Network Work group
Internet-Draft
Intended status: Standards Track
Expires: March 23, 2018

N. Kumar, Ed.
C. Pignataro, Ed.
Cisco
G. Swallow
Southend Technical Center
N. Akiya
Big Switch Networks
S. Kini
Individual
M. Chen
Huawei

September 19, 2017

Label Switched Path (LSP) Ping/Traceroute for Segment Routing IGP Prefix and Adjacency SIDs with MPLS Data-plane

draft-ietf-mpls-spring-lsp-ping-07

Abstract

Segment Routing architecture leverages the source routing and tunneling paradigms and can be directly applied to MPLS data plane. A node steers a packet through a controlled set of instructions called segments, by prepending the packet with a Segment Routing header.

The segment assignment and forwarding semantic nature of Segment Routing raises additional consideration for connectivity verification and fault isolation in LSP with Segment Routing architecture. This document illustrates the problem and describe a mechanism to perform LSP Ping and Traceroute on Segment Routing network over MPLS data plane.

Status of This Memo

This Internet-Draft will expire on March 23, 2018.

Copyright Notice

Table of Contents

1. Introduction

[I-D.ietf-spring-segment-routing] introduces and explains Segment Routing architecture that leverages the source routing and tunneling paradigms. A node steers a packet through a controlled set of instructions called segments, by prepending the packet with Segment Routing header. A detailed definition about Segment Routing architecture is available in [I-D.ietf-spring-segment-routing]

As defined in [I-D.ietf-spring-segment-routing] and [I-D.ietf-spring-segment-routing-mpls], the Segment Routing architecture can be directly applied to MPLS data plane in a way that, the Segment identifier (Segment ID) will be of 20-bits size and Segment Routing header is the label stack.

"Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures" [RFC8029] defines a simple and efficient mechanism to detect data plane failures in Label Switched Paths (LSP) by specifying information to be carried in an MPLS "echo request" and "echo reply" for the purposes of fault detection and isolation. Mechanisms for reliably sending the echo reply are defined. The functionality defined in [RFC8029] is modeled after the ping/traceroute paradigm (ICMP echo request [RFC0792]) and is typically referred to as LSP ping and LSP traceroute. [RFC8029] supports hierarchical and stitching LSPs.

Unlike LDP or RSVP which are the other well-known MPLS control plane protocols, the basis of segment ID assignment in Segment Routing architecture is not always on hop-by-hop basis. Depending on the type of segment ID, the assignment can be unique to the node or within a domain.

This nature of Segment Routing raises additional consideration for fault detection and isolation in Segment Routing network. This document illustrates the problem and describes a mechanism to perform LSP Ping and Traceroute on Segment Routing network over MPLS data plane.

2. Requirements notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

3. Terminology

This document uses the terminologies defined in [I-D.ietf-spring-segment-routing], [RFC8029], and so the readers are expected to be familiar with the same.

4. Challenges with Existing mechanism

This document defines sub-TLVs for the Target Forwarding Equivalence Class (FEC) Stack TLV and explains how they can be used to tackle below challenges.

4.1. Path validation in Segment Routing networks

[RFC8029] defines the OAM machinery that helps with fault detection and isolation in MPLS data-plane path with the use of various Target FEC Stack Sub-TLV that are carried in MPLS Echo Request packets and used by the responder for FEC validation. While it is obvious that new Sub-TLVs need to be assigned, the unique nature of Segment Routing architecture raises a need for additional machinery for path validation. This section discusses the challenges as below:

```
                      L1
                  +--------+
                  |   L2   |
                  R3-------R6
                 /           \
                /             \
        R1----R2               R7----R8
                \             /
                 \           /
                  R4-------R5
```

Figure 1: Segment Routing network

The Node Segment IDs for R1, R2, R3, R4, R5, R6, R7 and R8 are 5001, 5002, 5003, 5004, 5005, 5006, 5007, 5008 respectively.

```
9136 --> Adjacency Segment ID from R3 to R6 over link L1.
9236 --> Adjacency Segment ID from R3 to R6 over link L2.
9124 --> Adjacency segment ID from R2 to R4.
9123 --> Adjacency Segment ID from R2 to R3.
```

The forwarding semantic of Adjacency Segment ID is to pop the segment ID and send the packet to a specific neighbor over a specific link. A malfunctioning node may forward packets using Adjacency Segment ID

to incorrect neighbor or over incorrect link. Exposed segment ID (after incorrectly forwarded Adjacency Segment ID) might still allow such packet to reach the intended destination, although the intended strict traversal has been broken.

Assume in above topology, R1 sends traffic with segment stack as {9124, 5008} so that the path taken will be R1-R2-R4-R5-R7-R8. If the Adjacency Segment ID 9124 is misprogrammed in R2 to send the packet to R1 or R3, it will still be delivered to R8 but is not via the expected path.

MPLS traceroute may help with detecting such deviation in above mentioned scenario. However, in a different example, it may not be helpful. For example if R3, due to misprogramming, forwards packet with Adjacency Segment ID 9236 via link L1 while it is expected to be forwarded over Link L2.

5. Segment ID sub-TLV

The format of the following Segment ID sub-TLVs follows the philosophy of Target FEC Stack TLV carrying FECs corresponding to each label in the label stack. When operated with the procedures defined in [RFC8029], this allows LSP ping/traceroute operations to function when Target FEC Stack TLV contains more FECs than received label stack at responder nodes.

Three new sub-TLVs are defined for Target FEC Stack TLVs (Type 1), Reverse-Path Target FEC Stack TLV (Type 16) and Reply Path TLV (Type 21).

sub-Type	Value Field
34	IPv4 IGP-Prefix Segment ID
35	IPv6 IGP-Prefix Segment ID
36	IGP-Adjacency Segment ID

5.1. IPv4 IGP-Prefix Segment ID

The format is as below:

```
    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   |          Reserved             |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

IPv4 Prefix

This field carries the IPv4 prefix to which the Segment ID is assigned. In case of Anycast Segment ID, this field will carry IPv4 Anycast address. If the prefix is shorter than 32 bits, trailing bits SHOULD be set to zero.

Prefix Length

The Prefix Length field is one octet, it gives the length of the prefix in bits (values can be 1 - 32).

Protocol

Set to 1, if the Responder MUST perform FEC validation using OSPF as IGP protocol. Set to 2, if the Responder MUST perform Egress FEC validation using ISIS as IGP protocol. Set to 0, if Responder can use any IGP protocol for Egress FEC validation.

5.2. IPv6 IGP-Prefix Segment ID

The format is as below:

```
    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
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                                                               |
   |                          IPv6 Prefix                          |
   |                                                               |
   |                                                               |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |Prefix Length  |    Protocol   |              Reserved         |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

IPv6 Prefix

This field carries the IPv6 prefix to which the Segment ID is assigned. In case of Anycast Segment ID, this field will carry IPv4 Anycast address. If the prefix is shorter than 128 bits, trailing bits SHOULD be set to zero.

Prefix Length

The Prefix Length field is one octet, it gives the length of the prefix in bits (values can be 1 - 128).

Protocol

Set to 1, if the Responder MUST perform FEC validation using OSPF as IGP protocol. Set to 2, if the Responder MUST perform Egress FEC validation using ISIS as IGP protocol. Set to 0, if Responder can use any IGP protocol for Egress FEC validation.

5.3. IGP-Adjacency Segment ID

This Sub-TLV is applicable for any IGP-Adjacency defined in section 3.5 of [I-D.ietf-spring-segment-routing]. The format is as below:

```
 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   Adj. Type   |    Protocol   |          Reserved             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                Local Interface ID (4 or 16 octets)            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                Remote Interface ID (4 or 16 octets)           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
~                                                               ~
|          Advertising Node Identifier (4 or 6 octets)          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
~                                                               ~
|             Receiving Node Identifier (4 or 6 octets)         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Adj. Type (Adjacency Type)

Set to 1, when the Adjacency Segment is Parallel Adjacency as defined in Section 3.4.1 of [I-D.ietf-spring-segment-routing]. Set to 4, when the Adjacency segment is IPv4 based and is not a parallel adjacency. Set to 6, when the Adjacency segment is IPv6 based and is not a parallel adjacency. Set to 0, when the Adjacency segment is over unnumbered interface.

Protocol

Set to 1, if the Responder MUST perform FEC validation using OSPF as IGP protocol. Set to 2, if the Responder MUST perform Egress FEC validation using ISIS as IGP protocol. Set to 0, if Responder can use any IGP protocol for Egress FEC validation.

Local Interface ID

An identifier that is assigned by local LSR for a link on which Adjacency Segment ID is bound. This field is set to local link address (IPv4 or IPv6). Incase of unnumbered, 32 bit link identifier defined in [RFC4203], [RFC5307] is used. If the Adjacency Segment ID represents parallel adjacencies (Section 3.4.1 of [I-D.ietf-spring-segment-routing]) this field MUST be set to 4 bytes of zero.

Remote Interface ID

An identifier that is assigned by remote LSR for a link on which Adjacency Segment ID is bound. This field is set to remote (downstream neighbor) link address (IPv4 or IPv6). In case of unnumbered, 32 bit link identifier defined in [RFC4203], [RFC5307] is used. If the Adjacency Segment ID represents parallel adjacencies (Section 3.4.1 of [I-D.ietf-spring-segment-routing]) this field MUST be set to 4 bytes of zero.

Advertising Node Identifier

Specifies the advertising node identifier. When Protocol is set to 1, then the 32 rightmost bits represent OSPF Router ID and if protocol is set to 2, this field carries 48 bit ISIS System ID.

Receiving Node Identifier

Specifies the downstream node identifier. When Protocol is set to 1, then the 32 rightmost bits represent OSPF Router ID and if protocol is set to 2, this field carries 48 bit ISIS System ID.

6. Extension to Downstream Detailed Mapping TLV

In an echo reply, the Downstream Mapping TLV [RFC8029] is used to report for each interface over which a FEC could be forwarded. For a FEC, there are multiple protocols that may be used to distribute label mapping. The "Protocol" field of the Downstream Detailed Mapping TLV is used to return the protocol that is used to distribute the label carried in "Downstream Label" field. The following protocols are defined in Section 3.4.1.2 of [RFC8029]:

Protocol #	Signaling Protocol
0	Unknown
1	Static
2	BGP
3	LDP
4	RSVP-TE

With segment routing, OSPF or ISIS can be used for label distribution, this document adds two new protocols as follows:

Protocol #	Signaling Protocol
5	OSPF
6	ISIS

7. Procedures

This section describes aspects of LSP Ping and traceroute operations that require further considerations beyond [RFC8029].

7.1. FECs in Target FEC Stack TLV

When LSP echo request packets are generated by an initiator, FECs carried in Target FEC Stack TLV may need to have deviating contents. This document outlines expected Target FEC Stack TLV construction mechanics by initiator for known scenarios.

Ping

Initiator MUST include FEC(s) corresponding to the destination segment.

Initiator MAY include FECs corresponding to some or all of segments imposed in the label stack by the initiator to communicate the segments traversed.

Traceroute

Initiator MUST initially include FECs corresponding to all of segments imposed in the label stack.

When a received echo reply contains FEC Stack Change TLV with one or more of original segment(s) being popped, initiator MAY remove corresponding FEC(s) from Target FEC Stack TLV in the next (TTL+1) traceroute request as defined in Section 4.6 of [RFC8029].

When a received echo reply does not contain FEC Stack Change TLV, initiator MUST NOT attempt to remove FEC(s) from Target FEC Stack TLV in the next (TTL+1) traceroute request.

As defined in [I-D.ietf-ospf-segment-routing-extensions] and [I-D.ietf-isis-segment-routing-extensions], Prefix SID can be advertised as absolute value, index or as range. In any of these cases, Initiator MUST derive the Prefix mapped to the Prefix SID and use it in IGP-Prefix Segment ID defined in Section 5.1 and 5.2.

7.2. FEC Stack Change sub-TLV

Section 3.4.1.3 of [RFC8029] defines FEC Stack Change sub-TLV that a router must include when the FEC stack changes.

The network node which advertised the Node Segment ID is responsible for generating FEC Stack Change sub-TLV of pop operation for Node Segment ID, regardless of whether PHP is enabled or not.

The network node that is immediate downstream of the node which advertised the Adjacency Segment ID is responsible for generating FEC Stack Change sub-TLV for "POP" operation for Adjacency Segment ID.

7.3. Segment ID POP Operation

The forwarding semantic of Node Segment ID with PHP flag is equivalent to usage of implicit Null in MPLS protocols. Adjacency Segment ID is also similar in a sense that it can be thought of as locally allocated segment that has PHP enabled destined for next hop IGP adjacency node. Procedures described in Section 4.4 of [RFC8029] relies on Stack-D and Stack-R explicitly having Implicit Null value. It may simplify implementations to reuse Implicit Null for Node Segment ID PHP and Adjacency Segment ID cases.

7.4. Segment ID Check

This section updates the procedure defined in Steps 3 to 5 of Section 4.4.1 of [RFC8029]

If the Label-stack-depth is 0 and Target FEC Stack Sub-TLV at FEC-stack-depth is 34 (IPv4 IGP-Prefix Segment ID), the responder should set Best return code to 10, "Mapping for this FEC is not the given label at stack-depth <RSC>" if any below conditions fail:

/* The responder LSR is to check if it is the egress of the IPv4 IGP-Prefix Segment ID described in the Target FEC Stack Sub-TLV, and if the FEC was advertised with the PHP bit set.*/

* Validate that Node Segment ID is advertised for IPv4 Prefix by IGP Protocol{

  + When protocol field in received IPv4 IGP-Prefix Segment ID Sub-TLV is 0, Use any locally enabled IGP protocol.

  + When protocol field in received IPv4 IGP-Prefix Segment ID Sub-TLV is 1, Use OSPF as IGP protocol.

  + When protocol field in received IPv4 IGP-Prefix Segment ID Sub-TLV is 2, Use ISIS as IGP protocol.

  + When protocol field in received IPv4 IGP-Prefix Segment ID Sub-TLV is any other value, it MUST be treated as Protocol value of 0.

  }

* Validate that Node Segment ID is advertised with No-PHP flag {

  + When Protocol is OSPF, NP-flag defined in Section 5 of [I-D.ietf-ospf-segment-routing-extensions] MUST be set to 0.

  + When Protocol is ISIS, P-Flag defined in Section 2.1 of [I-D.ietf-isis-segment-routing-extensions] MUST be set to 0.

  }

If the Label-stack-depth is more than 0 and Target FEC Stack Sub-TLV at FEC-stack-depth is 34 (IPv4 IGP-Prefix Segment ID), the responder is to set Best return code to 10 if any below conditions fail:

* Validate that Node Segment ID is advertised for IPv4 Prefix by IGP Protocol{

  + When protocol field in received IPv4 IGP-Prefix Segment ID Sub-TLV is 0, Use any locally enabled IGP protocol.

  + When protocol field in received IPv4 IGP-Prefix Segment ID Sub-TLV is 1, Use OSPF as IGP protocol.

  + When protocol field in received IPv4 IGP-Prefix Segment ID Sub-TLV is 2, Use ISIS as IGP protocol.

  + When protocol field in received IPv4 IGP-Prefix Segment ID Sub-TLV is any other value, it MUST be treated as Protocol value of 0.

}

If the Label-stack-depth is 0 and Target FEC Sub-TLV at FEC-stack-depth is 35 (IPv6 IGP-Prefix Segment ID), set Best return code to 10 if any below conditions fail:

/* The LSR needs to check if its being a tail-end for the LSP and have the prefix advertised with PHP bit set*/

* Validate that Node Segment ID is advertised for IPv6 Prefix by IGP Protocol{

  + When protocol field in received IPv6 IGP-Prefix Segment ID Sub-TLV is 0, Use any locally enabled IGP protocol.

  + When protocol field in received IPv6 IGP-Prefix Segment ID Sub-TLV is 1, Use OSPF as IGP protocol.

  + When protocol field in received IPv6 IGP-Prefix Segment ID Sub-TLV is 2, Use ISIS as IGP protocol.

  + When protocol field in received IPv4 IGP-Prefix Segment ID Sub-TLV is any other value, it MUST be treated as Protocol value of 0.

  }

* Validate that Node Segment ID is advertised of PHP bit.

If the Label-stack-depth is more than 0 and Target FEC Sub-TLV at FEC-stack-depth is 35 (IPv6 IGP-Prefix Segment ID), set Best return code to 10 if any below conditions fail:

* Validate that Node Segment ID is advertised for IPv4 Prefix by IGP Protocol{

  + When protocol field in received IPv4 IGP-Prefix Segment ID Sub-TLV is 0, Use any locally enabled IGP protocol.

  + When protocol field in received IPv4 IGP-Prefix Segment ID Sub-TLV is 1, Use OSPF as IGP protocol.

  + When protocol field in received IPv4 IGP-Prefix Segment ID Sub-TLV is 2, Use ISIS as IGP protocol.

+ When protocol field in received IPv4 IGP-Prefix Segment ID Sub-TLV is any other value, it MUST be treated as Protocol value of 0.

}

If the Label-stack-depth is 0 and Target FEC sub-TLV at FEC-stack-depth is 36 (Adjacency Segment ID), set Best return code to TBD1 (Section 10.3) if any below conditions fail:

When the Adj. Type is 1 (Parallel Adjacency):

+ Validate that Receiving Node Identifier is local IGP identifier.

+ Validate that Adjacency Segment ID is advertised by Advertising Node Identifier of Protocol in local IGP database {

  - When protocol field in received IPv4 IGP-Prefix Segment ID Sub-TLV is 0, Use any locally enabled IGP protocol.

  - When protocol field in received IPv4 IGP-Prefix Segment ID Sub-TLV is 1, Use OSPF as IGP protocol.

  - When protocol field in received IPv4 IGP-Prefix Segment ID Sub-TLV is 2, Use ISIS as IGP protocol.

  - When protocol field in received IPv4 IGP-Prefix Segment ID Sub-TLV is any other value, it MUST be treated as Protocol value of 0.

  }

When the Adj. Type is 4 or 6 (IGP Adjacency or LAN Adjacency):

+ Validate that Remote Interface ID matches the local identifier of the interface (Interface-I) on which the packet was received.

+ Validate that Receiving Node Identifier is local IGP identifier.

+ Validate that IGP-Adjacency Segment ID is advertised by Advertising Node Identifier of Protocol in local IGP database {

- When protocol field in received IPv4 IGP-Prefix Segment ID Sub-TLV is 0, Use any locally enabled IGP protocol.

- When protocol field in received IPv4 IGP-Prefix Segment ID Sub-TLV is 1, Use OSPF as IGP protocol.

- When protocol field in received IPv4 IGP-Prefix Segment ID Sub-TLV is 2, Use ISIS as IGP protocol.

- When protocol field in received IPv4 IGP-Prefix Segment ID Sub-TLV is any other value, it MUST be treated as Protocol value of 0.

}

7.5. TTL Consideration for traceroute

LSP Traceroute operation can properly traverse every hop of Segment Routing network in Uniform Model described in [RFC3443]. If one or more LSRs employ Short Pipe Model described in [RFC3443], then LSP Traceroute may not be able to properly traverse every hop of Segment Routing network due to absence of TTL copy operation when outer label is popped. Short Pipe being the most commonly used model. The following TTL manipulation technique MAY be used when Short Pipe model is used.

When tracing a LSP according to the procedures in [RFC8029] the TTL is incremented by one in order to trace the path sequentially along the LSP. However when a source routed LSP has to be traced there are as many TTLs as there are labels in the stack. The LSR that initiates the traceroute SHOULD start by setting the TTL to 1 for the tunnel in the LSP's label stack it wants to start the tracing from, the TTL of all outer labels in the stack to the max value, and the TTL of all the inner labels in the stack to zero. Thus a typical start to the traceroute would have a TTL of 1 for the outermost label and all the inner labels would have TTL 0. If the FEC Stack TLV is included it should contain only those for the inner stacked tunnels. The Return Code/Subcode and FEC Stack Change TLV should be used to diagnose the tunnel as described in [RFC8029]. When the tracing of a tunnel in the stack is complete, then the next tunnel in the stack should be traced. The end of a tunnel can be detected from the "Return Code" when it indicates that the responding LSR is an egress for the stack at depth 1. Thus the traceroute procedures in [RFC8029] can be recursively applied to traceroute a source routed LSP.

8. Backward Compatibility with non Segment Routing devices

[I-D.ietf-spring-segment-routing-ldp-interop] describes how Segment Routing operates in network where SR-capable and non-SR-capable nodes coexist. In such networks, there may not be any FEC mapping in the responder when the Initiator is SR-capable while the responder is not (or vice-versa). But this is not different from RSVP and LDP interop scenarios. When LSP Ping is triggered, the responder will set the FEC-return-code to Return 4, "Replying router has no mapping for the FEC at stack-depth".

Similarly when SR-capable node assigns Adj-SID for non-SR-capable node, LSP traceroute may fail as the non-SR-capable node is not aware of "IGP Adjacency Segment ID" sub-TLV and may not reply with FEC Stack change. This may result in any further downstream nodes to reply back with Return-code as 4, "Replying router has no mapping for the FEC at stack-depth".

9. IANA Considerations

9.1. New Target FEC Stack Sub-TLVs

IANA is requested to assign three new Sub-TLVs from "Sub-TLVs for TLV Types 1, 16 and 21" sub-registry from the "Multi-Protocol Label Switching (MPLS) Label Switched Paths (LSPs) Ping Parameters" [IANA-MPLS-LSP-PING] registry.

Sub-Type	Sub-TLV Name	Reference
34	IPv4 IGP-Prefix Segment ID	Section 5.1 of this document
35	IPv6 IGP-Prefix Segment ID	Section 5.2 of this document
36	IGP-Adjacency Segment ID	Section 5.3 of this document

9.2. Protocol in Label Stack Sub-TLV of Downstream Detailed Mapping TLV

IANA is requested to create a new "Protocol" registry under the "Multi-Protocol Label Switching (MPLS) Label Switched Paths (LSPs) Ping Parameters" registry. Code points in the range of 0-250 will be assigned by Standards Action. The range of 251-254 are reserved for experimental use and will not be assigned. The initial entries into the registry will be.

Value	Meaning	Reference
0	Unknown	Section 3.4.1.2 of RFC8029
1	Static	Section 3.4.1.2 of RFC8029
2	BGP	Section 3.4.1.2 of RFC8029
3	LDP	Section 3.4.1.2 of RFC8029
4	RSVP-TE	Section 3.4.1.2 of RFC8029
5	OSPF	Section 6 of this document
6	ISIS	Section 6 of this document
7-250	Unassigned	
251-254	Experimental use	This document
255	Reserved	This document

9.3. Return Code

IANA is requested to assign a new Return Code from the "Multi-Protocol Label Switching (MPLS) Label Switched Paths (LSPs) Ping Parameters" in "Return Codes" Sub-registry.

Value	Meaning	Reference
TBD1	Mapping for this FEC is not associated with the incoming interface	Section 7.4 of this document

Note to the RFC Editor (please remove before publication): IANA has made early allocation for sub-type 34, 35 and 35. The early allocation expires 2017-09-15.

10. Security Considerations

This document defines additional Sub-TLVs and follows the mechanism defined in [RFC8029]. So all the security consideration defined in [RFC8029] will be applicable for this document and in addition it does not impose any security challenges to be considered.

11. Acknowledgement

The authors would like to thank Stefano Previdi, Les Ginsberg, Balaji Rajagopalan, Harish Sitaraman, Curtis Villamizar, Pranjal Dutta, Lizhong Jin, Tom Petch, Victor Ji and Mustapha Aissaoui, Tony Przygienda, Alexander Vainshtein for their review and comments.

The authors wold like to thank Loa Andersson for his comments and recommendation to merge drafts.

12. Contributors

The following are key contributors to this document:

Hannes Gredler, RtBrick, Inc.

Tarek Saad, Cisco Systems, Inc.

Siva Sivabalan, Cisco Systems, Inc.

Balaji Rajagopalan, Juniper Networks

Faisal Iqbal, Cisco Systems, Inc.

13. References

13.1. Normative References

[I-D.ietf-spring-segment-routing]
Filsfils, C., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", draft-ietf-spring-segment-routing-12 (work in progress), June 2017.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC3443] Agarwal, P. and B. Akyol, "Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks", RFC 3443, DOI 10.17487/RFC3443, January 2003, <https://www.rfc-editor.org/info/rfc3443>.

[RFC4203] Kompella, K., Ed. and Y. Rekhter, Ed., "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4203, DOI 10.17487/RFC4203, October 2005, <https://www.rfc-editor.org/info/rfc4203>.

[RFC5307] Kompella, K., Ed. and Y. Rekhter, Ed., "IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 5307, DOI 10.17487/RFC5307, October 2008, <https://www.rfc-editor.org/info/rfc5307>.

[RFC8029] Kompella, K., Swallow, G., Pignataro, C., Ed., Kumar, N., Aldrin, S., and M. Chen, "Detecting Multiprotocol Label Switched (MPLS) Data-Plane Failures", RFC 8029, DOI 10.17487/RFC8029, March 2017, <https://www.rfc-editor.org/info/rfc8029>.

13.2. Informative References

[I-D.ietf-isis-segment-routing-extensions]
Previdi, S., Filsfils, C., Bashandy, A., Gredler, H., Litkowski, S., Decraene, B., and j. jefftant@gmail.com, "IS-IS Extensions for Segment Routing", draft-ietf-isis-segment-routing-extensions-13 (work in progress), June 2017.

[I-D.ietf-ospf-segment-routing-extensions]
Psenak, P., Previdi, S., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF Extensions for Segment Routing", draft-ietf-ospf-segment-routing-extensions-19 (work in progress), August 2017.

[I-D.ietf-spring-segment-routing-ldp-interop]
Filsfils, C., Previdi, S., Bashandy, A., Decraene, B., and S. Litkowski, "Segment Routing interworking with LDP", draft-ietf-spring-segment-routing-ldp-interop-08 (work in progress), June 2017.

[I-D.ietf-spring-segment-routing-mpls]
Filsfils, C., Previdi, S., Bashandy, A., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with MPLS data plane", draft-ietf-spring-segment-routing-mpls-10 (work in progress), June 2017.

[IANA-MPLS-LSP-PING]
IANA, "Multi-Protocol Label Switching (MPLS) Label Switched Paths (LSPs) Ping Parameters", <http://www.iana.org/assignments/mpls-lsp-ping-parameters/mpls-lsp-ping-parameters.xhtml>.

[RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, DOI 10.17487/RFC0792, September 1981, <https://www.rfc-editor.org/info/rfc792>.

Authors' Addresses

Nagendra Kumar (editor)
Cisco Systems, Inc.
7200-12 Kit Creek Road
Research Triangle Park, NC 27709-4987
US

Email: naikumar@cisco.com

Carlos Pignataro (editor)
Cisco Systems, Inc.
7200-11 Kit Creek Road
Research Triangle Park, NC 27709-4987
US

Email: cpignata@cisco.com

George Swallow
Southend Technical Center

Email: swallow.ietf@gmail.com

Nobo Akiya
Big Switch Networks

Email: nobo.akiya.dev@gmail.com

Sriganesh Kini
Individual

Email: sriganeshkini@gmail.com

Mach(Guoyi) Chen
Huawei

Email: mach.chen@huawei.com