

Workgroup: Network Work group

Internet-Draft:

draft-nainar-mpls-spring-lsp-ping-sr-generic-sid-06

Published: 24 November 2021

Intended Status: Standards Track

Expires: 28 May 2022

Authors: N. Nainar, Ed.	C. Pignataro, Ed.	Z. Ali
Cisco	Cisco	Cisco
C. Filsfils	T. Saad	
Cisco	Juniper	

Segment Routing Generic TLV for MPLS Label Switched Path (LSP) Ping/ Traceroute

Abstract

RFC8402 introduces Segment Routing architecture that leverages source routing and tunneling paradigms and can be directly applied to the Multi Protocol Label Switching (MPLS) data plane. A node steers a packet through a controlled set of instructions called segments, by prepending the packet with Segment Routing header. SR architecture defines different types of segments with different forwarding semantics associated. SR can be applied to the MPLS directly and to IPv6 dataplane using a new routing header.

RFC8287 defines the extensions to MPLS LSP Ping and Traceroute for Segment Routing IGP-Prefix and IGP-Adjacency Segment Identifier (SIDs) with an MPLS data plane. Various SIDs are proposed as part of SR architecture with different associated instructions that raises a need to come up with new Target FEC Stack Sub-TLV for each such SIDs.

This document defines a new Target FEC Stack Sub-TLV that is used to validate the instruction associated with any SID.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 28 May 2022.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- 1. [Introduction](#)
 - 1.1. [Challenges with Existing Mechanism](#)
- 2. [Requirements notation](#)
- 3. [Terminology](#)
- 4. [Target FEC Stack sub-TLV for Segment Routing SID](#)
 - 4.1. [Segment Routing Generic Label](#)
 - 4.2. [FEC for Path validation](#)
- 5. [Procedures](#)
 - 5.1. [SID to Interface Mapping](#)
 - 5.2. [Initiator behavior](#)
 - 5.2.1. [SRGL in Target FEC Stack TLV](#)
 - 5.3. [Responder behavior](#)
 - 5.4. [PHP flag behavior](#)
- 6. [IANA Considerations](#)
 - 6.1. [New Target FEC Stack Sub-TLVs](#)
 - 6.2. [Security Considerations](#)
- 7. [Acknowledgement](#)
- 8. [Contributors](#)
- 9. [References](#)
 - 9.1. [Normative References](#)
 - 9.2. [Informative References](#)

[Authors' Addresses](#)

1. Introduction

[[RFC8402](#)] introduces and describes a 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 of the Segment Routing architecture is available in [[RFC8402](#)]

As described in [[RFC8402](#)] and [[I-D.ietf-spring-segment-routing-mpls](#)], the Segment Routing architecture can be directly applied to an MPLS data plane, the Segment identifier (Segment ID) will be of 20-bits size and the Segment Routing header is the label stack.

1.1. Challenges with Existing Mechanism

[[RFC8287](#)] defines the mechanism to perform LSP Ping and Traceroute for Segment Routing with MPLS data plane. [[RFC8287](#)] defines the Target FEC Stack Sub-TLVs for IGP-Prefix Segment ID and IGP-Adjacency Segment ID.

There are various other Segment IDs proposed by different documents that are applicable for SR architecture. [[I-D.ietf-idr-bgp-prefix-sid](#)] defines BGP Prefix Segment ID, [[I-D.ietf-idr-bgppls-segment-routing-epe](#)] defines BGP Peering Segment ID such as Peer Node SID, Peer Adj SID and Peer Set SID. [[I-D.sivabalan-pce-binding-label-sid](#)] defines Path Binding Segment ID. As SR evolves for different usecases, we may see more types of SIDs defined in the future. This raises a need to propose new Target FEC Stack Sub-TLV for each such Segment ID that may need specific or network wide software upgrade to support such new Target FEC Stack Sub-TLVs.

So instead of proposing different Target FEC Stack Sub-TLV for each SID, this document attempt to propose a SR Generic Label Sub-TLV for Target FEC Stack TLV with the procedure to validate the associated instruction.

This document describes the new Target FEC Stack Sub-TLV that carries the SID and the procedure to use LSP Ping and Traceroute using the new sub-tlv to support path validation and fault isolation for any SR Segment IDs. This document neither deprecates any existing Target FEC Stack Sub-TLVs nor precludes defining new Sub-TLVs in the future.

2. Requirements notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)] [RFC 8174](#) [[RFC8174](#)] when and only when, they appear in all capitals, as shown here.

3. Terminology

This document uses the terminologies defined in [[RFC8402](#)], [[RFC8029](#)], readers are expected to be familiar with it.


4. Target FEC Stack sub-TLV for Segment Routing SID

Following the procedure defined in [RFC8029], below defined Target FEC Stack Sub-TLV will be included for each labels in the stack. The below Sub-TLV is defined for Target FEC Stack TLV (Type 1), the Reverse-Path Target FEC Stack TLV (Type 16), and the Reply Path TLV (Type 21).

sub-Type	Value Field
TBD1	Segment Routing Generic Label (SRGL)

4.1. Segment Routing Generic Label

The format of the Sub-TLV is as specified below:



SR SID

Carries 20 bits of Segment ID that is used for validating the instruction.

4.2. FEC for Path validation

In SR architecture, any SID is associated with topology or service instruction. While the topology instruction steers the packet over best path or specific path, the service instruction instructs the type of service to be applied on the packet.

Figure 1: Segment Routing network

or can be based on a query to Path Computation Element (PCE) controller. The mapping can be persistent or on-demand triggered by receiving LSP Ping Request.

5.2. Initiator behavior

This section defines the Target FEC Stack TLV construction mechanism by an initiator when using SR Generic Label Sub-TLV.

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.

5.2.1. SRGL in Target FEC Stack TLV

When the last segment ID in the label stack is IGP Prefix SID, Adj-SID, Binding SID, BGP Prefix SID or BGP Peering SID, set the SR SID field to the Segment ID value advertised by the LSP End Point. When the SID is advertised as index, the Segment ID value MUST be derived based on the index and the SRGB advertised by the LSP End Point.

How the above values are derived is a local implementation matter. It can be manually defined using CLI knob while triggering the LSP Ping Request or can use other mechanisms like querying the local database.

5.3. Responder behavior

Step 4a defined in Section 7.4 of [\[RFC8287\]](#) is updated as below:

If the Label-stack-depth is 0 and Target FEC Stack Sub-TLV at FEC-stack-depth is TBD1 (SRGL) {

- Set the Best-return-code to 10 when the responding node is not the LSP End Point for SR SID.

- Set the Best-return-code to 35, if Interface-I does not match the SID to Interface mapping for the received SR SID.

- set FEC-Status to 1, and return.

}

If the Label-stack-depth is greater than 0 and Target FEC Stack Sub-TLV at FEC-stack-depth is TBD1 (SRGL), {

- If the Label at Label-stack-depth is Imp-null {

- oSet the Best-return-code to 10 when the responding node is not the LSP End Point for the SR SID.

- oSet the Best-return-code to 35, if Interface-I does not match the SID to Interface mapping for the received SR SID.

- oSet FEC-Status to 1, and return.

- }

- Else:

- oSet the Best-return-code to 10 when the index derived from the label at Label-stack-depth is not advertised by LSP End Point.

- oSet FEC-Status to 1, and return.

}

5.4. PHP flag behavior

Section 7.2 of [\[RFC8287\]](#) explains the behavior for FEC stack change for Adjacency Segment ID. The same procedure is applicable for BGP Peering SID as well.

6. IANA Considerations

6.1. 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
-----	-----	-----
TBD1	Segment Routing Generic Label	Section 4.1 of this document

6.2. Security Considerations

This document defines additional MPLS LSP Ping Sub-TLVs and follows the mechanisms defined in [[RFC8029](#)]. All the security considerations defined in [[RFC8029](#)] will be applicable for this document, and in addition, they do not impose any additional security challenges to be considered.

7. Acknowledgement

TBD

8. Contributors

Danial Johari, Cisco Systems

9. References

9.1. Normative References

[I-D.ietf-idr-bgp-prefix-sid]

Previdi, S., Filsfils, C., Lindem, A., Sreekantiah, A., and H. Gredler, "Segment Routing Prefix Segment Identifier Extensions for BGP", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-prefix-sid-27, 26 June 2018, <<https://www.ietf.org/archive/id/draft-ietf-idr-bgp-prefix-sid-27.txt>>.

[I-D.ietf-idr-bgppls-segment-routing-epe]

Previdi, S., Talaulikar, K., Filsfils, C., Patel, K., Ray, S., and J. Dong, "Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing BGP Egress Peer Engineering", Work in Progress, Internet-Draft, draft-ietf-idr-bgppls-segment-routing-epe-19, 16 May 2019, <<https://www.ietf.org/archive/id/draft-ietf-idr-bgppls-segment-routing-epe-19.txt>>.

[I-D.sivabalan-pce-binding-label-sid]

Sivabalan, S., Filsfils, C., Tantsura, J., Hardwick, J., Previdi, S., and C. Li, "Carrying Binding Label/Segment-ID in PCE-based Networks.", Work in Progress, Internet-Draft, draft-sivabalan-pce-binding-label-sid-07, 8 July 2019, <<https://www.ietf.org/archive/id/draft-sivabalan-pce-binding-label-sid-07.txt>>.

[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>>.

[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>>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

[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>>.

[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>>.

9.2. Informative References

[I-D.ietf-spring-segment-routing-mpls]

Bashandy, A., Filsfils, C., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with the MPLS Data Plane", Work in Progress, Internet-Draft, draft-ietf-spring-segment-routing-mpls-22, 1 May 2019, <<https://www.ietf.org/archive/id/draft-ietf-spring-segment-routing-mpls-22.txt>>.

[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>>.

Authors' Addresses

Nagendra Kumar Nainar (editor)
Cisco Systems, Inc.
7200-12 Kit Creek Road
Research Triangle Park, NC 27709-4987
United States of America

Email: naikumar@cisco.com

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

Email: cpignata@cisco.com

Zafar Ali
Cisco Systems, Inc.

Email: zali@cisco.com

Clarence Filsfils
Cisco Systems, Inc.

Email: cfilsfil@cisco.com

Tarek Saad
Juniper Networks

Email: tsaad@juniper.net