

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 1 September 2025

M. Shen
W. Li
L. Wang
G. Wang
Huawei Technologies
28 February 2025

BGP Flow-Spec Traffic Compress Action
draft-shen-idr-flowspec-traffic-compress-action-02

Abstract

Flow-spec is an extension to BGP that allows for the dissemination of traffic flow specification rules and traffic filtering actions. This document specifies a new traffic filtering action to support compressing traffic.

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 RFC 2119 [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 1 September 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 (<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. Traffic compress Extended Community	3
3. Error Handling	4
4. IANA Considerations	4
5. Security Considerations	4
6. Acknowledgements	4
7. References	4
7.1. Normative References	4
7.2. References	4
7.3. References	5
7.4. References	5
7.5. References	5
Authors' Addresses	5

1. Introduction

Flow-spec [RFC8955] [RFC8956] is an extension to BGP that allows for the dissemination of traffic flow specification rules and traffic filtering actions. The flow-spec standard defines widely-used filter actions such as discard, rate limit, traffic marking and so on.

Data transmitted on the enterprise network, data center, office, and external connections is inconsistent. The data may be text, audio, or video.

The BGP flowSpec extension allows traffic filters to be distributed to routers on the entire network. However, some nodes are suitable for compression and have a large compression ratio. Some nodes have no compression space. Some nodes have no compression or decompression capabilities. Therefore, a new traffic action needs to be added to selectively compress different traffic based on the existing traffic filters.

For example, database data is transmitted between financial data centers. We need to compress the data of different network nodes to save bandwidth.

This document specifies a new traffic filtering action that provides a method of traffic compressing. The details of the action, including compression algorithms, are encoded in newly defined BGP extended communities

2. Traffic compress Extended Community

This document proposes a new BGP extended community called the "flow-spec traffic compress action". It has a Generic Transitive Extended Community type "0x80". The sub-type value [to be assigned by IANA] indicates that the global administrator and local administrator fields encode a flow-spec traffic compress action. In the new extended community the 2-byte global administrator field encodes compression algorithms called Compression Parameter Index (CPI) [RFC-3173]. And the 4-byte local administrator field is reserved for future use.

This new BGP extended community is encoded as follows :

```

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
+-----+-----+-----+-----+-----+-----+-----+-----+
|   0x80   |   Sub-Type   |               CPI               |
+-----+-----+-----+-----+-----+-----+-----+-----+
|               CPV               |
+-----+-----+-----+-----+-----+-----+-----+

```

- * Sub-Type: TBD.
- * CPI: 2 octets. The values 0-63 designate well-known compression algorithms. The values 64-255 are reserved for future use. The values 256-61439 are negotiated between the two nodes in definition of an IPComp Association[RFC3173].
- * CPV: 4 octets. Some CPIs may include certain parameters, the significance of which, along with their implied meanings, varies from one CPI to another. Parameters consist of 4 octets, and depending on the CPI, the CPV can be divided into distinct segments, each with its own independent meaning, with the total length not exceeding 4 octets. All bits of segments that have no practical significance must be set to 0 by the originating BGP speaker and ignored by receiving.

For instance, when the compression algorithm is zstd, the CPV can be encoded as: <trainAppId (1 octet), CRC (2 octets), appId (1 octet)>

Generally, the receive end of Flowspec Traffic Compress Extended Community functions as the compressing end and the transmit end functions as the decompressing end. In addition, the CPI can be delivered through the controller or other devices. In this case, the traffic receiver must have the decompression capability corresponding to the traffic compression extended community.

3. Error Handling

If multiple traffic compression extended communities exist in the BGP route attribute, the value of the first traffic compression extended community must be used.

The validation of the TLVs/sub-TLVs introduced in this document and defined in their respective sub-sections of Section 2 MUST be performed to determine if they are malformed or invalid. In case of any error detected, the error handling should comply with [RFC7606].

4. IANA Considerations

This document requests the creation of a new registry called "Traffic compress Extended Community" under the "Extended Community" registry.

SubType	Description
TBD	Traffic compress Extended Community

5. Security Considerations

There are no additional security risks introduced by this design.

6. Acknowledgements

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>>.

7.2. References

- [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>>.

7.3. References

- [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>>.

7.4. References

- [RFC3173] Shacham, A., Monsour, B., Pereira, R., and M. Thomas, "IP Payload Compression Protocol (IPComp)", RFC 3173, DOI 10.17487/RFC3173, September 2001, <<https://www.rfc-editor.org/info/rfc3173>>.

7.5. References

- [RFC7606] Chen, E., Ed., Scudder, J., Ed., Mohapatra, P., and K. Patel, "Revised Error Handling for BGP UPDATE Messages", RFC 7606, DOI 10.17487/RFC7606, August 2015, <<https://www.rfc-editor.org/info/rfc7606>>.

Authors' Addresses

Ming Shen
Huawei Technologies
Huawei Campus, No. 156 Beiqing Road
Beijing
100095
China
Email: shenming2@huawei.com

Wenyan Li
Huawei Technologies
Huawei Campus, No. 156 Beiqing Road
Beijing
100095
China
Email: liwenyan@huawei.com

Lili Wang
Huawei Technologies
Huawei Campus, No. 156 Beiqing Road
Beijing
100095
China
Email: lily.wong@huawei.com

Guoqiang Wang
Huawei Technologies
Huawei Campus, No. 156 Beiqing Road
Beijing
100095
China
Email: wangguoqiang5@huawei.com