IDR X. Yi, Ed.
Internet-Draft M. Han, Ed.
Intended status: Standards Track China Unicom
Expires: 17 May 2026 G. Zeng, Ed.
Huawei Technologies
13 November 2025

Distribution of Service Metadata in BGP FlowSpec
draft-yi-idr-bgp-fs-edge-service-metadata-05

Abstract

In edge computing and distributed cloud environments, a service may be deployed on multiple instances across one or more sites, referred to as edge service. The edge service is typically associated with an ANYCAST IP address. With the emergence of Computing-Aware Traffic Steering (CATS) requirements, there is a growing need to consider both network and computing metrics when making traffic steering decisions. Traditional routing protocols lack the capability to convey compute-related information, necessitating extensions to existing protocols.

This draft defines a mechanism to distribute service routes along with computing-related metadata using BGP FlowSpec. The service metadata, including compute resource status and performance metrics, can be collected by a central controller, processed, and then distributed to ingress routers using BGP FlowSpec extensions. This enables ingress routers to make path selections based not only on routing cost but also on the running environment and resource availability of edge services, thereby optimizing Quality of Experience (QoE).

Status of This Memo

This Internet-Draft will expire on 17 May 2026.

Copyright Notice

Table of Contents

1. Introduction

The proliferation of edge computing and multi-cloud deployments has led to services being distributed across numerous geographically dispersed sites. These deployments support applications such as VR/AR, intelligent transportation, and distributed AI workloads, which require low latency and high reliability. In such environments, multiple service instances are replicated across various sites to ensure sufficient capacity and maintain the required QoE.

Computing-Aware Traffic Steering (CATS) [I-D.ietf-idr-5g-edge-service-metadata] has been proposed as a traffic engineering approach that optimizes traffic steering to service instances by considering both network and compute resources. However, existing routing protocols like BGP focus primarily on network-layer metrics (e.g., AS paths, hop count) and lack the ability to convey compute-related information such as CPU utilization, memory capacity, or service load.

This gap creates a critical challenge: without compute-aware metrics, networks cannot make optimal steering decisions. For example, a user might be routed to the nearest site based on network latency, only to find it overloaded, while a lighter-loaded site with slightly higher latency could provide better overall QoE. To address this, there is a need to extend BGP FlowSpec to carry both service routes and compute-related metadata, enabling ingress routers to make informed decisions based on a holistic view of network and compute resources.

This document defines an extension to BGP FlowSpec that allows the distribution of service metadata alongside service routes. The extension leverages the metrics framework defined in [I-D.ietf-cats-metric-definition], particularly the Level 2 (L2) normalized metrics, to provide a scalable and efficient way to convey compute-related information.

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 an 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.

```
      +------------+
      |  BGP FS    |
      | Controller |
      +------------+
         | FlowSpec route to Ingress:
         |   NLRI: Destination Prefix
         |   Redirect to IPv6 Nexthop: Egress's Address
         |   Policy Color: C1
         |   PrefixSID: End.X1
         |   Service Metadata: Compute metric
         |          .-----.
         |         (       )
         V     .--(         )--.
   +-------+  (                 )  +-------+          +---------+
   |       |_( SRv6 Core Network )_|       | (End.X1) |         |
   |Ingress| ( ================> ) |Egress|----------|   Site  |
   +-------+  (SR List<S1,S2,S3>)  +-------+          +---------+
               '--(         )--'
                   (       )
                    '-----'
```

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 in 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

To align with the metrics framework defined in [I-D.ietf-cats-metric-definition], this document introduces an Aggregated Metric Path Attribute TLV(See Figure 2) that carries a Level 2 (L2) normalized metric. The L2 metric is a single normalized value that represents the overall performance of a service instance, derived from lower-level metrics (L0 or L1) using aggregation and normalization functions.

```
 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|    Aggregated Metadata Type   |            Length             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|            Aggregated Metric Value (4 octets)                 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

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: 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-30, 18 September 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-5g-edge-service-metadata-30>.

[I-D.ietf-cats-metric-definition] Yao, K., Li, C., Contreras, L. M., Ros-Giralt, J., and H. Shi, "CATS Metrics Definition", Work in Progress, Internet-Draft, draft-ietf-cats-metric-definition-04, 20 October 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-cats-metric-definition-04>.

[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-07, 4 August 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-ts-flowspec-srv6-policy-07>.

[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>.

Contributors

Tao He
China Unicom
Email: het21@chinaunicom.cn

Hang Shi
Huawei Technologies
Email: shihang9@huawei.com

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

Xiangfeng Ding
Huawei Technologies
Email: dingxiangfeng@huawei.com

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

Authors' Addresses

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

Mengyao Han (editor)
China Unicom
Beijing
China
Email: hanmy12@chinaunicom.cn

Guanming Zeng (editor)
Huawei Technologies
Beijing
China
Email: zengguanming@huawei.com