

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 12, 2019

Z. Li
L. Li
Huawei
March 11, 2019

BGP Flow Specification for SRv6
draft-li-idr-flowspec-srv6-00

Abstract

This draft proposes BGP flow specification rules that are used to filter SRv6 packets.

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 September 12, 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	3
4. Security Considerations	5
5. IANA	5
6. Contributors	6
7. Acknowledgments	6
8. References	6
Authors' Addresses	6

1. Introduction

BGP Flow Specification (BGP-FS) [RFC5575] defines a new BGP NLRI to distribute traffic flow specification rules via BGP ([RFC4271]). BGP-FS policies have a match condition that may be n-tuple match in a policy, and an action that modifies the packet and forwards/drops the packet. Via BGP, new filter rules can be sent to all BGP peers simultaneously without changing router configuration, and the BGP peer can install these routes in the forwarding table. BGP-FS defines Network Layer Reachability Information (NLRI) format used to distribute traffic flow specification rules. NLRI (AFI=1, SAFI=133) is for IPv4 unicast filtering. NLRI (AFI=1, SAFI=134) is for BGP/MPLS VPN filtering. [I-D.ietf-idr-flowspec-l2vpn] [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.ietf-spring-srv6-network-programming] defined the SRv6 network programming concept and its most basic functions. SRv6 SID will have the form 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 new subset of BGP-FS component types to support Segment Routing over IPv6 data plane (SRv6) filtering.

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

This document proposes new flow specifications rules that is encoded in NLRI. The following new component types are defined

- o Whole SID

Type TBD1 - Whole SID

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

Contains a set 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:

0	1	2	3	4	5	6	7
+	+	+	+	+	+	+	+
e	a	lt	gt	eq	reserve		
+	+	+	+	+	+	+	+

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.

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 produce 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]

o Some bits of SID to match

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] defined the format of SID is LOC:FUNCT:ARGS::. In some scenarios, traffic packets can just match Locator, Function ID, Argument or combine of these different fields rather than whole 128 bits SID. This document defines a set of new component type TBD2 to reduce the length of matching.

Type TBD2 - Some bits of SID

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

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

The operator byte is encoded as:

```

      0   1   2   3   4   5   6   7
+---+---+---+---+---+---+---+---+
| e | a |           type           | reserve |
+---+---+---+---+---+---+---+---+

```

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 as SID with mask to match bits as type defined:



4. Security Considerations

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

5. IANA

IANA is requested to a new entry in "Flow Spec component types registry" with the following values:

Type	RFC or Draft	Description
TBD1	This Draft	SID
TBD2	This Draft	Some bits of SID

6. Contributors

TBD

7. Acknowledgments

TBD

8. References

- [I-D.filsfils-spring-srv6-network-programming]
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., Previdi, S., Leddy, J., Matsushima, S., and
d. daniel.voyer@bell.ca, "IPv6 Segment Routing Header
(SRH)", draft-ietf-6man-segment-routing-header-16 (work in
progress), February 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-09 (work in
progress), January 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>>.
- [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>>.

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