

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 10, 2015

B. Decraene
Orange
March 9, 2015

BGP Next-Hop Capabilities
draft-decraene-idr-next-hop-capability-00

Abstract

RFC 5492 defines capabilities advertisement for the BGP peer. In addition, it's useful to know the capabilities of the BGP Next-Hop, in particular for forwarding plane features. RFC 5492 is not applicable because the BGP peer may be different from the BGP Next-Hop, in particular when BGP Route Reflection is used. This document defines a mechanism to advertise such BGP Next Hop Capabilities.

This document defines a new BGP non-transitive attribute to carry Next-Hop Capabilities. This attribute is deleted when the BGP Next Hop is changed.

This document also defines a Next-Hop capability to advertise the ability to handle the Entropy Label defined in RFC 6790.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

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 <http://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 September 10, 2015.

Copyright Notice

Copyright (c) 2015 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 (<http://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.

1. Introduction

[RFC5492] defines capabilities advertisement for the BGP peer. It's also useful to know the capabilities of the BGP Next-Hop, in particular for forwarding plane features. RFC 5492 is not applicable because the BGP peer may be different from the BGP Next-Hop, in particular when BGP Route Reflection is used. This document defines a mechanism to advertise such BGP Next Hop Capabilities.

This document defines a new BGP non-transitive attribute to carry Next-Hop Capabilities. This attribute is deleted when the BGP Next Hop is changed.

This document also defines a first application to advertise the capability to handle the Entropy Label defined in [RFC6790]. Note that RFC 6790 had originally defined a BGP attribute for this but it has been latter deprecated in [RFC7447]

2. BGP Next-Hop Capability Attribute

The BGP Next-Hop Capabilities Attribute is an optional, non-transitive BGP Attribute, of value TBD1. The attribute consists of a set of Next-Hop Capabilities. Inclusion of a Next-Hop Capability "X" indicates that the BGP Next-Hop, encoded in either the NEXT_HOP attribute defined in [RFC4271] or the Network Address of Next Hop field of the MP_REACH_NLRI attribute defined in [RFC4760], supports the capability "X". This document do not make distinction between these two Next-Hop fields and refer to them as BGP Next-Hop.

A Next-Hop Capability is triple (Capability Code, Capability Length, Capability Value) aka a TLV:

A Next-Hop Capability.

```

+-----+
| Capability Code (1 octet) |
+-----+
| Capability Length (1 octet) |
+-----+
| Capability Value (variable) |
~                               ~
+-----+

```

Capability Code : a one-octet unsigned binary integer which indicates the type of "Next-Hop Capability" advertised and unambiguously identifies an individual capability.

Capability Field: a one-octet unsigned binary integer which indicates the length, in octets, of the Value Field. A length of 0 indicates that no Value Field is present.

Value Field: a variable-length field from 0 to 255 octets. It is interpreted according to the value of the Capability Code field.

BGP speakers SHOULD NOT include more than one instance of a Next-Hop capability with the same Capability Code, Capability Length, and Capability Value. Note, however, that processing of multiple instances of such capability does not require special handling, as additional instances do not change the meaning of the announced capability; thus, a BGP speaker MUST be prepared to accept such multiple instances.

BGP speakers MAY include more than one instance of a capability (as identified by the Capability Code) with non-zero Capability Length field, but with different Capability Value and either the same or different Capability Length. Processing of these capability instances is specific to the Capability Code and MUST be described in the document introducing the new capability.

3. BGP Next-Hop Capabilities Attribute Operation

The BGP Attribute being non-transitive, as per [RFC4271], a BGP speaker which does not understand it will quietly ignore it and not pass it along to other BGP peers.

A BGP speaker which understand the BGP Next-Hop Attribute and does not change the BGP Next-Hop, SHOULD NOT change the BGP Next-Hop Attribute and SHOULD pass the attribute unchanged along to other BGP peers.

A BGP speaker which understand the BGP Next-Hop Attribute and change the BGP Next-Hop, MUST remove the received BGP Next-Hop Attribute before propagating the BGP UPDATE other BGP peers. It MAY attach a new BGP Next-Hop attribute describing the capabilities of the new BGP Next-Hop.

4. BGP Next-Hop Attribute Error Handling

A BGP Next-Hop Capability Attribute is considered malformed if the length of the Attribute is not equal to the sum of all (BGP Hop Capability Length +2) of each capability carrier in this attribute. Note that "2" is the length of the fields "Type" and "Length" of each BGP Next Hop Capability.

A BGP UPDATE message with a malformed BGP Next-Hop Capability Attribute SHALL be handled using the approach of "attribute discard" defined in [I-D.ietf-idr-error-handling]. [Note: To be Discussed. Treat as withdraw would be safer if one implementation allow changing route preference based on BGP Next-Hop Capability. But this is the case of any attribute.]

If a Next-Hop Capability is malformed, this Next-Hop Capability Type MUST be ignored. Others Next-Hop Capabilities MUST be processed as usual.[Note: To be Discussed]

5. Entropy Label Next-Hop Capability

The Entropy Label Next-Hop Capability have type code 1 and a length of 0 octet.

The inclusion of the "Entropy Label" Next-Hop Capability indicates that the BGP Next-Hop can be sent packets with a MPLS entropy label as an egress LSR for all routes in that NLRI.

A BGP speaker S that originates an UPDATE MAY include the Entropy Label Next-Hop Capability only if either of the following is true:

A1: S sets the BGP NEXT_HOP attribute to itself AND S can process entropy labels. In other words, the BGP Next-Hop can process entropy labels.

A2: S sets the BGP NEXT_HOP attribute to itself AND S swaps the advertised label without popping the advertised label(s) stack AND S knows that the egress can process the entropy label (typically when redistributing a route received with the indication that the egress can be sent the entropy label (e.g. received via BGP with the "Entropy Label" Next-Hop Capability attached, or via LDP with the ELC

TLV...)). In other words, the BGP Next-Hop may or may not be able to process entropy labels, but it will not have to process it.

6. IANA Considerations

6.1. Next-Hop Capabilities Attribute

IANA is requested to allocate a new Path Attribute, called "Next-Hop Capabilities", type Code TBD1, from the "BGP Path Attributes" registry.

6.2. Next-Hop Capability registry

The IANA is requested to create and maintain a registry entitled "Next-Hop Capabilities".

The registration policies [RFC5226] for this registry are:

1-63	IETF Review
64-127	First Come First Served

IANA is requested to make the following initial assignments:

Registry Name: Next-Hop Capability.

Value	Meaning	Reference
0	Reserved	This document
1	Entropy Label	This document
2-127	Unassigned	
128-255	Private Use	This document

7. Security Considerations

This document does not introduce new security vulnerabilities in BGP. Please refer to the Security Considerations section of [RFC4271] for security mechanisms applicable to BGP.

8. Acknowledgement

The Entropy Label Next-Hop Capability defined in this document is based on the ELC BGP attribute defined in section 5.2 of [RFC6790].

9. References

9.1. Normative References

- [I-D.ietf-idr-error-handling]
Chen, E., Scudder, J., Mohapatra, P., and K. Patel,
"Revised Error Handling for BGP UPDATE Messages", draft-
ietf-idr-error-handling-18 (work in progress), December
2014.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway
Protocol 4 (BGP-4)", RFC 4271, January 2006.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter,
"Multiprotocol Extensions for BGP-4", RFC 4760, January
2007.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an
IANA Considerations Section in RFCs", BCP 26, RFC 5226,
May 2008.
- [RFC6790] Kompella, K., Drake, J., Amante, S., Henderickx, W., and
L. Yong, "The Use of Entropy Labels in MPLS Forwarding",
RFC 6790, November 2012.

9.2. Informative References

- [RFC5492] Scudder, J. and R. Chandra, "Capabilities Advertisement
with BGP-4", RFC 5492, February 2009.
- [RFC7447] Scudder, J. and K. Kompella, "Deprecation of BGP Entropy
Label Capability Attribute", RFC 7447, February 2015.

Author's Address

Bruno Decraene
Orange

Email: bruno.decraene@orange.com