

IDR Working Group
Internet-Draft
Intended status: Standards Track
Expires: 21 September 2026

T. Wu
H. Wang
S. Zhuang
J. Dong
Y. Huang
Huawei
20 March 2026

Source-IP-Community Filter for BGP Flow Specification
draft-wu-idr-flowspec-sip-community-filter-01

Abstract

BGP Flowspec mechanism (BGP-FS) propagates both traffic Flow Specifications and Traffic Filtering Actions by making use of the BGP NLRI and the BGP Extended Community encoding formats. This document specifies a new BGP-FS component type to support community-level filtering. The match field is the community of the source IP address that is encoded in the Flowspec NLRI. This function is applied in a single administrative domain.

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

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 21 September 2026.

Copyright Notice

Copyright (c) 2026 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. Definitions and Acronyms	3
3. The Flow Specification Encoding for Source-IP-Community Filter	3
4. Use Cases	4
5. IANA Considerations	6
6. Security Considerations	6
7. References	6
7.1. Normative References	6
7.2. Informative References	7
Authors' Addresses	7

1. Introduction

BGP Flow Specification (BGP-FS) [RFC8955] [RFC8956] 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] extends the flow-spec rules for layer 2 Ethernet packets.

This document specifies a new BGP-FS component type to support community-level filtering. The match field is the community of the source IP address that is encoded in the Flowspec NLRI. This function is applied in a single administrative domain.

2. Definitions and Acronyms

- * FS: Flow Specification
- * Source-IP-Community: The community of the source IP address

3. The Flow Specification Encoding for Source-IP-Community Filter

This document proposes a new flow specification component type that is encoded in the BGP Flowspec NLRI. The following new component type is defined.

- * Source-IP-Community

Type TBD1 - Source-IP-Community

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

Contains a set of {operator, value} pairs that are used to match the Source-IP-Community (i.e. the community of the source IP address).

The operator byte is encoded as:

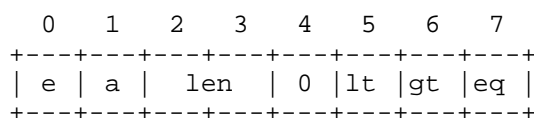


Figure 1: Numeric Operator (numeric_op)

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 MUST be unset in the Source-IP-Community filter.

len - The length of the value field for this operator given as (1 << len). This encodes 1 (len=00), 2 (len=01), 4 (len=10), and 8 (len=11) octets.

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 Source-IP-Community filter or a range of Source-IP-Community filter(e.g. less than community 1 and greater than community 2).

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
+-----+
~           Source-IP-Community  (4 octets)           ~
+-----+
```

Figure 2: Source-IP-Community

Per section 10 of [RFC8955] , If a receiving BGP speaker cannot support this new Flow Specification component type, it MUST discard the NLRI value field that contains such unknown components. Since the NLRI field encoding (Section 4 of [RFC8955]) is defined in the form of a 2-tuple <length, NLRI value>, message decoding can skip over the unknown NLRI value and continue with subsequent remaining NLRI.

4. Use Cases

This section describes how to use this function in a simple scenario. Considering the topology shown in Figure 3 ("Comm" is short for "Community"). In AS64597's R2, if the ISP AS64597 wants to redirect all packets originating from AS64598 to IP Prefix 61:

"first go to R3, then forward them to IP Prefix 61", the ISP AS64597 can use the traditional method or the method defining in this draft.

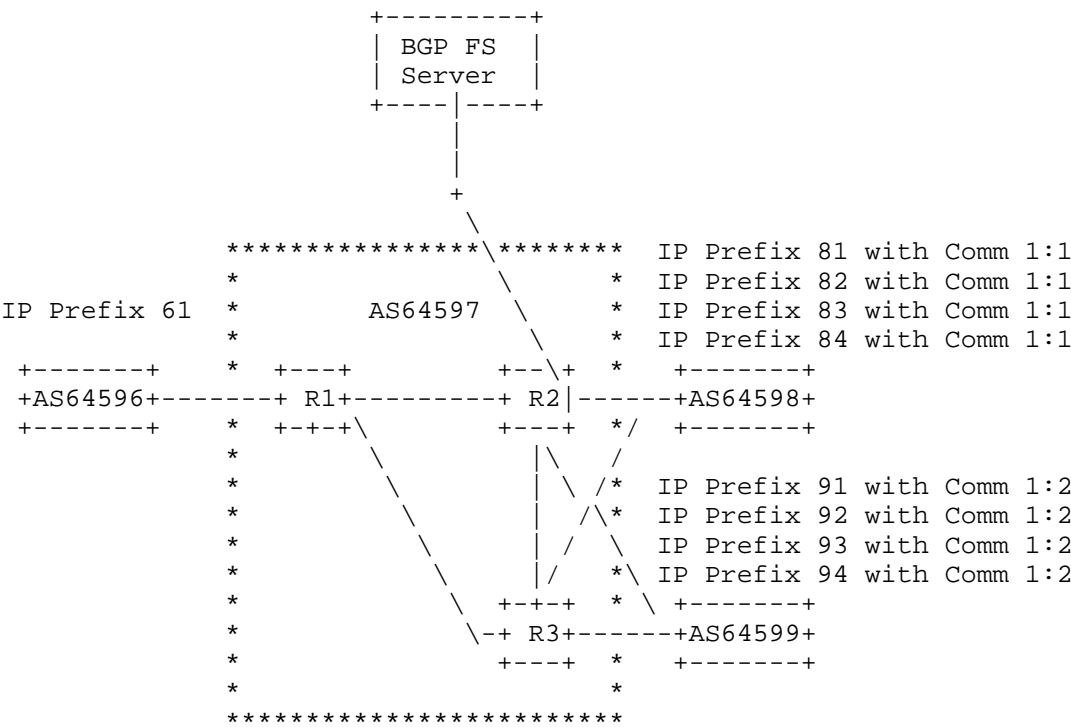


Figure 3: Redirect the traffic using Flowspec

Using the traditional method, the ISP AS64597 needs to setup multiple "Destination Prefix + Source Prefix" rules in Router R2 as following:

source Prefix	Destination Prefix	Redirect to IP Nexthop
IP Prefix 81	IP Prefix 61	R3
IP Prefix 82	IP Prefix 61	R3
IP Prefix 83	IP Prefix 61	R3
IP Prefix 84	IP Prefix 61	R3
More ...		

Figure 4: Using the traditional method to redirect the traffic

Using the method defining in this draft, the ISP AS64597 needs to setup only one "Source Community + Destination Prefix" rule in Router R2 as following:

Source Community	Destination Prefix	Redirect to IP Nexthop
1::1	IP Prefix 61	R3

Figure 5: Using the community-level filtering method to redirect the traffic

Obviously, the new method defining in this draft saves a lot of entry spaces on the control plane and forwarding plane, and it would greatly simplify the operation of the control plane, and the more source prefixes with the same community has, the more obvious the benefit.

5. IANA Considerations

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	Source-IP-Community

6. Security Considerations

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

7. References

7.1. Normative References

- [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>>.

- [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>>.

7.2. Informative References

- [I-D.ietf-idr-flowspec-l2vpn] Hao, W., 3rd, D. E. E., Litkowski, S., and S. Zhuang, "BGP Dissemination of L2 Flow Specification Rules", Work in Progress, Internet-Draft, draft-ietf-idr-flowspec-l2vpn-27, 16 March 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-flowspec-l2vpn-27>>.
- [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>>.
- [RFC8955] Loibl, C., Hares, S., Raszuk, R., McPherson, D., and M. Bacher, "Dissemination of Flow Specification Rules", RFC 8955, DOI 10.17487/RFC8955, December 2020, <<https://www.rfc-editor.org/info/rfc8955>>.
- [RFC8956] Loibl, C., Ed., Raszuk, R., Ed., and S. Hares, Ed., "Dissemination of Flow Specification Rules for IPv6", RFC 8956, DOI 10.17487/RFC8956, December 2020, <<https://www.rfc-editor.org/info/rfc8956>>.

Authors' Addresses

Tianhao Wu
Huawei
Email: wutianhao10@huawei.com

Haibo Wang
Huawei
Email: rainsword.wang@huawei.com

Shunwan Zhuang
Huawei
Email: zhuangshunwan@huawei.com

Jie Dong
Huawei
156 Beiqing Road

Beijing
100095
P.R. China
Email: jie.dong@huawei.com

Yang Huang
Huawei
156 Beiqing Road
Beijing
100095
P.R. China
Email: yang.huang@huawei.com