Inter-Domain Routing (IDR)
Internet-Draft
Intended status: Standards Track
Expires: June 5, 2017

W. Eddy
J. Dailey
G. Clark
MTI Systems

December 2, 2016

BGP Flow Specification Packet-Rate Action
draft-ietf-idr-flowspec-packet-rate-01

Abstract

This document defines a new type of traffic filtering action for the BGP flow specification. The new packet-rate action allows specifying a rate-limit in number of packets per second. This is intended to be used in combatting some types of denial of service attacks where the packet rate is more important than the raw bitrate of the attack traffic.

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 5, 2017.

Table of Contents

1. Introduction

The existing BGP flow specification [RFC5575] standard supports traffic-rate limits conveyed in bytes per second. In some cases, it may be easier, faster, or more relevant to perform accounting and decision-making based on quantities of packets per second. It is desirable to specify rate limits in terms of the number of packets per second, and not just the number of bytes per second.

As an example use case, there are several types of denial of service attacks that do not require large amounts of bandwidth, but operate based on a cadence of packets over time. TCP SYN flooding attacks are one well-known example [RFC4987], along with common packet-rate limits applied to ICMP messages and packets to unknown UDP ports.

1.1. 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 RFC 2119 [RFC2119].

2. Packet Rate Action

The traffic filtering actions pertaining to a matched flow specification are indicated using BGP extended communities [RFC7153]. Particular extended community values are defined in RFC 5575 for a

number of possible actions. New types of actions can be defined using additional extended community values. The value 0x8006 has been defined as the "traffic-rate" action, and specifies a rate-limit in a quantity of bytes per second. The new packet-rate extended

community described in this draft is similar, except the quantity is interpreted as packets per second.

type	extended community	encoding
TBD	packet-rate	2-byte ASN, 4-byte unsigned integer

Table 1

Packet-rate: The packet-rate extended community is a transitive extended community across the autonomous-system boundary and uses following extended community encoding:

The first two octets carry the 2-octet id, which can be assigned from a 2-byte AS number. When a 4-byte AS number is locally present, the 2 least significant bytes of such an AS number can be used. This value is purely informational and should not be interpreted by the implementation.

The remaining 4 octets carry the rate information as an unsigned integer in network byte order, with packets per second as the unit represented. A packet-rate of 0 should result on all traffic for the particular flow to be discarded.

Note that this is a transitive community type, as explained in RFC 7153 and not a non-transitive type as mentioned narratively in the RFC 5575 description of the traffic-rate action.

The packet-rate action SHOULD NOT be used together with a traffic-rate action within the same flow specification, due to unclear semantics. Implementations MUST be robust to receiving both actions and may choose to honor either one or the other, or the combination of both, depending on local implementation capabilities.

3. Discussion

Interaction between multiple actions in a flow specification is a matter of ongoing work in the IETF. There are potential semantic conflicts when a traffic-rate action is combined with a packet-rate action. This is legal syntactically, but devices are not universally capable of honoring both a packet per second and a bit per second rate limit, resulting in inability to meet the operator intention if both are used within a flow specification. Resolving this is potential future work, however, the present guidance is for implementations to avoid sending such flow specifications, while

still behaving in some reasonable manner when receiving them (e.g. honoring one or the other limit if it is not possible to honor both).

The traffic-rate action is specified as a floating point value. There can be discrepancies between what is configured and what is encoded, due to loss of precision in the floating point representation, which would be annoying for operators. Based on discussions in the IDR working group, for the packet-rate action, an unsigned integer value is used, which allows unambiguous packet per second rates to be specified. This limits rates represented to under 4.3 Gpps (4.3 billion packets per second), which is believed to be sufficient for the present and foreseeable future networks.

4. Acknowledgements

The initial idea to add a packet rate action was encouraged by comments from Donald Smith at a DHS meeting in 2015. The content of this document was shaped by discussion on the IETF IDR mailing list, including comments from Aseem Choudhary, Acee Lindem, Jeff Haas, John Schiel, Robert Raszuk, Kirill Kasavchenko. Special thanks go to the IDR chairs Sue Hares, and John Scudder for their leadership and shepherding of the many other BGP enhancements concurrently underway.

5. IANA Considerations

If accepted for publication, IANA will need to allocate a BGP extended community value for the "packet-rate" action from the "Generic Transitive Experimental Use Extended Community Sub-Types" registry.

6. Security Considerations

No security considerations are raised by this document.

7. References

7.1. Normative References

[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>.

[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>.

[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>.

7.2. Informative References

[RFC4987] Eddy, W., "TCP SYN Flooding Attacks and Common Mitigations", RFC 4987, DOI 10.17487/RFC4987, August 2007, <http://www.rfc-editor.org/info/rfc4987>.

Authors' Addresses

Wesley Eddy
MTI Systems

Email: wes@mti-systems.com

Justin Dailey
MTI Systems

Email: justin@mti-systems.com

Gilbert Clark
MTI Systems

Email: gclark@mti-systems.com