

SRv6 BGP based overlay services

draft-dawra-bess-srv6-services-02

Author(s):

Gaurav Dawra, LinkedIn
Clarence Filsfils, Cisco Systems
Swadesh Agarwal, Cisco Systems
Satoru Matsushima, Softbank
Daniel Voyer, Daniel Bernier, Bell Canada
Bruno Decraene, Orange
Dirk Steinberg, Steinberg Consulting
Robert Raszuk, Bloomberg LP
John Leddy, Comcast
Shunwan Zhuang, Huawei
Jorge Rabadan, Nokia

Presenter :

Gaurav Dawra, LinkedIn

IETF105, July/2019

Montreal, Canada



High Level Updates

- Presented L3VPNoSRv6 in IDR at IETF98
- Presented EVPNoSRv6/Global AFs in IDR at IETF101
- Presented SRv6 Services Draft in BESS at IETF104
- SRv6 Services document have matured significantly:
 - Multiple implementations across multiple vendors
 - Multiple customer deployments across the globe
- Minor updates to this document to cover SRv6 BGP Services:
 - BGP update packaging issues in certain cases has been addressed(ex. EVPN VPWS Service)
 - Introduction of *optional* SRv6 SID Structure Sub-Sub TLV
 - Addressed comments from WG members
 - Includes contributions from Ali Sajassi and Jorge Rabadan

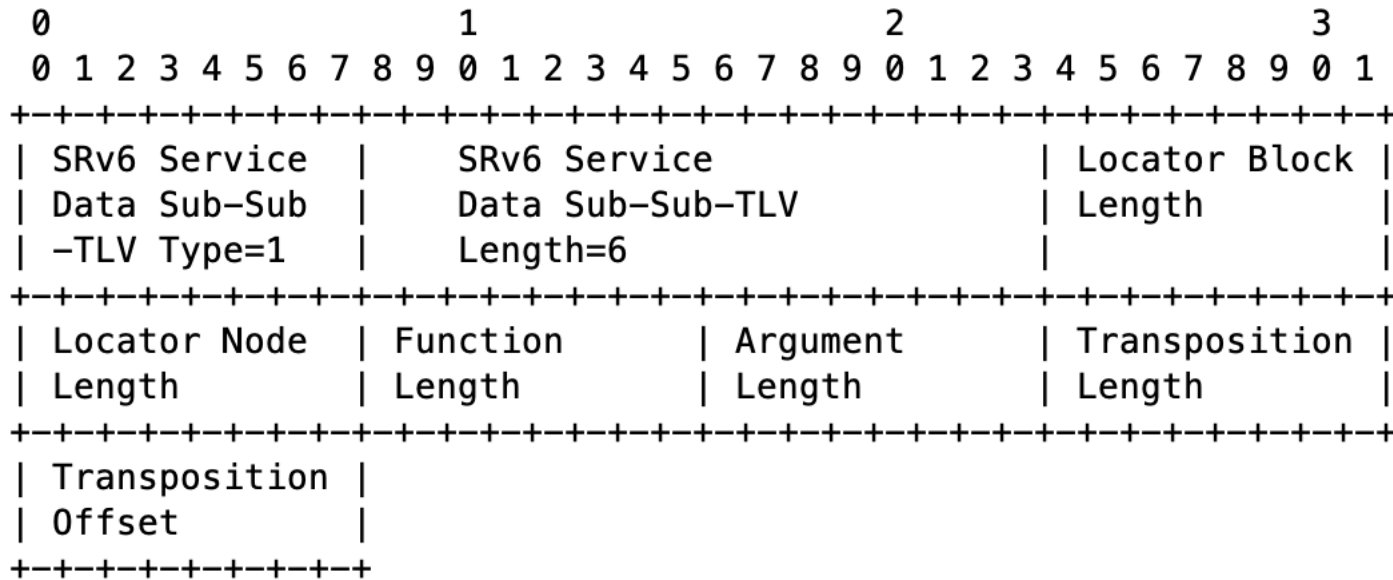
SRv6 Services Update

Introduced further optimization to BGP packaging in certain cases:

- **Problem:** When Function part of SRv6 SID is unique per NLRI(example EVPN VPWS service). SID Encoding packing may be sub-optimal when other BGP attributes are shared across NLRI.
- **Solution:**
 - Introduced flexibility of advertising Function and Arguments of SRv6 SID in label field of route NLRI
 - Advertisement of SRv6 SID structure sub-sub tlv in SID attribute to advertise the lengths of each individual parts of the SRv6 SID
 - Fully backward compatible to current implementations



Optional SRv6 SID Structure Sub-Sub-TLV



- **SRv6 Service Data Sub-Sub-TLV Type (1 octet):** Set to 1 to represent SRv6 SID Structure Sub-Sub TLV
- **TLV Length:** 16bit field. The total length of the value portion of the TLV.
- **Locator Block Length:** 8 bit field. Length of SRv6 SID locator Block in bits.
- **Locator Node Length:** 8 bit field. Length of SRv6 SID locator Node in bits.
- **Function Length:** 8 bit field. Length of SRv6 SID Function in bits.
- **Transposition Length:** 8 bit field. Size in bits for the part of SID shifted into label
- **Transposition Offset:** 8 bit field. The offset position in bits for SID shifted into label



Implementation Status

- Customer deployments: Softbank, China Telecom, Iliad etc (multiple more in progress...)
- Three Cisco Hardware-forwarding platforms: ASR 1K, ASR 9k and NCS 5500
- Two Cisco network operating systems: IOS XE and IOS XR
- Huawei Hardware-forwarding platforms: ATN, CX, ME, NE5000E, NE9000, NG-OLT
- Huawei network operating systems: VRPv8
- Barefoot Networks Tofino on OCP Wedge-100BF
- Linux Kernel officially up streamed in 4.10, fd.io
- More details: [draft-matsushima-spring-srv6-deployment-status/](#)



Draft: Next Steps

- Seeking WG adoption!
- Suggestions/comments are welcome!

