

Global Routing Operations
Internet-Draft
Updates: 7854 (if approved)
Intended status: Standards Track
Expires: 7 March 2026

P. Lucente
NTT
Y. Gu
Huawei
3 September 2025

BMP v4: TLV Support for BGP Monitoring Protocol (BMP) Route Monitoring
and Peer Down Messages
draft-ietf-grow-bmp-tlv-18

Abstract

Most of the BGP Monitoring Protocol (BMP) message types make provision for data in Type, Length, Value (TLV) format. However, Route Monitoring messages (which provide a snapshot of the monitored Routing Information Base) and Peer Down messages (which indicate that a peering session was terminated) do not. Supporting (optional) data in TLV format across all BMP message types provides consistent and extensible structures that would be useful among the various use-cases where conveying additional data to a monitoring station is required. This document updates RFC 7854 [RFC7854] to support TLV data in all message types.

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 7 March 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. Terminology	3
3. Message version	3
4. TLV Encoding	4
5. BMP Message Format	5
5.1. Common Header	5
5.2. TLV Data in Route Monitoring	5
5.2.1. Group TLV	6
5.2.2. VRF/Table Name TLV	6
5.2.3. Stateless Parsing TLV	7
5.3. TLV Data in Peer Down	8
5.4. TLV Data in Other BMP Messages	8
6. Mandatory Timestamp	8
7. Error Handling	8
8. Security Considerations	9
9. Operational Considerations	9
10. IANA Considerations	9
11. References	9
11.1. Normative References	10
11.2. Informative References	11
Appendix A. Wire-format Example	11
Acknowledgements	13
Authors' Addresses	13

1. Introduction

The BGP Monitoring Protocol (BMP) version 3 is defined in RFC 7854 [RFC7854].

The Route Monitoring message consists of:

- * Common Header
- * Per-Peer Header
- * BGP Update PDU

The Peer Down Notification message consists of:

- * Common Header
- * Per-Peer Header
- * Reason
- * Data (only if Reason code is 1, 2 or 3)
- * TLV (only if Reason code is 6)

This means that both Route Monitoring and Peer Down messages have a non-extensible format (except for the specific case of Peer Down Reason Code 6 as specified in Section 5.3 of [RFC9069]). In the Route Monitoring case, this prevents the transmission of parsing characteristics of transported NLRIs (e.g. ADD-PATH, Multi Labels, etc.), RIB status of a path (e.g. primary, backup, unused, etc.) or of vendor-specific data. In the Peer Down case, this prevents matching with TLVs previously sent with the Peer Up message. This document:

- * Bumps the BMP version for all message types defined in RFC 7854 [RFC7854] for backward compatibility
- * Changes the structure of Route Monitoring message type so that the BGP Update PDU is enclosed in a TLV. The BGP Message PDU TLV is mandatory to be included
- * Allows all defined BMP message types to make provision for optional TLV data.

2. Terminology

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.

The document uses the terms defined in RFC 7854 [RFC7854].

3. Message version

For an exporter to flag a receiver that it does comply with this specification, the Version field of the BMP Common header, documented in Section 4.4 of [RFC7854], MUST be set to 4. This applies to every BMP message type.

If a BMP station does not support the version indicated in the message, it SHOULD close the session and take the procedures described in Error Handling (Section 7)

4. TLV Encoding

The TLV data type (Information TLV) is defined in Section 4.4 of [RFC7854] for the Initiation and Peer Up message types. A TLV object consists of:

- * 2 octets of TLV Type,
- * 2 octets of TLV Length, and
- * 0 or more octets of TLV Value.

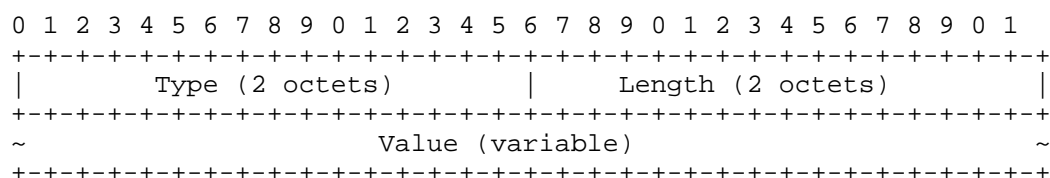


Figure 1

TLVs SHOULD be sorted by the sender by their type. Multiple TLVs of the same type can be repeated as part of the same message; it is left to the specific use-cases whether all, any, the first or the last TLV should be considered as well as whether ordering matters and repeating is allowed.

Route Monitoring messages may require per-NLRI TLVs. That is, there may be a need to map TLVs to NLRIs contained in the BGP Update message, for example, to express additional characteristics of a specific NLRI. For this purpose, TLVs enclosed in a Route Monitoring message MUST be indexed, with the index starting at one (1) to refer to the first NLRI. Index zero (0) specifies that a TLV does apply to all NLRIs contained in the BGP Update message. The Index field is 2-byte long of which the top-most bit, G-bit, is used to flag a Group Index (more in Section 5.2.1). TLVs of the same type and with the same index can be repeated as part of the same message, unless specified otherwise by the definition of the specific TLV. Indexed TLVs are encoded as in the following figure:

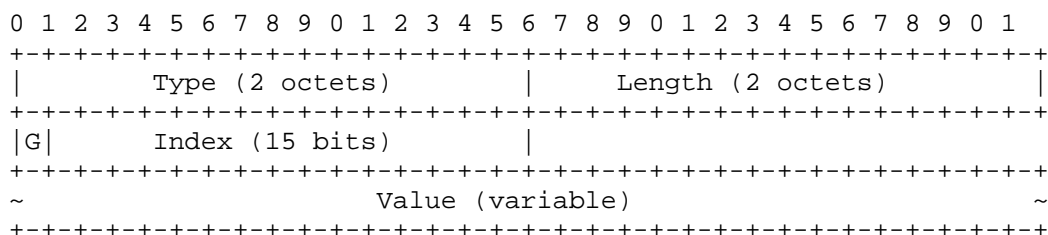


Figure 2

Indexed TLVs SHOULD be sorted by the sender by their type and index value. The reported length in indexed TLVs refers to the total encoded TLV value (ie. with the length of the index field excluded).

A monitoring station can properly match indexed TLVs to the corresponding NLRI only if - or as long as - NLRIs are decoded successfully. In case of any parsing or error condition that prevents full decoding of the BGP PDU, the station MUST stop matching indexed TLVs to NLRIs.

Of the BMP message types defined so far, indexed TLVs apply only to Route Monitoring messages. For example, they do not apply to Route Mirroring messages because a sender may not be aware of the payload of the transported BGP Update message.

5. BMP Message Format

5.1. Common Header

While the structure of the Common header remains unaltered, the following two definitions are changed compared to Section 4.1 of [RFC7854]:

- * Version: Indicates the BMP version. This is set to '4' for all message types defined in RFC 7854 [RFC7854].
- * Message Length: Total length of the message in bytes (including headers, encapsulated BGP Message PDU TLV and optional TLV data).

5.2. TLV Data in Route Monitoring

For consistency with the Route Mirroring type defined in Section 4.7 of [RFC7854], this document extends the encoding of the Route Monitoring message type where the Per-peer header is followed by mandatory and optional TLVs.

The BGP Update PDU (Section 4.3 of [RFC4271]) is encoded itself as part of a BGP Message TLV with code point 4 and index set to zero. A Route Monitoring message MUST contain one BGP Message TLV which may be preceded or followed by other optional TLV data.

Corollary, the BGP Update PDU is not encoded as part of the message as it was the case for BMPv3 (RFC 7854 [RFC7854]) but it is rather enclosed in a TLV.

5.2.1. Group TLV

In a Route Monitoring message where a BGP Update PDU carries N NLRIs, indexed TLVs do allow to handle the cases of 1:1 and N:1 relationship among TLVs and NLRIs (ie. one TLV applies to one NLRI, N TLVs apply to one same NLRI). The cases of 1:N and M:N relationships (i.e., one TLV applies to N NLRIs and M TLVs apply to N NLRIs) can benefit by a form of grouping. For that purpose, a Group TLV is defined with the aim to limit both verbosity and repetitions.

The Group TLV value MUST contain:

- * A 2-byte Group Index where the top-most bit (G-bit) MUST be set to one (1). The full 2-byte value, that is including the G-bit, MUST be unique to the message
- * Two or more 2-byte NLRI indexes whose values MUST be less or equal to the amount of NLRIs packed in the BGP Update PDU.

An NLRI index can be listed as part of multiple Group TLVs within the same message. NLRI indexes within a Group TLV SHOULD be sorted by the sender. A Group Index MUST NOT reference an NLRI index 0. A Group TLV MUST NOT include its own or another Group Index. Multiple non-Group TLVs MAY point to the same Group Index, i.e., a group can be reused within the same Route Monitoring message.

The Group TLV type is 2. It is RECOMMENDED that this TLV is encoded first in order to ease parsing of the Route Monitoring message at the BMP station side.

5.2.2. VRF/Table Name TLV

The Information field contains a UTF-8 string whose value MUST be equal to the value of the VRF or table name (i.e., RD instance name) being conveyed. The string size MUST be within the range of 1 to 255 bytes. This is in line with Section 5.2.1 of [RFC9069].

The VRF/Table Name TLV type is 3

5.2.3. Stateless Parsing TLV

Stateless parsing helps scaling the amount of Route Monitoring messages that can be processed at collection time, avoiding to have to correlate them to BGP capabilities received as part of the Peer Up message, for example.

Some BGP capabilities are not per AFI/SAFI, like 4-byte ASN RFC 6793 [RFC6793], and hence these can potentially be part of the BMP Peer flags [IANA-BPPF] of a Route Monitoring message. Those that are, instead, per AFI/SAFI require finer granularity and hence the need to use an indexed TLV. To maintain the consistency between BGP Capabilities, they are all carried in Stateless Parsing TLVs, regardless of their type.

The encoding of BGP Capabilities is already generically defined in Section 4 of [RFC5492]. It includes every information needed for its understanding, including AFI/SAFI codes. This encoding can be reused in the Stateless Parsing TLV, leveraging already existing BGP encoders and decoders for implementers.

The Stateless Parsing TLV type is 1 and its Value is the BGP Capability encoded as exactly as it would be in the BGP OPEN of the session. It is thus made of all three Capability Code, Capability Length, and Capability Value fields defined in [RFC5492].

For example, an ADD-PATH capability, as defined by RFC 7911 [RFC7911], for IP/Unicast with value Send/Receive would be encoded in the Capability Value as:

- * Capability Code, 1 byte, value=69
- * Capability Length, 1 byte, value=4
- * AFI, 2 bytes, value=1
- * SAFI, 1 byte, value=1
- * Value, 1 byte, value=3

The index of the Stateless Parsing TLV MUST be set to zero.

If no Stateless Parsing TLV is present in a Route Monitoring message, the receiver MUST fall back to use capabilities present in the BGP Open PDU contained in the relevant BMP Peer Up message in order to properly parse BGP Update PDUs. Each BGP capability is to be encoded in a separate Stateless Parsing TLV.

It is RECOMMENDED that the Stateless Parsing TLV is encoded preceding the BGP Message TLV in order to ease parsing of the Route Monitoring message at the BMP station side.

5.3. TLV Data in Peer Down

The Peer Down Notification message type (Section 4.9 of [RFC7854]) is extended following a consistent approach with the Peer Up type (Section 4.10 of [RFC7854]). That is, the message is extended so that optional TLVs are placed at the end of the message.

This means for Reason codes 1 or 3, a BGP Notification PDU follows; the PDU MAY be further followed by TLV data. For Reason code 2, a 2-byte field follows to provide additional Finite State Machine (FSM) info; this field MAY be followed by TLV data. For all other Reason codes, TLV data MAY follow the Reason field.

5.4. TLV Data in Other BMP Messages

All other message types defined in RFC7854 [RFC7854] do already provision for TLV data. It is RECOMMENDED that all future defined BMP message types will also provide for optional TLV data following a consistency model for encoding with existing message types.

6. Mandatory Timestamp

As defined by Section 4.2 of [RFC7854], it is not required to specify a timestamp as part of any BMP message that includes a Per-peer header. This is not ideal, as expressed in Section 4.10 of [I-D.boucadair-nmop-rfc3535-20years-later]. An event timestamp MUST always be defined. The observation timestamp SHOULD always be preferred as reference for its inherent maximum accuracy in reporting a given event; would that not be available in an implementation, the next accurate timestamp SHOULD be picked up to, as a last resort, the time at which the information was exported.

7. Error Handling

RFC8654 [RFC8654] permits BGP Update and other messages to grow to a length of 65535 octets. This may cause a BMP PDU that attempts to encapsulate such long messages to overflow.

A BMP exporter and a BMP station may not support the same version of the protocol; being BMP uni-directional, with data flowing only from the exporter to the station, the station SHOULD close the BMP session and log the condition as a warning; the exporter SHOULD retry to connect with a non-aggressive timer.

A BMP station may not support some of the TLVs encoded by the exporter; the station MUST ignore unsupported TLV types; additionally, in case of indexed TLVs, if the index is invalid (i.e. out of bounds), the TLV MUST be ignored. The station SHOULD log the condition as a warning.

8. Security Considerations

It is not believed that this document adds any additional security considerations compared to RFC7854 [RFC7854].

9. Operational Considerations

In Route Monitoring messages, the number of TLVs can be bound to the amount of NLRIs carried in the BGP Update message. This may degrade the packing of information in such messages and have specific impacts on the memory and CPU used in a BMP implementation. As a result of that it should always be possible to disable such features to mitigate their impact.

10. IANA Considerations

This document requests IANA to rename of the "BMP Peer Up Message TLVs" registry defined by BMP Peer Up Message Namespace [RFC9736] into "BMP Peer Up and Peer Down TLVs" and the definition of one new registry "BMP Route Monitoring TLVs". The new "BMP Route Monitoring TLVs" registry is seeded with the following new TLV types (Section 5.2):

- * Type = 1: Support for Stateless Parsing TLV. The value field is defined in Section 5.2.3.
- * Type = 2: Support for grouping of TLVs. The value field is defined in Section 5.2.1.
- * Type = 3: Support for VRF/Table Name TLV. The value field is defined in Section 5.2.2
- * Type = 4: Support for BGP Message TLV. The value field is defined in Section 5.2

Values 0 through 16383 MUST be assigned using the Standards Action policy as defined in Section 4.9 of [RFC8126]; values 16384 through 32767 MUST be assigned using the First Come First Served policy as defined in Section 4.4 of [RFC8126]. The upper bound of the registry is 65535. Value 65535 is Reserved.

11. References

11.1. Normative References

- [I-D.boucadair-nmop-rfc3535-20years-later]
Boucadair, M., Contreras, L. M., de Dios, O. G., Graf, T., Rahman, R., and L. Tailhardat, "RFC 3535, 20 Years Later: An Update of Operators Requirements on Network Management Protocols and Modelling", Work in Progress, Internet-Draft, draft-boucadair-nmop-rfc3535-20years-later-08, 12 May 2025, <<https://datatracker.ietf.org/doc/html/draft-boucadair-nmop-rfc3535-20years-later-08>>.
- [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>>.
- [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>>.
- [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>>.
- [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>>.
- [RFC8654] Bush, R., Patel, K., and D. Ward, "Extended Message Support for BGP", RFC 8654, DOI 10.17487/RFC8654, October 2019, <<https://www.rfc-editor.org/info/rfc8654>>.
- [RFC9069] Evens, T., Bayraktar, S., Bhardwaj, M., and P. Lucente, "Support for Local RIB in the BGP Monitoring Protocol (BMP)", RFC 9069, DOI 10.17487/RFC9069, February 2022, <<https://www.rfc-editor.org/info/rfc9069>>.
- [RFC9736] Scudder, J. and P. Lucente, "The BGP Monitoring Protocol (BMP) Peer Up Message Namespace", RFC 9736, DOI 10.17487/RFC9736, March 2025, <<https://www.rfc-editor.org/info/rfc9736>>.

11.2. Informative References

- [IANA-BCC] IANA, "Capabilities Codes", 2025,
<<https://www.iana.org/assignments/capability-codes/>>.
- [IANA-BPPF] IANA, "BMP Peer Flags", 2024,
<<https://www.iana.org/assignments/bmp-parameters/bmp-parameters.xhtml#peer-flags>>.
- [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>>.
- [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/info/rfc6793>>.
- [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>>.

Appendix A. Wire-format Example

The diagram in Figure 3 shows an example of a Route Monitoring message carrying a BGP UPDATE containing 10 NLRI's. The TLVs are comprised of:

1. a Group TLV with index 0x000b, pointing to NLRI 1, 2, 3 and 10
2. a Group TLV with index 0x000c, pointing to NLRI 4, 5 and 6
3. a Stateless Parsing TLV with index 0x0000, stating that the Add-Path Capability is set to Both(Snd/Rcv) for the IPv4 Unicast address family in the BGP UPDATE.
4. a TLV pertaining to NLRI 7
5. a TLV pertaining to the NLRI's listed in the Group TLV defined in 1
6. a TLV pertaining to the NLRI's listed in the Group TLV defined in 2

[Page 12]

```
+-----+
| 0 |           index=0x0007           |
+-----+
|                                     value={8 bytes}                                     ~
~                                     |
+-----+
```

Figure 3

Acknowledgements

The authors would like to thank Jeff Haas, Camilo Cardona, Thomas Graf, Pierre Francois, Ben Maddison, Tim Evens, Luuk Hendriks, Maxence Younsi, Ahmed Elhassany, Colin Petrie, Dhananjay Pakti and Shunwan Zhuang for their valuable input. The authors would also like to thank Greg Skinner, Zongpeng Du and Mohamed Boucadair for their review.

Authors' Addresses

Paolo Lucente
NTT
Veemweg 23
3771 Barneveld
Netherlands
Email: paolo@ntt.net

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