Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 6, 2020

C. Cardona
P. Lucente
NTT
P. Francois
INSA-Lyon
Y. Gu
Huawei
T. Graf
Swisscom

January 03, 2020

BMP Extension for Path Marking TLV
draft-cppy-grow-bmp-path-marking-tlv-02

Abstract

The BGP Monitoring Protocol (BMP) provides an interface for obtaining BGP Path information. BGP Path Information is conveyed within BMP Route Monitoring (RM) messages. This document proposes an extension to BMP to convey the status of a BGP path after being processed by the BGP best-path selection algorithm. This extension makes use of the TLV mechanims described in draft-lucente-bmp-tlv [I-D.lucente-bmp-tlv].

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 July 6, 2020.

Table of Contents

1. Introduction

For a given prefix, multiple paths with different path status, e.g., the "best-path", "back-up path" and so on, may co-exist in the BGP RIB after being processed by the local policy and the BGP decision process. The path status information is currently not carried in the BGP Update Message RFC4271 [RFC4271] or in the BMP Update Message RFC7854 [RFC7854].

External systems can use the path status for various applications. The path status is commonly checked by operators when performing troubleshooting. Having such status stored in a centralized system can enable the development of tools facilitating this process.

Optimisation systems can include the path status in their process, and also use the status as a validation source (since it can compare the calculated state to the actual outcome of the network, such as primary and backup path). As a final example, path status information can complement other centralized sources of data, for example, flow collectors.

This document defines a so-called Path Marking TLV to convey the BGP path status information to the BMP server. The BMP Path Marking is defined to be prepended in the BMP Route Monitoring (RM) Message.

2. Path Marking TLV for the RM Message

As per RFC4271 [RFC4271], the BMP RM Message consists of the Common Header, Per-Peer Header, and the BGP Update PDU. According to draft-lucente-bmp-tlv [I-D.lucente-bmp-tlv] , optional trailing data in TLV format is allowed in the BMP RM Message to convey characteristics of transported NLRIs (i.e. to help stateless parsing) or vendor-specific data. Such TLV types are to be defined for each application.

To include the path status along with each BGP path, we define the Path Marking TLV, shown as follows.

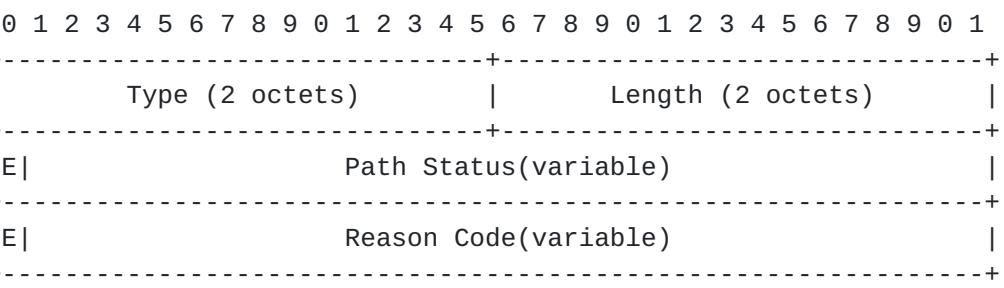

Figure 1: Path Marking TLV

- Type = TBD1 (2 Octets): indicates that it's the Path Marking TLV.
- Length (2 Octets): indicates the length of the value field of the Path Marking TLV. The value field further consists of the Path-Status field and Reason Code field.
- E bit (1 Bit) for Path Status: indicates if any enterprise-specific path status is used after the IANA-registered path status code.
- Path Status (4 Octets): indicates the path status of the BGP Update PDU encapsulated in the RM Message. Currently 7 types of path status are defined, as shown in Table 1.

- E bit (1 Bit) for Reason Code: indicates if any enterprise-specific reason code is used after the IANA-registered reason code.

- Reason Code (Variable): indicates the reasons/explanations of the path status indicated in the Path Type field. The detailed Reason Code field is defined in Section 2.2.

2.1. Path Status

The Path Status field contains a bit field where each bit encodes a specific role of the path. Multiple bits may be set when multiple path status apply to a path.

Two encoding options for Path Status are described in the following two sections.

2.1.1. IANA-registered Path Status Encoding

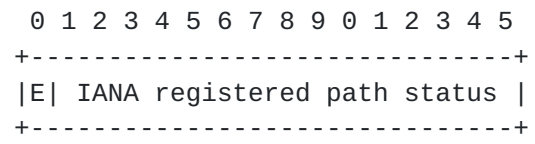

Figure 2: IANA-registered encoding of Path Status

- E bit (1 Bit): set to 0, indicating that only IANA-registered path status is used in this TLV.

- IANA-registered Path Status (2 octets): indicates the IANA-registered path status, as specified in Table 1.

Value	Path type
0x0000	Unknown
0x0001	Invalid
0x0002	Best
0x0004	Non-selected
0x0008	Primary
0x0010	Backup
0x0020	Non-installed
0x0040	Best external
0x0080	Add-Path

Table 1: Path Type

The best-path is defined in RFC4271 [RFC4271] and the best-external path is defined in draft-ietf-idr-best-external [I-D.ietf-idr-best-external].

An invalid path is a route that does not enter the BGP decision process.

A non-selected path is a route that is not selected in the BGP decision process. In other words, Best route and ECMP routes are not considered as non-selected.

A primary path is a recursive or non-recursive path whose nexthop resolution ends with an adjacency draft-ietf-rtgwg-bgp-pic [I-D.ietf-rtgwg-bgp-pic]. A prefix can have more than one primary path if multipath is configured draft-lapukhov-bgp-ecmp-considerations [I-D.lapukhov-bgp-ecmp-considerations]. A best-path is also considered as a primary path.

A backup path is also installed in the RIB, but it is not used until some or all primary paths become unreachable. Backup paths are used for fast convergence in the event of failures.

A non-installed path refers to the route that is not installed into the IP routing table.

For the advertisement of multiple paths for the same address prefix without the new paths implicitly replacing any previous ones, the add-path status is applied RFC7911 [RFC7911].

2.1.2. Enterprise-specific Path Status Encoding

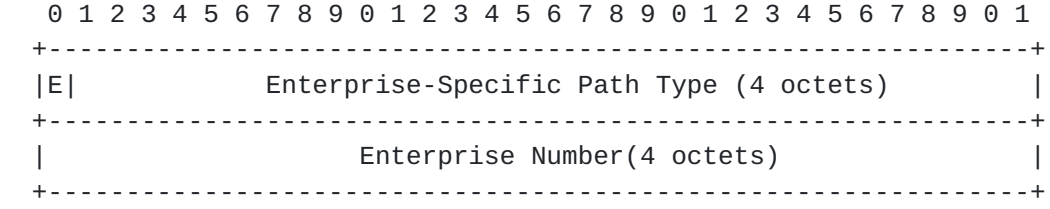

Figure 3: Enterprise-specific encoding of Path Status

- E bit (1 Bit): set to 1, indicating enterprise-specific path status is used in this TLV.
- Enterprise-specific Path Type (4 octets): indicates enterprise-specific path status, which remains to be defined.
- Enterprise Number (4 octets): indicates the IANA enterprise number IANA-PEN.

2.2. Reason Code

The Reason Code field contains a bit field where each bit encodes a specific reason. Multiple bits may be set when multiple reasons apply to a path.

2.2.1. IANA-registered Reason Code Encoding

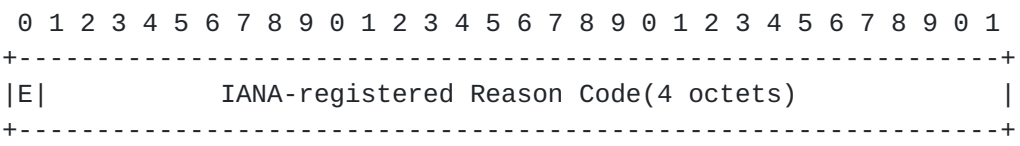

Figure 4: IANA-registered encoding of Reason Code

- E bit (1 Bit): set to 0, indicating that only IANA-registered reason code is used in this TLV. With the E bit set to 0, the Length field of the Path Marking TLV SHOULD be set to 8.

- IANA-registered Reason Code (4 octets): indicates the IANA-registered reason code of the path status.

2.2.2. Enterprise-specific Reason Code Encoding

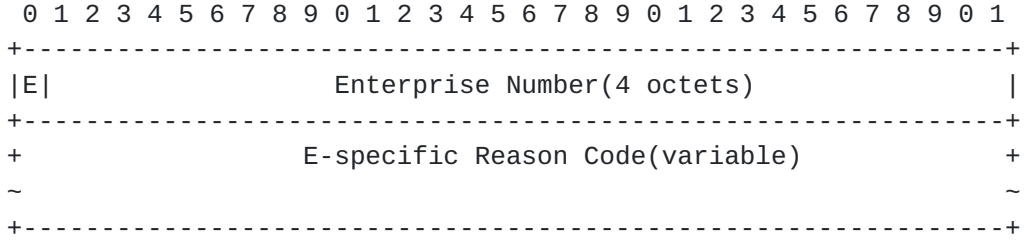

Figure 5: Enterprise-specific encoding of Reason Code

- E bit (1 Bit): set to 1, indicating enterprise-specific reason code is also used in this TLV.

- IANA-registered Reason Code (4 octets): indicates the IANA-registered reason code of the path status.

- Enterprise Number (4 octets): indicates the IANA enterprise number IANA-PEN.

- E-specific Reason Code (Variable): indicates enterprise-specific reason code of the path status.

3. Acknowledgements

We would like to thank Jeff Haas for his valuable comments.

4. IANA Considerations

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

4.1. Path Marking TLV

This document defines the Path Marking TLV with Type = TBD1: Path Marking (Section 2).

4.2. Path Marking TLV Reason Code

5. Security Considerations

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

6. Normative References

[I-D.ietf-idr-best-external]
Marques, P., Fernando, R., Chen, E., Mohapatra, P., and H. Gredler, "Advertisement of the best external route in BGP", draft-ietf-idr-best-external-05 (work in progress), January 2012.

[I-D.ietf-rtgwg-bgp-pic]
Bashandy, A., Filsfils, C., and P. Mohapatra, "BGP Prefix Independent Convergence", draft-ietf-rtgwg-bgp-pic-10 (work in progress), October 2019.

[I-D.lapukhov-bgp-ecmp-considerations]
Lapukhov, P. and J. Tantsura, "Equal-Cost Multipath Considerations for BGP", draft-lapukhov-bgp-ecmp-considerations-03 (work in progress), November 2019.

[I-D.lucente-bmp-tlv]
Lucente, P., Gu, Y., and H. Smit, "TLV support for BMP Route Monitoring and Peer Down Messages", draft-lucente-bmp-tlv-00 (work in progress), July 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>.

[RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <https://www.rfc-editor.org/info/rfc4271>.

[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>.

Authors' Addresses

Camilo Cardona
NTT
164-168, Carrer de Numancia
Barcelona 08029
Spain

Email: camilo@ntt.net

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

Email: paolo@ntt.net

Pierre Francois
INSA-Lyon
Lyon
France

Email: Pierre.Francois@insa-lyon.fr

Yunan Gu
Huawei
Huawei Bld., No.156 Beiqing Rd.
Beijing 100095
China

Email: guyunan@huawei.com

Thomas Graf
Swisscom
Binzring 17
Zurich 8045
Switzerland

Email: thomas.graf@swisscom.com