

Internet Engineering Task Force (IETF)
Request for Comments: 7674
Updates: 5575
Category: Standards Track
ISSN: 2070-1721

J. Haas, Ed.
Juniper Networks
October 2015

Clarification of the Flowspec Redirect Extended Community

Abstract

This document updates RFC 5575 ("Dissemination of Flow Specification Rules") to clarify the formatting of the BGP Flowspec Redirect Extended Community.

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 5741.

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/rfc7674>.

Copyright Notice

Copyright (c) 2015 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	3
2. IANA Considerations	5
2.1. BGP Transitive Extended Community Types	5
2.2. Update to BGP Generic Transitive Experimental Use Extended Community Sub-Types	5
2.3. Generic Transitive Experimental Use Extended Community Part 2 Sub-Types	5
2.4. Generic Transitive Experimental Use Extended Community Part 3 Sub-Types	6
3. Security Considerations	6
4. Normative References	7
Acknowledgements	7
Author's Address	7

1. Introduction

"Dissemination of Flow Specification Rules" [RFC5575], commonly known as BGP Flowspec, provided for a BGP Extended Community [RFC4360] that served to redirect traffic to a Virtual Routing and Forwarding (VRF) instance that matched the flow specification's Network Layer Reachability Information (NLRI). In RFC 5575, the Redirect Extended Community was documented as follows:

```
: +-----+-----+-----+-----+
: | type   | extended community | encoding                |
: +-----+-----+-----+-----+
: | 0x8008 | redirect           | 6-byte Route Target    |
: +-----+-----+-----+-----+
:
: [...]
:
: Redirect: The redirect extended community allows the traffic to be
: redirected to a VRF routing instance that lists the specified
: route-target in its import policy. If several local instances
: match this criteria, the choice between them is a local matter
: (for example, the instance with the lowest Route Distinguisher
: value can be elected). This extended community uses the same
: encoding as the Route Target extended community [RFC4360].
: [...]
:
: 11. IANA Considerations
: [...]
:
: The following traffic filtering flow specification rules have been
: allocated by IANA from the "BGP Extended Communities Type -
: Experimental Use" registry as follows:
: [...]
:
: 0x8008 - Flow spec redirect
```

The IANA registry of BGP Extended Communities clearly identifies communities of specific formats. For example, "Two-octet AS Specific Extended Community" [RFC4360], "Four-octet AS Specific Extended Community" [RFC5668], and "IPv4 Address Specific Extended Community" [RFC4360]. Route Targets [RFC4360] identify this format in the high-order (Type) octet of the Extended Community and set the value of the low-order (Sub-Type) octet to 0x02. The Value field of the Route Target Extended Community is intended to be interpreted in the context of its format.

Since the Redirect Extended Community only registered a single codepoint in IANA's BGP Extended Community registry, a common interpretation of the Redirect Extended Community's "6-byte Route Target" has been to look, at a receiving router, for a Route Target value that matches the Route Target value in the received Redirect Extended Community and import the advertised route to the corresponding VRF instance subject to the rules defined in [RFC5575]. However, because the Route Target format in the Redirect Extended Community is not clearly defined, the wrong match may occur.

This "value wildcard" matching behavior, which does not take into account the format of the Route Target defined for a local VRF and may result in the wrong matching decision, does not match deployed implementations of BGP Flowspec. Deployed implementations of BGP Flowspec solve this problem by defining different Redirect Extended Communities that are specific to the format of the Route Target value. This document defines the following Redirect Extended Communities:

type	extended community	encoding
0x8008	redirect AS-2byte	2-octet AS, 4-octet Value
0x8108	redirect IPv4	4-octet IPv4 Address, 2-octet Value
0x8208	redirect AS-4byte	4-octet AS, 2-octet Value

It should be noted that the low-order nibble of the Redirect's Type field corresponds to the Route Target Extended Community format field (Type). (See Sections 3.1, 3.2, and 4 of [RFC4360] plus Section 2 of [RFC5668].) The low-order octet (Sub-Type) of the Redirect Extended Community remains 0x08, in contrast to 0x02 for Route Targets.

The IANA registries for the BGP Extended Communities document [RFC7153] was written to update the previously mentioned IANA registries to better document BGP Extended Community formats. The IANA Considerations section below further amends those registry updates in order to properly document the Flowspec redirect communities.

2. IANA Considerations

2.1. BGP Transitive Extended Community Types

IANA has updated the "BGP Transitive Extended Community Types" registry as follows:

0x81 - Generic Transitive Experimental Use Extended Community Part 2
(Sub-Types are defined in the "Generic Transitive
Experimental Extended Community Part 2 Sub-Types" Registry)

0x82 - Generic Transitive Experimental Use Extended Community Part 3
(Sub-Types are defined in the "Generic Transitive
Experimental Use Extended Community Part 3 Sub-Types"
Registry)

2.2. Update to BGP Generic Transitive Experimental Use Extended Community Sub-Types

IANA has updated the "BGP Generic Transitive Experimental Use
Extended Community Sub-Types" registry as follows:

0x08 - Flow spec redirect AS-2byte format [RFC5575] [RFC7674]

2.3. Generic Transitive Experimental Use Extended Community Part 2 Sub-Types

IANA has created the "Generic Transitive Experimental Use Extended
Community Part 2 Sub-Types" registry. This has been created under
the "Border Gateway Protocol (BGP) Extended Communities" registry and
contains the following note:

This registry contains values of the second octet (the "Sub-Type"
field) of an extended community when the value of the first octet
(the "Type" field) is 0x81.

Registry Name: Generic Transitive Experimental Use Extended Community
Part 2 Sub-Types

RANGE		REGISTRATION PROCEDURE	
0x00-0xbf		First Come First Served	
0xc0-0xff		IETF Review	
SUB-TYPE VALUE		NAME	REFERENCE
0x00-0x07		Unassigned	
0x08		Flow spec redirect IPv4 format	[RFC7674]
0x09-0xff		Unassigned	

2.4. Generic Transitive Experimental Use Extended Community Part 3 Sub-Types

IANA has created the "Generic Transitive Experimental Use Extended Community Part 3 Sub-Types" registry. This registry has been created under the "Border Gateway Protocol (BGP) Extended Communities" registry and contains the following note:

This registry contains values of the second octet (the "Sub-Type" field) of an extended community when the value of the first octet (the "Type" field) is 0x82.

Registry Name: Generic Transitive Experimental Use Extended Community Part 2 Sub-Types

RANGE	REGISTRATION PROCEDURE	
0x00-0xbf	First Come First Served	
0xc0-0xff	IETF Review	
SUB-TYPE VALUE	NAME	REFERENCE
0x00-0x07	Unassigned	
0x08	Flow spec redirect AS-4byte format	[RFC7674]
0x09-0xff	Unassigned	

3. Security Considerations

This document introduces no additional security considerations than those already covered in [RFC5575]. It should be noted that if the wildcard behavior were actually implemented, this ambiguity may lead to the installation of Flowspec rules in an incorrect VRF and may lead to traffic to be incorrectly delivered.

4. Normative References

- [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>>.
- [RFC5575] Marques, P., Sheth, N., Raszuk, R., Greene, B., Mauch, J., and D. McPherson, "Dissemination of Flow Specification Rules", RFC 5575, DOI 10.17487/RFC5575, August 2009, <<http://www.rfc-editor.org/info/rfc5575>>.
- [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>>.
- [RFC7153] Rosen, E. and Y. Rekhter, "IANA Registries for BGP Extended Communities", RFC 7153, DOI 10.17487/RFC7153, March 2014, <<http://www.rfc-editor.org/info/rfc7153>>.

Acknowledgements

The content of this document was raised as part of implementation discussions of the BGP Flowspec with the following individuals:

Andrew Karch (Cisco)

Robert Raszuk

Adam Simpson (Alcatel-Lucent)

Matthieu Texier (Arbor Networks)

Kaliraj Vairavakkalai (Juniper)

Author's Address

Jeffrey Haas (editor)
Juniper Networks

Email: jhaas@juniper.net

