

IDR Working Group
Internet-Draft
Intended status: Standards Track
Expires: December 2, 2016

L. Yong
S. Hares
Q. Liang
J. You
Huawei
May 31, 2016

BGP Flow Specification Filter for MPLS Label
draft-ietf-idr-flowspec-mpls-match-00.txt

Abstract

This draft proposes BGP flow specification rules that are used to filter MPLS labeled packets.

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 December 2, 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. The Flow Specification Encoding for MPLS Match	3
3. Deployment Example: DDoS Traffic	5
4. Security Considerations	5
5. IANA Considerations	6
6. References	6
6.1. Normative References	6
6.2. Informative References	7
Authors' Addresses	7

1. Introduction

BGP Flow Specification (BGP-FS) [RFC5575] is an extension to that allows for the dissemination of traffic flow specification rules via BGP ([RFC4271]). BGP-FS policies have a match condition that may be n-tuple match in a policy, and an action that modifies the packet and forwards/drops the packet. Via BGP, new filter rules can be sent to all BGP peers simultaneously without changing router configuration, and the BGP peer can install these routes in the forwarding table. The typical application of BGP-FS is to automate the distribution of traffic filter lists to routers for DDOS mitigation.

[RFC5575] defines a new BGP Network Layer Reachability Information (NLRI) format used to distribute traffic flow specification rules. NLRI (AFI=1, SAFI=133) is for IPv4 unicast filtering. NLRI (AFI=1, SAFI=134) is for BGP/MPLS VPN filtering. [I-D.ietf-idr-flow-spec-v6] defines flow-spec extension for IPv6 data packets. [I-D.ietf-idr-flowspec-l2vpn] extends the flow-spec rules for layer 2 Ethernet packets (AFI=25, SAFI=133, SAFI=134). All these flow specifications match parts only reflect single layer IP (source/destination IP prefix, protocol type, ports, etc.) and Ethernet information with matches for source/destination MAC

MPLS technologies [RFC3031] have been widely deployed in WAN networks. MPLS label stack [RFC3032] is the foundation for label switched data plane. A label on a label stack may represent a label switch path (LSP), application identification such as Pseudo Wire (PW), a reserved label that triggers a specific data plane action, or etc. The data plane label switching operations includes pop, push, or swap label on the label stack.

For value added services, it is valuable for a MPLS network to have BGP-FS policy filter that matches on the MPLS portion of a packet and an action to modify the MPLS packet header and/or monitor the packets that match the policy. This document specifies an MPLS match filter.

[I-D.liang-idr-bgp-flowspec-label] specifies a BGP action to modify the MPLS label.

[I-D.hares-idr-flowspec-combo] describes the following two options for extending [RFC5575]:

- o Option 1: Extend [RFC5575] with new filters, match filters and actions. Extend the match default order by type and require that all matches be combined with an "AND". Extend the actions and define a default order and the resolution of conflicts.
- o Option 2: Create a version 2 of BGP flow Specification which can run in parallel to Option 1 which supports explicit ordering of match filters and actions. Option 2 will also refine the BGP-FS security to optionally include ROAs between ASes, and other mechanisms ([I-D.ietf-idr-bgp-flowspec-oid])

2. The Flow Specification Encoding for MPLS Match

This document proposes new flow specifications rules that is encoded in NLRI.

Type TBD1- MPLS Match1

Function: The match1 applies to MPLS Label field on the label stack.

Encoding: <type(1 octet), length(1 octet), [operator,value]+>.

It contains a set of {operator, value} pairs that are used for matching filter.

The operator byte is encoded as:

```

      0   1   2   3   4   5   6   7
+---+---+---+---+---+---+---+
| e | a | i | pos | Resv |
+---+---+---+---+---+---+---+

```

where:

e - end of list bit: Set in the last {op, value} pair in the list.

a - AND bit: If unset, the previous term is logically ORed with the current one. If set, the operation is a logical AND. It should be unset in the first operator byte of a

sequence. The AND operator has higher priority than OR for the purposes of evaluating logical expressions.

i - before bit: If unset, apply matching filter before MPLS label data plane action; if set, apply matching filter after MPLS label data plane action.

pos - the label position indication bits: where:

00: any position on the label stack - the presented label value is used to match any label on the label stack. When apply it, at least one label on the stack match the value

01: top label indication- the presented label value MUST be used to match the top label on the label stack.

10: bottom label indication- If it is set, the presented label value MUST match the bottom label on the label stack. When it is clear, the present label value can match to any label on the label stack

11: (for reserved labels)

The value field is encoded as:

```

0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                     |
|                                     Label                                     |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Type TBD2 - MPLS Match2

Function: MPLS Match2 applies to MPLS Label experiment bits (EXP) on the top label in the label stack.

Encoding: <type (1 octet), [op, value]+>


[op,value] - Defines a list of {operation, value} pairs used to match 3-bit exp field on the top label of packets [RFC3032].

Values are encoded using a single byte, where the five most significant bits are zero and the three least significant bits contain the exp value.

3. Deployment Example: DDoS Traffic

In this example, 5 local policy rules in the filter-based RIBs (FB-RB, aka Policy Routing) will match n-tuples (destination IP address, Destination Port, source IP address, Source IP Address, protocols (ICMP and STCP)). These policy rules can be created by standard yang modules for filter-based RIBS (configuration, and ephemeral configuration) or ACLs, or vendor based policy. These policies will put the DDoS attack data onto one LSP (LSP1) in order to send the DDoS traffic to the IDS/IPS processing attached to PE2.

The MPLS Filter allows the BGP Flow specification to match on the LSP label rather than the IP address so that PE2 (with the FB-RIBs on PE2) can forward the traffic to a set of IDS/IPS machines. The BGP Flow Specification (BGP-FS) can forward this simple match policy along with an action policy that constraints the traffic on this Flow to a certain rate (bytes/second).



BGP Flow Specification Filter 1

```

BGP Flow Specification
Match Policy
  Destination IP address (0/0) [Required by RFC5575]
  MPLS Label match (label-1)
Action Policy
  Traffic-rate (n bytes)
  
```

4. Security Considerations

The validation of BGP Flow Specification policy is considered in [I-D.hares-idr-flowspec-combo] for option 1, and for option 2. For Option 1, the MPLS Match can be one of the match filtes, and and the final match is an "AND" of all the filters. Match filters are tested in the order specified in [I-D.hares-idr-flowspec-combo] and/or an RFC5575bis document.

The traffic rate action described above is described in [RFC5575]. [I-D.hares-idr-flowspec-combo] suggests a default order for filters and for the BGP-FS action proposed after [RFC5575], and this document discusses how conflicts between action are handled.

5. IANA Considerations

This section complies with [RFC7153]

IANA is requested to a new entry in "Flow Spec component types registry" with the following values:

Value Name:	Value	Reference
=====	=====	=====
MPLS-Match1	TBD1	[This Document]
MPLS-Match2	TBD2	[This Document]

6. References

6.1. 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>>.
- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, DOI 10.17487/RFC3031, January 2001, <<http://www.rfc-editor.org/info/rfc3031>>.
- [RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, DOI 10.17487/RFC3032, January 2001, <<http://www.rfc-editor.org/info/rfc3032>>.
- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <<http://www.rfc-editor.org/info/rfc4271>>.
- [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>>.

6.2. Informative References

[I-D.hares-idr-flowspec-combo]

Hares, S., "An Information Model for Basic Network Policy and Filter Rules", draft-hares-idr-flowspec-combo-01 (work in progress), March 2016.

[I-D.ietf-idr-bgp-flowspec-oid]

Uttaro, J., Filsfils, C., Smith, D., Alcaide, J., and P. Mohapatra, "Revised Validation Procedure for BGP Flow Specifications", draft-ietf-idr-bgp-flowspec-oid-03 (work in progress), March 2016.

[I-D.ietf-idr-flow-spec-v6]

McPherson, D., Raszuk, R., Pithawala, B., Andy, A., and S. Hares, "Dissemination of Flow Specification Rules for IPv6", draft-ietf-idr-flow-spec-v6-07 (work in progress), March 2016.

[I-D.ietf-idr-flowspec-l2vpn]

Weiguo, H., Litkowski, S., and S. Zhuang, "Dissemination of Flow Specification Rules for L2 VPN", draft-ietf-idr-flowspec-l2vpn-04 (work in progress), May 2016.

[I-D.liang-idr-bgp-flowspec-label]

Hares, S., You, J., Raszuk, R., and d. danma@cisco.com, "Carrying Label Information for BGP FlowSpec", draft-liang-idr-bgp-flowspec-label-02 (work in progress), March 2016.

Authors' Addresses

Lucy Yong
Huawei

Email: lucy.yong@huawei.com

Susan Hares
Huawei
7453 Hickory Hill
Saline, MI 48176
USA

Email: shares@ndzh.com

Qiandeng Liang
Huawei
101 Software Avenue, Yuhuatai District
Nanjing 210012
China

Email: liangqiandeng@huawei.com

Jianjie You
Huawei
101 Software Avenue, Yuhuatai District
Nanjing 210012
China

Email: youjianjie@huawei.com