

Signaling RSVP-TE tunnels on a shared MPLS forwarding plane

draft-sitaraman-mpls-rsvp-shared-labels-00

Harish Sitaraman (hsitaraman@juniper.net)

Vishnu Pavan Beeram (vbeeram@juniper.net)

Tejal Parikh (tejal.parikh@verizon.com)

Motivation

- Couple the feature benefits of the RSVP-TE control plane with the simplicity of the SR MPLS forwarding plane.
- Prevent the platform specific label space limit on an LSR from being a constraint to pushing the limits of control plane scaling on that node.
 - [ietf-teas-rsvp-te-scaling-rec] proposes guidelines for scaling RSVP-TE control plane.

RSVP-TE + SR MPLS forwarding plane

- Shared forwarding plane
 - Transit label on a TE link is shared among RSVP-TE tunnels traversing the link.
- Forwarding plane state at an LSR
 - Reduced to roughly = number of RSVP neighbors
- Tunnel setup time
 - Forwarding plane state is not programmed during LSP setup and teardown.
- Share labels across MBB
 - Labels and routes do not change during make-before-break (MBB) on path overlap.
- Use all existing RSVP-TE control plane features
 - FRR, admission control, preemption, auto-bandwidth etc.

Allocation of pop labels

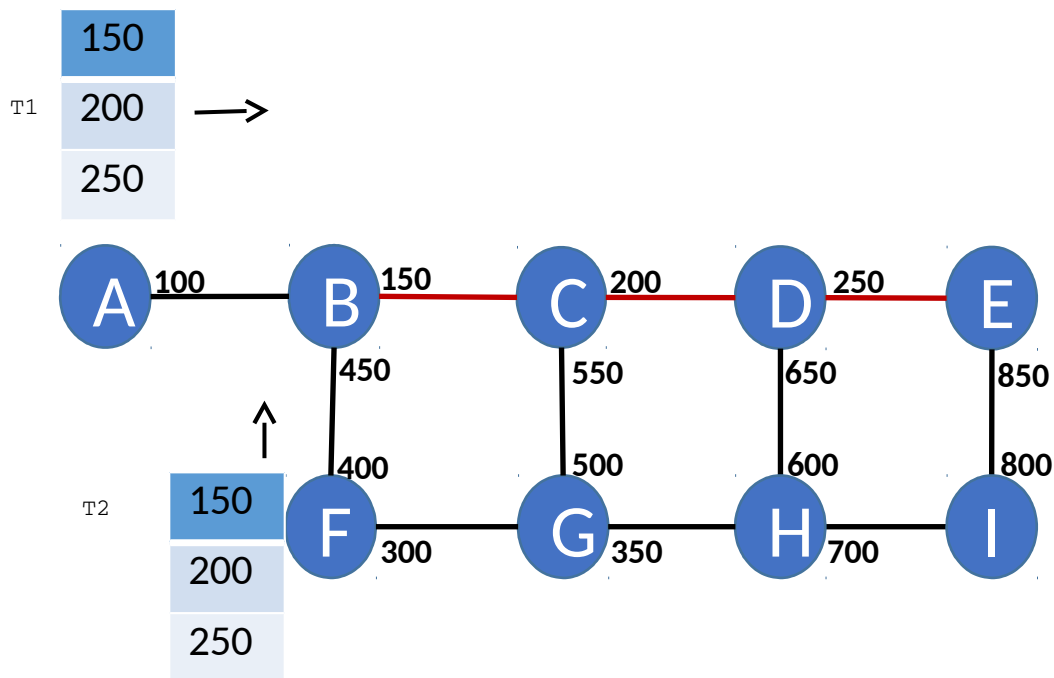
- Pop and forward label action
 - Existing operation
- LSR SHOULD allocate a unique pop label for each TE link
 - Install the label route
 - Label forwarding action: pop and forward to that neighbor
- Additional label per TE link necessary for facility backup link protection.

RSVP-TE pop and forward tunnel

RSVP-TE tunnel T1: From A to E on path A-B-C-D-E

RSVP-TE tunnel T2: From F to E on path F-B-C-D-E

} Shares all labels



LSP_ATTRIBUTES/LSP_REQUIRED_ATTRIBUTES in Path message

- Attributes Flags TLV: Pop Label

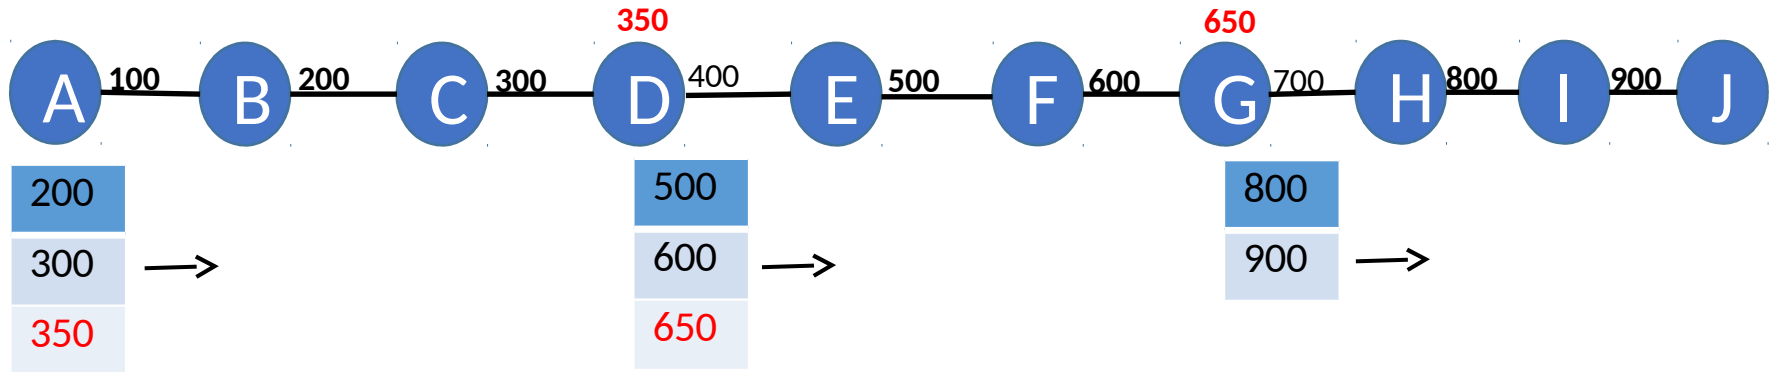
RRO Label Subobject Flag

- Pop Label

Delegation of label stack imposition

RSVP-TE pop and forward tunnel from A to J

Explicit/automatic selection of Label Stack delegation points: D and G




HOP_ATTRIBUTES subobject

- Attributes TLV: Label Stack Imposition TLV

Mixing pop and swap labels

RSVP-TE pop and forward tunnel from A to J

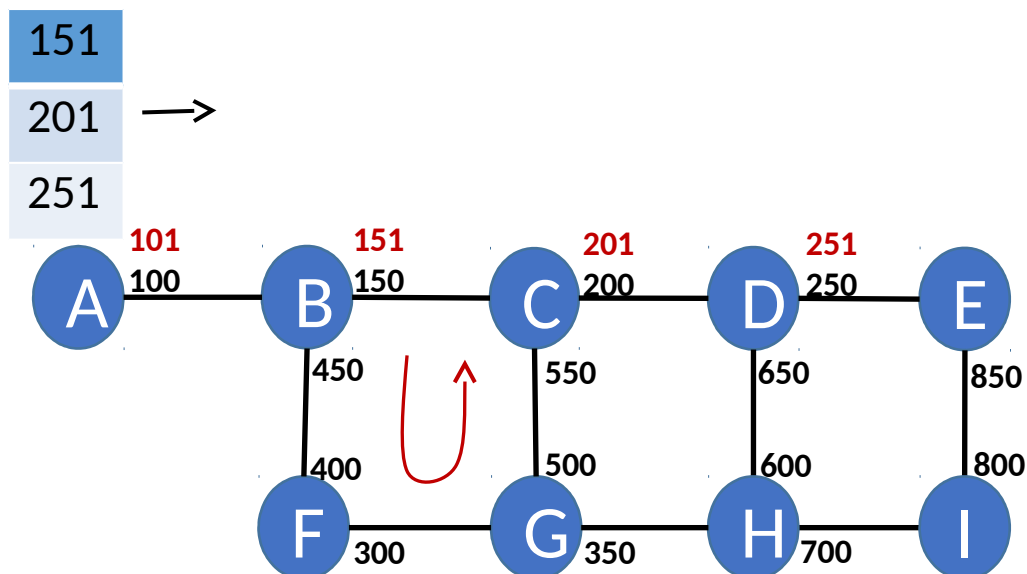
Hops D and G do not support allocation of pop labels



Local logic at ingress to construct label stack based on type of label received (in RRO) from each transit hop

Facility backup link protection

RSVP-TE pop and forward tunnel T1: From A to E on path A-B-C-D-E



At LSR B:

- Primary: Pop label 151 and forward via B-C link
- Backup: Pop label 151 and send via facility backup to C

Next Steps

- Request WG feedback
- Update the draft with feedback and experience from prototype
- Add sections for
 - Automatic delegation of label stack imposition
 - Node protection