

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 23 April 2026

L. Dunbar
A. Retana
Futurewei
20 October 2025

Metadata Constrained Distribution
draft-dunbar-idr-metadata-constrained-dist-00

Abstract

This document defines the Metadata-Filter (MDF) Subsequent Address Family Identifier (SAFI). A BGP speaker uses MDF to signal its lack of interest in receiving service metadata attributes on routes that carry specified Route Targets (RTs), while leaving reachability unchanged.

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 23 April 2026.

Copyright Notice

Copyright (c) 2025 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

| | |
|--|---|
| 1. Introduction | 2 |
| 2. Conventions used in this document | 3 |
| 3. Summary of Operation | 3 |
| 4. Capability Negotiation | 3 |
| 5. Metadata Filter NLRI Wire Format | 3 |
| 6. Error Handling | 4 |
| 7. Example | 4 |
| 8. Relationship to RTC (RFC 4684) | 4 |
| 9. Manageability and Operational Guidance | 4 |
| 9.1. Service-Class RT Plan | 5 |
| 9.2. Interplay with RTC and Import/Export Policy | 5 |
| 9.3. Migration and Staging | 5 |
| 9.4. Operational Telemetry (Recommended) | 6 |
| 10. IANA Considerations | 6 |
| 11. Security Considerations | 6 |
| 12. Normative References | 6 |
| Appendix A. Using Metadata-Filter in 5G Environments | 7 |
| A.1. Service-Class RTs and MDF | 8 |
| A.2. Operational Procedure (Example) | 8 |
| A.3. Benefits in 5G | 9 |
| A.4. Interoperability Notes | 9 |
| Authors' Addresses | 9 |

1. Introduction

Service Metadata can be added to BGP UPDATES to make path selections not only based on the routing cost but also the running environment of the edge services [I-D.ietf-idr-5g-edge-service-metadata]. These attributes can proliferate per service and churn rapidly. Ingress nodes may have constrained control-plane resources, and may not need nor be able to efficiently process large, fast-changing metadata on every route.

In typical iBGP topologies with Route Reflectors, edge nodes can advertise routes with service metadata without direct knowledge of which ingress nodes will receive and use information. Receivers therefore need a way to decline metadata for uninterested classes while continuing to receive reachability. The Metadata-Filter (MDF) SAFI provides a mechanism for a receiver to signal its lack of interest in service metadata for routes that carry specified Route Targets (RTs). In turn, the RR removes the specific attributes from the affected UPDATES and continues to process the message accordingly.

The signaling of a node's interest (or lack thereof) in metadata is dynamic. MDF allows ingress nodes adjust metadata delivery as service placement changes, while keeping reachability intact.

2. Conventions used in this document

The reader is expected to be familiar with the terminology defined in [I-D.ietf-idr-5g-edge-service-metadata].

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.

3. Summary of Operation

A node that doesn't need to receive service metadata for routes tagged with a given Route Target (RT) signals its peer by advertising an MDF NLRI (AFI=IPv4/IPv6, SAFI=TBD-MDF) that encodes the specific RT. Upon receiving this MDF NLRI, the receiver removes the service metadata attributes from any UPDATE that carries the RT when advertising to that peer, while leaving reachability and other attributes unchanged. The MDF state persists until the sender withdraws the MDF NLRI or the BGP session resets.

Support for Enhanced Route Refresh [RFC7313] is RECOMMENDED to facilitate on-demand resynchronization.

4. Capability Negotiation

A BGP speaker that is able to receive the MDF NLRI MUST use the Multiprotocol Extensions Capability Code [RFC4760] to advertise the corresponding (AFI, SAFI) pair (AFI=1|2, SAFI=TBD-MF).

5. Metadata Filter NLRI Wire Format

The MDF NLRI is encoded in MP_REACH_NLRI and MP_UNREACH_NLRI attributes [RFC4760] with (AFI=IPv4|IPv6, SAFI=TBD-MDF). The Length of Next Hop Network Address MUST be set to 0. The MDF NLRI is formatted as follows:

Length (1 octet): Number of octets that follow.

Flags (1 octet):

- * Reserved: MUST be transmitted as 0 and ignored when received.

Entity (12 octets): Route Target value, identical to the prefix

structure of the Route Target membership NLRI [RFC4684]: origin as (4) | route target (8).

6. Error Handling

Malformed MDF NLRI MUST be treated-as-withdraw [RFC7606]. The session MUST NOT be reset because of a malformed MDF NLRI. If the errors are persistent, an implementation MAY use the AFI/SAFI disable approach [RFC7606].

7. Example

A peer signals "omit metadata for RT=64512:100". The receiving peer advertises routes carrying RT=64512:100 to this peer without service metadata.

```
UPDATE
  Path Attributes:
    ORIGIN: IGP
    AS_PATH: (iBGP)
    MP_REACH_NLRI (AFI=IPv4, SAFI=TBD-MDF)
    NLRI:
      Metadata Filter NLRI:
        Length: 13
        Flags: 0
        Entity: RT 64512:100
```

8. Relationship to RTC (RFC 4684)

RTC [RFC4684] constrains which routes are sent to a peer based on RT interest. The MDF SAFI mirrors this architectural pattern but constrains which `_service metadata attributes_` are attached to routes that are sent. The two are complementary and can be deployed together: RTC limits route flooding and MDF limits metadata propagation.

9. Manageability and Operational Guidance

In this specification, Route Targets (RTs) can be used to delineate `_service classes_`, not merely membership. Each group of routes can have multiple RTs to, for example, identify a VPN or customer grouping, indicate reachability or constrain membership, or identify routes carrying particular service characteristics. MDF then keys on the service class to control delivery of service metadata.

9.1. Service-Class RT Plan

Operators SHOULD define a small, stable set of service classes per customer, application, or administrative domain. Advertised routes may be tagged with both:

- * a `_base RT_` (for normal reachability and membership), and
- * the `_service class_` that denotes the class whose routes may carry service metadata.

Nodes that `_use_` metadata simply do not install MDF and therefore receive the service metadata on those routes. Nodes that `_do not use_` metadata install `_MDF_` so the sender omits the metadata while leaving reachability unchanged.

`_Example (illustrative):_` A DC exporter advertises `_10.0.0.0/24_` with `_RT-VPN=64512:100_` and `_RT-ULL=64512:200_`, attaching the Metadata Path Attribute. The RR advertises the route `_with_` metadata to peers that do not have MDF[64512:200] and advertises the same route `_without_` metadata to peers that do.

9.2. Interplay with RTC and Import/Export Policy

If multiple RTs are used (as above), RTC SHOULD remain keyed to the base RT(s) so that reachability distribution is unaffected by MDF usage. The service class RT is used for MDF matching.

9.3. Migration and Staging

A pragmatic introduction plan is:

1. `_Define service class RTs_` and add them to exporter policy alongside the other existing RTs.
2. `_Enable MDF on RRs_`; verify that peers receive metadata unchanged.
3. `_Opt-out ingress nodes_` that do not use metadata by installing MDF; validate that reachability persists and metadata is omitted.
4. `_Broaden coverage_` to additional service classes only as needed; keep the service-class RT set small and well documented.

9.4. Operational Telemetry (Recommended)

While MDF focuses on RT assignment and filtering of metadata, implementations should expose minimal telemetry for validation: count of MDF entries per peer, advertisements with metadata omitted, and last-change timestamps.

10. IANA Considerations

IANA is requested to allocate a new SAFI from the "SAFI Values" registry:

Name: Metadata-Filter (MDF)
Reference: This document

11. Security Considerations

This document introduces no new security vulnerabilities beyond those discussed in [RFC4684] and [I-D.ietf-idr-5g-edge-service-metadata].

The Metadata Filter can reveal preference intent or limitations of specific nodes. To limit exposure, deployment should primarily occur in iBGP, and the peers that may advertise MDF entries should be limited. Omission of metadata does not affect reachability but can affect path selection at the receiver; operators should audit Metadata Filter policy and enable change logging. Ignoring an MDF NLRI may result in processing unnecessary metadata and cause undue resource consumption.

12. Normative References

[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>>.

[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>>.

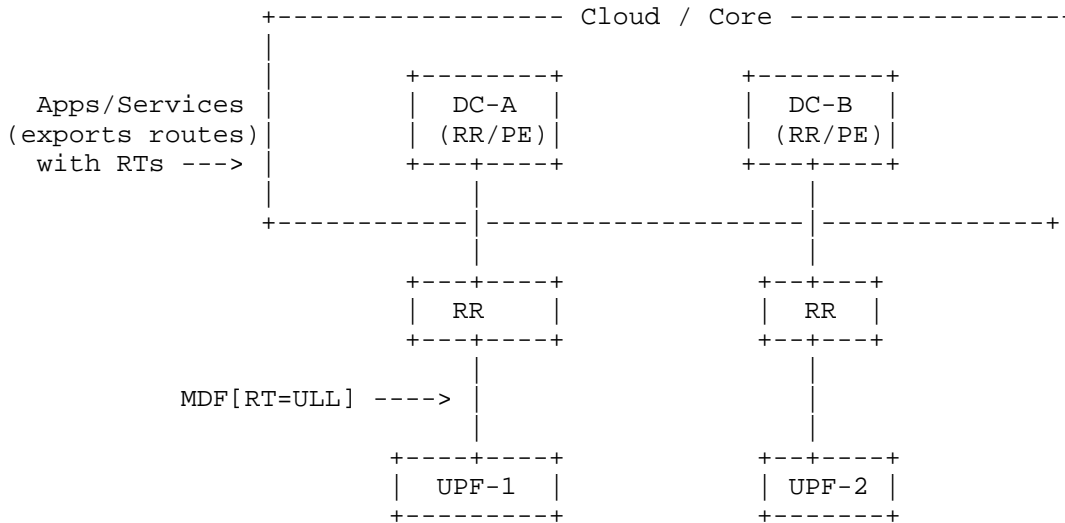
- [RFC4684] Marques, P., Bonica, R., Fang, L., Martini, L., Raszuk, R., Patel, K., and J. Guichard, "Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)", RFC 4684, DOI 10.17487/RFC4684, November 2006, <<https://www.rfc-editor.org/info/rfc4684>>.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, DOI 10.17487/RFC4760, January 2007, <<https://www.rfc-editor.org/info/rfc4760>>.
- [RFC7313] Patel, K., Chen, E., and B. Venkatachalapathy, "Enhanced Route Refresh Capability for BGP-4", RFC 7313, DOI 10.17487/RFC7313, July 2014, <<https://www.rfc-editor.org/info/rfc7313>>.
- [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>>.
- [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>>.

Appendix A. Using Metadata-Filter in 5G Environments

In 5G deployments, multiple User Plane Functions (UPFs) or edge gateways anchor PDU sessions near users while distributed edge data centers host application workloads. BGP advertises reachability for these workloads, and may carry service metadata so ingress nodes can consider service conditions in addition to routing metrics when selecting paths.

Service metadata can churn rapidly and inflate UPDATES if flooded to all peers. Many ingress nodes (e.g., UPFs handling best-effort traffic) neither need nor can efficiently process such rapidly changing attributes. The Metadata-Filter (MDF) SAFI allows a receiver to signal its lack of interest in metadata associated with specific Route Targets (RTs). Upon receiving an MDF NLRI, a Route Reflector (RR) omits the matching metadata for that peer while preserving reachability.

MDF signaling is dynamic: ingress nodes can add or withdraw MDF entries as service placement evolves, allowing networks to adapt metadata delivery without impacting route availability.



- DC-A/DC-B advertise routes tagged with a base RT (VPN/customer) and a service-class RT (e.g., RT=ULL).
- UPF-1 sends MDF NLRI for RT=ULL: RR omits metadata on routes with RT=ULL when advertising to UPF-1.
- UPF-2 does not send MDF: it receives routes with metadata unchanged.

Figure 1: Illustrative 5G Topology with MDF-Scoped Metadata Delivery

A.1. Service-Class RTs and MDF

Operators define a small, stable set of `_service-class RTs_` to delineate which groups of routes may carry service metadata (e.g., ultra-low-latency vs. best-effort). Exporters tag routes with both a base RT (for reachability/membership) and a service-class RT. MDF then keys on the service-class RT to control metadata delivery, while RTC (RFC 4684) and normal import/export policy remain keyed to the base RT so reachability is unaffected.

A.2. Operational Procedure (Example)

1. **Define service classes:* Choose a minimal set of RTs representing classes that may carry metadata (e.g., RT-ULL, RT-VID, RT-ML). Document ownership and intended use.
2. **Tag exports:* Data center exporters attach a base RT (VPN/customer) and a service-class RT to the same NLRI when metadata may accompany the route.
3. **Enable MDF on RRs:* RRs support the MDF SAFI and omit metadata per-peer when an MDF NLRI matches the service-class RT.
4. **Ingress selection:* UPFs/ingress nodes that do not use metadata advertise MDF NLRIs for the relevant service-class RTs; nodes that need metadata do not send MDF.

5. **Adjust dynamically:* As UE placement or service location changes, ingress nodes add/withdraw MDF entries to tune metadata reception over time.
6. **Telemetry (recommended):* Expose per-peer MDF entry counts, "metadata omitted" counters, and last-change timestamps to validate behavior.

A.3. Benefits in 5G

- * **Control-plane scale:* Limits fast-changing metadata propagation to UPFs and routers directly attached to UPFs, reducing UPDATE size and processing load on Route Reflectors and Provider Edge routers while preserving full reachability.
- * **Service agility:* Supports dynamic changes in metadata subscription as new UEs (for example, drones or autonomous vehicles) roam into or away from UPFs. When a UE moves to a new UPF, that UPF can dynamically express interest in receiving metadata needed for optimal path selection; when the UE leaves, the UPF withdraws its interest, preventing unnecessary metadata distribution.
- * **Operational safety:* Receiver-driven and RT-scoped; enables incremental rollout without impacting route propagation for other peers or service classes.

A.4. Interoperability Notes

MDF complements RTC: RTC constrains `_route_` flooding; MDF constrains `_metadata_` attachment to those routes for specific peers `:contentReference[oaicite:7]{index=7}`. MDF is carried in `MP_REACH/MP_UNREACH` with `AFI=IPv4|IPv6` and a dedicated SAFI; the Next Hop length is 0 for MDF NLRIs.

Authors' Addresses

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

Alvaro Retana
Futurewei
United States of America
Email: aretana@futurewei.com