

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: May 7, 2020

Z. Li
L. Li
Huawei
H. Chen
Futurewei
C. Loibl
Next Layer Communications
Y. Zhu
China Telecom
L. Liu
Fujitsu
X. Liu
Volta Networks
November 4, 2019

BGP Flow Specification for SRv6
draft-li-idr-flowspec-srv6-01

Abstract

This draft proposes extensions to BGP to distribute traffic Flow Specifications for SRv6 for filtering SRv6 packets that match a sequence of conditions.

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 [[RFC2119](#)].

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 7, 2020.

Internet-Draft

BGP Flow Specification for SRv6

November 2019

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.	Introduction	2
2.	Definitions and Acronyms	3
3.	The Flow Specification Encoding for SRv6	4
3.1.	Type TBD1 - Whole SID	4
3.2.	Type TBD2 - Some bits of SID	5
4.	Security Considerations	6
5.	IANA	6
6.	Acknowledgments	7
7.	References	7
7.1.	Normative References	7
7.2.	Informative References	7
	Authors' Addresses	8

[1.](#) Introduction

BGP Flow Specification (BGP-FS) [[RFC5575](#)] defines a new BGP NLRI to distribute a flow specification, which is an n-tuple comprising a sequence of matching criteria that can be applied to IP traffic. It defines a minimum set of filtering actions as BGP extended community values [[RFC4360](#)] that modify the traffic packet and forwards/drops the packet. [[I-D.ietf-idr-rfc5575bis](#)] gives more details about them. The NLRI (AFI=1, SAFI=133) is for IPv4 unicast flow specification and NLRI (AFI=1, SAFI=134) is for BGP/MPLS VPNv4 flow specification. [[I-D.ietf-idr-flow-spec-v6](#)] redefines the [[RFC5575](#)] SAFIs to make them AFI specific and applicable to both IPv4 and IPv6 applications. [[I-D.ietf-idr-flowspec-l2vpn](#)] extends the flow-spec rules for layer 2

Ethernet packets.

Segment Routing (SR) for unicast traffic has been proposed to cope with the usecases in traffic engineering, fast re-reroute, service chain, etc. SR architecture can be implemented over an IPv6 data

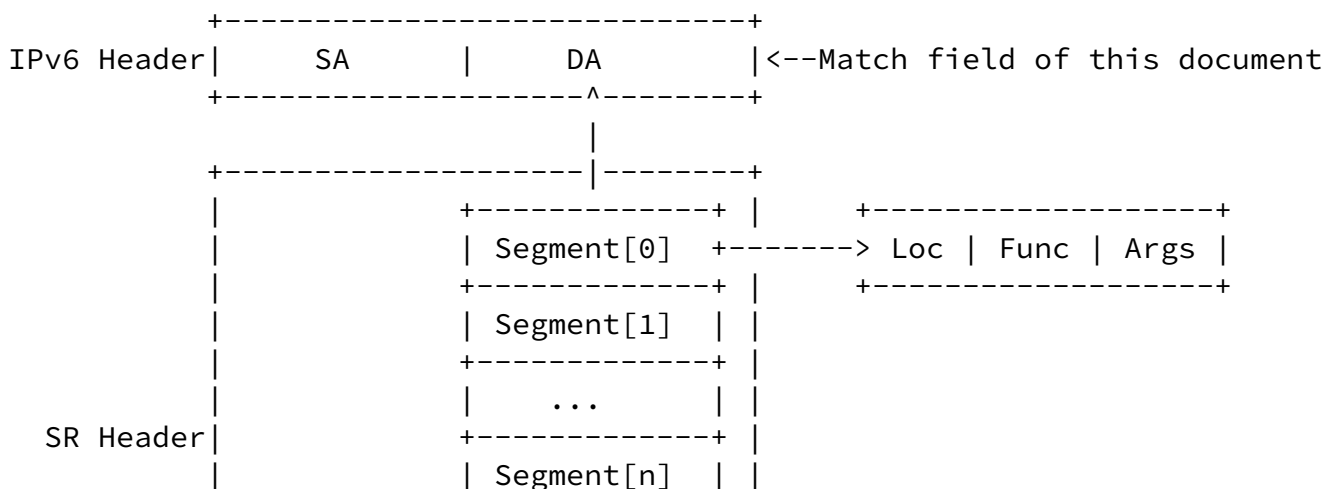
plane using a new type of Segment Routing Header (SRH) [[I-D.ietf-6man-segment-routing-header](#)]. SRv6 Network Programming [[I-D.filsfils-spring-srv6-network-programming](#)] defines the SRv6 network programming concept and its most basic functions. SRv6 SID may have the form of LOC:FUNCT:ARGS::.


LOC: Each operator is free to use the locator length it chooses. Most often the LOC part of the SID is routable and leads to the node which instantiates that SID.

FUNCT: The FUNCT part of the SID is an opaque identification of a local function bound to the SID. (e.g. End: Endpoint, End.X, End.T, End.DX2 etc.).

ARGS: A function may require additional arguments that would be placed immediately after the FUNCT.

This document specifies a couple of new BGP-FS component types to support Segment Routing over IPv6 data plane (SRv6) filtering. The match field is destination address of IPv6 header, but it's a SID copy from SRH rather than a traditional IPv6 address (refer to the figure below).





2. Definitions and Acronyms

- o FS: Flow Specification
- o SR: Segment Routing

- o SRv6: IPv6 Segment Routing, SRv6 is a method of forwarding IPv6 packets on the network based on the concept of source routing.
- o SID: Segment Identifier
- o BSID: Binding SID

3. The Flow Specification Encoding for SRv6


The Flow Specification NLRI-type consists of several optional components, each of which begins with a type field (1 octet) followed by a variable length parameter. [\[RFC5575\]](#) defines 12 component types for IPv4. IPv6 NLRI component types are described in [\[I-D.ietf-idr-flow-spec-v6\]](#). This document defines two new component types for SRv6.

3.1. Type TBD1 - Whole SID

Encoding: <type (1 octet), length(1 octet), [op, value]+>

Contains a list of {operator, value} pairs that are used to match the SID/binding SID or a range of whole SID.

The operator byte is encoded as:



Where:

e - end-of-list bit. Set in the last {op, value} pair in the sequence.

a - AND bit. If unset, the previous term is logically ORed with the current one. If set, the operation is a logical AND. It should be unset in the first operator byte of a sequence. The AND operator has higher priority than OR for the purposes of evaluating logical expressions.

lt - less than comparison between data and value.

gt - greater than comparison between data and value.

eq - equality between data and value.

The bits lt, gt, and eq can be combined to match the SID or a range of SID (e.g. less than SID1 and greater than SID2).

The value field is encoded as:

```

      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
+-----+
~                SID(128bits)                ~
+-----+
```

The format of SID is described in [\[I-D.ietf-6man-segment-routing-header\]](#) and [\[I-D.filsfils-spring-srv6-network-programming\]](#)

[3.2.](#) Type TBD2 - Some bits of SID

For some scenarios route policy with the whole 128 bits SID matching is too long and not necessary.


[\[I-D.filsfils-spring-srv6-network-programming\]](#) defines the format of SID is LOC:FUNCT:ARGS::. In some scenarios, traffic packets can just match Locator, Function ID, Argument or some combinations of these

different fields rather than whole 128 bits SID. The new component type TBD2 defined below is for matching some bits of SID.

Encoding: <type (1 octet), length(1 octet), [op, value]+>

Contains a list of {operator, value} pairs that are used to match some bits of SID.

The operator byte is encoded as:



Where:

e - end-of-list bit. Set in the last {op, value} pair in the list.

a - AND bit. If unset, the previous term is logically ORed with the current one. If set, the operation is a logical AND. It should be unset in the first operator byte of a sequence. The AND operator has higher priority than OR for the purposes of evaluating logical expressions.

type:


0000 : SID's LOC bits

0001 : SID's FUNCT bits

0010 : SID's LOC:FUNCT bits

0011 : SID's FUNCT:ARGS bits

The value field is encoded below as the lengths in bits of LOC, FUNCT and ARGS followed by the SID rounding up to bytes:



```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
~                               SID(continue)                               ~
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

Where:

LOC Length : 1-octet field indicating the length in bits of LOC in SID.

FUNCT Length : 1-octet field indicating the length in bits of FUNCT in SID.

ARGS Length : 1-octet field indicating the length in bits of ARGS in SID.

SID : the SID containing LOC, FUNCT and ARGS, and rounding up to bytes.

4. Security Considerations

No new security issues are introduced to the BGP protocol by this specification.

5. IANA

Under "Flow Spec Component Types" registry, IANA is requested to assign the following values:

Value	Name	Reference
TBD1 (15)	Whole SID	This Document
TBD2 (16)	Some bits of SID	This Document

6. Acknowledgments

The authors would like to thank Shunwan Zhuang and Rainsword Wang for their valuable suggestions and comments on this draft.

7. References

7.1. Normative References

- [I-D.ietf-idr-flow-spec-v6]
McPherson, D., Raszuk, R., Pithawala, B., akarch@cisco.com, a., and S. Hares, "Dissemination of Flow Specification Rules for IPv6", [draft-ietf-idr-flow-spec-v6-09](#) (work in progress), November 2017.
- [I-D.ietf-idr-rfc5575bis]
Loibl, C., Hares, S., Raszuk, R., McPherson, D., and M. Bacher, "Dissemination of Flow Specification Rules", [draft-ietf-idr-rfc5575bis-17](#) (work in progress), June 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>>.
- [RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", [RFC 4360](#), DOI 10.17487/RFC4360, February 2006, <<https://www.rfc-editor.org/info/rfc4360>>.
- [RFC5575] Marques, P., Sheth, N., Raszuk, R., Greene, B., Mauch, J., and D. McPherson, "Dissemination of Flow Specification Rules", [RFC 5575](#), DOI 10.17487/RFC5575, August 2009, <<https://www.rfc-editor.org/info/rfc5575>>.

7.2. Informative References

Filsfils, C., Camarillo, P., Leddy, J., daniel.voyer@bell.ca, d., Matsushima, S., and Z. Li, "SRv6 Network Programming", [draft-filsfils-spring-srv6-network-programming-07](#) (work in progress), February 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-idr-flowspec-l2vpn]

Weiguo, H., Eastlake, D., Uttaro, J., Litkowski, S., and S. Zhuang, "BGP Dissemination of L2VPN Flow Specification Rules", [draft-ietf-idr-flowspec-l2vpn-11](#) (work in progress), July 2019.

Authors' Addresses

Zhenbin Li
Huawei
156 Beiqing Road
Beijing, 100095
P.R. China

Email: lizhenbin@huawei.com

Lei Li
Huawei
156 Beiqing Road
Beijing 100095
P.R. China

Email: lily.lilei@huawei.com

Huaimo Chen
Futurewei
Boston, MA
USA

Email: Huaimo.chen@futurewei.com

Christoph Loibl
Next Layer Communications
Mariahilfer Guertel 37/7
Vienna 1150
AT

Email: cl@tix.at

Yongqing
China Telecom
109, West Zhongshan Road, Tianhe District
Guangzhou 510000
China

Email: zhuyq.gd@chinatelecom.cn

Lei Liu
Fujitsu
USA

Email: liulei.kddi@gmail.com

Xufeng Liu
Volta Networks
McLean, VA
USA

Email: xufeng.liu.ietf@gmail.com

