

Internet Engineering Task Force (IETF)
Request for Comments: 8092
Category: Standards Track
ISSN: 2070-1721

J. Heitz, Ed.
Cisco
J. Snijders, Ed.
NTT
K. Patel
Arrcus
I. Bagdonas
Equinix
N. Hilliard
INEX
February 2017

BGP Large Communities Attribute

Abstract

This document describes the BGP Large Communities attribute, an extension to BGP-4. This attribute provides a mechanism to signal opaque information within separate namespaces to aid in routing management. The attribute is suitable for use with all Autonomous System Numbers (ASNs) including four-octet ASNs.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc8092>.

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.

Table of Contents

1. Introduction	2
2. Requirements Language	3
3. BGP Large Communities Attribute	3
4. Aggregation	4
5. Canonical Representation	4
6. Error Handling	5
7. Security Considerations	5
8. IANA Considerations	6
9. References	6
9.1. Normative References	6
9.2. Informative References	6
Acknowledgments	7
Contributors	7
Authors' Addresses	8

1. Introduction

BGP [RFC4271] implementations typically support a routing policy language to control the distribution of routing information. Network operators attach BGP communities to routes to associate particular properties with these routes. These properties may include information such as the route origin location, or specification of a routing policy action to be taken, or one that has been taken, and is applied to all routes contained in a BGP Update Message where the Communities Attribute is included. Because BGP communities are optional transitive BGP attributes, BGP communities may be acted upon or otherwise used by routing policies in other Autonomous Systems (ASes) on the Internet.

A BGP Communities attribute is a variable-length attribute consisting of a set of one or more four-octet values, each of which specify a community [RFC1997]. Common use of the individual values of this attribute type split this single 32-bit value into two 16-bit values. The most significant word is interpreted as an Autonomous System Number (ASN), and the least significant word is a locally defined value whose meaning is assigned by the operator of the AS in the most significant word.

Since the adoption of four-octet ASNs [RFC6793], the BGP Communities attribute can no longer accommodate the above encoding, as a two-octet word cannot fit a four-octet ASN. The BGP Extended Communities attribute [RFC4360] is also unsuitable. The six-octet length of the Extended Community value precludes the common operational practice of encoding four-octet ASNs in both the Global Administrator and the Local Administrator sub-fields.

To address these shortcomings, this document defines a BGP Large Communities attribute encoded as an unordered set of one or more twelve-octet values, each consisting of a four-octet Global Administrator field and two four-octet operator-defined fields, each of which can be used to denote properties or actions significant to the operator of the AS assigning the values.

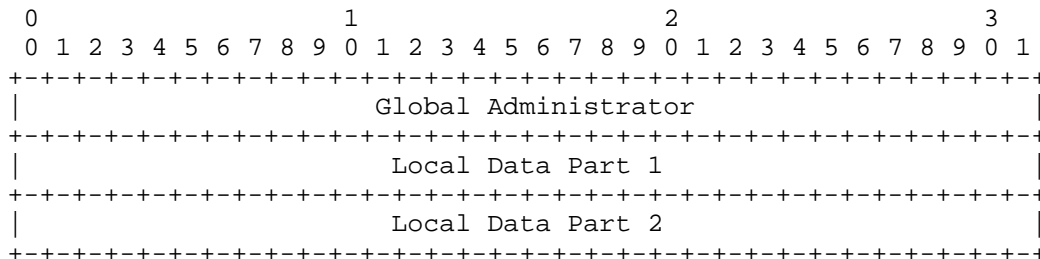
2. Requirements Language

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 [RFC2119].

3. BGP Large Communities Attribute

This document defines the BGP Large Communities attribute as an optional transitive path attribute of variable length. All routes with the BGP Large Communities attribute belong to the communities specified in the attribute.

Each BGP Large Community value is encoded as a 12-octet quantity, as follows:



Global Administrator: A four-octet namespace identifier.

Local Data Part 1: A four-octet operator-defined value.

Local Data Part 2: A four-octet operator-defined value.

The Global Administrator field is intended to allow different ASes to define BGP Large Communities without collision. This field SHOULD be an ASN, in which case the Local Data Parts are to be interpreted as defined by the owner of the ASN. The use of Reserved ASNs (0 [RFC7607], 65535 and 4294967295 [RFC7300]) is NOT RECOMMENDED.

There is no significance to the order in which twelve-octet Large Community Attribute values are encoded in a Large Communities attribute, A BGP speaker can transmit them in any order.

Duplicate BGP Large Community values MUST NOT be transmitted. A receiving speaker MUST silently remove redundant BGP Large Community values from a BGP Large Community attribute.

4. Aggregation

If a range of routes is aggregated, then the resulting aggregate should have a BGP Large Communities attribute that contains all of the BGP Large Communities attributes from all of the aggregated routes.

5. Canonical Representation

The canonical representation of BGP Large Communities is three separate unsigned integers in decimal notation in the following order: Global Administrator, Local Data 1, Local Data 2. Numbers MUST NOT contain leading zeros; a zero value MUST be represented with a single zero. Each number is separated from the next by a single colon. For example: 64496:4294967295:2, 64496:0:0.

BGP Large Communities SHOULD be represented in the canonical representation.

6. Error Handling

The error handling of BGP Large Communities is as follows:

- o A BGP Large Communities attribute SHALL be considered malformed if the length of the BGP Large Communities Attribute value, expressed in octets, is not a non-zero multiple of 12.
- o A BGP Large Communities attribute SHALL NOT be considered malformed due to presence of duplicate Large Community values.
- o A BGP UPDATE message with a malformed BGP Large Communities attribute SHALL be handled using the approach of "treat-as-withdraw" as described in Section 2 of [RFC7606].

The BGP Large Communities Global Administrator field may contain any value, and a BGP Large Communities attribute MUST NOT be considered malformed if the Global Administrator field contains an unallocated, unassigned, or reserved ASN.

7. Security Considerations

This document does not change any underlying security issues associated with any other BGP Communities mechanism. Specifically, an AS relying on the BGP Large Communities attribute carried in BGP must have trust in every other AS in the path, as any intermediate AS in the path may have added, deleted, or altered the BGP Large Communities attribute. Specifying the mechanism to provide such trust is beyond the scope of this document.

BGP Large Communities do not protect the integrity of each community value. Operators should be aware that it is possible for a BGP speaker to alter BGP Large Community Attribute values in a BGP Update Message. Protecting the integrity of the transitive handling of BGP Large Community attributes in a manner consistent with the intent of expressed BGP routing policies falls within the broader scope of securing BGP, and is not specifically addressed here.

Network administrators should note the recommendations in Section 11 of "BGP Operations and Security" [RFC7454].

8. IANA Considerations

IANA has assigned the value 32 (LARGE_COMMUNITY) in the "BGP Path Attributes" subregistry under the "Border Gateway Protocol (BGP) Parameters" registry.

9. References

9.1. 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>>.

9.2. Informative References

- [RFC1997] Chandra, R., Traina, P., and T. Li, "BGP Communities Attribute", RFC 1997, DOI 10.17487/RFC1997, August 1996, <<http://www.rfc-editor.org/info/rfc1997>>.
- [RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", RFC 4360, DOI 10.17487/RFC4360, February 2006, <<http://www.rfc-editor.org/info/rfc4360>>.
- [RFC6793] Vohra, Q. and E. Chen, "BGP Support for Four-Octet Autonomous System (AS) Number Space", RFC 6793, DOI 10.17487/RFC6793, December 2012, <<http://www.rfc-editor.org/info/rfc6793>>.
- [RFC7300] Haas, J. and J. Mitchell, "Reservation of Last Autonomous System (AS) Numbers", BCP 6, RFC 7300, DOI 10.17487/RFC7300, July 2014, <<http://www.rfc-editor.org/info/rfc7300>>.
- [RFC7454] Durand, J., Pepelnjak, I., and G. Doering, "BGP Operations and Security", BCP 194, RFC 7454, DOI 10.17487/RFC7454, February 2015, <<http://www.rfc-editor.org/info/rfc7454>>.

[RFC7607] Kumari, W., Bush, R., Schiller, H., and K. Patel,
"Codification of AS 0 Processing", RFC 7607,
DOI 10.17487/RFC7607, August 2015,
<<http://www.rfc-editor.org/info/rfc7607>>.

Acknowledgments

The authors would like to thank Ruediger Volk, Russ White, Acee Lindem, Shyam Sethuram, Jared Mauch, Joel M. Halpern, Jeffrey Haas, Gunter van de Velde, Marco Marzetti, Eduardo Ascenco Reis, Mark Schouten, Paul Hoogsteder, Martijn Schmidt, Greg Hankins, Bertrand Duvivier, Barry O'Donovan, Grzegorz Janoszka, Linda Dunbar, Marco Davids, Gaurab Raj Upadhaya, Jeff Tantsura, Teun Vink, Adam Davenport, Theodore Baschak, Pier Carlo Chiodi, Nabeel Cocker, Ian Dickinson, Jan Baggen, Duncan Lockwood, David Farmer, Randy Bush, Wim Henderickx, Stefan Plug, Kay Rechthien, Rob Shakir, Warren Kumari, Gert Doering, Thomas King, Mikael Abrahamsson, Wesley Steehouwer, Sander Steffann, Brad Dreisbach, Martin Millnert, Christopher Morrow, Jay Borkenhagen, Arnold Nipper, Joe Provo, Niels Bakker, Bill Fenner, Tom Daly, Ben Maddison, Alexander Azimov, Brian Dickson, Peter van Dijk, Julian Seifert, Tom Petch, Tom Scholl, Arjen Zonneveld, Remco van Mook, Adam Chappell, Jussi Peltola, Kristian Larsson, Markus Hauschild, Richard Steenbergen, David Freedman, Richard Hartmann, Geoff Huston, Mach Chen, and Alvaro Retana for their support, insightful review, and comments.

Contributors

The following people contributed significantly to the content of the document:

John Heasley
NTT Communications
Email: heas@shrubbery.net

Adam Simpson
Nokia
Email: adam.1.simpson@nokia.com

Authors' Addresses

Jakob Heitz (editor)
Cisco
170 West Tasman Drive
San Jose, CA 95054
United States of America

Email: jheitz@cisco.com

Job Snijders (editor)
NTT Communications
Theodorus Majofskistraat 100
Amsterdam 1065 SZ
The Netherlands

Email: job@ntt.net

Keyur Patel
Arrcus, Inc.

Email: keyur@arrcus.com

Ignas Bagdonas
Equinix
80 Cheapside
London EC2V 6EE
United Kingdom

Email: ibagdona.ietf@gmail.com

Nick Hilliard
INEX
4027 Kingswood Road
Dublin 24
Ireland

Email: nick@inex.ie