Routing area
Internet-Draft
Intended status: Standards Track
Expires: November 7, 2020

S. Hegde
K. Arora
Juniper Networks Inc.
S. Ninan
Individual Contributor
M. Srivastava
Juniper Networks Inc.
N. Kumar
Cisco Systems, Inc.
May 6, 2020

PMS/Head-end based MPLS Ping and Traceroute in Inter-domain SR Networks

draft-ninan-mpls-spring-inter-domain-oam-00

Abstract

Segment Routing (SR) architecture leverages source routing and tunneling paradigms and can be directly applied to the use of a Multiprotocol Label Switching (MPLS) data plane. A network may consist of multiple IGP domains or multiple ASes under the control of same organization. It is useful to have the LSP Ping and traceroute procedures when an SR end-to-end path spans across multiple ASes or domains. This document describes mechanisms to facilitae LSP ping and traceroute in inter-AS/inter-domain SR networks in an efficient manner with simple OAM protocol extension which uses dataplane forwarding alone for sending Echo-Reply.

Requirements Language

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 RFC 2119 [RFC2119].

This Internet-Draft will expire on November 7, 2020.

Copyright Notice

Table of Contents

1. Introduction

```
                       +-----------------+
                       | Controller/PMS |
                       +-----------------+



    |---------AS1----------|         |--------AS2----------|

                         ASBR2-------ASBR3
                        /             \
                       /               \
    PE1------P1------P2                 P3------P4------PE4
                       \               /
                        \             /
                         ASBR1-------ASBR4
```

Figure 1: Inter-AS Segment Routing topology

Many network deployments have built their networks consisting of multiple Autonomous Systems either for ease of operations or as a result of network mergers and acquisitions. Segment Routing can be deployed in such scenarios to provide end to end paths, traversing multiple Autonomous systems(AS). These paths consist of Segment Identifiers(SID) of different type as per [RFC8402].

[I-D.ietf-spring-segment-routing-mpls] specifies the forwarding plane behaviour to allow Segment Routing to operate on top of MPLS data plane. [I-D.ietf-spring-segment-routing-central-epe] describes BGP peering SIDs, which will help in steering packet from one Autonomous system to another. Using above SR capabilities, paths which span across multiple Autonomous systems can be created.

For example Figure 1 describes an inter-AS network scenario consisting of ASes AS1 and AS2. Both AS1 and AS2 are Segment Routing enabled and the EPE links have EPE labels configured and advertised via [I-D.ietf-idr-bgpls-segment-routing-epe]. Controller or head-end can build end-to-end Traffic-Engineered path Node-SIDs, Adjacency-SIDs and EPE-SIDs. It is advantageous for operations to be able to perform LSP ping and traceroute procedures on these inter-AS SR paths. LSP ping/traceroute procedures use ip connectivity for Echo-reply to reach the head-end. In inter-AS networks, ip connectivity may not be there from each router in the path.For example in Figure 1 P3 and P4 may not have ip connectivity for PE1.

[RFC8403] describes mechanisms to carry out the MPLS ping/traceroute from a PMS. It is possible to build GRE tunnels or static routes to each router in the network to get IP connectivity for the reverse path. This mechanism is operationally very heavy and requires PMS to be capable of building huge number of GRE tunnels, which may not be feasible.

It is not possible to carry out LSP ping and Traceroute functionality on these paths to verify basic connectivity and fault isolation using existing LSP ping and Traceroute mechanism([RFC8287] and [RFC8029]). This is because, there exists no IP connectivity to source address of ping packet, which is in a different AS, from the destination of Ping/Traceroute.

[RFC7743] describes a Echo-relay based solution based on advertising a new Relay Node Address Stack TLV containing stack of Echo-relay ip addresses. That mechanism requires the return ping packet to reach the control plane on every relay node.

This document describes a mechanism which is efficient and simple and can be easily deployed in SR networks. This mechanism uses a new Reverse Path Segment List TLV to convey the reverse path. The TLV can either be derived by a smart application/controller which has a full topology view or by the help of intermediate nodes.

2. Inter domain networks with multiple IGPs

```
 |-Domain 1|-------Domain 2-----|--Domain 3-|


 PE1------ABR1--------P--------ABR2------PE4
  \        / \                 /\        /
   --------   -----------------   -------
    BGP-LU         BGP-LU          BGP-LU
```

Figure 2: Inter-domain networks with multiple IGPs

When the network consists of large number of nodes, the nodes are seggregated into multiple IGP domains. The connectivity to the remote PEs can be achieved using BGP-LU [RFC3107] or by stacking the labels for each domain as described in [RFC8604]. It is useful to support mpls ping and traceroute mechanisms for these networks. The procedures described in this document for constructing reverse path TLV and its use in echo-reply is equally applicable to networks consisting of multiple IGP domains that use BGP-LU or label stacking.

3. Reverse Path Segment List TLV

Segment Routing networks statically assign the labels to nodes and PMS/Head-end may know the entire database. The reverse path can be built from PMS/Head-end by stacking segments for the reverse path. A new TLV "Reverse Path Segment List TLV" is defined. Each TLV contains a list of segment sub-TLVs which may be a prefix/adjacency/ binding SID/EPE SID. MPLS Echo -request should contain this TLV, which defines reverse path to reach source from the destination.

The new Reverse Path Segment List TLV is an optional TLV. This TLV is carried in the Echo-Request message. This optional TLV MAY appear in the Echo-request message in any order before or after Target FEC Stack TLV. The Reverse Path Segment List TLV is defined as below. Each MPLS Echo-request SHOULD contain this TLV in inter-AS/inter-domain cases, which will enable remote end(egress/transit routers) to send the reply to source.

In some cases, the head-end may not have complete visibility. In such cases, it can rely on downstream routers to build the reverse path. For this purpose, the TLV is carried in the Echo-Reply message. Section 6 describes one basic idea in this direction.

The Reverse Path Segment List TLV contains different types of segments. The type of segment that the head-end chooses to send in the Reverse Path Segment List TLV is geoverned by local policy. In certain scenarios the head-end may choose to send Type 2/Type 3 segments consisting of IPV4 address and SID. In such cases that node sending the echo-reply MUST derive the MPLS labels from SIDs based on SRGB and encode the echo-reply with MPLS labels.

Some networks may consist of pure IPV4 domains and Pure IPv6 domains. Handling end-to-end MPLS OAM for such networks is out of scope for this document. It is recommended to use dual stack in such cases and use end-to-end IPv6 addresses for MPLS ping and trace route procedures.

3.1. Reverse Path Segment List TLV definition

```
 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 = TBD                     |          Length               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                Segment sub TLV                                |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              ...                              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 3: Reverse Path Segment List TLV

Type: TBD

Length: Length of TLV including TLV header and length of sub TLV.

There can be one or more segment sub-TLVs in a Reverse Path Segment List TLV. The applicable segment types are described in Section 3.1.1. The Segment type in a Reverse Path Segment List TLV MAY be same or different.

3.1.1. Segment sub-TLV

[I-D.ietf-spring-segment-routing-policy] defines various types of segments. These segment types are applicable here. One or more segment sub-TLV can be included. The segment sub-TLVs included MAY be of different types.

Below types of segment sub-TLVs are applicable for the Reverse Path Segment List Tlv.

Type 1: SID only, in the form of MPLS Label

Type 3: IPv4 Node Address with optional SID

Type 4: IPv6 Node Address with optional SID for SR MPLS

3.1.1.1. Type 1: SID only, in the form of MPLS Label

The Type-1 Segment Sub-TLV encodes a single SID in the form of an MPLS label. The format is as follows:

```
     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      |    Length     |     Flags     |   RESERVED    |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |          Label                        | TC  |S|       TTL     |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 4: Type 1 Segment sub-TLV

where:

Type: 1 (to be assigned by IANA from the registry "SR Policy List Sub-TLVs" defined in [I-D.ietf-idr-segment-routing-te-policy]).

Length is 6.

Flags: 1 octet of flags as defined in Section Section 3.1.1.4.

RESERVED: 1 octet of reserved bits. SHOULD be unset on transmission and MUST be ignored on receipt.

Label: 20 bits of label value.

TC: 3 bits of traffic class

S: 1 bit of bottom-of-stack.

TTL: 1 octet of TTL.

The following applies to the Type-1 Segment sub-TLV:

The S bit SHOULD be zero upon transmission, and MUST be ignored upon reception.

If the originator wants the receiver to choose the TC value, it sets the TC field to zero.

If the originator wants the receiver to choose the TTL value, it sets the TTL field to 255.

If the originator wants to recommend a value for these fields, it puts those values in the TC and/or TTL fields.

The receiver MAY override the originator's values for these fields. This would be determined by local policy at the receiver. One

possible policy would be to override the fields only if the fields have the default values specified above.

3.1.1.2. Type 3: IPv4 Node Address with optional SID for SR-MPLS

The Type-3 Segment Sub-TLV encodes an IPv4 node address, SR Algorithm and an optional SID in the form of an MPLS label. The format is as follows:

```
     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      |   Length      |     Flags     |  SR Algorithm |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                 IPv4 Node Address (4 octets)                  |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                SID (optional, 4 octets)                       |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 5: Type 3 Segment sub-TLV

where:

Type: 3 (to be assigned by IANA from the registry "SR Policy List Sub-TLVs" defined in [I-D.ietf-idr-segment-routing-te-policy]).

Length is 6 or 10.

Flags: 1 octet of flags as defined in Section Section 3.1.1.4.

SR Algorithm: 1 octet specifying SR Algorithm as described in section 3.1.1 in [RFC8402], when A-Flag as defined in Section Section 3.1.1.4is present. SR Algorithm is used by SRPM as described in section 4 in [I-D.ietf-spring-segment-routing-policy]. When A-Flag is not encoded, this field SHOULD be unset on transmission and MUST be ignored on receipt.

IPv4 Node Address: a 4 octet IPv4 address representing a node.

SID: 4 octet MPLS label.

The following applies to the Type-3 Segment sub-TLV:

The IPv4 Node Address MUST be present.

The SID is optional and specifies a 4 octet MPLS SID containing label, TC, S and TTL as defined in Section Section 3.1.1.1.

If length is 6, then only the IPv4 Node Address is present.

If length is 10, then the IPv4 Node Address and the MPLS SID are present.

3.1.1.3. Type 4: IPv6 Node Address with optional SID for SR MPLS

The Type-4 Segment Sub-TLV encodes an IPv6 node address, SR Algorithm and an optional SID in the form of an MPLS label. The format is as follows:

```
 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      |   Length      |     Flags     |  SR Algorithm |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
//                IPv6 Node Address (16 octets)                //
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|               SID (optional, 4 octets)                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 6: Type 4 Segment sub-TLV

where:

Type: 4 (to be assigned by IANA from the registry "SR Policy List Sub-TLVs" defined in [I-D.ietf-idr-segment-routing-te-policy]).

Length is 18 or 22.

Flags: 1 octet of flags as defined in Section Section 3.1.1.4.

SR Algorithm: 1 octet specifying SR Algorithm as described in section 3.1.1 in [RFC8402], when A-Flag as defined in Section Section 3.1.1.4 is present. SR Algorithm is used by SRPM as described in section 4 in [I-D.ietf-spring-segment-routing-policy]. When A-Flag is not encoded, this field SHOULD be unset on transmission and MUST be ignored on receipt.

IPv6 Node Address: a 16 octet IPv6 address representing a node.

SID: 4 octet MPLS label.

The following applies to the Type-4 Segment sub-TLV:

The IPv6 Node Address MUST be present.

The SID is optional and specifies a 4 octet MPLS SID containing label, TC, S and TTL as defined in Section Section 3.1.1.1 .

If length is 18, then only the IPv6 Node Address is present.

If length is 22, then the IPv6 Node Address and the MPLS SID are present.

3.1.1.4. Segment Flags

The Segment Types described above MAY contain following flags in the "Flags" field (codes to be assigned by IANA from the registry "SR Policy Segment Flags" defined in[I-D.ietf-idr-segment-routing-te-policy])

```
 0 1 2 3 4 5 6 7
+-+-+-+-+-+-+-+-+
|V|A|           |
+-+-+-+-+-+-+-+-+
```

Figure 7: Flags

where:

V-Flag: This flag is used by SRPM for the purpose of "SID verification" as described in Section 5.1 in [I-D.ietf-spring-segment-routing-policy].

A-Flag: This flag indicates the presence of SR Algorithm id in the "SR Algorithm" field applicable to various Segment Types. SR Algorithm is used by SRPM as described in section 4 in [I-D.ietf-spring-segment-routing-policy].

Unused bits in the Flag octet SHOULD be set to zero upon transmission and MUST be ignored upon receipt.

The following applies to the Segment Flags:

V-Flag is applicable to all Segment Types.

A-Flag is applicable to Segment Types 3, 4 and 9. If A-Flag appears with any other Segment Type, it MUST be ignored.

3.2. SRv6 Dataplane

SRv6 dataplane is not in the scope of this document and will be addressed in a separate document.

4. Detailed Procedures

4.1. Sending an Echo-Request

In the inter-AS scenario when there is no reverse path connectivity, LSP ping initiator MUST add a Reverse Path Segment List TLV in the Echo-request message. The reverse Segment List MUST correspond to the return path from the egress. The Reverse Path Segment List TLV is an ordered list of Segments. The first Segment corresponds to the top Segment in MPLS header that the responder MUST use while sending the Echo-reply.

4.2. Receiving an Echo-Request

When a receiver does not understand the Reverse Path Segment List TLV, it SHOULD silently ignore the TLV and proceed with normal processing as described in [RFC8029].When a Reverse Path Segment List TLV is received, and the responder supports processing it, it MUST use the Segments in Reverse Path Segment List TLV to build the echo-reply. The responder MUST follow the normal FEC validation procedures as described in [RFC8029] and [RFC8287] and this document does not suggest any change to those procedures. When the Echo-reply has to be sent out the Reverse Path Segment List TLV is used to construct the MPLS packet to send out.

4.3. Sending an Echo-Reply

The Echo-Reply message is sent as MPLS packet with a MPLS label stack. The Echo-Reply message MUST be constructed as described in the [RFC8029]. An MPLS packet is constructed with Echo-reply in the payload. The top label MUST be constructed from the first Segment from the Reverse Path Segment List TLV. The remaining labels MUST follow the order from the Reverse Path Segment List TLV. The responder MAY check the reachability of the top label in its own LFIB before sending the Echo-Reply.

5. Detailed Example

An example topology is given in Figure 1 . This will be used in below sections to explain LSP Ping and Traceroute procedures. The PMS/ Head-end has complete view of topology. PE1, P1, P2, ASBR1 and ASBR2 are in AS1. Similarly ASBR3, ASBR4, P3, P4 and PE4 are in AS2.

AS1 and AS2 have Segment Routing enabled. IGPs like OSPF/ISIS are used to flood SIDs in each Autonomous System. The ASBR1, ASBR2, ASBR3, ASBR4 advertise BGP EPE SIDs for the inter-AS links. Topology of AS1 and AS2 are advertised via BGP-LS to the controller/PMS or

Head-end node. The EPE-SIDs are also advertised via BGP-LS as described in [I-D.ietf-idr-bgpls-segment-routing-epe]

The description in the document uses below notations for Segment Identifiers(SIDs).

Node SIDs : N-PE1, N-P1, N-ASBR1 etc.

Adjacency SIDs : Adj-PE1-P1, Adj-P1-P2 etc.

EPE SIDS : EPE-ASBR2-ASBR3, EPE-ASBR1-ASBR4, EPE-ASBR3-ASBR2 etc.

Let us consider a traffic engineered path built from PE1 to PE4 with Segment List stack as below. N-P1, N-ASBR1, EPE-ASBR1-ASBR4, N-PE4 for following procedures. This stack may be programmed by controller/PMS or Head-end router PE1 may have imported the whole topology information from BGP-LS and computed the inter-AS path.

5.1. Procedures for Segment Routing LSP ping

To perform LSP ping procedure on an SR-Path from PE1 to PE4 consisting of label stacks [N-P1,N-ASBR1,EPE-ASBR1-ASBR4, N-PE4], The remote end(PE4) needs IP connectivity to head end(PE1) for the Segment Routing ping to succeed, because Echo-reply needs to travel back to PE1 from PE4. But in typical deployment scenario there will be no ip route from PE4 to PE1 as they belong to different ASes.

PE1 adds Reverse Path from PE4 to PE1 in the MPLS Echo-request using multiple Segments in "Reverse Path Segment List TLV" as defined above. An example reverse path Segment List for PE1 to PE4 for LSP ping is [N-ASBR4, EPE-ASBR4-ASBR1, N-PE1]. An implementation may also build a Reverse Path Segment List consisting of labels to reach its own AS. Once the label stack is popped-off the Echo-reply message will be exposed.The further packet forwarding will be based on ip lookup. An example Reverse Path Segment List for this case could be [N-ASBR4, EPE-ASBR4-ASBR1].

On receiving MPLS Echo-request PE4 first validates FEC in the Echo-request. PE4 then builds label stack to send the response from PE4 to PE1 by copying the labels from "Reverse Path Segment List TLV". PE4 builds the Echo-reply packet with the MPLS label stack constructed and imposes MPLS headers on top of Echo-reply packet and sends out the packet towards PE1. This Segment List stack can successfully steer reply back to Head-end node(PE1).

5.2. Procedures for Segment Routing LSP Traceroute

As described in the procedures for LSP ping, the reverse Segment List may be sent from head-end in which case the LSP Traceroute procedures are similar to LSP ping. The head-end constructs the Reverse Path Segment List TLV and the egress node uses the Reverse Path Segment List to construct the Echo-reply packet header. Head-end/PMS is aware of the reverse path from every node visited in the network and builds the Reverse Path Segment List for every visited node accordingly.

For Example:

For the same traffic engineered path PE1 to PE4 mentioned in above sections, let us assume there is no reverse path available from the nodes ASBR4 to PE1. During the Traceroute procedure, when PE1 has to visit ASBR4, it builds reverse Path Label Stack TLV and includes label to the border-node which has the route to, PE1. In this example the Reverse Path Segment List TLV will contain [EPE-ASBR4-ASBR1]. Further down the traceroute procedure when P3 or P4 node is being visited, PE1 build the Reverse Path Segment List TLV containing [N-ASBR4, EPE-ASBR4-ASBR1]. The Echo-reply will be an MPLS packet with this label stack and will be forwarded to PE1.

6. Building Reverse Path Segment List TLV dynamically

In some cases, the head-end may not have complete visibility of Inter-AS topology. In such cases, it can rely on downstream routers to build the reverse path for mpls traceroute procedures. For this purpose, the Reverse Path Segment List TLV is carried in the Echo-Reply.

6.1. The procedures to build the reverse path

When an ASBR receives an echo-request from another AS, and ASBR is configured to build the Reverse Path dynamically, ASBR MUST build a Reverse Path Segmnet List TLV and add it in echo-reply. ASBR MUST locally decide the outgoing interface for the echo-reply packet. Generally, remote ASBR will choose interface on which the incoming OAM packet was receieved to send the echo-reply out. Reverse Path Segment List TLV is built by adding two segment sub TLVs. The top segment sub TLV consists of the ASBR's Node SID and second segment consists of the EPE SID in the reverse direction to reach the AS from which the OAM packet was received.The type of segment chosen to build Reverse Path Segment List TLV is implementation dependent. In cases where the AS is configured with different SRGBs, the Node SID of the ASBR should be represented using type 3 segment so that all the nodes inside the AS can correctly translate the Node-SID to a label.

Irrespective of which type of segment is included in the Reverse Path Segment List TLV, the responder of echo-request always translates the Reverse Path Segment List TLV to a label stack and builds MPLS header for the the echo-reply packet. This procedure can be applied to an end-to-end path consisting of multiple ASes. Each ASBR at the border adds its Node-SID and EPE-SID on top of existing segments in the Reverse Path Segment List.

6.2. Details with example

Let us consider a traffic engineered path built from PE1 to PE4 with a label stack as below. N-P1, N-ASBR1, EPE-ASBR1-ASBR4, N-PE4 for the following procedures. This traceroute doesn't need any Reverse Path Segment List TLV till it leaves AS1, because IP connectivity will be there to send echo-reply. But this traceroute requires Reverse Path Segment List TLV once it starts probing AS2 routers. According to this procedure, ASBR4 should add Reverse Path Segment List TLV in its echo-reply. ASBR4 should form this Reverse Path Segment List TLV using its own Node SID(N-ASBR4) and EPE SID (EPE-ASRB4-ASBR1) labels. Then PE1 should use this Reverse Path Segment List TLV in subsequent echo-requests. In this example, when the subsequent echo-request reaches P3, it should use this Reverse Path Segment List TLV for sending the echo-reply. The same Reverse Path Segment List TLV is enough for any router in AS2 to send the reply. Because the first label(N-ASBR4) can direct echo-reply to ASBR4 and second one (EPE-ASBR4-ASBR1) to direct echo-reply to AS1. Once echo reply reaches AS1, normal IP forwarding helps it to reach PE1 or the head-end.

7. Security Considerations

TBD

8. IANA Considerations

Multiprotocol Label Switching (MPLS) Label Switched Paths (LSPs) Ping Parameters TLVs Registry

Reverse Path Segment List TLV : TBD

9. Contributors

1.Carlos Pignataro

Cisco Systems, Inc.

cpignata@cisco.com

2. Zafar Ali

Cisco Systems, Inc.

zali@cisco.com

10. Acknowledgments

Thanks to Bruno Decreane for suggesting use of generic Segment sub-TLV.

11. References

11.1. Normative References

[I-D.ietf-idr-segment-routing-te-policy]
Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., Rosen, E., Jain, D., and S. Lin, "Advertising Segment Routing Policies in BGP", draft-ietf-idr-segment-routing-te-policy-08 (work in progress), November 2019.

[I-D.ietf-spring-segment-routing-central-epe]
Filsfils, C., Previdi, S., Dawra, G., Aries, E., and D. Afanasiev, "Segment Routing Centralized BGP Egress Peer Engineering", draft-ietf-spring-segment-routing-central-epe-10 (work in progress), December 2017.

[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>.

[RFC8287] Kumar, N., Ed., Pignataro, C., Ed., Swallow, G., Akiya, N., Kini, S., and M. Chen, "Label Switched Path (LSP) Ping/Traceroute for Segment Routing (SR) IGP-Prefix and IGP-Adjacency Segment Identifiers (SIDs) with MPLS Data Planes", RFC 8287, DOI 10.17487/RFC8287, December 2017, <https://www.rfc-editor.org/info/rfc8287>.

11.2. Informative References

[I-D.ietf-idr-bgpls-segment-routing-epe]
Previdi, S., Talaulikar, K., Filsfils, C., Patel, K., Ray, S., and J. Dong, "BGP-LS extensions for Segment Routing BGP Egress Peer Engineering", draft-ietf-idr-bgpls-segment-routing-epe-19 (work in progress), May 2019.

[I-D.ietf-mpls-interas-lspping]
Nadeau, T. and G. Swallow, "Detecting MPLS Data Plane Failures in Inter-AS and inter-provider Scenarios", draft-ietf-mpls-interas-lspping-00 (work in progress), March 2007.

[I-D.ietf-spring-segment-routing-mpls]
Bashandy, A., Filsfils, C., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with MPLS data plane", draft-ietf-spring-segment-routing-mpls-22 (work in progress), May 2019.

[I-D.ietf-spring-segment-routing-policy]
Filsfils, C., Sivabalan, S., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", draft-ietf-spring-segment-routing-policy-06 (work in progress), December 2019.

[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>.

[RFC3107] Rekhter, Y. and E. Rosen, "Carrying Label Information in BGP-4", RFC 3107, DOI 10.17487/RFC3107, May 2001, <https://www.rfc-editor.org/info/rfc3107>.

[RFC7743] Luo, J., Ed., Jin, L., Ed., Nadeau, T., Ed., and G. Swallow, Ed., "Relayed Echo Reply Mechanism for Label Switched Path (LSP) Ping", RFC 7743, DOI 10.17487/RFC7743, January 2016, <https://www.rfc-editor.org/info/rfc7743>.

[RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <https://www.rfc-editor.org/info/rfc8402>.

[RFC8403] Geib, R., Ed., Filsfils, C., Pignataro, C., Ed., and N. Kumar, "A Scalable and Topology-Aware MPLS Data-Plane Monitoring System", RFC 8403, DOI 10.17487/RFC8403, July 2018, <https://www.rfc-editor.org/info/rfc8403>.

[RFC8604] Filsfils, C., Ed., Previdi, S., Dawra, G., Ed., Henderickx, W., and D. Cooper, "Interconnecting Millions of Endpoints with Segment Routing", RFC 8604, DOI 10.17487/RFC8604, June 2019, <https://www.rfc-editor.org/info/rfc8604>.

Authors' Addresses

Shraddha Hegde
Juniper Networks Inc.
Exora Business Park
Bangalore, KA 560103
India

Email: shraddha@juniper.net

Kapil Arora
Juniper Networks Inc.

Email: kapilaro@juniper.net

Samson Ninan
Individual Contributor

Email: samson.cse@gmail.com

Mukul Srivastava
Juniper Networks Inc.

Email: msri@juniper.net

Nagendra Kumar
Cisco Systems, Inc.

Email: naikumar@cisco.com