

RSVP-TE P2MP Signaling Optimization for RMR

draft-zhang-mpls-rmr-rsvp-p2mp-00

MPLS WG

IETF102# Montreal

Zhaohui (Jeffrey) Zhang

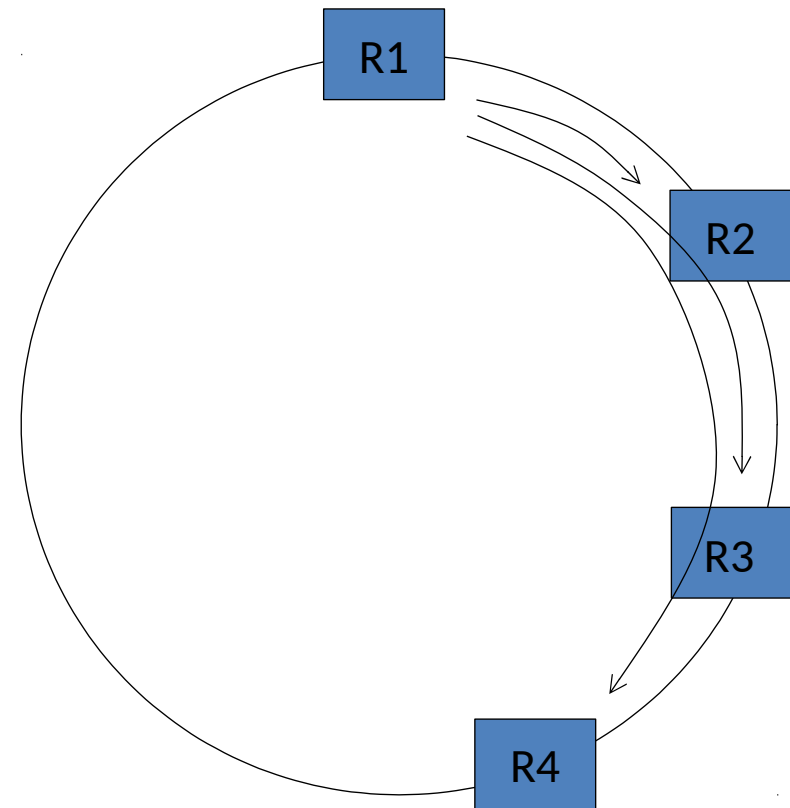
Abhishek Deshmukh

Ravi Singh

Traditional RSVP-TE P2MP Signaling

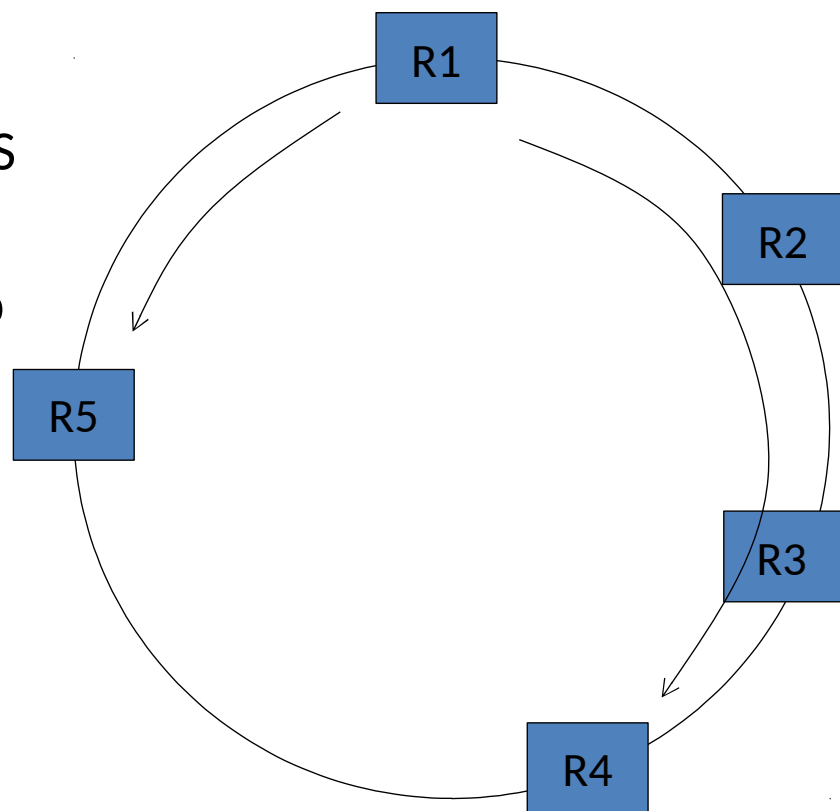
- One sub-LSP for each leaf
 - Lots of redundant PATH/RESV state near the ingress
 - Each leaf is explicitly listed
- Each sub-LSP optionally has its own Explicit PATH
- Extra state for tunnel protection

All these could be optimized away in case of RMR




Optimizations for RMR

- A single LSP
 - A single pair of PATH/RESV state on each node of the tunnel
 - Ingress could decide to use a single LSP in one direction for all leaves
- Or optionally two sub-LSPs in opposite directions
 - To reach different set of leaves
 - Not for protection purposes
- No explicit path needed
 - Just send along the ring in the specified direction




Optimizations for RMR

- Implicitly allowed leaves
 - PATH messages sent along the ring back to ingress
 - Ingress itself listed as a leaf
 - Leaves decide by themselves
 - Send RESV to PHOP
- Both explicitly listed leaves and implicitly allowed leaves are allowed
- Traffic stops at the last leaf in each direction
 - The last leaf does not have RESV state from downstream



Optimizations for RMR

- No additional signaling or state for protection
- Before global repair finishes after a failure:
 - Don't send RESV tear on failure
 - On link failure, PLR tunnels traffic to next node via a unicast ring LSP in the other direction
 - On node failure, PLR tunnel traffic s traffic to next next node
 - Traffic then continues from there on



Live-live Protection

- Live-live protection not needed in most situations
 - Traffic tunneled via Ring LSP upon failure
- Live-live protection can be easily achieved for mission-critical scenarios
 - If duplication removal is done by application
 - Just set up two opposite-direction sub-LSPs to reach all leaves and send traffic in both directions
 - Each leaf will deliver duplicate traffic (received in two directions) to application
 - No switchover upon failure detection; just global repair

MP2MP with RMR

- PATH message could carry a label used for downstream nodes to send traffic upstream
- Ingress node sends received upstream traffic downstream in the other direction
 - If two sub-LSPs in different directions are used

Related RSVP Objects

- RMR Object in PATH messages indicating RMR optimization is used:
 - Ring ID
 - Ring direction
- <S2L Sub-LSP Descriptor List> lists:
 - Explicit leaves
 - Ingress itself in case of implicit leaves
- PATH messages could carry a label object for MP2MP tunnels

Next Steps

- Seek comments
- Polish the document then request adoption
 - TEAS?