

IDR Working Group
Internet Draft
Intended status: Standards Track
Expires: September 01, 2026

J. Song
China Mobile
C. Lin
New H3C Technologies
March 01, 2026

BGP Flow Specification Filter by ClassID
draft-song-idr-flowspec-classid-filter-00

Abstract

BGP Flowspec mechanism (BGP-FS) [RFC8955] [RFC8956] 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 named ClassID to support ClassID filtering.

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 01 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 (<http://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. Requirements Language.....	4
3. The Flow Specification Encoding for ClassID.....	4
4. Use Cases.....	5
5. IANA Considerations.....	6
5.1. FSv2 IP Basic TLV Components.....	6
6. Security Considerations.....	7
7. Informative References.....	7
Authors' Addresses.....	9

1. Introduction

BGP Flowspec mechanism (BGP-FS) [RFC8955] [RFC8956] propagates both traffic Flow Specifications and Traffic Filtering Actions by making use of the BGP NLRI and the BGP Extended Community encoding formats.

As shown in Figure 1, PE2 advertises routes Prefix21, Prefix22, ..., Prefix2n, and PE3 advertises routes Prefix31, Prefix32, ..., Prefix3n. The server sends Flow-Spec rules to PE1 to redirect the traffic for routes Prefix2x and Prefix3x respectively.

When there are a large number of routes, a significant amount of BGP Flow-Spec rules need to be deployed.




Figure 1

To address this issue, when PE2 and PE3 advertise routes, they can include specific routing attributes.

For the included routing attributes, various combinations of BGP properties can be used, such as:

BGP Standard Community

BGP Large Community

Extended Community

Color

AS Path, etc.

By extending the BGP-FS component type, filtering of various BGP route attributes can be supported. However, the following issues exist:

- * Various attributes need to be planned on the device, transmitted across devices, and then traffic steering must be performed using BGP-FS rules on the corresponding devices.
- * Due to the large number of BGP attributes, extending corresponding BGP-FS component types for each attribute may generate a large volume of BGP FS routes.

*BGP needs to update these attributes to the forwarding table. The extensive attribute content results in the forwarding table maintaining a complex BGP attribute list.

This document specifies a new BGP-FS component type named ClassID to support ClassID filtering. The device maps various attributes into ClassIDs and updates the corresponding ClassIDs to the forwarding table. The match field is the ClassID of the destination IP address encoded in the Flowspec NLRI.

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 RFC 2119 [RFC2119] RFC 8174 [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. The Flow Specification Encoding for ClassID

[I-D. draft-ietf-idr-fsv2-ip-basic] defines the Components in the IP Basic TLV. This document proposes a new Component for ClassID information.

When filtering using FlowSpec rules, the condition can be specified as the ClassID.

This is primarily used for traffic diversion based on the ClassID.

Sub-TLV Definition

TBD - ClassID

The ClassID Components has following format:

```
+-----+
| SubTLV type = TBD (1 octet) |
+-----+
| length (1 octet)           |
+-----+
```

```
+-----+
| value (variable) |
+-----+
```

Type: ClassID(TBD)

Filter defines: match for ClassID

length: variable

value: [numeric_op, value]+

where the value following each numeric_op is a two octets ClassID.

4. Use Cases

As shown in Figure 2, on the ingress PE, Route-Policy is configured to map various route attributes into ClassID. These route attributes include Community, Large Community, Extended Community, Color, etc.

The BGP Flow-Spec Server deploys BGP Flow-Spec rules by specifying ClassID as the matching condition and Redirect as the action.

For traffic steering:

First, route attributes are mapped to a unified ClassID through Route-Policy based on traffic characteristics.

Then, traffic is redirected to the specified destination router by matching the ClassID against BGP Flow-Spec rules.

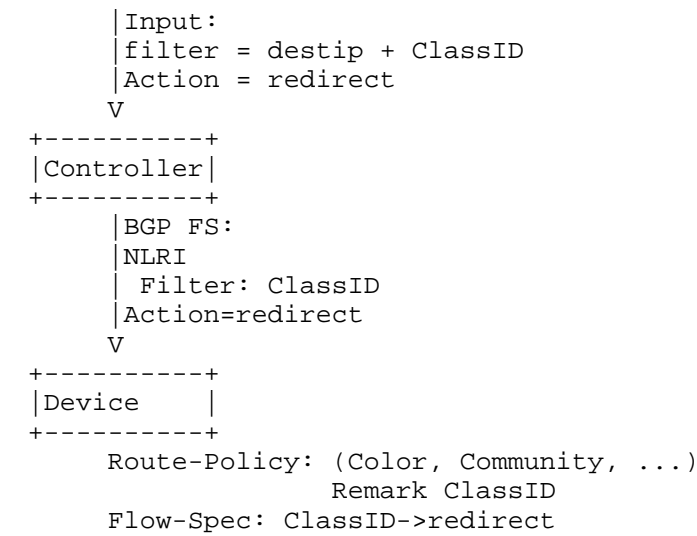


Figure 2

5. IANA Considerations

5.1. FSv2 IP Basic TLV Components

[I-D. draft-ietf-idr-fsv2-ip-basic] defines the Components in the IP Basic TLV.

This document requested to assign a new type code point from "FSv2 IP Basic TLV Components" registry for ClassID.

Sub-TLV	Definition
-----	-----
0 -	Reserved
1 -	IP Destination prefix
2 -	IP Source prefix
3 -	IPv4 Protocol / IPv6 Upper Layer Protocol
4 -	Port

5 -	Destination Port	
6 -	Source Port	
7 -	ICMPv4 type / ICMPv6 type	
8 -	ICMPv4 code / ICPv6 code	
9 -	TCP Flags	
10 -	Packet length	
11 -	DSCP	
12 -	Fragment	
13 -	Flow Label	
TBD1-	ClassID	This document
14-255	Reserved	

6. Security Considerations

TBD

7. Informative References

[I-D. ietf-idr-flowspec-v2] Hares, S., Eastlake, D. E., Yadlapalli, C., and S. Maduschke, "BGP Flow Specification Version 2", Work in Progress, Internet-Draft, draft-ietf-idr-flowspec-v2-04, 28 April 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-flowspec-v2-04>>.

[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>>.
- [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

Song Jian
China Mobile
China
Email: songjianyjy@chinamobile.com

Changwang Lin
New H3C Technologies
Email: linchangwang.04414@h3c.com

