from that data, but users can reposition nodes on the network map as required. When Save Map changes is selected, the new node positions are persisted only in the path simulation network map.


Restore Common Map Layout

Restore Common Map Layout can be selected from the contextual menu when the **More**  button is clicked on the path simulation network map view. When the NSP's Map Layout and Groups views are part of the deployment, its path simulation function fetches and utilizes the node positions from that data, but users can reposition nodes on the network map as required. Users can then revert to the map configuration reflected in the May Layout and Groups views by selecting Restore Common Map Layout.


Delete Topology and LSPs

Delete Topology and LSPs can be selected from the contextual menu when the **More**  button is clicked on the path simulation network map view. This option deletes the existing network topology, and all associated LSPs from the path simulation views - allowing for a new network to be imported.


Import Network

Import Network can be selected from the contextual menu when the **More**  button is clicked on the path simulation network map view. See [2.5 "How do I import a network?" \(p. 38\)](#) for more information about this option.


Delete All LSPs and SR Policies

Delete All LSPs and SR Policies can be selected from the contextual menu when the **More**  is clicked on the path simulation network map view. This option deletes all existing LSPs and SR policies from the path simulation views.


Rebuild

Rebuild can be selected from the contextual menu when the **More**  button is clicked on the path simulation network map view. This options deletes, then rebuilds the network map to ensure synchronization with the IGP topology layer.


Auto Layout

Auto Layout can be selected from the contextual menu when the **More**  button is clicked on the path simulation network map view and Map Actions is chosen. Runs layout algorithm on the server for the current map state and attempts to distribute the coordinates of the nodes' X and Y positions in an organic manner.

Clean-up References

Clean-up References can be selected from the contextual menu when the **More**  button is clicked on the path simulation network map view and Map Actions is chosen. When network elements go down, NSP does not automatically delete topology data. This allows NBI configurations to be preserved. However, when Clean-up References is selected, any such objects found in the IGP topology layer will be deleted. This prevents the accumulation of stale data. The network map is simultaneously rebuilt to reflect these changes.

Toggle Hidden Devices

Toggle Hidden Devices can be selected from the contextual menu when the **More**  button is clicked on the path simulation network map view and Map Actions is chosen. Depending on supported feature sets, some network elements may not appear on the network map. This option enables visibility of those entities.

2.3 How does path simulation implement user access control?


2.3.1 User access control for path simulation

NSP's path simulation function grants or restricts access to specific capabilities based on a user's permissions, as defined in NSP's Users and Security views. Customers upgrading to NSP Release 23.11 or later from any earlier NSP release must manually edit their existing roles in order to align with the new permissions defined below. Previous permissions are no longer enforced by the path simulation UI and APIs. See the *NSP System Administrator Guide* for more information.

In NSP Release 23.11, the following new permissions are introduced to further refine a user's access to path simulation functions:

- **Read** - View all available information. Cannot make any changes, except those allowed through Map Palette actions, which are applied only for the user that performed the action.
- **Read/Write** - Complete access to all path simulation functions.

The above permissions can be used to create the following roles from NSP's Users and Security views:

 **Note:** These roles must be created by the customer and are not provided as pre-configured roles.

| Role | Permissions |
|-----------|-------------|
| Simulator | Read/Write |
| Read-Only | Read |

2.4 How do I access path simulation APIs?

2.4.1 Path simulation APIs

NSP's path simulation functions are available for OSS using programmable APIs. For general information about developer support, visit the [Network Developer Portal](#).


For information about the specific REST APIs used by NSP's path simulation function, append **/sdn/doc** to the server URL. For example: `https://<NSP_cluster>:9543/sdn/doc`.

Where *NSP_cluster* is the IP address of the NSP cluster.

2.5 How do I import a network?


2.5.1 Steps

1

If the network will be imported from a file, the file must first be created navigating to the path control network map, clicking the **More**  button, and choosing **Export Network**. The network is exported as a .zip file.

2

Go to **Path Simulation, Network Map** and perform one of the following:

- a. If opening the path simulation network map for the first time, the Import Network form will appear.
- b. Click the **More**  button on the path simulation network map view and choose **Import Network**. The Import Network form appears.

3

Perform one of the following:

- a. Enable the **From Live Network** radio button, then either the **Network** radio button or the **Network and LSPs** radio button to specify which elements of the network to import, then configure the parameters:

| Parameter | Description |
|------------------------------|--|
| Remote CAS server IP address | The IP address of the server from which to retrieve a security token |
| Username | The username with which to log in to the server |
| Password | The password to be used with the specified username |

- b. Enable the **From File** radio button, then either click **browse** to select the file created in [Step 1](#), or drag and drop the file into the Import Network form.

4

Click **IMPORT**. The Network topology is imported.



Note: The default maximum file size for an imported network topology is 5 M. This can be modified by editing the arm-system.conf file as follows:

```
nrcp {  
  file-import  
  { max_file_size_in_byte = <desired_size> }  
}
```

END OF STEPS

2.6 How do I add an admin domain?

2.6.1 Steps

1

From **Path Simulation, Network Map**, click  and choose **Add Domain**. The Create Admin Domain form appears.

2

Configure the parameters:

| Parameter | Description |
|------------|--|
| Protocol | Specifies the protocol, IGP (OSPF/ISIS) or BGP |
| Network ID | Specifies the network identifier |
| AS number | Specifies the AS number |
| BGPLS ID | Specifies the BGP-LS identifier |

3

Click **CREATE**. The admin domain is added to the network topology.



Note: An admin domain is automatically deleted if not populated with network elements.




END OF STEPS

2.7 How do I run a worst case failure scenario?

2.7.1 Purpose

This procedure can be used to simulate a worst case failure scenario, in which multiple operational links are turned down.

2.7.2 Steps

- 1 _____
From the **Path Simulation, Network Map**, click **Run Scenario** and then choose one of the following:
 - a. Choose Run Worst Case Failure on All Operational Links. Go to [Step 3](#).
 - b. Choose Run Worst Case Failure on Select Links. A list of available operational links appears. Continue to [Step 2](#).
- 2 _____
Select one or more links from the list and click **RUN SCENARIO**.
- 3 _____
A dialog box appears, displaying the simulation's progress. When complete, click **VIEW SIMULATION RESULTS**.
- 4 _____
A list of links is displayed. Perform any of the following:
 - a. Click **Show on Map**  inline with any link to highlight that link on the path simulation network map.
 - b. Click **More**  inline with any link and choose LSP Paths Impacted to view a list of LSP paths that were affected by the selected link's failure.
 - c. Click **More**  inline with any link and choose Links Impacted to view a list of Links that were affected by the selected link's failure.

END OF STEPS

3 LSPs

3.1 What types of LSPs does NSP support?

3.1.1 PCC-initiated LSPs

NSP supports the creation of PCC-initiated Segment-Routed TE LSPs, as well as PCC-initiated RSVP LSPs. NSP's service management function sends a service creation request, through NFM-P, to the PCC router(s) endpoints which are part of the service. An empty LSP path is created on each of the PCC router(s). The PCC router(s) then send an LSP path request to NSP (PCE). NSP (PCE) computes the LSP path and sends the path to the PCC. The PCC then installs the computed path and informs NSP (PCE) that the path is ready. This is reported to NSP's service management function, where the path is attached to the service.

3.1.2 PCE-initiated LSPs

NSP supports the creation of PCE-initiated Segment-Routed TE LSPs. Operators can specify the LSP parameters and PCC address using an LSP creation form within NSP's path control function, or by using the NSP API. Operators can also select a path profile policy to associate to the LSP. There is also an NBI. PCE-initiated LSPs are deployed through PCEP.

In order to create a , the following commands must be executed on the node:

1. `config>router>pcep>pcc`
`max-srte-pce-init-lsps <max-number>`
 Where *max-number* is a number between 0 and 8191, which can only be modified when `config>router>pcep>pcc` is shutdown.
2. `config>router>mpls# lsp-template template-name pce-init-p2p-srte`
`{default | template-id } default-path {pathname}`
`path "P"`
`no shutdown`
`exit`
`lsp-template "test" pce-init-p2p-srte template-id default`
`default-path "P"`
`cspf`
`pce-report enable`
`no shutdown`
`exit`
3. `config>router>mpls>`
`pce-initiated-lsp sr-te`
`no shutdown`


3.1.3 RSVP LSPs

Any PCC node intending to request a path computation from NSP must first set the PCE computation option in the LSP definition. The PCC then assigns a unique PLSP-ID to the LSP. This uniquely identifies the LSP within a PCEP session and is maintained for the lifetime of the LSP. The PLSP-ID is also associated to the tunnel and path ID.


Once the PLSP-ID is assigned, the PCC sends a PCReq message to NSP PCE, requesting a path for the LSP. This request includes the LSP parameters in the METRIC object, the LSPA object, and the Bandwidth object. It also includes the LSP object with the selected PLSP-ID. NSP is now able to compute a new path, to check the bandwidth, and to return the path in a PCRep message with the computed Explicit Router Object (ERO) in the ERO object. It also includes the LSP object with the unique PLSP-ID, the METRIC object with the computed metric value (if any), and the Bandwidth object.

NSP does not keep track of the LSP yet. At this point, it has simply returned the ERO. The PCC has yet to confirm that the path was signaled. If the path was locally signaled, and the local TE database (TEDB) was updated, NSP receives the updates via BGP-LS and update its TEDB.

For stateful operation, which allows NSP to track the LSP path and bandwidth (among other constraints), the PCE report option must be set in the LSP definition. When this option is set, the PCC sends both a PCRpt message to update NSP with the state of UP, and the Record Route Object (RRO) object as confirmation. The RRO object now includes the LSP object with the unique PLSP-ID. With this, NSP is able to display the LSP, as well as its hops and constraints on the path control network map. The RRO also contains information about the protection that is enabled on the signaled path. Therefore, NSP is aware of the protection at the hops, but not aware of the detour/bypass tunnel details. If a local failure causes the LSP on the PCC to switch to a detour or bypass, a PCE report is sent to NSP, and NSP becomes aware that the LSP is using a detour or bypass.

 **Note:** In the VSR-NRC, the PCE reporting option can either be set globally, or on a per LSP basis.

The PCC can also delegate control of the LSP to NSP for either active control or LSP optimization. This is known as active stateful behavior. The delegation is awarded using the PCE control option. Once NSP is controlling the LSP, the operator can manually re-signal/re-optimize the LSP. Re-signalling routes the LSP using its original constraints, while re-optimizing routes the LSP using an optimization algorithm. NSP also reroutes LSPs automatically on resource failures, or when calculating disjoint paths.

 **Note:** When the PCC has delegated control of the LSP to NSP, any change to the LSP definition (such as changes in constraints) requires the PCC to first revoke the delegation via the PCE report option, and then to issue a new request to NSP.

Secondary path behavior

The PCC sends PCE requests for standby secondary paths. A new PLSP-ID is used for these paths over the PCEP session, and is associated to the LSP path ID and the LSP tunnel ID. When a secondary path is not in standby, the PCE request is not sent until the primary path is down, or in FRR. However, if the path is delegated to NSP, this results in a PCE update from NSP. The LSP may switch to the secondary path in the interim, but will switch back to the primary path as soon as possible.

NSP maintains the active path in case both the primary and secondary paths are signaled, and also when the primary path is down. NSP also maintains the shared explicit behavior when the primary and secondary paths share common link resources.

NSP also indicates the active path between the primary and secondary pair.

FRR notification

Fast reroute (FRR) is signaled locally, with locally-created detour tunnels. These tunnels are not reported to NSP, and therefore NSP is not aware of the detours and bypass. However, the types of node and/or link protection are communicated to NSP via the PCE report.

i **Note:** All the RSVP-TE LSPs created by NSP have FRR enabled by default. The FRR method used is “facility”.

Bandwidth management

NSP manages the LSP bandwidth consumption on the TE links for both stateless and stateful PCC configurations. In a stateless configuration, NSP receives TE updates from the network as LSPs are signaled, thereby mimicking the TE DB bandwidth consumption on the nodes. This allows for accurate LSP path computation without maintaining state on NSP. In a stateful case, wherein the reports are sent to NSP from the PCC, the bandwidth is again communicated by the PCC to NSP via the bandwidth object. Here, NSP will reconcile the TE update with the specific LSP bandwidth update via the report. Therefore, NSP maintains full LSP state along with the consumption on the TE links for these LSPs only.

It is possible that existing brownfield LSPs will not request paths from NSP, and therefore, will have no state within NSP. NSP will not display these LSP reservations on the TE links. For a mixture of LSPs that are PCE-reported and non-PCE-reported, NSP will track and show the actual TE consumption on a TE link in addition to the LSP reservation for PCE-reported LSPs.

Although bandwidth is not tracked until reported, bandwidth is reserved for one (1) minute when a request is made. Therefore, if multiple requests are made in quick succession, subsequent requests will be impacted, even though reports have not yet been received.

3.1.4 Segment-routed TE LSPs

Any PCC node intending to request a path computation from NSP must first set the PCE computation option in the LSP definition. The PCC then assigns a unique Path LSP-ID (PLSP-ID) to the LSP. This uniquely identifies the LSP within a PCEP session and is maintained for the lifetime of the LSP. The PLSP-ID is also associated to the tunnel and path ID.

Once the PLSP-ID is assigned, the PCC sends a PCReq message to NSP PCE, requesting a path for the LSP. This request includes the LSP parameters in the SRP object, the METRIC object, the LSPA object, and the Bandwidth object. It also includes the LSP object with the selected PLSP-ID. NSP will reserve bandwidth for the path to be returned, but will not keep track of the operational status or other requirements for the LSP yet. At this point, bandwidth is consumed and an ERO is returned. The PCC has yet to confirm that the path was signaled. If the path was locally signaled, and the local TEDB has been updated, NSP will receive a REPORT from the PCC and the updates via BGP-LS and update its TEDB. If the PCC fails to send a report, after a period of time the bandwidth reserved will be released from NSP. The path computed by NSP is specified explicitly with the next hop interfaces and the adjacency SIDs encoded in the SR ERO sub-object.

When the PCE report option is set in the LSP definition, the PCC sends both a PCRpt message to update NSP with the state of UP, and the RRO object as confirmation. The RRO object now includes the LSP object with the unique PLSP-ID. With this, NSP is able to display the LSP, as well as its hops and constraints. The RRO also contains information about the protection that is enabled on the signaled path. Therefore, NSP is aware of the protection at the hops, but not aware of the detour/bypass tunnel details. If a local failure causes the LSP on the PCC to switch to a detour or bypass, a PCE report is sent to NSP, and NSP becomes aware that the LSP is using a detour or bypass.

i **Note:** In the VSR-NRC, the PCE reporting option can either be set globally, or on a per LSP basis.

The PCC can also delegate control of the LSP to NSP for either active control or LSP optimization. This is known as active stateful behavior. The delegation is awarded using the PCE control option. Once NSP is controlling the LSP, the operator can manually re-signal/re-optimize the LSP. Re-signalling routes the LSP using its original constraints, while re-optimizing routes the LSP using an optimization algorithm. NSP also reroutes LSPs automatically on resource failures, or when calculating disjoint paths.

i **Note:** When the PCC has delegated control of the LSP to NSP, any change to the LSP definition (such as changes in constraints), requires the PCC to first revoke the delegation via the PCE report option, and then issue a new request to NSP.

SR-TE LSPs with PCE delegation support primary and secondary LSP paths. Switching between the LSP paths is accomplished on the node using sBFD to provide quick switches at the head-end when TI-LFA is unable to solve all network faults. NSP can force the primary path to be the most optimal path of a diverse pair by choosing a path-profile with diversity set and by setting the priority <setup-priority> appropriately on the primary and secondary paths.

Bandwidth management

A bandwidth value that is specified on an LSP has no significance on the PCC/router because the SR TE does not maintain any state on the intermediate or destination routers. Therefore, no bandwidth tracking is done in the local TE DB. The bandwidth has to be tracked by NSP if the LSP is configured to report bandwidth. Bandwidth tracking on NSP is done only after a valid PCE report message is generated by the PCC. NSP tracks the bandwidth reservation for SR TE LSPs separate from RSVP TE LSPs.

i **Note:** A loose hop SR LSP whose bandwidth is specified and computed locally will not be tracked by NSP, even with the PCE report option enabled. NSP only tracks SR TE LSP paths computed by NSP itself.

Although bandwidth is not tracked until reported, bandwidth is reserved for one (1) minute when a request is made. Therefore, if multiple requests are made in quick succession, subsequent requests will be impacted, even though reports have not yet been received.

Failure detection

The head end router for an SR TE path, or an SR path, has no indication when a downstream link failure has impacted traffic for that SR TE or SR path. For a stateless and stateful application without PCE control, the SR TE tunnel on the head end router will remain up, as it receives no

notification from the control plane either locally, or via NSP. For an LSP with delegated control to NSP, NSP will react to the topology change and issue a new ERO update to the PCC via PCE update.

3.1.5 TE-ECMP routing

Traffic Engineered Equal Cost Multi-path routing (TE-ECMP) enables users to create multiple equal-cost paths that NSP controls as a single LSP, thereby achieving the same protection as a pair of disjoint services. TE-ECMP is ideally-suited to leaf-spine architectures, whereby load balancing can be accomplished by a leaf connecting to multiple spine switches, and/or multiple parallel links being created between spine switches. The flexibility of SR networks further enhances this solution, as – in addition to node SIDs – traffic can be directed to either Anycast SIDs (which can be associated with multiple spine nodes) or Adjacency Set SIDs (which are shared by the parallel links that comprise the set).

When TE-ECMP routing is used, NSP identifies multiple equal-cost paths between a source and destination based on an objective metric, as known as an 'ECMP Tree'. NSP then creates an Explicit Route Object (ERO) that captures how best to implement this tree using a combination of the available SID types. The path information is then sent to the originating node via PCEP. Capacity can then be added to the network seamlessly (for example, by introducing a new link to an Adjacency Set, or by adding a new spine switch into an Anycast SID group), because the EROs which define the path do not have to change. As the network evolves and the configuration of the SR fabric changes, NSP remains synchronized with this data and adjusts the ERO accordingly.

Path diversity (SRLG only)

Multiple SR TE LSPs, each with their own set of ECMP paths, can exist simultaneously and remain disjoint from one another. This is accomplished by specifying that the paths remain SRLG diverse (Shared Risk Link Group). As the paths traverse the links between data centers, connections can pass through the same conduit, making them subject to a single fiber cut. However, when the LSPs are specified as being SRLG diverse, only the links for which SRLG values have been configured are considered by the paths.

Blackholing prevention

When a single link in an adjacency set goes down, traffic is forwarded over the remaining links in the set – however, if all links in the set were to fail, blackholing could occur. To prevent this, NSP has the ability to learn topology changes and diagnose problems. It would calculate a new ERO in order to avoid failed links. For example, if a leaf switch were sending traffic to two spine switches using their shared Anycast SID, and one of those spine switches was forwarding traffic over the failed link, the ERO would begin sending traffic to the other spine switch using its unique Node SID instead, thereby bypassing the failed link.

Path persistency

When two LSPs are deployed with diversity, the goal is often to ensure that at least one of them will remain operationally up when failures occur in the network. However, these dual LSPs can serve other purposes, such as load balancing, or designating specific traffic for specific paths.

In these scenarios, the primary LSP typically follows the shortest path and has the lowest latency, while the secondary LSP typically takes a longer path and has a higher latency. As such, it is often necessary for services that are latency sensitive to be routed over the primary LSP. In the event that both LSPs go down, it is therefore important that the paths are not swapped during restoration.

For this reason, NSP supports path persistency. This ensures that, if the secondary LSP comes up before the primary, it will be rerouted off the shortest path once the primary LSP is again available.

SID types

Node SIDs and Adjacency SIDs both identify entities through which traffic must be routed, however, Adjacency SIDs identify an exact path, whereas Node SIDs follow the IGP shortest path. Therefore, when a link fails while Adjacency SIDs are being used, the path is down (unless Loop-Free Alternative kicks in). But when a link fails while Node SIDs are being used, traffic is rerouted (as long as the next node is still resolvable).

By extension, Anycast SIDs behave similar to Node SIDs, and Adjacency Set SIDs behave similar to Adjacency SIDs. Operators may have different preferences in terms of SID usage. One may prefer using Node SIDs because of their inherent load balancing and resiliency features, while another may prefer using Adjacency SIDs because the resulting path is strictly deterministic. TE-ECMP routing allows the operator to specify a preference in terms of SID usage.

Segment label depth

SR TE LSPs are limited by hardware support for the number of labels in the stack of a segment-routed path. In order to minimize the number of labels, the path must be comprised of a mixture of adjacency SIDs (which adhere to the TE path, but require labels for every hop), and node SIDs (which do not adhere to the TE path, but require fewer labels). An algorithm is introduced to employ this technique when the Explicit Route Strategy of the LSP's path profile policy is set to "Compressed".

i **Note:** To ensure that the node SID hops continue to adhere to the TE path, and that the path is valid, the controller must support IP path monitoring.

3.1.6 Anycast and loopback for LSPs

NSP supports path computation requests that include anycast or loopback addresses as destinations.

When inter-domain with multiple instances on routers are supported, NSP can specify a loose hop ERO with anycast loopbacks as intermediate hops. This allows for the generation of an inter-domain ERO between domains when domain boundary routers have anycast loopbacks configured.

The ERO generation is controlled via a path profile policy with a new ERO specification option field. If the specification is *anycast preferred*, then the inter-domain computed path will consist of border routers which have the anycast configuration as loopback addresses with identical anycast SIDs. If the specification is *loose hop preferred*, then the inter-domain computed path will consist of the best loose hop border routers with node SIDs.

i **Note:** Anycast SIDs are node SIDs that are associated to the loopback addresses instead of the system address. In SROS, there is no specific designation for anycast SIDs.

i **Note:** The ERO specification default is the complete path with Adjacency SIDs, however, in the inter-domain cases, the number of Adjacency SIDs will most likely exceed the MSD.

i **Note:** When the *anycast preferred* ERO specification is used and the inter-domain border routers do not have anycast SIDs, the best loose hop node SID among the inter-domain border routers will be selected.

3.1.7 BGP-LS Application Specific Link Attributes

VSR-NRC learns link attributes from the network using BGP-LS and/or IGP and then uses advertise these to NSP using BGP-LS. This information can then be used by NSP when performing path calculations. These link attributes are displayed within the path control network map and can be configured using the NSP's APIs. The BGP-LS Application Specific Link Attributes (ASLA) TLV is used to advertise the link attributes, including link delay attributes such as Unidirectional Link Delay and Min/Max Unidirectional Link Delay. These two attributes can be measured using TWAMP-light delay in order to provide latency information that can be used for rerouting LSPs. This advertisement is application specific, which allows for segment routing to be enabled and used without having to enabling RSVP in the network.

3.2 How does latency-based LSP rerouting work?

3.2.1 Latency-based LSP rerouting

NSP can optimize and reroute LSPs based on live link latency telemetry from TWAMP-light or EthCFM tests on the NE. NSP measures latency per link, and the end-to-end latency of an LSP is assumed to be the sum of the latencies of the links through which it is routed. If the latency of the LSP increases beyond the latency threshold or maximum latency specified in the LSP's path profile policy, NSP will re-signal the LSP, calculate a path with reduced latency, and reroute the LSP. Latency changes to links not on the LSP's path will not immediately trigger a reroute, however, LSPs are re-signaled at regular intervals to ensure that the paths they take are relatively optimal - in terms of latency - at any point in time. When an LSP is routed off its original path due to an increase in latency, it will only return back to its original path if the latency of that path is optimal at the time of the next re-signal. A re-signal can be triggered by:

- The end-to-end latency increases beyond either the latency threshold or the maximum latency configurations on the path profile policy associated with the LSP
- The LSP is configured to use a latency threshold of 0 and the end-to-end latency increases
- A re-signal timer fires (re-signal timers regularly fire on LSPs)

i **Note:** Changes to latency can make it necessary for many LSPs to be rerouted at once. To reduce the load on NSP - and to potentially reduce LSP reroutes - LSPs that cross their latency threshold constraint are randomly queued for rerouting within a one minute period by default. This may cause a delay in LSP movement after latency thresholds are initially crossed.

i **Note:** A delay may occur between the time the link latency measurement is collected and the time it appears on the link. Link latency updates are throttled in order to prevent excessive noise in large scale setups.

3.2.2 Modes

i **Note:** "modelDriven" latency collection mode should be used for both classic routers and ModelDriven routers as "classic" latency collection mode was deprecated in NSP Release 23.4.

The following API is used to enable latency collection:

PATCH /sdn/api/v4/nsp/configuration/latency

```
{  
  
  "data": {  
  
    "modelDriven": true,  
  
  }  
  
}
```

The model-driven option must be enabled in order to collect latency via either the MD-OAM application or the NFM-P (for 7210 and 7705 NEs). To do this, all actions to create, modify, or execute tests must be performed within either the MD-OAM application or the NFM-P (depending on NE type). When results are available for a given test, MD-OAM or NFM-P obtain the results, adds all stored test information to the results, and publishes the result to one of the following Kafka topics:

- `oam.test_execution` (for TWAMP or EthCFM using MD OAM)
- `oam.pm_results` (for TWAMP using NFM-P)
- `oam.test_results` (for EthCFM using NFM-P)

i **Note:** When using `modelDriven` mode to collect EthCFM latency telemetry from the MD-OAM app, the session mappings must still be defined under `ethcfmTests`.

i **Note:** If link latency measurements are not being applied on the link automatically, confirm the measurements are being produced to the corresponding Kafka topic above. If the results are showing up on the Kafka topic but not being applied to the link, the session name and IP link mapping must be configured explicitly using the latency configuration API below:

PATCH /sdn/api/v4/nsp/configuration/latency

```
{
  "data": {
    "twampTests": [
      {
        "session": "<TWAMP-light-session-name>", // Session name of the
TWAMP-Light test
        "source": "x.x.x.x", // Source interface IP
        "destination": "x.x.x.x" // Destination interface IP
      }
    ]
    "ethcfmTests": [
      {
        "session": "<EthCFM-session-name>", // Name of CFM Two Way Delay
Test
        "source": "x.x.x.x", // Source interface IP
        "destination": "x.x.x.x" // Destination interface IP
      }
    ]
  }
}
```

```
    }  
  ]  
}  
}
```

3.2.3 Latency timeout

Latency timeout values can be configured using the latency configuration API below:

PATCH /sdn/api/v4/nsp/configuration/latency.

```
"timeout": {  
  "enabled": true,  
  "staleTtl": <stale_timeout>  
  "expiredTtl" <expired_timeout>  
}
```

The timeout specified for “staleTtl” indicates, in seconds, when link latency will be considered stale. Links with stale latency are identified by an orange dot within the NSP’s path control views. Links with stale latency values are still considered during latency-based path optimization and are treated the same as links with valid latency values, however, stale links should be investigated by an operator. An alarm is raised in NSP when a link’s latency becomes stale.

The timeout specified for “expiredTtl” indicates, in seconds, when link latency will be considered expired. Links with expired latency are identified by a red dot within the NSP’s path control views. The path-optimization algorithm will attempt to route traffic away from links with an expired latency value, however, expired links can still be taken if they are the only remaining option. This functionality can be disabled by setting the parameter to 0. Links with no active latency are marked as undefined, and are identified by a dash (-) within the NSP’s path control views.

i **Note:** Latency timeout functionality is exclusive to telemetry sourced (MDM or NFM-P) OAM latency. If modelDriven mode is disabled, all MDM-sourced OAM latency values will be reset to 0, latency timestamps will be reset to null, and latency states will be reset to undefined. Links with API-configured latency values will not be affected.

3.2.4 Streaming template configuration

The type of latency measurements MD-OAM collects from nodes is based on the value of the fd-avg parameter that is configured within the node’s streaming template. The options are:

- *Round-trip:* NSP collects the round-trip latency value. This value is then divided by two and subsequently applied to the link, or link set, between the source and destination nodes in both directions. The latency is applied to a link set between the nodes if system IPs are used in the REST API and multiple links exist between the nodes. If round-trip is used, only one session needs to exist between the pair of nodes. Round-trip does not require any clock synchronization between the two nodes.

-
- *Forward:* NSP collects the forward latency value. This value is applied to the link only in the forward direction. The system clock needs to be synchronized between the nodes for the value to be accurate. This can be done using NTP.
 - *Backward:* NSP collects the backward latency value. This value is applied to the link only in the backward direction. The system clock needs to be synchronized between the nodes for the value to be accurate. This can be done using NTP.
- i** **Note:** Configuring the `fd-avg` parameter to collect the forward latency value, and exclusively supplying interface IPs (both in `config>oam-pm>session>ip config` and in the NSP REST API) will ensure that the TWAMP packets take the directly-connected link and report accurate forward latency to NSP, which will apply the forward latency to that link.

3.3 How do I view LSP path history?

3.3.1 LSP path history

Users of the NSP's path control function can view historical PCEP (RSVP or SR-TE) LSP data - including path history, IGP topology changes, and changes to bandwidth and latency - using REST APIs. To enable this functionality, the following APIs must be executed:


```
PATCH /sdn/api/v4/nsp/configuration/nrcp-historical
```

```
GET /sdn/api/v4/nsp/configuration/nrcp-historical
```

Where *NSP_cluster* is the IP address of the NSP cluster.

Once executed, historical LSP data is sent to the NSP's historical application and stored in a database. Additional APIs are then used to retrieve historical LSP data, as well as configure data retention policies. For more information, see the [Network Developer Portal](#).

By default, the database used by the NSP's historical application to store historical LSP data is the PostgreSQL database, however - if an auxiliary database configuration is present in the `nsp-config.yml` file - installation scripts will detect this and instead store the historical data in the auxiliary database. This can be overruled by setting the `app.auxdb_enabled` parameter in the `auxdb.conf` file to `false` and restarting the NSP server. The `auxdb.conf` file is located at `/opt/nsp/configure/config/` where the NSP historical application is deployed.

 **Note:** For large networks, Nokia recommends that historical LSP data be sent to the auxiliary database.

3.4 What are association groups?

3.4.1 Association groups

Association groups use a specific criteria to identify LSPs with some shared attribute and organize those LSPs into a group. In the NSP implementation, NSP can group LSPs with one of two common attributes.

PCE-initiated LSPs can be assigned to an association group of the subtype 'Policy' during creation (see [3.6 “How do I create PCE-initiated LSPs?” \(p. 55\)](#)). PCC-initiated LSPs can also be discovered with 'Policy' association groups already configured on the PCC. A 'Policy' association group is used to tag LSPs with traffic engineering criteria and policy behavior. The NSP will interpret this discovered policy object by retrieving an NSP-created path profile object definition that matches the same ID. The path profile traffic engineering criteria and behavior will be invoked on the LSP, however, the diversity and bi-directionality configurations specified inside of the path profile object will be ignored.

PCC-initiated LSPs can also be discovered with association groups of the subtype 'Disjoint' configured. This is used to group LSPs that must remain diverse from one another during path calculation based on shared constraints.

All existing association groups and their members can be viewed from the **Path Control, Policies** view.

3.5 How are BSIDs used to reduce MSD?

3.5.1 Reducing MSD with BSIDs

The NSP supports draft-ietf-pce-binding-label-sid for the advertisement of binding segment identifiers (BSIDs) for both PCC-initiated and PCE-initiated SR-TE LSPs in order to avoid exceeding their maximum stack depth (MSD). If and how this is accomplished is determined by the configuration of the path profile policy associated with a given LSP. See [5.1.2 “Path profile policy parameters” \(p. 87\)](#) for more information.

The NSP is able to discover LSPs with pre-existing BSIDs, and can use these BSIDs when attempting to reduce MSD. Alternatively, BSIDs can be created during either PCE-initiated LSP creation, or SR policy creation. See [3.6 “How do I create PCE-initiated LSPs?” \(p. 55\)](#) or [5.7 “How do I create an SR policy?” \(p. 100\)](#) for more information.

The NSP is also capable of auto-generating BSIDs of either type (PCEP LSP or BGP SR policy) when the configuration of a path profile policy indicates that they are required but none are available. Whether or not BSIDs can be auto-generated - and which type will be auto-generated - is specified via API. Visit the [Network Developer Portal](#) for more information.


3.6 How do I create PCE-initiated LSPs?

3.6.1 Purpose

i **Note:** Before a PCE-initiated LSP can be successfully created, certain commands must be executed on the source node. See the *NSP Release Description* for more information.

3.6.2 Steps

1

From the **Path Control, LSPs** view, click on **Create PCE-Init LSP** . The Create PCE-Init LSP form opens with the Identification panel displayed.

2

Configure the required parameters:

| Parameter | Description |
|------------------|--|
| Path Name | The name of the PCE-initiated LSP |
| PCC Address | The address of the PCC |
| Source | Specifies the source node for the path |
| Destination | Specifies the destination node for the path |
| Administration | Specifies the desired administrative state |
| Template ID | Specifies the ID of the template to be applied |
| Path Type | Specifies the type of path (must be Segment Routing) |
| Objective | Specifies the primary goal when identifying path resources |
| Bandwidth (Mbps) | Specifies the bandwidth required for the LSP |
| Setup Priority | Specifies a diversity-grouped LSP's priority access to the shortest path. Value 0 is the highest priority. |
| Assign BSID | Specifies whether or not a BSID will be assigned to the LSP |

3

In the Constraints panel, configure the required parameters:

| Parameter | Description |
|---------------------|--|
| Max Hops (Span) | Specifies the maximum number of hops to consider |
| MSD | Specifies the maximum SID depth to consider |
| Max Cost | Specifies the maximum cost to consider |
| Max TE Metric | Specifies the maximum TE metric to consider |
| Max Latency | Specifies the maximum latency to consider |
| Include Any Bit Pos | Specifies any bit between 0 and 31 to include |
| Exclude Any Bit Pos | Specifies any bit between 0 and 31 to exclude |
| Include All Bit Pos | Specified all bits between 0 and 31 to include |

4

In the Association Groups panel click **+ ADD**. The Create Association Groups form opens. Configure the required parameters:

| Parameter | Description |
|-------------------------------|--|
| Association Type | Specifies the association type. PCE-initiated LSPs only support the 'Policy' type. |
| Association Group ID | Specifies the association group identifier. When the Association Type is set to 'Policy', the Association Group ID is used to perform a lookup against a Path Profile definition. An ID is limited to one per association type. If multiple association groups exist, the association with the lowest ID will be used. |
| Association Source IP Address | Specifies the IP address of the association source. Both IPv4 and IPv6 addresses are supported. |

Click **ADD**. The Create Association Group form closes.

5


In the Path Profiles panel, configure the required parameters:

| Parameter | Description |
|------------|--|
| Profile ID | Specifies the identifier of the path profile policy to apply |
| Group ID | Specifies the identifier of the path profile group to which this LSP belongs |

6

Click **CREATE**. The PCE-initiated LSP is created.

7

From the LSPs view, a variety of actions can be performed on individual PCE-initiated LSPs by clicking **More**  and choosing the desired action from the contextual menu. These actions include show on map, re-signal, optimize, edit, shutdown, no shutdown, delete, and configure path profile override. See [5.3 "How do I apply a path profile override?" \(p. 95\)](#) for more information.



Note: When an LSP is re-signaled, a search is conducted for the constrained shortest path, and the LSP is rerouted to a new, best path (if one is found). The best path is determined by the LSP's objective. When all LSPs are re-signaled, they are rerouted one by one. Therefore, each previous reroute affects the viability of future reroutes. LSPs can be re-signaled on-demand at any time, though NSP periodically re-signals LSPs at set intervals.




Note: When LSPs are optimized, a set of LSPs are analyzed simultaneously, and paths are routed with the goal of producing a network with minimal bandwidth utilization on each link. Because multiple LSPs are optimized at once, there is a significant computation cost and impact to the network, however, the paths are routed with awareness and consideration of each other.

END OF STEPS

3.7 How do I create an LSP?

3.7.1 Steps

1

From the **Path Simulation, LSPs** view, click on **Create LSP** . The Create LSP form opens with the Identification panel displayed.

2

Configure the required parameters:

| Parameter | Description |
|------------------|--|
| Path Name | The name of the PCE-initiated LSP |
| Source | Specifies the source node for the path |
| Destination | Specifies the destination node for the path |
| Path Type | Specifies the type of path, SR or RSVP |
| Administration | Specifies the desired administrative state |
| Template ID | Specifies the ID of the template to be applied |
| Objective | Specifies the primary goal when identifying path resources |
| Bandwidth (Mbps) | Specifies the bandwidth required for the LSP |
| Setup Priority | Specifies a diversity-grouped LSP's priority access to the shortest path. Value 0 is the highest priority. |
| Assign BSID | Specifies whether or not a BSID will be assigned to the LSP |

3

In the Constraints panel, configure the required parameters:

| Parameter | Description |
|-----------------|--|
| Max Hops (Span) | Specifies the maximum number of hops to consider |
| MSD | Specifies the maximum SID depth to consider |
| Max Cost | Specifies the maximum cost to consider |

| Parameter | Description |
|---------------------|--|
| Max TE Metric | Specifies the maximum TE metric to consider |
| Max Latency | Specifies the maximum latency to consider |
| Include Any Bit Pos | Specifies any bit between 0 and 31 to include |
| Exclude Any Bit Pos | Specifies any bit between 0 and 31 to exclude |
| Include All Bit Pos | Specified all bits between 0 and 31 to include |

4

In the Association Groups panel click **+ ADD**. The Create Association Groups form opens. Configure the required parameters:

| Parameter | Description |
|-------------------------------|--|
| Association Type | Specifies the association type, Policy or Disjoint |
| Association Group ID | Specifies the association group identifier. When the Association Type is set to 'Policy', the Association Group ID is used to perform a lookup against a Path Profile definition. An ID is limited to one per association type. If this is exceeded, an error will occur and the LSP will become operationally down. |
| Association Source IP Address | Specifies the IP address of the association source. Both IPv4 and IPv6 addresses are supported. |
| Diversity | When the association type is 'Disjoint', specifies the diversity criteria |

i **Note:** The Association Group ID and Association Source IP Address parameters combine to form a unique key that is particularly useful for enforcing diversity between two nodes with same configured source.

Click **ADD**. The Create Association Group form closes.

5


In the Path Profiles panel, configure the required parameters:

| Parameter | Description |
|------------|--|
| Profile ID | Specifies the identifier of the path profile policy to apply |
| Group ID | Specifies the identifier of the path profile group to which this LSP belongs |

6

Click **CREATE**. The LSP is created.

7

From the LSPs view, a variety of actions can be performed on individual LSPs by clicking **More**  and choosing the desired action from the contextual menu. These actions include show on map, re-signal, optimize, edit, delete, and configure path profile override. See [5.3 "How do I apply a path profile override?" \(p. 95\)](#) for more information.



Note: When an LSP is re-signaled, a search is conducted for the constrained shortest path, and the LSP is rerouted to a new, best path (if one is found). The best path is determined by the LSP's objective. When all LSPs are re-signaled, they are rerouted one by one. Therefore, each previous reroute affects the viability of future reroutes. LSPs can be re-signaled on-demand at any time, though NSP periodically re-signals LSPs at set intervals.



Note: When LSPs are optimized, a set of LSPs are analyzed simultaneously, and paths are routed with the goal of producing a network with minimal bandwidth utilization on each link. Because multiple LSPs are optimized at once, there is a significant computation cost and impact to the network, however, the paths are routed with awareness and consideration of each other.

END OF STEPS

3.8 How do I collect statistics?

3.8.1 Purpose

NSP can collect LSP, LSP path, IP interface, and MPLS interface statistics. This is accomplished using GRPC from the node. NSP calculates a rolling average over a number of periods and uses data such as rising/falling thresholds and number of periods to determine when to trigger potential LSP reroute actions.

i **Note:** For RSVP and SR-TE LSPs that originate on 7250 IXR nodes, the NSP Flow Collector must be configured for statistics collection. Using the Flow Collector GUI, the Estimated Bitrate parameter must be set to a value of '5 min'.

i **Note:** When collecting LSP statistics on 7705 nodes, a 15 minute averaging period will be in effect. 7705 nodes only support the collection of RSVP LSP statistics. SR-TE LSP statistics collection is not supported.

3.8.2 Steps

1

Using the **PATCH /sdn/api/v4/nsp/configuration/traffic-data-collection** API call, set 'enabled' to true.

2

If collecting statistics on 7705 nodes - which do not support GRPC - you must perform one of more of the following to enable MIB entry policies, which will allow for statistics collection using SNMP:

a. To enable RSVP-TE LSP statistics collection on 7705 nodes, perform the following:

1. From the NFM-P GUI toolbar, click Tools and choose Statistics > MIB Policies > Default Stats Policy > Properties > MIB Entry Policies > Search from the contextual menu.
The NE MIB Statistics Policy form opens.
2. Click on the MIB Entry Policies tab and configure the filters as follows:
 - Product Name: 7705 SAR
 - MIB Entry Name: Lsp
3. Enable the MIB entry policies that are returned.

b. To enable IP Interface statistics collection on 7705 nodes, perform the following:

1. From the NFM-P GUI toolbar, click Tools and choose Statistics > MIB Policies > Default Stats Policy > Properties > MIB Entry Policies > Search from the contextual menu.
The NE MIB Statistics Policy form opens.
2. Click on the MIB Entry Policies tab and configure the filters as follows:
 - Product Name: 7705 SAR
 - MIB Entry Name: vrtrifstatsentry
3. Enable the MIB entry policies that are returned.

-
- c. To enable MPLS Interface statistics collection on 7705 nodes, perform the following:
1. From the NFM-P GUI toolbar, click Tools and choose Statistics > MIB Policies > Default Stats Policy > Properties > MIB Entry Policies > Search from the contextual menu.
The NE MIB Statistics Policy form opens.
 2. Click on the MIB Entry Policies tab and configure the filters as follows:
 - Product Name: 7705 SAR
 - MIB Entry Name: vrtrMplsIf
 3. Enable the MIB entry policies that are returned.



Note: Enabling MIB entry policies for other nodes - even those that support GRPC - will allow them to have their statistics collected using SNMP.

3

If collecting LSP statistics, perform the following:


1. Create a path profile policy, as described in [5.2 “How do I create a path profile policy?” \(p. 92\)](#), ensuring that 'Bandwidth Strategy' is set to telemetry.
2. Create or modify an LSP, ensuring that the path profile policy created in the previous step is assigned to the LSP. See [3.6 “How do I create PCE-initiated LSPs?” \(p. 55\)](#) for more information.

END OF STEPS

3.9 How do I view an LSP's computed path?

3.9.1 Purpose

When an LSP uses a node SID, NSP can display a computed path in order to visualize how the traffic is being routed.

 **Note:** Computed paths are only viable if an LSP has been replied to, or updated by NSP.

3.9.2 Steps

- 1 _____
From the **Path Control, Network Map** view, click on an LSP. The Info panel opens.
- 2 _____
Within the Info panel, enable the Computed Path radio button. The LSP's computed path is displayed on the map.

END OF STEPS _____


3.10 How do I view the LSPs associated with a path profile policy?

3.10.1 Purpose

You can view the LSPs that share a path profile policy and group, or just the path profile policy.


3.10.2 Steps

1

From either the **Path Control, LSPs** view or the **Path Simulation, LSPs** view, select an LSP and click . NSP displays a list of LSPs that share both the path profile policy and group of the selected LSP.

The path profile policy is displayed as Path Profile *profile_id/group_id*.

2

From either the **Path Control, Policies** view or the **Path Simulation, Policies** view, click **More** , **View associated LSPs** inline with a policy. NSP displays a list of LSPs that share this specific path profile policy.

END OF STEPS

3.11 How do I modify the demand bandwidth of an LSP?


3.11.1 Purpose

This procedure can be used to manually modify the demand bandwidth of one or more LSPs. An LSP's demand bandwidth is the aggregate bandwidth requested by its source and destination.

3.11.2 Steps



1

From the **Path Simulation, LSP Demand Matrix** view, a list of all demands associated with all existing LSPs is displayed.

To view a list of LSPs affected by a specific demand, click **More** , **Show in LSP List** inline with any demand. A list of the LSPs affected by the selected demand is displayed.

2

Return to the list of demands, then perform one of the following:

- a. To edit the bandwidth of a single demand, click **More** , **Edit Bandwidth** inline with any demand. The Edit Demand Bandwidth form opens.
- b. To edit the bandwidth of multiple demands, select multiple demands and click **Edit Bandwidth** . The Edit Demand Bandwidth form opens.

3

Enable the radio button next to one of the following parameters to provide a value:

| Parameter | Description |
|---------------|---|
| Set to (Mbps) | Specifies a new demand bandwidth value in Megabits per second. |
| Increase by % | Specifies an increase in demand bandwidth relative to the existing value. |
| Decrease by % | Specifies a decrease in demand bandwidth relative to the existing value. |

4

Click **SAVE**. The Edit Demand Bandwidth form closes.

5

To import LSP bandwidth, click **IMPORT LSP BANDWIDTH**. The Import LSP bandwidth form opens.

6

Configure the required parameters:

| Parameter | Description |
|---|--|
| Remote CAS Authentication Server IP Address | The IP address of the server from which to retrieve LSP bandwidth |
| Username | The username with which to log in to the server |
| Password | The password to be used with the specified username |
| LSP AGGREGATION PERIOD | The interval at which all LSP bandwidths will be updated with the most recently-available aggregations |

7

Click **IMPORT**. The Import LSP bandwidth form closes.



END OF STEPS

3.12 How do I highlight a potential path on the network map?

3.12.1 Purpose

This procedure can be used to highlight a potential path that meets defined criteria on the path control or path simulation network map.

3.12.2 Steps

- 1 _____
From either the **Path Control, Network map** view or the **Path Simulation, Network map** view, click on the **Highlight**  tab in the info panel.
- 2 _____
Click on the  button in the Path Finder section of the panel. The Path Finder form opens.
- 3 _____
Configure the parameters, as required:

| Parameter | Description |
|---------------------------|--|
| Name | Specifies the name of the potential path |
| Source | Specifies the source of the potential path |
| Destination | Specifies the destination of the potential path |
| Secondary Source | Specifies the secondary source of the potential path |
| Secondary Destination | Specifies the secondary destination of the potential path |
| Permitted IP Address Type | Specifies the IP address type(s) to be permitted by the potential path |
| Objective | Specifies the primary goal when identifying potential paths |
| Path Type | Specifies the type of potential path to find |
| Disjoint | Specifies the Disjoint mode to be used in potential path computation, if any |
| Bi-directional | Specifies whether or not the potential path will be bi-directional |
| Max Hops (Span) | Specifies the maximum number of hops to consider |
| Max Cost | Specifies the maximum cost to consider |

| Parameter | Description |
|----------------------------------|--|
| Max Latency (microseconds) | Specifies the maximum latency to consider |
| Latency Threshold (microseconds) | Specifies when to re-signal the potential path, if optimized on latency |
| Max TE Metric | Specifies the maximum TE metric to consider |
| MSD | Specifies the maximum SID depth to consider |
| Bandwidth (Mbps) | Specifies the maximum bandwidth to consider |
| Reverse Bandwidth (Mbps) | Specifies the maximum reverse bandwidth to consider |
| Include Any Bit Pos | Specifies any bit between 0 and 31 to include |
| Exclude Any Bit Pos | Specifies any bit between 0 and 31 to exclude |
| Include All Bit Pos | Specifies all bits between 0 and 31 to include |
| Profile ID | Specifies the Profile ID of the paths to be included in potential path computation |
| Algo Number | Specifies the Flex Algo Definition to be used in potential path computation |
| Explicit Route Strategy | Specifies the strategy to use when calculating the explicit route object (ERO) |
| Generation: Stop on first found | Only applicable when Explicit Route Strategy is set to Standard BSID Preferred or Loose Hop BSID Preferred. Specifies whether the BSID generation algorithm will stop after a result is found, or continue attempting additional strategies for a potentially more optimal result. |
| Generation: Run permutations | Only applicable when Explicit Route Strategy is set to Standard BSID Preferred or Loose Hop BSID Preferred. In a multi-area topology with an area splitting strategy, specifies whether BSIDs will be programmed in all feasible areas (false) or if the computation will attempt different permutations to achieve label stack reduction. |

| Parameter | Description |
|---|--|
| Generation: Emplacement Preference | Only applicable when Explicit Route Strategy is set to Standard BSID Preferred or Loose Hop BSID Preferred and when the Run Permutations parameter is set to true. Specifies where in the network-relative topology the BSID will be placed, Core or Edge. |
| Compressed Fallback IGP TE | Only applicable when Explicit Route Strategy is set to Compressed. If maximum stack depth is exceeded, specifies whether LSPs will instead adhere to IGP shortest path (with traffic engineering) to allow for compression. |
| Compressed Fallback IGP No TE | Only applicable when Explicit Route Strategy is set to Compressed. If maximum stack depth is exceeded, specifies whether LSPs will instead adhere to IGP shortest path (without traffic engineering) allow for compression. |
| Compressed Fallback Loose Hop | Only applicable when Explicit Route Strategy is configured with a value of Compressed. If maximum stack depth is exceeded, specifies whether LSPs will instead use only node SIDs in order to stay operationally up. This parameter is disabled for disjoint LSPs. |
| Explicit Route Strategy ECMP Preference | Only applicable when Explicit Route Strategy is configured with a value of ECMP. Specifies the type of SID that will be preferred when calculating the ERO. |

4

Click **FIND**. The potential path is highlighted on the network map.

END OF STEPS

3.13 How do I collect bit rate statistics using 7250 IXR?

3.13.1 Purpose

NSP can use Flow Collectors to collect and process 7250 IXR gen 1 bit rate telemetry over SR-TE and RSVP-TE LSP hops. Bit rate is calculated on LSPs for transit routers only at the next-hop ingress, and not on edge or head-end routers. The BR statistics are viewable in **Path Control, LSPs** in the Bandwidth column.

The following limitations affect this use-case:

- Flow collection must be configured and running in the network.
- Strict hop adjacency SID LSPs only (RRO must be only ADJ-SIDs). Supports SR-TE policies.
- Only Adjacency SIDs and Binding SIDs are supported in the LSP Path. The use of Node SIDs is not supported. The Explicit Route Strategy of Compressed can not be used for LSP Paths originating on a 7250 IXR gen 1 that require statistics collection.
- Parallel LSP paths cannot have their bandwidth determined, they must either be set to zero or split evenly.
- Due to the fact that interface stats are being used, the label stack cannot be sampled on the final hop or for single hop LSPs.
- The 7250 IXR does not support outputting egress MPLS Cflowd Stats from the head-end. Cflowd must be collected on the head-end neighbor ingress.
- IPFIX standard is limited to 10 labels, so ADJ-SIDs and service/entropy labels cannot exceed 10.
- Collection of flow records from multi-vendor nodes is not supported. It may be possible to collect statistics further downstream where there is a higher chance of not being able to distinguish parallel LSP Paths and incorrectly accounting for the proportion of traffic carried by each LSP Path.

3.13.2 Steps

1

For use-cases with an NSP-only deployment (no NFM-P), the following configuration must be present in the FCC config.yml file:

```
fcc:  
  nfmp:  
    enabled: false
```

2

Aggregations must to be enabled and configured from FCC WebUI:

1. Navigate to NSP Flow Collector→Aggregation Policy tab.
2. Set "Aggregation Intervals" to 5.
3. Navigate to NSP Flow Collector→Aggregation Policy tab→Estimated Bitrate tab.
4. Enable the "LSP Stack" parameter.

-
5. Set "Estimated Bitrate" to 5.

3

Provide 7250 IXR routers with the required configuration, as follows:

```
*A:SR_12# configure cflowd
*A:SR_12>config>cflowd# info
```

```
-----
active-timeout 1
inactive-timeout 59
use-vrtr-if-index
rate x
collector x.x.x.x version 10
    template-set mpls-ip
exit
-----
```

```
*A:SR_12>config>router# interface "abcd"
*A:SR_12>config>router>if# info
```

```
-----
cflowd-parameters
    sampling unicast type interface direction ingress-only
exit
no shutdown
-----
```

The following parameter values are important:

- **active-timeout:** the window of time used by the router to calculate flows (in minutes). In this example, with it set to 1, each minute the router will send its flow sample for that 1 minute to the collector. This value needs to be smaller or equal to the Estimated Bitrate set on the FC.
The values presented for active-timeout and inactive-timeout are suggested values, but they may need to be tuned based on an analysis of Cflowd statistics on the node.
- **sample-rate:** the rate at which the router samples flows. With a rate of 1, each sample is used and sent to the collector. With a rate of 100, 1 out of 100 samples will be used. When running lower amounts of traffic and flows, a smaller sample rate is ideal. With more traffic a higher sample rate is recommended to not use an excessive amount of processing power and affect routing.
- **use-vrtr-if-index:** this instructs Cflowd export to use the virtual IF index rather than the global IF index. This configuration is required.
- **rate:** determine an appropriate rate based on the flow rate supported by your platform of 7250 IXR gen 1 that can be obtained from Nokia Support.

See the nodal documentation for more information about these configuration parameters.

4

Configure Cflowd to use virtual interface index.

```
cflowd
use-vrtr-if-index
```

5

Perform the following configuration in the nrcp block of the /opt/nsp/configure/config/arm-system.conf file:

```
nrcp {
    telemetry {
        sr-te {
            stats_wipe_interval_in_seconds = 120
            stats_time_to_live_in_seconds = 600
        }
    }
}
```

END OF STEPS

4 NEs and links

4.1 What are BGP EPE links?

4.1.1 BGP EPE links

BGP Egress Peer Engineering (EPE) allows for the traffic-engineered creation of egress links from a source domain to an external domain. These links function as a BGP peer between two ASBR routers, which may be configured as E-BGP or I-BGP. When I-BGP is used, NSP does not track or know the underlying path the I-BGP peer session is representing and forwarding over. The BGP peer can also be associated with a label and used within a segment-routed path. When used in this way, the policy that controls the links that are used is defined at the source router as part of the label stack in the path, which can be computed by NSP. EPE links are represented as a dashed line within the path control network map, as they are virtual links that can span over multiple hops. The links are automatically collapsed for bi-directional representation if the reverse link direction is known.

All EPE links must contain a peer node SID for the remote peer, but may also contain an adjacency SID or peer SID set. Since NSP does not support SID sets, the priorities for specifying the hop in the ERO are:

1. Adjacency SID, if it exists
2. Peer node SID

If a peer node SID is not found, the link is considered invalid.

i **Note:** Due to limitations in current router implementations, NSP assumes that the EPE SID cannot be used as the top SID in a label stack.

4.1.2 Node configuration

The following is an example of the required node configuration:

```
A:s111_11_111_Both>config>router>bgp# info
```

```
-----  
family ipv4 bgp-ls  
min-route-advertisement 3  
link-state-import-enable  
egress-peer-engineering  
no shutdown  
exit  
group "abr"  
peer-as 400  
egress-engineering
```

```
no shutdown
exit
neighbor x.x.x.x
    egress-engineering
        no shutdown
    exit
exit
```

4.1.3 Bandwidth management

The bandwidth capacity of an EPE link is not advertised. If the bandwidth capacity of an EPE link can be known (for example, if it is provided in the BGP-LS advertisement or is NBI configured), NSP will manage and track bandwidth against it as it does other links, but it will not be mapped to any underlying IGP. If the bandwidth capacity of an EPE link cannot be known, NSP will not enforce bandwidth constraints on the link, nor will it perform any optimizations.

4.1.4 TE attributes

BGP-LS does not provide any TE attributes for EPE links. Users can provide these using the following REST API:

```
https://<NSP_cluster>:443/sdn/api/v4/nsp/net/l3/link/<Link_ID>
```

Where

NSP_cluster is the IP address of the NSP cluster

Link_ID is the ID of the EPE link

For more information, see the [Network Developer Portal](#).

4.2 How do I place a link set into maintenance mode?


4.2.1 Purpose

When a link set is placed into maintenance mode, the LSPs riding the link set must be rerouted. This can be done manually or automatically.

In order for the NSP host server to reroute these LSPs automatically, the `nrcp` block of the `/opt/nsp/configure/config/arm-system.conf` file must be modified as follows:

```
nrcp {  
  
nrcp_link_maintenance_policy="swift"  
  
bgpLs  
{  
    isTopoSourceBgpLS=true  
}  
}
```


When maintenance mode is deactivated for a link set, and the above modification has been made, the LSPs will automatically return to their original link set.

 **Note:** Link sets will be automatically placed into maintenance mode if either of their source nodes are placed into maintenance mode. This can be done using an API. For more information, see the [Network Developer Portal](#)

4.2.2 Steps

1 _____

Perform one of the following:


- a. From the **Path Control, Network Map** view, select an IGP link on the map and click **Info** .
- b. From the **Path Control, Links** view, select an IGP link from the list.

2 _____


Click **More** , **Activate Maintenance Mode for Link Set**. A confirmation window opens.

3 _____

Click **OK**. The confirmation window closes and the link set is placed into maintenance mode.

 **Note:** If the NSP host server has not been configured to automatically reroute LSPs whose link set has been placed into maintenance mode, these LSPs will need to be manually rerouted.

4

To deactivate maintenance mode for a link set, click **More**  , **Deactivate Maintenance Mode for Link Set**. A confirmation window opens.

5

Click **OK**. The confirmation window closes and maintenance mode is deactivated for the link set.



Note: If the NSP host server has not been configured to automatically reroute LSPs whose link set has been placed into maintenance mode, these LSPs will need to be manually returned to their original link set.

END OF STEPS

4.3 How do I turn down a link set?

4.3.1 Steps

1

Perform one of the following:

- a. To turn down a link set from the **Path Simulation, Network Map** view, perform the following:
 1. Select an IGP link on the map and click **Info** ⓘ.
 2. Click **More** ⋮ and choose **Turn Link Set Down**. The link set is turned down.
- b. To turn down a link set from the **Path Simulation, Links** view, perform the following:
 1. Select an IGP link from the list.
 2. Click **Turn Link Down** ↓ inline with the desired link. The link, and the other member of its set, is turned down.

2

Click **SIMULATE**. The Simulation Results form opens.

i **Note:** The **SIMULATE** button is only available when a traffic-affecting network enhancement or degradation has occurred. Some actions, such as Create LSP or Re-signal, are disabled when the **SIMULATE** button is available.

i **Note:** Simulation results are only available for the most recent simulated activity and will be overwritten when the **SIMULATE** button is next clicked.

3


Click **DETAILS** to view a list of LSPs affected by the configuration.

i **Note:** This list can be revisited by clicking on the Simulation Results tab. Only the results of the most recent configuration are displayed.

4

As required, perform one of the following:

- a. To return a link set to an operational state from the **Path Simulation, Network Map** view, perform the following:
 1. On the map, select an IGP link that is operationally down and click **Info** ⓘ.
 2. Click **More** ⋮ and choose **Turn Link Set Up**. The link set is turned down.
- b. To return a link set to an operational state from the **Path Simulation, Links** view, perform the following:
 1. Select an operationally down IGP link from the list.

-
2. Click **Turn Link Up**  inline with the desired link. The link, and the other member of its set, is turned up.


END OF STEPS

4.4 How do I create an IGP link?

4.4.1 Purpose

i **Note:** Nodes must have been imported or added before a link can be successfully added. See [2.5 “How do I import a network?” \(p. 38\)](#) or [4.7 “How do I add a node?” \(p. 84\)](#) for more information.

4.4.2 Steps

1 _____
From either the **Path Simulation, Network Map** view or the **Path Simulation, Links** view, click  and choose **Create Point-to-point Link**, then **Create IGP Link** from the contextual menu. A creation form appears.

2 _____
In the Identification panel, configure the required parameters:

| Parameter | Description |
|--------------|---|
| Admin Domain | Specifies the administrative domain in which to create the link |
| Protocol | Specifies the protocol to apply to the link |
| Source | Specifies the link source |
| Destination | Specifies the link destination |

3 _____
In the Forward Direction panel, configure the required parameters:

| Parameter | Description |
|---|---|
| IP Address | Specifies the IP address of the destination node (forward direction) or source node (reverse direction) |
| Cost | Specifies the maximum cost to consider |
| Latency | Specifies the maximum latency to consider |
| Traffic engineering parameters The below parameters are only configurable when the Enable Traffic Engineering slider is enabled | |
| Bandwidth (Mbps) | Specifies the maximum bandwidth to consider |


| Parameter | Description |
|-------------|---|
| Admin Group | Specifies the admin group to which the link will belong |
| SRLG | Specifies the SRLG to which the link will belong |
| TE Metric | Specifies the maximum TE metric to consider |
| RSVP | Specifies whether or not the link will be RSVP |
| SR | Specifies whether or not the link will be SR |

4

In the Reverse Direction panel, configure the parameters as described in step [Step 3](#).

5

Perform one of the following:

- a. Click on the **CREATE** button that features a refresh icon . The link is added to the network topology and the form remains open for further additions.
- b. Click **CREATE**. The link is added to the network topology.

END OF STEPS

4.5 How do I create a BGP link?

4.5.1 Purpose

i **Note:** Nodes must have been imported or added before a link can be successfully added. See [2.5 “How do I import a network?” \(p. 38\)](#) or [4.7 “How do I add a node?” \(p. 84\)](#) for more information.

4.5.2 Steps

1

From either the **Path Simulation, Network Map** view or the **Path Simulation, Links** view, click **+** and choose **Create Point-to-point Link**, then **Create BGP Link** from the contextual menu. A creation form appears.

2

Configure the parameters:

| Parameter | Description |
|-----------------------|---|
| Local BGP Router ID | Specifies the local BGP router ID |
| Local AS | Specifies the local autonomous system number |
| Local Peer Interface | Specifies the local peer interface |
| Remote BGP Router ID | Specifies the remote BGP router ID |
| Remote AS | Specifies the remote autonomous system number |
| Remote Peer Interface | Specifies the remote peer interface |
| Peer Node SID | Specifies the segment identifier representing the remote peer node |
| Peer Adjacency SID | Specifies the segment identifier representing the link to the remote peer |
| Link Bandwidth | Specifies the link bandwidth in Mbps |
| Link Metric | Specifies the numerical cost of the link |

3

Click **CREATE**. The link is added to the network topology.

END OF STEPS

4.6 How do I create a broadcast link?

4.6.1 Purpose

i **Note:** Nodes must have been imported or added before a link can be successfully added. See [2.5 “How do I import a network?” \(p. 38\)](#) or [4.7 “How do I add a node?” \(p. 84\)](#) for more information.

4.6.2 Steps

1 _____
From either the **Path Simulation, Network Map** view or the **Path Simulation, Links** view, click  and choose **Create Broadcast Link** from the contextual menu. A creation form appears.

2 _____
In the Identification panel, configure the required parameters:

| Parameter | Description |
|--------------|---|
| Admin Domain | Specifies the administrative domain in which to create the link |
| Protocol | Specifies the protocol to apply to the link |
| Source | Specifies the link source |
| New Subnet | Specifies whether or not a new subnet will be created |
| Destination | Specifies the link destination |


3 _____
In the Link Characteristics panel, configure the required parameters:

| Parameter | Description |
|---|---|
| IP Address | Specifies the IP address of the destination node (forward direction) or source node (reverse direction) |
| Cost | Specifies the maximum cost to consider |
| Latency | Specifies the maximum latency to consider |
| Traffic engineering parameters The below parameters are only configurable when the Enable Traffic Engineering slider is enabled | |
| Bandwidth (Mbps) | Specifies the maximum bandwidth to consider |

| Parameter | Description |
|-------------|---|
| Admin Group | Specifies the admin group to which the link will belong |
| SRLG | Specifies the SRLG to which the link will belong |
| TE Metric | Specifies the maximum TE metric to consider |
| RSVP | Specifies whether or not the link will be RSVP |
| SR | Specifies whether or not the link will be SR |

4

Perform one of the following:

- a. Click on the **CREATE** button that features a refresh icon . The link is added to the network topology and the form remains open for further additions.
- b. Click **CREATE**. The link is added to the network topology.

END OF STEPS


4.7 How do I add a node?

4.7.1 Purpose

i **Note:** An admin domain must have been imported or added before a node can be successfully added. See [2.5 “How do I import a network?” \(p. 38\)](#) or [2.6 “How do I add an admin domain?” \(p. 39\)](#) for more information.

4.7.2 Steps

1 _____

From either the **Path Simulation, Network Map** view or the **Path Simulation, Routers** view, click  and choose **Add Node**. The Create Node form appears.

2 _____

Configure the parameters:

| Parameter | Description |
|-----------|---|
| Router ID | Specifies the router identifier |
| Topology | Specifies the domain to which the router belongs |
| Protocol | Specifies the protocol to be used by the router; OSPF, ISIS, or BGP |

3 _____

Enable the desired checkboxes to specify the node's protocol. The options are:

- OSPF
- ISIS
- BGP

i **Note:** OSPF and ISIS can be enabled simultaneously.


4 _____

If the ISIS checkbox was enabled in [Step 3](#), perform one of the following:

- Add the node to an existing area(s).
 - Click **+ ADD**. A form opens with a list of existing areas displayed.
 - Select one or more areas from the list and click **DONE**. The form closes.
- Add the node to a new area(s).
 - Enable the **Create New Area(s)** checkbox.
 - Populate the **New Areas** field, clicking the **+** button after each required entry.

5

Perform one of the following:

- a. Click on the **CREATE** button that features a refresh icon . The node is added to the network topology and the form remains open for further additions.
- b. Click **CREATE**. The node is added to the network topology.



Note: Nodes can be deleted from the network entirely or removed on a per-instance basis.

END OF STEPS

4.8 How do I group NEs by region?

4.8.1 Purpose

This procedure can be used to group NEs by region on either the path control or path simulation network map.

4.8.2 Steps

1

From the NSP's Map Layouts and Groups views, add a region to a layout, and populate it with NEs. See the *NSP System Administrator Guide* for specific instructions.



Note: When regions are to be used, it is recommended that all NEs in the network are added to a region.

2

From the either the **Path Control, Network Map** view or the **Path Simulation, Network Map** view, click **Clustering Controls**  and enable Regions.



Note: When regions are enabled, Auto Layout functionality is disabled.

3

Click on a region to populate the Info panel with details. Double click to expand the region, displaying all nodes and associated subnets.



Note: Equally-weighted subnets may move between regions when the map is refreshed.

END OF STEPS

5 Policies

5.1 What is a path profile policy?

5.1.1 Path profile policies

A path profile policy is configured to provide a set of rules that determine how LSPs will be routed through the network. These rules ensure that the paths of the LSPs are optimized in order to meet specified criteria. Pairs of LSPs whose paths will have an affect on one another should have the same path profile policy applied to them, and be made part of the same path profile group. These include bi-directional LSPs (which follow the same path in opposing directions) and disjoint LSPs (which follow paths that must be diverse from one another).

i **Note:** LSPs can follow paths that are both bi-directional and disjoint from one another.

For information about creating a path profile policy, see [5.2 “How do I create a path profile policy?” \(p. 92\)](#).

For information about applying a path profile policy and a path profile group to an LSP, see [3.6 “How do I create PCE-initiated LSPs?” \(p. 55\)](#).

i **Note:** When an existing path profile policy is modified, the changes are not automatically applied to the LSPs using that policy. In order for the changes to be applied, those LSPs must first be re-signaled. By default, this occurs automatically every 30 minutes, though they can be re-signaled manually at any time.

i **Note:** Some third-party LSPs do not support the direct application of path profile policies. In this case, the user can force the LSP to inherit the set of rules configured in a path profile policy by applying a path profile override. For more information, see [5.3 “How do I apply a path profile override?” \(p. 95\)](#).

5.1.2 Path profile policy parameters

The following section describes some of the options available when configuring certain path profile policy parameters.

Bi-directional — this parameter specifies whether or not a pair of LSPs must follow the same path (in opposing directions). The options are:

- *No*: LSPs do not have to follow the same path
- *Symmetric Loose*: LSPs should follow the same path, unless impossible
- *Symmetric Strict*: LSPs must follow the same path or else fail

Disjoint — this parameter specifies whether or not a pair of LSPs must follow paths that are diverse from one another. The options are:

- *No*: LSPs do not have to follow paths that are diverse from one another

-
- *Node Strict*: LSPs must not traverse the same nodes, unless they are source or destination nodes
 - *Link Strict*: LSPs must not traverse the same links
 - *SRLG*: LSPs must not have overlapping SRLG values. LSPs may run over the same links, provided those links don't have SRLG values assigned to them.
 - *Node Strict and SRLG*: LSPs must not traverse the same nodes, unless they are source or destination nodes, and must not have overlapping SRLG values
 - *Link Strict and SRLG*: LSPs must not traverse the same links, and must not have overlapping SRLG values

Optimize On (Objective) — this parameter specifies the primary goal when identifying potential paths for LSPs to follow. The options are:

- *Hops*: LSPs should follow the path that requires traversing the fewest links
- *Cost*: LSPs should follow the path with the lowest sum, as determined by IGP Link metrics
- *TE Metric*: LSPs should follow the path with the lowest sum, as determined by TE metrics
- *Star Weight*: LSPs should follow the path with the least value, as determined by the STAR algorithm
- *Latency (microseconds)*: LSPs should follow the path with the lowest latency

Bandwidth Strategy — this parameter specifies the strategy that will be used to determine interface and LSP bandwidth. The options are:

- *Standard*: specifies that the LSP bandwidth used during path calculation for the LSP will be the requested bandwidth value configured for the LSP (LSP reservation bandwidth configuration)
- *Telemetry*: LSP bandwidth statistics are obtained by measuring dynamically. When selected, NSP will attempt to mitigate congestion by rerouting LSPs so as to avoid links that have exceeded their specified utilization threshold. NSP will still attempt to calculate a path that best meets an LSP's routing objectives while ensuring congested links are avoided.

Explicit Route Strategy — this parameter specifies the strategy to use when calculating the explicit route object (ERO). The options are:

- *Standard*: the ERO will contain end-to-end strict hops for the path
- *Standard BSID Preferred*: the ERO will contain end-to-end strict hops for the path, with binding SIDs preferred along the way to reduce label stack depth. If required, new binding SIDs will be generated and injected into the topology. The generation and compression of binding SIDs only occurs when necessary to avoid exceeding the MSD.
- *Loose Hop*: the ERO will contain the border routers through which traffic will be routed, which are identified as loose hops. As the exact path is generally not known (due to ECMP), NSP will not provide a computed path for loose hop-routed LSPs.
- *Loose Hop AnyCast Preferred*: the ERO will contain the border routers through which traffic will be routed, which are identified as loose hops. Manually-configured node SIDs will be preferred at parallel exit points to add resiliency. As the exact path is generally not known (due to ECMP), NSP will not provide a computed path for loose hop-routed LSPs.
- *Loose Hop BSID Preferred*: the ERO will contain loose hops along the borders of the path, with binding SIDs preferred along the way to reduce label stack depth. When a path crosses multiple

ASes or areas, NSP identifies the border router between the two areas as a loose hop and uses a Binding SID to specify the path through the next area.

- *Compressed*: the ERO will contain a combination of adjacency SIDs and node SIDs for the path to reduce label stack depth
- *ECMP*: the ERO will contain a combination of adjacency SIDs and node SIDs along a path comprised of parallel links and nodes. when this strategy is used, telemetry and bandwidth management are not available.

i **Note:** The traffic route between loose hops is not tracked as it is generally not known (due to ECMP) and is therefore not maintained by NSP. As a result, NSP will not use telemetry to keep track of utilization for loose hop-routed LSPs. It does, however, use telemetry to determine the utilization of all other LSPs that have been configured to use telemetry as their bandwidth strategy.

i **Note:** If the Explicit Route Strategy is set to Compressed, the label stack of the computed path will only be compressed if it exceeds the maximum stack depth that has been requested or configured.

BSID generation parameters — these parameters control BSID selection and/or BSID generation rules when computing a path result. The options are:

- *Generation: Stop on first found*: The BSID generation algorithm computes using a sequence of strategies. Specifies whether the calculation will stop after a result is found, or continue attempting additional strategies for a potentially more optimal result.
- *Generation: Run permutations*: A strategy may have various potential permutations in terms of where a BSID can be placed to achieve label stack reduction. For example, in a multi-area topology with an area splitting strategy, the algorithm may have an option to place a BSID in Area 1, or Area 0, or both. When the value of this parameter is set to false, the computation will program a BSID in all feasible areas. When the value is set to true, the computation will attempt different permutations and determine a subset result to achieve label stack reduction. The emplacement preference parameter setting will be used to assign preference for the final result, given multiple equal options.
- *Generation: Emplacement Preference*: This setting is applicable when the Run Permutations parameter is set to true. Given a potential combination of results, this option controls where in the network-relative topology the BSID will be placed. For example, the CORE option will prefer to place BSIDs in the center of the topology (such as Area 0), whereas the EDGE option will prefer to place BSIDs closer to the edge/destination of the path, such as a non-backbone area.

Compressed Fallback parameters — these parameters specify one or more actions that can be taken in order to keep an LSP with the Explicit Route Strategy of Compressed from going down when its label stack cannot be compressed below the configured maximum stack depth. When multiple options are enabled, they will each be attempted in sequence, until a suitable alternate path is found. The options are:

- *Compression Fallback IGP TE*: if maximum stack depth is exceeded, LSPs will instead adhere to IGP shortest path (with traffic engineering) to allow for compression. When enabled, this compression fallback action takes precedence over all others. It is enabled by default, and it is recommended that it remain enabled.
- *Compression Fallback IGP No TE*: if maximum stack depth is exceeded, LSPs will instead

adhere to IGP shortest path (without traffic engineering) allow for compression. When enabled, this compression fallback action takes precedence over all others, with the exception of Compression Fallback IGP TE. It is disabled by default, and it is recommended that it only be enabled if it is preferred that a tunnel remain operationally up rather than adhere to all configured traffic engineering requirements.

- *Compression Fallback Status Quo*: if the prior attempts fail to compute a compressed label stack path, the LSP will be permitted to adhere to its current path, if that path is deemed healthy. When enabled, this compression fallback action only takes precedence over Compression Fallback Loose Hop. It is enabled by default, and it is recommended that it remain enabled.
- *Compression Fallback Loose Hop*: if maximum stack depth is exceeded, LSPs will instead only use node SIDs in order to stay operationally up. This parameter is disabled for disjoint LSPs. When enabled, this compression fallback action is the last to be taken. It is disabled by default, and it is strongly recommended that it only be enabled if Compression Fallback Status Quo is also enabled.

i **Note:** If the Compressed Fallback Loose Hop parameter is used to find a path, all path tracking and traffic engineering will be disabled.

i **Note:** The last calculation behavior decision that was made can be viewed under the LSP details.

Explicit Route Strategy ECMP Preference — this parameter specifies the type of SID that will be preferred when calculating the ERO for an LSP that has been configured with the Explicit Route Strategy of ECMP. The options are:

- *Adjacency SID*: adjacency SIDs (including adjacency set SIDs) are preferred
- *Node SID*: node SIDs (including Anycast node SIDs) are preferred

Control Route Strategy — this parameter specifies whether an LSP is able to reroute, or must remain on its current path. The options are:

- *Standard*: LSPs are rerouted based on network changes or manual re-signaling
- *Loose*: LSPs are rerouted based on manual re-signaling only
- *Strict*: LSPs are not rerouted, regardless of network changes or manual re-signaling

SID Protection Strategy — this parameter specifies the extent of SID protection to be used when routing a path. The options are:

- *Standard*: LSPs prefer routes that include links with SID protection
- *Protected Only*: LSPs must take routes comprised exclusively of links with SID protection
- *Unprotected Only*: LSPs must take routes comprised exclusively of links without SID protection
- *Unprotected Preferred*: LSPs prefer routes that include links without SID protection

i **Note:** The SID Protection Strategy parameter is applied as a constraint when evaluating adjacency SID eligibility during path calculation. The 'Backup Flag' identified in the IGP information is evaluated to determine whether a SID is protected or not.


Latency Threshold — this parameter specifies when to re-signal an LSP that is optimized on latency. If the parameter is set to 0, indicating no change in latency, the LSP is automatically re-

signaled when its end-to-end latency (the sum of all link latencies along the path) increases. If the parameter is set to a value less than 0, this automatic re-signal does not occur. If the parameter is set to a value greater than 0, the LSP is re-signalized when its end-to-end latency is greater than the defined threshold. If a path cannot be found that is below the Latency Threshold, the LSP will not be brought down, but an alarm will be raised. The LSP is brought down when no path that satisfies the max latency constraint can be found. The default value is -1.

5.2 How do I create a path profile policy?

5.2.1 Steps

1

From either the **Path Control, Policies** view or the **Path Simulation, Policies** view, choose Path Profiles from the drop-down menu and click **Create Policy** . The Create Path Profile policy form opens.

2

Configure the required parameters:

| Parameter | Description |
|---------------------------------|--|
| Reserved Profile ID | Specifies whether the path profile policy will assume the name and role of the default path profile policy |
| Name | Specifies the name of the path profile policy |
| Profile ID | Specifies the Profile ID of the paths to be included in path computation |
| Bidirectional | Specifies the bidirectional mode to be used in path computation, if any |
| Disjoint | Specifies the Disjoint mode to be used in path computation, if any |
| Optimize on (Objective) | Specifies the primary goal when identifying paths for path computation |
| Bandwidth Strategy | Specifies the strategy to be used for bandwidth collection |
| Explicit Route Strategy | Specifies the explicit route strategy for the service |
| Generation: Stop on first found | Only applicable when Explicit Route Strategy is set to Standard BSID Preferred or Loose Hop BSID Preferred. Specifies whether the BSID generation algorithm will stop after a result is found, or continue attempting additional strategies for a potentially more optimal result. |

| Parameter | Description |
|------------------------------------|--|
| Generation: Run permutations | Only applicable when Explicit Route Strategy is set to Standard BSID Preferred or Loose Hop BSID Preferred. In a multi-area topology with an area splitting strategy, specifies whether BSIDs will be programmed in all feasible areas (false) or if the computation will attempt different permutations to achieve label stack reduction. |
| Generation: Emplacement Preference | Only applicable when Explicit Route Strategy is set to Standard BSID Preferred or Loose Hop BSID Preferred and when the Run Permutations parameter is set to true. Specifies where in the network-relative topology the BSID will be placed, Core or Edge. |
| Compressed Fallback IGP TE | Only applicable when Explicit Route Strategy is set to Compressed. If maximum stack depth is exceeded, specifies whether LSPs will instead adhere to IGP shortest path (with traffic engineering) to allow for compression. |
| Compressed Fallback IGP No TE | Only applicable when Explicit Route Strategy is set to Compressed. If maximum stack depth is exceeded, specifies whether LSPs will instead adhere to IGP shortest path (without traffic engineering) allow for compression. |
| Compressed Fallback Status Quo | Only applicable when Explicit Route Strategy is set to Compressed. If prior path calculations found potential paths with label stacks that could not be compressed below the configured maximum stack depth, specifies whether the LSP will be permitted to adhere to that path if it is otherwise deemed healthy. |
| Compressed Fallback Loose Hop | Only applicable when Explicit Route Strategy is configured with a value of Compressed. If maximum stack depth is exceeded, specifies whether LSPs will instead use only node SIDs in order to stay operationally up. This parameter is disabled for disjoint LSPs. |

| Parameter | Description |
|---|---|
| Explicit Route Strategy ECMP Preference | Only applicable when Explicit Route Strategy is configured with a value of ECMP. Specifies the type of SID that will be preferred when calculating the ERO. |
| Control Route Strategy | Specifies the strategy to be used when rerouting a path |
| SID Protection Strategy | Specifies the extent of SID protection to be used when routing a path |
| Max Hops (Span) | Specifies the maximum number of hops (nodes) to consider when performing path computation |
| Max Cost | Specifies the maximum sum, as determined by IGP link metric, to consider when performing path computation |
| Max TE Metric | Specifies the maximum sum, as determined by TE metric, to consider when performing path computation |
| Max Latency (microseconds) | Specifies the maximum latency to consider when performing path computation |
| Latency Threshold | Specifies when to re-signal an LSP that is optimized on latency |
| Description | Describes the path profile policy |

i **Note:** For more information about many of the above parameters, including their available options, refer to [5.1.2 “Path profile policy parameters” \(p. 87\)](#).

3

As required, Exclude Route Objects by adding the IP address(es) of the object(s) to be excluded.

4

As required, Include Route Objects by adding the IP address(es) of the object(s) to be included. You must also specify Hop Type.

5

Click **CREATE**. The Path Profile policy is created.


END OF STEPS

5.3 How do I apply a path profile override?

5.3.1 Purpose

Some non-Nokia LSPs do not support the direct application of path profile policies. In this case, the user can force the LSP to inherit the set of rules configured in a path profile policy by applying a path profile override, as described in this procedure.

5.3.2 Steps

- 1 _____
Navigate to either the **Path Control, LSPs** view or the **Path Simulation, LSPs** view.
- 2 _____
Click **More** , **Configure Path Profile Override** in-line with any third-party LSP. The Configure Path Profile Override form opens.
- 3 _____
Configure the required parameters:

| Parameter | Description |
|-----------------|---|
| Profile ID | Specifies the ID of the path profile policy to apply |
| Profile Ext. ID | Specifies the extended path profile policy ID, for association/grouping |
- 4 _____
Click SAVE. The path profile override is applied.

END OF STEPS _____


5.4 How do I create a router ID mapping policy?

5.4.1 Purpose

Use this procedure to create a router ID mapping policy. NSP is able to discover and display multiple IGP instances (OSPF and ISIS), which are each discovered as a unique domain. These domains are interconnected on the same routers, which themselves have multiple instances defined. If the Router IDs for these instances are the same, they will be displayed as a single router on service management's multi-domain topology maps. If the Router IDs are different, a Router ID Mapping policy must be provisioned in order for the instances to be displayed as a single router on service management's multi-domain topology maps.

5.4.2 Steps

1

From either the **Path Control, Policies** view or the **Path Simulation, Policies** view, choose Router ID Mapping from the drop-down menu and click **Create Policy** . The Create Router ID Mapping Policy form opens.

2

Configure the required parameters:

| Parameter | Description |
|-------------------|---|
| Name | Specifies the name of the Router ID Mapping template |
| System IP Address | Specifies the system IP address of the router |
| System Name | Specifies the router system name |
| PCC Address | Specifies the address of the PCC associated with the router |
| Description | Specifies the router description |
| Router Info | Click ADD to add as many Router Info entries, as required. For each Router Info entry, you must specify the following information: <ul style="list-style-type: none">• Network Identifier• AS Number• BGP-LS ID (topology identifier)• Router ID• Protocol (the protocol that the IGP router is using) |

3

Click **CREATE**. The Router ID Mapping policy is created.

END OF STEPS

5.5 How do I modify the system IP MPLS configuration policy?

5.5.1 Purpose



Use this procedure to modify the IP/MPLS Configuration policy. The IP/MPLS Configuration policy allows you to configure the maintenance mode of an IP link. You can choose one of the following maintenance modes:

- **Manual**
You must manually resignal LSPs, trigger GCO, or wait for network changes to move resources from the affected link and back onto the best path.
- **Automatic**
NSP automatically moves LSPs from the affected resources, usually using re-signalling.

5.5.2 Steps

1

Perform one of the following:

- From the **Path Control, Policies** view, choose System IP MPLS Configuration from the drop-down menu and click **Create Policy** . The Edit System IP MPLS Configuration form opens.
- From the **Path Control, Policies** view, click **Edit**  in-line with the System IP MPLS Configuration policy. The Edit System IP MPLS Configuration form opens.

2

Configure the required parameters:

| Parameter | Description |
|------------------|---|
| Description | Describes the System IP MPLS Configuration policy |
| Maintenance Mode | Specifies whether or not links are placed into maintenance mode automatically or manually |

3

Click **SAVE**. The System IP MPLS Configuration policy is modified.

END OF STEPS

5.6 What are SR policies?

5.6.1 SR policies

An SR policy serves as instructions to route packets through the network on a specified path. These instructions are embedded at the headend. The other nodes remain stateless. The headend of an SR policy binds a segment identifier (called a Binding SID - or BSID) to its policy. When the headend receives a packet whose active segment matches the BSID of a local SR policy, it steers the packet into the associated SR policy. This can be used to reduce the maximum stack depth (MSD) of an SR-TE LSP. NSP distributes SR policies to nodes using BGP exclusively.



Note: MSD discovery is supported from BGP-LS for nodes not using PCEP.



Note: When distributing SR policies, NSP does not receive feedback from routers.

5.7 How do I create an SR policy?

5.7.1 Purpose

SR policies can be manually created, or automatically generated. If, for example, the creation of a new LSP (with an Explicit Route Strategy of 'Standard BSID Preferred') causes the maximum stack depth (MSD) to be exceeded due to the number of strict hops through the network, NSP checks if an existing SR policy can be reused. If the existing policies are not sufficient, NSP will automatically generate a new SR policy using a backwards-recursive algorithm. SR Policies are preferred at domain boundaries.

This procedure is used to manually create an SR policy.

i **Note:** Before creating an SR policy, the BGP of all eligible nodes must be configured to support sr-policy-ipv4 in family.

i **Note:** Before creating an SR policy, open the sros-vm.conf file and ensure that it reads as follows:

```
sros-vm {  
    vms =[  
        {  
            rom=true  
        }  
    ]  
}
```

SR policies are imported into path simulation when importing a network from a file, or if 'Network and LSPs' is selected when importing a live network.

i **Note:** Impacted SR policies are not included in the results of simulations or worst case failure scenarios.

5.7.2 Steps

1

Navigate to either the **Path Control, SR Policies** view or the **Path Simulation, SR Policies** view. A list of existing SR policies is displayed.

2

Click on the  button. The Create SR policy form opens.

3

Configure the Policy parameters, as required:

| Parameter | Description |
|-------------|-------------------------------------|
| Policy Name | Specifies the name of the SR policy |

| Parameter | Description |
|----------------|---|
| Description | Describes the SR policy |
| Color | Specifies the color to be associated with the SR policy. This parameter is required. |
| Headend | Specifies the IP address of a known router for the SR policy headend. This parameter is required. |
| Endpoint | Specifies the IP address of the SR policy endpoint. This parameter is required. |
| Administration | Specifies the administrative state of the SR policy |

4

Configure the Candidate Path parameters, as required:

| Parameter | Description |
|---------------------|--|
| Candidate Path Name | Specifies the name of the candidate path |
| Distinguisher | Specifies the unique distinguisher in the context of BGP, when combined with endpoint and color. This parameter is required. |
| Description | Describes the candidate path |
| Administration | Specifies the administrative state of the candidate path |
| Type | Specifies the type of the candidate path; Dynamic or Static. This parameter is required. |
| Binding SID | Specifies the binding SID of the candidate path. This parameter is required. |
| Preference | Specifies preference of the candidate path when selecting the best candidate path for the SR policy. This parameter is required. |
| Bandwidth (Mbps) | Specifies the bandwidth capacity, in Mbps, of the candidate path (static only) |

i **Note:** If not manually specified, binding SIDs will only be generated if the label stack depth of the computed TE path exceeds the maximum stack depth requested or configured.

i **Note:** The following REST API can be used to enable binding SID generation, and to configure binding SID and color range:
PATCH /sdn/api/v4/nsp/configuration/sr-policy-config
Where *NSP_cluster* is the IP address of the NSP cluster.
If the Binding SID parameter is manually configured with a value of -1, the binding SID will inherit configuration from the above REST API.
For more information, see the [Network Developer Portal](#).

i **Note:** When updating an SR policy, if a candidate path that is administratively 'Up' is modified - but remains administratively 'Up' - no changes to that candidate path will be processed.

5 Perform one of the following:

- a. If the Type parameter was set to Dynamic in [Step 4](#), continue to [Step 6](#).
- b. If the Type parameter was set to Static in [Step 4](#), go to [Step 7](#).

6 Perform one of the following:

- a. Using the Profile ID parameter, specify the identifier of a path profile policy to associate with the SR policy. The SR policy will automatically inherit the path profile policy's predefined constraints.
- b. Manually configure the constraint parameters, as required:

| Parameter | Description |
|-------------------------------|---|
| Max Cost | Specifies the maximum cost to consider |
| Max Latency (microseconds) | Specifies the maximum latency, in microseconds, to consider |
| Max Hops (Span) | Specifies the maximum number of hops to consider |
| MSD | Specifies the maximum SID depth to consider |
| Max TE Metric | Specifies the maximum TE metric to consider |
| Include Any Bit Pos[0,...,31] | Specifies any bit between 0 and 31 to include |
| Exclude Any Bit Pos[0,...,31] | Specifies any bit between 0 and 31 to exclude |
| Include All Bit Pos[0,...,31] | Specifies all bits between 0 and 31 for inclusion |

Go to [Step 10](#).

7

Click + **SEGMENT LIST** and configure the parameters:

| Parameter | Description |
|-------------------|---|
| Segment List Name | Specifies the name of the segment list |
| Weight | Specifies the segment list's weighted loadshare. Default weight is 1. |

8

Click + **SID** and provide a unique identifier. Repeat as required.

9

Repeat [Step 7](#) as required to create additional segment lists.

10

Repeat [Step 4](#) as required to create additional candidate paths.

11

Click **CREATE**. The SR policy is created.

END OF STEPS

