

Inter-Domain Routing
Internet-Draft
Intended status: Standards Track
Expires: 16 November 2025

X. Yi, Ed.
T. He, Ed.
China Unicom
C. Li, Ed.
H. Shi
X. Ding
H. Wang
Huawei Technologies
Z. Wang
Inspur
15 May 2025

Distribution of Service Metadata in BGP FlowSpec
draft-yi-idr-bgp-fs-edge-service-metadata-04

Abstract

In edge computing, a service may be deployed on multiple instances within one or more sites, called edge service. The edge service is associated with an ANYCAST IP address. Its routes along with service metadata can be collected by a central controller. The controller may process the metadata and distribute the result to ingress routers using BGP FlowSpec. The service metadata can be used by ingress routers to make path selections not only based on the routing cost but also on the running environment of the edge services. This document describes a mechanism to distribute the information of the service routes and related service metadata using BGP FlowSpec.

Discussion Venues

This note is to be removed before publishing as an RFC.

Discussion of this document takes place on the Inter-Domain Routing Working Group mailing list (idr@ietf.org), which is archived at <https://mailarchive.ietf.org/arch/browse/idr/>.

Source for this draft and an issue tracker can be found at <https://github.com/VMatrix1900/draft-bgp-fs-edge-service-metadata>.

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 16 November 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
1.1. Terminology	3
1.2. Requirements Language	3
2. BGP FlowSpec Extension for Service Metadata	3
2.1. Metadata Path Attribute TLV	4
2.2. Aggregated Metric Path Attribute TLV	4
3. Security Considerations	5
4. IANA Considerations	5
5. Normative References	5
Authors' Addresses	6

1. Introduction

Many modern services deploy their service instances in multiple sites to get better response time and resource utilization. These sites are often geographically distributed to serve the user demand. For some services such as VR/AR and intelligent transportation, the QoE will depend on both the network metrics and the compute metrics. For example, if the nearest site is overloaded due to the demand fluctuation, then steering the user traffic to another light-loaded

site may improve the QoE. To steer the traffic to the best site, the computing metadata of the site needs to be collected.

[I-D.ietf-idr-5g-edge-service-metadata] describes the BGP extension of distributing service routes with network and computing-related metrics. The router connected to the site will receive the service routes and service metadata sent from devices inside the edge site, and then generate the corresponding routes and distribute them to ingress routers. However, the route with service metadata on the router connected to the site can be also collected by a central controller using BGP LS. Then the central controller may process the metadata and distribute the result to the ingress router using BGP FlowSpec.

This document defines an extension of BGP FlowSpec to carry the service metadata along with the service route which is received from the controller. Using the service metadata and the service route, the ingress router can calculate the best site for the traffic, giving each user the best QoE.

1.1. Terminology

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

2. BGP FlowSpec Extension for Service Metadata

The goal of the BGP FlowSpec extension is to distribute the information of the service route and metadata. A service is identified by a prefix and this information is carried using the existing Destination Prefix Component specified in [RFC8955] and [RFC8956]. [I-D.ietf-idr-ts-flowspec-srv6-policy] defines that the Color Extended Community and BGP Prefix-SID attribute is carried in the context of the FlowSpec NLRI.

In addition to that, this document proposes to carry the service metadata attribute (See Figure 1). The ingress router can compare the compute metric of different sites and steer the traffic into the best one using the SR policy. The metadata can be original values defined in [I-D.ietf-idr-5g-edge-service-metadata] or an aggregated one calculated using original values.

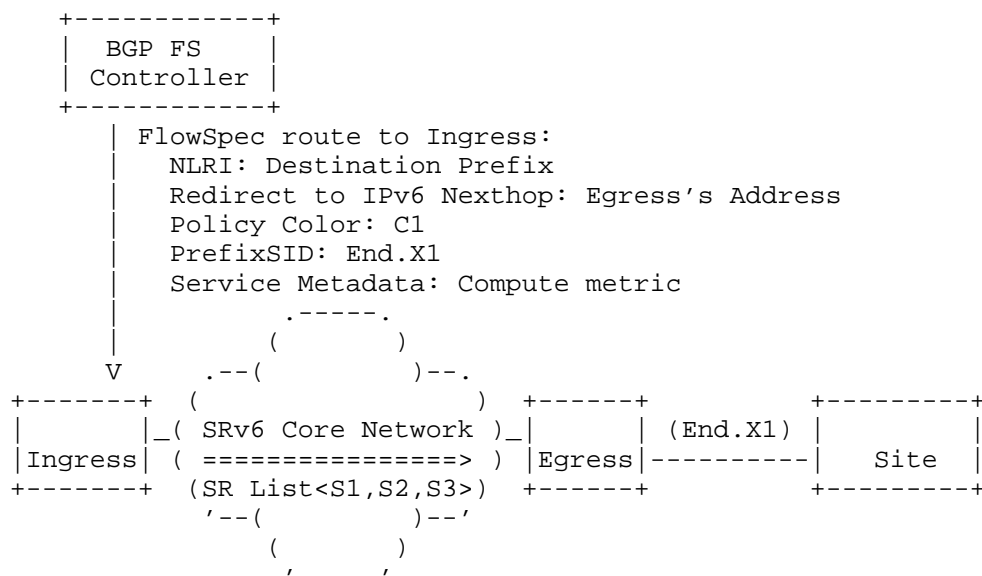


Figure 1: Example of using BGP FlowSpec to distribute the service route and metadata

2.1. Metadata Path Attribute TLV

The Metadata Path Attribute TLV is the same as defined in [I-D.ietf-idr-5g-edge-service-metadata], including the following three sub-TLVs:

1. Site Preference Index sub-TLV indicates the preference to choose the site.
2. Capacity Index sub-TLV indicates the capability of a site. One Edge Site can be at full capacity, reduced capacity, or completely out of service.
3. Load Measurement sub-TLV indicates the load level of the site.

2.2. Aggregated Metric Path Attribute TLV

The Aggregated Metric Path Attribute is a newly defined TLV(See Figure 2). It contains a single aggregated value which is calculated by the controller using the original metrics such as site preference, capacity, and load measurement.

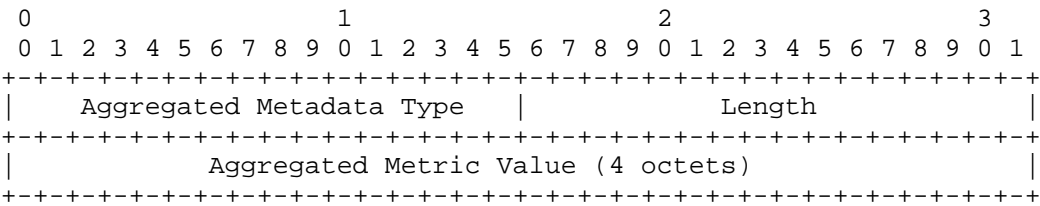


Figure 2: Aggregated Metric Path Attribute TLV format

- * Type: identify the Aggregated Metadata Attribute, to be assigned by IANA.
- * Length: the total number of the octets of the value field.
- * Value: the value of Aggregated Computing metric.

3. Security Considerations

TBD

4. IANA Considerations

This document requires IANA to assign the following code points from the registry called "BGP Path Attributes":

Value	Description	Reference
TBD1	Aggregated Metadata Type	Section 2.2

Table 1

5. Normative 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>>.

[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>>.

[I-D.ietf-idr-5g-edge-service-metadata]

Dunbar, L., Majumdar, K., Li, C., Mishra, G. S., and Z. Du, "BGP Extension for 5G Edge Service Metadata", Work in Progress, Internet-Draft, draft-ietf-idr-5g-edge-service-metadata-29, 28 April 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-5g-edge-service-metadata-29>>.

[I-D.ietf-idr-ts-flowspec-srv6-policy]

Wenying, J., Liu, Y., Zhuang, S., Mishra, G. S., and S. Chen, "Traffic Steering using BGP FlowSpec with SR Policy", Work in Progress, Internet-Draft, draft-ietf-idr-ts-flowspec-srv6-policy-05, 6 January 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-ts-flowspec-srv6-policy-05>>.

[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>>.

Authors' Addresses

Xinxin Yi (editor)
China Unicom
Beijing
China
Email: yixx3@chinaunicom.cn

Tao He (editor)
China Unicom
Beijing
China
Email: het21@chinaunicom.cn

Cheng Li (editor)
Huawei Technologies
China
Email: c.l@huawei.com

Hang Shi
Huawei Technologies
Beijing
China
Email: shihang9@huawei.com

Xiangfeng Ding
Huawei Technologies
Beijing
China
Email: dingxiangfeng@huawei.com

Haibo Wang
Huawei Technologies
Beijing
China
Email: rainsword.wang@huawei.com

Zicheng Wang
Inspur
Beijing
China
Email: wangzicheng01@inspur.com