

Global Routing Operations
Internet-Draft
Updates: 7854 (if approved)
Intended status: Standards Track
Expires: 2 September 2025

P. Lucente
C. Cardona
C. Petrie
NTT
L. Hendriks
NLnet Labs
1 March 2025

BMP State Summaries
draft-lucente-grow-bmp-offline-01

Abstract

BMP (BGP Monitoring Protocol) is perfectly suited for real-time consumption but less ideal in stream processing and off-wire historical scenarios. The main issue is that the ability to correctly parse BGP Update PDUs, carried in BMP Route Monitoring messages, depends on the BGP Capabilities exchanged during the establishment of the BGP session between the peers via BGP Open PDUs. The BGP Open PDUs, carried in BMP Peer Up Notification messages, are exported at the establishment of the BMP session. Similar to BGP, BMP sessions are typically long-lived, so the crucial information to correctly parse subsequent messages of such sessions was possibly sent a relatively long time ago (days, weeks, months).

This document introduces the concept of Summaries. It defines a new optional BMP message type, called State Summary, and a new TLV, called Summary Id. A Summary is similar to the initial synchronisation performed upon establishment of the BMP session: all BGP session information is exported in Peer Up Notification messages, and all RIB contents are exported in Route Monitoring messages. All the messages carry the new Summary Id TLV, containing an ID uniquely identifying the summary these messages belong to. The messages are preceded by a State Summary message carrying the same Summary Id TLV, as well as meta-data describing the Summary.

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 2 September 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
2. Terminology	4
3. State Summary message format	4
3.1. Summary Information TLVs	4
3.1.1. Summary Id TLV	4
3.1.2. Optional meta-data TLVs	4
3.2. Third party off-wire encoding formats	6
4. Operational Considerations	6
4.1. Summary Id scheme	6
4.1.1. Global uniqueness	6
4.1.2. Increasing identifiers	6
5. Security Considerations	6
6. IANA Considerations	7
7. Normative References	7
Acknowledgements	8
Authors' Addresses	8

1. Introduction

Correctly parsing BGP Update PDUs included in BMP messages (ie. Route Monitoring) does require a stateful approach by keeping track of capabilities exchanged in the BGP Open PDUs as reported by BMP Peer Up Notification messages.

Peer Up Notification messages are sent only during the initial synchronisation at the start of the BMP session, or, whenever a BGP session is established on the monitored router. The necessary information in the BGP Open (and Peer Up Notification) messages may thus have been received long before its needed to parse incoming BGP Updates in Route Monitoring messages. This inevitable stateful approach might be challenging in certain scenarios, such as consuming archived BMP messages or deployments where BMP Stations or processing nodes not necessarily see the (complete) start of every BMP stream.

TLV support for BMP Route Monitoring and Peer Down Messages [I-D.ietf-grow-bmp-tlv] defines a Stateless Parsing TLV aimed at including relevant capabilities that have an impact in BGP Update message parsing as part of optional informational TLV in Route Monitoring messages. While the method is valid, in fact it does allow with minor effort to encapsulate BMP in MRT format for offline consumption as documented by Storing BMP messages in MRT Format [I-D.petrie-grow-mrt-bmp], it comes with some drawbacks like extra verbosity and increased correlation effort at a BMP exporter, where resources may be limited.

Similar to how the necessary information carried in the BGP Open messages is sent from the exporter to the station only in the beginning of a BMP session (in forms of Peer Up Notifications), the contents of the RIBs is also only sent at the beginning (in forms of Route Monitoring messages). In other words, to make correct and complete sense of offline BMP data, both current state (RIB contents) and session details (BGP Open Capabilities) to parse future state changes are required.

This document introduces Summaries, enabling synchronisation of both BGP session information and RIB contents anywhere in a BMP session. Summaries are enabled by the new optional State Summary message type and the new Summary Id TLV, both introduced in this document. A Summary is a collection of Peer Up Notification messages, Route Monitoring messages and a single State Summary message, all carrying the newly introduced Summary Id TLV containing the unique identifier for that Summary. The State Summary message further contains TLVs with meta-data describing when the Summary was created, information on the exporting side such as sysName and IP address, and/or information on the collecting side such as BMP station software name and version, etc.

The new concepts described in this document are not restricted to either the BMP exporter or the BMP station. By building upon TLVs, supported in all BMP message types from BMPv4, the Summary approach imposes minimal requirements over the initial synchronisation in BMP today. Furthermore, if at any point in the future another message

type needs to be incorporated in a Summary, it will be a simple matter of attaching the Summary Id TLV to those messages. No existing message types have to be adapted to support Summaries.

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. State Summary message format

The State Summary message starts with a BMP Common Header as defined in Section 4.1 of [RFC7854]. The Common Header is directly followed by TLVs. The Summary Id TLV defined in Section 3.1.1 MUST be present and SHOULD be the first of the TLVs. All additional TLVs listed in Section 3.1.2 are optional.

The State Summary message can be generated by a BMP exporter or by a BMP station, as long as all required data for the Peer Up Notification and Route Monitoring messages pertaining to the Summary are available. The State Summary message MUST be followed by those messages, with the Summary Id TLV attached.

3.1. Summary Information TLVs

3.1.1. Summary Id TLV

The Summary Id TLV is the only mandatory TLV of a State Summary message as defined in Section 3. It is an indexed TLV (as it will be included in Route Monitoring messages, with index zero), structured as defined in Section 4 of [I-D.ietf-grow-bmp-tlv], with a fixed value length of 16 bytes. This allows the use of UUID identifiers, or provides sufficient space for alternative schemes. Different approaches for schemes are discussed in Section 4.

3.1.2. Optional meta-data TLVs

The Peer Summary message SHOULD carry TLVs providing additional information on the BMP session being summarized. These Summary Information TLVs describe the BMP exporter and station involved, and the date and time the summary was generated. By embedding these TLVs in the offline file, a consumer of the file does not have to rely on the filename or other external data to get these types of information. All TLVs are non-indexed.

- * Type = TBD3: Datetime of summary
 - Length: 8 bytes
 - Value: 64bit UNIX epoch, in seconds
- * Type = TBD4: Exporter IP address
 - Length: 16 bytes
 - Value: IPv6 or IPv4-mapped IPv6 address
- * Type = TBD5: Exporter sysName
 - Length: variable, non-zero, describing the number of bytes
 - Value: UTF-8 string
- * Type = TBD6: Exporter sysDesc
 - Length: variable, non-zero, describing the number of bytes
 - Value: UTF-8 string
- * Type = TBD7: Station IP address
 - Length: variable, non-zero, describing the number of bytes
 - Value: IPv6 or IPv4-mapped IPv6 address
- * Type = TBD8: Station sysName
 - Length: variable, non-zero, describing the number of bytes
 - Value: UTF-8 string
- * Type = TBD9: Station sysDesc
 - Length: variable, non-zero, describing the number of bytes
 - Value: UTF-8 string

3.2. Third party off-wire encoding formats

While this document does define a way to facilitate stream processing, replay and, more in general, consumption of raw BMP data offline, similar benefits may be harnessed by third party off-wire formats in replay and, more in general, consumption of raw BMP data offline, similar benefits may be harnessed by third party off-wire formats in which BMP can be encapsulated into, for example MRT (Multi-Threaded Routing Toolkit) as defined by RFC 6396 [RFC6396]. As a result of that, this document does not recommend a preferred way to stream process or store BMP data offline.

4. Operational Considerations

4.1. Summary Id scheme

The generation and form of the Summary Ids introduced in this document is left to implementations. This document does not enforce any specific approach, though at least the following points should be considered. Note that implementations are not limited to supporting only one Id scheme, but ideally support multiple schemes via local configuration.

4.1.1. Global uniqueness

In deployments where information is received from multiple BMP vantage points, unique Summary Ids might prove handy or even crucial in order to distinguish Summary A originally sent by BMP exporter X, from Summary B sent by exporter Y. If all exporting processes rely on an algorithm producing globally unique identifiers, e.g. UUID version 4, they all can send out Summaries without possibly using an identical Summary Id generated by another exporter.

4.1.2. Increasing identifiers

Generating (linearly) increasing identifiers enable the BMP station to order Summaries, and, to spot any missing Summaries. Furthermore, in a (long running) BMP session where the exporter generates Summaries, the Summary Id doubles as a counter signalling how many Summaries have been sent so far. Note that some of these can be deduced via other means: ordering of Summaries can be done based on the Timestamp TLV (TBD3), and the number of sent Summaries could be included in the Stats Report message (Section 4.8 of [RFC7854]).

5. Security Considerations

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

6. IANA Considerations

IANA is asked to allocate a new Peer Summary message type in the BMP Message Types registry with value TBD1. IANA is also asked to to create a registry within the BMP group, named "BMP Peer Summary Message TLVs".

Registration procedures for this registry are:

Range		Registration Procedures
0-32767		Standards Action
32768-65530		First Come, First Served
65531-65534		Experimental
65535		Reserved

Table 1

Initial values for this registry are:

Type	Description	Reference
TBD2	Peer	this document

Table 2

7. Normative References

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

Lucente, P. and C. Cardona, "Logging of routing events in BGP Monitoring Protocol (BMP)", Work in Progress, Internet-Draft, draft-ietf-grow-bmp-rel-02, 8 July 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-grow-bmp-rel-02>>.

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

Lucente, P. and Y. Gu, "BMP v4: TLV Support for BGP Monitoring Prtoocol (BMP) Route Monitoring and Peer Down Messages", Work in Progress, Internet-Draft, draft-ietf-grow-bmp-tlv-16, 24 February 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-grow-bmp-tlv-16>>.

[I-D.petrie-grow-mrt-bmp]

Petrie, C., "Storing BMP messages in MRT Format", Work in Progress, Internet-Draft, draft-petrie-grow-mrt-bmp-00, 1 November 2019, <<https://datatracker.ietf.org/doc/html/draft-petrie-grow-mrt-bmp-00>>.

[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>>.

[RFC6396] Blunk, L., Karir, M., and C. Labovitz, "Multi-Threaded Routing Toolkit (MRT) Routing Information Export Format", RFC 6396, DOI 10.17487/RFC6396, October 2011, <<https://www.rfc-editor.org/info/rfc6396>>.

[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>>.

[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>>.

Acknowledgements

TBD

Authors' Addresses

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

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

Colin Petrie
NTT
Veemweg 23
3771 MT Barneveld
Netherlands
Email: colin@ntt.net

Luuk Hendriks
NLnet Labs
Science Park 400
1098 XH Amsterdam
Netherlands
Email: luuk@nlnetlabs.nl