GROW
Internet-Draft
Intended status: Standards Track
Expires: 23 April 2026

D. Patki
P. Narasimha
Cisco Systems, Inc.

20 October 2025

Common BMP Route-Monitoring Messages for Routes Unchanged by Policy
draft-patki-grow-bmp-common-updates-02

Abstract

A route unmodified by the inbound policy on a monitored router is included both in Pre-Policy Adj-RIB-In as well as Post-Policy Adj-RIB-In Route-Monitoring messages when both the Pre-Policy and Post-Policy Route-Monitoring modes are enabled. Similarly, a route unmodified by the outbound policy is included in Pre-Policy Adj-RIB-Out as well as Post-Policy Adj-RIB-Out Route-Monitoring messages. This document defines a method to avoid duplicate inclusion of routes unmodified by policy either in Adj-RIB-In or Adj-RIB-Out.

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

[RFC7854] defined Pre-Policy and Post-Policy Adj-RIB-In Route-Monitoring messages, whereas [RFC8671] defined Pre-Policy and Post-Policy Adj-RIB-Out Route-Monitoring messages. If both Pre-Policy and Post-Policy Route-Monitoring modes are enabled on a device for a RIB (Adj-RIB-In or Adj-RIB-Out), the routes are included in both Pre-Policy and Post-Policy Route-Monitoring messages, even if the routes remains unmodified as a result of the application of policy.

The optimization proposed in this document will help improve the BMP convergence as described in the section below.

1.1. BMP Convergence

The monitored routers may have policies that modify none, some or several attributes of prefixes learnt from a few to many BGP peers. For example, a Route Reflector Inbound policy may modify very few of the received attributes. Whereas, a Provider Edge router Inbound policy may modify more attributes in the prefixes learnt across several peers. Consider a monitored router that learns 1,000,000 prefixes from various peers and, in different cases, 100%, 50%, 10% and none of the prefixes are modified by the policies. For the sake of simplicity, consider that 10 prefixes are packed in a single Route-Monitoring message and the average size of Route-Monitoring messages is 200 bytes. The following illustration shows the number of Route-Monitoring messages sent in each of these cases.

Prefixes modified by inbound policy	Pre-Policy Messages	Post-Policy Messages	Common Messages	Total Messages Transmitted	Total Bytes Transmitted
100% = 1,000,000	100,000	100,000	0	200,000	40 MB
50% = 500,000	50,000	50,000	50,000	150,000	30 MB
10% = 100,000	10,000	10,000	90,000	110,000	22 MB
None	0	0	100,000	100,000	20 MB

Table 1: Route-Monitoring messages generated for inbound policy variations

While there can be multi-dimensional variations that determine the number of messages sent, the above simplified cases broadly illustrates that the number or Route-Monitoring messages can be reduced by a factor of two in the best case. This can therefore reduce the transmission processing, number of transmit buffers required for sending the BMP updates and internal queuing delays in the monitored router and load on the network connecting to the monitoring station; thereby improving the overall BMP convergence. This can also reduce the number of messages processed by the monitoring station.

1.2. Solution

To avoid sending duplicate unmodified routes in the Post-Policy Route-Monitoring messages, we introduce in this document a method based on Common Update TLV as defined in [I-D.ietf-grow-bmp-tlv].

1.3. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2. Common Update TLV

Here we define a new TLV named Common Update TLV using the TLV construct defined in Section 4.2 of [I-D.ietf-grow-bmp-tlv]. In addition to allowing sharing a common BGP Update PDU between Pre-Policy and Post-Policy modes of Adj-RIB-In and the same for Adj-RIB-Out, this method is extensible in allowing sharing across Adj-RIB-In and Adj-RIB-Out views, though we see it being used rarely.

The TLV has Index zero (0) which specifies that a TLV applies to all NLRIs contained in the BGP Update PDU. The value of the TLV defines flags that are described below.

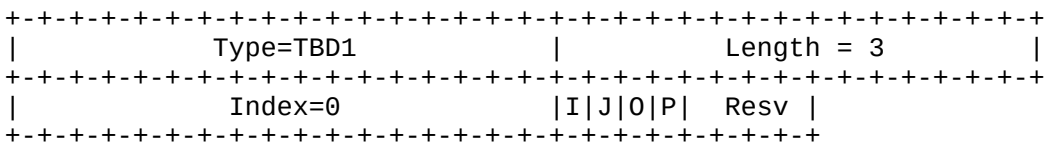

Figure 1: Common Update TLV

* When I flag is set, it indicates that the BGP Update PDU reflects the Pre-Policy Adj-RIB-In view of all contained NLRIs

* When J flag is set, it indicates that the BGP Update PDU reflects the Post-Policy Adj-RIB-In view of all contained NLRIs

* When O flag is set, it indicates that the BGP Update PDU reflects the Pre-Policy Adj-RIB-Out view of all contained NLRIs

* When P flag is set, it indicates that the BGP Update PDU reflects the Post-Policy Adj-RIB-Out view of all contained NLRIs

* The remaining bits are reserved for future use. They MUST be transmitted as 0 and their values MUST be ignored on receipt.

2.1. Examples of the Common Update TLV

* When I=1, J=1, O=0, P=0 it indicates that the BGP Update PDU is the same for Pre-Policy and Post-Policy Adj-RIB-In views.

* When I=0, J=0, O=1, P=1 it indicates that the BGP Update PDU is the same for Pre-Policy and Post-Policy Adj-RIB-Out views.

The following examples demonstrate sharing across Adj-RIB-In and Adj-RIB-Out views as well, but we anticipate this not to be used

* When I=0, J=1, O=1, P=0 it indicates that the BGP Update PDU is the same for Post-Policy Adj-RIB-In and Pre-Policy Adj-RIB-Out views.

* When I=0, J=1, O=1, P=1 it indicates that the BGP Update PDU is the same for Post-Policy Adj-RIB-In, and Pre-Policy and Post-Policy Adj-RIB-Out views.

3. BMP Messages

The Common Update TLV is used in the context of BGP Update PDU, and has no significance for Peer-Up, Peer-Down, Initiation, Termination and Statistics Report messages. Though the Route Mirroring message contains a BGP Update PDU, as there is no policy execution involved in its transmission, the Common Update TLV has no significance. In all messages except the Route-Monitoring message, the Common Update TLV MUST NOT be included during transmission and MUST be ignored if found on reception.

3.1. Route Monitoring

The Common Update TLV is of relevance only in the Adj-RIB-In and Adj-RIB-Out Route-Monitoring messages and not in the Loc-RIB Route-Monitoring messages.

A Route-Monitoring Update message containing the Common Update TLV MAY also include Adj-In Time and Adj-Out Time Timestamp TLVs defined in [I-D.younsi-grow-bmp-snts]. The Adj-In Timestamp TLV MUST NOT be included if the Common Update TLV does not have I flag or J flag set, and, the Adj-Out Timestamp TLV MUST NOT be included if the Common Update TLV does not have O flag or P flag set. If included, the same must be ignored by the receiver.

3.2. Aggregated Route Monitoring

The Common Update TLV can also be used in the Aggregated BMP Route-Monitoring Message defined in [I-D.liu-grow-bmp-rm-aggregated]. This TLV is of relevance only in the Adj-RIB-In and Adj-RIB-Out Aggregated Route-Monitoring messages and not in the Loc-RIB Aggregated Route-Monitoring messages.

An Aggregated Route-Monitoring Update message containing the Common Update TLV MAY also include Adj-In Time and Adj-Out Time Timestamp TLVs defined in [I-D.younsi-grow-bmp-snts]. The Adj-In Timestamp TLV MUST NOT be included if the Common Update TLV does not have I flag or J flag set, and, the Adj-Out Timestamp TLV MUST NOT be included if the Common Update TLV does not have O flag or P flag set. If included, the same must be ignored by the receiver.

3.3. Statistics Report

This document defines new statistics types that use the following bitmap which is used to indicate a combination of Route-Monitoring views for which routes are the same, i.e. unmodified by policy.

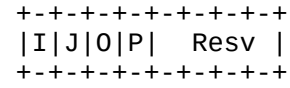

Figure 2: Bitmap of Route-Monitoring views

* I bit - Pre-Policy Adj-RIB-In
* J bit - Post-Policy Adj-RIB-In
* O bit - Pre-Policy Adj-RIB-Out
* P bit - Post-Policy Adj-RIB-Out
* The remaining bits are reserved for future use. They MUST be transmitted as 0 and their values MUST be ignored on receipt.

The following new statistics types are defined.

* Stat Type = TBD2: Number of routes common across a combination of Route-Monitoring views. The value is structured as follows: Bitmap of Route-Monitoring views, Number of routes (64-bit Gauge) common between the views indicated by the bitmap. Multiple instances of this statistics type MAY be included in the same Statistics Report message, each for a unique value of the bitmap.

* Stat Type = TBD3: Number of routes common across a combination of Route-Monitoring views per-AFI/SAFI. The value is structured as follows: 2-byte Address Family Identifier (AFI), 1-byte Subsequent Address Family Identifier (SAFI), Bitmap of Route-Monitoring views, Number of routes (64-bit Gauge) common between the views indicated by the bitmap. Multiple instances of this statistics type MAY be included in the same Statistics Report message, each for a unique value of AFI/SAFI and the bitmap.

4. IANA Considerations

IANA needs to assign the following new parameters to the "BGP Monitoring Protocol (BMP) Parameters" registry (https://www.iana.org/assignments/bmp-parameters/).

4.1. Addition to BMP Route Monitoring TLVs

IANA needs to make the following assignment for the "Common Update TLV" in the "BMP Route Monitoring TLVs" registry.

Type = TBD1 (15 Bits): Common Update TLV

4.2. Additions to BMP Statistics Types Registry

IANA needs to make the following assignment for the statistics types defined in Section 3.3 of this document:

Stat Type	Description
TBD2	Number of routes common across a combination of Route-Monitoring views.
TBD3	Number of routes common across a combination of Route-Monitoring views per-AFI/SAFI.

Table 2: Additions to the "BMP Statistics Types" Registry

5. Security Considerations

This document does not add any additional security considerations. The considerations in Section 11 of [RFC7854] apply to this document.

6. Acknowledgements

TBD

7. 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/info/rfc2119>.

[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>.

[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>.

[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-19, 10 October 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-grow-bmp-tlv-19>.

[I-D.liu-grow-bmp-rm-aggregated] Liu, Y., Lin, C., Srivastava, M., and N. Prasad, "Definition for Aggregated BMP Route Monitoring Message", Work in Progress, Internet-Draft, draft-liu-grow-bmp-rm-aggregated-03, 18 June 2025, <https://datatracker.ietf.org/doc/html/draft-liu-grow-bmp-rm-aggregated-03>.

[I-D.younsi-grow-bmp-snts] Younsi, M. and P. Francois, "BMP Sequence Number, Timestamp and Flags TLVs", Work in Progress, Internet-Draft, draft-younsi-grow-bmp-snts-01, 17 October 2025, <https://datatracker.ietf.org/doc/html/draft-younsi-grow-bmp-snts-01>.

Authors' Addresses

Dhananjay Patki
Cisco Systems, Inc.
Cessna Business Park SEZ, Kadubeesanahalli
Bangalore 560103
Karnataka
India

Email: dhpatki@cisco.com

Prasad S. Narasimha
Cisco Systems, Inc.
Cessna Business Park SEZ, Kadubeesanahalli
Bangalore 560103
Karnataka
India
Email: snprasad@cisco.com