


SRv6, traffic engineering, production roadmap

Agenda

Part 1	Why Choose SRv6?
Part 2	SRv6 Transport Technology
Part 3	Traffic Engineering
Part 4	VNPT Roadmap

1. Why Choose SRv6?

A single IPv6 dataplane, replacing multiple MPLS layers



Simpler

Eliminate LDP / RSVP-TE / MPLS labels

Programmable

Ready for SDN & automation

Highly Scalable

One dataplane for entire Network

Smooth migration

Coexists, migrating gradually from MPLS

1. Why Choose SRv6?

IETF: Continuous standardization from architecture to SID compression (2018-2025)

2018	2019	2020	2021	2022	2023	06/2025
RFC 8402	RFC 8660	RFC 8754	RFC 8986	RFC 9256/9252	RFC 9350	RFC 9800
Segment Routing Architecture	SR-MPLS	SRH, SRv6 data plane	SRv6 Network Programming	SR Policy; BGP service over SRv6	Flex-Algo	SID Compression: uSID / cSID

2. SRv6 Transport Technology

Three fundamental principles of SRv6

Source routing

The Headend selects the path, the core is stateless

Native IPv6


No MPLS labels, forwarding is based on IPv6

Network programming

One SID = one Node + one action

2. SRv6 Transport Technology

SID – Segment Identifier: Each segment is assigned a numeric value. This value is called a SID.




- Each segment is assigned a unique SID
- The SID tells the device what action to perform next.
- A sequence of SIDs forms the path (Segment List) the packet must traverse

2. SRv6 Transport Technology

Full SID: Each segment is a full 128-bit IPv6 address


- Each SID is a 128-bit IPv6 address: Locator (routes to the node) + Function (action at the node)
- Headend encodes the segment list [R4, R6, R8] in the SRH and sets the DA to the first SID
- Each hop: decrement Segments Left, copy the next SID into the DA
- Cost: 16 bytes per SID (3 SIDs = 56-byte SRH), limited by hardware MSD



BGP: service SID 5f00:0:8::d4 (End.DT4) advertised R8 → R1

2. SRv6 Transport Technology


uSID: Compresses each segment to 16 bits.



- Each node removes its own uSID from the Destination Address, shifts the remaining uSID left, and forwards the packet to the next hop.
- Up to ~6 16-bit uSIDs in one IPv6 address (container)
- Paths of ≤6 hops: only the IPv6 header, no SRH


2. SRv6 Transport Technology

SRH: when more than 6 uSIDs are required.




2. SRv6 Transport Technology

SRH: when more than 6 uSIDs are required.



2. SRv6 Transport Technology

SRH: when more than 6 uSIDs are required.




2. SRv6 Transport Technology

SRv6 Function (Endpoint Behavior)

Codename	uSID	Behavior
End	uN	Endpoint
End.X	uA	Endpoint with Layer-3 cross-connect
End.B6.Insert	uB6.Insert	Endpoint bound to an SRv6 policy
End.B6.Encap	uB6.Encaps	Endpoint bound to an SRv6 encapsulation policy
End.DX6	uDX6	Endpoint with decapsulation and IPv6 cross-connect
End.DX4	uDX4	Endpoint with decapsulation and IPv4 cross-connect
End.DT6	uDT6	Endpoint with decapsulation and specific IPv6 table lookup
End.DT4	uDT4	Endpoint with decapsulation and specific IPv4 table lookup
End.DX2	uDX2	Endpoint with decapsulation and L2 cross-connect
End.DT2U/M	uDT2U/M	Endpoint with decapsulation and L2 unicast lookup / flooding
End.DTM	uDTM	Endpoint with decapsulation and MPLS table lookup
H.Insert / H.Encaps		Headend with Insertion / Encapsulation of / into an SRv6 policy
H.Encaps.L2		H.Encaps applied to received L2 frames
H.Encaps.M		H.Encaps applied to MPLS label stack

3. Traffic Engineering

SRv6 BE: best-effort by default



- The Headend (PE1) sets one SID (egress), no SRH is required
- The Core (P1) only performs normal IGP table lookup
- The Egress (PE2) performs End.DT4/DT6: decapsulation, then VRF lookup and forward to CE2
- Traffic follows the IGP shortest path; no steering policy or SLA guarantees

3. Traffic Engineering

SRv6 TE Policy: Intent-based traffic engineering is performed through a SID list


Simple, automated and highly scalable

- No state is maintained on core nodes
- No pre-configuration is required on the devices
- No pre-steering is required at the headend: automated steering

An SRv6-TE policy is identified by:

- Head-end
- Color
- End-point

SR Policy
(1, green, 4) Head-end: 1 Color: green End-point: 4




3. Traffic Engineering

SRv6 TE Policy


SRv6-TE policy - Candidate Path (CP):

- An SR-TE policy has one or more CPs (only the active CP forwards traffic based on the preference attribute).
- If multiple CPs have the same preference, selection follows the order: protocol-origin, Existing Installed Path, originator, Discriminator.
- A CP can contain one or more segment lists (weight: for traffic load sharing).



SRv6-TE policy – CP types:

- Explicit Path
- Dynamic Path
- Composite




3. Traffic Engineering

ODN: On-Demand Nexthop

Automatically generate SRv6 TE policy

- There is an SRv6 TE Policy configuration template for each Color.
- When the headend receives a BGP route with a color community and next-hop, it creates the corresponding SRv6 TE Policy for that color and next-hop.
- The headend can compute the path itself or delegate it to an SDN Controller (PCEP/BGP SR).




3. Traffic Engineering

PCE/Controller

The controller steers SRv6 traffic (instead of routers deciding on their own):

- Collect and discover the network: PCEP, BGP-LS, telemetry.
- Compute paths: CSPF, disjoint, bandwidth
- Create, update, delete and optimize policies and push them to routers: PCEP, BGP SR, NETCONF/YANG



3. Traffic Engineering

Binding SID

One BSID identifies one SRv6-TE Policy

- Hierarchical networking
- Scalability
- Service steering

First DA at PE-A

fcbb:bb00:10a1:b0:30b2:d4::


10a1 uN of ABR-A

b0 BSID at ABR-A (maps to the policy)

30b2 uN of PE-B


d4 uDT4 at PE-B (VRF decap)

Data forwarding with BSID:



3. Traffic Engineering


TI-LFA/Microloop Avoidance



- The PLR pre-computes a post-convergence repair path for each SID
- The repair path is guaranteed loop-free
- Activated locally when a link/node fails, without waiting for IGP convergence
- Default protection for SRv6 BE, recovery in under 50 ms

3. Traffic Engineering

sBFD: Failure detection and path protection for SRv6 TE



sBFD

End-to-end monitoring along the exact segment list

Path protection

Switch to the disjoint backup CP when the primary fails; revert when stable

3. Traffic Engineering

Flex-Algo: Multiple planes on a single topology

Each Flex-Algo is defined by an FAD:


Metric	IGP · TE · Delay (measured by TWAMP/STAMP)
Constraints	Include / exclude affinity · exclude SRLG

Flex-Algo characteristics:

- Each node advertises a separate locator per algo; SPF and TI-LFA are computed separately within each plane, and nodes compute independently in a distributed manner.
- A single SID is enough to force traffic onto a specific path/plane.
- No automatic fallback to algo 0.

Algorithm values

0	SPF based on IGP metric (default)
1	Strict SPF (cannot be overridden by local policy)
2 - 127	Unassigned
128 - 255	Reserved for operator-defined Flex-Algo



4. VNPT Roadmap

Addressing and Flex-Algo Planning for the VNPT Network

SRv6 SID

5f00::/16 (RFC 9602)

Flex-Algo

Minimum 3 algorithms: Default · Delay · TE metric

Compression mechanism

Prioritize 16-bit SRv6 uSID design

BSID


Improves scale and enables easier interoperability across multi-domain.

SDN

Enforce TE, network slicing, SLA assurance, seamless

4. VNPT Roadmap

Locator Planning:



4. VNPT Roadmap

Three phases: Lab → Pilot metro → network-wide seamless



Thank you
Q & A