Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: May 25, 2016

W. Eddy
J. Dailey
G. Clark
MTI Systems

November 22, 2015

BGP Flow Specification Packet-Rate Action
draft-eddy-idr-flowspec-packet-rate-00

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.

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 May 25, 2016.

Table of Contents

1. Introduction 2
1.1. Requirements Language 2
2. Packet Rate Action 2
3. Discussion 3
4. IANA Considerations 3
5. Security Considerations 4
6. Normative References 4
Authors' Addresses 4

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.

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 as#, 4-byte float

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 in IEEE floating point [IEEE.754.1985] format, units being packets per second. 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.

3. Discussion

Although a floating-point value for packets per second may seem odd or unnatural compared to an integer value, the motivations for this are:

The maximum value that a 32-bit unsigned integer could hold would limit to specifying under 2.15 Gpps (2.15 billion packets per second). For large or high-performance networks especially in the future, this may not be sufficient. The maximum floating point value is much higher (on the order of 10^38) and should be future-proof.

The reduced precision of the floating-point limit that can be specified compared to an integer encoding does not seem to be a major concern.

This maintains consistency with the present syntax for bytes per second rate limits.

4. 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.

5. Security Considerations

No security considerations are raised by this document.

6. Normative References

[IEEE.754.1985]
Institute of Electrical and Electronics Engineers, "Standard for Binary Floating-Point Arithmetic", IEEE Standard 754, August 1985.

[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>.

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