

Network Working Group
Internet Draft
Updates: 7854 (if approved)
Intended status: Standards Track
Expires: December 16, 2025

Y. Liu
China Mobile
C. Lin
New H3C Technologies
M. Srivastava
Juniper Networks
P. Narasimha
Cisco Systems, Inc.
June 16, 2025

Definition for Aggregated BMP Route Monitoring Message
draft-liu-grow-bmp-rm-aggregated-03

Abstract

This document proposes an aggregated BMP route monitoring message. It can compress multiple BMP route monitoring messages into one aggregated BMP route monitoring message to reduce the amount of reported BMP route monitoring messages and reduce network overhead. This document updates RFC 7854 by adding new BMP Messages type.

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 December 16, 2025.

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
1.1. Requirements Language.....	3
2. Solution.....	3
2.1. Current BGP Group Packaging.....	3
2.2. BMP send batching optimization.....	4
3. Aggregated Route Monitoring Definition.....	6
4. Aggregated Route Monitoring Message Format.....	6
4.1. Route Monitoring Message.....	7
4.2. Support for various RIB-Views.....	8
4.3. Examples of multi-peer headers.....	8
5. Security Considerations.....	10
6. IANA Considerations.....	10
7. Normative References.....	11
Authors' Addresses.....	12

1. Introduction

[RFC7854] defines BMP route monitoring message, which is used to send incremental routes advertised and withdrawn by peers to the monitoring station. BMP route monitoring message consists of Common Header, Per-Peer Header and BGP Update PDU. Among them, Common Header and Per-Peer Header are defined in [RFC7854], and the BGP Update PDU contains the BGP PATH attribute and prefix, as shown in Figure 1.

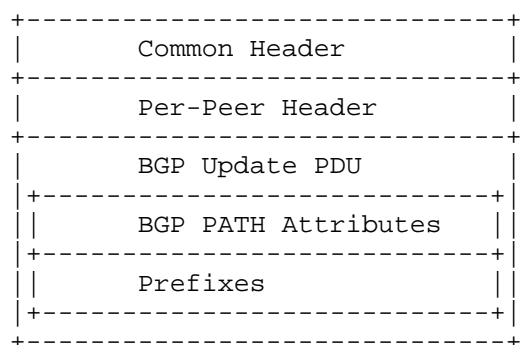


Figure 1: BMP Route Monitoring Message Structure

Currently, a piece of BMP route monitoring message can only contain one Common Header, one Per-Peer Header and one BGP Update PDU, and the BGP Update PDU can contain multiple non-repeatable BGP PATH attributes and prefixes.

This document proposes an aggregated BMP route monitoring message. It can compress multiple BMP route monitoring messages into one aggregated BMP route monitoring message to reduce the amount of reported BMP route monitoring messages and reduce network overhead.

1.1. 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 RFC 2119 [RFC2119].

2. Solution

2.1. Current BGP Group Packaging

The rapid growth of existing network routing tables and the complexity of network topology have led to the need for BGP to support more peers. In particular, in scenarios with a large number of peers and a large amount of routes, the router needs to send routes to a large number of BGP peers, and most of the peers have the same configuration, which requires higher packaging and sending performance.

The BGP group packaging technology treats all BGP peers with common configurations as a packaging group. In this way, each route to be sent is packaged only once and then sent to all neighbors in the packaging group, which exponentially improves the packaging efficiency.

Before the group packaging feature was supported, each route to be sent had to be packaged separately for each peer. Group packaging enables unified packaging and separate sending, that is, each route to be sent is packaged only once and then sent to all peers in the packaging group, which exponentially improves the efficiency of packaging and sending. In scenarios with a large number of neighbors and a large amount of routes, as shown in Figure 2, group packaging greatly improves the performance of BGP packaging and sending.



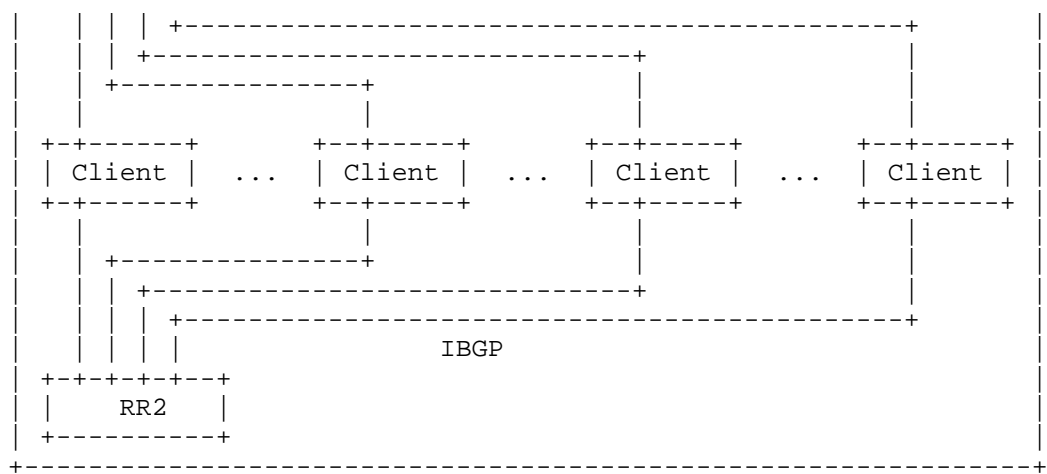


Figure 2: Typical application scenarios for group packaging

Figure 2 is a typical network diagram of reflectors with multiple clients, RR routers need to send routes to a large number of BGP client peers, and most of the client peers have the same configuration. Suppose an RR reflector has 100 clients and 100,000 routes to be reflected. If each neighbor is packaged separately, when the reflector RR sends routes to 100 clients, the total number of times all routes are packaged is 100,000 x 100. The group packaging technology reduces this process to 100,000 x 1, which is equivalent to a 100-fold improvement in performance.

The comparison of packet assembly by peer and peer packaging group is shown in Figure 3. It can be clearly seen that assembling packets by peer packaging group can greatly improve packet performance.

Encapsulation by Peer	Encapsulation by Peer Packaging Group
N Peers	N Peers of Peer Packaging Group
N Times Packaging	1 Times Packaging
N Times Sending	N Times Sending

Figure 3: Comparison of Two Packaging Methods

2.2. BMP send batching optimization

Currently, network devices have implemented BGP send batching technology. This allows messages with the same attributes to be batched together and sent to multiple neighbors simultaneously.

However, BMP messages are handled on a per-neighbor basis, which means they cannot take advantage of the BGP send batching functionality.

From the perspective of BMP packet format, if BMP packets are also assembled according to peers, they need to be assembled once for each peer, and the assembly performance will also be limited. Moreover, the BMP message information is different depending on the peer, and needs to be sent to the monitoring server multiple times, which increases network overhead.

In multiple BMP route monitoring messages, if the prefixes are the same but the Per-Peer Header and BGP PATH attributes are different, according to the way of packet assembly by peer packaging group, the different BGP attributes can be extracted, combined with the corresponding Per-Peer Header, and reuse of the same BGP PATH attributes, which together form a aggregated BMP route monitoring message. See section 4 for detailed format of the aggregated BMP route monitoring message. As shown in Figures 4 and 5, compared with the original multiple BMP route monitoring message, the aggregated BMP route monitoring message exponentially reduces the Common Header and the same BGP PATH attributes and prefixes, and is only assembled once, which not only effectively the network overhead is reduced and the assembly performance is further improved.

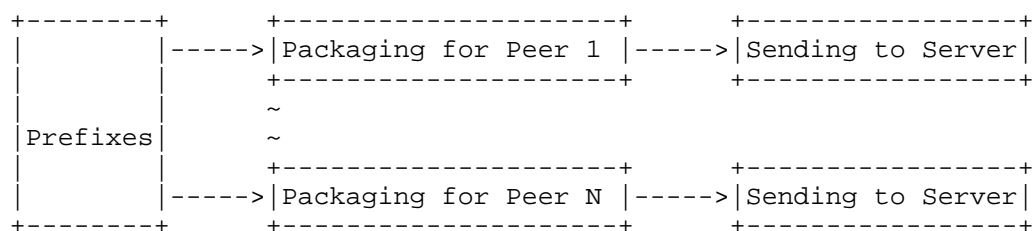


Figure 4: BMP Encapsulation by Peer

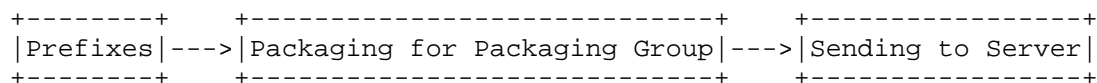


Figure 5: BMP Encapsulation by Peer packaging Group

This document defines this aggregated BMP routing monitoring message, and its format has been greatly adjusted based on the original BMP routing monitoring message format, see section 4 for its detailed format.

3. Aggregated Route Monitoring Definition

This section adds a new BMP message type for BMP route monitoring, which is populated in the message type field of the Common Header.

Message Type = TBD: Aggregated Route Monitoring, Recommended value 7.

4. Aggregated Route Monitoring Message Format

This section defines the aggregated BMP route monitoring message format, as shown in Figure 6, including Common Header, Multi-Peer Header and BGP Update PDU. Among them, the Common Header is the same as that defined in [RFC7854], the BGP Update PDU contains the same BGP PATH attribute and prefixes, and the Multi-Peer Header will be defined below.

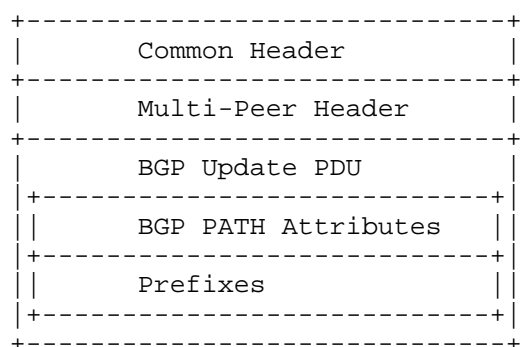
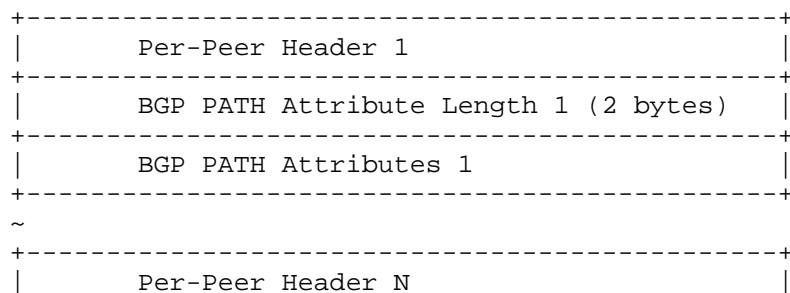


Figure 6: BMP Aggregated Route Monitoring Message Format

As shown in Figure 7, the format of the Multi-Peer Header in the BMP aggregated route monitoring message is defined, which contains multiple Per-Peer Headers. Each Per-Peer Header could carry the unique BGP PATH attribute of the corresponding peer route. If no BGP PATH attribute is carried, the corresponding BGP PATH attribute length must be 0.



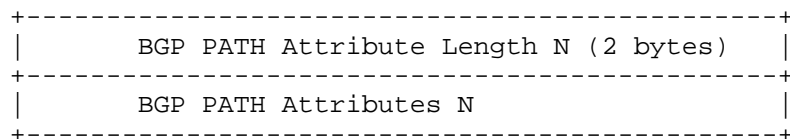


Figure 7: Multi-Peer Header Format

In the Multi-Peer Header format, the Per-Peer Header format is the same as that defined in [RFC7854].

BGP PATH Attribute Length (2 bytes) indicates the length of the BGP PATH attribute in the Multi-Peer Header.

The format of the BGP PATH Attribute in the Multi-Peer Header is the same as that in the BGP Update PDU. The BGP PATH Attribute area does not need to be filled in when the route attributes corresponding to each peer are exactly the same. Only when the routing attributes corresponding to the peers are different to a certain extent, different parts of the routing attributes need to be filled in BGP PATH attribute area according to the peers. The same parts of the routing attributes are filled in the BGP Update PDU. If the routing attributes corresponding to each peer are completely different, the routing attributes in the BGP Update PDU will be empty.

If the BGP PATH Attribute in Multi-Peer Header exists, it means that the BGP PATH Attribute of the BGP Update PDU needs to be integrated with the BGP PATH Attribute of Multi-Peer Header to obtain the complete BGP PATH Attribute of BGP UPDATE PDU sent or received to the corresponding peer. Otherwise, the BGP PATH Attribute of the BGP UPDATE PDU is the complete BGP PATH Attribute.

The BGP PATH Attributes in Multi-Peer Header is optional. A multi-peer header may just include the Per-peer header(s) without any PATH attributes following them.

Multiple Aggregated Route Monitoring messages can be used for same or different set of BGP Update prefixes.

4.1. Route Monitoring Message

Sender MUST send either the Route Monitoring Message (type 0) as defined in [RFC7854] OR the Aggregate Route Monitoring Message and MUST NOT combine the two message formats in the updates

4.2. Support for various RIB-Views

Aggregated Route Monitoring messages can be used for any of the RIB-views (Adj-RIB-In pre, Adj-RIB-In post, Adj-RIB-Out pre, Adj-RIB-Out post & Local-RIB) when batching is feasible.

Batching across RIB-views with Aggregated Route Monitoring Message can be used leveraging methods defined in [I-D. patki-grow-bmp-common-updates]

4.3. Examples of multi-peer headers

Figure 8: Multiple peers (3) with no per-peer header PATH attributes along with all PATH attributes from BGP Update PDU being shared.

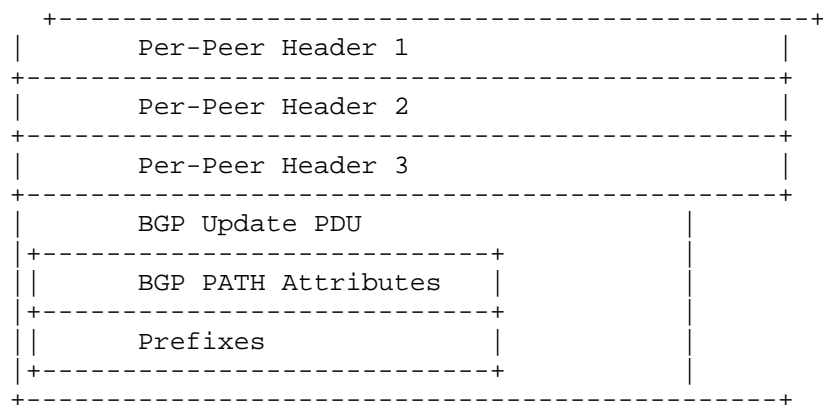
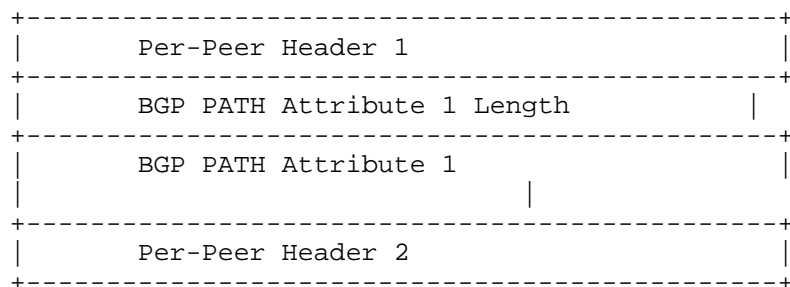


Figure 8: Multiple peers (3) with no per-peer header

Figure 9: Multiple peers (3) with one of the peers having different PATH attribute along with shared PATH attributes in the BGP Update PDU.



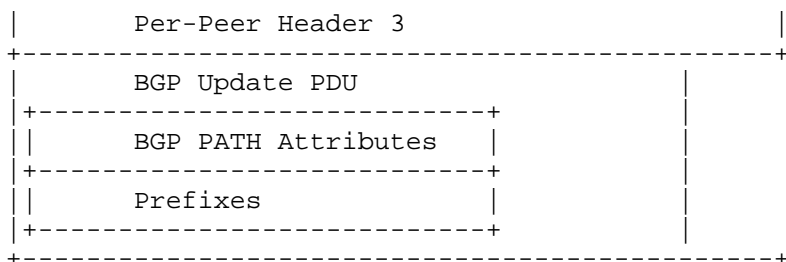


Figure 9: Multiple peers (3) with one of the peers having different PATH attribute

Figure 10: Multiple peers (3) with all the peers having different PATH attributes along with shared PATH attributes in BGP Update PDU.

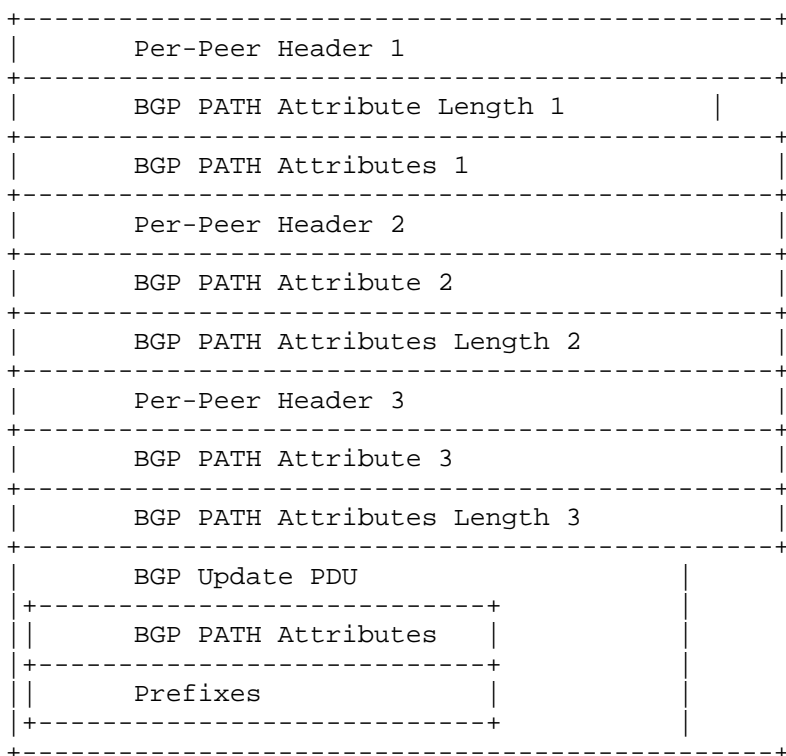


Figure 10: Multiple peers (3) with all the peers having different PATH attributes along with share BGP Path

Figure 11: Multiple peers (3) with all the peers having different PATH attributes along with no shared PATH attributes in BGP Update PDU

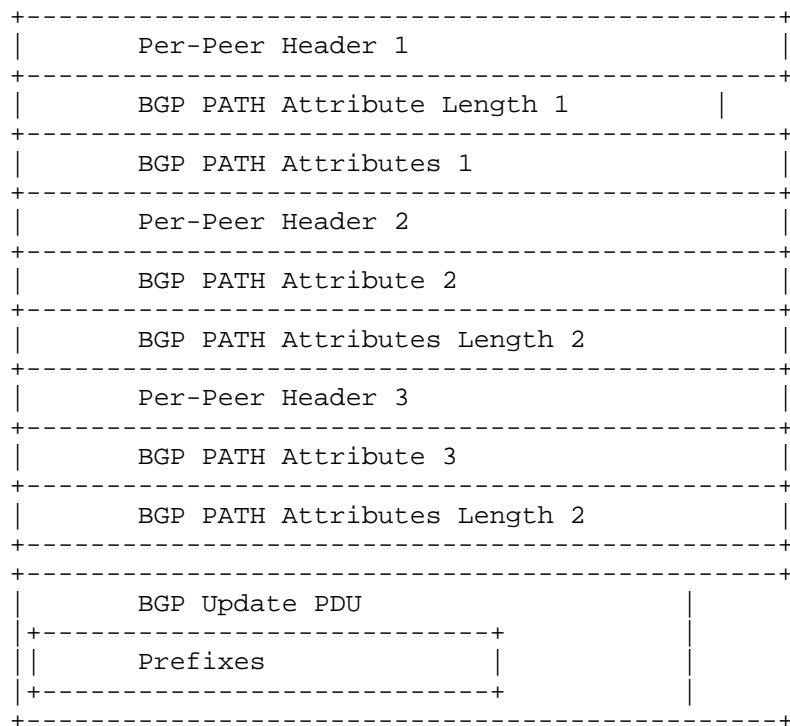


Figure 11: Multiple peers (3) with all the peers having different PATH attributes along with no share BGP Path

5. Security Considerations

The considerations in Section 11 of [RFC7854] apply to this document. It is also believed that this document does not add any additional security considerations.

6. IANA Considerations

This document requests that IANA assign the following new parameters to the BMP parameters name space (<https://www.iana.org/assignments/bmp-parameters/bmp-parameters.xhtml>).

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/rfc/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>>.
- [I-D. patki-grow-bmp-common-updates] Patki D. and Narasimha P., "Common BMP Route-Monitoring Messages for Routes Unchanged by Policy" Work in Progress, Internet-Draft, draft-patki-grow-bmp-common-updates-01, 28 February 2025, <<https://datatracker.ietf.org/doc/html/draft-patki-grow-bmp-common-updates-01>>

Authors' Addresses

Yisong Liu
China Mobile
China
Email: liuyisong@chinamobile.com

Changwang Lin
New H3C Technologies
China
Email: linchangwang.04414@h3c.com

Mukul Srivastava
Juniper Networks
10 Technology Park Dr
Westford, MA 01886
United States of America
Email: msri@juniper.net

Prasad S. Narasimha
Cisco Systems, Inc.
Email: snprasad@cisco.com

