NOKIA

NSP
Network Services Platform
Release 22.6

IP/MPLS Optimization Application Help

3HE-18132-AAAB-TQZZA

Issue 1

June 2022

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.

Contents

1 IP/MPLS Optimization

1.1 Why use IP/MPLS Optimization?

1.1.1

The IP/MPLS Optimization application 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.

1.2 IP/MPLS Optimization API support

IP/MPLS Optimization 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 the IP/MPLS Optimization application, append **/sdn/doc** to the server URL. For example: https://<NSP_cluster>:8543**/sdn/doc**.

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

1.3 What does the IP/MPLS Optimization application do?

The IP/MPLS Optimization application 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 the Service Fulfillment application, 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.

IP/MPLS Optimization application is 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. The IP/MPLS Optimization application 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.

1.3.1 Multi-domain path computation

The IP/MPLS Optimization application 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. Any optimization limiting domain traversals is not considered. Both Segment-Routed TE paths and RSVP TE paths are supported and deployed. Existing constraints such as the Max SID label depth apply.

1.3.2 Northbound interface for topology retrieval

The IP/MPLS Optimization application 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 the NSP, including additional information (such as area number/level instance and node/link/prefix SIDs). Northbound applications and controllers, such as the Cross Domain Coordinator application, can also modify the following TE attributes: SRLG, TE metric, IGP metric, Latency, and admin group.

1.3.3 Unnumbered interface support

The IP/MPLS Optimization application 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 IP/MPLS Optimization application. 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.

1.4 How does the IP/MPLS Optimization application function as a PCE?

The IP/MPLS Optimization application is the 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 the IP/MPLS Optimization application. The IP/MPLS Optimization application 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 the IP/MPLS Optimization application 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.

Note: Only the VSR-NRC supports topology discovery for the IP/MPLS Optimization application. 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 the IP/MPLS Optimization application, see the *NSP Deployment and Installation Guide*.

1.5 What types of LSPs does the IP/MPLS Optimization application support?

1.5.1 PCC-initiated LSPs

The IP/MPLS Optimization application supports the creation of PCC-initiated Segment-Routed TE LSPs, as well as PCC-initiated RSVP LSPs. The Service Fulfillment application 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 the IP/MPLS Optimization application (PCE). The IP/MPLS Optimization application (PCE) computes the LSP path and sends the path to the PCC. The PCC then installs the computed path and informs the IP/MPLS Optimization application (PCE) that the path is ready. This is reported to the Service Fulfillment application, where the path is attached to the service.

1.5.2 PCE-initiated LSPs

The IP/MPLS Optimization application 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 the IP/MPLS Optimization application, 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 PCE-initiated LSP, 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

1.5.3 RSVP LSPs

Any PCC node intending to request a path computation from the IP/MPLS Optimization application 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 the IP/MPLS Optimization application 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. The IP/MPLS Optimization application 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.

The IP/MPLS Optimization application 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, the IP/MPLS Optimization application receives the updates via BGP-LS and update its TEDB.

For stateful operation, which allows the IP/MPLS Optimization application 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 the IP/MPLS Optimization application 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, the IP/MPLS Optimization application 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, the IP/MPLS Optimization application 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 the IP/MPLS Optimization application, and the IP/MPLS Optimization application 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 the IP/MPLS Optimization application for either active control or LSP optimization. This is known as active stateful behavior. The delegation is awarded using the PCE control option. Once the IP/MPLS Optimization application 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. The IP/MPLS Optimization application also reroutes LSPs automatically on resource failures, or when calculating disjoint paths.

Note: When the PCC has delegated control of the LSP to the IP/MPLS Optimization application, 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 the IP/MPLS Optimization application.

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 the IP/MPLS Optimization application, this results in a PCE update from the IP/MPLS Optimization application. The LSP may switch to the secondary path in the interim, but will switch back to the primary path as soon as possible.

The IP/MPLS Optimization application maintains the active path in case both the primary and secondary paths are signaled, and also when the primary path is down. The IP/MPLS Optimization application also maintains the shared explicit behavior when the primary and secondary paths share common link resources.

The IP/MPLS Optimization application 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 the IP/MPLS Optimization application, and therefore the IP/MPLS Optimization application is not aware of the detours and bypass. However, the types of node and/or link protection are communicated to the IP/MPLS Optimization application via the PCE report.

Note: All the RSVP-TE LSPs created by the the IP/MPLS Optimization application have FRR enabled by default. The FRR method used is “facility”.

RSVP LSP bandwidth management

The IP/MPLS Optimization application manages the LSP bandwidth consumption on the TE links for both stateless and stateful PCC configurations. In a stateless configuration, the IP/MPLS Optimization application 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 the IP/MPLS Optimization application. In a stateful case, wherein the reports are sent to the IP/MPLS Optimization application from the PCC, the bandwidth is again communicated by the PCC to the IP/MPLS Optimization application via the bandwidth object. Here, the IP/MPLS Optimization application will reconcile the TE update with the specific LSP bandwidth update via the report. Therefore, the IP/MPLS Optimization application 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 the IP/MPLS Optimization application, and therefore, will have no state on the IP/MPLS Optimization application. The IP/MPLS Optimization application will not show these LSP reservations on the TE links. For a mixture of LSPs that are PCE-reported and non-PCE-reported, the IP/MPLS Optimization application 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.

1.5.4 Segment-routed TE LSPs

Any PCC node intending to request a path computation from the IP/MPLS Optimization application 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 the IP/MPLS Optimization application 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. The IP/MPLS Optimization application 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, the IP/MPLS Optimization application 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 the IP/MPLS Optimization application. The path computed by the IP/MPLS Optimization application 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 the IP/MPLS Optimization application 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, the IP/MPLS Optimization application 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, the IP/MPLS Optimization application 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 the IP/MPLS Optimization application, and the IP/MPLS Optimization application 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 the IP/MPLS Optimization application for either active control or LSP optimization. This is known as active stateful behavior. The delegation is awarded using the PCE control option. Once the IP/MPLS Optimization application 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. the IP/MPLS Optimization application also reroutes LSPs automatically on resource failures, or when calculating disjoint paths.

Note: When the PCC has delegated control of the LSP to the IP/MPLS Optimization application, 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 the IP/MPLS Optimization application.

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 the IP/MPLS Optimization application if the LSP is configured to report bandwidth. Bandwidth tracking on the IP/MPLS Optimization application is done only after a valid PCE report message is generated by the PCC. The IP/MPLS Optimization application tracks the bandwidth reservation for SR TE LSPs separate from RSVP TE LSPs.

Note: A loose hop SR LSP whose bandwidth is specified and computed locally will not be tracked by the IP/MPLS Optimization application, even with the PCE report option enabled. The IP/MPLS Optimization application only tracks SR TE LSP paths computed by the IP/MPLS Optimization application 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 the IP/MPLS Optimization application. For an LSP with delegated control to the IP/MPLS Optimization application, the IP/MPLS Optimization application will react to the topology change and issue a new ERO update to the PCC via PCE update.

1.5.5 TE-ECMP routing

Traffic Engineered Equal Cost Multi-path routing (TE-ECMP) enables users to create multiple equal-cost paths that the IP/MPLS Optimization application 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, the IP/MPLS Optimization application identifies multiple equal-cost paths between a source and destination based on an objective metric, as known as an 'ECMP Tree'. The IP/MPLS Optimization application 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, the IP/MPLS Optimization application remain 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, the IP/MPLS Optimization application 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 it's 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, the IP/MPLS Optimization application 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”.

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.

1.5.6 Anycast and loopback for LSPs

The IP/MPLS Optimization application supports path computation requests that include anycast or loopback addresses as destinations.

When inter-domain with multiple instances on routers are supported, the IP/MPLS Optimization application 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.

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.

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.

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.

1.5.7 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 IP/MPLS Optimization application 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.

1.6 How does latency-based LSP rerouting work?

The IP/MPLS Optimization application can optimize and reroute LSPs based on network latency measurements collected by MDM from OAM TWAMP tests 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, the IP/MPLS Optimization application 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 trigger immediately 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 a 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)

The following configurations are required on all nodes that will use latency-based LSP rerouting functionality:

- A gRPC connection must be established between the routers and MDM
- A streaming policy must exist on the node
- The required TWAMP or EthCFM tests can be defined using either CLI, NFM-P, or the Insights Administrator application (or its APIs).

Note: Changes to latency can make it necessary for many LSPs to be rerouted at once. To reduce the load on the IP/MPLS Optimization application - 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.

1.6.1 TWAMP test definition

The IP/MPLS Optimization application can use TWAMP tests as its source of latency input. To assist with the definition of the required TWAMP tests using CLI, use the following as an example when configuring node 1:

```
A:nodeA_1>config>oam-pm# info

----------------------------------------------

session "<session_name>" test-family ip session-type proactive create

  ip

   dest-udp-port <reflector_port>

   destination <node2_ip>

   forwarding-interface "<to_interface>

   source <node1_ip>

   twamp-light test-id <id> create
```

```
        delay-template "<stream_name>"
        no shutdown
       exit
      exit
    exit
    streaming
      delay-template "<stream_name>" create
       fd-avg forward
       sample-window <stat_frequency>
       window-integrity 80
       no shutdown
      exit
    exit
    ----------------------------------------------
```

Use the following example when configuring node 2:

```
A:nodeA_2>config>router>twamp-light# info
----------------------------------------------
reflector udp-port <reflector_port> create
  prefix <node1_ip> create
  exit
  no shutdown
exit
----------------------------------------------
```

Where

session_name is a unique name for the TWAMP session

reflector_port is the port to be used on the reflector

node2_ip is the system or interface IP address of node or interface 2

to_interface is the outgoing interface to be used

node1_ip is the system or interface IP address of node or interface 1

id is a unique identifier assigned to the TWAMP test

stream_name is the name of the streaming policy

stat_frequency is the interval, between 10 and 60 seconds, at which statistics are retrieved

Note: Users can also define TWAMP tests using NFM-P, however, CLI must be used to instantiate the streaming policy.

1.6.2 EthCFM test definition

The IP/MPLS Optimization application can use EthCFM tests as its source of latency input. To assist with the definition of the required EthCFM tests using CLI, use the following as an example when configuring node 1:

```
(ro)[/configure eth-cfm]

A:admin@nodeA_1# info

----------------------------------------------

    domain "<domain_name>" {

        level 2

        dns "<domain_name>"

        md-index 1

        association "<domain_association>" {

            icc-based "<association_icc_name>"

            ma-index 1

            remote-mep 1000 {

           }

        }

    }

(ro)[/configure router "Base" interface "to_2"]

A:admin@nodeA_1# info
```

```
-------------------------------------------

    admin-state enable

    port 1/1/3

    eth-cfm {

        mep md-admin-name "<domain_name>" ma-admin-name "<domain_
association>" mep-id 1001 {

            admin-state enable

        }

    }

    ipv4 {

        primary {

            address <interface1_ip>

            prefix-length 24

        }

    }

(ro)[/configure oam-pm streaming delay-template "eth-cfm-stream-node-1"]

A:admin@nodeA_1# info

-------------------------------------------

    admin-state enable

    sample-window 10

    window-integrity 80

    fd-avg round-trip { }

(ro)[/configure oam-pm session "<session_name>"]

A:admin@nodeA_1# info

-------------------------------------------
```

```
        ethernet {
            dest-mac <mac_address>
            source {
                mep 1001
                md-admin-name "<domain_name>"
                ma-admin-name "<domain_association>"
            }
            dmm {
                admin-state enable
                test-id 12345
                delay-template "<template_name>"
            }
        }
-------------------------------------------
```

Use the following example when configuring node 2::

```
(ro)[/configure eth-cfm]
A:admin@nodeB_2# info
-------------------------------------------
    domain "<domain_name>" {
        level 2
        dns "<domain_name>"
        md-index 1
        association "<domain_association>" {
            icc-based "<association_icc_name>"
            ma-index 1
```

```
                remote-mep 1001 {

               }

            }

        }

(ro)[/configure router "Base" interface "to_1"]

A:admin@nodeB_2# info

----------------------------------------------

    admin-state enable

    port 1/1/3

    eth-cfm {

        mep md-admin-name "<domain_name>" ma-admin-name "<domain_
association>" mep-id 1000 {

            admin-state enable

        }

    }

    ipv4 {

        primary {

            address <interface2_ip>

            prefix-length 24

        }

    }

----------------------------------------------
```

Where

domain_name is a unique name for the EthCFM domain

domain_association is a unique name for the EthCFM domain assocation

association_icc_name is the maintenance association name in ICC-based format

interface1_ip is the IP address of interface 1

session_name is a unique name for the session

mac_address is the MAC address of the destination node or interface

template_name is the name of the delay template to be used

test_id is the identifier of the EthCFM test

interface2_ip is the IP address of interface 2

Note: Users can also define EthCFM tests using NFM-P, however, CLI must be used to instantiate the streaming policy.

1.6.3 Modes

Two modes are available for TWAMP or EthCFM test configuration, either of which can be turned on or off independently by setting them to 'true' or 'false' when using the following API command: **PATCH https://{{server}}:8543/sdn/api/v4/nsp/configuration/latency**.

The model-driven option collects latency via the MD OAM application. To do this, all actions to create, modify, or execute tests must be performed within the MD OAM application. The MD OAM application can discover tests from the node automatically. When results are available for a given test, the application obtains the results, adds all stored test information to the results, and publishes them to oam.test_execution kafka topic.

The classic option collects latency through MDM, leveraging user-configured test information provided via API. The source and destination of each test session to be collected must be provided manually in the twampTests or ethcfmTests section of the above API call, as follows:

```
{

"data": {

  "classic": true,

  "modelDriven": false,

  "<test_type>": [

   {

    "session": "<session_name>",

    "source": "<node1_ip>",

    "destination": "<node2_ip>"

   }

  ]
```

```
}

}
```

Where *test_type* is twampTests or ethcfmTests.

Tests can also be discovered using the MDM Datasync app. If a particular session appears in the MDOAM application, then the modelDriven latency collection mode is the preferred method of collecting statistics from that session for the following reasons:

- If modelDriven mode is enabled, the session does not need to be defined under twampTests
- If modelDriven mode is enabled along with classic, then double latency application will occur

Note: When using modelDriven mode to collect EthCFM latency telemetry, the session mappings must still be defined under ethcfmTests.

To avoid double latency application to the same link (the same statistics coming from both classic and modelDriven mode), delete the session using the following REST API:

DELETE https://<NSP_cluster>:8543/sdn/api/v4/nsp/configuration/latency/<session>

Where

NSP_cluster is the IP address of the NSP cluster

Session is the TWAMP or EthCFM session

For more information, see the Network Developer Portal.

This will delete the entry from the twampTests or ethcfmTests config, ensuring that classic mode will no longer listen to telemetry messages for that session.

1.6.4 Latency timeout

Latency timeout values can be specified during TWAMP or EthCFM test configuration when “timeout” is enabled using the following API command: **PATCH https://{{server}}:8543/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 IP/MPLS Optimization application. 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 the Fault Management application 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 IP/MPLS Optimization

application. 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 IP/MPLS Optimization application.

Note: Latency timeout functionality is exclusive to MDM-sourced OAM latency. If both classic and modelDriven mode are 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.

1.6.5 Streaming template configuration

The type of latency measurements NSP 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.

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 the IP/MPLS Optimization application, which will apply the forward latency to that link.

1.7 How can I view LSP path history?

Users of the IP/MPLS Optimization application 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 https://<NSP_cluster>:
8543/sdn/api/v4/nsp/configuration/nrcp-historical
```

```
GET https://<NSP_cluster>:
8543/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, and stored in, the NSP's historical application. 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.

1.8 How are IRO objects used for path calculation?

The IP/MPLS Optimization application supports the IRO object specification within a PCC request. The IP/MPLS Optimization application 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.

1.9 What algorithms does the IP/MPLS Optimization application use?

1.9.1 STAR algorithm

The IP/MPLS Optimization application 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 in the IP/MPLS Optimization application and requires setting the objective to use STAR WEIGHT. The path profile policy is specified with the LSP definition and is conveyed to the IP/MPLS Optimization application via a PCE request message.

1.9.2 Disjoint optimal path computation algorithm

The IP/MPLS Optimization application 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.

Note: The IP/MPLS Optimization application 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.

1.9.3 Global concurrent optimization algorithm

The IP/MPLS Optimization application 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 the IP/MPLS Optimization application 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 on from the Service Fulfillment application.

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

1.10 What are 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 network map of the IP/MPLS Optimization application, 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.

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.

1.10.1 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
```

1.10.2 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.

1.10.3 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>:8543/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.

1.11 Navigating the IP/MPLS Optimization application's dashboard

The IP/MPLS Optimization dashboard is the first page of the application that a user will see after logging in. The dashboard provides an at-a-glance summary of the state of the application and provides links to points of interest. After navigating away from the dashboard, the user can return at

any time by clicking **Dashboard** .

The dashboard is comprised of the following sections:

1.11.1 Network Summary

The Network Summary section of the dashboard is visible when Network Info & Statistics is selected from the drop-down menu, and 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 Node List page of the application.

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 Path List page of the application.

1.11.2 System status

The System status section of the dashboard is visible when Network Info & Statistics is selected from the drop-down menu, and 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 Link List page of the application.

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 Path List page of the application.

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 Link List page of the application.

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 Path List page of the application, 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.

1.11.3 System automation activities

The System automation activities section of the dashboard is visible when Network Info & Statistics is selected from the drop-down menu, and 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.

1.11.4 System Activity Logging

The System Activity Logging section of the dashboard is visible when System Activity Logging is selected from the drop-down menu.

This section 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.

1.12 How do I modify the IP/MPLS Optimization application's network map?

The IP/MPLS Optimization application's 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 IP/MPLS Optimization application's network map:

1.12.1 Refresh

Refresh is located on the network map page of the IP/MPLS Optimization application. When clicked, the latest version of the map is fetched.

1.12.2 Search In Map

Search In Map is located on the network map page of the IP/MPLS Optimization application. 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.

1.12.3 Export Network

Export Network can be selected from the contextual menu when **More** is clicked on the network map page of the IP/MPLS Optimization application. 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.

1.12.4 Save Map changes

Save Map changes can be selected from the contextual menu when **More** is clicked on the network map page of the IP/MPLS Optimization application. When the Group Manager application is part of the deployment, the IP/MPLS Optimization application fetches and utilizes the node positions from this application, 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 IP/MPLS Optimization application.

1.12.5 Restore Common Map Layout

Restore Common Map Layout can be selected from the contextual menu when **More** is clicked on the network map page of the IP/MPLS Optimization application. When the Group Manager application is part of the deployment, the IP/MPLS Optimization application fetches and utilizes the node positions from this application, but users can reposition nodes on the network map as required. Users can then revert to the map configuration reflected in the Group Manager application by selecting Restore Common Map Layout.

1.12.6 Rebuild

Rebuild can be selected from the contextual menu when **More** is clicked on the network map page of the IP/MPLS Optimization application and Map Actions is chosen. This options deletes, then rebuilds the network map to ensure synchronization with the IGP topology layer.

1.12.7 Auto Layout

Auto Layout can be selected from the contextual menu when **More** is clicked on the network map page of the IP/MPLS Optimization application 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.

1.12.8 Clean-up References

Clean-up References can be selected from the contextual menu when **More** is clicked on the network map page of the IP/MPLS Optimization application and Map Actions is chosen. When network elements go down, the IP/MPLS Optimization application 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.

1.12.9 Toggle Hidden Devices

Toggle Hidden Devices can be selected from the contextual menu when **More** ⋮ is clicked on the network map page of the IP/MPLS Optimization application 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.

1.13 What is a path profile policy?

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).

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

For information about creating a path profile policy, see 1.14 “How do I create a path profile policy?” (p. 33).

For information about applying a path profile policy and a path profile group to an LSP, see 1.15 “How do I create PCE-initiated LSPs?” (p. 36).

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.

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 1.17 “How do I apply a path profile override?” (p. 39).

1.13.1 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:* LSP bandwidth statistics are retrieved from BGP-LS discovered values
- *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.

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.

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.

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

Note: The last calculation behavior decision that was made can be viewed under the LSP details using the IP/MPLS Optimization application.

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 protectiong

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-signaled 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.

1.14 How do I create a path profile policy?

1

From the Policy List page of the application, 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

Parameter	Description
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.
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.

Parameter	Description
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.
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

Note: For more information about many of the above parameters, including their available options, refer to 1.13.1 “Path profile policy parameters” (p. 29).

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

1.15 How do I create PCE-initiated LSPs?

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.

1

From the Path List page of the application, 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.

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 Path List page, 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 1.17 “How do I apply a path profile override?” (p. 39) 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 the IP/MPLS Optimization application 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

1.16 What are 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, the IP/MPLS Optimization application 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 1.15 “How do I create PCE-initiated LSPs?” (p. 36)). 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 Policy List page of the IP/MPLS Optimization application.

1.17 How do I apply a path profile override?

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.

1

Navigate to the Path List page of the application.

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

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

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

1

Perform one of the following:

a. From the Network Map page of the application, select an IGP link on the map and click **Info** ⓘ .

b. From the Link List page of the application, 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

1.19 How do I create a router ID mapping policy?

Use this procedure to create a router ID mapping policy. The IP/MPLS Optimization application 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 the Service Fulfillment application'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 the Service Fulfillment application's multi-domain topology maps.

1

From the Policy List page of the application, 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: • 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

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

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

  The IP/MPLS Optimization application automatically moves LSPs from the affected resources, usually using re-signalling.

1

Perform one of the following:

a. From the Policy List page of the application, choose System IP MPLS Configuration from the drop-down menu and click **Create Policy** ⊕ . The Edit System IP MPLS Configuration form opens.

b. From the Policy List page of the application, 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

1.21 What is an SR policy?

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 SID (called a Binding segment, 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.

1.22 How do I create an SR policy?

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.

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

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

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

1

From the Path List page of the application, choose SR Policy from the drop-down. 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
Description	Describes the SR policy
Color	Specifies the color to be associated with the SR policy. This parameter is required.

Parameter	Description
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)

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.

Note: The following REST API can be used to enable binding SID generation, and to configure binding SID and color range:
`PATCH https://<NSP_cluster>:8543/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.

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

1.23 How do I display link utilization for IP and MPLS interfaces?

Configure statistics collection on NFM-P

1

From the NFM-P GUI toolbar, click Tools and choose TCA Policies from the contextual menu. The TCA Policies form opens.

2

Click Create. The TCA Policy (Create) form opens.

3

Configure the required parameters as described in the *NFM-P User Guide*.

4

Click Select next to the Monitored Object Type field. A form opens.

5

From the list, choose the Network Interface (Routing Management: General) monitored object. The Class Internal Name is rtr.NetworkInterface. Click OK. The form closes.

6

Click OK. The TCA Policy (Create) form closes.

7

Select the newly-created TCA Policy from the list. The TCA Policy (Edit) form opens.

8

On the General tab, expand the Custom panel and click Create. The Custom TCA (Create) form opens.

9

Click Select next to the Profile Id field. The Select custom profile TCA form opens.

10

Click Create. The Custom Profile TCA (Create) form opens.

11

Click on the Stats Type drop-down menu and choose IP Interface Stats (Routing Management: General).

12

Click Build Formula next to the Formula field. The Build Formula form opens.

13

Click on the Counter Type drop-down menu and choose txV4BytesPeriodic. Click Add.

14

Modify the formula to read as follows:

```
((txV4BytesPeriodic*8)/1024)/(
```

15

Click on the Counter Type drop-down menu and choose periodicTime. Click Add.

16

Modify the formula to read as follows:

```
((txV4BytesPeriodic*8)/1024)/(periodicTime/1000)
```

17

Click OK. The Build Formula form closes.

18

Click OK. The Custom Profile TCA (Create) form closes.

19

Select the newly-created custom TCA from the list. The Select custom profile TCA form closes and the Custom TCA (Create) form reappears.

20

With the Rules panel expanded, click Create. The Custom TCA Rule (Create) form opens.

21

Configure the desired Threshold value, then click on the Threshold Direction drop-down menu and choose Rising Above.

22

Disable the Alarm check box and click OK. The Custom TCA Rule (Create) form closes.

23

With the Rules panel expanded, click Create. The Custom TCA Rule (Create) form opens.

24

Configure the desired Threshold value, then click on the Threshold Direction drop-down menu and choose Falling Below.

25

Disable the Alarm check box and click OK. The Custom TCA Rule (Create) form closes.

26

Click OK. The Custom TCA (Create) form closes.

27

Click on the Monitored Object tab, then click Add. A form opens.

28

Click Search to populate the list, then select one or more objects and click OK. The form closes.

29

Select one or more policies from the list, then click Statistics Policies and choose MIB Entry Policy from the contextual menu. The Stats Classes form opens.

30

Choose IP Interface Stats (Routing Management General) from the list of Statistics Types and click OK. The Stats Classes form closes and the MIB Entry Policy (Edit) form opens.

31

On the General tab, click on the Polling Interval drop-down menu and choose 1 minute.

32

Click on the Administrative State drop-down menu and choose Up.

33

Click OK and confirm that you want to proceed. The MIB Entry Policy (Edit) form closes.

34

Click Apply and confirm that you want to proceed.

35

To verify that the interfaces are collecting statistics, select an object from the list and click Properties. The Network Interface (Edit) form opens.

36

Click on the Statistics tab.

37

Click on the Select Object Type drop-down menu and choose IP Interface Stats (Routing Management: General).

38

Click Search. If the list is populated, statistics are being collected. Close the form.

39

From the NFM-P GUI toolbar, click Application and choose Copy Property Form Identifier from the contextual menu.

END OF STEPS

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

This procedure is used to associate workflows from the Workflow Manager application with nodes, links, or LSPs in the IP/MPLS Optimization application.

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 IP/MPLS Optimization application.

Workflows must contain the following inputs:

- **payload** - includes maintenance mode, which is defined by the IP/MPLS Optimization application. The IP/MPLS Optimization application 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

From the Workflow Manager application, import or create a workflow to associate with a node, link, or LSP. See the Workflow Manager application help for more information.

2

From the IP/MPLS Optimization application, perform one of the following:

a. From the Network Map page of the application, select a node or link from the map, then click **More** ⋮ , **Show workflows**.

b. From the Router List page of the application, click on the Show workflows button in-line with any node.

c. From the Link List page of the application, click on the Show workflows button in-line with any link.

d. From the Path List page of the application, select LSP from the drop-down menu, then 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

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

This procedure can be used to enable automatic swtichover 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 alarms 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.

Create a Kafka Trigger

1

From the Workflow Manager application, choose Kafka Triggers from the drop down menu and 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 Workflow Manager application, choose Environment from the drop-down menu and 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 Workflow Manager application, choose Workflows from the drop-down menu and click **Import from Filesystem** .

b. From the Workflow Manager application, choose Workflows from the drop-down menu and click **Import from Git** .

c. From the Workflow Manager application, choose Workflows from the drop-down menu and click **CREATE WORKFLOW**.

7

See the Workflow Manager application help for more information about creating workflows.

END OF STEPS

1.26 How do I group NEs by region?

This procedure can be used to group NEs by region on the IP/MPLS Optimization application's network map.

1

In the Group Manager application, add a region to a layout, and populate it with NEs. See the Group Manager application help 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 Network Map page of the IP/MPLS Optimization application, 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

1.27 How do I collect statistics using MDM telemetry?

The IP/MPLS Optimization application can use MDM telemetry to collect LSP, LSP path, IP interface, and MPLS interface statistics. This is accomplished using GPRC 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.

Note: For RSVP and SR-TE LSPs that originate on 7250 IXR nodes, the NSP Flow Collector must be configured for statistics collection.

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.

1

Using the **PATCH <server>:8543/sdn/api/v4/nsp/configuration/traffic-data-collection** API call, set 'enabled' to true.

2

If collecting LSP statistics, perform the following:

1. Create a path profile policy, as described in 1.14 “How do I create a path profile policy?” (p. 33), 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.

End of steps

1.28 How can I view an LSP's computed path?

When an LSP uses a node SID, the IP/MPLS Optimization application 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 the IP/MPLS Optimization application.

1

From the Network Map page of the application, 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