Distributed Bandwidth Management in Segment Routing Networks using RSVP-TE

draft-beeram-mpls-rsvp-sr

Signaling RSVP-TE Tunnels on an SR-MPLS Forwarding Plane Using Adjacency Segment Identifiers

Vishnu Pavan Beeram HPE · *Colby Barth* HPE · *Andy Smith* Arrcus

Status: Adoption call in progress in the MPLS WG

draft-beeram-teas-rsvp-srv6

Signaling RSVP-TE Tunnels on an SRv6 Forwarding Plane Using End.X Segment Identifiers

Vishnu Pavan Beeram HPE · *Colby Barth* HPE · *Andy Smith* Arrcus
Sven Maduschke Verizon

Status: SRv6 interop demonstrated at the IETF 127 Hackathon (HPE, Arrcus)

Presentation Scope

In scope for this slot

- The why: the toolkit, and why SR operators want it now
- How this fits the SR architecture
- Open points where SPRING input helps
- SR-MPLS first, then SRv6 as the same story on a different data plane

THE ONE-LINE VERSION

Give SR operators the proven RSVP-TE bandwidth toolkit, with no orchestrator, on the SIDs they already have.

Not in this slot

- Encodings and procedures: MPLS WG (10 min) and TEAS (15 min), later this week

THREE QUESTIONS FOR SPRING TO MULL OVER

1. **Do we need bandwidth management in SR networks?**
2. **Do we need distributed bandwidth management in SR networks?**
3. **Can we reuse the tried and tested RSVP-TE bandwidth management toolkit in SR networks?**

The Distributed Bandwidth Management Toolkit

Per-hop bandwidth admission control

Bandwidth is checked and reserved at every hop, against the current state of the link.

Priority-based preemption

Setup and holding priorities let important tunnels displace less important ones when a link is full.

Distributed auto-bandwidth

Each ingress resizes its tunnel, or container-tunnel, to measured traffic, and every hop re-admits the new size.

Bandwidth-aware local protection

Fast reroute that accounts for bandwidth, so protection path does not run into congestion when in use.

Widely deployed in traditional MPLS networks.
RFC 8577 brought the toolkit to a shared MPLS forwarding plane, using TE link labels.

Why Now? Operators Want a Choice

Early View

Sophisticated bandwidth engineering is not needed in SR networks.

Controller plane

The need was acknowledged. The answer was a controller plane that holds and manages the bandwidth state.

Today

Several SR operators, including those with RSVP-TE practice, want the toolkit without an off-box centralized reservation orchestrator.

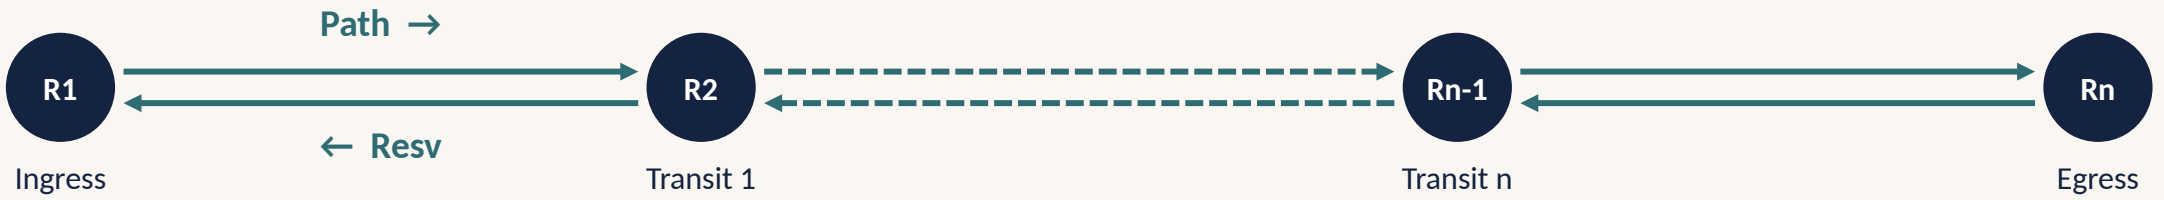
A choice, not a replacement: centralized, hybrid or distributed.
These drafts add the distributed option and remove nothing.

The Basic Idea

The signaling sequence not only takes care of admission control and per-link bandwidth reservation; **it also picks or validates the SIDs the ingress uses in its segment list.**

SID hop 1
SID hop 2
⋮
SID hop n
service SID

Segment list at the ingress
SR-MPLS: Adj-SIDs · SRv6: End.X SIDs



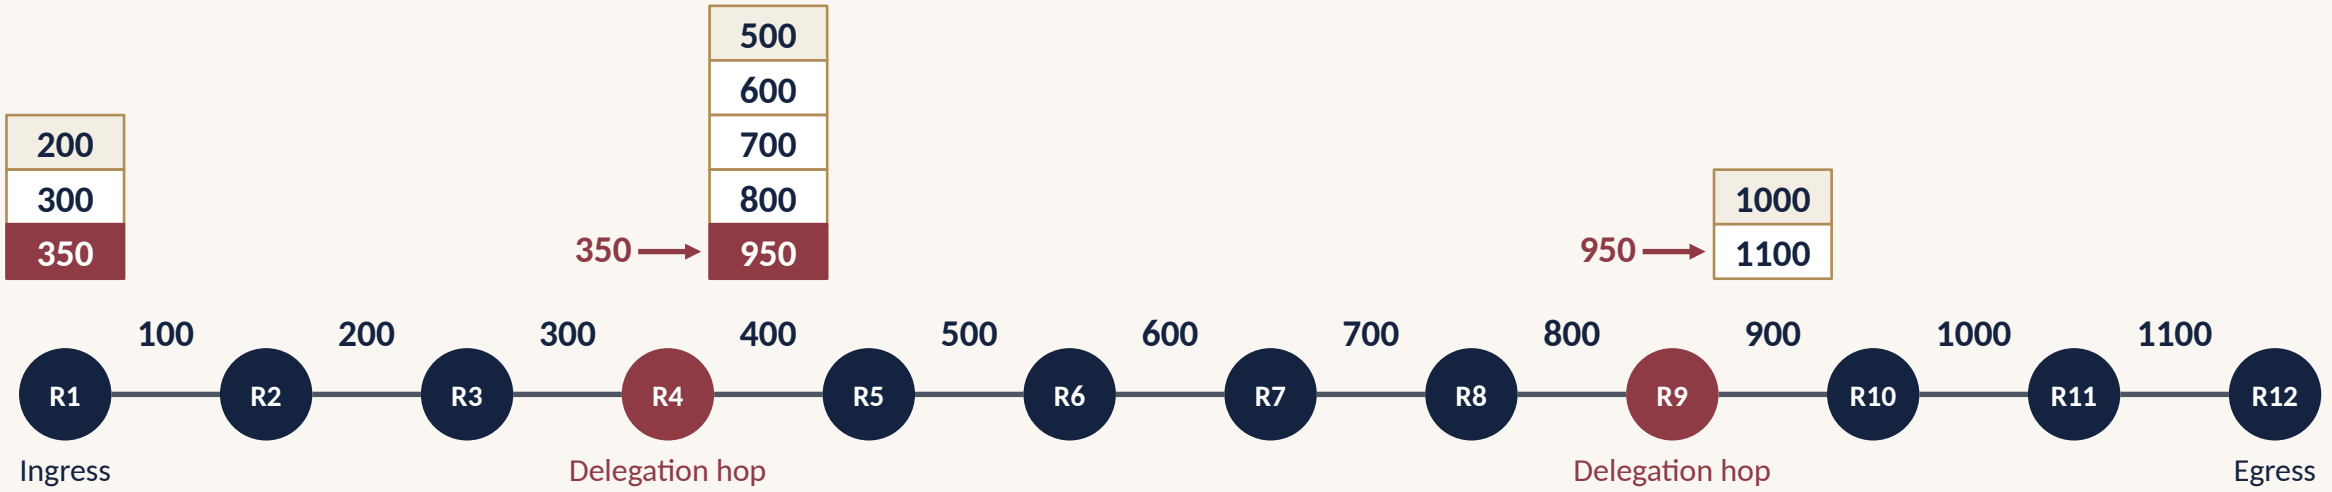
- Leverages existing SID types advertised by the IGP

Beyond Bandwidth: Automatic Segment List Optimization

Without delegation, the ingress pushes 10 SIDs. With two signaled Binding SIDs, it pushes 3.

Binding SIDs, instantiated by signaling at the delegation hops
Adj-SIDs (End.X SIDs in SRv6) of the outgoing link
Top of the stack

Assumptions (all nodes)
MSD 5 · RLD 8 · 2 service SIDs at the ingress



- Delegation hops R4 and R9 instantiate Binding SIDs 350 and 950, each standing for the SIDs that follow
- No node pushes more SIDs than it can, and no transit node receives more than it can read

A controller can do this too. With a signaling plane, the network does it automatically.

SRv6: The Same Why, a Larger Gap

No RFC 8577 equivalent for SRv6

- No distributed bandwidth management toolkit today
- No controller-less way to instantiate Binding SIDs dynamically

Why the gap is larger

- 128-bit SIDs: long segment lists strain headend and MTU limits
- Micro-SIDs (NEXT-CSID, RFC 9800) ease this, but operators want to stay within one 128-bit carrier: six micro-SIDs in the common format
- Binding SIDs keep longer paths within that limit

What the draft covers

- Compressed segment lists accommodated
- Endpoint behavior flavors (PSP, USD) requested by the ingress

Details: TEAS session

Architectural Stance: What We Reuse, and What It Costs

What we reuse

- Existing SID types: Adj-SIDs and End.X SIDs
- Existing forwarding plane, with no new per-tunnel forwarding state
- Minor tweaks to accommodate SID protection mechanisms
- Scope: all four toolkit items; protection is link protection only for now

The cost

- RSVP session state at each hop (control plane only)
- It buys distributed admission control, priorities and preemption
- Scales with the number of tunnels, not forwarding entries

Architectures evolve with operator needs.
Our role is to widen choices, not to pick one model.

Open Points Where SPRING Input Helps

Coexistence with SR-native traffic (RFC 8426)	Recommended: selected SIDs carry only RSVP-signaled traffic. Otherwise static partitioning. A marked or dedicated SID, if wanted, is a separate document.
Headend construct: TE tunnel or SR Policy?	Draft is agnostic. A separate SPRING document could specify an SR Policy knob if the WG wants it.
Protection flavors beyond TI-LFA	A signaling plane brings new Adj-SID protection flavors, such as RSVP-signaled facility backup. Telling them apart from TI-LFA needs an IGP extension. Protection text to be completed.
Fixes planned in the next revision	Scope wording, naming, an operational section on non-TE traffic, CSPF/TED text, tooling-impact note.

The Need, the Viability, and the Ask

The need

- SR operators need bandwidth management
- Some want it distributed, without an orchestrator
- A proven MPLS toolkit, not yet available for SR

The viability

- Reuses Adj-SIDs and End.X SIDs
- SRv6 interop shown at IETF 127 Hackathon
- Remaining open points are known and scoped

THE ASK

Should SR evolve to give operators a distributed bandwidth management option, by signaling RSVP-TE over IGP-advertised SIDs?

If you see the need and the viability, please say so on the MPLS adoption call (ends 17 Nov) and in TEAS. If not, tell us what would have to change.

Thank You

Questions and comments welcome

draft-beeram-mpls-rsvp-sr
draft-beeram-teas-rsvp-srv6