

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 26, 2019

J. Borkenhagen
AT&T
R. Bush
Internet Initiative Japan
R. Bonica
Juniper Networks
S. Bayraktar
Cisco Systems
January 22, 2019

Well-Known Community Policy Behavior
draft-ietf-grow-wkc-behavior-02

Abstract

Well-Known BGP Communities are manipulated inconsistently by current implementations. This results in difficulties for operators. Network operators are encouraged to deploy consistent community handling across their networks, taking the inconsistent behaviors from the various bgp implementations they operate into consideration. Also, bgp implementors are expected to not create any further inconsistencies from this point forward.

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 [RFC 2119](#) [[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 <https://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 July 26, 2019.

Copyright Notice

Copyright (c) 2019 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 (<https://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.	Manipulation of Communities by Policy	3
3.	Community Manipulation Policy Differences	3
4.	Documentation of Vendor Implementations	3
4.1.	Note on an Inconsistency	4
5.	Note for Those Writing RFCs for New Community-Like Attributes	4
6.	Action Items	5
7.	Security Considerations	5
8.	IANA Considerations	5
9.	Acknowledgements	5
10.	Normative References	5
	Authors' Addresses	6

[1.](#) Introduction

The BGP Communities Attribute was specified in [[RFC1997](#)] which introduced the concept of Well-Known Communities. In hindsight, [[RFC1997](#)] did not prescribe as fully as it should have how Well-Known Communities may be manipulated by policies applied by operators. Currently, implementations differ in this regard, and these differences can result in inconsistent behaviors that operators find difficult to identify and resolve.

This document describes the current behavioral differences in order to assist operators in generating consistent community-manipulation policies in a multi-vendor environment, and to prevent the introduction of additional divergence in implementations.

2. Manipulation of Communities by Policy

[RFC1997] says:

"A BGP speaker receiving a route with the COMMUNITIES path attribute may modify this attribute according to the local policy."

A basic operational need is to add or remove one or more communities to the received set. Another common need is to replace all received communities with a new set. To simplify the second case, most BGP policy implementations provide syntax to "set" community that operators use to mean "remove any/all communities present on the update received from the neighbor, and apply this set of communities instead."

Some operators prefer to write explicit policy to delete unwanted communities rather than using "set;" i.e. using a "delete community *:*" and then "add community x:y ..." configuration statements in an attempt to replace all received communities. The same community manipulation policy differences described in the following section exist in both "set" and "delete community *:*" syntax. For simplicity, the remainder of this document refers only to the "set" behaviors.

3. Community Manipulation Policy Differences

Vendor implementations differ in the treatment of certain Well-Known communities when modified using the syntax to "set" the community. Some replace all communities including the Well-Known ones with the new set, while others replace all non-Well-Known Communities but do not modify any Well-Known Communities that are present.

These differences result in what would appear to be identical policy configurations having very different results on different platforms.

4. Documentation of Vendor Implementations

In this section we document the syntax and observed behavior of the "set" directive in several popular bgp implementations.

In Juniper Networks' JunOS, "community set" removes all received communities, Well-Known or otherwise.

In Cisco Systems' IOS-XR, "set community" removes all received communities except for the following:

Numeric	Common Name
0:0	internet
65535:0	graceful-shutdown
65535:1	accept-own rfc7611
65535:65281	NO_EXPORT
65535:65282	NO_ADVERTISE
65535:65283	NO_EXPORT_SUBCONFED (or local-AS)

Communities not removed by Cisco IOS/XR

Table 1

IOS-XR does allow Well-Known communities to be removed one at a time by explicit policy; for example, "delete community accept-own". Operators are advised to consult IOS-XR documentation and/or Cisco Systems support for full details.

On Brocade NetIron: "set community X" removes all communities and sets X.

In Huawei's VRP product, "community set" removes all received communities, well-known or otherwise.

In OpenBSD's OpenBGPD, "set community" does not remove any communities, Well-Known or otherwise.

[4.1.](#) Note on an Inconsistency

The IANA publishes a list of Well-Known Communities [[IANA-WKS](#)].

IOS-XR's set of well-known communities that "set community" will not overwrite diverges from IANA's list. Quite a few well-known communities from IANA's list do not receive special treatment in IOS-XR, and at least one specific community on IOS-XR's special treatment list (internet == 0:0) is not really on IANA's list -- it's taken from the "Reserved" range [0x00000000-0x0000FFFF].

This merely notes an inconsistency. It is not a plea to 'protect' the entire IANA list from "set community."

[5.](#) Note for Those Writing RFCs for New Community-Like Attributes

Care should be taken when establishing new [[RFC1997](#)]-like attributes (large communities, wide communities, etc) to avoid repeating this mistake.

6. Action Items

Unfortunately, it would be operationally disruptive for vendors to change their current implementations.

Vendors SHOULD clearly document the behavior of "set" directive in their implementations.

Vendors MUST ensure that any Well-Known Communities specified after this document's publication are removed by their "set" directive.

Given the implementation inconsistencies described in this document, network operators are urged never to rely on any implicit understanding of a neighbor ASN's bgp community handling. I.e., before announcing prefixes with NO_EXPORT or any other community to a neighbor ASN, the operator should confirm with that neighbor how the community will be treated.

Network operators are encouraged to limit their use of the "set" directive (within reason), to improve the readability of their configurations and hopefully to achieve behavioral consistency across platforms.

7. Security Considerations

Surprising defaults and/or undocumented behaviors are not good for security. This document attempts to remedy that.

8. IANA Considerations

This document has no IANA Considerations other than to be aware that any future Well-Known Communities will be subject to the policy treatment described here.

9. Acknowledgements

The authors thank Martijn Schmidt, Qin Wu for the Huawei data point, Job Snijders, David Farmer, John Heasley, and Jakob Heitz.

10. Normative References

[IANA-WKS]

"IANA Well-Known Communities",
<<https://www.iana.org/assignments/bgp-well-known-communities/bgp-well-known-communities.xhtml>>.

[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>>.

Authors' Addresses

Jay Borkenhagen
AT&T
200 Laurel Avenue South
Middletown, NJ 07748
United States of America

Email: jayb@att.com

Randy Bush
Internet Initiative Japan
5147 Crystal Springs
Bainbridge Island, WA 98110
United States of America

Email: randy@psg.com

Ron Bonica
Juniper Networks
2251 Corporate Park Drive
Herndon, VA 20171
US

Email: rbonica@juniper.net

Serpil Bayraktar
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95134
United States of America

Email: serpil@cisco.com

