

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 26 February 2026

C. Cardona
P. Lucente
NTT
P. Francois
INSA-Lyon
Y. Gu
Huawei
T. Graf
Swisscom
25 August 2025

BMP Extension for Path Status TLV
draft-ietf-grow-bmp-path-marking-tlv-04

Abstract

The BGP Monitoring Protocol (BMP) provides an interface for obtaining BGP path information, which is conveyed through BMP Route Monitoring (RM) messages. This document specifies a BMP extension to convey the status of a path after being processed by the BGP process.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 RFC 2119 [RFC2119] RFC 8174 [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 26 February 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. Path Status TLV encoding	3
3. IANA encoding of Path Status TLV	4
3.1. IANA path status codes	4
3.2. IANA reason codes	5
4. Implementation notes	6
4.1. Configuration of BMP path marking	7
4.2. Scalability and churn	7
4.3. Paths with no markings	7
4.4. Path markings applicability and consistency	8
4.4.1. Significance of status and origin RIBs	8
4.5. Multiple TLVs assigned to the same route.	8
4.6. Enterprise-specific status	8
4.7. Multiple reason codes	9
5. Acknowledgments	9
6. IANA Considerations	9
7. Security Considerations	9
8. References	9
8.1. Normative References	9
8.2. Informative References	11
Authors' Addresses	11

1. Introduction

Multiple paths with different path status, (e.g., the "best path", "backup path", "invalid", and so on), may co-exist for a given prefix in the BGP RIBs after being processed by the BGP decision process. The path status information is not carried in the BGP Update Message RFC4271 [RFC4271] or in the BMP Route Monitoring Message RFC7854 [RFC7854].

External systems can use the path status for various applications. For example, operators commonly use path status during troubleshooting. Having such status stored and tracked enables the development of tools that facilitate this process. Optimization systems can consider path status in their process, e.g., as a validation source (since it can compare the calculated state to the actual outcome of the network, such as primary and backup path). Moreover, path status information can complement other centralized sources of data. For example, flow collectors can leverage it to identify the exact forwarding paths, yielding more accurate traffic matrices than existing methods.

This document defines a Path Status TLV to convey the BGP path status to a BMP server. The BMP Path Status TLV is carried in the BMP Route Monitoring (RM) Message RFC7854 [RFC7854].

2. Path Status TLV encoding

The path status TLV follows the common header defined in [I-D.ietf-grow-bmp-tlv] and [I-D.ietf-grow-bmp-tlv-ebit].




Figure 2: Encoding of Path Status TLV

- * The common TLV header that can encode IANA-registered TLV or Enterprise-specific markings using [I-D.ietf-grow-bmp-tlv-ebit].
- * Path Status (4 Octets): indicates the path status of the BGP Update PDU encapsulated in the RM Message. Path status are encoded using a bitmap where each bit position encodes a specific role of the path. Multiple bits may be set when multiple path status apply to a path. All zeros are reserved for paths with no marking.
- * Reason Code (2 Octets, optional): indicates the reason of the path status indicated in the Path Status field. The reason code field is optional. If no reason code is carried, this field is not included. If a reason code is carried, the reason code is indicated by a two-byte value.

3. IANA encoding of Path Status TLV

3.1. IANA path status codes

Value	Path Type
0x00000001	Invalid
0x00000002	Best
0x00000004	Nonselected
0x00000008	Primary
0x00000010	Backup
0x00000020	Non-installed
0x00000040	Best-external
0x00000080	Add-Path
0x00000100	Filtered in inbound policy
0x00000200	Filtered in outbound policy
0x00000400	Stale
0x00000800	Suppressed

Table 1: IANA-Registered Path Status

Table 1 includes a list of IANA status codes. This list might be extended. An explanation of each of the types is included next:

- * An invalid route is a route that does not enter the BGP decision process as indicated in Section 9.1.2 of RFC4271 [RFC4271].
- * The best route is defined in Section 9.1 of RFC4271 [RFC4271].
- * Nonselected routes are those that are not selected in the BGP decision process. Backup routes are considered nonselected, while the best and primary routes are not considered as nonselected.

- * A primary route is a path used for traffic forwarding. See draft-ietf-rtgwg-bgp-pic [I-D.ietf-rtgwg-bgp-pic]. A prefix can have more than one primary path when multipath is configured draft-lapukhov-bgp-ecmp-considerations [I-D.lapukhov-bgp-ecmp-considerations]. The best path is also a primary path.
- * A backup path is installed in the RIB, but it is not used until some or all primary paths become unreachable. Backup paths are used for fast convergence in the event of primary path failures.
- * A non-installed path refers to the route that is not installed into the IP routing table.
- * The best external path is defined in draft-ietf-idr-best-external [I-D.ietf-idr-best-external].
- * The advertisement of multiple paths for the same address prefix without the new paths implicitly replacing any previous ones, the add-path status is applied [RFC7911].
- * Filtered in inbound policy routes are those that are filtered in the Adj-RIB-In policy
- * Filtered in outbound policy routes are those that are filtered in the Adj-RIB-Out policy
- * Stale routes refer to paths which have been declared stale by the BGP Graceful Restart mechanism, as described in Section 4.1 of [RFC4724].
- * Suppressed routes refer to a path which has been declared suppressed by the BGP Route Flap Damping mechanism as described in Section 2.2 of [RFC2439].

3.2. IANA reason codes

Table 2 includes a list of IANA reason codes. This list can be extended in future documents. This document includes a brief explanation of each code and the path status they are directed to explain. Please see Section 4.4 for notes on potentially inconsistencies on the path marking data

- * Invalid routes due to AS loop and unresolvable nexthop are defined in Section 9.1.2 of [RFC4271]. These codes target routes of type "Invalid".

- * The reason codes starting with 'not preferred' are aimed at paths not selected as best, and describe the reason they were ranked lower in the decision process. AIGP is explained in RFC7311 [RFC7311]. The rest of the codes are described in Section 9.1.2.2. of [RFC4271]

Value	Reason Code
0x0001	Invalid due to AS loop
0x0002	Invalid due to unresolvable nexthop
0x0003	Not preferred for local preference
0x0004	Not preferred for AS Path Length
0x0005	Not preferred for origin
0x0006	Not preferred for MED
0x0007	Not preferred for peer type
0x0008	Not preferred for IGP cost
0x0009	Not preferred for router ID
0x000A	Not preferred for peer address
0x000B	Not preferred for AIGP

Table 2: IANA-Registered Reason Codes

4. Implementation notes

The BMP path marking TLV remains optional within BMP implementations.

An implementation of the BMP path marking TLV may not fully support marking of all status defined in Table 1 or any future extensions. Similarly, an implementation may choose to support the inclusion of the reason code (for which support is also optional), without necessarily incorporating any of the reason codes defined in Table 2 or future extensions.

This document refrains from defining mechanisms for signaling the status or reason codes an implementation supports. This could be established through external means (e.g. documentation) or potentially addressed in a subsequent document.

The remainder of this section covers additional points related to the implementation of the BMP Path marking TLV.

4.1. Configuration of BMP path marking

Implementations supporting the BMP path marking TLV should provide an option for enabling/disabling the Path Marking TLV over BMP monitoring sessions. Furthermore, the configuration options for this TLV should provide the means to enable/disable the transmission of reason codes, if the reason codes are supported by the implementation.

4.2. Scalability and churn

The Path Marking TLV introduces metadata on the routes, which could increase the churn (Section 8.1.6 of RFC4098 [RFC4098]) of paths within the BMP session. For instance, if path marking is configured, and a non-installed path changes status to a backup route, the device should send an update about this path with the new markings, even if its BGP attributes remain unchanged. Enabling reason codes could additionally increase the churn. Churn could be more pronounced during the start of a BGP session, where the device is processing all available routes.

If churn is undesired, an implementation could make use of "state compression" to hide state until paths converge (Section 5 of [RFC7854]). It could also initially send BMP routes without the path marking TLV, even if it were configured, and then add them once the implementation considers the path to be stable enough. This document does not provide a definitive solution for churn since it depends on the capabilities of an implementation and the requirements of an operator.

4.3. Paths with no markings

Some BGP routes might not require any type of status or reasons. For example, a path in Adj-RIB-In where the BGP best path decision has not been applied yet, falls under this category, since there is really nothing to mark for that path. This document suggests applying an explicit marking of this route, by attaching a BMP path marking TLV with no bits set. This will help BMP monitor stations to differentiate this case from those in which markings are not configured, or not yet attached by the device.

4.4. Path markings applicability and consistency

The status and reason codes from Table 1 and Table 2 are included based on use cases from network operators and defined following the most relevant protocol references available. While implementations are strongly encouraged to align with these code definitions, this document does not enforce strict validity rules for code combinations to accommodate the diversity of BGP implementations.

The experience during testing of this TLV revealed scenarios where implementations might combine codes differently than originally anticipated. For example, one test implementation marked routes with both 'Invalid' and 'Best' status bits set, which is contradictory from the point of view of [RFC4271], but made sense for their specific implementation.

Operators should apply their own validation checks on the data from TLVs and discuss potential inconsistencies with their vendors, and raise bugs if applicable.

4.4.1. Significance of status and origin RIBs

This document refrains from imposing on any implementation the requirement to mark specific status from specific RIBs. Some implementations might be able to mark some status over one RIB while others do it on others. For instance, some might be able to mark Adj-RIB-In filtered routes when obtained from the Adj-RIB-In pre, while other could do it only from the Adj-RIB-In post. To remove ambiguities in implementations, it is recommended that the meaning of status (and reason codes) does not depend on the origin RIB of a route.

4.5. Multiple TLVs assigned to the same route.

We advocate for the use of TLV grouping wherever feasible (Section 5.2.1. of [I-D.ietf-grow-bmp-tlv]). The inclusion of all marking information within a single message is recommended. In situations where multiple TLVs are associated with a single route, all markings and reasons will be applicable to that route.

4.6. Enterprise-specific status

Implementations introducing their own status and reason codes are advised to adhere to [I-D.ietf-grow-bmp-tlv-ebit] and use the enterprise-bit (ebit) for vendor-specific status and reasons.

For scenarios where a path state combines a standard status with an enterprise-specific reason code (or vice versa), the following alternatives are presented:

- * Replication of the standard definitions within the enterprise-specific space, thus permitting direct marking within the same packet using the ebit.
- * Assigning two TLVs to the same path(s): one containing the standard part and another housing the vendor-specific part.

4.7. Multiple reason codes

The path marking TLV was not designed to optimally hold more than one reason code per path. However, if needed by a specific use case, the implementation can use two or more path markings TLVs for the same path listing the multiple reasons that apply to it.

5. Acknowledgments

We would like to thank Jeff Haas and Maxence Younsi for their valuable comments.

6. IANA Considerations

This document requests that IANA assign the following new TLV type to the BMP Route Monitoring TLVs.

Type = 5: indicates that this is the IANA-registered Path marking TLV. The value field is defined in Section 2.

RFC Editor and IANA registry note: The registry is created with Section 10 of [I-D.ietf-grow-bmp-tlv] and populated with initial values 1-4. This document adds value 5 to the registry. Please remove this sentence before publishing the document as RFC.

7. Security Considerations

Using the path status information may affect other applications which rely on this information for operational decisions. Operators should secure BMP sessions and control access to TLV data to mitigate these risks.

8. References

8.1. Normative References

[I-D.ietf-grow-bmp-tlv]

Lucente, P. and Y. Gu, "BMP v4: TLV Support for BGP Monitoring Protocol (BMP) Route Monitoring and Peer Down Messages", Work in Progress, Internet-Draft, draft-ietf-grow-bmp-tlv-17, 7 July 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-grow-bmp-tlv-17>>.

[I-D.ietf-grow-bmp-tlv-ebit]

Lucente, P. and Y. Gu, "Support for Enterprise-specific TLVs in the BGP Monitoring Protocol", Work in Progress, Internet-Draft, draft-ietf-grow-bmp-tlv-ebit-06, 17 January 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-grow-bmp-tlv-ebit-06>>.

[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>>.

[RFC2439] Villamizar, C., Chandra, R., and R. Govindan, "BGP Route Flap Damping", RFC 2439, DOI 10.17487/RFC2439, November 1998, <<https://www.rfc-editor.org/info/rfc2439>>.

[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>>.

[RFC4724] Sangli, S., Chen, E., Fernando, R., Scudder, J., and Y. Rekhter, "Graceful Restart Mechanism for BGP", RFC 4724, DOI 10.17487/RFC4724, January 2007, <<https://www.rfc-editor.org/info/rfc4724>>.

[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>>.

[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/info/rfc7854>>.

[RFC7911] Walton, D., Retana, A., Chen, E., and J. Scudder, "Advertisement of Multiple Paths in BGP", RFC 7911, DOI 10.17487/RFC7911, July 2016, <<https://www.rfc-editor.org/info/rfc7911>>.

- [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>>.

8.2. Informative References

- [I-D.ietf-idr-best-external]
Marques, P., Fernando, R., Chen, E., Mohapatra, P., and H. Gredler, "Advertisement of the best external route in BGP", Work in Progress, Internet-Draft, draft-ietf-idr-best-external-05, 3 January 2012, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-best-external-05>>.
- [I-D.ietf-rtgwg-bgp-pic]
Bashandy, A., Filsfils, C., and P. Mohapatra, "BGP Prefix Independent Convergence", Work in Progress, Internet-Draft, draft-ietf-rtgwg-bgp-pic-22, 20 April 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-rtgwg-bgp-pic-22>>.
- [I-D.lapukhov-bgp-ecmp-considerations]
Lapukhov, P. and J. Tantsura, "Equal-Cost Multipath Considerations for BGP", Work in Progress, Internet-Draft, draft-lapukhov-bgp-ecmp-considerations-13, 6 January 2025, <<https://datatracker.ietf.org/doc/html/draft-lapukhov-bgp-ecmp-considerations-13>>.
- [RFC4098] Berkowitz, H., Davies, E., Ed., Hares, S., Krishnaswamy, P., and M. Lepp, "Terminology for Benchmarking BGP Device Convergence in the Control Plane", RFC 4098, DOI 10.17487/RFC4098, June 2005, <<https://www.rfc-editor.org/info/rfc4098>>.

Authors' Addresses

Camilo Cardona
NTT
164-168, Carrer de Numancia
08029 Barcelona
Spain
Email: camilo@ntt.net

Paolo Lucente
NTT
Siriusdreef 70-72
2132 Hoofddorp
Netherlands
Email: paolo@ntt.net

Pierre Francois
INSA-Lyon
Lyon
France
Email: Pierre.Francois@insa-lyon.fr

Yunan Gu
Huawei
Huawei Bld., No.156 Beiqing Rd.
Beijing
100095
China
Email: gyunan@huawei.com

Thomas Graf
Swisscom
Binzring 17
CH-8045 Zurich
Switzerland
Email: thomas.graf@swisscom.com