

Network Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: 30 November 2026

L. Dunbar  
Futurewei  
K. Majumdar  
Oracle  
C. Li  
Huawei Technologies  
G. Mishra  
Verizon  
Z. Du  
China Mobile  
29 May 2026

BGP Extension for 5G Edge Service Metadata  
draft-ietf-idr-5g-edge-service-metadata-33

## Abstract

This draft describes a new Edge Metadata Path Attribute and some Sub-TLVs for egress routers to advertise the Edge Metadata about the attached edge services (ES). The edge service Metadata can be used by the ingress routers in the 5G Local Data Network to make path selections not only based on the routing cost but also the running environment of the edge services. The goal is to improve latency and performance for 5G edge services.

The extension enables an edge service at one specific location to be more preferred than the others with the same IP address (ANYCAST) to receive data flow from a specific source, like a specific User Equipment (UE).

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

## 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 November 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 (<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 . . . . .	3
2. Conventions used in this document . . . . .	4
3. Edge Metadata Influenced Ingress Node Behavior . . . . .	5
3.1. Edge Metadata Influenced BGP Path Selection . . . . .	6
3.2. Ingress Router Forwarding Behavior . . . . .	6
3.3. Forwarding Behavior when UEs Move . . . . .	7
3.4. Relationship to Existing BGP Metric/Characteristic Attributes . . . . .	7
4. Edge Service Metadata Encoding . . . . .	7
4.1. Edge Metadata Path Attribute . . . . .	7
4.1.1. Edge Metadata Path Attribute Characteristics . . . . .	8
4.1.2. Propagation and Attribute Level Processing . . . . .	8
4.1.3. Sub-TLV Handling Rules . . . . .	9
4.2. The Site Preference Index Sub-TLV . . . . .	11
4.3. Site Physical Availability Index Metadata Sub-TLV . . . . .	12
4.3.1. Site Index Associated to Routes . . . . .	14
4.3.2. BGP UPDATE with standalone Site Availability Index . . . . .	14
4.4. Service Delay Prediction Sub-TLV . . . . .	14
4.5. Raw Measurement Sub-TLV . . . . .	16
4.6. Service-Oriented Capability Sub-TLV . . . . .	19
4.7. Service-Oriented Available Resource Sub-TLV . . . . .	20
5. Edge Metadata Processing Capability in BGP OPEN Message . . . . .	21
6. Service Metadata Propagation Scope . . . . .	23
6.1. AS-Scope SubTLV . . . . .	24
6.1.1. AS-Scope Value Checking Procedure . . . . .	25

7.	Policy Based Metadata Integration . . . . .	26
7.1.	Policy Application Order . . . . .	26
7.2.	Metadata Selection by Local Policy . . . . .	26
7.3.	Policy-Based Preference Computation . . . . .	26
7.4.	Policy Treatment of Routes with Degraded Metadata . . . . .	27
7.5.	Tie Breaking and ECMP . . . . .	27
8.	Minimum Interval for Metrics Change Advertisement . . . . .	28
9.	Validation and Error Handling . . . . .	29
10.	Manageability Considerations . . . . .	30
11.	Security Considerations . . . . .	30
12.	IANA Considerations . . . . .	31
12.1.	Edge Metadata Path Attribute . . . . .	31
12.2.	Edge Metadata Capability Code . . . . .	32
12.3.	Edge Metadata Path Attribute Sub-Types . . . . .	32
13.	Contributors . . . . .	33
14.	Acknowledgements . . . . .	33
15.	References . . . . .	33
15.1.	Normative References . . . . .	33
15.2.	Informative References . . . . .	35
Appendix A.	Service Delay Prediction Based on Load Measurement . . . . .	36
Appendix B.	Service Metadata Influenced Decision Process . . . . .	37
B.1.	Egress Router Behavior . . . . .	37
B.2.	Integrating Network Delay with the Service Metrics . . . . .	38
B.3.	Integrating with BGP Route Selection . . . . .	38
Appendix C.	Deployment Examples for Metadata-Aware Route Selection . . . . .	39
C.1.	Centralized RR Model . . . . .	40
C.2.	Ingress-Node Decision Model . . . . .	40
C.3.	Consistent Distributed Model . . . . .	40
C.4.	Example Policy Weighting Approaches . . . . .	41
Authors' Addresses	. . . . .	41

## 1. Introduction

This document describes a new Edge Metadata Path Attribute added to a BGP UPDATE message [RFC4271] for egress routers to advertise the Metadata about 5G low latency edge services directly attached to the egress routers. 5G [TS.23.501-3GPP] is characterized by having edge services closer to the Cell Towers reachable by Local Data Networks (LDN). From an IP network perspective, the 5G LDN is a limited domain [RFC8799] with edge services a few hops away from the ingress nodes. Only selective UE services are considered as 5G low latency edge services.

Note: The proposed edge service Metadata Path Attribute are not intended for the best-effort services reachable via the public Internet. The information carried by the Edge Metadata Path

Attribute can be used by the ingress routers to make path selections for selective low latency services based on not only the network distance but also the running environment of the edge cloud sites. The goal is to improve latency and performance for 5G ultra-low latency services.

This extension is targeted for a single domain with a BGP Route Reflector (RR) [RFC4456] controlling the propagation of the BGP UPDATES. The edge service Metadata Path Attribute is only attached to the low latency services (routes) hosted in the 5G edge cloud sites. These routes are only a small subset of services initiated from UEs, not for UEs accessing many internet sites.

While the proposed Edge Metadata Path Attribute is particularly beneficial for low latency services, the Edge Metadata Path Attributes can be expanded to propagate information about GPU availability, power, or other resources necessary for compute-intensive services such as AI and machine learning. This flexibility makes it a valuable tool for a wide range of applications beyond just low latency services when used within a limited domain network.

## 2. Conventions used in this document

The following conventions are used in this document.

Edge DC: Edge Data Center, which provides the hosting environment for the edge services. An Edge DC might host 5G core functions in addition to the frequently used edge services.

Edge Service Management Function: A local or centralized function that interprets Edge Metadata Path Attribute values and applies deployment-specific policy to assist in selecting the preferred path or egress router for an edge service. How this function is implemented is outside the scope of this document.

gNB: next generation Node B [TS.23.501-3GPP]

RTT: Round-trip Time

PSA: PDU Session Anchor (UPF) [TS.23.501-3GPP]

UE: User Equipment

UPF: User Plane Function [TS.23.501-3GPP]

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

### 3. Edge Metadata Influenced Ingress Node Behavior

The goal of this edge service Metadata Path Attribute is for egress routers to propagate the metrics about the running environment for a subset of edge services to ingress routers so that the ingress routers can make path selections based on not only the routing cost but also the running environment for those edge services. The BGP speakers that do not support the Edge Metadata Path Attribute can ignore the Edge Metadata Path Attribute in a BGP UPDATE Message. All intermediate nodes can forward the entire BGP UPDATE as it is. Multiple metrics can be attached to one Metadata Path Attribute. One Metadata Path Attribute can contain computing service capability information, computing service states, computing resource states of the corresponding edge site, or more. Computing service capability information can be used to record information of the computing power node or initialization deployment information for computing service initialization. Computing service states can include one of the service connection numbers, service duration, and so on. Computing resource states can be detailed information on computing resources such as CPU/GPU. They can also be an abstract metric from these detailed parameters to indicate the resource status of the edge site. There could be more metrics about the running environment being attached to the Metadata Path Attribute; e.g., some of the metrics being discussed by the IETF CATS Working Group. This document illustrates a few examples of Sub-TLVs of the metrics under the edge service Metadata Path Attribute:

- the site physical availability index,
- the site preference index,
- the service delay predication index x, and
- the raw load measurement.

This section specifies how those Metadata impact the ingress node's path selections.

### 3.1. Edge Metadata Influenced BGP Path Selection

When an ingress router receives BGP UPDATES for the same IP prefix from multiple egress routers, all these egress routers' loopback addresses are considered as the next hops for the IP prefix. For the selected low latency edge services, the ingress router BGP engine would call an edge service Management function that can select paths based on the edge service Metadata received. Section 5.1 has an exemplary algorithm to compute the weighted path cost based on the edge service Metadata carried by the Sub-TLV(s) specified in this document.

Section 5 has the detailed description of the edge service Metadata influenced optimal path selection.

### 3.2. Ingress Router Forwarding Behavior

When the ingress router receives a packet and does a lookup on the route in the FIB, it determines the destination prefix's entire path including the optimal egress node. The ingress router encapsulates the packet destined towards the optimal egress router. For routes that carry the Metadata Path Attribute but lack the Tunnel Encapsulation Path Attribute [RFC9012], it is recommended that the ingress router encapsulate the original packet using an IP-in-IP encapsulation appropriate for the address families involved. This encapsulation ensures that intermediate nodes not supporting the Metadata Path Attribute do not forward the packet to unintended destinations. The outer header SHOULD set the destination address to the optimal egress router and the source address to the ingress router.

For routes without the Metadata Path Attribute, no changes are required. Packets are forwarded according to existing behavior: encapsulation is applied when Tunnel Attributes are present, and packets are forwarded without encapsulation when they are not.

For subsequent packets belonging to the same flow, the ingress router needs to forward them to the same egress router unless the selected egress router is no longer reachable. Forwarding packets for a particular flow to the same egress router, also known as Flow Affinity, is supported by many commercial routers. Most registered EC services have relatively short-lived flows.

How Flow Affinity is implemented is out of the scope for this document.

### 3.3. Forwarding Behavior when UEs Move

When a UE moves to a new 5G gNB which is anchored to the same UPF, the packets from the UE traverse to the same ingress router. Path selection and forwarding behavior are same as before.

If the UE maintains the same IP address when anchored to a new UPF, the directly connected ingress router might use the information passed from a neighboring router to derive the optimal BGP Next Hop for this route. The detailed algorithm is out of the scope of this document.

### 3.4. Relationship to Existing BGP Metric/Characteristic Attributes

AIGP [RFC7311] carries accumulated IGP metric information for use in BGP path selection within an administrative domain. Other IDR work, such as NHC, has discussed carrying characteristics associated with the BGP next hop. The Edge Metadata Path Attribute defined in this document carries different information: service- and site-specific metadata associated with selected edge-service routes, such as site preference, site availability, service delay prediction, and service-oriented resource information.

These metadata values are not accumulated network path metrics and are not only characteristics of the BGP next hop. Different services reachable through the same egress router or at the same site may have different metadata values. Therefore, this document defines a separate attribute to carry edge-service metadata, while allowing local policy to combine that metadata with other BGP attributes when applicable.

## 4. Edge Service Metadata Encoding

### 4.1. Edge Metadata Path Attribute

The Edge Metadata Path Attribute is an optional non-transitive BGP Path attribute that carries metadata associated with edge services attached to the egress router. The attribute consists of one or more Edge Metadata Sub-TLVs, where each Sub-TLV encodes one specific metadata item associated with the advertised route or service.

The Edge Metadata Path Attribute MAY be included in a BGP UPDATE together with other BGP Path Attributes, such as Communities, NEXT\_HOP, Tunnel Encapsulation Path Attribute, and other applicable attributes. The choice of which routes carry the Edge Metadata Path Attribute, and which Sub-TLVs are included for those routes, is determined by local policy. The fields within the Edge Metadata Path Attribute and all included Sub-TLVs MUST use network byte order.

Boundary filtering SHOULD be applied at the administrative boundary to prevent the Edge Metadata Path Attribute from being distributed beyond its intended scope.

#### 4.1.1. Edge Metadata Path Attribute Characteristics

The Edge Metadata Path Attribute has the following characteristics:

- The attribute is optional non-transitive.
- The attribute MUST contain at least one Edge Metadata Sub-TLV.
- A single Edge Metadata Path Attribute MAY carry multiple Sub-TLVs.
- The attribute MAY be attached to UPDATE messages for any supported AFI/SAFI as allowed by local policy and by the capability negotiation specified in Section 5.
- Only a subset of routes are expected to carry the Edge Metadata Path Attribute. Which routes carry the attribute is deployment specific.

#### 4.1.2. Propagation and Attribute Level Processing

A BGP speaker that receives a BGP UPDATE containing the Edge Metadata Path Attribute and readvertises that route within the same metadata distribution domain SHOULD propagate the Edge Metadata Path Attribute without modification, unless local policy explicitly requires otherwise.

When advertising the route to a peer outside the intended metadata distribution domain, the speaker SHOULD remove the Edge Metadata Path Attribute.

If a BGP speaker originates or modifies a route and is configured to attach Edge Metadata, it MAY add the Edge Metadata Path Attribute to the UPDATE message, subject to local policy and the capability negotiation specified in Section 5.

A BGP speaker that receives a malformed Edge Metadata Path Attribute that cannot be parsed according to the attribute format and length rules MUST handle the error as specified in Section 9.

The Edge Metadata Path Attribute is intended to describe the specific edge service route to which it is attached. This document does not define aggregation of Edge Metadata Path Attribute values across multiple BGP routes. A BGP speaker that locally originates an aggregate route MUST NOT infer or copy Edge Metadata from component routes unless it explicitly originates new Edge Metadata for the aggregate route by local policy.

#### 4.1.3. Sub-TLV Handling Rules

When a BGP speaker receives a well-formed Edge Metadata Path Attribute, it MUST process each included Sub-TLV independently.

An implementation MUST enforce an upper bound on the total number of Sub-TLVs accepted in a single Edge Metadata Path Attribute. The bound MAY be configurable and MUST have an implementation-defined default value. If the number of Sub-TLVs in an Edge Metadata Path Attribute exceeds the configured or implementation-defined bound, the receiver MUST treat the Edge Metadata Path Attribute as unusable for metadata-based route selection and handle the attribute according to Section 9.

If the attribute contains one or more Sub-TLVs whose types are recognized by the receiving speaker, the receiving speaker SHOULD process those recognized Sub-TLVs according to their definitions in this document and according to local policy.

If the attribute contains a Sub-TLV whose type is unknown or unsupported by the receiving speaker, the speaker MUST ignore that Sub-TLV and MUST continue processing the remaining Sub-TLVs. The presence of an unknown or unsupported Sub-TLV MUST NOT by itself cause the entire Edge Metadata Path Attribute to be considered malformed.

If a Sub-TLV type is recognized by the receiving speaker, but the value carried in that Sub-TLV is invalid according to the definition of that Sub-TLV, the speaker MUST treat that Sub-TLV as unusable and MUST ignore it for metadata-based route selection. The speaker SHOULD continue processing the remaining Sub-TLVs. An invalid value in one recognized Sub-TLV MUST NOT by itself cause the entire Edge Metadata Path Attribute to be considered malformed unless the corresponding Sub-TLV definition explicitly states otherwise.

If the Length field of a Sub-TLV is inconsistent with the encoding defined for that Sub-TLV, or if the Sub-TLV cannot be fully parsed based on the encoded length, the Edge Metadata Path Attribute MUST be treated as malformed, and error handling MUST follow the procedures specified in Section 9.

If a recognized Sub-TLV appears more times than allowed by its definition, the receiver SHOULD use only the first occurrence unless the specific Sub-TLV definition states otherwise, and SHOULD ignore the additional occurrences.

A BGP speaker that propagates the Edge Metadata Path Attribute SHOULD NOT delete unrecognized Sub-TLVs solely because they are unrecognized. If the route is propagated with the Edge Metadata Path Attribute, unrecognized Sub-TLVs SHOULD remain unchanged in the propagated attribute unless local policy requires removal of the entire attribute.

If some Sub-TLVs are absent, the receiving speaker MUST treat the attribute as carrying only the metadata explicitly present. The absence of a particular Sub-TLV MUST NOT be interpreted as a zero value, an infinite value, a degraded condition, or any other inferred semantic value unless the specific Sub-TLV definition explicitly states such behavior.

If none of the included Sub-TLVs are recognized by the receiving speaker, the speaker MUST treat the Edge Metadata Path Attribute as present but unusable for local metadata-based route selection. In that case, the speaker SHOULD fall back to route selection based on other applicable BGP attributes and local policy.

The same treatment applies when the only recognized Sub-TLVs present contain invalid values and are ignored according to their Sub-TLV definitions.

By default, a BGP speaker is not required to report unknown, unsupported, or unusable Sub-TLVs to its peer. Logging or notification to a local management system is OPTIONAL.

Ingress nodes that use Edge Metadata for route selection SHOULD apply a deployment-specific algorithm to the set of recognized Sub-TLVs that are present and usable in the received attribute. To ensure consistent route selection, nodes participating in the same deployment SHOULD use consistent policy regarding which Sub-TLVs are considered and how their values are incorporated into route selection.

#### 4.2. The Site Preference Index Sub-TLV

Different services might have different preference index values configured for the same site. For example, Service-A requires high computing power, Service-B requires high bandwidth among its microservices, and Service-C requires high volume storage capacity. For a DC with relatively low storage capacity but high bisectonal bandwidth, its preference index value for Service-B is higher and lower for Service-C. Site Preference Index can also be used to achieve stickiness for some services.

It is out of the scope of this document how the preference index is determined or configured.

The Site Preference Index is not intended to replace LOCAL\_PREF or change the LOCAL\_PREF semantics defined by BGP. LOCAL\_PREF expresses operator routing policy according to the BGP deployment in which it is used. By contrast, the Site Preference Index carries service-specific site metadata associated with an edge-service route. Different services at the same site can have different Site Preference Index values, and local policy may use those values, together with LOCAL\_PREF and other BGP attributes, when selecting among candidate routes for metadata-aware services.

The Site Preference Index Sub-TLV has the following format:

```

 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
+-----+-----+-----+-----+-----+-----+-----+-----+
|Site-Preference-Index Sub-Type | Length           | Reserved       |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               Site Preference Index value
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 1: Site Preference Index Sub-TLV

- Site-Preference-Index Sub-Type (16 bits): 1 (specified in this document).
- Length (8 bits): Specifies the total length in octets of the value field (not including the Type and Length fields). For the Site-Preference-Index Sub-Type, the length SHOULD be set to 5.
- Reserved: Reserved for future use. In this version of the document, the Reserved field MUST be set to zero and MUST be ignored upon receipt. Received values MUST be propagated without change.

- Site Preference Index value: 1 .. ( $2^{32}-1$ ); the higher the value, the more preference for the site. Site Preference Index value == 0 is reserved, and the Site-Preference-Index Sub-TLV SHOULD be ignored when 0 is received..

#### 4.3. Site Physical Availability Index Metadata Sub-TLV

The Site Physical Availability Index indicates the percentage of impact on a group of routes associated with a common physical characteristic, for example, a pod, a row of server racks, a floor, or an entire DC. The purpose is to use one UPDATE message to indicate a group of routes of different NLRIs impacted by a physical event. For example, a power outage to a pod can cause the Site Physical Availability Index to be 0% for all the routes in the pod. Partial fiber cut to a row of shelves can cause the Site Physical Availability Index to be 50% for all the routes in those shelves. The value is 0-100, with 100% indicating the site is fully functional, 0% indicating the site is entirely out of service, and 50% indicating the site is 50% degraded.

It is recommended to assign each route with one Site-ID. When a route is associated with multiple Site-IDs, the latest BGP UPDATE will override any previous associations. For example, one DC can use POD number as Site-ID, another DC can use Row of Shelves as the Site-ID.

Cloud Site/Pod failures and degradation include, but are not limited to, a site degradation or an entire site going down caused by a variety of reasons. Examples include fiber cuts impacting a site or among pods, cooling failures, insufficient backup power, cyber threats attacks, too many changes outside of the maintenance window, etc. Fiber-cuts are not uncommon within a Cloud site or between sites.

When a physical failure occurs at an edge site (or a pod), many instances can be affected, and the associated routes (i.e., IP addresses) may not be easily aggregated. Instead of sending numerous BGP UPDATE messages to ingress routers for each impacted instance, the egress router can send a single BGP UPDATE to indicate the site's physical capacity availability. Based on this update, ingress routers can decide to reroute all or some of the affected instances, depending on the extent of the site's degradation. This approach significantly improves efficiency, particularly when fault detection within an edge site relies on proprietary or deployment-specific mechanisms.

The BGP UPDATE for the individual instances (i.e., the routes) can include the Capacity Availability Index solely for ingress routers to associate the routes with the Side-ID. The actual Capacity Availability Index value, i.e., the percentage for all the routes associated with the Side-ID, is generated by the egress routers with the egress routers' loopback address as the NLRI.

The Site Physical Availability Index Sub-TLV has fixed length of 8 Octets, including the Type field.

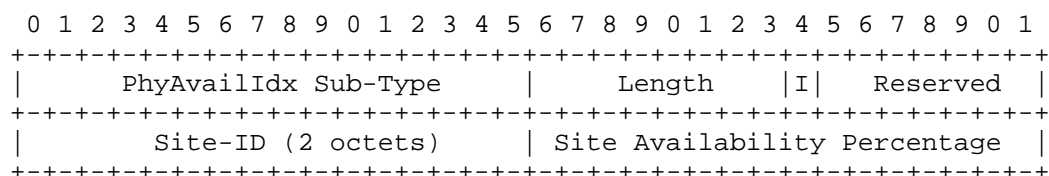


Figure 2: Site Physical Availability Index Sub-TLV

- PhyAvailIdx Sub-Type (16 bits): Indicates the Site-Physical-Availability-Index Sub-Type=2 (Specified in this document).
- Length (8 bits): Specifies the total length in octets of the value field (not including the Type and Length fields). For the PhyAvailIdx Sub-Type, the length SHOULD be set to 5.

Route-Flag I (1 bit): is a flag bit. When set to 1, the Site Availability Index is for BGP speakers (receivers) to associate the routes with the Site-ID. The Site Availability Percentage value is ignored. When set to 0, the BGP speakers (receivers) SHOULD apply the Site Availability Index value to all the routes associated with the Site-ID.

Reserved (7 bits): Reserved for future use. The bits are set to zero upon transmission, and ignored upon reception.

- Site ID (16 bits): is an identifier for a group of routes associated with a common physical characteristic, for example, a pod, a row of server racks, a floor, or an entire DC. The purpose is to use one UPDATE message to indicate a group of routes impacted by a physical event. Those routes might be from different address families or NLRIs. There could be multiple sites connected to one egress router (a.k.a. Edge DC GW).
- Site Availability Percentage (16 bits): When the RouteFlag-I is 1,

the Site Availability Percentage is ignored by the Ingress routers. When the RouteFlag I is set to 0, the Site Availability Percentage represents the percentage of the site availability for all the routes associated with the Site-ID; e.g., 100%, 50%, or 0%. When a site goes dark, the Index is set to 0. 50 means 50% functioning. When the value is outside the 0-100% range, the value carried in this Sub-TLV is ignored.

#### 4.3.1. Site Index Associated to Routes

An egress router sets itself as the next hop for a BGP peer before sending an UPDATE with the Edge Metadata Path Attribute that includes the Site Physical Availability Index Sub-TLV. The Site Physical Availability Index Sub-TLV (with RouteFlag-I=1) is for ingress routers to associate the Site Identifier with the prefixes.

#### 4.3.2. BGP UPDATE with standalone Site Availability Index

A BGP UPDATE that includes the Site Availability Index Sub-TLV without specifying attached routes in the NLRI, but instead using the egress router's loopback address in the NLRI, is referred to as a standalone Site Availability Index BGP UPDATE. When an ingress router receives such a BGP UPDATE containing the Edge Metadata Path Attribute with the standalone Site Physical Availability Index Sub-TLV from Router-X or its RR with the Originator-ID equal to Router-X, the ingress router SHOULD use the site availability index to efficiently reduce or increase the preference for all BGP routes attached to Router-X.

The BGP UPDATE with a standalone Site Availability Index is NOT intended for resolving NextHop.

#### 4.4. Service Delay Prediction Sub-TLV

It is desirable for an ingress router to select a site with the shortest processing time for an ultra-low latency service. However, it is not easy to predict which site has "the fastest processing time" or "the shortest processing delay" for an incoming service request because:

- The given service instance shares the same physical infrastructure with many other applications and service instances. Service requests by other applications, UEs, or applications running behavior can impact the processing time for the given service instance.
- The given service instance can be served by a cluster of servers

behind a Load Balancer. To the network, the service is identified by one service ID.

- The service complexity is different. One service may call many microservices, need to access multiple backend databases, and need to go through sophisticated security scrubbing functions, etc. Another service can be processed by a few simple steps. Without the application internal logic, it is not easy to estimate the processing time for future service requests.

Even though utilization measurements, like those below, are collected by most data centers, they cannot indicate which site has the shortest processing time. A service request might be processed faster on Site-A even if Site-A is overutilized.

- Server utilization for the server where the instance is instantiated.
- The network utilization for the links to the server where the instance is instantiated.
- The number of databases that the service instance will access.
- The memory utilization of the databases.

The remaining available resource at a site is a more reasonable indication of process delay for future service requests.

- The remaining available Server resources.
- The remaining available network utilization for the links to the server where the instance is instantiated.
- The number of databases that the service instance will access.
- The remaining storage available for the databases.

The Service Delay Prediction Index is a value that predicts processing delays at the site for future service requests. The higher the value, the longer of the delay.

While out of scope, we assume there is an algorithm that can derive the Service Delay Prediction Index that can be assigned to the egress router. When the Service Delay Prediction value is updated, which can be triggered by the available resources change, etc., the egress router can attach the updated Service Delay Predication value in a Sub-TLV under the Edge Metadata Path Attribute of the BGP Route UPDATE message to the ingress routers.

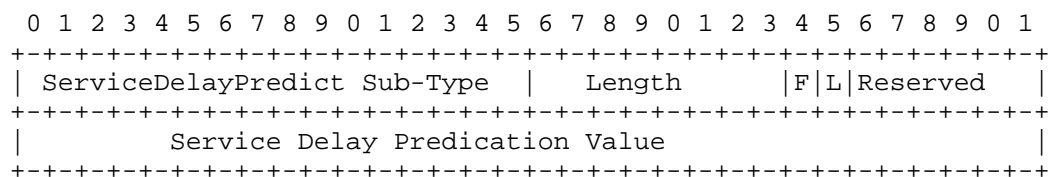


Figure 3: Service Delay Prediction Index Sub-TLV

- ServiceDelayPredict Sub-Type (16 bits): 3 (specified in this document).
- Length (8 bits): specifies the total length in octets of the value field, not including the sub-Type and Length field. The value of Length can be 5 or 9 depends on what format the Service Delay Prediction Vlaue uses.
- Flag (F) (1 bit): Indicates whether the Service Delay is a timer value (F=0) or a relative value (F=1) where a higher value represents a longer delay
- Flag (L) (1 bit): Indicates the unit of measurement for the Service Delay Prediction Value. When the F-flag is set to 0, L=0 specifies the 64-bit NTP Timestamp format, and L=1 indicates milliseconds. If the F-flag is set to 1, the L-flag value is ignored.
- Reserved (6 bits): These bits are reserved for future use and MUST be set to zero. Future documents may specify different uses for these bits.
- Service Delay Predication Value (when the Flag bit is set to 1): an integer in the range of 0-100, with 0 indicating that the service delay is negligible and 100 indicating that the site has the most significant delay compared to all other sites for the same service. When the value is outside the 0-100 range, the value carried in this Sub-TLV is ignored.
- Service Delay Predication Value (when the Flag bit is set to 0): the estimated delay time encoded in the NTP Format as defined in [RFC5905]. When the L-flag is 1, then it is a 64-bit format, otherwise it is a 32-bit short format.

#### 4.5. Raw Measurement Sub-TLV

When ingress routers have embedded analytics tool relying on the raw measurements, it is useful for the egress router to send the raw measurement.

Raw Measurement Sub-TLV has the following format:

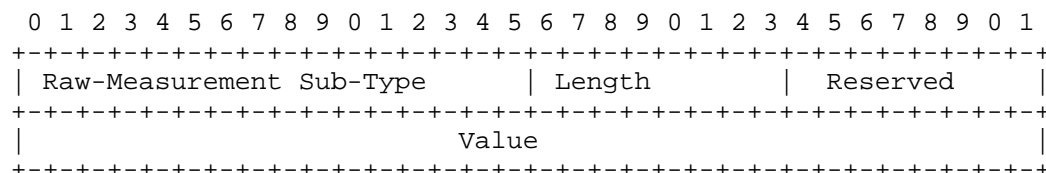


Figure 4: Service Delay Prediction Raw Measurements Sub-TLV

- Raw-Measurement Sub-Type (16 bits): 4 (specified in this document). Indicating raw measurements Metadata associated with the edge service address.
- Length (8 bits): specifies the total length, in octets, of the value field, excluding the Sub-Type and the Length fields. For the Raw-Measurement Sub-Type, the length is determined by the Value field, which carries one or more types of raw measurement.
- Reserved (8 bits): These bits are reserved for future use and MUST be set to zero. Future documents may specify different uses for these bits.
- Value: The value field can contain multiple types of raw measurements, each represented as a Sub-Sub-TLV.

One example of a raw measurement Metadata Sub-sub-TLV is defined below to convey the total number of packets or bytes transmitted over a specified period for a particular edge service address. When a Data DC GW router cannot directly access the internal state of an edge service, the volume of incoming traffic can be a reliable indicator of its load. A sudden increase in packets or bytes can signal a surge in requests, potentially leading to performance issues or resource constraints on the service side.

To differentiate this measurement from others that may be defined in the future, this document assigns a Sub-sub-Type value of 1 to represent the total packets or bytes transmitted to an edge service address.

Future documents may define additional Sub-sub-types of raw measurement metadata. Each type of raw measurement will have a unique Sub-sub-type value assigned at the time of its specification.

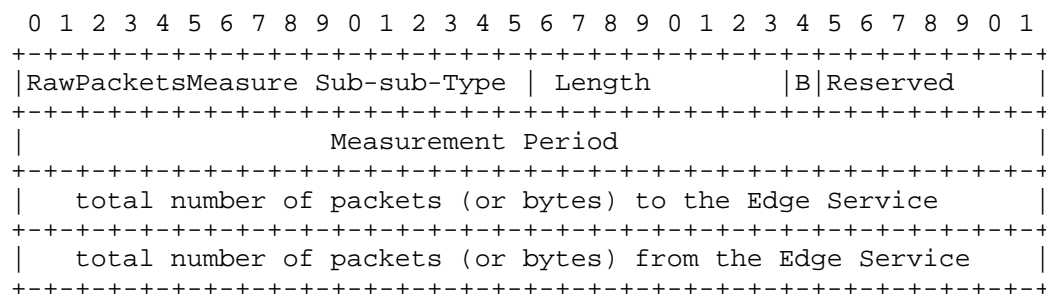


Figure 5: Packets or Bytes Measurements Sub-TLV

- RawPacketsMeasure Sub-sub-Type (8 bits): 1 (specified in this document). Indicating raw measurements of packets or bits transmitted to or from the edge service address.
- Length (8 bits): specifies the total length in octets of the value field, excluding the Sub-sub-Type and the Length fields. For the raw measurements of packets transmitted to or from the edge service address Sub-sub-Type, the length SHOULD be 22.
- B flag (1 bit): If set to 0, the raw measurement is the number of packets. If set to 1, the raw measurement is the number of bytes.
- Reserved (7 bits): These bits are reserved for future use and MUST be set to zero.
- Measurement Period: BGP Update period in Seconds or user-specified period.
- Total number of packets to the Edge Service (32 bits): This field specifies the total number of packets transmitted to the edge service address over the specified measurement period.
- Total number of packets from the Edge Service (32 bits): This field specifies the total number of packets from the edge service address over the specified measurement period.

The receiver nodes can compute the needed metrics, such as the Service Delay Prediction, for the service based on the raw measurements sent from the egress router and preconfigured algorithms.

#### 4.6. Service-Oriented Capability Sub-TLV

The service-oriented capability Sub-TLV is for distributing information regarding the capabilities of a specific service in a deployment environment. Depending on the deployment, a deployment environment can be an edge site or other types of environments. This information provides ingress routers or controllers with the available resources for the specific service in each deployment environment. It enables them to make well-informed decisions for the optimal paths to the selected deployment environment.

Currently, the Sub-TLV only has an abstract value derived from various metrics, although the specifics of this derivation are beyond the scope of this document. Importantly, this value is significant only when comparing multiple data center sites for the same service. This value is not meaningful when comparing different services, meaning the capability value relevant to Service A cannot be directly compared with that for Service B. Future enhancements may expand this sub-TLV to include more types of metrics or even raw data that represents direct metrics. This information is important in 5G network environments where efficient resource utilization is crucial for enhancing performance and service quality.

```

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
+-----+-----+-----+-----+-----+-----+-----+-----+
| ServiceOriented Cap Sub-Type | Length | Res | MT |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               SO-CapValue                               |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 6: Service-Oriented Capability Sub-TLV

- ServiceOriented Cap Sub-Type (16 bits): 5 (specified in this document).
- Length (8 bits): Specifies the total length in octets, excluding the sub-Type and Length fields. For the ServiceOriented Cap Sub-Type, the Length SHOULD be 5.
- Res (4 bits): These bits are reserved for future use and MUST be set to zero.
- MT (Metric Type)(4 bits): An unsigned 4 bits integer. When the MT

value is set to 0, it indicates the SoCapValue field contains a normalized metric derived from multiple metric types. The rules for deriving this normalized metric are out of scope of this document and defined by per-service. Additional metric types may be defined in future documents.

- SO-CapValue (32 bits): The Service-Oriented Capability Abstract Value is an integer between 0 and  $2^{32}-1$ . A larger number means higher capability, and a value of 0 indicates the site has the lowest relative capability for the service. The method used to derive this value is beyond the scope of this document.

Multiple Service-Oriented Capability Sub-TLVs with different metric types can be encoded in a Edge Metadata Path Attribute, indicating that multiple metrics are carried. However, if more than one Service-Oriented Capability Sub-TLVs with the same metric type are encoded in a Edge Metadata Path Attribute, only the first one will be processed and the others will be ignored in processing.

#### 4.7. Service-Oriented Available Resource Sub-TLV

The "Service-Oriented Available Resource Sub-TLV" is for distributing a metric that measures the real-time available resources allocated for processing specific services or applications at an edge site. This Sub-TLV complements the "Service-Oriented Capability Sub-TLV" described in Section 4.6, which addresses the static resource capability of a site for a service. While the Capability Abstract Value provides a baseline understanding of a site's potential to handle a service, the Available Resource metric offers a dynamic perspective by quantifying how much of this capacity is currently available. This distinction is crucial for managing resource efficiency and responsiveness in network operations, ensuring that capabilities are not only available but also optimally used to meet the actual service demands.

```

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
+-----+-----+-----+-----+-----+-----+-----+-----+
|ServiceOriented Avail Sub-Type |   Length       |P| Res |  MT  |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               SO-AvailRes                               |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 7: Service-Oriented Available Resource Sub-TLV

- ServiceOriented Avail (Service-Oriented Available Resource) Sub-Type: 6 (specified in this document).

- Length (8 bits) Specifies the total length in octets, excluding the sub-Type and the length field. For the ServiceOriented Available Resource Sub-Type, the Length SHOULD be 5.
- Flag (P): Is a single-bit Percentage flag. When it is set to 1, it indicates the value is the Service-Oriented Available Resource in percentage. When the "P" flag is set to 0, the value in this Sub-TLV is the abstract value of the available resource.
- Res (3 bits): These bits are reserved for future use and MUST be set to zero.
- MT (4 bits) Metric Type. This document defines a default metric type as value 0, indicating this is the normalized metric derived by multiple type of metrics. The rules to derive the normalized metric are out of scope of this document and defined by the service. Other Metric Types could be defined by other documents in the future.
- SO-AvailRes (32 bits): When the P-Flag bit is set to 1, Service-Oriented Available Resource Value is a percentage (0-100), with 0 indicating that 0% of the capability is available and 100 indicating that 100% of the capability is available. When the value is outside the 0-100 range, the value carried in this Sub-TLV is ignored. For example, Capacity value is 50 and the SO-AvailRes is 50 when P-flag is set, it means 50% of 50 unit of resource is available, while 25 unit of resource is available in this site for the service. When the P-flag is 0, then the value of this field is the abstract value of the available resource. For example, When the capacity value is 50, and the SO-AvailRes is 50, it means all the resource is available.

Multiple Service-Oriented Available Resource Sub-TLVs with different metric types can be encoded in a Edge Metadata Path Attribute, indicating that multiple metrics are carried. However, if more than one Service-Oriented Available Resource Sub-TLVs with the same metric type are encoded in a Edge Metadata Path Attribute, only the first one will be processed and the others will be ignored in processing.

## 5. Edge Metadata Processing Capability in BGP OPEN Message

The BGP Capabilities Optional Parameter allows a BGP speaker to advertise, during the BGP OPEN message exchange, the set of capabilities supported on a session. As specified in [RFC5492], each capability is encoded as a Capability Code, a Capability Length, and a Capability Value.

To enable the exchange of the Edge Metadata Path Attribute on a BGP session, this document defines a new Edge Metadata Processing Capability (=78). This capability is used by a BGP speaker to indicate that it can send and receive the Edge Metadata Path Attribute for one or more AFI/SAFI pairs on that session.

The Value Field of the Edge Metadata Processing Capability has the following format:

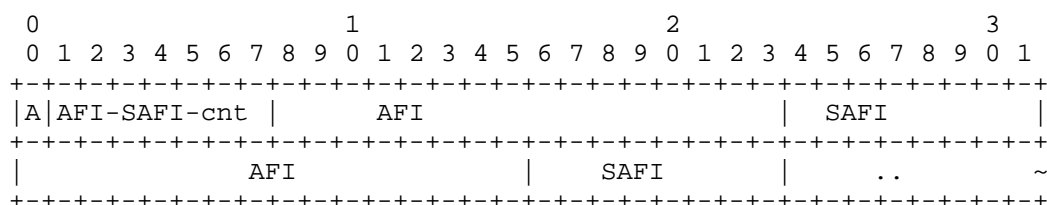


Figure 8: Edge Metadata Capability Value Field

Where:

- A Flag(1 bit): When set to 1, this flag indicates that the sender is willing to send and receive the Edge Metadata Path Attribute for any AFI/SAFI enabled on the BGP session. When set to 0, the capability applies only to the AFI/SAFI pairs explicitly listed in the Capability Value.
- AFI-SAFI-CNT (7 bits): Indicates the number of AFI/SAFI pairs encoded in the Capability Value.
- AFI (16 bits): Address Family Identifier.
- SAFI (8 bits): Sub-address Family identifier.

When the A Flag is set to 1, the capability applies to any AFI/SAFI enabled on the BGP session. In this case, AFI-SAFI-CNT SHOULD be set to 0 and no AFI/SAFI tuples need be present in the Capability Value. If one or more AFI/SAFI tuples are present when the A Flag is set to 1, the receiver SHOULD ignore those tuples and process the capability as applying to all AFI/SAFI enabled on the session.

When the A Flag is set to 0, AFI-SAFI-CNT indicates the exact number of AFI/SAFI pairs listed in the Capability Value, and the capability applies only to those listed AFI/SAFI pairs.

A BGP speaker MUST NOT attach the Edge Metadata Path Attribute to any UPDATE message sent on a BGP session unless both peers have advertised the Edge Metadata Processing Capability for the corresponding AFI/SAFI on that session. If one peer has advertised the capability with the A Flag set to 1, that advertisement is considered to cover any AFI/SAFI enabled on the session for the purpose of this check.

If a BGP speaker has not advertised the Edge Metadata Processing Capability on a session, or has not received this capability from its peer on that session, the speaker MUST NOT send any UPDATE on that session that carries the Edge Metadata Path Attribute.

If a BGP speaker receives an UPDATE carrying the Edge Metadata Path Attribute on a session for which the corresponding Edge Metadata Processing Capability was not successfully advertised by both peers for that AFI/SAFI, the receiver SHOULD ignore the Edge Metadata Path Attribute and process the remainder of the UPDATE according to local policy and the error-handling procedures specified in Section 9.

If a BGP speaker does not include the Edge Metadata Processing Capability in its BGP OPEN message for a specific BGP session, or if it does not receive the Edge Metadata Processing Capability from its peer on that session, it MUST NOT send any BGP UPDATE message on that session that bind the Edge Metadata Path Attribute to any prefix.

## 6. Service Metadata Propagation Scope

The propagation scope of the Edge Metadata Path Attribute needs careful consideration to ensure it does not inadvertently leak to other BGP domains. According to Section 3 of [ATTRIBUTE-ESCAPE], it is necessary for the Route Reflector (RR) to be upgraded to constrain the propagation scope when propagating the metadata path attributes. Therefore, the Edge Metadata Path Attribute originator sets the attribute as Non-transitive when sending the BGP UPDATE message to its corresponding RR. Non-transitive attributes are only guaranteed to be dropped during BGP route propagation by implementations that do not recognize them, ensuring that the Edge Metadata path attributes do not propagate beyond the intended scope.

The RR can append the NO-ADVERTISE well-known community to the BGP UPDATE message with the Edge Metadata Path Attribute when forwarding it to the ingress routers. This signals to the ingress nodes that the associated route's Metadata Path Attribute SHOULD NOT be further advertised beyond their scope. This precautionary measure ensures that the receiver of the BGP UPDATE message refrains from forwarding the received update to its peers, preventing the undesired propagation of the information carried by the Metadata Path Attribute.

### 6.1. AS-Scope SubTLV

To address the potential issue where the NO-ADVERTISE well-known community of the BGP UPDATE message can be dropped by some routers, a new AS-Scope Sub-TLV can be included in the Metadata Path Attribute to prevent the Metadata Path Attribute from being leaked to unintended Autonomous Systems (ASes). The AS-Scope Sub-TLV will enforce stricter control over the propagation of the metadata by associating it with specific AS numbers.

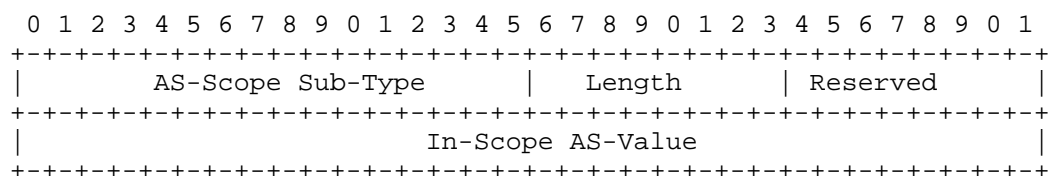


Figure 9: AS-Scope Sub-TLV

- AS-Scope Sub-Type (16 bits): 7 (specified in this document).
- Length (8 bits) Specifies the total length in octets of the value field, excluding the Sub-Type and Length fields. For the AS-Scope Sub-TLV, the value field consists of the Reserved field followed by one or more 32-bit In-Scope AS values. Therefore, the Length field is  $1 + 4*N$ , where N is the number of In-Scope AS values and N MUST be at least 1.
- Reserved (8 bits): These bits are reserved for future use and MUST be set to zero.
- In-Scope AS-Value (32 bits): One or more 32-bit AS values that

identify ASes within the intended metadata distribution domain. When multiple ASes are included, they are encoded as consecutive 32-bit AS values within the same AS-Scope Sub-TLV. A receiver considers the AS-Scope check successful if the local AS, or an AS recognized by local configuration as part of the same metadata distribution domain, matches any of the included AS values.

#### 6.1.1. AS-Scope Value Checking Procedure

When a router receives a BGP UPDATE message containing the AS-Scope Sub-TLV, it must perform the following steps to process the AS-Scope value:

An AS-Scope value of 0 is invalid and MUST NOT be used. A receiver that encounters an AS-Scope value of 0 MUST ignore that AS value. If all AS values carried in the AS-Scope Sub-TLV are invalid, the receiver MUST treat the AS-Scope check as failed.

When BGP confederations [RFC5065] are used, AS-Scope values refer to member-AS numbers for sessions within the confederation, and to the confederation identifier for sessions outside the confederation.

AS Recognition: The router will check the AS value in the AS-Scope Sub-TLV. If the AS value matches the local AS or a recognized AS in its configuration, the router will process the update as usual. If the AS value does not match or is not recognized, the router SHOULD NOT process the Edge Metadata Path Attribute values in the BGP UPDATE and SHOULD NOT propagate the received BGP UPDATE to other nodes. I.e., treat-as-withdraw behavior will be used.

#### Example Usage:

Consider a scenario where a router in AS 65001 advertises a BGP UPDATE message with the AS-Scope Sub-TLV set to AS 65001. When another router in AS 65002 receives this UPDATE, it will check the AS-Scope Sub-TLV value:

Since AS 65002 does not match the AS value 65001, the router in AS 65002 will drop the UPDATE, preventing the metadata from leaking into AS 65002.

This mechanism ensures that the metadata remains confined to the intended ASes, enhancing the security and control over the propagation of BGP metadata.

## 7. Policy Based Metadata Integration

This section describes how the information carried in the Edge Metadata Path Attribute can be incorporated into BGP route selection by local policy. The procedures in this section do not modify the base BGP decision process defined in [RFC4271]. Instead, they describe how local policy can use recognized Edge Metadata values when comparing candidate routes for services configured for metadata-aware route selection.

### 7.1. Policy Application Order

To remain consistent with Section 9.1.1 of [RFC4271], metadata-aware policy evaluation **MUST** be applied after LOCAL\_PREF has been set for iBGP routes, or after equivalent inbound policy has been applied for eBGP routes.

The use of Edge Metadata does not replace existing BGP routing policy. Rather, the Edge Metadata Path Attribute provides additional inputs that local policy **MAY** use when comparing candidate routes for selected services.

### 7.2. Metadata Selection by Local Policy

A deployment **MAY** use only a subset of the metadata attributes carried in the Edge Metadata Path Attribute. Which metadata attributes are considered, and for which services they are considered, is determined by local policy.

For example, one deployment may consider only the Service Delay Prediction Sub-TLV for latency-sensitive services, while another deployment may consider only availability-related or service-capability-related Sub-TLVs. A route that carries additional recognized metadata does not require all such metadata to be used in route selection.

If none of the recognized metadata carried by a route are selected by local policy for preference computation, the route is evaluated using ordinary BGP policy and tie-breaking procedures.

### 7.3. Policy-Based Preference Computation

For services configured for metadata-aware route selection, local policy **MAY** use one or more recognized metadata values carried in the Edge Metadata Path Attribute, together with other routing attributes, to derive a preference for each candidate route.

The procedure for combining recognized metadata with traditional BGP attributes is deployment specific and outside the scope of this document. The preference computation MAY be performed at a Route Reflector (RR), at an ingress node, or at another policy decision point within the same administrative domain.

When metadata-aware policy is applied to a set of candidate routes, the route with the most preferred policy outcome is selected. If two or more routes remain equally preferred after metadata-aware policy evaluation, the normal BGP tie-breaking procedures defined in [RFC4271] apply.

This document does not define a default ordering or weighting among Edge Metadata Sub-TLVs. When more than one recognized metadata Sub-TLV is present, the local policy MUST specify which Sub-TLVs are used and how their values are ordered, weighted, or otherwise combined for preference computation. In the absence of such local policy, a BGP speaker MUST NOT use the Edge Metadata Path Attribute to alter best-path selection and MUST evaluate the route using ordinary BGP policy and tie-breaking procedures.

#### 7.4. Policy Treatment of Routes with Degraded Metadata

Local policy MAY define threshold conditions for one or more metadata types. When the recognized metadata associated with a route indicates that such a threshold has been crossed, local policy MAY reduce the preference of that route or MAY treat the route as ineligible for metadata-aware service steering.

This document does not mandate a specific action for degraded metadata values. The action taken, if any, is determined by local policy. For example, local policy may de-prefer a route whose Service Delay Prediction exceeds a configured threshold, or a route whose availability-related metadata falls below a configured level.

If local policy excludes a route from metadata-aware service steering, the route MAY still remain valid for ordinary BGP reachability unless separate policy removes or suppresses that route.

#### 7.5. Tie Breaking and ECMP

After metadata-aware policy evaluation, if multiple candidate routes remain equally preferred, BGP tie-breaking proceeds according to [RFC4271].

If the decision process results in multiple equally preferred paths and the deployment permits Equal Cost Multi Path (ECMP), those paths MAY be installed in the forwarding plane according to existing BGP procedures and platform capabilities.

Different local policies may result in different selected paths; deployments that require consistent route selection across multiple decision points SHOULD configure consistent metadata-aware policy on those decision points.

## 8. Minimum Interval for Metrics Change Advertisement

### Route Churn Considerations

While the mechanism detailed in this document aims to provide dynamic metrics like Capacity Availability Index, Site Delay Prediction Index, Service Delay Prediction Index, and Raw Measurement to optimize path selection, it is essential to consider the broader implications of metric-induced churn. Particularly, in the context of routes used for BGP nexthop resolution (e.g., labeled unicast), frequent changes in these metrics can lead to significant churn not only for the prefixes carrying the data but also for dependent routes.

In normal operation, the metadata associated with a prefix is propagated along with BGP UPDATE messages as per standard BGP behavior. The advertisement interval is governed by the underlying BGP mechanisms and implementation behavior, including any configured MRAI timer. This document does not rely on a specific MRAI value for metadata update pacing. This document does not propose a new periodic advertisement mechanism independent of routing updates. If metadata attributes (e.g., compute availability, service locality) change, a BGP UPDATE is triggered accordingly. If there is no change to the advertised metadata, no additional UPDATE is sent, in order to avoid unnecessary update churn and to comply with BGP best practices. Any active or proactive refresh mechanisms for metadata would require explicit triggers and change detection mechanisms, which are outside the scope of this document.

Measurement-bearing Sub-TLVs, such as the Service Delay Prediction Sub-TLV and Raw Measurement Sub-TLV, can be derived from local measurements that change more frequently than BGP UPDATES should be advertised. The measurement collection interval is deployment specific and is independent of the interval at which changed metadata is advertised in BGP. Implementations SHOULD apply dampening, hysteresis, or threshold-based change detection before advertising updated Edge Metadata Path Attribute values, so that small or short-lived measurement fluctuations do not cause excessive BGP UPDATE churn.

This behavior is analogous to the impacts observed with RSVP auto-bandwidth, which can introduce considerable instability within a network. Such route churn can propagate through the network, causing a cascade of UPDATES and potential route flaps, thereby affecting overall network stability and performance.

To mitigate these effects, network operators SHOULD carefully manage the advertisement intervals of these dynamic metrics, ensuring they are set to avoid unnecessary churn. A default minimum interval of 30 seconds for advertising Edge Metadata Path Attribute changes is RECOMMENDED to balance responsiveness with stability, unless a deployment configures a different interval. However, in scenarios with higher sensitivity to route stability, operators may consider increasing this interval further to reduce the frequency of UPDATES.

Significant load changes at EC data centers can be triggered by short-term gatherings of UEs, like conventions, lasting a few hours or days. Therefore, a high metrics change rate can persist for hours or days.

## 9. Validation and Error Handling

The Edge Metadata Path Attribute is an optional non-transitive BGP Path attribute that carries metrics and metadata about the edge services attached to the egress router. The Edge Metadata Path Attribute, to be assigned by IANA, consists of a set of Sub-TLVs, and each Sub-TLV contains information for specific metrics of the edge services.

When more than one sub-TLV is present in a Metadata Path Attribute, they are processed independently. Suppose a Edge Metadata Path Attribute can be parsed correctly but contains a Sub-TLV whose type is not recognized by a particular BGP speaker; that BGP speaker MUST NOT consider the attribute malformed. Instead, it MUST interpret the attribute as if that Sub-TLV had not been present. Logging the error locally or to a management system is optional. If the route carrying the Edge Metadata path attribute is propagated with the attribute, the unrecognized Sub-TLV remains in the attribute.

If an Edge Metadata Path Attribute does not have any valid Sub-TLVs, or its Attribute Flags are inconsistent with the optional non-transitive definition of this attribute, the "attribute discard" procedure of [RFC7606] is applied.

## 10. Manageability Considerations

The edge service Metadata described in this document are only intended for propagating between ingress and egress routers of one single BGP Administrative Domain [RFC1136]. A single BGP Administrative Domain can consist of one AS or multiple ASes.

Only a small subset of services are expected to require the Edge Metadata Path Attribute. These are typically services for which metadata-aware route selection is beneficial. The domain in which such metadata is propagated is typically operated under a common administrative policy, even when the routers are supplied by different vendors.

Additional non-normative examples of deployment models and metadata-aware route-selection procedures are provided in Appendix C.

## 11. Security Considerations

The proposed edge service Metadata are advertised within the trusted domain of 5G LDN's ingress and egress routers. The ingress routers SHOULD NOT propagate the edge service Metadata to any nodes that are not within the trusted domain.

The Sub-TLV count limit specified in Section 4.1.3 helps reduce exposure to resource-exhaustion attacks caused by excessively large Edge Metadata Path Attributes.

To prevent the BGP UPDATE receivers (a.k.a. ingress routers in this document) from leaking the Edge Metadata Path Attribute by accident to nodes outside the trusted domain [ATTRIBUTE-ESCAPE], the following practice SHOULD be enforced:

- The Edge Metadata Path Attribute is optional non-transitive. Per [RFC4271], non-transitive Path Attributes are dropped during BGP route propagation by implementations that do not recognize them.
- Route Reflectors can append the NO-ADVERTISE well-known community to the BGP UPDATE message with Edge Metadata Path Attribute when forwarding to the ingress routers. By doing so, the Route Reflector signals to ingress nodes that the routes with the Edge Metadata Path Attribute SHOULD NOT be further advertised beyond their scope. This precautionary measure ensures that the receiver of the BGP UPDATE message refrains from forwarding the received UPDATE to its peers, preventing the undesired propagation of the information carried by the Edge Metadata Path Attribute.

BGP Route Filtering or BGP Route Policies [RFC5291] can also be used to ensure that BGP UPDATE messages with Edge Metadata Path Attribute attached do not get forwarded out of the administrative domain. BGP route filtering [RFC5291] allows network administrators to control the advertisements and acceptance of BGP routes, ensuring that specific routes do not leak outside the intended administrative domain. Here are the steps to achieve this:

- Use Route Filtering: Implement route filtering policies on ingress routers to restrict the propagation of BGP UPDATE messages carrying the Edge Metadata Path Attribute beyond the administrative domain. Access control lists (ACLs), prefix lists, or route maps can be used to filter the corresponding BGP routes for which the Edge Metadata Path Attribute is distributed from egress routers to ingress routers.
- Filter by Prefix: Use prefix filtering to specify which IP prefixes SHOULD be advertised to peers and which SHOULD be suppressed. This step ensures that only authorized routes are sent to external peers.
- Use Route Maps: Route maps provide a flexible way to filter and manipulate BGP route advertisements. You can create route maps to match specific conditions and then apply them to the BGP configuration.

## 12. IANA Considerations

### 12.1. Edge Metadata Path Attribute

IANA has done early allocation [RFC7120] of the codepoint 42 to the "Edge Metadata Path Attribute" in the "BGP Path Attributes" registry in the BGP Parameters registry group. The reference for this assignment is [this document].

Value	Description	Reference
42	Edge Metadata Path Attribute	[this document]

## 12.2. Edge Metadata Capability Code

IANA has assigned a Capability Code of 78 from the "BGP Capability Codes" registry in "Capability Codes registry group" for the Edge Metadata Capability in the BGP OPEN message.

Value	Description	Reference
78	Edge Metadata Capability	[This document]

## 12.3. Edge Metadata Path Attribute Sub-Types

IANA is requested to create a new sub-registry under the Edge Metadata Path Attribute registry as follows:

Name: Sub-TLVs under the "Edge Metadata Path Attribute"

Registration Procedure: Expert Review [RFC8126].

Detailed Expert Review procedure will be added per [RFC8126].

Reference: [this document]

Sub-Type	Description	Reference
0	reserved	[this document ]
1	Site Preference Index	[this document:4.2]
2	Site Physical Avail Index	[this document:4.3]
3	Service Delay Predication	[this document:4.4]
4	Raw Measurement	[this document:4.5]
5	Service-Oriented Capability	[this document:4.6]
6	Service-Oriented Available Resource	[this document:4.7]
7	AS-Scope	[this document:6.1]
8-65534	unassigned	
65535	reserved	[this document]

### 13. Contributors

Changwang Lin

New H3C Technologies

China

Email: linchangwang.04414@h3c.com

### 14. Acknowledgements

Acknowledgements to Jeff Haas, Tom Petch, Adrian Farrel, Alvaro Retana, Robert Raszuk, Sue Hares, Shunwan Zhuang, Donald Eastlake, Dhruv Dhody, Cheng Li, DongYu Yuan, and Vincent Shi for their suggestions and contributions.

### 15. References

#### 15.1. 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>>.
- [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>>.
- [RFC4786] Abley, J. and K. Lindqvist, "Operation of Anycast Services", BCP 126, RFC 4786, DOI 10.17487/RFC4786, December 2006, <<https://www.rfc-editor.org/info/rfc4786>>.
- [RFC5065] Traina, P., McPherson, D., and J. Scudder, "Autonomous System Confederations for BGP", RFC 5065, DOI 10.17487/RFC5065, August 2007, <<https://www.rfc-editor.org/info/rfc5065>>.
- [RFC5291] Chen, E. and Y. Rekhter, "Outbound Route Filtering Capability for BGP-4", RFC 5291, DOI 10.17487/RFC5291, August 2008, <<https://www.rfc-editor.org/info/rfc5291>>.
- [RFC5492] Scudder, J. and R. Chandra, "Capabilities Advertisement with BGP-4", RFC 5492, DOI 10.17487/RFC5492, February 2009, <<https://www.rfc-editor.org/info/rfc5492>>.
- [RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, DOI 10.17487/RFC5905, June 2010, <<https://www.rfc-editor.org/info/rfc5905>>.
- [RFC7120] Cotton, M., "Early IANA Allocation of Standards Track Code Points", BCP 100, RFC 7120, DOI 10.17487/RFC7120, January 2014, <<https://www.rfc-editor.org/info/rfc7120>>.
- [RFC7311] Mohapatra, P., Fernando, R., Rosen, E., and J. Uttaro, "The Accumulated IGP Metric Attribute for BGP", RFC 7311, DOI 10.17487/RFC7311, August 2014, <<https://www.rfc-editor.org/info/rfc7311>>.
- [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>>.

- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [RFC9012] Patel, K., Van de Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", RFC 9012, DOI 10.17487/RFC9012, April 2021, <<https://www.rfc-editor.org/info/rfc9012>>.
- [RFC9107] Raszuk, R., Ed., Decraene, B., Ed., Cassar, C., ㄒㄢ, E., and K. Wang, "BGP Optimal Route Reflection (BGP ORR)", RFC 9107, DOI 10.17487/RFC9107, August 2021, <<https://www.rfc-editor.org/info/rfc9107>>.

## 15.2. Informative References

- [ATTRIBUTE-ESCAPE] J. Haas, "BGP Attribute Escape", July 2023, <<https://datatracker.ietf.org/doc/draft-haas-idr-bgp-attribute-escape/>>.
- [RFC1136] Hares, S. and D. Katz, "Administrative Domains and Routing Domains: A model for routing in the Internet", RFC 1136, DOI 10.17487/RFC1136, December 1989, <<https://www.rfc-editor.org/info/rfc1136>>.
- [RFC4456] Bates, T., Chen, E., and R. Chandra, "BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)", RFC 4456, DOI 10.17487/RFC4456, April 2006, <<https://www.rfc-editor.org/info/rfc4456>>.
- [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>>.
- [RFC8799] Carpenter, B. and B. Liu, "Limited Domains and Internet Protocols", RFC 8799, DOI 10.17487/RFC8799, July 2020, <<https://www.rfc-editor.org/info/rfc8799>>.
- [TS.23.501-3GPP] 3rd Generation Partnership Project (3GPP), "System Architecture for 5G System; Stage 2, 3GPP TS 23.501 v2.0.1", December 2017.

## Appendix A. Service Delay Prediction Based on Load Measurement

When data centers detailed running status are not exposed to the network operator, historic traffic patterns through the egress routers can be utilized to predict the load to a specific service. For example, when traffic volume to one service at one data center suddenly increases a huge percentage compared with the past 24 hours average, it is likely caused by a larger than normal demand for the service. When this happens, another data center with lower-than-average traffic volume for the same service might have a shorter processing time for the same service.

Here are some measurements that can be utilized to derive the Service Delay Predication for a service ID:

- Total number of packets to the attached service instance (ToPackets);
- Total number of packets from the attached service instance (FromPackets);
- Total number of bytes to the attached service instance (ToBytes);
- Total number of bytes from the attached service instance (FromBytes);
- The actual load measurement to the service instance attached to an egress router can be based on one of the metrics above or including all four metrics with different weights applied to each, such as:

$$\text{LoadIndex} = w1 * \text{ToPackets} + w2 * \text{FromPackets} + w3 * \text{ToBytes} + w4 * \text{FromBytes}$$

Where  $w1/w2/w3/w4$  are between 0-1.  $w1 + w2 + w3 + w4 = 1$ ;

The weights of each metric contributing to the index of the service instance attached to an egress router can be configured or learned by self-adjusting based on user feedbacks.

The Service Delay Prediction Index can be derived from  $\text{LoadIndex}/24\text{Hour-Average}$ . A higher value means a longer delay prediction. The egress router can use the ServiceDelayPred sub-TLV to indicate to the ingress routers of the delay prediction derived from the traffic pattern.

Note: The proposed IP layer load measurement is only an estimate based on the amount of traffic through the egress router, which might not truly reflect the load of the servers attached to the egress

routers. They are listed here only for some special deployments where those metrics are helpful to the ingress routers in selecting the optimal paths.

## Appendix B. Service Metadata Influenced Decision Process

### B.1. Egress Router Behavior

Multiple instances of the same service could be attached to one egress router. When all instances of the same service are grouped behind one application layer load balancer, they appear as one single route to the egress router, i.e., the application loader balancer's prefix. Under this scenario, the compute metrics for all those instances behind one application layer balancer are aggregated under the application load balancer's prefix. In this case, the compute metrics aggregated by the Load Balancer are visible to the egress router as associated with the Load Balancer's prefix. However, how the application layer Load Balancers distribute the traffic among different instances is out of the scope of this document. When multiple instances of the same service have different paths or links reachable from the egress router, multiple groups of metrics from respective paths could be exposed to the egress router. The egress router can have preconfigured policies on aggregating various metrics from different paths and the corresponding policies in selecting a path for forwarding the packets received from ingress routers. The aggregated metrics can be carried in the BGP UPDATE messages instead of detailed measurements to reduce the entries advertised by the control plane and dampen the routes update in the forwarding plane. Upon receiving packets from ingress routers, the egress router can use its policies to choose an optimal path to one service instance. It is out of the scope of this document how the measurements are aggregated on egress routers and how ingress routers are configured with the algorithms to integrate the aggregated metrics with network layer metrics.

Many measurements could impact and correspondingly reflect service performance. In order to simplify an optimal selection process, egress routers can have preconfigured policies or algorithms to aggregate multiple metrics into one simple one to ingress routers. Though out of the scope of this document, an egress router can also have an algorithm to convert multiple metrics to network metrics, an IGP cost for each instance, to pass to ingress nodes. This decision-making process integrates network metrics computed by traditional IGP/BGP and the service delay metrics from egress routers to achieve a well-informed and adaptive routing approach. This intelligent orchestration at the edge enhances the service's overall performance and optimizes resource utilization across the distributed infrastructure. When the egress has merged the compute metrics from

the local sites behind it, it can include one or more aggregated compute metrics in the Metadata Path Attribute in the BGP UPDATE to the Ingress. Also, an identifier or flag can be carried to indicate that the metrics are merged ones. After receiving the routes for the Service ID with the identifier, the ingress would do the route selection based on pre-configured algorithms (see Section 3 of this document).

## B.2. Integrating Network Delay with the Service Metrics

As the service metrics and network delays are in different units, here is an exemplary algorithm for an ingress router to compare the cost to reach the service instances at Site-i or Site-j.

$$\text{Cost-i} = \min\left(w * \frac{\text{ServD-i} * \text{CP-j}}{\text{ServD-j} * \text{CP-i}} + (1-w) * \frac{\text{Pref-j} * \text{NetD-i}}{\text{Pref-i} * \text{NetD-j}}\right)$$

CP-i: Capacity Availability Index at Site-i. A higher value means higher capacity available.

NetD-i: Network latency measurement (RTT) to the Egress Router at the site-i.

Pref-i: Preference Index for Site-i, a higher value means higher preference.

ServD-i: Service Delay Predication Index at Site-i for the service, i.e., the ANYCAST address [RFC4786] for the service.

w: Weight is a value between 0 and 1. If smaller than 0.5, Network latency and the site Preference have more influence; otherwise, Service Delay and capacity availability have more influence.

When a set of service Metadata is converted to a simple metric, a decision process is determined by the metric semantics and deployment situations. The goal is to integrate the conventional network decision process with the service Metadata into a unified decision-making process for path selection.

## B.3. Integrating with BGP Route Selection

Not all metadata attributes specified in this document are intended for use in every deployment. Each deployment may choose to consider only a subset of the available metadata attributes based on its specific service requirements.

- Deployment-Specific Attribute Selection:

A deployment may prioritize only certain metadata attributes relevant to its operational needs. For example, one deployment might only use the Service Delay Prediction Index for latency-sensitive applications, while another might focus solely on the Capacity Availability Index to manage resource availability.

- Influence on BGP Decision Process:

The edge service Metadata influences next-hop selection differently from traditional BGP metrics (e.g., Local Preference, MED). Unlike a general next-hop metric that can affect many routes, edge service Metadata selectively impacts optimal next-hop selection for specific routes configured to consider these service-specific attributes. This targeted influence allows for optimized path selection without disrupting broader route decisions.

- Handling Degraded Metrics (Policy-Based):

If a service-specific metric degrades beyond a configured threshold (e.g., the Service Delay Prediction Index exceeds an acceptable delay threshold or the Capacity Availability Index drops below a required level), the ingress router will treat that route as ineligible for traffic steering. This is similar to a BGP route withdrawal, where the degraded route is deprioritized or ignored, even if traditional BGP attributes would otherwise favor it. This ensures that traffic is directed only to service instances that meet the defined performance criteria.

- Fallback to Non-Metadata Routes:

If no suitable routes with the required metadata are available, the BGP decision process defaults to traditional attribute evaluation [RFC4271], ensuring consistent routing even when metadata-specific paths are absent.

This approach provides flexibility and adaptability in routing decisions, allowing each deployment to apply relevant metadata attributes and enforce performance thresholds for improved service quality.

## Appendix C. Deployment Examples for Metadata-Aware Route Selection

This appendix provides non-normative examples of how a deployment may apply the procedures described in Section 7.

### C.1. Centralized RR Model

In a deployment where the Route Reflector (RR) is the primary policy decision point, the RR may apply metadata-aware local policy when selecting routes for reflection. In such a deployment, routers that rely on the RR for best-path selection receive routes that already reflect the policy outcome.

In deployments where the RR is responsible for pre-selecting routes, the RR may combine recognized Edge Metadata with traditional BGP attributes when determining the preferred route for a service. The RR can then reflect only the selected route to its client routers, such as ingress PEs, in accordance with local policy. This can help align reflected routes with service-specific requirements while limiting the number of routes distributed to clients.

Deployments using this model SHOULD consider Optimal Route Reflection (ORR) [RFC9107] so that route selection reflects the perspective of ingress routers rather than the physical location of the RR.

### C.2. Ingress-Node Decision Model

In some deployments, the RR may reflect multiple candidate routes, for example by using Add-Paths. In such a deployment, the ingress node receives those candidate routes and applies local metadata-aware policy to determine the preferred route for the selected service.

The ingress node may combine recognized metadata values with traditional BGP attributes when deriving route preference. This allows the ingress node to make service-specific routing decisions based on its local policy and on the metadata available for the candidate routes.

### C.3. Consistent Distributed Model

In a deployment where routers exchange iBGP routes directly in addition to receiving reflected routes, all participating nodes, including any RR, should apply consistent metadata-aware policy so that route selection remains aligned across the administrative domain.

In this model, the RR is not the sole policy decision point. Instead, each node that performs metadata-aware preference computation applies consistent policy to the same set of recognized metadata and routing attributes. This helps reduce the risk of inconsistent route selection among nodes that receive the same candidate routes.

#### C.4. Example Policy Weighting Approaches

A deployment may choose to assign greater weight to recognized metadata values than to traditional routing attributes, may weigh them equally, or may treat metadata only as a secondary refinement after traditional routing considerations. The weighting method is deployment specific and is not specified by this document.

For example, one deployment may emphasize service-delay-related metadata for latency-sensitive services, while another may emphasize availability-related or resource-related metadata. Another deployment may use metadata only after candidate routes have already been narrowed by traditional BGP policy. These examples are illustrative only and do not impose any required computation method.

#### Authors' Addresses

Linda Dunbar  
Futurewei  
Dallas, TX,  
United States of America  
Email: ldunbar@futurewei.com

Kausik Majumdar  
Oracle  
California,  
United States of America  
Email: kausik.majumdar@oracle.com

Cheng Li  
Huawei Technologies  
Beijing  
China  
Email: c.l@huawei.com

Gyan Mishra  
Verizon  
United States of America  
Email: gyan.s.mishra@verizon.com

Zongpeng Du  
China Mobile  
Beijing  
China

Email: [duzongpeng@chinamobile.com](mailto:duzongpeng@chinamobile.com)