IP RSVP-TE: Extensions to RSVP for P2P IP-TE LSP Tunnels

<draft-saad-teas-rsvpte-ip-tunnels-00>

Tarek Saad, Juniper Networks

Vishnu Pavan Beeram, Juniper Networks

IETF-105, July 2019, Montreal

Outline

- Motivation
- IP RSVP-TE Tunnels
  - Creation and management
  - FRR protection
  - Shared forwarding state
- Signaling extensions
- Next steps

What problem are we solving?

- Ubiquitous deployments of IP forwarding networks
  - DC fabric networks
  - Cores for DCI connectivity
- Motivation for migrating to IPv4/v6 data plane
  - Reduced Capex/Opex – no need to worry about supporting multiple data encapsulation types
  - Focus on operational simplicity - adopting a new data plane technology like SRv6 is expensive
- Need for Traffic Engineering in native IP forwarding networks
  - Stringent SLAs for guaranteed network services (5G requirements)
  - Steering over non shortest paths: bandwidth, latency, disjointness, SRLG aware path(s), etc.
  - Tactical/Strategic network flow placement
  - Optimal use of network resources

IP RSVP-TE: what does it offer?

- Natively supports TE for IPv6 and IPv4 forwarding
- Forwarding state sharing for Multi-point-to-point (see further slides)

- Features borrowed from MPLS RSVP-TE
  - Auto-Bandwidth and container tunnel(s)
  - Per IP Path bandwidth reservations on shared resources
  - Soft and hard per IP Path priority-based preemption
  - Distributed or centralized (PCE) path computation for IP Path(s)
  - Make-before-break for IP Path changes
  - Failure protection support
    - Link/node+SRLG Fast-reroute bypass protection (auto-bypass)
    - End-to-end path protection with secondary path(s)

IP RSVP-TE: how does it work?

- Pre-Requisite for setting up IP RSVP-TE Tunnels:
  - Egress router allocates Egress Address Block (EAB)
    - Addresses managed by RSVP at the egress.
    - Addresses not advertised in IGP; not globally routable
- Setting up IP RSVP-TE Tunnels:
  - Ingress initiates IP RSVP-TE signaling:
    - Signals reservations along the explicitly-specified IP Path
    - Request egress router to assign an EA from the EAB for the given IP Path
      - Same EA can be shared among multiple IP Paths
  - Ingress/Transit/Egress routers:
    - program the EA in their forwarding
  - Once setup completes
    - Ingress steers traffic over the IP tunnel
    - Ingress optionally inserts per flow Flow-label for proper ECMP hashing at transit routers

IP RSVP-TE Paths: signaling

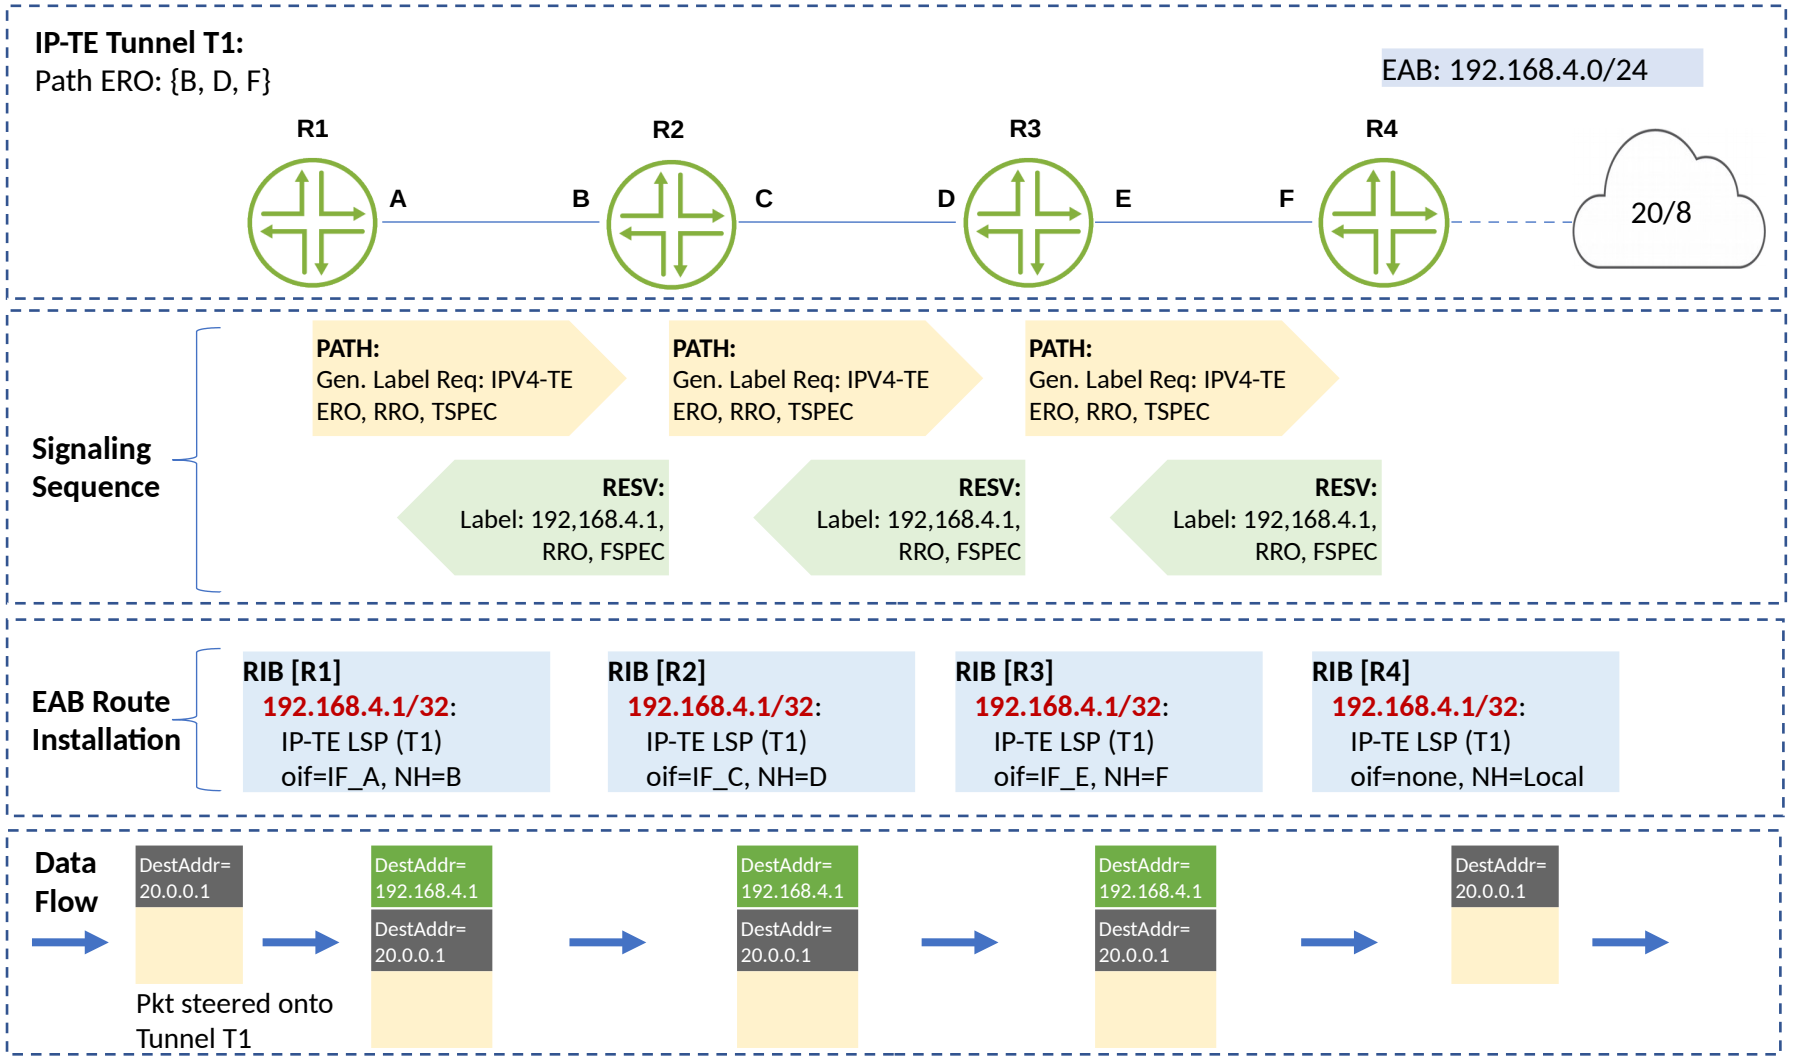

Signaling extensions

The Generalized Label Request Object RFC3471 is used to request an EA address

The Generalized Label object is used to carry the EA address

```
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      | LSP Enc. Type |Switching Type |             G-PID             |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

   To request an IPv4 or IPv6 binding to an IP-TE LSP tunnel, the
   Generalized Label object carries the following specifics:



      1.  the LSP encoding type is set to Packet (1) [RFC3471].

      2.  the LSP switching type is set to "IPv4-TE" (TBD1), or IPv6-TE
          (TBD2)

      3.  the Generalized Payload Identifier (G-PID) MAY be set to All
          (0) or in some cases to the specific payload type if known,
          e.g.  Ethernet (33) [RFC3471].
```

IP RSVP-TE Paths: fastreroute

T100 (bypass) protects T1 IP Path (protected) against Link CD failure

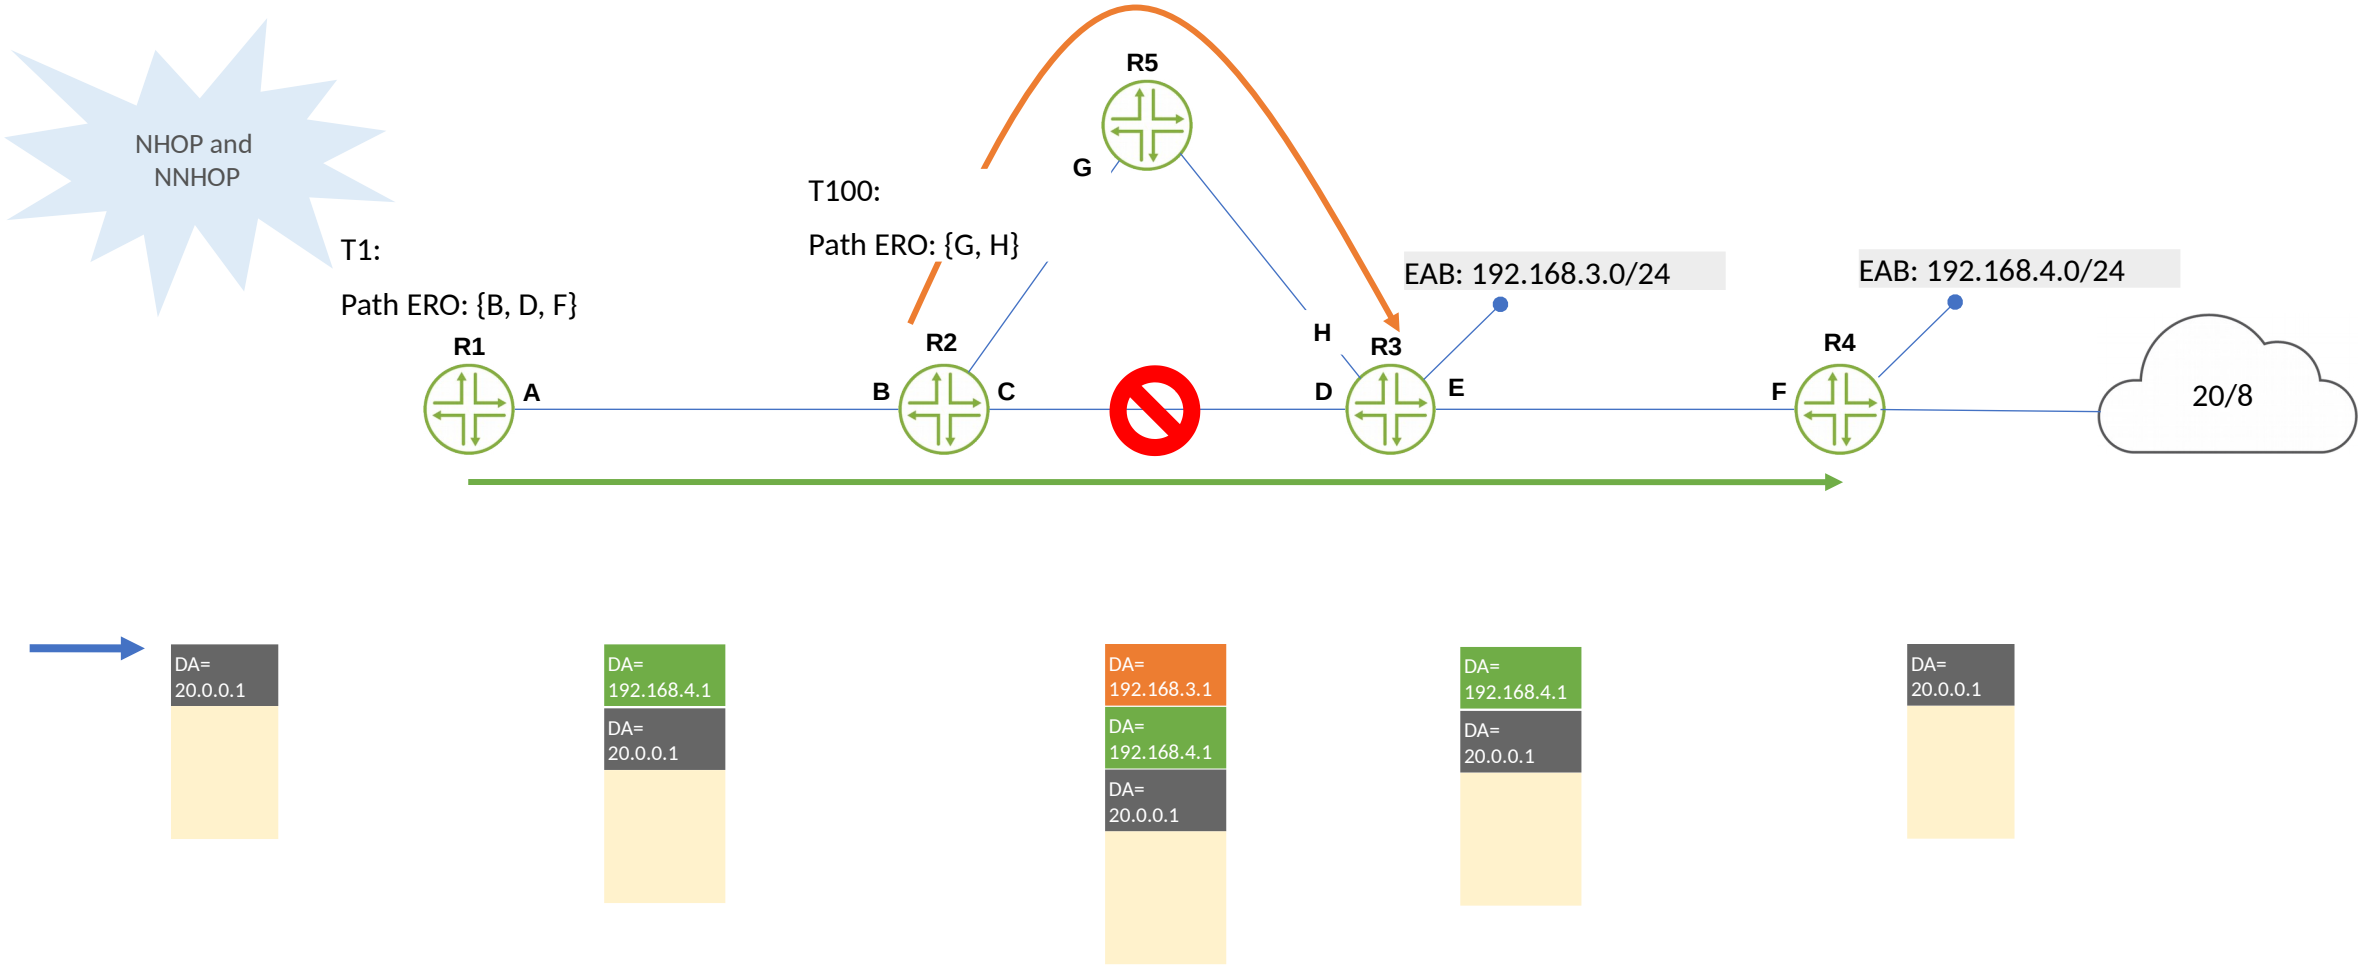

IP RSVP-TE Paths: forwarding state sharing

MP2P path sharing

RIB (R1):

192.168.5.1/32: IP-PATH(GREEN)

oif=if1, NH=R2

RIB (R2):

192.168.5.1/32: I-PATH(GREEN, BLUE)

oif=if1, NH=R3

RIB (R3):

192.168.5.1/32: I-PATH(GREEN, BLUE, RED)

oif=if1, NH=R4

RIB (R4):

192.168.5.1/32: IP-PATH(GREEN,BLUE,RED)

oif=if1, NH=R5

RIB (R6):

192.168.5.1/32: IP-PATH(RED)

oif=if1, NH=R3

Record Route Object (RRO) is used by egress to detect if it is possible that IP-Path(s) can share same EAB address

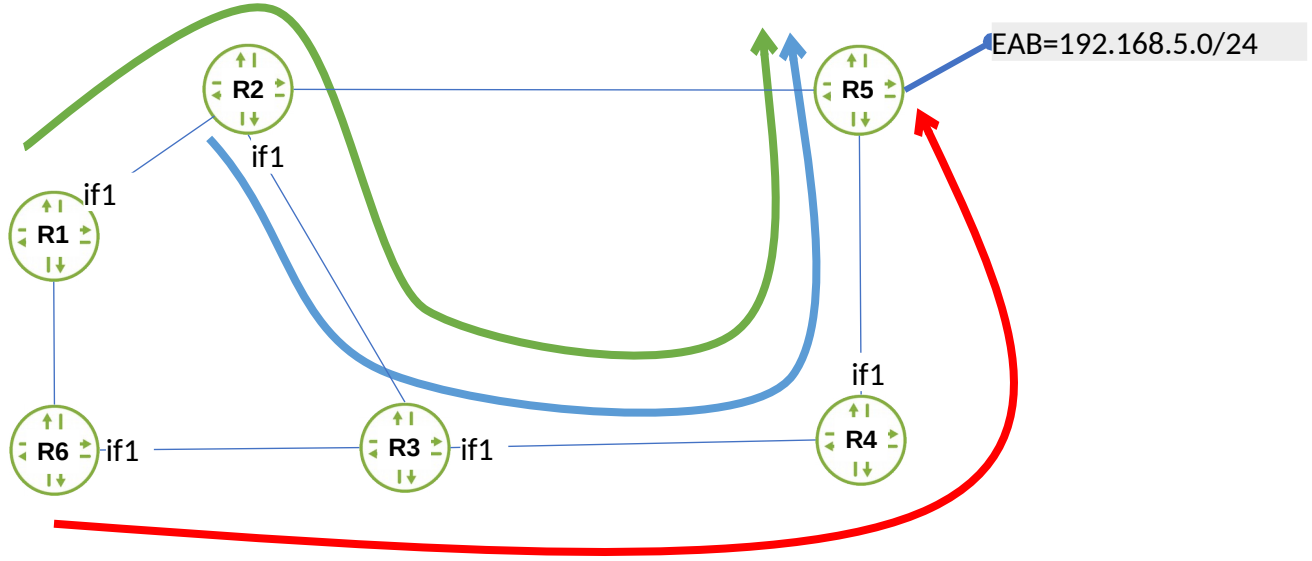

IP Paths can share the forwarding state (same EAB address) if they share **the full path** towards the destination after merge happens

Example: On R3, IP-Path (BLUE), IP-Path (RED) and IP-Path (GREEN) share the path (R3, R4, R5) after they merge on R3

Next steps

- Solicit more input from the WG
- Additional features attributes are under consideration