Workgroup: IDR Working Group
Internet-Draft:
draft-ietf-idr-flowspec-l2vpn-23
Published: 15 April 2024
Intended Status: Standards Track
Expires: 17 October 2024

Authors: W. Hao D. Eastlake
Huawei Technologies Independent
S. Litkowski S. Zhuang
Cisco Systems, Inc. Huawei Technologies

BGP Dissemination of L2 Flow Specification Rules

Abstract

This document defines a Border Gateway Protocol (BGP) Flow Specification (flowspec) extension to disseminate Ethernet Layer 2 (L2) and Layer 2 Virtual Private Network (L2VPN) traffic filtering rules either by themselves or in conjunction with L3 flowspecs. AFI/SAFI 6/133 and 25/134 are used for these purposes. New component types and two extended communities are also defined.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

Border Gateway Protocol (BGP) Flow Specification [RFC8955] (flowspec) is an extension to BGP that supports the dissemination of traffic flow specifications and resulting actions to be taken on packets in a specified flow. It leverages the BGP Control Plane to simplify the distribution of ACLs (Access Control Lists). Using the Flow Specification extension new filter rules can be injected to all BGP peers simultaneously without changing router configuration. A typical application is to automate the distribution of traffic filter lists to routers for DDoS (Distributed Denial of Service) mitigation, access control, and similar applications.

BGP Flow Specification [RFC8955] defines a BGP Network Layer Reachability Information (NLRI) format used to distribute traffic flow specification rules. The NLRI for (AFI=1, SAFI=133) specifies IPv4 unicast filtering. The NLRI for (AFI=1, SAFI=134) specifies IPv4 BGP/MPLS VPN filtering [RFC7432]. The Flow Specification match part defined in [RFC8955] only includes L3/L4 information like IPv4 source/destination prefix, protocol, ports, and the like, so traffic flows can only be filtered based on L3/L4 information. This has been extended by [RFC8956] to cover IPv6 (AFI=2) L3/L4.

Layer 2 Virtual Private Networks (L2VPNs) have been deployed in an increasing number of networks. Such networks also have requirements to deploy BGP Flow Specification to mitigate DDoS attack traffic. Within an L2VPN network, both IP and non-IP Ethernet traffic may exist. For IP traffic filtering, the VPN Flow Specification rules defined in [RFC8955] and/or [RFC8956], which include match criteria and actions, can still be used. For non-IP Ethernet traffic filtering, Layer 2 related information like source/destination MAC and VLAN must be considered.

There are different kinds of L2VPN networks like EVPN [RFC7432], BGP VPLS [RFC4761], LDP VPLS [RFC4762] and border gateway protocol (BGP) auto discovery [RFC6074]. Because the Flow Specification feature relies on the BGP protocol to distribute traffic filtering rules, it can only be incrementally deployed in those L2VPN networks where BGP has already been used for auto discovery and/or signaling purposes such as BGP-based VPLS [RFC4761], EVPN, and LDP-based VPLS [RFC4762] with BGP auto-discovery [RFC6074].

This document defines new flowspec component types and two new extended communities to support L2 and L2VPN flowspec applications. The flowspec rules can be enforced on all border routers or on some interface sets of the border routers. SAFI=133 in [RFC8955] and [RFC8956] is extended for AFI=6 as specified in Section 2 to cover L2 traffic filtering information and in Section 3 SAFI=134 is extended for AFI=25 to cover the L2VPN environment.

1.1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

The following acronyms and terms are used in this document:

AFI - Address Family Identifier

ACL - Access Control List

DDoS - Distributed Denial of Service

DEI - Drop Eligible Indicator

EVPN - Ethernet VPN [RFC7432]

flowspec - BGP Flow Specification

L2 - Layer 2

L2VPN - Layer 2 VPN

L3 - Layer 3

L3VPN - Layer 3 VPN

NLRI - Network Layer Reachability Information

PCP - Priority Code Point [IEEE802.1Q]

SAFI - Subsequent Address Family Identifier

TPID - Tag Protocol ID, typically a VLAN ID

VLAN - Virtual Local Area Network

VPLS - Virtual Private Line Service [RFC4762]

VPN - Virtual Private Network

2. Layer 2 Flow Specification Encoding

[RFC8955] defines SAFI 133 and SAFI 134, with AFI=1, for "dissemination of IPv4 flow specification rules" and "dissemination of VPNv4 flow specification rules", respectively. [RFC8956] extends [RFC8955] to also allow AFI=2 thus making it applicable to both IPv4 and IPv6 applications. This document further extends SAFI=133 for AFI=6 and SAFI=134 for AFI=25 to make them applicable to L2 and L2VPN applications. This document also provides for the optional combination of L3 flow specifications with these L2 flow specifications.

This section specifies the L2 flowspec for AFI=6/SAFI=133. To simplify assignments, a new registry is used for L2 flowspec. Since it is frequently desirable to also filter on L3/L4 fields, provision is made for their inclusion along with an indication of the L3 protocol involved (IPv4 or IPv6).

The NLRI part of the MP_REACH_NLRI and MP_UNREACH_NLRI is encoded as a 1- or 2-octet total NLRI length field followed by several fields as described below.

```
+-------------------------------+
| total-length (0xnn or 0xfnnn) |  2 or 3 octets
+-------------------------------+
|             L3-AFI            |  2 octets
+-------------------------------+
|  L2-length (0xnn or 0xfnnn)   |  2 or 3 octets
+-------------------------------+
|           NLRI-value          |  variable
+-------------------------------+
```

Figure 1: Flow Specification NLRI for L2

The fields show in Figure 1 are further specified below:

total-length: The length of the subsequent fields (L3 AFI, L2-length, and NRLI-value) encoded as provided in Section 4.1 of [RFC8955]. If this field is less than 4, which is the minimum valid value, then the NLRI is malformed in which case a NOTIFICATION message is sent and the BGP connection closed as provided in Section 6.3 of [RFC4271].

L3-AFI: If no L3/L4 filtering is desired, this two octet field MUST be zero which is a reserved AFI value. Otherwise L3-AFI indicates the L3 protocol involved by giving its AFI (0x0001 for IPv4 or 0x0002 for IPv6). If the receiver does not understand the value of the L3-AFI field, the MP_REACH or MP_UNREACH attribute is ignored.

L2-length: The length of the L2 components at the beginning of the NLRI-value field encoded as provided in Section 4.1 of [RFC8955]. If the value of this field indicates that the L2 components extend beyond the total-length, the NLRI is malformed in which case a NOTIFICATION message is sent and the BGP connection closed as provided in Section 6.3 of [RFC4271]. N2-length MAY be zero although, in that case, it would have been more efficient to encode the attribute as an L3 Flow spec unless it is desired to apply an L2 action (see Section 4). A null L2 flowspec always matches.

NLRI-value: This consists of the L2 flowspec, of length L2-length, followed by an optionally present L3 flowspec. The result can be treated in most ways as a single flowspec, matching the intersection (AND) of all the components except that the components in the initial L2 region are interpreted as L2 components and the remainder as L3 components per the L3-AFI field. This is necessary because there are different registries for the L2, L3 IPv4, and L3 IPv6 component types. If the L3 flowspec is null (length zero), it always matches.

2.1. L2 Component Types

The L2 flowspec portion of the NLRI-value consists of flowspec components as in [RFC8955] but using L2 components and types as specified below. All components start with a type octet followed by a length octet followed by any additional information needed. The length octet gives the length, in octets, of the information after the length octet. This structure applies to all new components to be defined in the L2 Flow-spec Component Registry (see Section 6) and to all existing components except Types 2 and 3 where the length is in bits.

2.1.1. Type 1 - Ethernet Type (EtherType)

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

Defines a list of {operation, value} pairs used to match the two-octet EtherType field. op is encoded as specified in Section 4.2.1.1 of [RFC8955]. Values are encoded as 2-octet quantities. Ethernet II framing defines the two-octet Ethernet Type (EtherType) field in an Ethernet frame, preceded by destination and source MAC addresses, that identifies an upper layer protocol encapsulating the frame data. The match fails if LLC encoding is being used rather than EtherType encoding.

2.1.2. Type 2 - Source MAC

Encoding: <type (1 octet), MAC Prefix length (1 octet), MAC Prefix>

Defines the source MAC Address prefix to match encoded as in BGP UPDATE messages [RFC4271]. Prefix length is in bits and the MAC Prefix is fill out with from 1 to 7 padding bits so that it is an integer number of octets. These padding bits are ignored for matching purposes.

2.1.3. Type 3 - Destination MAC

Encoding: <type (1 octet), MAC Prefix length (1 octet), MAC Prefix>

Defines the destination MAC Address to match encoded as in BGP UPDATE messages [RFC4271]. Prefix length is in bits and the MAC Prefix is fill out with from 1 to 7 padding bits so that it is an integer number of octets. These padding bits are ignored for matching purposes.

2.1.4. Type 4 - DSAP (Destination Service Access Point)

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

Defines a list of {operation, value} pairs used to match the 1-octet DSAP in the IEEE 802.2 LLC (Logical Link Control Header). Values are encoded as 1-octet quantities. op is encoded as specified in Section 4.2.1.1 of [RFC8955]. The match fails if EtherType L2 header encoding is being used rather than LLC encoding.

2.1.5. Type 5 - SSAP (Source Service Access Point)

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

Defines a list of {operation, value} pairs used to match the 1-octet SSAP in the IEEE 802.2 LLC. Values are encoded as 1-octet quantities. op is encoded as specified in Section 4.2.1.1 of [RFC8955]. The match fails if EtherType L2 header encoding is being used rather than LLC encoding.

2.1.6. Type 6 - Control field in LLC

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

Defines a list of {operation, value} pairs used to match the 1-octet control field in the IEEE 802.2 LLC. Values are encoded as 1-octet quantities. op is encoded as specified in Section 4.2.1.1 of [RFC8955]. The match fails if EtherType L2 header encoding is being used rather than LLC encoding.

2.1.7. Type 7 - SNAP

Encoding: <type (1 octet), length (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 8-octet quantities with the zero padded SNAP left justified. op is encoded as specified in Section 4.2.1.1 of [RFC8955]. The match fails if EtherType L2 header encoding is being used rather than LLC encoding.

2.1.8. Type 8 - VLAN ID

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

Defines a list of {operation, value} pairs used to match VLAN ID. Values are encoded as 2-octet quantities, where the four most significant bits are set to zero and ignored for matching and the 12 least significant bits contain the VLAN value. op is encoded as specified in Section 4.2.1.1 of [RFC8955].

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

2.1.9. Type 9 - VLAN PCP

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

Defines a list of {operation, value} pairs used to match 3-bit VLAN PCP (priority code point) fields [IEEE802.1Q]. Values are encoded using a single octet, where the five most significant bits are set to zero and ignored for matching and the three least significant bits contain the VLAN PCP value. op is encoded as specified in Section 4.2.1.1 of [RFC8955].

In the virtual local-area network (VLAN) stacking case, the VLAN PCP is part of the outer VLAN tag.

2.1.10. Type 10 - Inner VLAN ID

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

Defines a list of {operation, value} pairs used to match the inner VLAN ID for virtual local-area network (VLAN) stacking or Q-in-Q use. Values are encoded as 2-octet quantities, where the four most significant bits are set to zero and ignored for matching and the 12 least significant bits contain the VLAN value. op is encoded as specified in Section 4.2.1.1 of [RFC8955].

In the single VLAN case, this component type MUST NOT be used. If it appears the match will fail.

2.1.11. Type 11 - Inner VLAN PCP

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

Defines a list of {operation, value} pairs used to match 3-bit inner VLAN PCP fields [IEEE802.1Q] for virtual local-area network (VLAN) stacking or Q-in-Q use. Values are encoded using a single octet, where the five most significant bits are set to zero and ignored for matching and the three least significant bits contain the VLAN PCP value. op is encoded as specified in Section 4.2.1.1 of [RFC8955].

In the single VLAN case, this component type MUST NOT be used. If it appears the match will fail.

2.1.12. Type 12 - VLAN DEI

Encoding: <type (1 octet), length (1 octet), op (1 octet)>

This type tests the DEI (Drop Eligible Indicator) bit in the VLAN tag. If op is zero, it matches if and only if the DEI bit is zero. If op is non-zero, it matches if and only if the DEI bit is one.

In the virtual local-area network (VLAN) stacking case, the VLAN DEI is part of the outer VLAN tag.

2.1.13. Type 13 - Inner VLAN DEI

Encoding: <type (1 octet), length (1 octet), op (1 octet)>

This type tests the DEI bit in the inner VLAN tag. If op is zero, it matches if and only if the DEI bit is zero. If op is non-zero, it matches if and only if the DEI bit is one.

In the single VLAN case, this component type MUST NOT be used. If it appears the match will fail.

2.1.14. Type 14 - Source MAC Special Bits

Encoding: <type (1 octet), length (1 octet), op (1 octet)>

This type tests the bottom nibble of the top octet of the Source MAC address. The two low order bits of that nibble have long been the local bit (0x2) and the group addressed bit (0x1). However, recent changes in IEEE 802 have divided the local address space into 4 quadrants specified by the next two bits (0x4 and 0x8) [RFC7042bis]. This flowspec component permits testing, for example, that a MAC is group addressed or is a local address in a particular quadrant. The encoding is as given in Section 4.2.1.2 of [RFC8955].

2.1.15. Type 15 - Destination MAC Special Bits

Encoding: <type (1 octet), length (1 octet), op (1 octet)>

As discussed in Section 2.1.14 but for the Destination MAC Address special bits.

2.2. Order of Traffic Filtering Rules

The existing rules in Section 5.1 of [RFC8955] and in [RFC8956] for the ordering of traffic filtering are extended as follows:

L2 flowspecs (AFI = 6, 25) take precedence over L3 flowspecs (AFI = 1, 2). Between two L2 flowspecs, precedence of the L2 portion is determined as specified in this section after this paragraph. If the L2 flowspec L2 portions are the same and the L3-AFI is nonzero, then the L3 portions are compared as specified in [RFC8955] or [RFC8956] as appropriate. Note: if the L3-AFI fields are different between two L2 flowspecs, they will never match the same packet so it will not be necessary to prioritize two flowspecs with different L3-AFI values.

The original definition for the order of traffic filtering rules can be reused for L2 with new consideration for the MAC Address offset. As long as the offsets are equal, the comparison is the same, retaining longest-prefix-match semantics. If the offsets are not equal, the lowest offset has precedence, as this flow matches the most significant bit.

Pseudocode:

```
flow_rule_L2_cmp (a, b)
{
    comp1 = next_component(a);
    comp2 = next_component(b);
    while (comp1 || comp2) {
        // component_type returns infinity on end-of-list
        if (component_type(comp1) < component_type(comp2)) {
            return A_HAS_PRECEDENCE;
        }
        if (component_type(comp1) > component_type(comp2)) {
            return B_HAS_PRECEDENCE;
        }

        if (component_type(comp1) == MAC_DESTINATION || MAC_SOURCE) {
            common = MIN(MAC Address length (comp1),
                         MAC Address length (comp2));
            cmp = MAC Address compare(comp1, comp2, common);
            // not equal, lowest value has precedence
            // equal, longest match has precedence
        } else {
            common =
               MIN(component_length(comp1), component_length(comp2));
            cmp = memcmp(data(comp1), data(comp2), common);
            // not equal, lowest value has precedence
            // equal, longest string has precedence
        }
    }
    return EQUAL;
}
```

3. L2VPN Flow Specification Encoding in BGP

The NLRI format for AFI=25/SAFI=134 (L2VPN), as with the other VPN flowspec AFI/SAFI pairs, is the same as the non-VPN Flow-Spec but with the addition of a Route Distinguisher to identify the VPN to which the flowspec is to be applied.

In addition, the IANA entry for SAFI 134 is slightly generalized as specified at the beginning of Section 6.

The L2VPN NLRI format is as follows:

```
+-------------------------------+
| total-length (0xnn or 0xfnnn) |  2 or 3 octets
+-------------------------------+
|      Route Distinguisher      |  8 octets
+-------------------------------+
|            L3-AFI             |  2 octets
+-------------------------------+
|  L2-length (0xnn or 0xfnnn)   |  2 or 3 octets
+-------------------------------+
|           NLRI-value          |  variable
+-------------------------------+
```

Figure 2: Flow Specification NLRI for L2VPN

The fields in Figure 2, other than the Route Distinguisher, are encoded as specified in Section 2 except that the minimum value for total-length is 12.

Flow specification rules received via this NLRI apply only to traffic that belongs to the VPN instance(s) into which it is imported. Flow rules are accepted as specified in Section 5.

3.1. Order of L2VPN Filtering Rules

The order between L2VPN filtering rules is determined as specified in Section 2.2. Note that if the Route Distinguisher is different between two L2VPN filtering rules, they will never both match the same packet so they need not be prioritized.

4. Ethernet Flow Specification Traffic Actions

The default action for an L2 traffic filtering flowspec is to accept traffic that matches that particular rule. The following extended community values per [RFC8955] can be used to specify particular actions in an L2 VPN network:

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

Table 1

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 specifies the following BGP extended communities for Ethernet flows to extend [RFC8955]:

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

Table 2

4.1. VLAN-action

The VLAN-action extended community, as shown in the diagram below, consists of 6 octets that include action Flags, two VLAN IDs, and the associated PCP and DEI values. The action Flags fields are further divided into two parts which correspond to the first action and the second action respectively. Bit 0 to bit 7 give the first action while bit 8 to bit 15 give the second action. The bits of PO, PU, SW, RI and RO in each part represent the action of Pop, Push, Swap, Rewrite inner VLAN and Rewrite outer VLAN respectively. Through this method, more complicated actions also can be represented in a single VLAN-action extended community, such as SwapPop, PushSwap, etc. For example, SwapPop action is the sequence of two actions, the first action is Swap and the second action is Pop.

```
  0   1   2   3   4   5   6   7   8   9   10  11  12  13  14  15
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|PO1|PU1|SW1|RI1|RO1| Resv      |PO2|PU2|SW2|RI2|RO2| Resv      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| VLAN ID1                                      | PCP1      |DE1|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| VLAN ID2                                      | PCP2      |DE2|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

PO1: Pop action. If the PO1 flag is one, it indicates the outmost VLAN should be removed.

PU1: Push action. If PU1 is one, it indicates VLAN ID1 will be added, the associated PCP and DEI are PCP1 and DE1.

SW1: Swap action. If the SW1 flag is one, it indicates the outer VLAN and inner VLAN should be swapped.

PO2: Pop action. If the PO2 flag is one, it indicates the outmost VLAN should be removed.

PU2: Push action. If PU2 is one, it indicates VLAN ID2 will be added, the associated PCP and DEI are PCP2 and DE2.

SW2:
Swap action. If the SW2 flag is one, it indicates the outer VLAN and inner VLAN should be swapped.

RI1 and RI2: Rewrite inner VLAN action. If the RIx flag is one (where "x" is "1" or "2"), it indicates the inner VLAN should be replaced by a new VLAN where the new VLAN is VLAN IDx and the associated PCP and DEI are PCPx and DEx. If the VLAN IDx is 0, the action is to only modify the PCP and DEI value of the inner VLAN.

RO1 and RO2: Rewrite outer VLAN action. If the ROx flag is one (where "x" is "1" or "2"), it indicates the outer VLAN should be replaced by a new VLAN where the new VLAN is VLAN IDx and the associated PCP and DEI are PCPx and DEx. If the VLAN IDx is 0, the action is to only modify the PCP and DEI value of the outer VLAN.

Resv: Reserved for future use. MUST be sent as zero and ignored on receipt.

Giving an example below: if the action of PUSH Inner VLAN 10 with PCP value 5 and DEI value 0 and PUSH Outer VLAN 20 with PCP value 6 and DEI value 0 is needed, the format of the VLAN-action extended community is as follows:

```
  0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0 |1 |0 |0 |0 |0 |0 |0 |0 |1 |0 |0 |0 |0 |0 |0 |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| 10                                |1 |0 |1 |0 |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| 20                                |1 |1 |0 |0 |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

4.2. TPID-action

The TPID-action extended community consists of 6 octets which includes the fields of action Flags, TP ID1 and TP ID2.

```
  0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|TI|TO|                       Resv               |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                    TP ID1                      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                    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.

TO: Mapping outer TP ID action. If the TO 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. MUST be sent as zero and ignored on receipt.

5. Flow Spec Validation

Flow Specifications received over AFI=25/SAFI=134 are validated against routing reachability received over AFI=25/SAFI=128 as modified to conform to [RFC9117].

6. IANA Considerations

IANA is requested to change the description for SAFI 134 [RFC8955] to read as follows and to change the reference for it to [this document]:

134 VPN dissemination of flow specification rules

IANA is requested to create an L2 Flow Specification Component Type registry on the Flow Spec Component Types registries web page as follows:

Name: L2 Flow Specification Component Types
Reference: [this document]
Registration Procedures:

0	Reserved
1-127	Specification Required
128-255	First Come First Served

Initial contents:

type	Reference	description
0	[this document]	Reserved
1	[this document]	Ethernet Type
2	[this document]	Source MAC
3	[this document]	Destination MAC
4	[this document]	DSAP in LLC
5	[this document]	SSAP in LLC
6	[this document]	Control field in LLC
7	[this document]	SNAP
8	[this document]	VLAN ID
9	[this document]	VLAN PCP
10	[this document]	Inner VLAN ID
11	[this document]	Inner VLAN PCP
12	[this document]	VLAN DEI
13	[this document]	Inner VLAN DEI
14	[this document]	Source MAC Special Bits
15	[this document]	Destination MAC Special Bits
16-254	[this document]	unassigned
255	[this document]	Reserved

Table 3

IANA is requested to assign two values from the "BGP Extended Communities Type - extended, transitive" registry [suggested value provided in square brackets]:

Type value	Name	Reference
TBD1[0x080A]	Flow spec VLAN action	[this document]
TBD2[0x080B]	Flow spec TPID action	[this document]

Table 4

7. Security Considerations

For General BGP Flow Specification Security Considerations, see [RFC8955].

VLAN tagging identifies Layer 2 communities which are commonly expected to be isolated except when higher layer connection is provided, such as Layer 3 routing. Thus, the ability of the flowspec VLAN action to change the VLAN ID in a frame might compromise security.

8. 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, <https://www.rfc-editor.org/info/rfc2119>.

[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, <https://www.rfc-editor.org/info/rfc4271>.

[RFC4761] Kompella, K., Ed. and Y. Rekhter, Ed., "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", RFC 4761, DOI 10.17487/RFC4761, January 2007, <https://www.rfc-editor.org/info/rfc4761>.

[RFC4762] Lasserre, M., Ed. and V. Kompella, Ed., "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", RFC 4762, DOI 10.17487/RFC4762, January 2007, <https://www.rfc-editor.org/info/rfc4762>.

[RFC6074] Rosen, E., Davie, B., Radoaca, V., and W. Luo, "Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs)", RFC 6074, DOI 10.17487/RFC6074, January 2011, <https://www.rfc-editor.org/info/rfc6074>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8955] Loibl, C., Hares, S., Raszuk, R., McPherson, D., and M. Bacher, "Dissemination of Flow Specification Rules", RFC 8955, DOI 10.17487/RFC8955, December 2020, <https://www.rfc-editor.org/info/rfc8955>.

[RFC8956] Loibl, C., Ed., Raszuk, R., Ed., and S. Hares, Ed., "Dissemination of Flow Specification Rules for IPv6", RFC 8956, DOI 10.17487/RFC8956, December 2020, <https://www.rfc-editor.org/info/rfc8956>.

[RFC9117] Uttaro, J., Alcaide, J., Filsfils, C., Smith, D., and P. Mohapatra, "Revised Validation Procedure for BGP Flow Specifications", RFC 9117, DOI 10.17487/RFC9117, August 2021, <https://www.rfc-editor.org/info/rfc9117>.

9. Informative References

[IEEE802.1Q] IEEE 802, "IEEE Standard for Local and metropolitan area networks - Media Access Control (MAC) Bridges and Virtual Bridge Local Area Networks", IEE Std 802.1Q-2014, 3 November 2014.

[RFC7042bis]
Eastlake, D., Abley, J., and Y. Li, "OUI Registry Restructuring", work in Progress, 14 April 2023, <https://www.ietf.org/archive/id/draft-intarea-rfc7042bis-02.txt>.

[RFC7432] Sajassi, A., Ed., Aggarwal, R., Bitar, N., Isaac, A., Uttaro, J., Drake, J., and W. Henderickx, "BGP MPLS-Based Ethernet VPN", RFC 7432, DOI 10.17487/RFC7432, February 2015, <https://www.rfc-editor.org/info/rfc7432>.

Acknowledgements

The authors wish to acknowledge the important contributions and suggestions of the following:

Hannes Gredler, Xiaohu Xu, Zhenbin Li, Lucy Yong, and Feng Dong.

Contributors

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

Email: liangqiandeng@huawei.com

Authors' Addresses

Weiguo Hao
Huawei Technologies
101 Software Avenue
Nanjing
Jiangsu, 210012
China

Email: haoweiguo@huawei.com

Donald E. Eastlake, 3rd
Independent
2386 Panoramic Circle
Apopka, Florida 32703
United States of America

Phone: +1-508-333-2270
Email: d3e3e3@gmail.com

Stephane Litkowski

Cisco Systems, Inc.

Email: slitkows.ietf@gmail.com

Shunwan Zhuang
Huawei Technologies
Huawei Building, No.156 Beiqing Road
Beijing
100095
China

Email: zhuangshunwan@huawei.com