

LSR Working Group
Internet-Draft
Intended status: Standards Track
Expires: 28 August 2022

P. Kaneriya
R. Shetty
S. Hegde
R. Bonica
Juniper Networks
24 February 2022

IS-IS Extensions To Support The IPv6 Compressed Routing Header (CRH)
draft-bonica-lsr-crh-isis-extensions-06

Abstract

Source nodes can use the IPv6 Compressed Routing Header (CRH) to steer packets through a specified path. This document defines IS-IS extensions that support the CRH.

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 August 2022.

Copyright Notice

Copyright (c) 2022 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	2
2. Requirements Language	2
3. Advertising The CRH Capability	2
4. Advertising Prefix Segment Identifiers	3
5. Advertising Adjacency Segment Identifiers	4
6. Advertising Adjacency Segment Identifiers Into LANs	5
7. IANA Considerations	7
7.1. The CRH Sub-TLV	7
7.2. Prefix SID Sub-TLV	8
7.3. Adjacency SID Sub-TLV	8
8. Security Considerations	9
9. Acknowledgements	9
10. References	9
10.1. Normative References	9
10.2. Informative References	11
Authors' Addresses	11

1. Introduction

Source nodes can use the IPv6 Compressed Routing Header (CRH) [I-D.bonica-6man-comp-rtg-hdr] to steer packets through a specified path. This document defines IS-IS extensions that support the CRH.

2. Requirements Language

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 BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Advertising The CRH Capability

The Router CAPABILITY TLV [RFC7981] MAY contain exactly one CRH sub-TLV. The CRH sub-TLV indicates that the advertising node can process the CRH.

The CRH sub-TLV MAY contain sub-sub-TLVs. No sub-sub-TLVs are currently defined.

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	Max CRH Len	Reserved
+	+	+	+

Figure 1: CRH Sub-TLV

Figure 1 depicts the CRH sub-TLV. The CRH sub-TLV contains the following fields:

- * Type: 8 bits. CRH (value TBD by IANA. Suggested value is 30.)
- * Length: 8 bits. Length of TLV data excluding the TLV header. MUST be equal to 2 plus the length of sub-sub-TLVs (if any).
- * Max CRH Len: 8 bits. Maximum CRH length supported by the advertising node, measured in 8-octet units, not including the first 8 octets. See Note 1.
- * Reserved: 8 bits. SHOULD be set to zero by sender. MUST be ignored by receiver.

Note 1: According to [RFC8200], all IPv6 Routings header include a "Hdr Ext Len" field. That field specifies the length of the Routing header in 8-octet units, not including the first 8 octets. The same unit of measure was chosen for the "Max CRH Len" field in the CRH sub-TLV.

4. Advertising Prefix Segment Identifiers

The following TLVs MAY contain one or more Prefix SID sub-TLVs:

- * TLV-236 (IPv6 IP Reachability) [RFC5308].
- * TLV-237 (Multitopology IPv6 IP Reachability) [RFC5120].

The Prefix SID sub-TLV is valid only when its parent TLV specifies a prefix length of 128. In this case, it binds the SID that it contains to the prefix (i.e., IPv6 address) that its parent TLV contains. This information is used to construct the mapping table described in [I-D.bonica-6man-comp-rtg-hdr].

When the parent TLV is propagated across level boundaries, the Prefix SID sub-TLV SHOULD be kept.




Figure 2: Prefix SID sub-TLV

Figure 2 depicts the Prefix SID sub-TLV. It contains the following fields:

- * Type: 8 bits. Prefix SID sub-TLV (Value TBD by IANA. Suggested value is 33.)
- * Length: 8 bits. Length of TLV data excluding the TLV header, measured in bytes.
- * Reserved: 16 bits. SHOULD be set to zero by the sender. MUST be ignored by the receiver.
- * SID - Variable length. Segment Identifier.

5. Advertising Adjacency Segment Identifiers

The following TLVs can contain one or more Adjacency SID sub-TLVs:

- * TLV-22 (Extended IS reachability) [RFC5305]
- * TLV-222 (Multitopology IS) [RFC5120]
- * TLV-23 (IS Neighbor Attribute) [RFC5311]
- * TLV-223 (Multitopology IS Neighbor Attribute) [RFC5311]
- * TLV-141 (inter-AS reachability information) [RFC5316]

The Adjacency SID sub-TLV is valid only when its parent TLV also contains an IPv6 Neighbor Address sub-TLVs [RFC6119]. In this case, the SID contained by the Adjacency SID sub-TLV is bound to the IPv6 address contained by the IPv6 Neighbor Address sub-TLV. This information is used to construct the mapping table described in [I-D.bonica-6man-comp-rtg-hdr].




Figure 3: Adjacency SID Sub-TLV

Figure 3 depicts the Adjacency SID sub-TLV. It contains the following fields:

- * Type: 8 bits. Adjacency SID sub-TLV (Value TBD by IANA. Suggested value is 45.)
- * Length: 8 bits. Length of TLV data excluding the TLV header, measured in bytes.
- * Flags: 8 bits. See below.
- * Weight: 8 bits. The value represents the SID weight for the purpose of load balancing.
- * SID - Variable length. Segment Identifier.

```

      0 1 2 3 4 5 6 7
      +--+--+--+--+--+--+
      |B|S|P| Reserved|
      +--+--+--+--+--+--+

```

Figure 4: Adjacency SID Sub-TLV Flags

Figure 4 depicts Adjacency SID Sub-TLV flags. They include the following:

- * B-Flag: Backup flag. If set, the SID is eligible for protection.
- * S-Flag: Set flag. When set, the S-Flag indicates that the SID refers to a set of adjacencies (and therefore MAY be assigned to other adjacencies as well).
- * P-Flag: Persistent flag. When set, the P-Flag indicates that the SID is persistently allocated, i.e., the SID value remains consistent across router restart and/or interface flap.)

6. Advertising Adjacency Segment Identifiers Into LANs

In LAN subnetworks, the Designated Intermediate System (DIS) is elected and originates the Pseudonode-LSP (PN-LSP) including all neighbors of the DIS.

When the CRH is used, each router in the LAN MAY advertise its Adjacency SIDs of each of its neighbors. Since, on LANs, each router only advertises one adjacency to the DIS (and doesn't advertise any other adjacency), each router advertises the set of Adjacency SIDs (for each of its neighbors) inside a newly defined sub-TLV part of the TLV advertising the adjacency to the DIS (e.g.: TLV-22).

The following TLVs can contain one or more LAN Adjacency SID sub-TLVs:

- * TLV-22 (Extended IS reachability) [RFC5305]
- * TLV-222 (Multitopology IS) [RFC5120]
- * TLV-23 (IS Neighbor Attribute) [RFC5311]
- * TLV-223 (Multitopology IS Neighbor Attribute) [RFC5311]

The LAN Adjacency SID sub-TLV binds an IPv6 address to a SID. The sub-TLV contains both the IPv6 address and the SID. This information is used to construct the mapping table described in [I-D.bonica-6man-comp-rtg-hdr].

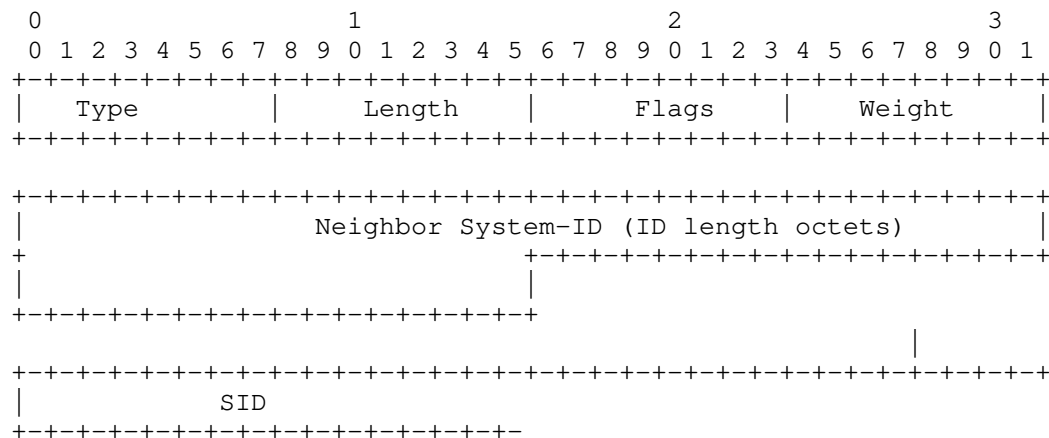


Figure 5: LAN Strictly Routed SID Sub-TLV

Figure 5 depicts the Adjacency SID sub-TLV. It contains the following fields:

- * Type: 8 bits. Adjacency SID sub-TLV (Value TBD by IANA. Suggested value is 46.)
- * Length: 8 bits. Length of TLV data excluding the TLV header, measured in bytes.
- * Flags: 8 bits. See below.
- * Weight: 8 bits. The value represents the SID weight for the purpose of load balancing.

- * Neighbor System-ID: 6 bytes. IS-IS System-ID of length "ID Length" as defined in [ISO10589].
- * SID - Variable length. Segment Identifier.

```

    0 1 2 3 4 5 6 7
    +-+-+-+-+-+-+-+-+
    |B|S|P| Reserved|
    +-+-+-+-+-+-+-+-+

```

Figure 6: Adjacency SID Sub-TLV Flags

Figure 6 depicts Adjacency SID Sub-TLV flags. They include the following:

- * B-Flag: Backup flag. If set, the SID is eligible for protection.
- * S-Flag: Set flag. When set, the S-Flag indicates that the SID refers to a set of adjacencies (and therefore MAY be assigned to other adjacencies as well).
- * P-Flag: Persistent flag. When set, the P-Flag indicates that the SID is persistently allocated, i.e., the SID value remains consistent across router restart and/or interface flap.)

7. IANA Considerations

7.1. The CRH Sub-TLV

IANA is requested to add a new sub-TLV in the Sub-TLVs for TLV 242 (IS-IS Router CAPABILITY TLV) Registry [capreg].

- * Value - TBD by IANA. (Suggested value is 30).
- * Description - CRH

This document requests the creation of a new IANA managed registry for sub-sub-TLVs of the CRH sub-TLV. The registration procedure is "Expert Review" as defined in [RFC7370]. Suggested registry name is "sub-sub-TLVs for CRH sub-TLV". No sub-sub-TLVs are defined by this document except for the reserved value.

- * 0 - Reserved
- * 1 - 255 Unassigned

7.2. Prefix SID Sub-TLV

IANA is requested to add a new entry in the Sub-TLVs for TLVs 135, 235, 236, and 237 (Extended IP reachability, MT IP. Reach, IPv6 IP. Reach, and MT IPv6 IP. Reach TLVs) Registry [loosereg].

- * Value - TBD by IANA. (Suggested value is 33)
- * Description - Prefix SID
- * 135 - N
- * 136 - N
- * 236 - Y
- * 237 - Y
- * Reference - This document.

7.3. Adjacency SID Sub-TLV

IANA is requested to add the following entries in the Sub-TLVs for TLVs 22, 23, 25, 141, 222, and 223 (Extended IS reachability, IS Neighbor Attribute, L2 Bundle Member Attributes, inter-AS reachability information, MT-ISN, and MT IS Neighbor Attribute TLVs) Registry [strictreg].

The first entry follows:

- * Value - TBD by IANA (Suggested value is 45).
- * Description - Adjacency SID
- * 22 - Y
- * 23 - Y
- * 25 - N
- * 141 - Y
- * 222 - Y
- * 223 - Y
- * Reference - This document.

The second entry follows:

- * Value - TBD by IANA (Suggested value is 46)
- * Description - LAN Adjacency SID
- * 22 - Y
- * 23 - Y
- * 25 - N
- * 141 - N
- * 222 - Y
- * 223 - Y
- * Reference - This document.

8. Security Considerations

Security concerns for IS-IS are addressed in [ISO10589], [RFC5304], and [RFC5310].

9. Acknowledgements

Thanks to Ram Santhanakrishnan for his comments on this document.

10. References

10.1. Normative References

[I-D.bonica-6man-comp-rtg-hdr]

Bonica, R., Kamite, Y., Alston, A., Henriques, D., and L. Jalil, "The IPv6 Compact Routing Header (CRH)", Work in Progress, Internet-Draft, draft-bonica-6man-comp-rtg-hdr-27, 15 November 2021, <<https://www.ietf.org/archive/id/draft-bonica-6man-comp-rtg-hdr-27.txt>>.

[ISO10589] IANA, "Intermediate system to Intermediate system routing information exchange protocol for use in conjunction with the Protocol for providing the Connectionless-mode Network Service (ISO 8473)", August 1987, <ISO/IEC 10589:2002>.

- [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>>.
- [RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", RFC 5120, DOI 10.17487/RFC5120, February 2008, <<https://www.rfc-editor.org/info/rfc5120>>.
- [RFC5304] Li, T. and R. Atkinson, "IS-IS Cryptographic Authentication", RFC 5304, DOI 10.17487/RFC5304, October 2008, <<https://www.rfc-editor.org/info/rfc5304>>.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <<https://www.rfc-editor.org/info/rfc5305>>.
- [RFC5308] Hopps, C., "Routing IPv6 with IS-IS", RFC 5308, DOI 10.17487/RFC5308, October 2008, <<https://www.rfc-editor.org/info/rfc5308>>.
- [RFC5310] Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", RFC 5310, DOI 10.17487/RFC5310, February 2009, <<https://www.rfc-editor.org/info/rfc5310>>.
- [RFC5311] McPherson, D., Ed., Ginsberg, L., Previdi, S., and M. Shand, "Simplified Extension of Link State PDU (LSP) Space for IS-IS", RFC 5311, DOI 10.17487/RFC5311, February 2009, <<https://www.rfc-editor.org/info/rfc5311>>.
- [RFC5316] Chen, M., Zhang, R., and X. Duan, "ISIS Extensions in Support of Inter-Autonomous System (AS) MPLS and GMPLS Traffic Engineering", RFC 5316, DOI 10.17487/RFC5316, December 2008, <<https://www.rfc-editor.org/info/rfc5316>>.
- [RFC6119] Harrison, J., Berger, J., and M. Bartlett, "IPv6 Traffic Engineering in IS-IS", RFC 6119, DOI 10.17487/RFC6119, February 2011, <<https://www.rfc-editor.org/info/rfc6119>>.
- [RFC7370] Ginsberg, L., "Updates to the IS-IS TLV Codepoints Registry", RFC 7370, DOI 10.17487/RFC7370, September 2014, <<https://www.rfc-editor.org/info/rfc7370>>.

- [RFC7981] Ginsberg, L., Previdi, S., and M. Chen, "IS-IS Extensions for Advertising Router Information", RFC 7981, DOI 10.17487/RFC7981, October 2016, <<https://www.rfc-editor.org/info/rfc7981>>.
- [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>>.
- [RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.

10.2. Informative References

- [capreg] IANA, "Sub-TLVs for TLV 242 (IS-IS Router CAPABILITY TLV)", August 1987, <<https://www.iana.org/assignments/isis-tlv-codepoints/isis-tlv-codepoints.xhtml#isis-tlv-codepoints-242>>.
- [loosereg] IANA, "Sub-TLVs for TLVs 135, 235, 236, and 237 (Extended IP reachability, MT IP. Reach, IPv6 IP. Reach, and MT IPv6 IP. Reach TLVs)", August 1987, <<https://www.iana.org/assignments/isis-tlv-codepoints/isis-tlv-codepoints.xhtml#isis-tlv-codepoints-135-235-236-237>>.
- [strictreg] IANA, "Sub-TLVs for TLVs 22, 23, 25, 141, 222, and 223 (Extended IS reachability, IS Neighbor Attribute, L2 Bundle Member Attributes, inter-AS reachability information, MT-ISN, and MT IS Neighbor Attribute TLVs)", August 1987, <<https://www.iana.org/assignments/isis-tlv-codepoints/isis-tlv-codepoints.xhtml#isis-tlv-codepoints-22-23-25-141-222-223>>.

Authors' Addresses

Parag Kaneriya
Juniper Networks
Elnath-Exora Business Park Survey
Bangalore 560103
Karnataka
India
Email: pkaneria@juniper.net

Rejesh Shetty
Juniper Networks
Elnath-Exora Business Park Survey
Bangalore 560103
Karnataka
India
Email: mrajesh@juniper.net

Shraddha Hegde
Juniper Networks
Elnath-Exora Business Park Survey
Bangalore 560103
Karnataka
India
Email: shraddha@juniper.net

Ron Bonica
Juniper Networks
2251 Corporate Park Drive
Herndon, Virginia 20171
United States of America
Email: rbonica@juniper.net