

IDR
Internet-Draft
Intended status: Standards Track
Expires: April 11, 2017

J. Heitz
Cisco
K. Patel
Arrcus
J. Snijders
NTT
I. Bagdonas
Equinix
A. Simpson
Nokia
October 8, 2016

Large BGP Communities
draft-ietf-idr-large-community-02

Abstract

This document describes the Large BGP Community 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 in 4-octet ASNs.

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](#)].

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 April 11, 2017.

Internet-Draft

Large BGP Communities

October 2016

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) 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.	Large BGP Communities Attribute	3
3.	Aggregation	4
4.	Textual Representation	4
5.	Reserved Large BGP Community values	4
6.	Error Handling	5
7.	Security Considerations	5
8.	Implementation status - RFC EDITOR: REMOVE BEFORE PUBLICATION	5
9.	IANA Considerations	6
10.	Acknowledgments	6
11.	References	7
11.1.	Normative References	7
11.2.	Informative References	7
11.3.	URIs	7
	Authors' Addresses	7

[1.](#) Introduction

BGP implementations typically support a routing policy language to control the distribution of routing information. Network operators attach BGP communities to routes to identify intrinsic properties of 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 may apply to an individual route or to a group of routes. Because BGP communities are optional transitive BGP attributes, BGP communities may be acted upon or

Global Administrator: A four-octet namespace identifier. This SHOULD be an Autonomous System Number.

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 Autonomous Systems to define Large Communities without collision. Implementations MUST allow the operator to specify any value for the Global Administrator field.

There is no significance to the order in which Large Communities are encoded in a path attributes field and a receiving speaker MAY retransmit them in an order different from which it received them.

Duplicate Large Communities SHOULD NOT be transmitted. A receiving speaker SHOULD silently remove duplicate Large Communities from a BGP UPDATE message.

3. Aggregation

If a range of routes is aggregated and the resulting aggregates attribute section does not carry the ATOMIC_AGGREGATE attribute, then the resulting aggregate should have a Large Communities path attribute which contains all of the large communities from all of the aggregated routes.

4. Textual Representation

BGP Communities [[RFC1998](#)] are usually represented in routing policy languages as two individual two-octet unsigned integers separated by a colon; for example, 64496:12345.

BGP Large Communities implementations MUST represent Large Communities in a manner similar to their representation of BGP Communities [[RFC1998](#)]. Large Communities MUST be represented as three separate four-octet unsigned integers in decimal format with no leading zeros. These integers MUST NOT be omitted, even when zero. For example, 64496:4294967295:2 or 64496:0:0.

Vendors MAY provide other textual representations. For example, a vendor's routing policy language may use a separator other than a colon or may require keywords or characters prepending or postpending the Large Communities attribute. Such differences are permitted. However, each implementation MUST make a representation available that depicts the integers in decimal and in the following order: Global Administrator, Local Data Part 1, Local Data Part 2.

[5.](#) Reserved Large BGP Community values

The Large BGP Community attribute values in the following ranges are reserved:

```
0:0:0 - 0:4294967295:4294967295
65535:0:0 - 65535:4294967295:4294967295
4294967295:0:0 - 4294967295:4294967295:4294967295
```

[6.](#) Error Handling

The error handling of Large Communities is as follows:

- o A Large Communities BGP Path Attribute with a length of zero MUST be ignored upon receipt and removed when sending.
- o A Large Communities attribute SHALL be considered malformed if its length is not a non-zero multiple of 12 bytes.
- o A BGP UPDATE message with a malformed Large Communities attribute SHALL be handled using the approach of "treat-as-withdraw" as described in [section 2 \[RFC7606\]](#).

The BGP Large Communities Global Administrator field may contain any value, and a Large Communities attribute MUST NOT be considered malformed if the Global Administrator field contains an unallocated, unassigned or reserved ASN or is set to one of the reserved Large BGP Community values defined in [Section 5](#).

[7.](#) Security Considerations

This extension to BGP has similar security implications as BGP Communities [[RFC1997](#)].

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

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

8. Implementation status - RFC EDITOR: REMOVE BEFORE PUBLICATION

This section records the status of known implementations of the protocol defined by this specification at the time of posting of this Internet-Draft, and is based on a proposal described in [[RFC7942](#)]. The description of implementations in this section is intended to assist the IETF in its decision processes in progressing drafts to RFCs. Please note that the listing of any individual implementation here does not imply endorsement by the IETF. Furthermore, no effort has been spent to verify the information presented here that was supplied by IETF contributors. This is not intended as, and must not be construed to be, a catalog of available implementations or their

Heitz, et al.

Expires April 11, 2017

[Page 5]

Internet-Draft

Large BGP Communities

October 2016

features. Readers are advised to note that other implementations may exist.

As of today these vendors have produced an implementation of Large BGP Community:

- o Cisco IOS XR
- o ExaBGP
- o GoBGP
- o BIRD
- o OpenBGPD

The latest implementation news is tracked at <http://largebgpcommunities.net/> [1].

9. IANA Considerations

IANA has made an Early Allocation of the value 30 (LARGE_COMMUNITY) in the "BGP Path Attributes" registry under the "Border Gateway Protocol (BGP) Parameters" group and is now asked to make that Permanent.

10. Acknowledgments

The authors would like to thank Ruediger Volk, Russ White, Acee Lindem, Shyam Sethuram, Jared Mauch, Joel M. Halpern, Nick Hilliard, Jeffrey Haas, John Heasley, Gunter van de Velde, Marco Marzetti, Eduardo Ascenco Reis, Mark Schouten, Paul Hoogsteder, Martijn Schmidt, Greg Hankins, Acee Lindem, 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 and Tom Scholl for their support, insightful review and comments.

11. References

11.1. Normative 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>>.

- [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>>.
- [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>>.
- [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>>.

11.2. Informative References

- [RFC1998] Chen, E. and T. Bates, "An Application of the BGP Community Attribute in Multi-home Routing", [RFC 1998](#), DOI 10.17487/RFC1998, August 1996, <<http://www.rfc-editor.org/info/rfc1998>>.
- [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>>.
- [RFC7942] Sheffer, Y. and A. Farrel, "Improving Awareness of Running Code: The Implementation Status Section", [BCP 205](#), [RFC 7942](#), DOI 10.17487/RFC7942, July 2016, <<http://www.rfc-editor.org/info/rfc7942>>.

11.3. URIs

- [1] <https://largebgpcommunities.net>

Authors' Addresses

Heitz, et al.

Expires April 11, 2017

[Page 7]

Internet-Draft

Large BGP Communities

October 2016

Jakob Heitz

Cisco
170 West Tasman Drive
San Jose, CA 95054
USA

Email: jheitz@cisco.com

Keyur Patel
Arrcus, Inc

Email: keyur@arrcus.com

Job Snijders
NTT Communications
Theodorus Majofskistraat 100
Amsterdam 1065 SZ
NL

Email: job@ntt.net

Ignas Bagdonas
Equinix
London
UK

Email: ibagдона.ietf@gmail.com

Adam Simpson
Nokia
600 March Road
Ottawa Ontario K2K 2E6
Canada

Email: adam.1.simpson@nokia.com