

Inter-Domain Routing Group
Internet-Draft
Intended status: Standards Track
Expires: May 21, 2020

A. Alston
D. Henriques
Liquid Telecom
R. Bonica
Juniper Networks
November 18, 2019

**BGP Extensions for IPv6 Compressed Routing Header (CRH)
draft-alston-idr-crh-bgp-signalling-00**

Abstract

This document describes a new BGP extension for signalling the mapping between Segment Identifiers (SID's), as used by a SRv6 Compressed Routing Header (CRH) and the IPv6 Addresses they represent. The extension defines both a new optional BGP attribute to signal the Maximum SID Value (MSV) and a new Sub-Address Family (SAFI) of the IPv6 Address family.

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 May 21, 2020.

Copyright Notice

Copyright (c) 2019 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

- [1. Overview](#) [2](#)
- [2. Requirements Language](#) [3](#)
- [3. SID Signalling](#) [3](#)
 - [3.1. BGP Attribute](#) [3](#)
 - [3.2. MP Reach Attribute](#) [3](#)
 - [3.3. NLRI Format](#) [4](#)
- [4. IANA Considerations](#) [5](#)
 - [4.1. BGP Path Attributes](#) [5](#)
 - 4.2. Subsequent Address Family Identifiers (SAFI) Parameters . 5
- [5. Security Considerations](#) [5](#)
- [6. Acknowledgements](#) [5](#)
- [7. References](#) [5](#)
 - [7.1. Normative References](#) [5](#)
 - [7.2. Informative References](#) [6](#)
- Authors' Addresses [7](#)

1. Overview

The SRm6 Compressed Routing Header uses an ordered sequence of segment identifiers (SID) to specify the end to end path a packet should follow through the network. This allows for much smaller header sizes than found in the SRH (Segment Routing Header), which utilizes an ordered sequence of 128 bit IPv6 address to achieve the same goal. In addition, this method prevents the overloading of the IPv6 address space.

This results in the need to signal the mapping between the SIDs used in the CRH and the IPv6 addresses they represent. While such signalling can be achieved through IGP extensions [[I-D.bonica-lsr-crh-isis-extensions](#)] in a single network domain, circumstances may dictate that the SID to address mapping be signaled both to systems that do not partake in the IGP used within that network domain, and between autonomous systems.

It is envisaged that such signalling will be required to signal, among other things, deep packet inspection systems and flow analysis systems that need the ability to see the full path a packet is traversing, while at the same time not necessarily partaking in the IGP which would normally be used for such signalling. This also allows signalling of SID to Address mapping in environments that do not run an IGP capable of such signalling.

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. SID Signalling

3.1. BGP Attribute

This document describes a new BGP extension for signalling the mapping between Segment Identifiers (SIDs) as used by an SRm6 Compressed Routing Header (CRH) and the IPv6 Addresses they represent.

The extension defines both a new optional transitive attribute to signal the SID size in Octets and a new Subsequent Address Family Identifier (SAFI) of the IPv6 Address family.

The document defines a new BGP attribute which signals the maximum size in octets of a CRH Segment Identifier (SID). The attribute MUST be included in any BGP session carrying the CRH signalling SAFI. Additionally a new attribute BGP Path Attribute code will need to be assigned by IANA.

The attribute is comprised of a single octet, which signals the maximum length in octets of the SIDs found in the Network Layer Reachability Information (NLRI) section of the relevant MP-BGP Attribute.

Since SIDs in the context of a compressed routing header can be either 16bit or 32bit, the attribute value MUST be either 2 or 4 and the attribute MUST be ignored if this is not the case. In the event of the CRH signalling attribute not being present in a BGP Update Message, any BGP Updates containing the CRH SAFI MUST be considered malformed and should be ignored.

3.2. MP Reach Attribute

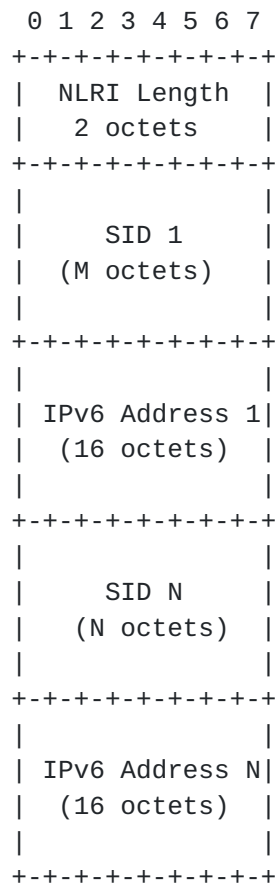
The format of the MP Reach Attribute utilized by the CRH SAFI is as per [[RFC4760](#)]. The AFI MUST be set to 2, and the SAFI is currently TBD (see [Section 4.2](#)).

The Nexthop field of the attribute contains no significance and is maintained purely for compatibility. For standardization purposes we maintain a next-hop length field of 16 which contains an arbitrary

value. Implementations MAY at their discretion use the originating router ID, or 128 bit mapped equivalent, in this field for simple identification of mapping source

3.3. NLRI Format

The format of the NLRI contained within the MP Reach Attribute is comprised of a 16bit Length (2 octets) field, followed by a series of repeating tuples. The length in octets of the first element of each tuple is determined by the SID Length specified in the CRH signalling attribute ([Section 3.1](#)). The second element of the tuple is an IPv6 address and MUST be 16 octets in length. The length of the NLRI can be calculated as (SID Length+16)*N where N is the number of tuples contained within the NLRI.



The CRH SAFI uses a new NLRI defined as follows: where M MUST be the size in octets of the MSV signalled via BGP upon session establishment and N specifies a given number of SID/IPv6 Address pairings.

4. IANA Considerations

This document defines new Sub-TLVs in the following existing registries:

- o BGP Path Attributes
- o Subsequent Address Family Identifiers (SAFI) Parameters

4.1. BGP Path Attributes

A new SAFI in the IANA registry for "Subsequent Address Family Identifiers (SAFI) Parameters" will be required:

| Codepoint | Description | Reference |
|-----------|-------------|---------------|
| TBD | SRm6 CRH | This document |

4.2. Subsequent Address Family Identifiers (SAFI) Parameters

A new SAFI in the IANA registry for "Subsequent Address Family Identifiers (SAFI) Parameters" will be required:

| Codepoint | Description | Reference |
|-----------|--------------------------|---------------|
| TBD | SRm6 CRH Signalling SAFI | This document |

5. Security Considerations

SRm6 CRH BGP Signalling is envisioned to be run within a trusted domain.

Further aspects of security are TBD.

6. Acknowledgements

The authors wish to acknowledge Ben Roberts for his support.

7. References

7.1. Normative References

[I-D.bonica-6man-comp-rtg-hdr]
 Bonica, R., Kamite, Y., Niwa, T., Alston, A., Henriques, D., Jalil, L., So, N., Xu, F., Chen, G., Zhu, Y., Yang, G., and Y. Zhou, "The IPv6 Compressed Routing Header (CRH)", [draft-bonica-6man-comp-rtg-hdr-08](#) (work in progress), October 2019.

[I-D.bonica-lsr-crh-isis-extensions]

Kaneriya, P., Shetty, R., Hegde, S., and R. Bonica, "IS-IS Extensions To Support The IPv6 Compressed Routing Header (CRH)", [draft-bonica-lsr-crh-isis-extensions-00](#) (work in progress), May 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>>.

[RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), DOI 10.17487/RFC4271, January 2006, <<https://www.rfc-editor.org/info/rfc4271>>.

[RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", [RFC 4760](#), DOI 10.17487/RFC4760, January 2007, <<https://www.rfc-editor.org/info/rfc4760>>.

[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>>.

[7.2.](#) Informative References

[I-D.alston-spring-crh-bgp-signalling]

Alston, A., Henriques, D., and R. Bonica, "BGP Extensions for IPv6 Compressed Routing Header (CRH)", [draft-alston-spring-crh-bgp-signalling-01](#) (work in progress), July 2019.

[I-D.ietf-6man-segment-routing-header]

Filsfils, C., Dukes, D., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", [draft-ietf-6man-segment-routing-header-26](#) (work in progress), October 2019.

[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-srv6-network-programming]

Filsfils, C., Camarillo, P., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "SRv6 Network Programming", [draft-ietf-spring-srv6-network-programming-05](#) (work in progress), October 2019.

[RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", [RFC 3031](#), DOI 10.17487/RFC3031, January 2001, <<https://www.rfc-editor.org/info/rfc3031>>.

Authors' Addresses

Andrew Alston
Liquid Telecom
Nairobi
Kenya

Email: Andrew.Alston@liquidtelecom.com

Daniam Henriques
Liquid Telecom
Johannesburg
South Africa

Email: daniam.henriques@liquidtelecom.com

Ron Bonica
Juniper Networks
Herndon, Virginia 20171
USA

Email: rbonica@juniper.net

