Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 13, 2021

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

July 12, 2020

BMP Extension for Path Status sub-TLV
draft-cppy-grow-bmp-path-marking-tlv-04

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 before and after being processed by the BGP best-path selection algorithm. This extension makes use of the Path Information TLV defined in draft-cppy-grow-bmp-path-info-tlv [I-D.cppy-grow-bmp-path-info-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 January 13, 2021.

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 that facilitate 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 Status sub-TLV to convey the BGP path status to the BMP server. The BMP Path Status sub-TLV is encapsulated within the BMP Path Information TLV carried in the BMP Route Monitoring (RM) Message draft-cppy-grow-bmp-path-info-tlv [I-D.cppy-grow-bmp-path-info-tlv].

2. Path Status sub-TLV

As stated in draft-cppy-grow-bmp-path-info-tlv [I-D.cppy-grow-bmp-path-info-tlv], the order of the sub-TLVs MUST be in accordance with the prefix order encapsulated in the Update PDU. This document defines two types of Path Status sub-TLVs: one is IANA-registered Path Status sub-TLV, and the other is Enterprise-specific Path Status sub-TLV.

2.1. IANA-registered Path Status sub-TLV

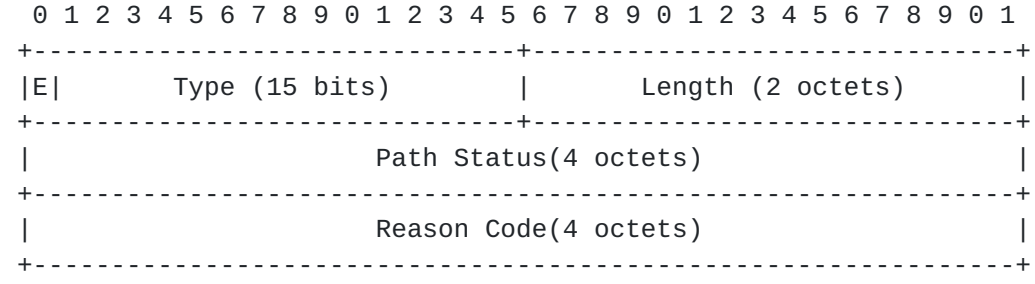

Figure 2:IANA-Registered Encoding of the path status sub-TLV

- E bit: For an IANA-registered sub-TLV, the E bit MUST be set to 0.
- Type = TBD2 (15 Bits): indicates that it is the IANA-registered Path Status sub-TLV.
- Length (2 Octets): indicates the length of the value field of the Path Status TLV. The value field further consists of the Path-Status field and Reason Code field.
- Path Status (4 Octets): indicates the path status of the BGP Update PDU encapsulated in the RM Message. Currently 9 types of path status are defined, as shown in Table 1.
- Reason Code (4 Octets): indicates the reasons/explanations of the path status indicated in the Path Type field. The currently defined reason codes are shown in Table 2.

Value	Path type
0x00000000	Unknown
0x00000001	Invalid
0x00000002	Best
0x00000004	Non-selected
0x00000008	Primary
0x00000010	Backup
0x00000020	Non-installed
0x00000040	Best-external
0x00000080	Add-Path

Table 1: IANA-Registered Path Type

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

- 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. Back-up routes are considered non-selected, while the best 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].

value	reason code
[0x00000000]	invalid for unknown,
[0x00000001]	invalid for super network,
[0x00000002]	invalid for dampening,
[0x00000003]	invalid for history,
[0x00000004]	invalid for policy deny,
[0x00000005]	invalid for ROA not validation,
[0x00000006]	invalid for interface error,
[0x00000007]	invalid for nexthop route unreachable,
[0x00000008]	invalid for nexthop tunnel unreachable,
[0x0000000f]	invalid for nexthop restrain,
[0x00000010]	invalid for relay BGP LSP,
[0x00000014]	invalid for being inactive within VPN instance
[0x00000015]	invalid for prefix-sid not exist,
[0x00000200]	not preferred for peer address,
[0x00000300]	not preferred for router ID,
[0x00000400]	not preferred for Cluster List,
[0x00000500]	not preferred for IGP cost,
[0x00000600]	not preferred for peer type,
[0x00000700]	not preferred for MED,
[0x00000800]	not preferred for origin,
[0x00000900]	not preferred for AS-Path,
[0x00000a00]	not preferred for route type,
[0x00000b00]	not preferred for Local_Pref,
[0x00000c00]	not preferred for PreVal,
[0x00000f00]	not preferred for not direct route,
[0x00001000]	not preferred for nexthop bit error,
[0x00001100]	not preferred for received path-id,
[0x00001200]	not preferred for validation,
[0x00001300]	not preferred for originate IP,
[0x00001500]	not preferred for route distinguisher,
[0x00001600]	not preferred for route-select delay,
[0x00001700]	not preferred for being imported route,
[0x00001800]	not preferred for med-plus-igp,
[0x00001c00]	not preferred for AIGP,
[0x00001d00]	not preferred for nexthop-resolved aigp,
[0x00002000]	not preferred for nexthop unreachable,
[0x00002100]	not preferred for nexthop IP,
[0x00002300]	not preferred for high-priority,
[0x00002400]	not preferred for nexthop-priority,
[0x00002500]	not preferred for process ID,
[0xFFFFFFFF]	no reason code

Table 2: IANA-Registered Reason Code

2.2. Enterprise-specific Path Status sub-TLV

The E-bit [I-D.lucente-grow-bmp-tlv-ebit] mechanism allows the usage of vendor-specific TLVs in addition to IANA-registered one. In this document, both encoding options for the Path Status sub-TLV are described.

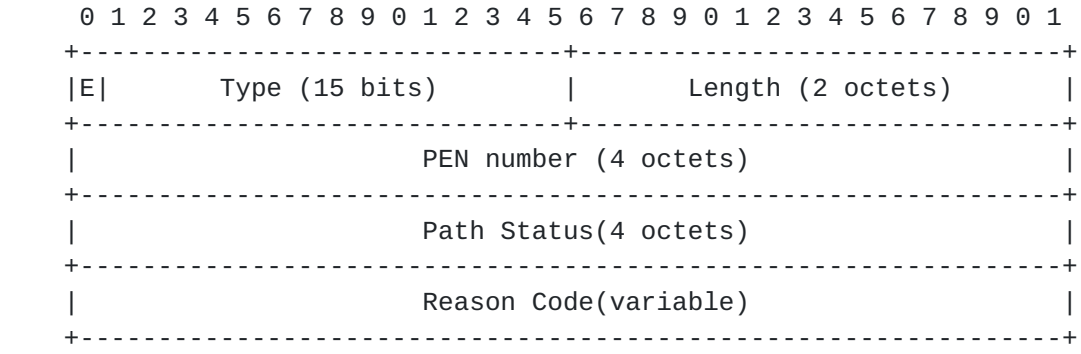

Figure 3: Enterprise-specific encoding of Path Status sub-TLV

- E bit: For an Enterprise-specific sub-TLV, the E bit MUST be set to 1.
- Type = 1 (15 Bits): indicates that it's the Enterprise-specific Path Status sub-TLV.
- Length (2 Octets): indicates the length of the value field of the Path Status sub-TLV. The value field further consists of the Path-Status field and Reason Code field.
- PEN Number (4 octets): indicates the IANA enterprise number IANA-PEN.
- Path Status (4 Octets): indicates the enterprise-specific path status.
- Reason Code (Variable): indicates the reasons/explanations of the path status indicated in the Path Type field.

3. Acknowledgments

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.

Type = TBD1 (15 Bits): indicates that it is the IANA-registered Path Status sub-TLV.

5. Security Considerations

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

6. Normative References

[I-D.cppy-grow-bmp-path-info-tlv]
Cardona, C., Lucente, P., Francois, P., Gu, Y., and T. Graf, "BMP Extension for Path Information TLV", draft-cppy-grow-bmp-path-info-tlv-00 (work in progress), July 2020.

[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-02 (work in progress), March 2020.

[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-11 (work in progress), February 2020.

[I-D.lapukhov-bgp-ecmp-considerations]
Lapukhov, P. and J. Tantsura, "Equal-Cost Multipath Considerations for BGP", draft-lapukhov-bgp-ecmp-considerations-04 (work in progress), May 2020.

[I-D.lucente-grow-bmp-tlv-ebit]
Lucente, P. and Y. Gu, "Support for Enterprise-specific TLVs in the BGP Monitoring Protocol", draft-lucente-grow-bmp-tlv-ebit-01 (work in progress), May 2020.

[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