idr
Internet-Draft
Intended status: Standards Track
Expires: 21 February 2026

J. Haas
J. Scudder
HPE

20 August 2025

BGP-4 Path Attribute Filtering Capability
draft-haas-idr-path-attribute-filtering-02

Abstract

Path Attributes in the BGP-4 protocol (RFC 4271) carry data associated with BGP routes. Many of the Path Attributes carried in BGP are intended for limited scope deployment. However, the extension mechanism defined by BGP that carries these attributes often carries them farther than necessary, sometimes with unfortunate results.

This document defines a mechanism using BGP Capabilities (RFC 5492) that permits eBGP speakers to determine what Path Attributes should be permitted to cross external BGP routing boundaries.

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 21 February 2026.

Copyright Notice

Table of Contents

1. Introduction

A BGP Route (Section 1.1 of [RFC4271]) is a tuple consisting of a set of Path Attributes (Section 5 of [RFC4271]) and sets of network reachability carried as Network Layer Reachability Information (NLRI). Some of these Path Attributes are defined as part of the core BGP-4 protocol. Path Attributes are the main extension mechanism defined by BGP, and may be scoped as "transitive" or "non-transitive."

Non-Transitive Path Attributes require the BGP speaker to understand the attribute in order to determine if it will be locally used and perhaps later propagated to additional BGP speakers. Unrecognized non-transitive Path Attributes are discarded by the receiving BGP speaker.

Transitive Path Attributes, when not understood by the receiving BGP speaker, are required to be propagated to other BGP speakers.

Some Path Attributes defined by BGP extensions are intended to be used in limited scopes, such as a single BGP Autonomous System (AS). When such attributes are distributed beyond the expected scope, this is called an "Attribute Escape" [I-D.haas-idr-bgp-attribute-escape].

Such attribute escapes may lead to improper BGP protocol behavior when received outside of their expected scope, and may lead to incorrect forwarding, or be a serious security consideration.

This document defines a mechanism exchanged through BGP Capabilities [RFC5492] where BGP speakers can more appropriately scope both Path Attributes to prevent escape, and to limit the distribution of routes that carry escaped Path Attributes.

1.1. Specification of Requirements

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. BGP Path Attribute Filtering Capability

The BGP Path Attribute Filtering Capability is encoded as follows:

* Capability Code of (TBD).

* Capability Length of 0..32 octets.

* Capability Value contains a bit-string padded to an octet boundary where a bit is set if this BGP speaker considers the corresponding BGP Path Attribute from the remote BGP speaker to be unwanted. Bit number 0 is the most significant bit of the first octet, bit number 1 is the second most significant bit of the first octet, and so on.

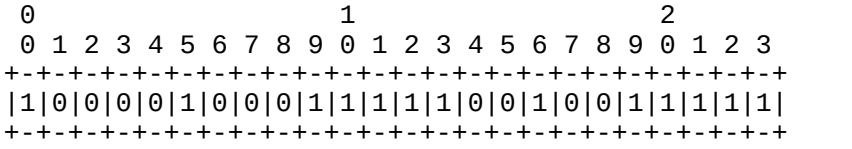

In this example, bits are clear for Path Attributes:

Origin (1),
AS_PATH (2),
NEXT_HOP (3),
MULTI_EXIT_DISCR (4),
ATOMIC_AGGREGATE (6),
AGGREGATOR (7),
COMMUNITIES (8),
MP_REACH_NLRI (14),
MP_UNREACH_NLRI (15),
AS4_PATH (17),
AS4_AGGREGATOR (18).

Other path attributes through attribute 23 are unwanted.
Path attributes 24 and beyond are accepted.

Figure 1: Example encoding for Capability Value

Bits 1, 2, 3, 6, 7, 14, 15, 17, 18 MUST be clear (value 0), because support for [RFC4271], [RFC4760], and [RFC6793] procedures are required when this specification is in use.

Any bit not explicitly represented (e.g., bits 24 and beyond in the above example) is deemed to be clear (value 0). That is, the default is to accept any path attribute not explicitly unwanted.

3. Operation

A clear (value 0) bit in the Path Attribute Filtering capability indicates that the BGP speaker advertising it is willing to accept the corresponding Path Attribute and will process it according to the normal rules of the BGP protocol and the attribute in question.

A set (value 1) bit in the Path Attribute Filtering capability indicates that the BGP speaker advertising it is not willing to accept the corresponding Path Attribute. We refer to such Path Attributes as "unwanted Path Attributes".

A BGP speaker MUST NOT send an unwanted Path Attribute to its peer. (Of course, this expectation will be met only by BGP speakers that support this specification; therefore a BGP speaker that implements this specification SHOULD be prepared for the possibility it will receive unwanted Path Attributes; this is discussed below.)

One strategy to accomplish the above requirement is for the BGP speaker to not advertise BGP routes containing the unwanted Path Attribute in question. This might require a withdraw to be sent instead. This is similar to treat-as-withdraw as defined in [RFC7606].

Another strategy that could be used, when appropriate for the procedure covering a given BGP Path Attribute, is for the BGP speaker to remove the unwanted Path Attributes when it distributes the route to the remote BGP speaker. This is similar to the Attribute Discard behavior defined in [RFC7606].

Receiving BGP speakers SHOULD filter routes or discard unwanted Path Attributes if they are incorrectly sent by the remote BGP speaker. Minimally, a receiving BGP speaker receiving an unwanted Path Attribute SHOULD use treat-as-withdraw procedures. Receiving BGP speakers MAY accept the route and discard the unwanted Path Attribute if permitted to by local configuration.

4. Selecting supported Path Attributes

Implementations MUST, as described in Figure 1, clear the bits covering required core eBGP Path Attributes.

Common BGP features that are defined for Internet use SHOULD be clear by default between two BGP speakers. These include:

* Communities (8)
* Extended Communities (16)
* Large BGP Communities (32)
* BGPsec_Path (33)
* Only To Customer (OTC) (35)

BGP features required to support a given AFI/SAFI MUST also be clear when that address family is configured. An example of this is the BGP-LS attribute (29) when the BGP-LS feature is in use.

5. Error Handling

If the received Capability Length for the Path Attributes Filtering Capability is greater than 32, the filtering capability MUST be ignored and treated as if not received by the BGP speaker.

To support core BGP features, Section 2 requires that bits 1, 2, 3, 6, 7, 14, 15, 17, 18 be clear (value 0). A BGP speaker receiving a Path Attribute Filtering Capability with these bits set (value 1) MAY reject the BGP session. If it does so, it uses a NOTIFICATION message with an error subcode of "Unsupported Capability".

6. Operational Considerations

6.1. Impacts on Path Attribute Transitivity

The feature described in this document does not change the semantics of when a Path Attribute is intended to be transitive per RFC-4271 definition. However, it does act as a policy to limit the distribution of routes containing a transitive Path Attribute, or may cause that attribute to be filtered.

6.2. Impacts on Incremental Deployment of New Features

eBGP speakers using this features must be cognizant of the impact their filtering policies will have on the incremental deployment of new BGP features.

6.3. Path Attribute Content Filtering Considerations

Path Attributes originally defined for use solely for one AFI/SAFI may later be updated to be applicable for other AFI/SAFIs, including those used by the Internet. Similarly, some Path Attribute features may be internally extensible in a way where filtering choices may be difficult to characterize using the coarse filtering feature defined by this document.

Documents defining new BGP Path Attributes MUST discuss their filtering, transitivity, and attribute escape considerations, and adopt strategies to limit unintentional escape of Path Attributes, or scoped sub-features of a Path Attribute. At a minimum, the new attribute’s registration must provide values for the "Should Filter By Default" and "Filtering Profile" columns (see Section 7). This is, however, a minimum. If the attribute has needs more nuanced propagation control than the all-or-nothing (per attribute type code) behavior offered by the present specification, it must be specified and handled on a case-by-case basis.

6.4. Operational Visibility

eBGP speakers that do not propagate a route on transmit or do not accept a route on reception due to unwanted Path Attributes SHOULD provide operationally visible feedback for this filtering behavior. Using terminology consistent with BMP ([RFC7854], et seq.) or the BGP YANG model ([I-D.ietf-idr-bgp-model]), routes that are not sent may have visibility in the Pre-Policy Adj-Ribs-Out view, but not be available in the Post-Policy Adj-Ribs-Out view. Updates to BMP or BGP YANG may provide operational visibility to this filtering. eBGP speakers that do not discard received routes with unwanted Path Attributes may be able to provide similar visibility in the Pre- and Post-Policy Adj-Ribs-In views.

Since utilizing treat-as-withdraw for routes received with unwanted Path Attributes is a likely implementation choice, implementations SHOULD provide operational visibility when treat-as-withdraw is done, via logging or aggregate statistics, on a per-BGP neighbor basis. Such logging or aggregate statistics should include the unwanted Path Attribute codes that cause routes to be discarded.

For some unwanted Path Attributes, an implementation may choose to use attribute discard for the unwanted Path Attributes rather than using treat-as-withdraw for the entire route. Implementations SHOULD provide operational visibility when attribute discard is done, via logging or aggregate statistics, on a per-BGP neighbor basis. Such logging or aggregate statistics should include the Path Attribute codes that are discarded.

Future extensions to BMP or BGP YANG may be proposed to support these operational considerations.

7. IANA Considerations

This document requests a new BGP Capability Code to be allocated from the First Come First Served range of the Capability Codes registry. The description should be "Path Attribute Filtering", and the reference should be this document.

This document requests that the "BGP Path Attributes" registry be updated to add two new columns, called "Should Filter By Default" and "Filtering Profile". These columns are to be seeded with the values shown in Table 1. This document should be added as a reference for the registry. Authors seeking guidance for how to populate these columns should refer to Section 6.3 and Section 10.

8. Security Considerations

The motivation for this feature is to attempt to address the numerous BGP security implications where BGP Path Attributes propagate beyond their intended scope.

The definition of a feature that limits the distribution of BGP Path Attributes unfortunately moves BGP's default behavior away from "distribute unknown things easily" and thus hampers incremental deployment of new features. However, operators have already begun indiscriminate filtering of Path Attributes they do not themselves require. This feature attempts to provide a more flexible negotiated mode to permit such filtering while at the same time not completely precluding incremental deployment of new features.

9. Recommendations for Future BGP Extensions

When a new BGP extension defines a new BGP Path Attribute, the specifying document MUST define its expected filtering profile for routes containing the new Path Attribute. The following are suggested propagation policies for new features:

* Default deny: Without explicit configuration from the sending BGP speaker and the receiving BGP speaker, routes containing the new Path Attribute are considered unwanted. Implementations supporting the Path Attribute Filtering Capability SHOULD set the bit for this Path Attribute to 1 (unwanted) by default. The sending BGP speaker MUST NOT advertise routes containing such a Path Attribute to the remote BGP speaker. BGP speakers receiving routes containing the unwanted Path Attribute SHOULD NOT accept the route into its Adj-Rib-In for that peer (treat-as-withdraw).

* Default discard: Without explicit configuration from the sending BGP speaker and the receiving BGP speaker, routes containing the new Path Attribute are considered unwanted. Implementations supporting the Path Attribute Filtering Capability SHOULD set the bit of this Path Attribute to 1 (unwanted) by default.

  The sending BGP speaker MUST use one of the following two options to avoid sending the unwanted Path Attribute: it can either send the route to the remote speaker after discarding the offending Path Attribute (attribute discard) or it can avoid sending the route to such a peer.

The receiving BGP speaker SHOULD NOT accept routes with the offending Path Attribute into its Adj-Rib-In for that peer (treat-as-withdraw) or MAY discard the offending Path Attribute (attribute-discard). The choice of the desired behavior to discard or to not propagate will depend on the operational expectations for the feature.

* AFI/SAFI conditional: When the specified AFI/SAFI is configured for the BGP session and the Path Attribute Filtering Capability is supported, the filtering capability bit SHOULD be set to 0 (wanted) by default. When that AFI/SAFI is not configured, the defining specification should detail whether the default filtering policy is wanted and the filtering capability bit for that Path Attribute MUST be set accordingly.

* Default permit: Without explicit configuration from the sending BGP speaker and the receiving BGP speaker, routes containing the new Path Attribute are considered wanted. Implementations supporting the Path Attribute Filtering Capability SHOULD set the bit of this Path Attribute to 0 (wanted) by default.

10. Recommendations for Filtering Path Attributes

The table below provides suggested default filtering behaviors for many BGP Path Attributes. In addition to "yes" and "no", which indicate that an implementation SHOULD either filter, or not filter, the given attribute by default, "never" indicates that an implementation MUST NOT filter the given attribute under any circumstances, and "-" indicates that no recommendation is made.

The table also recommends filtering profiles, as defined in Section 9. Where no filtering recommendation is made, no profile is suggested either.

It is not recommended that BGP Path Attributes that are deprecated should be filtered by default. This permits their long term reassignment and re-use. Operators with a strong filtering policy may consider filtering these to block routes from older implementations of these features that may still be deployed.

Upon publication of this document as an RFC, this table will become historic and the authoritative source of recommendations will be the relevant IANA registry (Section 7).

Code Point	Name	Should Filter By Default	Filtering Profile
0	Reserved	Yes	Default deny
1	ORIGIN [RFC4271]	Never	Default permit
2	AS_PATH [RFC4271]	Never	Default permit
3	NEXT_HOP [RFC4271]	Never	Default permit
4	MULTI_EXIT_DISC [RFC4271]	No	Default permit
5	LOCAL_PREF [RFC4271]	Yes	Default discard
6	ATOMIC_AGGREGATE [RFC4271]	No	Default permit
7	AGGREGATOR [RFC4271]	No	Default permit
8	COMMUNITIES [RFC1997]	No	Default permit
9	ORIGINATOR_ID [RFC4456]	Yes	Default discard
10	CLUSTER_LIST [RFC4456]	Yes	Default discard
11	DPA (deprecated) [RFC6938]	-	
12	ADVERTISER (historic) (deprecated) [RFC1863][RFC4223][RFC6938]	-	
13	RCID_PATH / CLUSTER_ID (Historic) (deprecated) [RFC1863][RFC4223][RFC6938]	-	

14	MP_REACH_NLRI [RFC4760]	Never	Default permit
15	MP_UNREACH_NLRI [RFC4760]	Never	Default permit
16	EXTENDED COMMUNITIES [Eric_Rosen][draft-ramachandra-bgp-ext-communities-00][RFC4360]	No	Default permit
17	AS4_PATH [RFC6793]	Never	Default permit
18	AS4_AGGREGATOR [RFC6793]	Never	Default permit
19	SAFI Specific Attribute (SSA) (deprecated) [Gargi_Nalawade][draft-kapoor-nalawade-idr-bgp-ssa-00][draft-nalawade-idr-mdt-safi-00][draft-wijnands-mt-discovery-00]	-	
20	Connector Attribute (deprecated) [RFC6037]	-	
21	AS_PATHLIMIT (deprecated) [draft-ietf-idr-as-pathlimit-02]	-	
22	PMSI_TUNNEL [RFC6514]	Yes	Default deny
23	Tunnel Encapsulation [RFC9012]	Yes	Default deny
24	Traffic Engineering [RFC5543]	Yes	Default discard
25	IPv6 Address Specific Extended Community [RFC5701]	No	Default permit
26	AIGP [RFC7311]	Yes	Default discard
27	PE Distinguisher Labels [RFC6514]	Yes	Default deny

28	BGP Entropy Label Capability Attribute (deprecated) [RFC6790][RFC7447]	-	
29	BGP-LS Attribute [RFC9552]	Yes	AFI/SAFI Conditional
30	Deprecated [RFC8093]	-	
31	Deprecated [RFC8093]	-	
32	LARGE_COMMUNITY [RFC8092]	No	Default permit
33	BGPsec_Path [RFC8205]	Never	Default permit
34	BGP Community Container Attribute (TEMPORARY - registered 2017-07-28, extension registered 2024-08-22, expires 2025-07-28) [draft-ietf-idr-wide-bgp-communities-11]	No	Default permit
35	Only to Customer (OTC) [RFC9234]	Never	Default permit
36	BGP Domain Path (D-PATH) (TEMPORARY - registered 2019-07-08, extension registered 2025-06-06, expires 2026-07-08) [draft-ietf-bess-evpn-ipvpn-interworking-10]	Yes	Default deny
37	SFP attribute [RFC9015]	Yes	Default deny
38	BFD Discriminator [RFC9026]	Yes	Default discard
39	BGP Next Hop Dependent Capabilities (NHC) (TEMPORARY - registered 2022-12-20, extension registered 2024-12-10, expires 2025-12-20) [draft-ietf-idr-entropy-label-13]	Yes	Default discard
40	BGP Prefix-SID [RFC8669]	Yes	Default

			discard
41	BIER [RFC9793]	Yes	Default deny
42	Edge Metadata Path Attribute (TEMPORARY - registered 2025-04-23, expires 2026-04-23) [draft-ietf-idr-5g-edge-service-metadata-27]	Yes	
128	ATTR_SET [RFC6368]	Yes	Default deny
129	Deprecated [RFC8093]	-	
241	Deprecated [RFC8093]	-	
242	Deprecated [RFC8093]	-	
243	Deprecated [RFC8093]	-	
255	Reserved for development [RFC2042]	Yes	Default discard

Table 1

Acknowledgements

The authors would like to thank the following individuals whose comments contributed to the refinement of this document: Donatas Abraitis, Ignas Bagdonas, Bruno Decraene, David Lamparter, Robert Raszuk.

References

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/rfc/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/rfc/rfc4271>.

[RFC5492] Scudder, J. and R. Chandra, "Capabilities Advertisement with BGP-4", RFC 5492, DOI 10.17487/RFC5492, February 2009, <https://www.rfc-editor.org/rfc/rfc5492>.

[RFC6793] Vohra, Q. and E. Chen, "BGP Support for Four-Octet Autonomous System (AS) Number Space", RFC 6793, DOI 10.17487/RFC6793, December 2012, <https://www.rfc-editor.org/rfc/rfc6793>.

[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/rfc/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/rfc/rfc8174>.

Informative References

[I-D.haas-idr-bgp-attribute-escape]
Haas, J., "BGP Attribute Escape", Work in Progress, Internet-Draft, draft-haas-idr-bgp-attribute-escape-03, 9 April 2025, <https://datatracker.ietf.org/doc/html/draft-haas-idr-bgp-attribute-escape-03>.

[I-D.ietf-idr-bgp-model]
Jethanandani, M., Patel, K., Hares, S., and J. Haas, "YANG Model for Border Gateway Protocol (BGP-4)", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-model-18, 21 October 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-model-18>.

[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/rfc/rfc4760>.

[RFC7854] Scudder, J., Ed., Fernando, R., and S. Stuart, "BGP Monitoring Protocol (BMP)", RFC 7854, DOI 10.17487/RFC7854, June 2016, <https://www.rfc-editor.org/rfc/rfc7854>.

Authors’ Addresses

Jeffrey Haas
HPE
1133 Innovation Way
Sunnyvale, CA 94089
United States of America
Email: jhaas@juniper.net

John Scudder
HPE
1133 Innovation Way
Sunnyvale, CA 94089
United States of America
Email: jgs@juniper.net