

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 17, 2020

M. Srivastava
Juniper Networks
P. Lucente
NTT
October 15, 2019

BMP Compression
draft-msri-grow-bmp-compression-00

Abstract

This document provides specification for an optional compressed BMP Feed from a router to BMP station.

Requirements Language

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.

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 April 17, 2020.

Copyright Notice

Copyright (c) 2019 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Procedures	3
2.1. Starting Compressor Capability	3
2.2. Compression Information TLV	3
2.3. Compressed BMP Messages	4
2.4. Compressor Overflow	4
2.5. Error Handling	4
2.6. Processing of Compressed Route Monitoring messages	5
3. Acknowledgements	5
4. IANA Considerations	5
4.1. BMP Compression Information TLV	5
4.2. BMP Compression Route Monitoring message type	5
5. Security Considerations	5
6. Normative References	5
Authors' Addresses	6

1. Introduction

The BGP Monitoring Protocol (BMP) allows monitoring of Rib-in RFC7854 [RFC7854], Loc-Rib, BGP local-rib [I-D.ietf-grow-bmp-local-rib] and Rib-in and Rib-Out monitoring allows pre-policy and post-policy view of the prefix. Thus, for a scaled setup, with all these kinds of monitoring enabled, BMP will get a lot of back pressure in the protocol as it needs to dump a huge data for its monitored peers, through a single socket towards BMP station. BGP update PDU which is part of the BMP Route-monitoring (RM) message is also increasing. It is no more limited to 4K as noted in draft-ietf-idr-bgp-extended-messages-21. Essentially, BMP is heading towards becoming I/O bound monitoring protocol. This document proposes compression of BMP feed towards BMP station. Compression will ease the pressure on TCP socket between a router and BMP station. Such a scheme would be useful if a route can spare some extra CPU for BMP operation.

As it must be obvious, this scheme will require compressor mechanism at the BMP speaking router and a decompressor on the BMP station. The compression mechanism used at the BMP speaking is an implementation specific detail and is beyond the scope of this specification.

2. Procedures

2.1. Starting Compressor Capability

BMP compression feature on the router and BMP decompressor feature on the BMP station has to be present at the same time. Enabling compression feature at router end only will lead to incomprehensible data at the BMP station end. Also same technique should be used to compress and decompress the data on wire. Using different technique to compress and decompress would lead to incomprehensible data at the BMP station end.

BMP compression feature on the router and BMP decompressor feature on the BMP station can be enabled via configuraton. Once this feature is enabled between router and BMP station, the monitored router should indicate this to the BMP Station using new Compression Information TLV as described in following section.

From that point onwards, the router would send the compressed BMP feed towards BMP station. BMP session needs to be bounced every-time this feature is enabled on a current active BMP session.

2.2. Compression Information TLV

As noted in RFC7854 [RFC7854], the initiation message provides a means for the monitored router to inform the monitoring station of its vendor specific details. It can carry Information TLVs containing information about the monitored router.

The monitored router MUST communicate the compression capability to BMP staton using Compression Information TLV described below.




Figure 1: Compression Information TLV

- o Type = TDB1 (2 Octets): Compression Information TLV type.
- o Length (2 Octets): indicates the length of the value field of the Compression Information TLV. The value field further consists of the Compression string.

- o CM (4 bits): CM indicating DEFLATE compressed format value as specified in RFC1950.
- o CINFO (4 bits): INFO as specified in RFC1950. Invalid values MUST lead to the capability being ignored. The compressing peer MUST use this value for the parametrization of its algorithm.

2.3. Compressed BMP Messages

Following rules should be following for achieving BMP feed compression:

1. A new message type, Compressed Route Monitoring (CRM), MUST be used. This is to ensure backward compatibility with BMP stations that do not support the compression capability. The message type is same in structure as described by TLV support for BMP Route Monitoring and Peer Down Messages [I-D.ietf-grow-bmp-tlv]. Compression is to be applied only to this message type, all other BMP message types shall not be compressed.
2. Compression is applicable to all the payload following the Common Header, described in Section 4.1 of [RFC7854]. This allows to read the total BMP message length, i.e. to perform sanity checks against socket and compressor information.
3. Each compressed BMP message MUST be sent as a block, i.e. the decompression MUST be able to yield decompressed results of the without waiting for further compressed updates. This is different from the normally used stream compression mode.
4. The compressed message MAY exceed the maximum message size but in such case compressor overflow per Section 2.4 MUST be invoked.

2.4. Compressor Overflow

This should be handled in same was as described in draft-przygienda-idr-compressed-updates [I-D.przygienda-idr-compressed-updates].

2.5. Error Handling

If the decompression on the BMP station fails for any reason, it needs to bring down the BMP session.

If the compression on the monitoring router fails for any reason, it is at the discretion of the router to handle it. It may try it few more times. In the worse case it MAY bring down the BMP session

2.6. Processing of Compressed Route Monitoring messages

A BMP station receiving a compressed message SHOULD process it as follows:

1. Decode the BMP Common Header where message length is specified
2. Decompress remainder of the Compressed Route Monitoring message and determine the decompressed message size from the decompressor
3. Decode the BMP Per-peer header
4. Decode the BGP UPDATE PDU header to infer the presence of trailing TLVs
5. Decode the BMP message TLVs
6. Decode the actual BGP UPDATE PDU

3. Acknowledgements

TBD.

4. IANA Considerations

This document requests that IANA assign the following new parameters to the BMP parameters name space.

4.1. BMP Compression Information TLV

This document defines the BMP Compression Information TLV Header with Type = TBD (Section 2.2).

4.2. BMP Compression Route Monitoring message type

This document also defines the BMP Compressed Route Monitoring message type with Type = TBD (Section 2.3).

5. Security Considerations

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

6. Normative References

[I-D.ietf-grow-bmp-local-rib]

Evans, T., Bayraktar, S., Bhardwaj, M., and P. Lucente,
"Support for Local RIB in BGP Monitoring Protocol (BMP)",
draft-ietf-grow-bmp-local-rib-05 (work in progress),
August 2019.

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

Lucente, P., Gu, Y., and H. Smit, "TLV support for BMP
Route Monitoring and Peer Down Messages", draft-ietf-grow-
bmp-tlv-01 (work in progress), October 2019.

[I-D.przygienda-idr-compressed-updates]

Przygienda, T., Lingala, A., Mate, C., and J. Tantsura,
"Compressed BGP Update Message", draft-przygienda-idr-
compressed-updates-07 (work in progress), August 2019.

[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>>.

[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>>.

Authors' Addresses

Mukul Srivastava
Juniper Networks
10 Technology Park Drive
Westford MA 01886
USA

Email: msri@juniper.net

Paolo Lucente
NTT
Siriusdreef 70-72
Hoofddorp, WT 2132
Netherlands

Email: paolo@ntt.net