

PCEP for SR P2MP Policy

draft-hsd-pce-sr-p2mp-policy-00

Authors:

Hooman Bidgoli, Nokia

Daniel Voyer, Bell Canada

Siva Sivabalan, Cisco

Jayant Kotalwar, Nokia

S. Rajarathinam, Nokia

IETF105, July 2019

Montreal

Presenter Hooman Bidgoli



I E T F[®]

Multicast Evolution

- There is a desire to simplify Next generation complex networks (i.e. 5G transport) from administration and protocol point of view.
- The controller provides an end to end view of the network and simplifies traffic engineering, slicing and monitoring of the end to end SLAs for each slice
- Protocols like SR simplify the underlay by removing the need of LDP/RSVP-TE protocols and use IGP/BGP to signal segments.
- **Multicast needs to follow suite**
- SR P2MP Policy removes legacy P2MP MPLS protocols like mLDP/RSVP-TE while providing traffic engineering via SR Policy attributes

SR P2MP Segment

- A Point-to-Multipoint (P2MP) segment connects a Root node to a set of Leaf nodes in segment routing domain.
- A P2MP segment contains Replication Segments, each providing forwarding instructions at Root, Transit Nodes and Leaf Nodes.
- PCC Initiated: Root and Leaves can be discovered via multicast procedures like NG-MVPN (RFC 6514, 6513) or PIM (Protocol Independent Multicast) on PCC and the relevant information send to the PCE
- PCE Initiated: Root and Leaves can be configure explicitly on the PCE or controller and programmed on the PCC

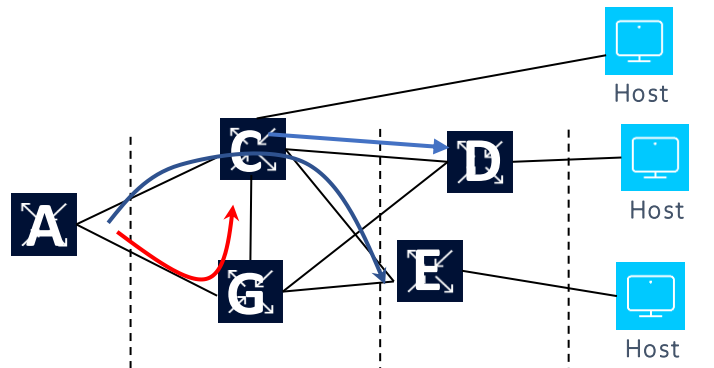
Replication Segment

- Is the forwarding instructions for the P2MP LSP
 - P2MP LSP redundancy concepts
 - Global optimization concepts
 - Label instructions
 - Fast Reroute instructions
- A Replication segment is defined via following
 - Node-ID: The node that the replication segment is for;
 - Root: The root of the P2MP segment that the replication segment is for;
 - Tree-ID: Tree that the replication segment is part of;
 - **Replication-SID: Segment ID for this Replication Segment.**
 - **Replicaition-SIDs can't be stacked as each replication segment can be a egress or transit.**
- Two Replication Segments can be connected directly via adjacent nodes or they can be non-adjacent and connected via a SID List (Unicast)

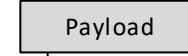
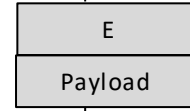
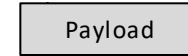
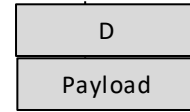
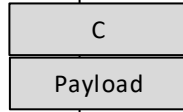
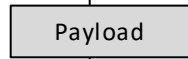
Packet Flow with FRR



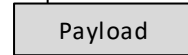
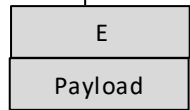
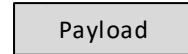
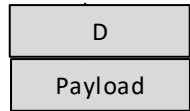
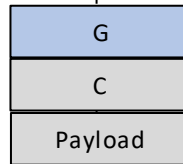
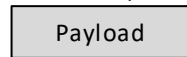
Source
(S1, G1)



Candidate Path 1
PLSP-1
LSP-ID 1

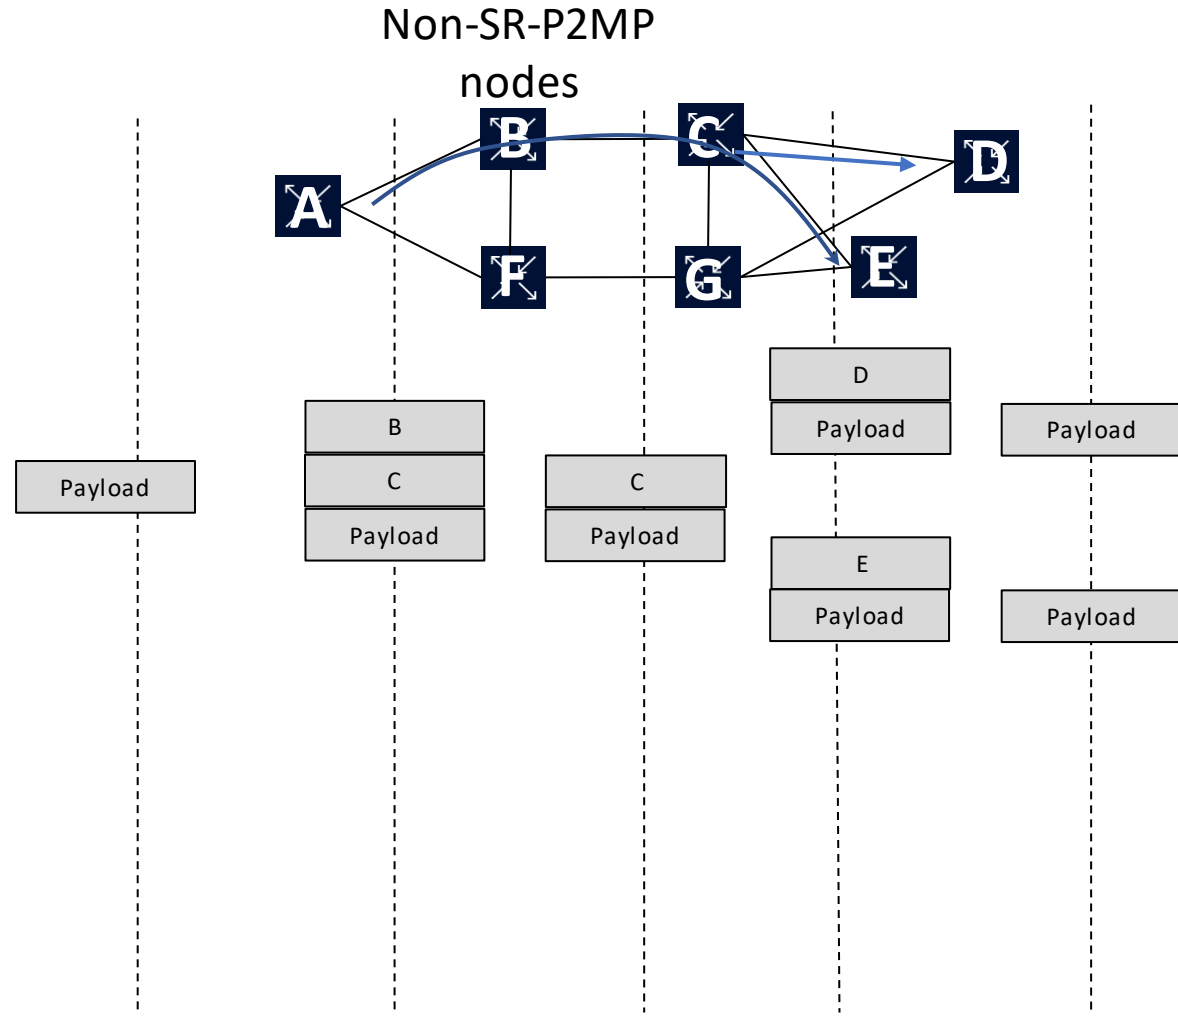


Facility Fast Reroute for
CP1
Link protect A-C
PLSP-1
LSP-ID 1



Packet Flow through a P2P SR Policy domain

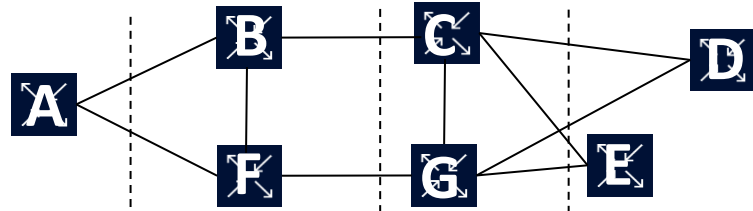
Candidate Path 1
PLSP-1
LSP-ID 1



Packet Flow LSP Red

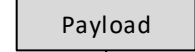
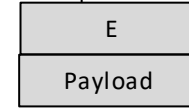
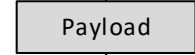
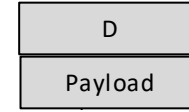
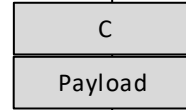
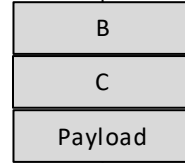
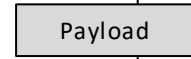
Non-SR-P2MP

nodes



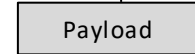
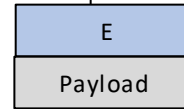
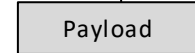
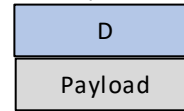
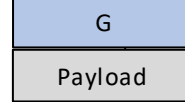
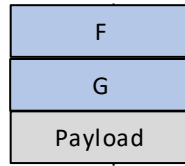
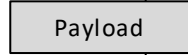
Candidate Path 1

PLSP-1
LSP-ID 1



Candidate path 2

PLSP-2
LSP-ID 1



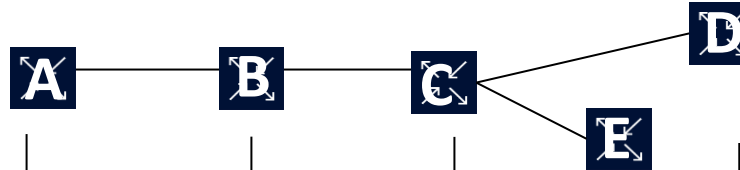
SR P2MP Objects

Non-SR-P2MP nodes



SR P2MP Policy

- ROOT Node
- Leaf Node
- Constrains
- Tree-ID



SR P2MP Policy

Replication Policy

- Node-ID
- Tree-ID
- Root

Unicast SR Policy

Replication Policy

- Node-ID
- Tree-ID
- Root

Replication Policy

- Node-ID
- Tree-ID
- Root

Candidate path 1

- Preference
- PLSP-ID = 1
- Incoming Label

Candidate path N

- Preference
- PLSP-ID = N
- Incoming Label

Candidate path 1

- Preference
- PLSP-ID = 1
- Incoming Label

Candidate path N

- Preference
- PLSP-ID = N
- Incoming Label

Candidate path 1

- Preference
- PLSP-ID + 1
- Incoming Label

Candidate path N

- Preference
- PLSP-ID + N
- Incoming Label

LSP-ID-1

LSP-ID-2

- Next-hop-group

LSP-ID-1

LSP-ID-2

- Next-hop-group

LSP-ID-1

LSP-ID-2

- Next-hop-group

LSP-ID-1

LSP-ID-2

- Next-hop-group

LSP-ID-1

LSP-ID-2

- Next-hop-group

LSP-ID-1

LSP-ID-2

- Next-hop-group

Forwarding Info

- Next-hop-group-id [nh-id] //array of nh
 - Next-hop-id <id>
 - Next-hop-add
 - Next-hop-int
 - Protect-nh <id>
 - P2mp-policy-label
 - Sid-list [list of outgoing labels]

SID, Forwarding instruction for this segment

P2MP LSP Redundancy

End to End Optimization

Forwarding info
Sid-List
Fast Reroute

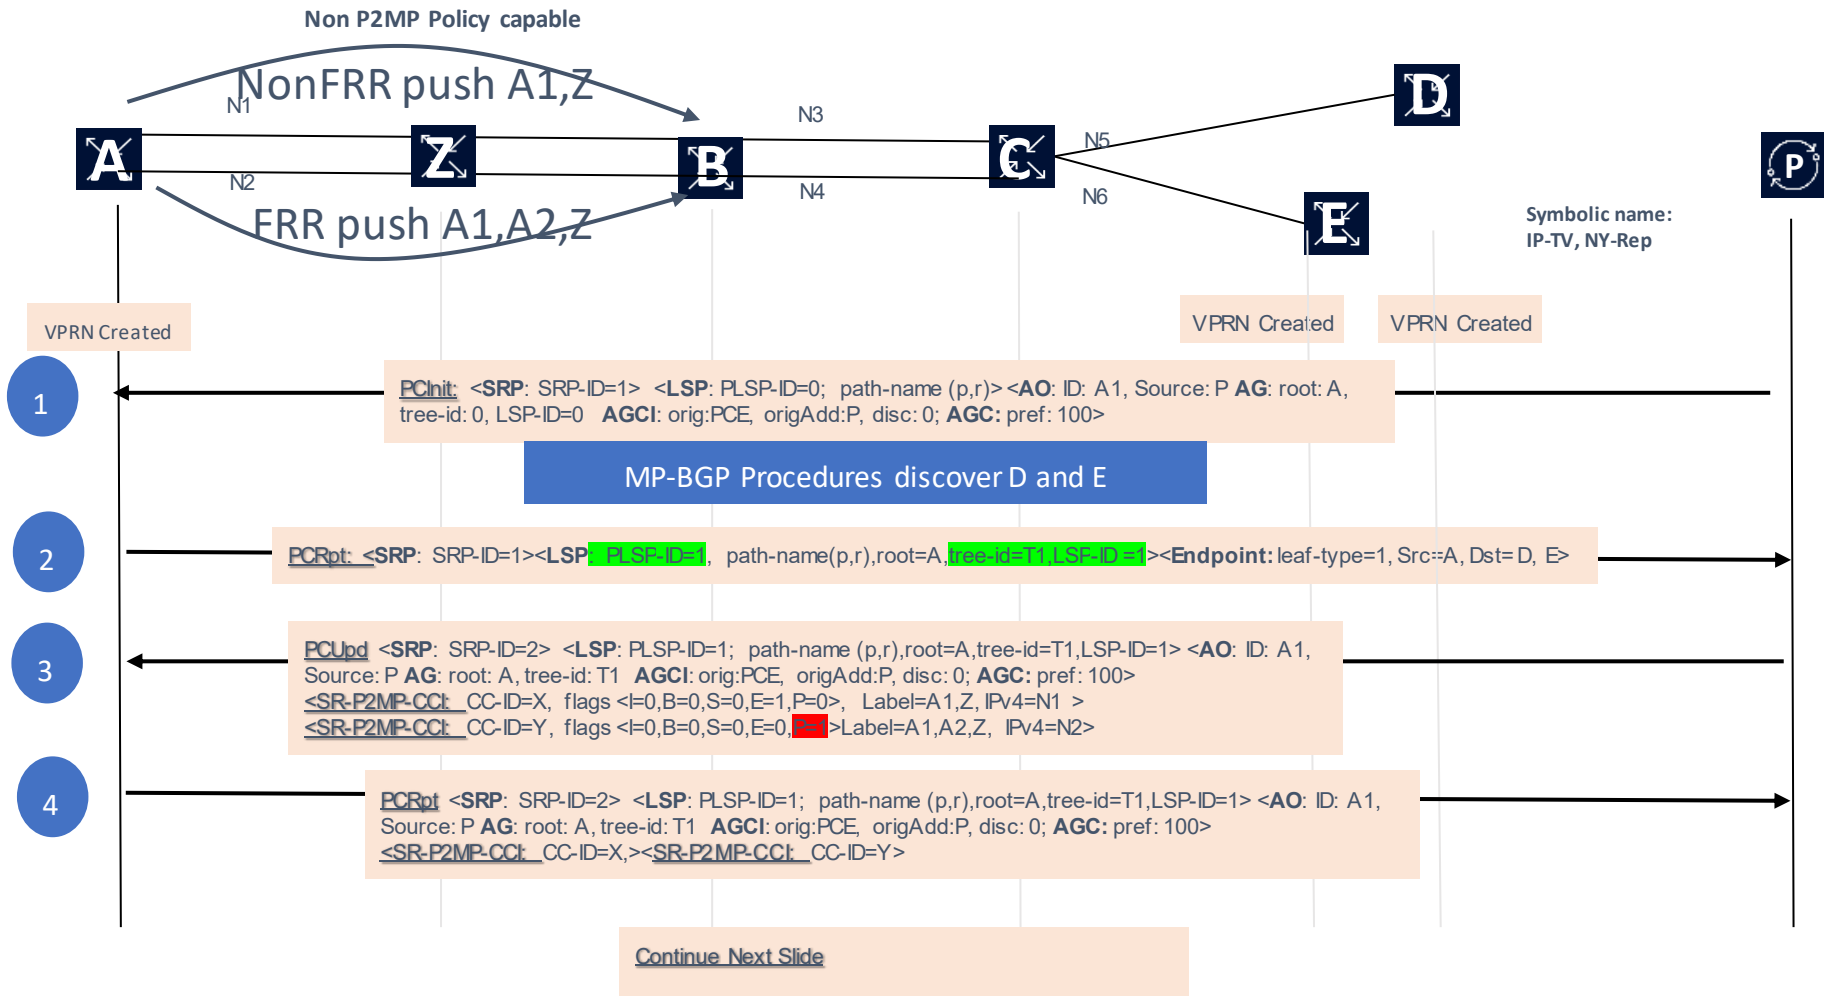
Drafts used to construct these objects

- ❑ RFC8236: Bases for stateful PCE
- ❑ RFC8236: P2MP Capabilities advertisement and a variation of the LSP Identifier
- ❑ RFC 8306: P2MP end point objects, used for PCC update to PCE
- ❑ RFC3209: LSP-ID for global optimization
- ❑ Draft-barth-pce-segment-routing-policy-cp: Candidate paths for P2MP LSP redundancy
- ❑ Draft-ietf-spring-segment-routing-policy: for segment-list to connect two non-adjacent replication policy
- ❑ Draft-ietf-pce-pcep-extensions-for-cpe-controller: Variant of CCI object used for downloading replication policy forwarding instructions.

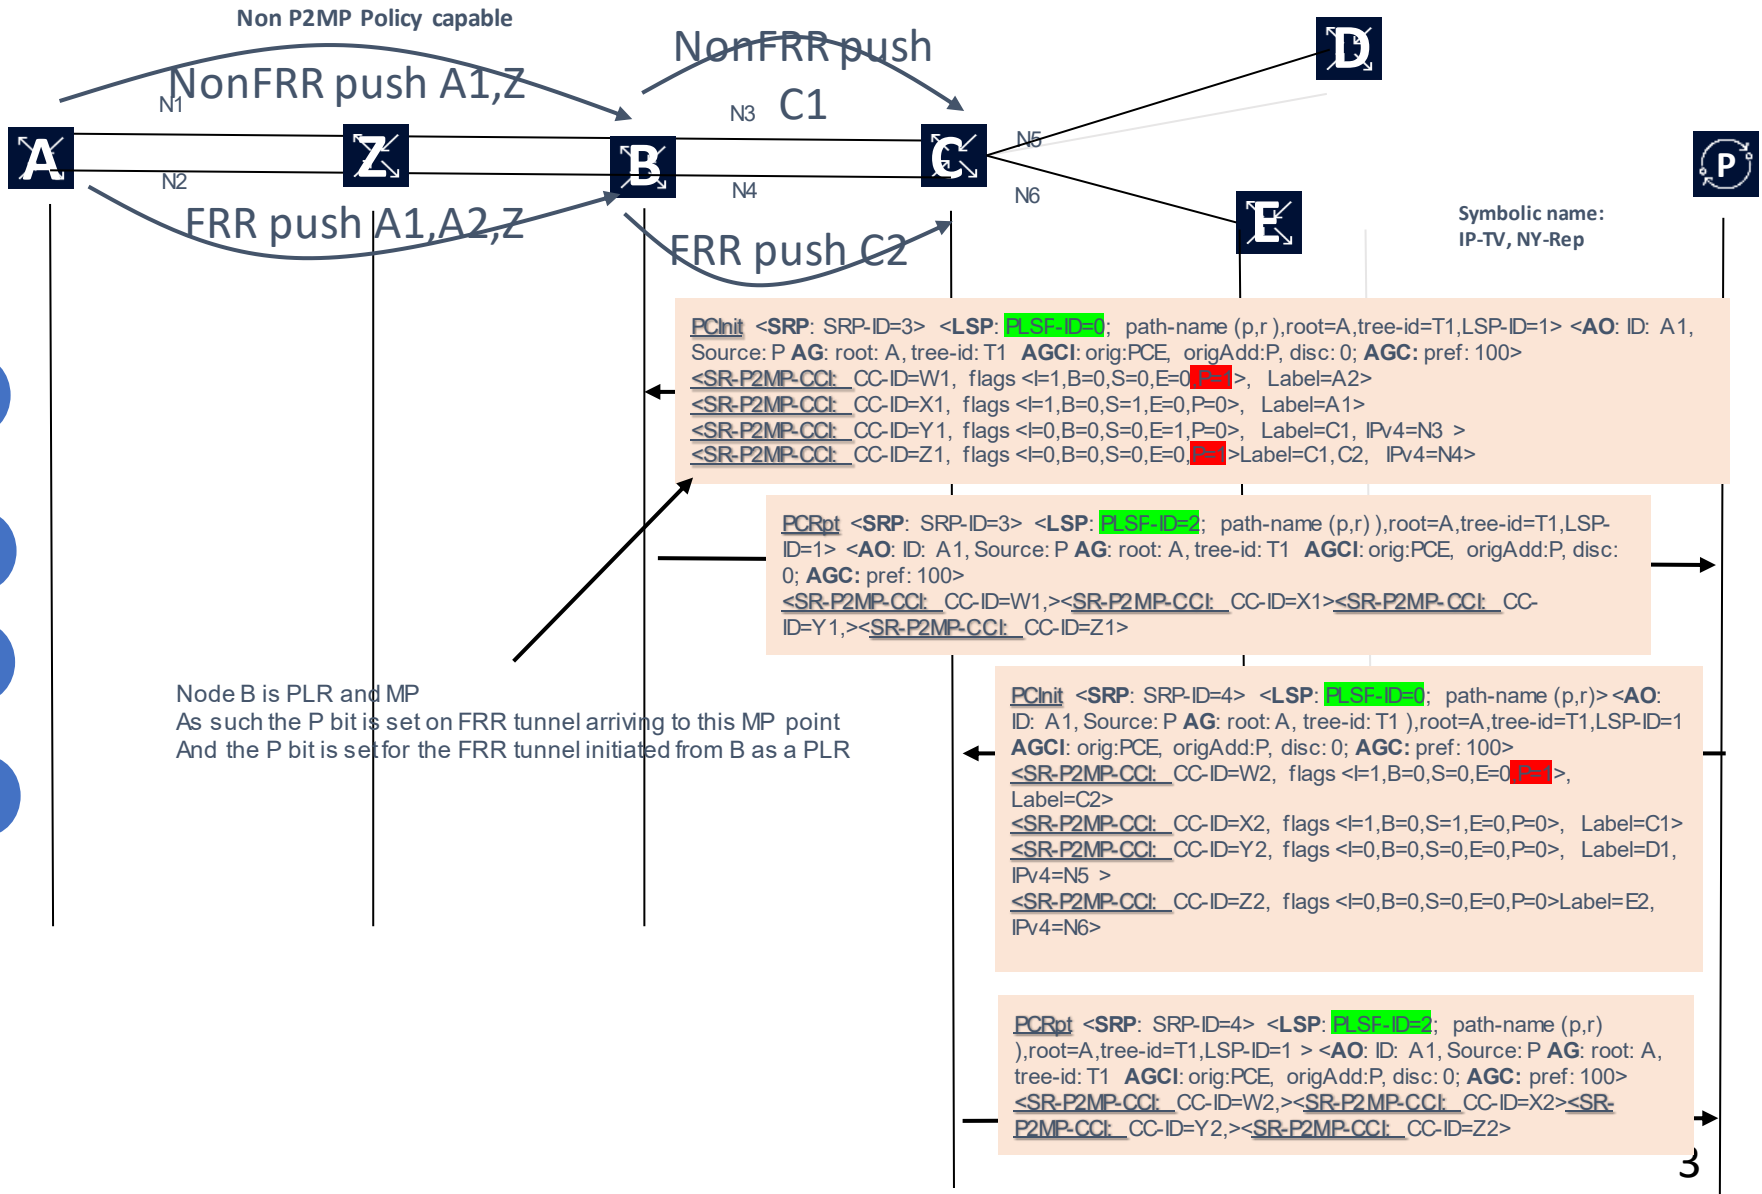
New TLVs and Objects defined

- ❑ SR-P2MP-CCI: for forwarding instructions, label actions, fast reroute, sid list
- ❑ New Association type (P2MP SR Policy): For Candidate paths
- ❑ New P2MP-LSP ID TLV: to identify the P2MP LSP <Root, Tree-ID, Lsp-ID>
- ❑ RFC3209: LSP-ID for global optimization

Example Work Flow



Example Work Flow (Cont)



5

6

7

8

Change: Authors and Contributors

Moved to contributor section:

Tanmoy Kundu Nokia

Stone Andrew Nokia

Next Steps