

IETF 109 (Virtual)
Nov 2020

SR Generic FEC TLV for LSP Ping

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

Nagendra Kumar Nainar, (Presenter)

Carlos Pignataro,

Zafar Ali,

Clarence Filsfils

(Cisco Systems, Inc.)

Tarek Saad,

(Juniper)

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

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          |
+-----+-----+-----+-----+

```

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)          |
+-----+-----+-----+-----+

```

```

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)

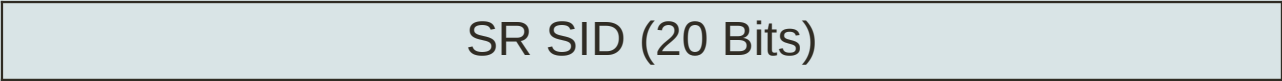
- Requires a lot of information to be derived by the Initiator to include in the Echo Request.
- Complex FEC filling procedures at Ingress (one for each Prefix SID type).
- Complex validation procedures at Egress (one for each Prefix SID type).

Solution

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

- FEC validation Procedure
 - Segment ID to Interface mapping is maintained by any node.
 - Local implementation matter
 - Initiator defines the SID value.
 - Responder validates if it is the LSP End Point and if the probe is received over the right incoming interface.
 - Respond based on the validation.

SR Generic Label Sub-TLV



SR SID (20 Bits)

- SR SID
 - Carries 20 bits of Segment ID used for validation.

Procedure

Prefix SID Validation

16000x □ Prefix SID for Rx for Algo 0
16128x □ Prefix SID for Rx for Algo 128



- Initiator (R1) triggers LSP Ping with below SR Generic Label Sub-TLV:
 - For Prefix SID 160008 {SID=160008}
 - For Prefix SID 161288 {SID=161288}
- R8 validates if LSP-EndPoint is self; and if 160008 is assigned locally.

Procedure

Parallel Adj-SID Validation

16000x □ Prefix SID for Rx for Algo 0

16128x □ Prefix SID for Rx for Algo 128

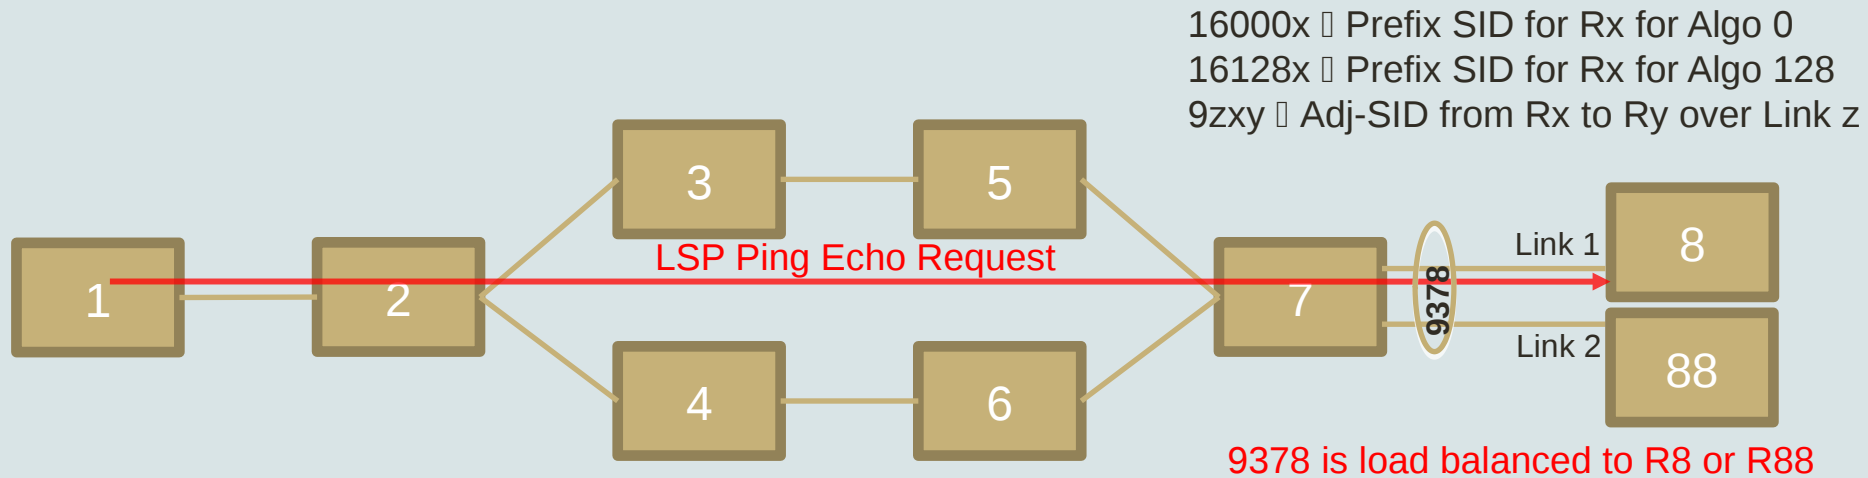
9zxy □ Adj-SID from Rx to Ry over Link z



- Initiator (R1) triggers LSP Ping with below SR Generic Label Sub-TLV:
 - For Parallel Adj SID 9378 {SID=9378}
- R8 validates if LSP-EndPoint == self; and if Inteface- matches interface for 9378.

Procedure

Parallel Adj-SID Validation



- Initiator (R1) triggers LSP Ping with below SR Generic Label Sub-TLV:
 - For Parallel Adj SID 9378 {SID=9378}
- Responder (R8 or R88) validates if SID is assigned by upstream; validates if Interface-I matches interface for 9378.


Procedure

Segment ID to Interface Mapping

16000x □ Prefix SID for Rx for Algo 0

16128x □ Prefix SID for Rx for Algo 128

9zxy □ Adj-SID from Rx to Ry over Link z



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

In a nut shell

- One Target FEC Stack Sub-TLV that covers multiple Segment IDs.
- 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 request
- Thank you!