

IETF 106 – Singapore
Nov 2019

SR Generic FEC TLV for LSP Ping

(draft-nainar-mpls-spring-lsp-ping-sr-generic-sid)

Nagendra Kumar Nainar, Ed.

Carlos Pignataro, Ed.

Zafar Ali (Presenter)

Clarence Filsfils

(Cisco Systems, Inc.)

Problem Statement

- Requires new target FEC Stack sub-TLV definition and standardization efforts for each new Segment ID defined.
 - Define new TLV.
 - Update FEC validation procedure of RFC-8029
- Requires domain/node wide software upgrade depending on the type of the Segment ID defined.
- Raises scalability challenges.

Problem Statement (A partial list of New SR FECs)

BGP Peer Node SID

```

0          1          2          3
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
+-----+-----+-----+-----+
|Type = TBD1          |          Length = x          |
+-----+-----+-----+-----+
|  AF.Type          |          Reserved          |
+-----+-----+-----+-----+
|          Local BGP Router ID (4 octets)          |
+-----+-----+-----+-----+
|          Local ASN (4 octets)          |
+-----+-----+-----+-----+
|          Peer BGP Router ID (4 octets)          |
+-----+-----+-----+-----+
|          Peer ASN (4 octets)          |
+-----+-----+-----+-----+
|          Local Interface address (4 or 16 octets)          |
+-----+-----+-----+-----+
|          Remote Interface address (4 or 16 octets)          |
+-----+-----+-----+-----+

```

BGP Peer Adj-SID

```

0          1          2          3
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
+-----+-----+-----+-----+
|          Type = TBD2          |          Length = 24          |
+-----+-----+-----+-----+
|          Local BGP Router ID (4 octets)          |
+-----+-----+-----+-----+
|          Local ASN (4 octets)          |
+-----+-----+-----+-----+
|          Peer BGP Router ID (4 octets)          |
+-----+-----+-----+-----+
|          Peer ASN (4 octets)          |
+-----+-----+-----+-----+
|          Local Link Identifier (4 octet)          |
+-----+-----+-----+-----+
|          Remote Link Identifier (4 octet)          |
+-----+-----+-----+-----+

```

BGP Peer Set SID

```

0          1          2          3
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
+-----+-----+-----+-----+
|Type = TBD3          |          Length = x          |
+-----+-----+-----+-----+
|          Local BGP Router ID (4 octets)          |
+-----+-----+-----+-----+
|          Local ASN (4 octets)          |
+-----+-----+-----+-----+
|          Peer Set Count          |          Reserved          |
+-----+-----+-----+-----+

```

List of Peer Set Sub-TLVs

BGP Peer Set SID Sub-TLVs

```

0          1          2          3
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
+-----+-----+-----+-----+
|Type = 1 (Peer)          |          Length = 8          |
+-----+-----+-----+-----+
|          Peer ASN (4 octets)          |
+-----+-----+-----+-----+
|          Peer BGP Router ID (4 octets)          |
+-----+-----+-----+-----+

```

FEC changes for Flex-Algo

```

0          1          2          3
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
+-----+-----+-----+-----+
|          IPv4 prefix          |
+-----+-----+-----+-----+
|Prefix Length |          Protocol          |          Algo          |          Reserved          |
+-----+-----+-----+-----+

```

```

0          1          2          3
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
+-----+-----+-----+-----+
|Type = 2 (Link Id)          |          Length = 12          |
+-----+-----+-----+-----+
|          Peer ASN (4 octets)          |
+-----+-----+-----+-----+
|          Local Link Identifier (4 octet)          |
+-----+-----+-----+-----+
|          Remote Link Identifier (4 octet)          |
+-----+-----+-----+-----+

```

Problem Statement (Cont'ed)


- Requires a lot of information to be derived by the Initiator to include in the Echo Request.
- Complex FEC filling procedures at Ingress (one for each Prefix SID type).
- Complex validation procedures at Egress (one for each Prefix SID type).

Solution

- SR SID data model is:
 - Segment ID (Label)
 - SID Assigner

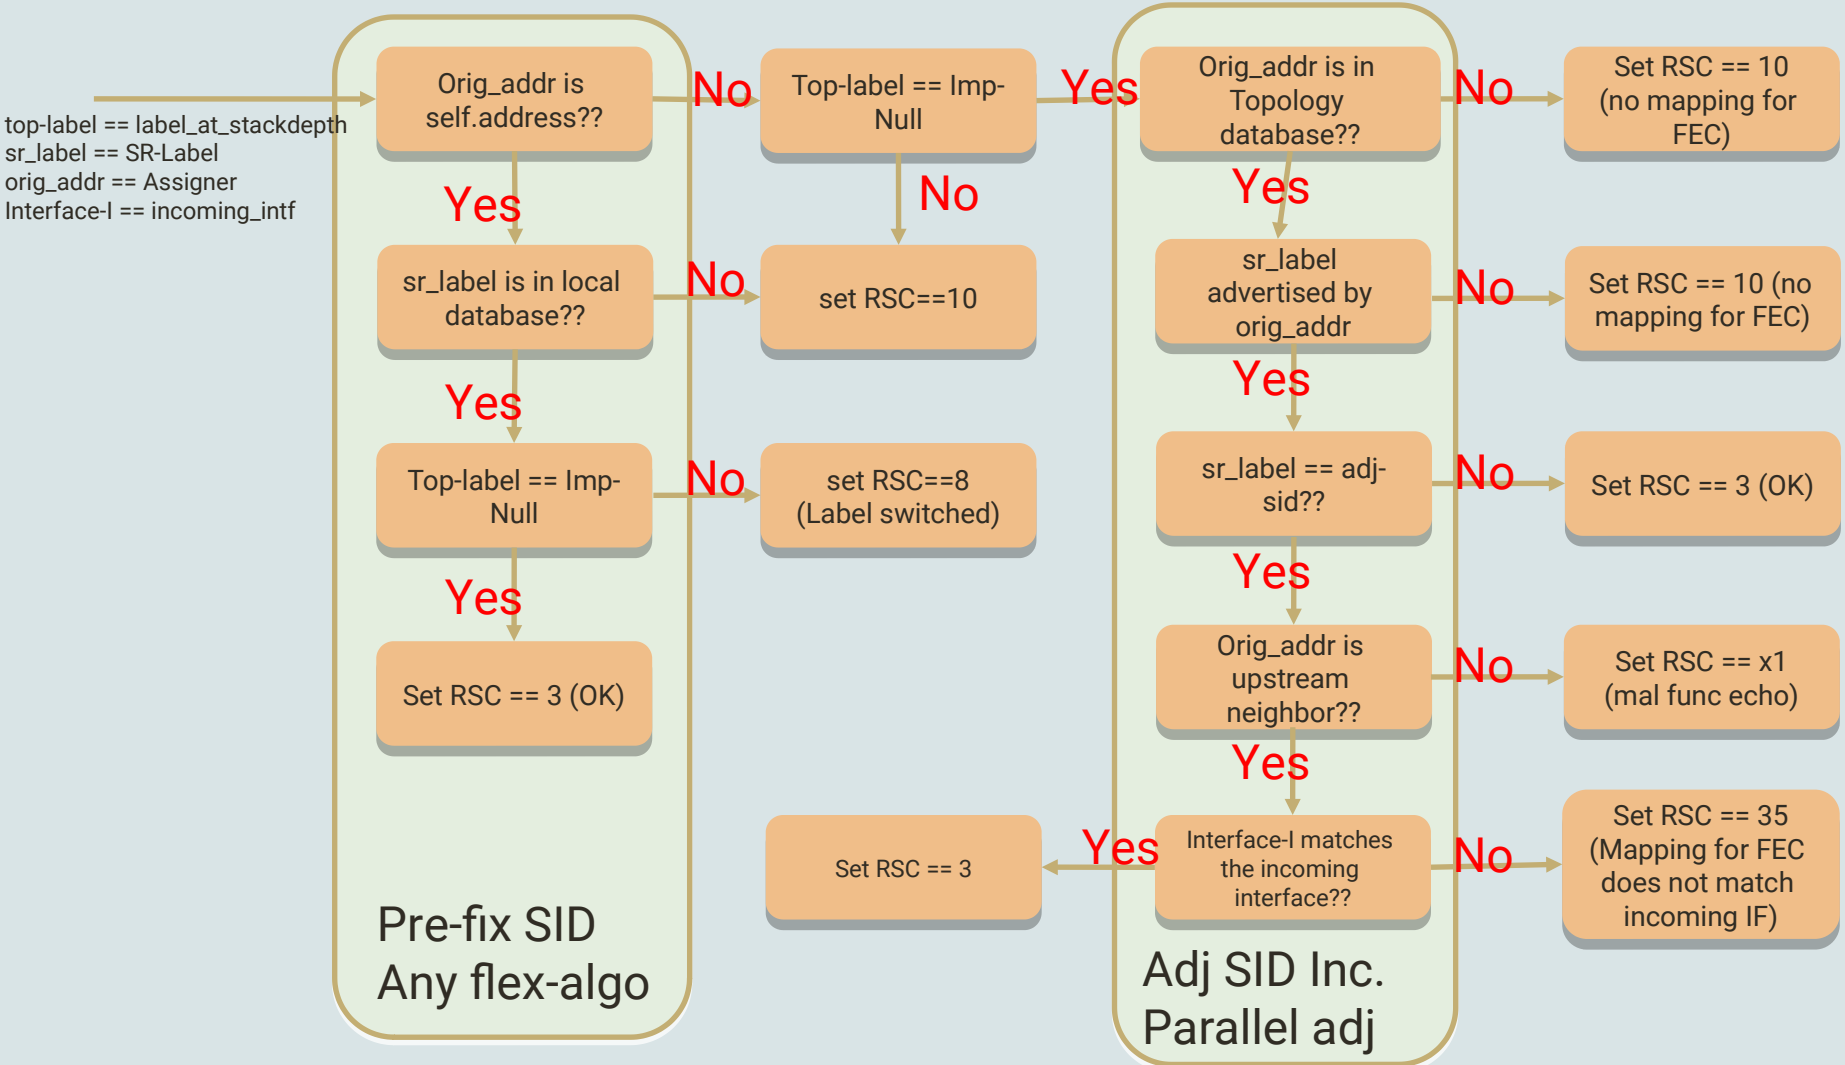
- FEC validation Procedure
 - Segment ID to Interface mapping is maintained by any node.
 - Local implementation matter
 - Initiator defines the SID value and LSP EndPoint while triggering LSP Ping
 - Manually defined via CLI or dynamic PCE query.
 - Responder validates the LSP End Point and incoming interface.
 - Respond based on the validation.

SR Generic Label Sub-TLV



- SR SID
 - Carries 20 bits of Segment ID used for validation.
- LSP End Point
 - Node address of the endpoint that terminates the LSP.
 - LSP End Point may be set to 0.0.0.0 by the initiator.
 - E.g., for parallel adjacency.
 - If LSP End Point address is set, the Egress MAY skip the SID assigner check.
 - E.g., for BSID


Responder behavior



Procedure

Segment ID to Interface Mapping

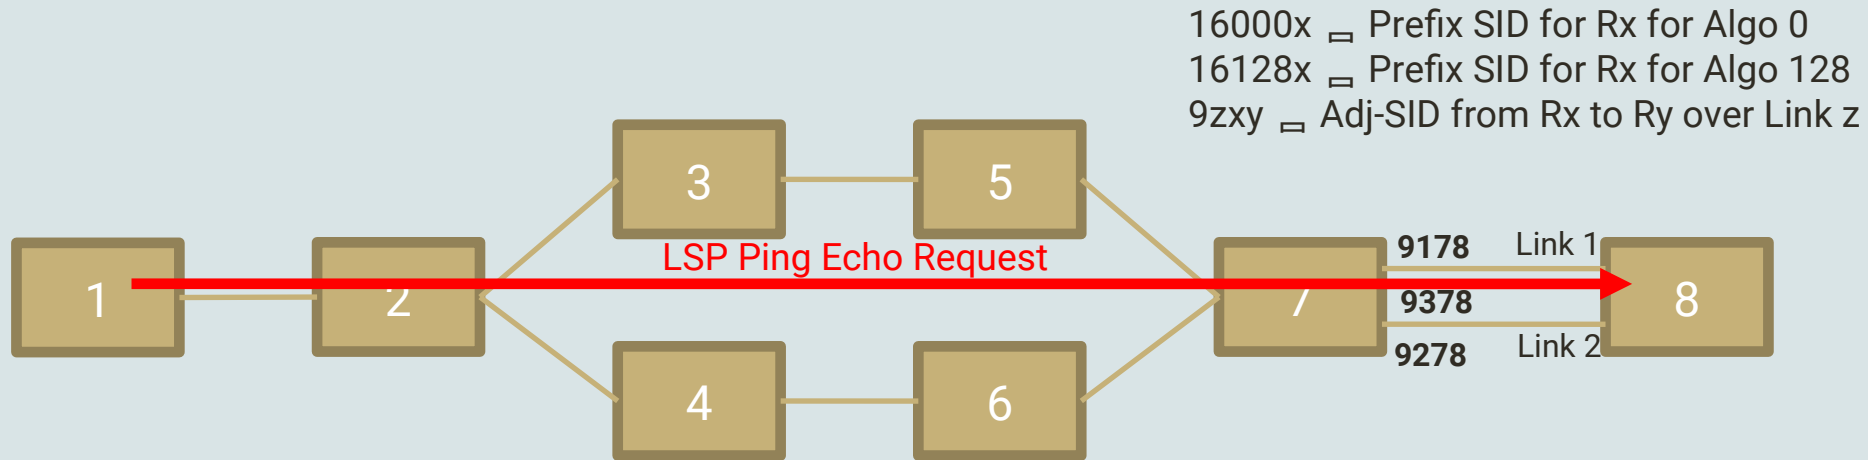
16000x \sqsubseteq Prefix SID for Rx for Algo 0
16128x \sqsubseteq Prefix SID for Rx for Algo 128
9zxy \sqsubseteq Adj-SID from Rx to Ry over Link z



- R8 maintains the below mapping:
- 160008 \sqsubseteq Incoming Interface: {Any}
 - 161288 \sqsubseteq Incoming Interface: ({Any})
 - 9178 \sqsubseteq Incoming Interface: {Link 1}
 - 9278 \sqsubseteq Incoming Interface: {Link 2}
 - 9378 \sqsubseteq Incoming Interface: {Link 1 or Link 2}

Procedure

Initiator Behavior




➤ Initiator (R1) triggers LSP Ping with below SR Generic Label Sub-TLV:

- For Prefix SID 160008 {SID=160008; LSP-EndPoint = R8}
- For Prefix SID 161288 {SID=161288; LSP-EndPoint = R8}
- For ADJ-SID 9178 {SID=9178; LSP-EndPoint = R8}
- For ADJ-SID 9278 {SID=9278; LSP-EndPoint = R8}
- For ADJ-SID 9378 {SID=9378; LSP-EndPoint = R8}

Procedure

Responder Behavior

16000x \sqsubseteq Prefix SID for Rx for Algo 0
16128x \sqsubseteq Prefix SID for Rx for Algo 128
9zxy \sqsubseteq Adj-SID from Rx to Ry over Link z



- Responder (R8) uses SID, LSP Endpoint info and the local mapping for validation.
 - LSP EndPoint == self?
 - Incoming_Interface matches the local table?
- Sends positive/negative response accordingly.

In a nut shell

- One Target FEC Stack Sub-TLV that covers multiple Segment IDs.
 - Prefix/Flex-Algo/Adj/ePE SIDs
- Drastically reduces the information required on the Initiator.
 - Ease of operation.
- Reduces the information to be processed by the responder.
- Extendable to accommodate future Segment IDs.

IANA Registry Allocation

- Request for a new Sub-TLV for TLV types 1, 16 and 21.
- Value from range 38-31743 (Unassigned range)
- Re-uses existing Return codes and Return Sub-codes

I-D Status

- **Next Steps:**
 - WG feedback sought
 - Textual Contributions Welcomed!
 - WG Adoption in MPLS WG
- **Thank you!**