

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

P. Lucente
C. Cardona
NTT
3 September 2025

Logging of routing events in BGP Monitoring Protocol (BMP)
draft-ietf-grow-bmp-rel-04

Abstract

The BGP Monitoring Protocol (BMP) does provision for BGP session event logging (Peer Up, Peer Down), state synchronization (Route Monitoring), debugging (Route Mirroring) and Statistics messages, among the others. This document defines a new Route Event Logging (REL) message type for BMP with the aim of covering use-cases with affinity to alerting, reporting and on-change analysis.

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. Route Event Logging (REL) message	3
3.1. Per-Peer Header	4
3.2. BGP Update PDU	4
3.3. Informational TLVs	5
3.3.1. Event Reason TLV	5
3.3.2. Log Action TLV	5
3.3.3. Policy Discard TLV	6
3.3.4. Validation Fail TLV	6
3.3.5. Malformed Packet TLV	7
3.4. Group TLV	7
3.5. Stateless Parsing TLV	7
4. Examples and use-cases	7
5. Operational Considerations	8
6. Security Considerations	8
7. IANA Considerations	9
7.1. BMP Route Event Logging TLVs Registry	9
7.2. Event Reason TLV Registry	10
7.3. Log Action TLV Registry	10
7.4. Policy Discard TLV Registry	10
7.5. Validation Fail TLV	11
7.6. Malformed Packet TLV Registry	11
8. References	11
Acknowledgements	13
Authors' Addresses	13

1. Introduction

As NLRIs are advertised and distributed, policies are applied and, as a result, actions are performed on them. Currently, in order to infer the outcome of an evaluation process, a comparative analysis needs to be performed between Route Monitoring data for two distinct observation points of interest, for example Adj-Rib-In pre-policy and post-policy. It would be instead more useful if a monitored router could export event-driven data with the relevant information.

The envisioned use-cases are the most diverse and range from logging route filtering to reporting the outcome of validation processes taking place on the monitored router, to isolating certain subsets of data to be validated offline, to report malformed BGP packets, to broader closed-loop operations.

This document defines a new Route Event Logging (REL) message type that is suitable to carry event-driven data and is extensible in nature. While the message format is similar to the Route Mirroring message type defined in RFC 7854 [RFC7854] and to the Route Monitoring message type as defined in [I-D.ietf-grow-bmp-tlv], the semantics are different.

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.

3. Route Event Logging (REL) message

In basic terms a REL message does carry Events. Each Event is logically composed by one Event Subject and one or more Event Attributes.

The structure of the Route Event Logging message is similar to the Route Monitoring message defined in [I-D.ietf-grow-bmp-tlv], as such all TLVs enclosed in the message are indexed.

The message is composed of the Per-Peer Header, one mandatory BGP Message TLV, one mandatory Informational TLV indicating the reason of the event and any further optional Informational TLVs to better characterize the nature of the event. An example of the indexed TLV structure is available in Appendix A of [I-D.ietf-grow-bmp-tlv].

One or more Event Subjects are packed as part of the BGP Update PDU. The BGP Update PDU Section 4.3 of [RFC4271] is encoded itself as part of the BGP Message TLV with code point 4 and index set to zero. Each Event Subject is represented by an NLRI carried in the PDU.

The BGP Message TLV may be preceded and/or followed by indexed Informational TLVs that carry Event Attributes, where attributes are bound to subjects referring to their positional index within the PDU or via a Group TLV as described in Section 5.2.1 of [I-D.ietf-grow-bmp-tlv]

Speaking comparatively to other existing message types, REL does not require an initial flooding of information as per the state synchronization nature of Route Monitoring and does not aim to provide a non-state-compressed full-fidelity view of all messages received as per the debugging nature of Route Mirroring.

In the context of BMP REL message, and hence in the reminder of this document, the term Event Subject and NLRI will be used interchangeably. Also the term Event Attribute and Informational TLV will be used interchangeably.

The following sections will describe each component of the REL message in more detail.

3.1. Per-Peer Header

The message does start with a BMP per-peer header as defined in RFC 7854 [RFC7854], subsequently extended by RFC 8671 [RFC8671] and RFC 9069 [RFC9069] allowing, among the other things, to timestamp an Event and set its observation point among those defined in BMP.

Because the main purpose of the REL message is to log events at the time of applying an action, the Peer Flags field - even if applied to Adj-Rib-In or Adj-Rib-Out does not have the concept of pre- and post-policy. The flags are hence defined as follows:

```

      0 1 2 3 4 5 6 7
      +---+---+---+---+
      |V|A| Reserved  |
      +---+---+---+---+

```

The V flag and A flag do carry the same meaning as originally defined by RFC 7854 [RFC7854]. The remaining bits are reserved for future use. They MUST be transmitted as 0 and their values MUST be ignored on receipt.

3.2. BGP Update PDU

The PDU enclosed as part of a BGP Message TLV can be either a verbatim copy or artificial, either packed from scratch or repacked starting from an existing BGP Update PDU to only contain the relevant NLRIs affected by an Event (one or multiple). The Event is going to be further described by means of Event Attributes by indexed Informational TLVs.

The choice of describing one or multiple Event Subjects via a BGP Update PDU is because, on one hand, this does allow to not have to invent new encodings for NLRIs, while on the other, to support all

types and encodings already supported by BGP. The advantage being that only minimal new code, on both the exporting and the receiving sides, will have to be produced.

3.3. Informational TLVs

Informational TLVs in BMP are formalized by the intersection of RFC 7854 [RFC7854], [I-D.ietf-grow-bmp-tlv] and [I-D.ietf-grow-bmp-tlv-ebit]. TLVs in a REL message are indexed.

Contrary to other BMP messages where all Informational TLVs are entirely optional, in order for a REL message to be meaningful, it MUST contain at least one Event Reason TLV and MAY contain other optional attribute TLVs to further characterize the Event.

A new registry called "Route Event Logging TLVs" is defined and is seeded with the TLVs detailed in the following sections.

3.3.1. Event Reason TLV

5 = Event Reason TLV (4 octets). Indicates the IANA-registered reason code describing the type of the event. The following reason codes are defined as part of the "Event Reason TLV" registry:

Value	Event Reason
0x0001	Log Action
0x0002	Policy Discard
0x0004	Validation Fail
0x0008	Malformed Packet

Table 1: IANA-Registered
Event reasons

3.3.2. Log Action TLV

6 = Log Action TLV. The length is variable. The first byte defines the nature of the logging, depending on the code point additional data may follow. the following code points are defined:

- * 1 = Config. Prefix is being logged due to a configuration statement. Data contains a UTF-8 string whose value can be organized freely by an implementation and is meant to give additional information about why the log was made.
- * 2 = Route unstable. Optional data contains a 4 bytes value representing the observed timeframe in seconds, followed by a 4 bytes value indicating the amount of times the event occurred within the timeframe.
- * 3 = Crossed Warning Bound. Prefix is over the warning threshold of the maximum number of prefixes that can be received from a BGP neighbor. Data contains a 4 bytes value representing the threshold number.
- * 4 = Crossed Upper Bound. Prefix is over the upper threshold of the maximum number of prefixes that can be received from a BGP neighbor. Data contains a 4 bytes value representing the threshold number.

3.3.3. Policy Discard TLV

7 = Policy Discard TLV. The length is variable. The first byte of the value field indicates how the rest is organized:

- * 1 = String. The value contains a UTF-8 string whose value can be organized freely by an implementation. For example, it may contain the routing policy name that caused the discard; or it may list a sequence of policies and policy nodes traversed; or, more simply, it could be a meaningful return code.
- * 2 = Structured. On the escort of Section 4 of [RFC9067] and YANG Model for Border Gateway Protocol (BGP-4) [I-D.ietf-idr-bgp-model] the value is organized as two consecutive null-terminated strings, the first indicating the policy name, the second the statement name within the policy.

3.3.4. Validation Fail TLV

8 = Validation Fail TLV. The length is set to 1 byte and the value represents a code giving more information about the specific validation fail. Following are the defined code points:

- * Code = 1: RPKI Invalid. The prefix is being marked as RPKI 'invalid' and either has no coverage or it is unknown whether it has coverage by a valid prefix.

- * Code = 2: RPKI Invalid with covering Valid prefix. The NLRI is being marked as RPKI 'invalid' but is covered by a Valid prefix.

3.3.5. Malformed Packet TLV

9 = Malformed Packet TLV. The length is set to 1 byte and the value represents a code giving more information about the specific malforming. Following are the defined code points:

- * Code = 1: Errored PDU. The BGP message was found to have some error that made it unusable, causing it to be treated-as-withdraw RFC7606 [RFC7606].

3.4. Group TLV

The Group TLV is to form N:M relationships among NLRIs in the BGP Update PDU and TLVs of the same Route Event Logging message. This TLV has code point 2 and follows the definition of Group TLV in [I-D.ietf-grow-bmp-tlv].

3.5. Stateless Parsing TLV

The Stateless Parsing TLV is to allow parsing of the BGP Update PDU independently from a Peer Up message previously received for the same BGP session. This TLV can be especially relevant to Route Event Logging where the BGP Update PDU is artificial. The TLV has code point 1, it follows the definition of Stateless Parsing TLV in [I-D.ietf-grow-bmp-tlv] and uses code point definitions in the Stateless Parsing Registry.

4. Examples and use-cases

REL can be used to send real-time notifications for specific routing events enabling rapid alerting of issues like policy discards, validation failures, or malformed packets to operators. For example, an operator is notified immediately when a route is discarded due to policy, assisting quick diagnosis and policy refinement.

By logging every routing event and the corresponding reason code, REL enables thorough audits of route changes and network behavior over time. For example, when a route fails validation, a log entry with the "Validation Fail" reason is stored for compliance checks and future forensics.

REL events, especially with machine-readable reason codes, can feed analytics engines and automated workflows to correlate events across the network and trigger remediation. For example, analytics dashboards continuously monitor for spikes in "Malformed Packet" events to detect possible protocol attacks or systemic misconfigurations.

5. Operational Considerations

Route Event Logging messages are event-driven in nature so the general recommendation is to use them to report on specific conditions of interest in order, for example, to facilitate data mining or avoid differential analysis. When the objective is to annotate every received or announced NLRI then the recommendation is to use Route Monitoring messages with BMP Path Marking [I-D.ietf-grow-bmp-path-marking-tlv]. As an example consider RPKI validation status: when the objective is to report on any validations tatus (ie. valid, invalid and unknown), BMP Path Marking should be used; when the objective is instead to report only invalids then Route Event Logging with Validation Fail Event Reason should be used.

There exist a definite overlap between Route Event Logging when used to report Malformed Packet and the use-cases for Route Mirroring where Errored PDUs may be sampled for reporting. From implementors perspective, if one wants to implement broader event-driven notifications and does not want to offer exact mirroring of monitored BGP sessions without state compression it may be adviceable to prefer implementing Route Event Logging message type over Route Mirroring. From a collector perspective, similarly, one may want to activate distinct BMP feeds for event logging and route collection and, also in this case, reporting malformed packets via Route Event Logging message type may be preferable over Route Mirroring.

Crossed warning bound and crossed upper bound events refer to the received route thresholds that can be configured according to Section 6.7 of [RFC4271]. Also the stats counters part of these events is being addressed by the Definition For New BMP Statistics Type [I-D.ietf-grow-bmp-bgp-rib-stats] document.

6. Security Considerations

It is not believed that this document adds any additional security considerations.

7. IANA Considerations

This document requests that IANA creates all the new registries in the following sections under the "BGP Monitoring Protocol (BMP) Parameters" group. The registries will record type code points for TLVs specific to the Route Event Logging (REL) message type, as defined in this document.

7.1. BMP Route Event Logging TLVs Registry

TLV Type consists of a code point (unsigned 16-bit value) and initial allocations are as follows:

- * Type = 0: Reserved for future use.
- * Type = 1: Support for Stateless Parsing TLV. The value is defined in Section 5.2.3 of [I-D.ietf-grow-bmp-tlv].
- * Type = 2: Support for grouping of TLVs. The value is defined in Section 5.2.1 of [I-D.ietf-grow-bmp-tlv].
- * Type = 3: Reserved for future use.
- * Type = 4: Support for BGP Message TLV. The value is defined in Section 3
- * Type = 5: Indicates IANA-registered reason code for event. The value is defined in Section 3.3.1.
- * Type = 6: Describes specific logging actions. The value is defined in Section 3.3.2.
- * Type = 7: Indicates NLRI discarded due to routing policy. The value is defined in Section 3.3.3.
- * Type = 8: Marks validation-related failure (e.g., RPKI invalidation). The value is defined in Section 3.3.4.
- * Type = 9: Reports a malformed BGP message indicating the reason. The value is defined in Section 3.3.5.

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.

7.2. Event Reason TLV Registry

TLV Type consists of a code point (unsigned 8-bit value) and is defined in Section 3.3.1. Initial allocations are as follows:

Type = 0x0001: Log Action reason.

Type = 0x0002: Policy Discard reason.

Type = 0x0004: Validation Fail reason.

Type = 0x0008: Malformed Packet reason.

Values 0 through 63 MUST be assigned using the Standards Action policy as defined in Section 4.9 of [RFC8126]; values 64 through 127 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 255. Value 255 is Reserved.

7.3. Log Action TLV Registry

TLV Type consists of a code point (unsigned 8-bit value) and is defined in Section 3.3.2. Initial allocations are as follows:

Type = 1: Config (prefix logged due to configuration).

Type = 2: Route unstable.

Type = 3: Crossed Warning Bound

Type = 4: Crossed Upper Bound

Values 0 through 63 MUST be assigned using the Standards Action policy as defined in Section 4.9 of [RFC8126]; values 64 through 127 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 255. Value 255 is Reserved.

7.4. Policy Discard TLV Registry

TLV Type consists of a code point (unsigned 8-bit value) and is defined in Section 3.3.3. Initial allocations are as follows:

Type = 1: String (UTF-8 policy name/reason)

Type = 2: Structured (policy and statement, null-terminated)

Values 0 through 63 MUST be assigned using the Standards Action policy as defined in Section 4.9 of [RFC8126]; values 64 through 127 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 255. Value 255 is Reserved.

7.5. Validation Fail TLV

TLV Type consists of a code point (unsigned 8-bit value) and is defined in Section 3.3.4. Initial allocations are as follows:

Type = 1: RPKI Invalid

Type = 2: RPKI Invalid with Covering Valid Prefix

Values 0 through 63 MUST be assigned using the Standards Action policy as defined in Section 4.9 of [RFC8126]; values 64 through 127 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 255. Value 255 is Reserved.

7.6. Malformed Packet TLV Registry

TLV Type consists of a code point (unsigned 8-bit value) and is defined in Section 3.3.5. Initial allocations are as follows:

Type = 1: Errored PDU (treated-as-withdraw per RFC 7606)

Values 0 through 63 MUST be assigned using the Standards Action policy as defined in Section 4.9 of [RFC8126]; values 64 through 127 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 255. Value 255 is Reserved.

8. References

[I-D.ietf-grow-bmp-bgp-rib-stats]

Srivastava, M., Liu, Y., Lin, C., and J. Li, "Definition For New BGP Monitoring Protocol (BMP) Statistics Types", Work in Progress, Internet-Draft, draft-ietf-grow-bmp-bgp-rib-stats-08, 24 April 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-grow-bmp-bgp-rib-stats-08>>.

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

Cardona, C., Lucente, P., Francois, P., Gu, Y., and T. Graf, "BMP Extension for Path Status TLV", Work in Progress, Internet-Draft, draft-ietf-grow-bmp-path-

marking-tlv-04, 25 August 2025,
<<https://datatracker.ietf.org/doc/html/draft-ietf-grow-bmp-path-marking-tlv-04>>.

[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>>.

[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>>.

[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>>.

[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>>.

[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>>.
- [RFC8671] Evens, T., Bayraktar, S., Lucente, P., Mi, P., and S. Zhuang, "Support for Adj-RIB-Out in the BGP Monitoring Protocol (BMP)", RFC 8671, DOI 10.17487/RFC8671, November 2019, <<https://www.rfc-editor.org/info/rfc8671>>.
- [RFC9067] Qu, Y., Tantsura, J., Lindem, A., and X. Liu, "A YANG Data Model for Routing Policy", RFC 9067, DOI 10.17487/RFC9067, October 2021, <<https://www.rfc-editor.org/info/rfc9067>>.
- [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>>.

Acknowledgements

The authors would like to thank Jeff Haas, Luuk Hendriks, Ruediger Volk, Ahmed Elhassany, Thomas Graf and Ben Maddison for their valuable input. The authors would also like to thank Mike Booth for his review.

Authors' Addresses

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

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