

IDR
Internet Draft
Intended status: Standards Track
Expires: December 30, 2025

C. Lin
New H3C Technologies
H. Yao
China Mobile
Z. Li
China Mobile
Q. Xiong
ZTE Corporation
June 30, 2025

BGP Flowspec for Computing-Aware Traffic Steering
draft-lin-idr-cats-flowspec-ts-03

Abstract

A BGP Flow Specification is an n-tuple consisting of several matching criteria that can be applied to IP traffic. Computing-Aware Traffic Steering (CATS) is a framework which optimizes traffic steering to a given service instance by taking into account the dynamic nature of both computing and network resources. This document specifies a new BGP Flow Spec Component Type in order to support CATS traffic forwarding.

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 30 December 2025.

Copyright Notice

Copyright (c) 2025 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.....	3
3. Terminology.....	3
4. Flow Specifications for CATS.....	3
4.1. Scenario 1.....	4
4.2. Scenario 2.....	5
4.3. Scenario 3.....	5
5. Extend of BGP-FS.....	6
5.1. Match by CS-ID.....	6
5.2. MARK CS-ID Action	8
5.3. Redirect to CSCI-ID Action.....	9
6. Error Handling.....	10
7. Security Considerations.....	10
8. IANA Considerations.....	10
8.1. FSv2 Component types.....	10
8.2. CATS Action.....	10
9. References.....	11
9.1. Normative References.....	11
9.2. Informative References.....	11
Authors' Addresses.....	12

1. Introduction

A Flow Specification (Flow Spec) is an n-tuple consisting of several matching criteria that can be applied to IP traffic [RFC8955]. The Flow Spec conveys match conditions (each may include several components) which are encoded using MP_REACH_NLRI and MP_UNREACH_NLRI attributes [RFC4760], while the associated actions such as redirect and traffic marking are encoded in BGP Extended Communities [RFC4360][RFC5701]. The IPv4 NLRI component types and traffic filtering actions sub-types are described in [RFC8955], while the IPv6 related are described in [RFC8956]. [I-D.ietf-flowspec-v2] specifies BGP Flow Specification Version 2.[draft-ietf-idr-fsv2-ip-basic] defines FSv2 Extended Community Actions in the IP Basic TLV.

Computing-Aware Traffic Steering (CATS) is introduced in [I-D.ietf-cats-framework]. The CS-ID is a globally unique object in the network, which can be a server's IP address, an SR label, or a simple unsigned numerical value. In CATS network, the C-PS component

performs path selection based on the CS-ID and forwards service traffic according to the selected path.

This document specifies a new BGP Flow Spec Component Type to support CATS traffic filtering. Traffic is classified and mapped to the corresponding CS-ID using BGP Flow Spec rules, and path selection is then performed based on the CS-ID. BGP Flow Spec implements functions similar to the C-TC component.

On the other hand, BGP Flow Spec can also use CS-ID as a filtering criterion. For the matched traffic, it forwards based on the actions specified by BGP Flow Spec, replacing the path selection function of C-PS. It also specifies traffic filtering actions to enable the creation of the CS-ID in the outer tunnel encapsulation when matched to the corresponding Flow Spec rules.

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. Terminology

This document makes use of the terms as defined in [I-D. draft-ietf-cats-framework].

Flow Spec: Flow Specification

BGP-FS: Border Gateway Protocol (BGP) Flow Specification (FS)

4. Flow Specifications for CATS

The Flow Spec for CATS is shown in Figure 1, that is, the Controller is used to set up BGP connection with the policy enforcement points in CATS network.

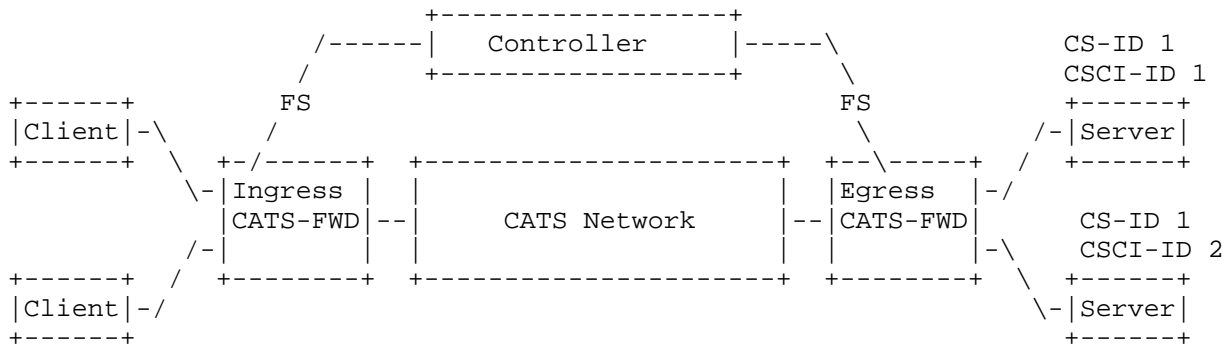


Figure 1. Flow Spec for CATS

CATS Framework Workflow is referenced in [I-D.ietf-cats-framework]. This document describes the main applications of BGP-FS in the CATS framework as follows:

4.1. Scenario 1

Flow entries are directly pushed to redirect traffic to the CSCI-ID based on flow characteristics, enabling subsequent forwarding.

Rule 1

Filter: Flow characteristics

Action: Redirect to IPv4/IPv6 CSCI-ID

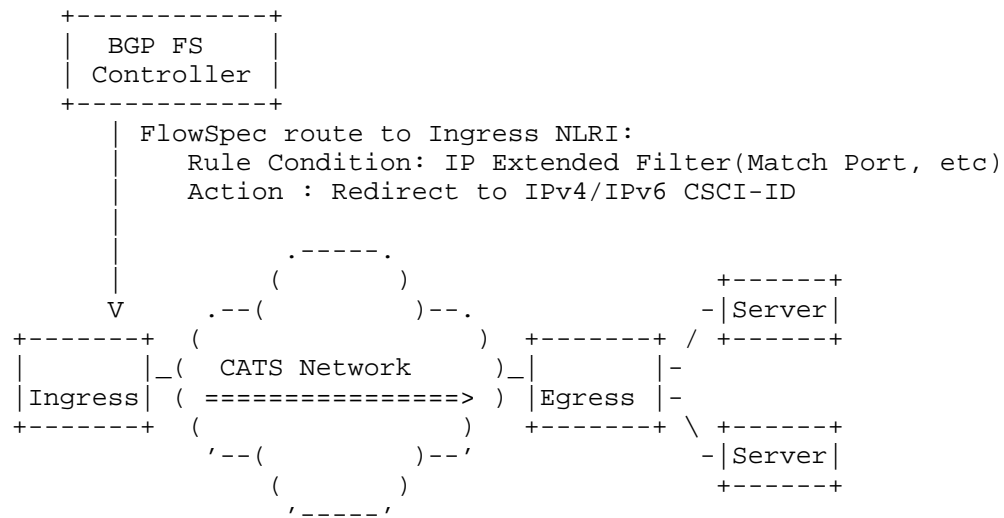


Figure 2: Scenario 1 Examples

4.2. Scenario 2

The data layer of the traffic includes a CS-ID field. Rules are pushed to filter based on the CS-ID, redirecting traffic to the CSCI-ID for forwarding.

Rule 1

Filter: IPv4/IPv6 CS-ID

Action: Redirect to IPv4/IPv6 CSCI-ID

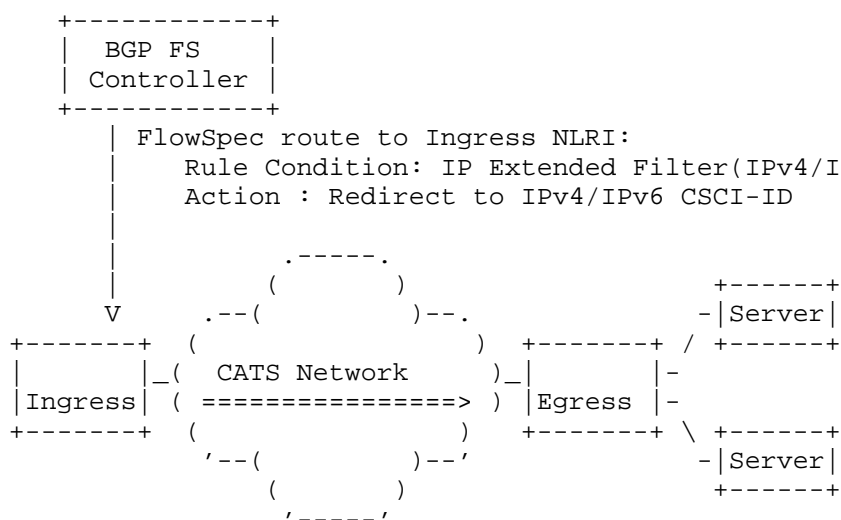


Figure 3: Scenario 2 Examples

4.3. Scenario 3

Flow characteristics mapped to the CS-ID; deploy rules to filter based on the CS-ID and redirect traffic to the CSCI-ID for forwarding..

Rule 1

Filter: Flow characteristics

Action: Mark IPv4/IPv6 CS-ID

Rule 2

Filter: IPv4/IPv6 CS-ID

Action: Redirect to IPv4/IPv6 CSCI-ID

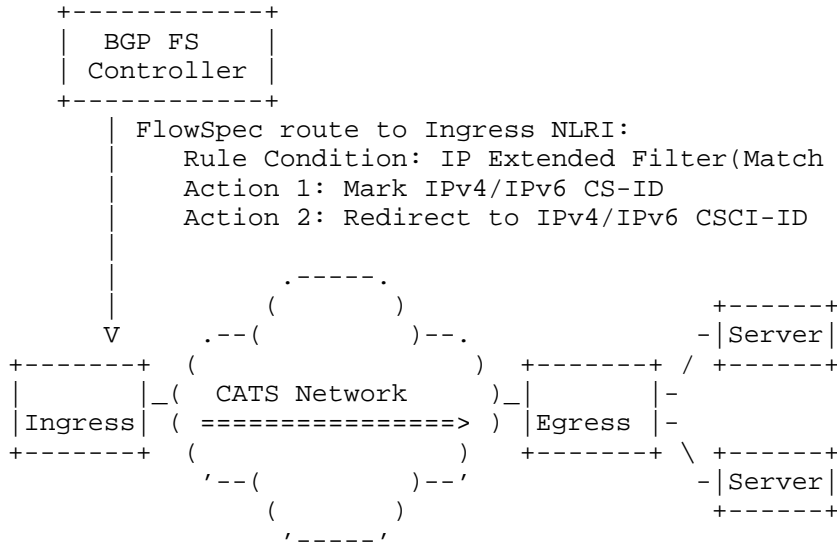


Figure 4: Scenario 3 Examples

5. Extend of BGP-FS

This document elaborates on the protocol enhancements grounded in the BGP Flowspec v2 format[I-D.ietf-idr-flowspec-v2]. It introduces one matching: matching CS-ID. Furthermore, it enriches the action repertoire with two actions: Mark CS-ID action, and Redirect to CSCI-ID action.

In scenario 4.1, the Redirect to CS-ID action is leveraged.

In scenario 4.2, matching CS-ID is deployed in tandem with the Redirect to CS-ID action.

In scenario 4.3, the Mark CS-ID action is implemented in conjunction with the Redirect to CS-ID action.

5.1. Match by CS-ID

The format of the FSv2 NLRI field for IP Filters is defined in [I-D.draft-ietf-idr-flowspec-v2]. This format includes a common header

with fields for user specified order, dependency filter chain, and a TLV for filter components (type, length, value).

This document proposes a new Component for defining CS-ID information from the Extended IP Filters Components for for IP Extended Filters version 2.

When filtering using FlowSpec rules, the condition can be specified as the CS-ID.

This is primarily used to achieve the C-PS function of CATS based on CS-ID for path selection.

The IPv4 CS-ID Components has following format:

```
+-----+
| SubTLV type = TBD1 (1 octet) |
+-----+
| length (1 octet)             |
+-----+
| value (variable)             |
+-----+
```

Filter defines: a list of match criteria for IPv4 CS-ID

Type: TBD1

length: variable

IPv4 value: [numeric_op, value]+

Each IPv4 CS-ID value is 4 octets.

The IPv6 CS-ID Components has following format:

```
+-----+
| SubTLV type = TBD2 (1 octet) |
+-----+
| length (1 octet)             |
+-----+
```

```

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

```

Filter defines: a list of match criteria for IPv6 CS-ID

Type: TBD2

length: variable

IPv6 value: [numeric_op, value]+

Each IPv6 CS-ID value is 16 octets.

5.2. MARK CS-ID Action

When implementing the C-TC function, the corresponding CS-ID is assigned based on traffic characteristics, which are specified according to the FSv2 IP Extend Filters.

If a rule is matched, execute the Mark CS-ID action.

This document defines a new traffic filtering action: "Mark CS-ID" Action. It is specifically encapsulated and carried through the BGP Community Container Attribute (also known as BGP Wide Communities) defined in [I-D.ietf-idr-wide-bgp-communities].

"Mark IPv4 CS-ID" Action SubTLV has the format:

```

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
+-----+-----+-----+-----+-----+-----+-----+-----+
| Action type(TBD3) | Length(4) |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               IPv4 CS-ID                               |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 5: "Mark IPv4 CS-ID" Action SubTLV

"Mark IPv6 CS-ID" Action SubTLV has the format:

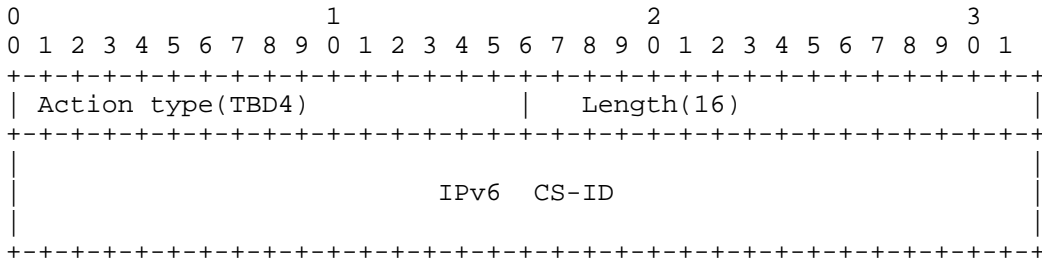


Figure 6: "Mark IPv6 CS-ID" Action SubTLV

5.3. Redirect to CSCI-ID Action

While specifying the path, C-PS designates the CSCI-ID information.

"Redirect to IPv4 CSCI-ID" Action SubTLV has the format:

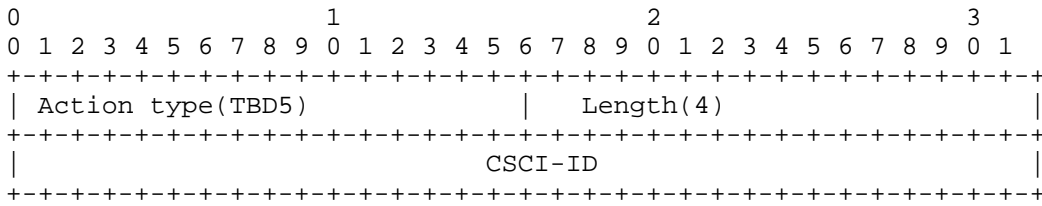


Figure 7: "Redirect to IPv4 CSCI-ID" Action SubTLV

"Redirect to IPv6 CSCI-ID" Action SubTLV has the format:

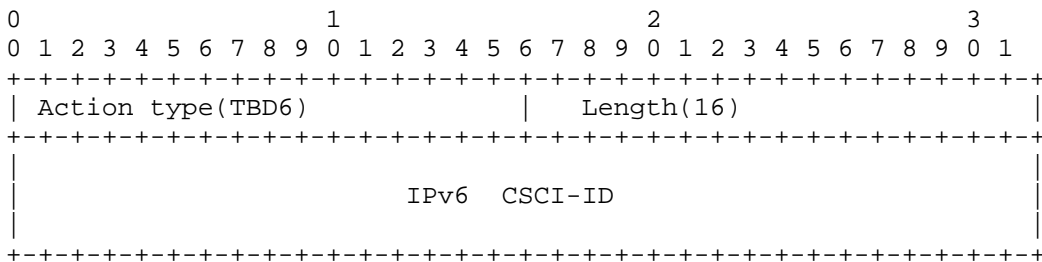


Figure 8: "Redirect to IPv6 CSCI-ID" Action SubTLV

6. Error Handling

The error handling and validation of BGP FSv2 follows section 4 of [I-D.ietf-idr-flowspec-v2]. When the Mark CS-ID Action and Redirect to CSCI-ID Action SubTLVs are missing at the same time, the FlowSpec steering falls back to the FSv2 IP Basic actions in [I-D.ietf-idr-fsv2-ip-basic].

7. Security Considerations

Security considerations for BGP FSv2 are covered in the [draft-ietf-idr-fsv2-ip-basic]. The security of CATS information distributed has been discussed in [draft-ietf-cats-framework]. They are also applicable to the extensions described in this document.

8. IANA Considerations

8.1. FSv2 Component types

This document requested to assign new type code point from FSv2 Component types for CS-ID.

Table 3 Non-IP Types for IP Filters[I-D.ietf-idr-flowspec-v2]

SubTLV

-type	Definition
-------	------------

=====	=====
-------	-------

TBD1 -	CATS Match 1: IPv4 CS-ID
--------	--------------------------

TBD2 -	CATS Match 1: IPv6 CS-ID
--------	--------------------------

8.2. CATS Action

IANA is requested to assign the following code points from the "BGP FSv2 Action types" Registry:

TBD3: MARK IPv4 CS-ID Action

TBD4: MARK IPv6 CS-ID Action

TBD5: Redirect to IPv4 CSCI-ID Action

TBD6: Redirect to IPv6 CSCI-ID Action

9. References

9.1. Normative References

- [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>>.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC4760, DOI 10.17487/RFC4760, January 2007, <<https://www.rfc-editor.org/info/rfc4760>>.
- [RFC5701] Rekhter, Y., "IPv6 Address Specific BGP Extended Community Attribute", RFC 5701, DOI 10.17487/RFC5701, November 2009, <<https://www.rfc-editor.org/info/rfc5701>>.
- [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>>.

9.2. Informative References

- [I-D.ietf-cats-framework] C. Li., Z. Du., M. Boucadair., L. M. Contreras., J. Drake., "A Framework for Computing-Aware Traffic Steering (CATS)", draft-ietf-cats-framework-10 (work in progress), June 2025.

Authors' Addresses

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

Huijuan Yao
China Mobile
No.32 XuanWuMen West Street
Beijing
100053
China
Email: yaohuijuan@chinamobile.com

Zhenqiang Li
China Mobile
China
Email: lizhenqiang@chinamobile.com

Quan Xiong
ZTE Corporation
Email: xiong.quan@zte.com.cn

