

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: November 19, 2017

K. Patel
Arrcus, Inc.
J. Uttaro
ATT
B. Decraene
Orange
W. Henderickx
Alcatel Lucent
J. Haas
Juniper Networks
May 18, 2017

Constrain Attribute announcement within BGP
draft-ietf-idr-bgp-attribute-announcement-01.txt

Abstract

[RFC4271] defines four different categories of BGP Path attributes. The different Path attribute categories can be identified by the attribute flag values. These flags help identify if an attribute is optional or well-known, Transitive or non-Transitive, Partial, or of an Extended length type. BGP attribute announcement depends on whether an attribute is a well-known or optional, and whether an attribute is a transitive or non-transitive. BGP implementations MUST recognize all well-known attributes. The well-known attributes are always Transitive. It is not required for BGP implementations to recognise all the Optional attributes. The Optional attributes could be Transitive or Non-Transitive. BGP implementations MUST store and forward any Unknown Optional Transitive attributes and ignore and drop any Unknown Optional Non-Transitive attributes.

Currently, there is no way to confine the scope of Path attributes within a given Autonomous System (AS) or a given BGP member-AS in Confederation. This draft defines attribute extensions that help confine the scope of Optional attributes within a given AS or a given BGP member-AS in Confederation

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-Draft Constrain Attribute announcement within BGP

May 2017

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 November 19, 2017.

Copyright Notice

Copyright (c) 2017 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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction	3
1.1.	Requirements Language	4
2.	Path Attribute Format	4
3.	Extended Path Attribute Flags	4
4.	Operation	6
5.	IANA Considerations	7
6.	Security Considerations	7

6.1. Acknowledgements	7
7. References	8
7.1. Normative References	8
7.2. Information References	8
Authors' Addresses	8

[1.](#) Introduction

[RFC4271] defines four different categories of BGP Path attributes. The different Path attribute categories can be identified by the attribute flag values. These flags help identify if an attribute is optional or well-known, Transitive or non-Transitive, Partial, or of an Extended length type. BGP attribute announcement depends on whether an attribute is a well-known or optional, and whether an attribute is a transitive or non-transitive. BGP implementations MUST recognize all well-known attributes. The well-known attributes are always Transitive. It is not required for BGP implementations to recognise all the Optional attributes. The Optional attributes could be Transitive or Non-Transitive. BGP implementations MUST store and forward any Unknown Optional Transitive attributes and ignore and drop any Unknown Optional Non-Transitive attributes.

Optional Transitive attributes help foster partial deployments of newer BGP features. Alternatively, Optional Non-Transitive attributes are drop by BGP speakers that do not recognise the attribute. The optional attributes in their current definition do not provide any automated attribute level filtering to control the scope of announcements within a given AS or a BGP member-AS in Confederation. Scoped announcements of attributes may be needed in certain scenarios. Announcing attributes beyond their intended scope MAY result in breakage of functionalities or leaking of any undesired information.

This draft defines new attribute extensions that help confine the scope of Path attributes; in particular Optional attributes within a given Autonomous System or a given BGP member-AS in confederation or a given Administrative domain. Note that "BGP Member-AS in Confederation" and "Member-AS" are used entirely interchangeably throughout this document.

As part of new attribute extensions, this draft defines a new attribute format to incorporate the scoping information. The new

attribute format applies to all the new attribute types that will be defined moving forward. The newly defined attribute scoping is specifically for newer attributes that explicitly state their use of such scoping bits. These newly defined attributes would be either an Optional transitive attributes (recognized and unrecognized) or any recognized optional non-transitive attributes. For any well-known attributes or unrecognized optional non-transitive attributes, the standard rules mentioned in [\[RFC4271\]](#) applies.

[1.1.](#) 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](#)].

[2.](#) Path Attribute Format

[RFC4271] defines path attribute format as a triple [attribute type, attribute length, attribute value] of a variable length. The attribute value field is of a variable length. This draft augments the path attribute value field and reserves first four bytes of path attribute value field as path attribute extended flags field. All the path attributes carrying extended flags field will have a minimum attribute length of 4 bytes. The augmented path attribute format applies to all the current undefined attributes types (30-39, 41-127, 129-254). Any attribute specific data follows the path attribute extended flags field.

[3.](#) Extended Path Attribute Flags

[RFC4271] defines four type of BGP Path attributes using the attribute Flags field as follows:

Path Attribute flags:

```

  0 1 2 3 4 5 6 7
+---+---+---+---+
|O|T|P|E|   R   |   (R = MUST Be Zero)
```


+--+--+--+--+--+--+--+

- O Optional or a Well-known as defined in [[RFC4271](#)]
- T Transitive or Non-Transitive as defined in [[RFC4271](#)]
- P Partial as defined in [[RFC4271](#)]
- E Extended Length type as defined in [[RFC4271](#)]

This draft introduces new Flags field known as Extended Path Attribute Flags. The Extended Path Attribute Flags field is defined as first 4 bytes of the Path Attribute's value field. This draft introduces three new Extended Path Attributes flags as follows:

Internet-Draft Constrain Attribute announcement within BGP May 2017

Path Attribute flags:



(R = MUST Be Zero)

- A AS Wide Scope
- C Member-AS in Confederation Scope
- M Multi-AS Scope

The first least significant bit ("A") is defined as the AS Wide Scope bit, which is used to indicate that an optional attribute cannot be announced outside a given AS boundary. When set, the given optional

attribute MUST be filtered by the sending BGP speaker at an AS boundary. If the "A" bit is set then the "O" bit defined in BGP Path Attribute Flag field MUST be set. Otherwise a BGP speaker MUST consider an attribute as an error and malformed.

The second least significant bit ("C") is defined as the Member-AS Scope bit, which is used to indicate that an optional attribute cannot be announced outside a given Member-AS boundary. When set, the given optional attribute MUST be filtered by the sending BGP speaker at a Member-AS boundary. If the "C" bit is set then the "O" bit defined in BGP Path Attribute Flag field MUST be set. Otherwise a BGP speaker MUST consider an attribute as an error and malformed. "C" bit SHOULD only be set when an Autonomous System is configured as a BGP Confederation. A BGP speaker MUST not transmit an attribute with "C" bit set to peers that are not members of the local confederation. Otherwise a BGP speaker MUST consider an attribute as an error and malformed.

Both the first and the second most least significant bit together is defined as the Multiple AS Scope within a Single Administration. When both the first and the second bits are set, optional attribute can be traversed across multiple AS and filtered by the sending BGP speaker at the Administration boundary.

The handling of malformed attributes SHOULD follow the procedures mentioned in [[RFC7606](#)]. For any malformed attribute that is handled

by the "attribute discard" instead of the "treat-as-withdraw" approach, it is critical to consider the potential impact. In particular, if the attribute has an impact on the route selection or installation process, then the presumption is that "attribute discard" is unsafe and "treat-as-withdraw" procedure SHOULD be considered. Otherwise, "attribute discard" procedure SHOULD be used.

[4.](#) Operation

When originating a well-known Path attribute, a BGP speaker MUST set both the AS Wide Scope and Member-AS Scope bit to 0. When originating an optional Path attribute, a BGP speaker SHOULD use and set AS Wide Scope bit if it wants to restrict the announcement within a AS. Similarly, when originating an optional Path attribute, a BGP speaker SHOULD use and set Member-AS Scope bit if it wants to

restrict the announcement with a Member-AS. When originating an optional Path attribute, a BGP speaker SHOULD use and set both Member-AS Scope bit and AS Wide Scope bit if it wants to restrict the announcement within a single administration composed of multiple ASes.

When a BGP speaker receives or originates a route that includes any well-known Path attribute with either a AS Wide Scope bit set or a Member-AS Scope bit set then it SHOULD consider the attribute as malformed. The handling of malformed attributes SHOULD follow the procedures mentioned in [[RFC7606](#)].

When a BGP speaker receives or originates a route that includes an optional Path attribute with a AS Wide Scope bit set and a Member-AS Scope bit cleared, it MUST remove that Path attribute when announcing the route to any of its EBGp speakers. To deal with partial deployments it is suggested that a BGP speaker SHOULD quietly ignore and not pass along to other BGP peers any Path attribute received from its EBGp peers with a AS Wide Scope bit set and a Member-AS Scope bit cleared unless configured explicitly using a policy.

When a BGP speaker receives or originates a route that includes an optional Path attribute with a Member-AS Scope bit set and a AS Wide Scope bit cleared, it MUST remove that Path attribute when announcing the route to any of its BGP speakers outside its Member-AS. To deal with partial deployments it is suggested that a BGP speaker SHOULD quietly ignore and not pass along to other BGP peers any Path attribute received from its BGP peers with a Member-AS Scope bit set and a AS Wide Scope bit cleared unless configured explicitly as a policy.

When a BGP speaker receives or originates a route with an optional path attribute that has both, the AS Wide Scope bit set and the

Member-AS Scope bit set, it MUST announce it to all its EBGp peers within its administrative domain. Such an attribute MUST be filtered when the attribute is announced outside its administrative domain. The BGP peering boundaries for an administrative domain is a matter of a policy and is set by the operators.

Any implementation that supports the extensions defined in this draft MUST support the Enhanced Error handling defined in [[RFC7606](#)].

Enhanced Error handling allows any error condition that MAY occur during the parsing and processing of new attribute flags to be treated according to the procedures of [RFC7606]. Furthermore, it is assumed that the BGP network is enabled with Enhanced Error Handling feature. This allows BGP speakers not implementing the draft extensions to apply the procedures of [RFC7606].

5. IANA Considerations

This draft define a new path attribute format for all undefined attribute types. We request IANA to record the use of new path attribute format for the following undefined attribute types (30-39, 41-127, 129-254).

This draft defines two new Extended Path attribute flags. We request IANA to create a new registry for BGP Extended Path Attribute Flags under BGP Path attributes as follows:

Under "Border Gateway Protocol (BGP) Parameters" registry, "BGP Extended Path Attributes Flags" Reference: [draft-keyupate-idr-bgp-attribute-annoucement-01](#) Registration Procedures as follows:

Bit Value (LSB)	Type	Reference
1	AS Wide Scope	Current Draft
2	Member-AS in Confederation	Current Draft

6. Security Considerations

This extension to BGP does not change the underlying security issues inherent in the existing [RFC4724] and [RFC4271].

6.1. Acknowledgements

The authors would like to thank John Scudder, Jakob Heitz, Shyam Seturam, Juan Alcaide and Acee Lindem for the review and comments.

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://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, <<http://www.rfc-editor.org/info/rfc4271>>.
- [RFC7606] Chen, E., Ed., Scudder, J., Ed., Mohapatra, P., and K. Patel, "Revised Error Handling for BGP UPDATE Messages", [RFC 7606](#), DOI 10.17487/RFC7606, August 2015, <<http://www.rfc-editor.org/info/rfc7606>>.

7.2. Information References

- [RFC3392] Chandra, R. and J. Scudder, "Capabilities Advertisement with BGP-4", [RFC 3392](#), DOI 10.17487/RFC3392, November 2002, <<http://www.rfc-editor.org/info/rfc3392>>.
- [RFC4486] Chen, E. and V. Gillet, "Subcodes for BGP Cease Notification Message", [RFC 4486](#), DOI 10.17487/RFC4486, April 2006, <<http://www.rfc-editor.org/info/rfc4486>>.
- [RFC4724] Sangli, S., Chen, E., Fernando, R., Scudder, J., and Y. Rekhter, "Graceful Restart Mechanism for BGP", [RFC 4724](#), DOI 10.17487/RFC4724, January 2007, <<http://www.rfc-editor.org/info/rfc4724>>.

Authors' Addresses

Keyur Patel
Arrcus, Inc.

Email: keyur@arrcus.com

James Uttaro
ATT
200 S. Laurel Ave
Middletown, NJ 07748
USA

Email: uttaro@att.com

Bruno Decraene
Orange

Email: bruno.decraene@orange.com

Wim Henderickx
Alcatel Lucent
Copernicuslaan 50
Antwerp 2018
Belgium

Email: wim.henderickx@alcatel-lucent.com

Jeff Haas
Juniper Networks
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
USA

Email: jhaas@juniper.net

