Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 20 April 2026

P. Francois
M. Younsi
INSA-Lyon
P. Lucente
NTT

17 October 2025

BMP Loc-RIB: Peer address
draft-ietf-grow-bmp-loc-peer-02

Abstract

BMP Loc-RIB [RFC9069] enforces that the BMP router sets the Peer Address value of a path information to zero. This document introduces the option to communicate the actual peer from which a path was received when advertising that path with BMP Loc-RIB.

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 [RFC2119] [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 20 April 2026.

Copyright Notice

Table of Contents

1. Introduction

Using BMP Loc-RIB [RFC9069], the Peer Address field of a Per-Peer header is Zero-filled. This prevents a collector from knowing from which peer a path selected as best was received. The next-hop attribute of a path is indeed not an identifier of the peer from which the path was received. Knowing the peer address is also especially useful when Loc-RIB paths come from Add-Path [RFC7911] enabled peers as the path identifier space of paths are defined per peer.

When virtual routing and forwarding (VRFs) are in use, the peer address information can only be interpreted in the VRF context within which the corresponding peering is taking place.

This document introduces a BMPv4 [I-D.ietf-grow-bmp-tlv] TLV describing the address of the peer that announced the path to the BMP router, and other TLVs describing the VRF context in which a path was received.

2. TLV Based Behavior

This section describes a solution based on BMPv4 TLVs. Section 2.1 describes a BMPv4 TLV used to convey the peer address. Section 2.2 introduces optional TLVs for the case of paths imported from another VRF.

2.1. Rx Peer-Address TLV

BMP TLVs can be used to provide optional information along with monitored paths. Peer Address information can be included using one such TLV.

The "Rx Peer-Address TLV" TLV type is TBD1 (see IANA section). The value of the TLV is the "Address Type" code followed by the address of the peer from which the monitored path was received. The address type 0 is reserved and MUST NOT used. A set of address types is described in the following subsections.

The length field is one (for the "Address Type" field) plus the length of the "Rx Peer Address" field. The "Index" field is, as described by [I-D.ietf-grow-bmp-tlv], not included in the length.

The TLV structure is illustrated in Figure 1.

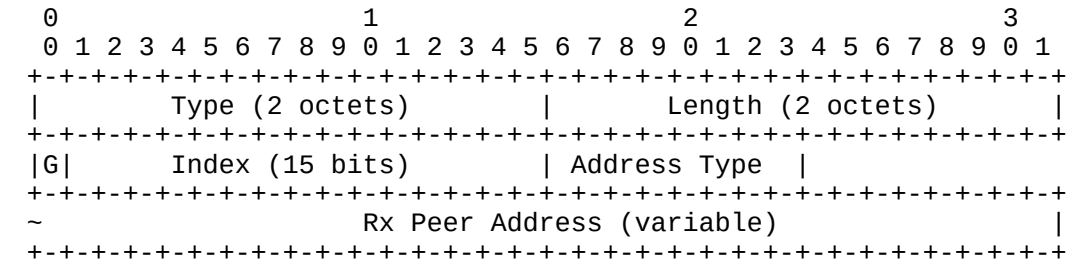

Figure 1: Rx Peer-Address TLV

2.1.1. Self-Originated

The "Rx Peer-Address TLV" may describe a self-originated path by setting the value of the "Address Type" to 1. The "Rx Peer Address" is empty. The "Length" is thus set to 1.

2.1.2. IPv4 Peer Address

In case of a BGP peering established using IPv4, the "Address Type" is set to 2. The "Rx Peer Address" is the 4 bytes IPv4 Address of the peer. The "Length" is thus set to 5.

2.1.3. IPv6 Global Unicast Address

In case of a BGP peering established using an IPv6 Global Link Address, the "Address Type" is set to 3. The "Rx Peer Address" is the 16 bytes IPv6 Global Link Address of the peer. The "Length" is thus set to 17.

2.1.4. IPv6 Address with Interface ID

In some scenarios, for example, in the case case of a BGP session established using IPv6 Link Local Addresses, an interface identifier is needed to disambiguate the address. The "Address Type" is set to 4. The "Rx Peer Address" is the 16 bytes IPv6 Address of the peer, followed by an interface ID of a variable size S. The "Length" is thus set to 1 + 16 + S.

2.1.5. IPv6 Address with Interface Name

Similar to Section 2.1.4 but with interfaces identified using a name instead of an ID, the "Address Type" is set to 5. The "Rx Peer Address" is the 16 bytes IPv6 Address of the peer, followed by an interface name of a variable size S, encoded in UTF-8 without specific termination characters. The "Length" is thus set to 1 + 16 + S.

2.2. VRF Import TLVs

Path information advertised through BMP Loc-RIB might be related to a path imported from another VRF. In such a scenario, the sole knowledge of the remote peer IP address is not sufficient to unambiguously determine the origin of this path.

2.2.1. Origin VRF TLV

The "Origin VRF TLV" describes the VRF context in which this path was received from a peer or where it was self-originated. It contains a variable-length value field matching the definition of VRF/Table name from [RFC9069]. The value of the type field of this TLV is TBD2. The length field of this BMPv4 TLV is the length, in bytes, of the UTF-8 string of the VRF name. When this TLV is present, the Rx Peer-Address TLV associated with that path refers to the IP address of the peer from which it was received, in the VRF context refered in this TLV. The format of the Origin VRF TLV is shown in Figure 2.

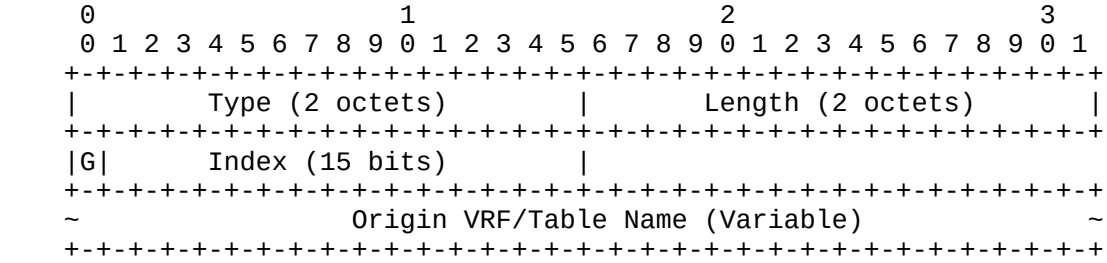

Figure 2: Origin VRF TLV Format

2.2.2. Previous VRF TLV

The "Previous VRF TLV" describes the VRF from which this path was imported. It contains a variable-length value field matching the definition of VRF/Table name from [RFC9069]. The value of the type field of this TLV is TBD3. The length field of this is the length, in bytes, of the UTF-8 string of the VRF name. The format of the "Previous VRF TLV" is shown in Figure 3.

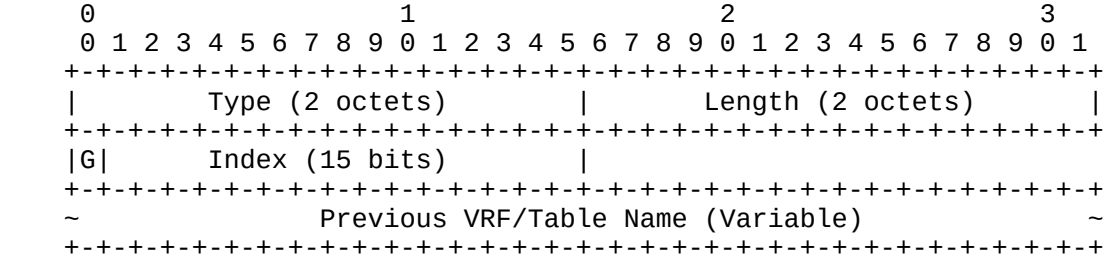

Figure 3: Previous VRF TLV Format

For example, if BMP Loc-RIB describes a path P in VRF C, which was received from a peer I in VRF A, imported into VRF B, and finally imported from VRF B into VRF C, the Origin VRF Name is A, the Previous VRF Name is B, the VRF/Table Name TLV (per [RFC9069] is C, and the Rx Peer-Adress TLV is I.

2.2.3. Previous VRF Sequence TLV

The "Previous VRF Sequence" describes the entire chain of VRFs through which this path was imported before landing in the current VRF. The list starts with the previous VRF, and ends with the Origin VRF in which this path was received or originated. The length field is an 8-bit value capturing the length, in bytes, of the Name field. The name field is the VRF name of the described VRF of the sequence, matching the definition of VRF/Table name from [RFC9069]. A complete Previous VRF Sequence TLV structure is shown in Figure 4.

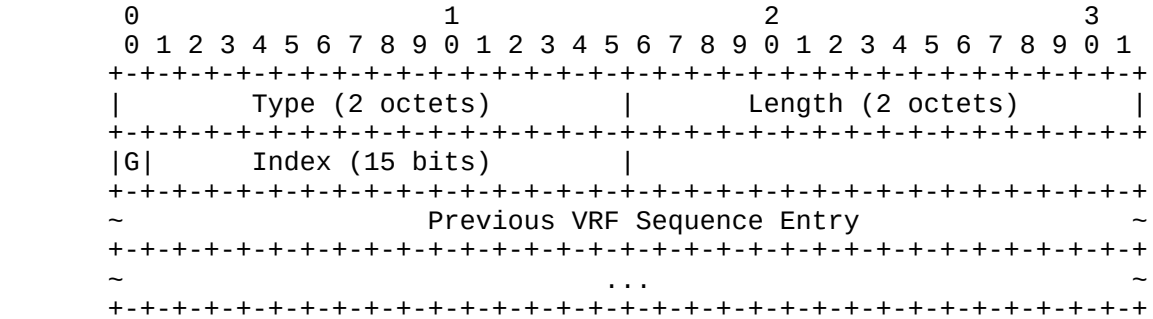

Figure 4: Previous VRF Sequence TLV Format

The format of each entry is shown in Figure 5.

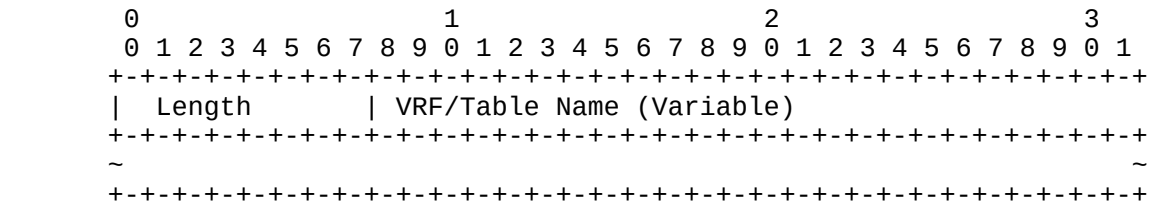

Figure 5: Previous VRF Sequence Entry Format

The value of the type field of this TLV is TBD4.

The length of a "Previous VRF Sequence" TLV is the sum of the total lengths of each VRF entry in the sequence (1 byte for the length field + the value of the length field). This does not include the length of the Index field as defined in [I-D.ietf-grow-bmp-tlv].

In the example of Section 2.2.2, the sequence listed in the Previous VRF sequence would be [B, A].

3. Adj-RIB-In and Adj-RIB-Out Considerations

This document defines multiple BMPv4 TLVs for BMP Route Monitoring Messages for Local-RIB peers. While this document focuses on this particular use case, these TLVs can also be used for both Adj-RIB-In and Adj-RIB-Out.

In Adj-RIB-In, the "Rx Peer Address TLV" (Section 2.1), is more expressive than the current Peer Address field of the Per-Peer Header [RFC7854] by allowing more address types to be specified, namely IPv6 Local Link Addresses. The "Self-Originated" address type defined in Section 2.1.1 is not applicable for Adj-RIB-In, and it SHOULD NOT be used in this case.

In Adj-RIB-Out, the Peer Address field of the Per-Peer Header [RFC8671] tells the operator to which peer the path is being exported. In this case, the "Rx Peer Address TLV" (Section 2.1) allows the operator to also know from which peer this path was received, without needing to correlate Adj-RIB-Out with Adj-RIB-In.

The VRF-related TLVs defined by this document in Section 2.2 are applicable to all RIBs. This allows tracking the flow of paths between VRFs, even if only one type of RIB is being exported.

4. IANA Considerations

This document requests that IANA assigns the following new parameters to the "BMP Route Monitoring TLVs" registry

* Type = TBD1: Rx Peer-Address TLV type. TLV containing the Peer Address of the peer from which a monitored path has been received from.

* Type = TBD2: Origin VRF TLV type. TLV containing the name of the VRF in which a monitored has been first received.

* Type = TBD3: Previous VRF TLV type. TLV containing the name of the VRF from which a monitored path has been imported.

* Type = TBD4: Previous VRF Sequence TLV type. TLV containing the sequence of names of VRFs through which a monitored path has been imported.

This document also requests the definition of a "BMP Local-RIB Peer Address Types" registry seeded as follows:

* Type = 1: Self-Originated address type. Set to 1 if the route described by the BGP PDU enclosed in the BMP Route Monitoring Message was originated from the BMP station (router).

* Type = 2: IPv4 address type. Set to 2 if the following Peer Address contained in the Rx Peer-Address TLV is an IPv4 address.

* Type = 3: Global Link IPv6 address type. Set to 3 if the following Peer Address contained in the Rx Peer-Address TLV is a Global Link IPv6 address.

* Type = 4: IPv6 + Interface ID address type. Set to 4 if the following Peer Address contained in the Rx Peer-Address TLV is an IPv6 address followed by a numerical interface ID of variable size.

* Type = 5: IPv6 + Interface Name address type. Set to 5 if the following Peer Address contained in the Rx Peer-Address TLV is an IPv6 address followed by an interface name encoded as an UTF-8 string of variable size.

Address Type values 1 through 255 MUST be assigned using the "Standards Action" policy and value 0 is Reserved.

5. Security Considerations

This document does not introduce new security considerations.

6. Acknowledgements

We would like to thank Camilo Cardona, Jeff Haas, Mohamed Boucadair for their valuable input on this document.

7. References

7.1. Normative References

[I-D.ietf-grow-bmp-tlv]
Lucente, P. and Y. Gu, "BMP v4: TLV support for BMP Route Monitoring and Peer Down Messages", Work in Progress, Internet-Draft, draft-ietf-grow-bmp-tlv-15, 17 January 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-grow-bmp-tlv-15>.

[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>.

[RFC7911] Walton, D., Retana, A., Chen, E., and J. Scudder, "Advertisement of Multiple Paths in BGP", RFC 7911, DOI 10.17487/RFC7911, July 2016, <https://www.rfc-editor.org/info/rfc7911>.

[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>.

[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>.

7.2. Informative References

Authors' Addresses

Pierre Francois
INSA-Lyon
Villeurbanne
France
Email: pierre.francois@insa-lyon.fr

Maxence Younsi
INSA-Lyon
Villeurbanne
France
Email: maxence.younsi@insa-lyon.fr

Paolo Lucente
NTT
Siriusdreef 70-72
Hoofddorp, WT 2132
Netherlands
Email: paolo@ntt.net