

IDR

Internet Draft

Intended status: Standards Track

W. Hao

Q. Liang

Huawei

Jim Uttaro

AT&T

S. Litkowski

Orange Business Service

S. Zhuang

Huawei

Expires: February 2016

August 12, 2015

Dissemination of Flow Specification Rules for L2 VPN  
draft-ietf-idr-flowspec-l2vpn-02.txt

Abstract

This document defines BGP flow-spec extension for Ethernet traffic filtering in L2 VPN network. SAFI=134 in [RFC5575] is redefined for dissemination traffic filtering information in an L2VPN environment. A new subset of component types and extended community also are defined.

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), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

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 ..... [2](#)
- [2.](#) Layer 2 Flow Specification encoding in BGP..... [3](#)
- [3.](#) Ethernet Flow Specification encoding in BGP..... [4](#)
- [4.](#) Ethernet Flow Specification Traffic Actions..... [6](#)
- [5.](#) Security Considerations..... [9](#)
- [6.](#) IANA Considerations ..... [9](#)
  - [6.1.](#) Normative References..... [10](#)
  - [6.2.](#) Informative References..... [11](#)
- [7.](#) Acknowledgments ..... [11](#)

[1.](#) Introduction

BGP Flow-spec is an extension to BGP that allows for the dissemination of traffic flow specification rules. It leverages the BGP Control Plane to simplify the distribution of ACLs, new filter rules can be injected to all BGP peers simultaneously without changing router configuration. The typical application of BGP Flow-spec is to automate the distribution of traffic filter lists to routers for DDOS mitigation, access control, etc.

[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. The Flow specification match part only includes L3/L4 information like source/destination prefix, protocol, ports, and etc, so traffic flows can only be selectively

filtered based on L3/L4 information.

Layer 2 Virtual Private Networks L2VPNs have already been deployed in an increasing number of networks today. In L2VPN network, we also have requirement to deploy BGP Flow-spec to mitigate DDoS attack traffic. Within L2VPN network, both IP and non-IP Ethernet traffic

maybe exist. For IP traffic filtering, the Flow specification rules defined in [\[RFC5575\]](#) which include match criteria and actions can still be used, flow specification rules received via new NLRI format apply only to traffic that belongs to the VPN instance(s) in which it is imported. For non-IP Ethernet traffic filtering, Layer 2 related information like source/destination MAC and VLAN should be considered. But the flow specification match criteria defined in [RFC5575](#) only include layer 3 and layer 4 IP information, layer 2 Ethernet information haven't been included.

There are different kinds of L2VPN networks like EVPN [\[EVPN\]](#), BGP VPLS [\[RFC4761\]](#), LDP VPLS [\[RFC4762\]](#) and border gateway protocol (BGP) auto discovery [\[RFC 6074\]](#). Because the flow-spec feature relies on BGP protocol to distribute traffic filtering rules, so it can only be incrementally deployed in those L2VPN networks where BGP is used for auto discovery and/or signaling purposes such as BGP-based VPLS [\[4761\]](#), EVPN and LDP-based VPLS [\[4762\]](#) with BGP auto-discovery [\[6074\]](#).

This draft proposes a new subset of component types and extended community to support L2VPN flow-spec application. The flow-spec rules can be enforced on all border routers or on some interface sets of the border routers. SAFI=134 in [\[RFC5575\]](#) is redefined for dissemination traffic filtering information in an L2VPN environment.

## 2. Layer 2 Flow Specification encoding in BGP

The [\[RFC5575\]](#) defines SAFI 133 and SAFI 134 for ''dissemination of IPv4 flow specification rules'' and ''dissemination of VPNv4 flow specification rules'' respectively. [\[draft-ietf-idr-flow-spec-v6-06\]](#) redefines the [\[RFC5575\]](#) SAFIs in order to make them applicable to both IPv4 and IPv6 applications. This document will further redefine the SAFI 134 in order to make them applicable to L2VPN applications.

The following changes are defined:

'SAFI 134 for dissemination of L3VPN flow specification rules' to now be defined as 'SAFI 134 for dissemination of VPN flow specification rules'

For SAFI 134 the indication to which address family it is referring to will be recognized by AFI value (AFI=1 for VPNv4, AFI=2 VPNv6 and AFI=25 for L2VPN). Such modification is fully backwards compatible with existing implementation and production deployments.

### 3. Ethernet Flow Specification encoding in BGP

The NLRI format for this address family consists of a fixed-length Route Distinguisher field (8 bytes) followed by a flow specification, following the encoding defined in this document. The NLRI length field shall include both the 8 bytes of the Route Distinguisher as well as the subsequent flow specification.

Flow specification rules received via this NLRI apply only to traffic that belongs to the VPN instance(s) in which it is imported. Flow rules are accepted by default, when received from remote PE routers.

Besides the component types defined in [[RFC5575](#)] and [[draft-ietf-idr-flow-spec-v6-06](#)], this document proposes the following additional component types for L2VPN Ethernet traffic filtering:

Type 14 - Source MAC

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

Defines a list of {operation(op), value} pairs used to match source MAC. Values are encoded as 6-byte quantities. The operation field is encoded as 'Numeric operator' defined in [[RFC5575](#)].

Type 15 - Destination MAC

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

Defines a list of {operation, value} pairs used to match destination MAC. Values are encoded as 6-byte quantities.

Type 16 - Ethernet Type

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

Defines a list of {operation, value} pairs used to match two-octet field. Values are encoded as 2-byte quantities.

Ethernet II framing defines the two-octet EtherType field in an Ethernet frame, preceded by destination and source MAC addresses, that identifies an upper layer protocol encapsulating the frame data.

Type 17 - DSAP(Destination Service Access Point) in LLC

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

Defines a list of {operation, value} pairs used to match 1-octet DSAP in the 802.2 LLC(Logical Link Control Header). Values are encoded as 1-byte quantities.

Type 18 - SSAP(Source Service Access Point) in LLC

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

Defines a list of {operation, value} pairs used to match 1-octet SSAP in the 802.2 LLC. Values are encoded as 1-byte quantities.

Type 19 - Control field in LLC

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

Defines a list of {operation, value} pairs used to match 1-octet control field in the 802.2 LLC. Values are encoded as 1-byte quantities.

Type 20 - SNAP

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

Defines a list of {operation, value} pairs used to match 5-octet SNAP(Sub-Network Access Protocol) field. Values are encoded as 5-byte quantities.

## Type 21 - VLAN ID

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

Defines a list of {operation, value} pairs used to match VLAN ID. Values are encoded as 1- or 2-byte quantities.

In virtual local-area network (VLAN) stacking case, the VLAN ID is outer VLAN ID.

## Type 22 - VLAN COS

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

Defines a list of {operation, value} pairs used to match 3-bit VLAN COS fields [802.1p]. Values are encoded using a single byte, where the five most significant bits are zero and the three least significant bits contain the VLAN COS value.

In virtual local-area network (VLAN) stacking case, the VLAN COS is outer VLAN COS.

## Type 23 - Inner VLAN ID

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

Defines a list of {operation, value} pairs used to match inner VLAN ID using for virtual local-area network (VLAN) stacking or Q in Q case. Values are encoded as 1- or 2-byte quantities.

In single VLAN case, the component type MUST not be used.

## Type 24 - Inner VLAN COS

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

Defines a list of {operation, value} pairs used to match 3-bit inner VLAN COS fields [802.1p] using for virtual local-area network (VLAN) stacking or Q in Q case. Values are encoded using a single byte, where the five most significant bits are zero and the three least

significant bits contain the VLAN COS value.

In single VLAN case, the component type MUST not be used.

#### 4. Ethernet Flow Specification Traffic Actions

The default action for a layer 2 traffic filtering flow specification is to accept traffic that matches that particular rule. The following extended community values per [RFC5575](#) can be used to specify particular actions in L2VPN network:

type	extended community	encoding
0x8006	traffic-rate	2-byte as#, 4-byte float
0x8007	traffic-action	bitmask
0x8008	redirect	6-byte Route Target
0x8009	traffic-marking	DSCP value

Redirect: The action should be redefined to allow the traffic to be redirected to a MAC or IP VRF routing instance that lists the specified route-target in its import policy.

Besides the above extended communities, this document also proposes the following BGP extended communities specifications for Ethernet flow to extend [[RFC5575](#)]:

type	extended community	encoding
TBD1	VLAN-action	bitmask
TBD2	TPID-action	bitmask

VLAN-action: The VLAN-action extended community consists of 6 bytes which include the fields of action Flags, and two VLAN ID/COS value. The action Flags field includes PO, PU, SW, RI, RO, CI and CO Flag to indicate the action type. The two VLAN ID/COS value are carried

in COS1/VLAN ID1 and COS2/VLAN ID2 field respectively.

```
0                                                                 15
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|P0|PU|P2|SW|RI|RO|CI|CO|           Resv           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| COS1   |R |           VLAN ID1                   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| COS2   |R |           VLAN ID2                   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

P0: Pop action. If the P0 flag is one it indicates the outmost VLAN should be removed.

PU: Push action. If PU is one it indicates VLAN ID1 will be added. If PU is one and P2 is set to zero, VLAN ID1 is the outmost VLAN. If PU and P2 are both set to one, VLAN ID1 will be the inner VLAN and VLAN ID2 will be the outer VLAN.

P2: Push action 2. If P2 is one it indicates VLAN ID2 will be added. If PU is zero and P2 is one, the format is invalid and no VLAN Push action should be taken.

SW: Swap action. If the SW flag is one it indicates the outer VLAN and inner VLAN should be swapped.

RI: Rewrite inner VLAN action. If the RI flag is one it indicates the inner VLAN should be replaced by a new VLAN, the new VLAN is VLAN ID1.

R0: Rewrite outer VLAN action. If the R0 flag is one it indicates the outer VLAN should be replaced by a new VLAN, the new VLAN is VLAN ID2.

CI: Mapping inner COS action. If the CI flag is one it indicates the inner COS should be replaced by a new COS, the new COS is COS1.

C0: Mapping outer COS action. If the C0 flag is one it indicates the outer COS should be replaced by a new COS, the new COS is COS2.

Resv: Reserved for future use.



COS1: 3 bits. COS value.


COS2: 3 bits. COS value.

VLAN ID1: 12 bits. VLAN ID value.

VLAN ID2: 12 bits. VLAN ID value.


R: Reserved for future use.

For example, if PUSH Inner VLAN 10 and Outer VLAN 20 action is needed, the format of the VLAN-action extended community is as follows:



In some cases, more complicated actions are needed, such as SwapPop, PushSwap, etc. These actions can't be represented by single one VLAN-action extended community and should be represented by the catenation of two VLAN-actions which are carried in two separate VLAN-action extended communities. The arrival order of the VLAN-actions decides the action order. For example, SwapPop action can be represented by the catenation of Swap action 1 and Pop action 2, PushSwap action can be represented by the catenation of Push action 1 and Swap action 2.

TPID-action: The TPID-action extended community consists of 6 bytes which includes the fields of action Flags, TPID1 and TPID2.



```

|                                     TP ID2                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

TI: Mapping inner TP ID action. If the TI flag is one it indicates the inner TP ID should be replaced by a new TP ID, the new TP ID is TP ID1.

T0: Mapping outer TP ID action. If the T0 flag is one it indicates the outer TP ID should be replaced by a new TP ID, the new TP ID is TP ID2.

Resv: Reserved for future use.

## 5. Security Considerations

No new security issues are introduced to the BGP protocol by this specification.

## 6. IANA Considerations

IANA is requested to rename currently defined SAFI 134 per [[RFC5575](#)] to read:

```

134      VPN dissemination of flow specification rules

```

IANA is requested to create and maintain a new registry for "Flow spec L2VPN Component Types". For completeness, the types defined in [[RFC5575](#)] and [[draft-ietf-idr-flow-spec-v6-06](#)] also are listed here.

```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| type   |           RFC or Draft           |           discription           |

```

1	<a href="#">RFC5575</a>	Destination Prefix
1	<a href="#">draft-ietf-idr-flow-spec-v6-06</a>	Destination IPv6 Prefix
2	<a href="#">RFC5575</a>	Source Prefix
2	<a href="#">draft-ietf-idr-flow-spec-v6-06</a>	Source IPv6 Prefix
3	<a href="#">RFC5575</a>	IP Protocol
3	<a href="#">draft-ietf-idr-flow-spec-v6-06</a>	Next Header
4	<a href="#">RFC5575</a>	Port
5	<a href="#">RFC5575</a>	Destination port
6	<a href="#">RFC5575</a>	Source port
7	<a href="#">RFC5575</a>	ICMP type
8	<a href="#">RFC5575</a>	ICMP code
9	<a href="#">RFC5575</a>	TCP flags
10	<a href="#">RFC5575</a>	Packet length
11	<a href="#">RFC5575</a>	DSCP
12	<a href="#">RFC5575</a>	Fragment
13	<a href="#">draft-ietf-idr-flow-spec-v6-06</a>	Flow Label
14	This draft	Source MAC
15	This draft	Destination MAC
16	This draft	Ethernet Type
17	This draft	DSAP in LLC
18	This draft	SSAP in LLC
19	This draft	Control field in LLC
20	This draft	SNAP
21	This draft	VLAN ID
22	This draft	VLAN COS
23	This draft	Inner VLAN ID
24	This draft	Inner VLAN COS

IANA is requested to update the reference for the following assignment in the "BGP Extended Communities Type - extended, transitive" registry:

Type value Name Reference

0x080A Flow spec VLAN action [this document]

0x080B Flow spec TPID action [this document]

## 6.1. Normative References

[1] [[RFC2119](#)] Bradner, S., "Key words for use in RFCs to Indicate

Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

- [2] [[RFC5575](#)] P. Marques, N. Sheth, R. Raszuk, B. Greene, J. Mauch, D. McPherson, "Dissemination of Flow Specification Rules", [RFC 5575](#), August 2009.
- [3] [[RFC4761](#)] K. Kompella, Ed., Y. Rekhter, Ed., "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", [RFC4761](#), January 2007.
- [4] [[RFC4762](#)] M. Lasserre, Ed., V. Kompella, Ed., "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", [RFC4762](#), January 2007.
- [5] [[RFC6074](#)] E. Rosen, B. Davie, V. Radoaca, "Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs)", [RFC6074](#), January 2011.

## [6.2](#). Informative References

- [1] [EVPN] Sajassi et al., "BGP MPLS Based Ethernet VPN", [draft-ietf-l2vpn-evpn-07.txt](#), work in progress, May, 2014.
- [2] [IEEE 802.1p] Javin, et.al. "IEEE 802.1p: LAN Layer 2 QoS/CoS Protocol for Traffic Prioritization", 2012-02-15

## [7](#). Acknowledgments

The authors wish to acknowledge the important contributions of Hannes Gredler, Xiaohu Xu and Lucy Yong.

---

Internet-Draft

EVPN Flow Spec

August 2015

Authors' Addresses

Weiguo Hao  
Huawei Technologies  
101 Software Avenue,  
Nanjing 210012  
China  
Email: haoweiguo@huawei.com

Qiandeng Liang  
Huawei Technologies  
101 Software Avenue,  
Nanjing 210012  
China  
Email: liangqiandeng@huawei.com

James Uttaro  
AT&T  
Email: uttaro@att.com

Stephane Litkowski  
Orange  
Email: stephane.litkowski@orange.com

Shunwan Zhuang  
Huawei Technologies  
Huawei Bld., No.156 Beiqing Rd.  
Beijing 100095  
China  
Email: zhuangshunwan@huawei.com

