

IDR
Internet-Draft
Intended status: Standards Track
Expires: January 6, 2017

J. Heitz
K. Patel
Cisco
J. Snijders
NTT
I. Bagdonas
Equinix
July 5, 2016

Large BGP Community
draft-heitz-idr-large-community-00

Abstract

A new type of BGP community attribute that contains communities that each hold a 4-octet AS number and a 6-octet opaque field is defined.

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 January 6, 2017.

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](#) 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 Community Attribute	2
3.	4-Octet AS Specific Large Community	4
3.1.	Textual Representation	5
4.	Compatibility with Extended Communities	5
5.	Large Regular Communities	6
6.	Security Considerations	6
7.	IANA Considerations	6
8.	Acknowledgements	7
9.	References	7
9.1.	Normative References	7
9.2.	Informative References	8
	Authors' Addresses	8

[1.](#) Introduction

A BGP Community attribute is defined that encodes 12 byte communities, suitable for 4-Octet Autonomous System Numbers that require a 6-Octet Local Administrator field.

The 2-octet AS Specific Extended Community defined in [[RFC4360](#)] has been widely used. 4-octet AS numbers as defined by [[RFC4893](#)] are unable to make use of this popular extended community. Subsequently, [[RFC5668](#)] defined a 4-octet AS Specific Extended community. However, to make room for the extra 2 octets of AS number, the Local Administrator field was shrunk from 4 octets to 2. This document defines a community to extend that to 6 octets.

To ensure rapid and smooth adoption of the new community attribute, it must be as similar to the extended community as possible, only bigger.

[2.](#) Large BGP Community Attribute

The Large Community Attribute is a transitive optional BGP attribute, with the Type Code (suggested 41) to be assigned by IANA. The attribute consists of a set of "Large Communities". All routes with

The organization identified by the Autonomous System number in the Global Administrator sub-field can encode any information in these sub-fields. The format and meaning of the value encoded in these sub-fields should be defined by the sub-type of the community.

3.1. Textual Representation

The textual representation of the 4-Octet AS Specific Large Community is A:B:C, where A is the Global Administrator, B is the Local Administrator 1 and C is the Local Administrator 2. A ranges from 0 to 4294967295. B ranges from 0 to 4294967295. C ranges from 0 to 65535. A, B and C are plain decimal numbers without leading zeroes. Each number must appear, even if it is 0. For example, "0:1:2" cannot be written as ":1:2".

4. Compatibility with Extended Communities

Any 2-octet AS Specific Extended Community [[RFC4360](#)] can be converted into a 4-octet AS Specific Large Community by copying:

- o bits 1 and 2 of the Type into the Transitivity, and
- o the Sub-Type unchanged, and
- o the 2-octet Global Administrator into the low order octets of the 4-octet Global Administrator, and
- o the 4-octet Local Administrator into the Local Administrator 1, and
- o setting the remaining octets to zero.

Notice that the Global Administrator and the Local Administrator fields in the Large community are in the reverse order compared to those in the Extended Community. This is done for better octet alignment.

If a community contains an Autonomous System Number less than 65536 and a Local Administrator 2 field of zero, then it can be represented either as a 4-Octet AS Specific Large Community or a 2-Octet AS Specific Extended Community. These communities would be treated as different, even though they hold the same information. To prevent such inconsistencies, such communities SHOULD be encoded as a 2-Octet Specific Extended Community.

Similarly, if a community contains an Autonomous System Number greater than 65535 and a Local Administrator field less than 65536,

then it SHOULD be encoded as a 4-Octet AS Specific Extended Community as per [[RFC5668](#)].

5. Large Regular Communities

The AS portion of BGP Communities described in [[RFC1997](#)] is too small to fit a 4-octet ASN.

[[I-D.ietf-idr-as4octet-extcomm-generic-subtype](#)] defines an Extended Community sub-type to perform the same function with a 4-octet ASN. Large Communities will provide the same functionality, but provide an extra 4 octets of Local Administrator space.

6. Security Considerations

TBD

7. IANA Considerations

IANA is requested to assign a BGP path attribute value for the Large community attribute.

IANA is requested to create and maintain a registry for the Type field of the Large Community. This document reserves the Type value 0 for the 4-Octet AS Specific Large Community.

IANA is requested to create and maintain a registry for the Sub-Type field of the 4-Octet AS Specific Large Community. The initial values in the registry should be the same as those in the registry for the 2-octet AS Specific Extended Community. These values are reproduced as follows:

0x02 Route Target [[RFC4360](#)]
0x03 Route Origin [[RFC4360](#)]
0x04 Link Bandwidth [[I-D.ietf-idr-link-bandwidth](#)]
0x05 OSPF Domain Identifier [[RFC4577](#)]
0x08 BGP Data Collection [[RFC4384](#)]
0x09 Source AS [[RFC6514](#)]
0x0a L2VPN Identifier [[RFC6074](#)]
0x10 Cisco VPN-Distinguisher [Eric_Rosen]

0x80 Virtual-Network Identifier Extended Community
[[I-D.drao-bgp-l3vpn-virtual-network-overlays](#)]

As the generic sub-type defined in
[[I-D.ietf-idr-as4octet-extcomm-generic-subtype](#)] is 4 and clashes with
the value for the Link Bandwidth, IANA is requested to assign a new
value.

8. Acknowledgements

Thanks to Russ White, Acee Lindem and Shyam Sethuram for insightful
review and comments.

9. References

9.1. Normative References

- [I-D.ietf-idr-as4octet-extcomm-generic-subtype]
Rao, D., Mohapatra, P., and J. Haas, "Generic Subtype for
BGP Four-octet AS specific extended community", [draft-ietf-idr-as4octet-extcomm-generic-subtype-08](#) (work in
progress), June 2015.
- [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>>.
- [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>>.
- [RFC4893] Vohra, Q. and E. Chen, "BGP Support for Four-octet AS
Number Space", [RFC 4893](#), DOI 10.17487/RFC4893, May 2007,
<<http://www.rfc-editor.org/info/rfc4893>>.
- [RFC5668] Rekhter, Y., Sangli, S., and D. Tappan, "4-Octet AS
Specific BGP Extended Community", [RFC 5668](#),
DOI 10.17487/RFC5668, October 2009,
<<http://www.rfc-editor.org/info/rfc5668>>.

9.2. Informative References

- [I-D.drao-bgp-l3vpn-virtual-network-overlays]
Rao, D., Mullooly, J., and R. Fernando, "Layer-3 virtual network overlays based on BGP Layer-3 VPNs", [draft-drao-bgp-l3vpn-virtual-network-overlays-03](#) (work in progress), July 2014.
- [I-D.ietf-idr-link-bandwidth]
Mohapatra, P. and R. Fernando, "BGP Link Bandwidth Extended Community", [draft-ietf-idr-link-bandwidth-06](#) (work in progress), January 2013.
- [RFC4384] Meyer, D., "BGP Communities for Data Collection", [BCP 114](#), [RFC 4384](#), DOI 10.17487/RFC4384, February 2006, <<http://www.rfc-editor.org/info/rfc4384>>.
- [RFC4577] Rosen, E., Psenak, P., and P. Pillay-Esnault, "OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)", [RFC 4577](#), DOI 10.17487/RFC4577, June 2006, <<http://www.rfc-editor.org/info/rfc4577>>.
- [RFC6074] Rosen, E., Davie, B., Radoaca, V., and W. Luo, "Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs)", [RFC 6074](#), DOI 10.17487/RFC6074, January 2011, <<http://www.rfc-editor.org/info/rfc6074>>.
- [RFC6514] Aggarwal, R., Rosen, E., Morin, T., and Y. Rekhter, "BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs", [RFC 6514](#), DOI 10.17487/RFC6514, February 2012, <<http://www.rfc-editor.org/info/rfc6514>>.

Authors' Addresses

Jakob Heitz
Cisco
170 West Tasman Drive
San Jose, CA, CA 95054
USA

Email: jheitz@cisco.com

Keyur Patel
Cisco
170 West Tasman Drive
San Jose, CA 95054
USA

Email: keyupate@cisco.com

Job Snijders
NTT Communications, Inc.
Theodorus Majofskistraat 100
Amsterdam 1065 SZ
NL

Email: job@ntt.net

Ignas Bagdonas
Equinix
London
UK

Email: ibagdona.ietf@gmail.com

