



NSP

Network Services Platform

Release 21.3

IP/MPLS Optimization

Application Help

3HE-17268-AAAA-TQZZA

Issue 1

March 2021

Legal notice

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

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

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

Nokia reserves the right to make changes to this document without notice.

No part of this document may be copied, reproduced, modified or transmitted.

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

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

© 2021 Nokia.

Contents

1	IP/MPLS Optimization	5
1.1	Why use IP/MPLS Optimization?	5
1.2	IP/MPLS Optimization API support	5
1.3	What does the IP/MPLS Optimization application do?	5
1.4	How does the IP/MPLS Optimization application function as a PCE?	6
1.5	What types of LSPs does the IP/MPLS Optimization application support?	7
1.6	Latency-based LSP rerouting	13
1.7	How are IRO objects used for path calculation?	16
1.8	What algorithms does the IP/MPLS Optimization application use?	17
1.9	Navigating the IP/MPLS Optimization application's dashboard	18
1.10	How do I modify the IP/MPLS Optimization application's network map?	20
1.11	How do I create PCE-initiated LSPs?	21
1.12	How do I place a link set into maintenance mode?	23
1.13	How do I create a path profile policy?	24
1.14	How do I create a router ID mapping policy?	26
1.15	How do I modify the system IP MPLS configuration policy?	27
1.16	How do I display link utilization for IP and MPLS interfaces?	28
1.17	How do I associate a workflow with a node?	31
1.18	How do I enable automatic VSR-NRC site switchover?	32
1.19	How do I group NEs by region?	34
1.20	How do I collect statistics using MDM telemetry?	34
1.21	How can I view an LSP's computed path?	35

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 Nokia Network Developer Portal at <https://network.developer.nokia.com/api-documentation>.

For specific information about the REST APIs for the IP/MPLS Optimization application, append `/sdn/doc` to the server URL. For example: <https://server-ip:8543/sdn/doc>.

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.

i **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 to associate to the LSP path. 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.

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

The PCC can also delegate control of the LSP to 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.

i **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.

i **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.

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

The PCC can also delegate control of the LSP to 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.

i **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.

i **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 its unique Node SID instead, thereby bypassing the failed link.

Path persistency

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

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

For this reason, 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 is set to "Compressed".

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

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

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

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

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

1.6 Latency-based LSP rerouting

The IP/MPLS Optimization application can optimize and reroute LSPs based on network latency measurements provided by MDM OAM tests. TWAMP can be used to measure link latency. If the latency of the LSP increases beyond the threshold specified in the path profile, the IP/MPLS Optimization application can calculate a path with reduced latency and reroute the LSP. An LSP is rerouted when its end-to-end latency changes; changes to the latency of other links will not trigger a reroute. When an LSP is routed off its original path due to an increase in latency, the LSP will only return to its original path after a re-signal.

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 tests must be defined in oam-pm

1.6.1 TWAMP test definition

The IP/MPLS Optimization application uses TWAMP tests as its source of latency input. To define the required TWAMP tests using CLI, configure node 1 as follows:

```
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
```

Configure node 2 as follows:

```
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 2

to_interface is the outgoing interface to be used

node1_ip is the system or interface IP address of node 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

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

1.6.2 Modes

Two modes are available for TWAMP 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 TWAMP tests must be performed within the MD OAM application. 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`.

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

```
{
  "data": {
    "classic": true,
    "modelDriven": false,
    "twampTests": [
      {
        "session": "<session_name>",
        "source": "<node1_ip>",
        "destination": "<node2_ip>"
      }
    ]
  }
}
```

1.6.3 Path profile policy parameters

In order for an LSP to optimize its path and reroute based on latency, the LSP must be assigned a path profile which has latency configured as its objective and the Latency Threshold parameter configured to specify when to reroute the LSP. If Latency Threshold is set to 0, the LSP is automatically re-signaled when its end-to-end latency (the sum of all link latencies along the path) increases. If Latency Threshold is set to less than 0 (the default is minus -1), there is no re-signal when the latency increases. If Latency Threshold is set to 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.

1.7 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 CSPF 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 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.

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

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


1.8.2 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.


i **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.

1.8.3 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.

1.9 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:

1.9.1 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 LSP list** will take you to the LSP List page of the application.

1.9.2 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 LSP list** will take you to the LSP 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 LSP list** will take you to the LSP 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.

1.9.3 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:

i **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.

1.9.4 System Activity Log


The System Activity Log section of the dashboard is visible when System Activity Log 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.


1.10 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:


1.10.1 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.


1.10.2 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.


1.10.3 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.


1.10.4 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.


1.10.5 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.


1.10.6 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.


1.10.7 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.

1.10.8 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.

1.10.9 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.

1.11 How do I create PCE-initiated LSPs?

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

1 _____
From the LSP List page of the application, click on **Create PCE LSP** .

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

Parameter	Description
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

3


As required, click **ADVANCED OPTIONS**. Configure the parameters:

Parameter	Description
Max Hops (Span)	Specifies the maximum number of hops 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	Specifies all bits between 0 and 31 to include
Profile ID	Specifies the identifier of the path profile to apply
Group ID	Specifies the identifier of the group to which this LSP belongs
Setup Priority	Specifies a diversity-grouped LSP's priority access to the shortest path. Value 0 is the highest priority.

4

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

5

From the LSP 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, resignal, optimize, edit, shutdown, no shutdown, delete, and configure path profile override.

i **Note:** When an LSP is resignalled, 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 resignalled, they are rerouted one by one. Therefore, each previous reroute affects the viability of future reroutes. LSPs can be resignalled on-demand at any time, though the IP/MPLS Optimization application periodically resignals LSPs at set intervals.

i **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

1.12 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.


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**  and choose **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.
- i** **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**  and choose **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.
- i** **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

1.13 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	When this parameter is enabled, the Path Profile template assumes the Name and role of the default Path Profile template
Name	The name of the Path Profile template
Profile ID	The Profile ID of the paths to be included in path computation
Bidirectional	The bidirectional mode to be used in path computation
Disjoint	The Disjoint mode to be used in path computation

Parameter	Description
Optimize on (Objective)	Specifies the primary goal when identifying paths for path computation
Bandwidth Strategy	Specifies the strategy to use for LSP bandwidth in the path computation
Explicit Route Strategy	Specifies the explicit route strategy for the service
Explicit Route Strategy ECMP Preference	Only applicable when Explicit Route Strategy is configured with a value of ECMP. Specifies the preference of SID selection when calculating an ECMP ERO.
Control Route Strategy	Specifies the strategy to use when recomputing the path. The options are: <ul style="list-style-type: none"> • Standard - LSPs are rerouted based on network changes or manual re-signaling • Loose - LSPs are rerouted based on manual re-signalling only • Strict - LSPs are not rerouted based on network changes or manual re-signaling
SID Protection Strategy	Specifies the strategy to use when computing a path in regards to segment ID protection. The options are: <ul style="list-style-type: none"> • Standard (Protected Preferred) - permits any link, regardless of protection state, but prefers protected SIDs • Protected Only - only resources with SID protection are permitted during path calculation • Unprotected Only - only resources without SID protection are permitted during path calculation • Unprotected Preferred - will permit any link, regardless of protection state, but prefers unprotected SIDs
Max Hops (span)	Specifies the Max Hops constraint to be used in path computation
Max Cost	Specifies the Max Cost constraint to be used in path computation
Max TE Metric	Specifies the Max TE Metric constraint to be used in path computation

Parameter	Description
Max Latency	Specifies the maximum latency to consider
Latency Threshold	Specifies when to re-signal an LSP that is optimized on latency. If no change in latency (0), re-signal is automatically triggered. If the change is less than 0, the LSP is not re-signalled. If the change is greater than 0, the LSP is re-signalled. The default value is -1.
Description	Describes the Path Profile template

i **Note:** The SID Protection Strategy parameter is applied as a constraint when evaluating link 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.

- 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 _____

1.14 How do I create a router ID mapping policy?

- 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

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

3



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

END OF STEPS

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

1

Perform one of the following:

- 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.
- 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

Parameter	Description
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 _____

1.16 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 _____

1.17 How do I associate a workflow with a node?


This procedure is used to associate workflows from the Workflow Manager application with nodes 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 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'.

- 1 _____
From the Workflow Manager application, import or create a workflow to associate with a node. 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 from the map, click **More**  in the Info panel, and choose **Show workflows** from the contextual menu.
- b. From the Node List page of the application, click on the Show workflows button in-line with any node.

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

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

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

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

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

1.19 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.

i **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.

i **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.

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

END OF STEPS

1.20 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.

1

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

2

If collecting LSP statistics, perform the following:

1. Create an LSP path profile, as described in [1.13 "How do I create a path profile policy?"](#) (p. 24), ensuring that 'Bandwidth Strategy' is set to telemetry.

-
2. Create or modify an LSP, ensuring that the new LSP path profile is selected as this LSP's profile.

END OF STEPS

1.21 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
