



IETF 116 – Yokohama
Mar 2023
BESS Working Group

draft-trr-bess-bgp-srv6-args-01

Ketan Talaulikar (Cisco)

Kamran Raza (Cisco)

Jorge Rabadan (Nokia)

Wen Lin (Juniper)

Background

- During implementation and interop of BGP Services for SRv6 (RFC 9252) we found that:
 - The specifications were not detailed enough for End.DT2M signalling
 - Certain ambiguities needed to be clarified
 - Examples needed for more clarity
- An incorrect assumption needed to be addressed
 - For End.DT2M imposition, the procedure specified in RFC9252 is to perform OR-ing of SID advertised in RT3 and ARG advertised in RT1
 - However, this is workable only when the SID structures signalled via the two routes are identical
 - Generic mechanism is to pick ARG from RT1 and put it after LOC+FUNC in the SID advertised via RT3

This draft proposes to update RFC9252 for the procedures related to the signalling of EVPN Route Types 1 & 3 with the End.DT2M behavior

Purpose of SID Advertisements with Route Types

- Ethernet ES-AD Route Type 1 (with End.DT2M)
 - To provide Arg.FE2 when using ESI Filtering and the size of the ARG
- IMET Route Type 3 (with End.DT2M)
 - To provide LOC+FUNC parts of the SID
 - To indicate support for receiving Arg.FE2, the ARG size, and ARG position in the SID to be constructed

Example: With ESI Filtering

Ethernet AD-ES Route Type 1

Route Type 1:

< NLRI >

< other attrs >

BGP Prefix SID Attr:

SRv6 L2 Service TLV:

SRv6 SID Information sub-TLV:

SID: 0:0:0:0:aaaa::

Behaviour: End.DT2M

SRv6 SID Structure sub-sub-TLV:

LBL: 32, LNL: 16, FL: 16, AL: 16, TPOS-L: 0, TPOS-O: 0

IMET Route Type 3

Route Type 3:

< NLRI >

< other attrs >

BGP Prefix SID Attr:

SRv6 L2 Service TLV:

SRv6 SID Information sub-TLV:

SID: 2001:db8:1:fd1::

Behaviour: End.DT2M

SRv6 SID Structure sub-sub-TLV:

LBL: 32, LNL: 16, FL: 16, AL: 16, TPOS-L: 0, TPOS-O: 0

Example: Without ESI Filtering (or no multi-homing)

Ethernet AD-ES Route Type 1

Route Type 1:

< NLRI >

< other attrs >

BGP Prefix SID Attr:

SRv6 L2 Service TLV:

SRv6 SID Information sub-TLV:

SID: ::

Behaviour: End.DT2M

SRv6 SID Structure sub-sub-TLV:

LBL: 32, LNL: 16, FL: 16, AL: 0, TPOS-L: 0, TPOS-O: 0

IMET Route Type 3

Route Type 3:

< NLRI >

< other attrs >

BGP Prefix SID Attr:

SRv6 L2 Service TLV:

SRv6 SID Information sub-TLV:

SID: **2001:db8:1:fb1::**

Behaviour: End.DT2M

SRv6 SID Structure sub-sub-TLV:

LBL: 32, LNL: 16, FL: 16, AL: 0, TPOS-L: 0, TPOS-O: 0

Processing on Ingress PE to form the End.DT2M Service SID used in the packet

- Received Route Type 3 gives the LOC+FUNC and the structure of the SID
 - Check if ARG is supported (e.g., when ESI filtering is in use) via the SID Structure in RT3
- IF no support for ARG:
 - Then the encoded SID is just LOC+FUNC signalled via RT3
- Else:
 - Check for matching Route Type 1 and if has been signalled with End.DT2M behavior
 - If yes, check that the ARG Length (AL in SID Structure) in Route Type 1 match to what is advertised in Route Type 3
 - If ARG length is consistent, encode the ARG value signalled via Route Type 1 into the position of ARG as indicated by the SID structure of Route Type 3 (i.e., after the LOC+FUNC)

Processing change from
RFC9252
(instead of OR-ing)

Backward Compatibility

- No encoding changes
- Mostly procedural clarifications
- Interoperates with RFC9252 procedure for OR-ing of SID values between Route Types 1 & 3 as long as their SID Structures are identical

Next Steps ...

- Seek WG review and inputs
- Can we do an expedited WG adoption?
 - The need for these clarifications have been acknowledged by many WG members