

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: June 4, 2017

D. Rao
Cisco Systems
P. Mohapatra
Sproute Networks
J. Haas
Juniper Networks, Inc.
December 1, 2016

Generic Subtype for BGP Four-octet AS specific extended community
draft-ietf-idr-as4octet-extcomm-generic-subtype-10

Abstract

Maintaining the current best practices with communities, ISPs and enterprises that are assigned a 4-octet AS number may want the BGP UPDATE messages they receive from their customers or peers to include a 4-octet AS specific BGP extended community. This document defines a new sub-type within the four-octet AS specific extended community to facilitate this practice.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in [[RFC2119](#)] only when they appear in all upper case. They may also appear in lower or mixed case as English words, without normative meaning.

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 June 4, 2017.

Internet-Draft

4octet extcomm generic sub-type

December 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](#) 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. | Generic Sub-type Definition | 3 |
| 3. | Deployment Considerations | 4 |
| 4. | Acknowledgments | 4 |
| 5. | IANA Considerations | 4 |
| 6. | Security Considerations | 5 |
| 7. | References | 5 |
| 7.1. | Normative References | 5 |
| 7.2. | Informative References | 5 |
| Appendix A. | Document History | 5 |
| | Authors' Addresses | 6 |

[1.](#) Introduction

Maintaining the current best practices with communities, ISPs and enterprises that are assigned a 4-octet AS number may want the BGP UPDATE messages they receive from their customers or peers to include a 4-octet AS specific extended community. This document defines a new sub-type within the four-octet AS specific extended community to facilitate this practice.

For example, [\[RFC1998\]](#) describes an application of BGP community attribute ([\[RFC1997\]](#)) to implement flexible routing policies for sites multi-homed to one or multiple providers. In a two-octet AS environment, the advertised routes are usually associated with a community attribute that encodes the provider's AS number in the


first two octets of the community and a LOCAL_PREF value in the second two octets of the community. The community attribute signals the provider edge routers connected to the site to set the corresponding LOCAL_PREF on their advertisements to the IBGP mesh.

In this way, customers can put into practice topologies like active-backup.

When such a provider is assigned a four-octet AS number, the existing mechanism of using communities is not sufficient since the AS portion of the [RFC 1997](#) community cannot exceed two bytes. The natural alternative is to extend the same mechanism using extended communities since it allows for encoding eight bytes of information.

[RFC5668] defines a format for a four-octet AS specific extended community with a designated type field. That document defines two sub-types: Four-octet specific Route Target extended community and Four-octet specific Route Origin extended community. This document specifies a generic sub-type for the four-octet AS specific extended community to provide benefits such as the one cited above as the Internet migrates to four-octet AS space.

2. Generic Sub-type Definition



This is an extended type with Type Field comprising of 2 octets and Value Field comprising of 6 octets.

The high-order octet of this extended type is set to either 0x02 (for transitive communities) or 0x42 (for non-transitive communities). The low-order octet or the sub-type is set to 0x04.

The Value Field consists of two sub-fields:

Global Administrator sub-field: 4 octets

This sub-field contains a four-octet Autonomous System number.

Local Administrator sub-field: 2 octets

This sub-field contains a value that can influence routing policies. This value has semantics that are of significance for the Autonomous System in the Global Administrator field.

[3.](#) Deployment Considerations

There are situations in peering where a 4-octet AS specific generic extended community cannot be used.

A speaker with a 4-octet AS may not support 4-octet extended communities; or the speaker may have a customer or peer that does not support 4-octet extended communities. In all such cases, the speaker may need to define an appropriate standard community value for the same purpose. As an example, a peer may tag its routes with a community that encodes AS_TRANS [[RFC4893](#)] as the first two octets.

Similarly, as per [[RFC4893](#)], a 2-octet Autonomous System number can be converted into a 4-octet Autonomous System number by setting the two high-order octets of the 4-octet field to zero. As a consequence, at least in principle, an Autonomous System that has a 2-octet AS number could use either a standard community or the 4-octet AS specific generic extended community. This is undesirable, as they would be treated as different communities, even if they had the same values.

Therefore, for backward compatibility with existing deployments and to avoid inconsistencies between standard communities and 4-octet extended communities, Autonomous Systems that use 2-octet Autonomous System numbers SHOULD use standard 2-octet communities as defined in [RFC1997](#) rather than the 4-octet AS specific extended community as defined in this document.

[4.](#) Acknowledgments

The authors would like to thank Paul Jakma, Bruno Decraene and Cayle Spandon for their useful comments on the document.

5. IANA Considerations

Prior revisions of this document requested IANA to make assignments from the Transitive Four-Octet AS Specific Extended Community Sub-Type registry and the Non-Transitive Four-Octet AS Specific Extended Community Sub-Type registry. The sub-type value of 0x04 in each of those registries was previously assigned:

| Name | Value |
|---|--------|
| ---- | ----- |
| transitive generic four-octet AS specific | 0x0204 |
| non-transitive generic four-octet AS specific | 0x4204 |

IANA is requested to deprecate these assignments.

6. Security Considerations

There are no additional security risks introduced by this design.

7. References

7.1. 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>>.
- [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>>.

[7.2.](#) Informative References

- [I-D.ietf-idr-large-community]
Heitz, J., Snijders, J., Patel, K., Bagdonas, I., and N.
Hilliard, "BGP Large Communities", [draft-ietf-idr-large-community-09](#) (work in progress), November 2016.
- [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>>.

[Appendix A.](#) Document History

This final version of the document exists only to request IANA to deprecate its prior Extended Community assignments and provide a historical record of the reason.

During the development of the BGP Four-octet feature [[RFC4893](#)], operators had offered their commentary that parity was needed with existing BGP Community practices similar to those defined in [[RFC1998](#)]. What became clear over time was that some operators

| | | |
|-------------|----------------------|----------|
| Rao, et al. | Expires June 4, 2017 | [Page 5] |
|-------------|----------------------|----------|

| | | |
|----------------|---------------------------------|---------------|
| Internet-Draft | 4octet extcomm generic sub-type | December 2016 |
|----------------|---------------------------------|---------------|

encoded an AS number as the second field of their community; essentially, as the "target".

Since an Extended Community's Local Administrator field cannot encode more than two octets of value, the Extended Community format was not appropriate for addressing parity of existing operational practices. The BGP Large Communities Feature [[I-D.ietf-idr-large-community](#)] supplanted the work begun in this document.

Authors' Addresses

Dhananjaya Rao
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95134

USA

Email: dhrao@cisco.com

Pradosh Mohapatra
Sproute Networks

Email: mpradosh@yahoo.com

Jeffrey Haas
Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, CA 94089
US

Email: jhaas@juniper.net