Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 11 November 2023

S. Peng
Z. Li
S. Fang
Huawei Technologies
Y. Cui
Tsinghua University

10 May 2023

Dissemination of BGP Flow Specification Rules for APN
draft-peng-apn-bgp-flowspec-03

Abstract

A BGP Flow Specification is an n-tuple consisting of several matching criteria that can be applied to IP traffic. Application-aware Networking (APN) is a framework, where APN data packets convey APN attribute including APN ID and/or APN Parameters. The dynamic Flow Spec mechanism for APN is designed for the new applications of traffic filtering in an APN domain as well as the traffic control and actions at the policy enforcement points in this domain. These applications require coordination among the ASes within a service provider.

This document specifies a new BGP Flow Spec Component Type in order to support APN traffic filtering. The match field is the APN ID. It also specifies traffic filtering actions to enable the creation of the APN ID in the outer tunnel encapsulation when matched to the corresponding Flow Spec rules.

Table of Contents

1. Introduction

A Flow Specification (Flow Spec) is an n-tuple consisting of several matching criteria that can be applied to IP traffic [RFC8955]. The Flow Spec conveys match conditions (each may include several components) which are encoded using MP_REACH_NLRI and MP_UNREACH_NLRI attributes [RFC4760], while the associated actions such as redirect and traffic marking are encoded in BGP Extended Communities [RFC4360][RFC5701]. The IPv4 NLRI component types and traffic filtering actions sub-types are described in [RFC8955], while the IPv6 related are described in [RFC8956]. [I-D.ietf-idr-flowspec-l2vpn] extends the flow-spec rules and actions for Ethernet Layer 2 and L2VPN. The corresponding (AFI, SAFI) pairs are defined by IANA, respectively. [I-D.hares-idr-flowspec-v2] specifies BGP Flow Specification Version 2.

Application-aware Networking (APN) is introduced in [I-D.li-apn-framework] and [I-D.li-apn-problem-statement-usecases]. APN data packets convey the APN attribute (incl. APN ID and/or APN Parameters). The APN ID is a structured value, treated as an opaque object in the network, to which the network operator applies policies in various nodes/service functions along the path so to provide corresponding services. For an IPv6 network, a design proposal of such structured value is provided by [I-D.li-apn-header][I-D.li-apn-ipv6-encap]. The APN attribute can be encapsulated in various data planes adopted within a Network Operator controlled limited domain, e.g. IPv6, MPLS, and other tunnel technologies, which wait to be further specified.

With APN, it becomes possible to apply various policies in different nodes along a network path onto a traffic flow overall in a more efficient way, that is, at the headend to steer into corresponding path, at the midpoint to collect corresponding performance measurement data, and at the service function to execute particular policies. Prior to APN, there was no efficient way to realize this composite network service provisioning along the path.

This document specifies a new BGP Flow Spec Component Type to support APN traffic filtering. The match field is the APN APN ID [I-D.li-apn-framework]. It also specifies traffic filtering actions to enable the creation of the APN ID in the outer tunnel encapsulation when matched to the corresponding Flow Spec rules.

Depends upon specific deployment requirements, the functions specified in this draft can also be applied on BGP Flow Specification Version 2, which will be specified in the future versions.

2. Requirements Language

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 RFC 2119 [RFC2119] RFC 8174 [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Terminologies

APN: Application-aware Networking

APN ID: APN Identifier

AS: Autonomous System

Flow Spec: Flow Specification

BGP-FS: Border Gateway Protocol (BGP) Flow Specification (FS)

4. Flow Specifications for APN

The APN framework is introduced in [I-D.li-apn-framework]. The Flow Spec for APN is shown in Figure 1, that is, the Controller is used to set up BGP connection with the policy enforcement points in an APN domain.

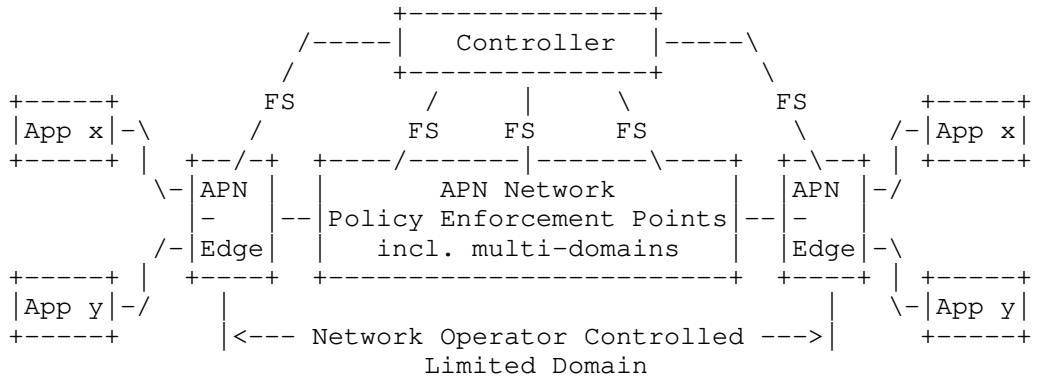

Figure 1. Flow Spec for APN

5. Component Type for APN

The IPv4 NLRI component types are defined in [RFC8955], while the IPv6 related are specified in [RFC8956]. This document defines a new component type for APN.

5.1. APN ID - Type TBD1

Encoding: <type (1 octet), length (1 octet), mask (variable), APN ID (variable)>

Defines the APN ID to match. The mask is used to indicate the bits of the APN ID carried in the packet which are used to match against the APN ID value in this Flow Spec component.

type (1 octet): This indicates the new component type TBD1.

length (1 octet): This indicates the length of the mask and the length of the APN ID. The mask and the APN ID have the same length.

mask (variable): This indicates the bits of the APN ID carried in the data packet which are used to match.

APN ID (variable): This indicates the APN ID that is used for the match.

5.2. Encoding Example

Since the APN ID is a structured value, the mask in the Flow Spec is used to enable flexible matching of the particular parts of the APN ID.

As an example, shown in Figure 2, the APN ID in the data packet contains two parts, the APP Group ID (0x300A) and User Group ID (0x0C08). In the Flow Spec, the mask is 0xFFFF0000 and the APN ID is 0x300A0000. Processing the match of the APN ID component is done by using the mask (0xFFFF0000) to indicate the bits of the APN ID carried in the packet to be matched against the one carried in the Flow Spec (0x300A0000). The result of this example is a successful match.

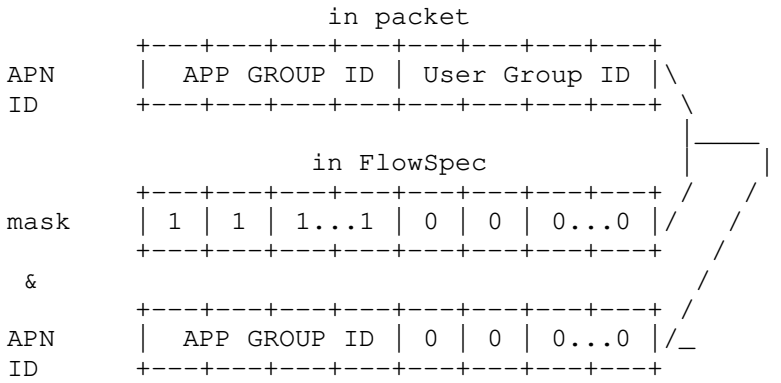

Figure 2. The match example of the APN ID component

6. Traffic Filtering

Traffic filtering policies have been traditionally considered to be relatively static. The dynamic Flow Spec mechanism for APN is designed for the new applications of traffic filtering in an APN domain as well as the traffic control and actions on the policy enforcement points in this domain. These applications require coordination among the ASes within a service provider. The new component and encoding are defined in Section 4. The actions are defined in this section.

6.1. Ordering of Flow Specifications

More than one Flow Specification rule may match a particular traffic flow at a node. The co-existing rules are mixed and need to be effectively organized. However, there is still no efficient way to achieve such classification. Thus, it is necessary to specify the grouping mechanism for the Flow Specification rules to be matched in a desired order as well as the actions being applied to a particular traffic flow. This ordering function is such that it does not depend on the arrival order of the Flow Specification via BGP and thus is consistent in the network [RFC8955].

The definition of this ordering is very important to the Flow Spec for APN because of the following reasons.

1. There can be other co-existing Flow Spec rules (e.g. based on 5-tuple) rather than only APN Flow Spec rules (i.e. based on APN ID).

2. The different parts of the APN ID can be determined by the different Flow Spec rules.

Therefore, the ordering of the Flow Spec rules for APN needs to be clearly specified.

6.2. Encoding format of the Grouping Identifier Extend Community Sub-Type TBD2

We define a Grouping Identifier Opaque Extend Community [RFC4360] (Sub-Type = TBD2) carrying both Group ID (2 octets) and Sub-group ID (2 octets) and indicating the grouping of the Flow Spec rules it accompanies.

The encoding format of the Grouping Identifier Opaque Extend Community is as follows.

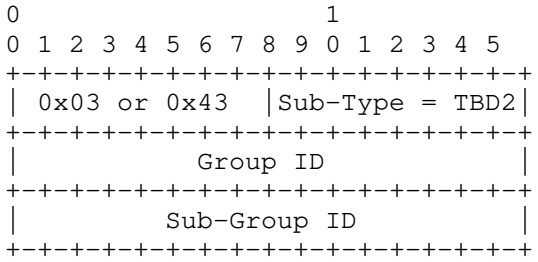

Figure 3: Encoding of the Grouping Identifier Extend Community

6.3. Usage Principles

The following principles are defined.

1. Within a sub-group, the order is the same as the previously defined.

* If the traffic-action Extended Community is carried and the Terminal Action (T, bit 47) [RFC8955] is not set, when one condition in this sub-group is matched, the evaluation of any subsequent flow specifications within this sub-group stops; if T is set, then the evaluation continues;

* If the traffic-action Extended Community is not carried, when one condition in this sub-group is matched, the evaluation of any subsequent flow specifications within this sub-group stops;

2. Between sub-groups, the sub-group is ordered by increasing Sub-group ID, when the evaluation in one sub-group stops or finishes, it will start the evaluation in the following sub-group if there are any sub-groups.

3. Between groups, the group is ordered by increasing Group ID, if at least one condition in this group is matched, when the evaluation of the flow specifications within the group reaches the end, the evaluation stops so the evaluation of the following group(s) will not start.

6.4. Usage example

At the APN Edge where the APN ID is created based on the Flow Specifications and encapsulated in the outer tunnel header [I-D.li-apn-framework], more than one Flow Specification rule condition may match a particular traffic flow. The different parts of the APN ID can be determined by the different Flow Spec rules. For example, as shown in Figure 4, the App Group ID is created by matching the 5-tuple components (e.g. destination IP address and transport layer ports), the User Group ID is created by matching the access ports, and the Reserved (R.) Group ID is created by matching the 5-tuple components.

Moreover, there are also other co-existing Flow Spec rules mixed at the node rather than only APN Flow Spec rules (i.e. based on APN ID). All the rules need to be effectively organized and applied to the particular traffic flow in a desired order.

In Figure 4, the Flow Specification rules for APN and other existing rules are categorized into two groups, and given Group ID = 1 and 2, respectively. The Flow Specification rules for creating different parts of the APN ID are categorized into three sub-groups, and given Sub-Group ID = 1, 2, and 3, respectively.

Based on the usage principles described in the above section, for the case of APN as shown in Figure 4, the usage principles are as follows,

1. Within a sub-group, the order is the same as the previously defined.

* If the traffic-action Extended Community is carried and the Terminal Action (T, bit 47) [RFC8955] is not set, when one condition in this sub-group is matched, the evaluation of any subsequent flow specifications within this sub-group stops and the App Group ID is created; if T is set, then the evaluation continues and the App Group ID will be created if there is a match within this sub-group;

* If the traffic-action Extended Community is not carried, when one condition in this sub-group is matched, the evaluation of any subsequent flow specifications within this sub-group stops and the App Group ID is created;

2. Between sub-groups, the sub-group is ordered with Sub-group ID, when the evaluation in the Sub-group ID = 1 stops or finishes, it will start the evaluation in the following Sub-group ID = 2 and create the User Group ID if matched, and then the Sub-group ID = 3 to create the R. Group ID if matched.

3. Between groups, the group is ordered with Group ID, if at least one condition in this Group ID = 1 is matched, when the evaluation of the flow specifications within the group reaches the end, the evaluation stops and the APN ID is created. The evaluation of the following group(s) will not start, that is, the Group ID = 0 will not be evaluated.

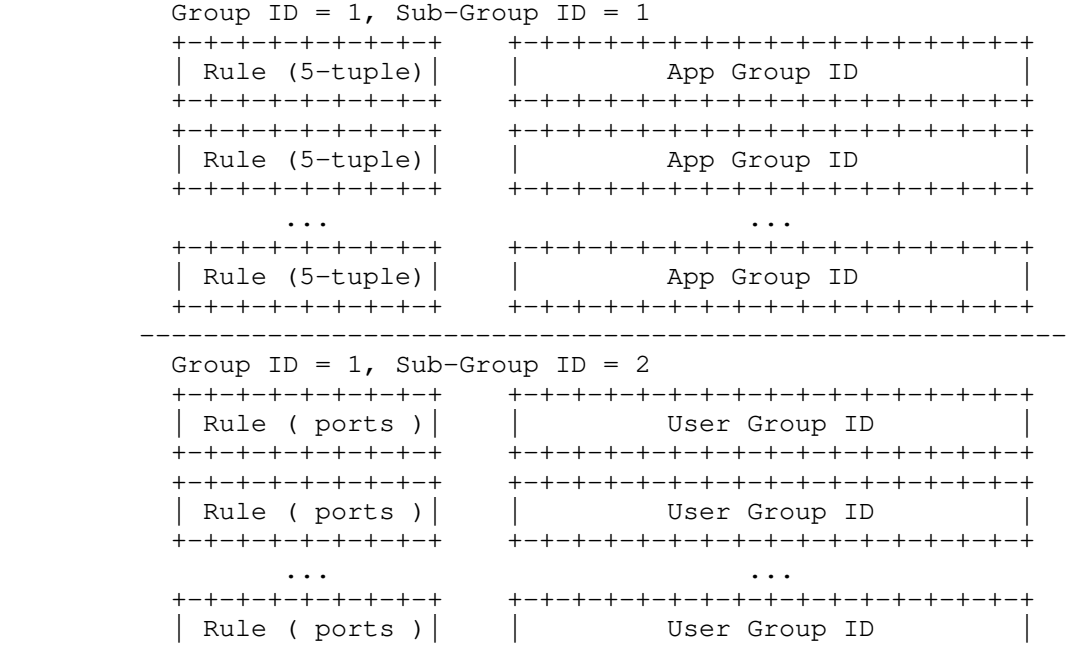

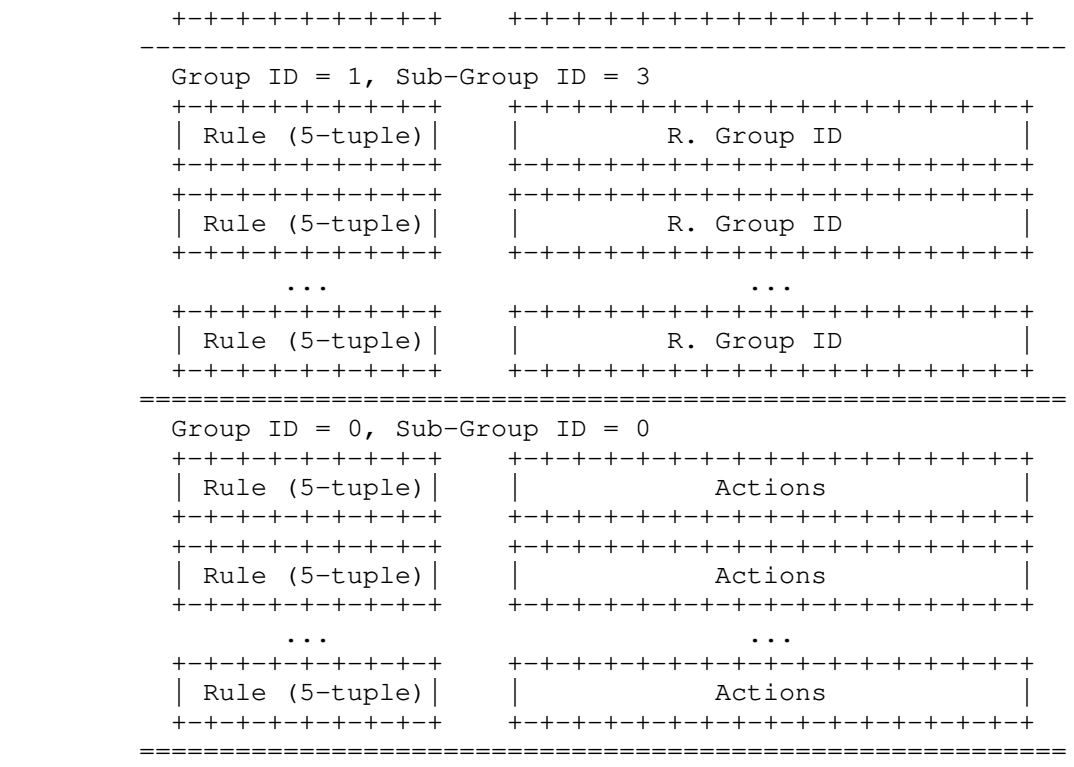

Figure 4: Usage of Grouping Identifer Extended Community for APN

7. Traffic Filtering Actions

Community 0xttss Sub-Type	action	encoding
TBD3	traffic-marking-apn (Section 7.1)	4/16-octet APN ID 1-octet IPv6 (ExH) Type 1-octet Reserved
TBD4	traffic-marking-apn-partial (Section 7.2)	4/16-octet Bitmask 4/16-octet APN ID 1-octet IPv6 (ExH) Type 1-octet Reserved
TBD5	inherit-apn (Section 7.3)	4/16-octet Bitmask 1-octet IPv6 (ExH) Type 1-octet Reserved
TBD6	stitch-apn (Section 7.2)	4/16-octet Bitmask 4/16-octet APN ID 1-octet IPv6 (ExH) Type 1-octet Reserved

7.1. Traffic Marking (traffic-marking-apn) Sub-Type TBD3

The traffic-marking-apn Extended Community instructs a system to create the APN ID and encapsulate it in the indicated outer tunnel header of a transiting IP packet.

In this case, the tunnel encapsulation header is IPv6, possibly followed by an extension header (ExH). The corresponding Extended Community [RFC5701] is encoded as follows:

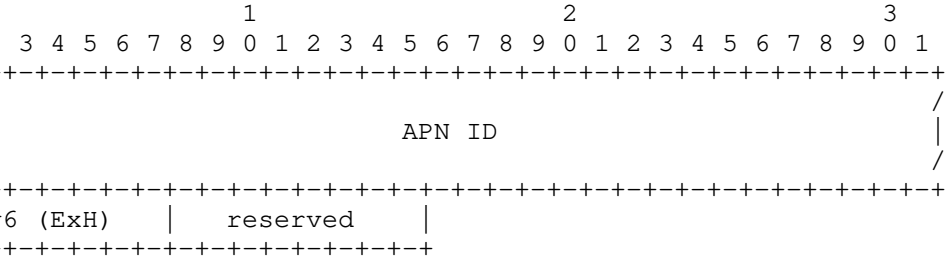

APN ID: 4/16 octets, APN ID value to be created and encapsulated in the indicated outer tunnel header of the transiting IP packet.

IPv6 (ExH): 1 octet, the type of each IPv6 extension header [RFC8200][RFC2780][RFC5871] is directly reused to indicate the outer tunnel to be used to encapsulate the APN ID.

reserved: 1 octet, MUST be set to 0 on encoding and MUST be ignored during decoding.

7.2. Traffic Marking (traffic-marking-apn-partial) Sub-Type TBD4

The traffic-marking-apn-partial Extended Community instructs a system to use the bitmask indicating the bits of the APN ID to be encapsulated in the indicated outer tunnel header of a transiting IP packet. The ultimately constructed APN ID may comprise of several parts obtained by the matches of different rules, and it is encapsulated in the indicated outer tunnel header.

In this case, the tunnel encapsulation header is IPv6, possibly followed by an extension header (ExH). The corresponding Extended Community [RFC5701] is encoded as follows:

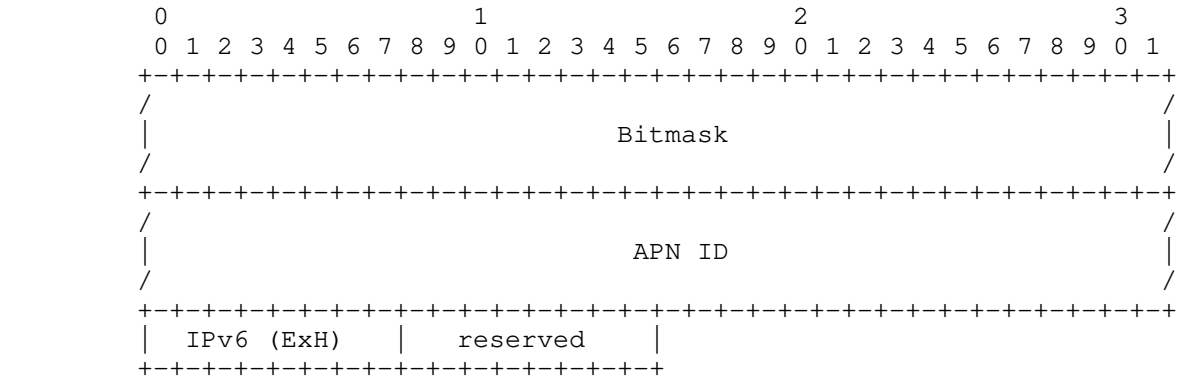

Bitmask: 4/16 octets, the same length as the APN ID, indicating the bits of the APN ID to be encapsulated in the indicated outer tunnel header.

APN ID: 4/16 octets, the APN ID value to be created and encapsulated in the indicated outer tunnel header of the transiting IP packet.

IPv6 (ExH): 1 octet, the type of each IPv6 extension header [RFC8200][RFC2780][RFC5871] is directly reused to indicate the outer tunnel to be used to encapsulate the APN ID.

reserved: 1 octet, MUST be set to 0 on encoding and MUST be ignored during decoding.

7.3. Inherit (inherit-apn) Sub-Type TBD5

The inherit-apn Extended Community instructs a system to use the Bitmask to "and" operate on the existing APN ID of a transiting IP packet and encapsulate the inherited APN ID in the indicated outer tunnel header.

In this case, the tunnel encapsulation header is IPv6, possibly followed by an extension header (ExH). The corresponding Extended Community [RFC5701] is encoded as follows:

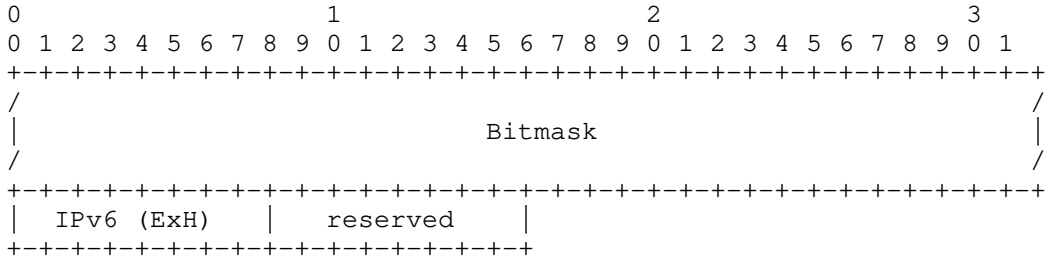

Bitmask: 4/16 octets, the same length as the APN ID, to "and" operate on the existing APN ID of a transiting IP packet.

IPv6 (ExH): 1 octet, the type of each IPv6 extension header [RFC8200][RFC2780][RFC5871] is directly reused to indicate the outer tunnel to be used to encapsulate the APN ID.

reserved: 1 octet, MUST be set to 0 on encoding and MUST be ignored during decoding.

7.4. Stitch (stitch-apn) Sub-Type TBD6

The stitch-apn Extended Community instructs a system to "and" the Bitmask with the existing APN ID of a transiting IP packet to get the part to be further encapsulated, and "and" the negation of the Bitmask with the APN ID in the Flow Spec and get the other part to be further encapsulated. The stitched APN ID is encapsulated in the indicated outer tunnel header. That is to say, the Bitmask specifies

the bits of the received APN ID to be replaced by the corresponding bits from the APN ID in the action sub-TLV value to produce a new outer APN ID. The other bits of the received APN ID are copied to the new outer AP ID.

In this case, the tunnel encapsulation header is IPv6, possibly followed by an extension header (ExH). The corresponding Extended Community [RFC5701] is encoded as follows:

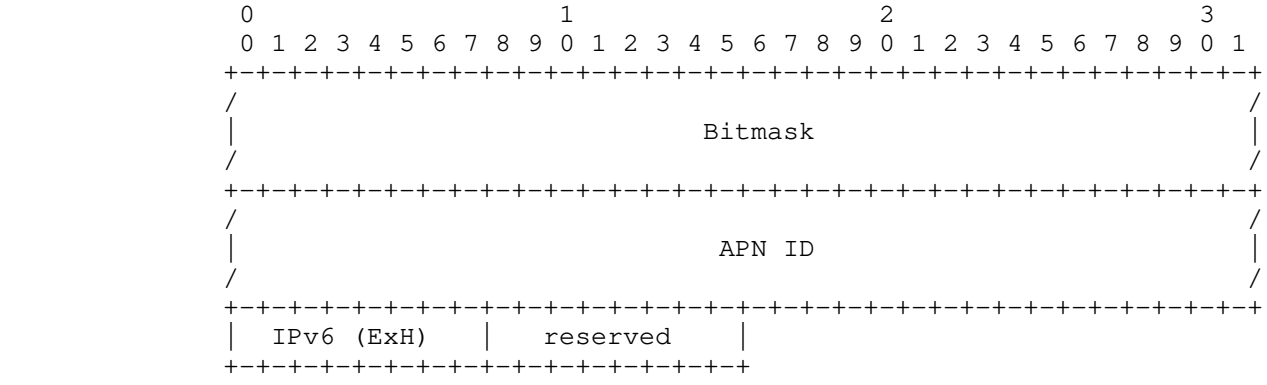

Bitmask: 4/16 octets, the same length as the APN ID, used to operate on the APN ID (both carried in the transiting IP packet and in the Flow Spec).

APN ID: 4/16 octets, the APN ID value to be created and encapsulated in the indicated outer tunnel header of the transiting IP packet.

IPv6 (ExH): 1 octet, the type of each IPv6 extension header [RFC8200][RFC2780][RFC5871] is directly reused to indicate the outer tunnel to be used to encapsulate the APN ID.

reserved: 1 octet, MUST be set to 0 on encoding and MUST be ignored during decoding.

8. IANA Considerations

8.1. Flow Spec Component - APN ID

IANA is requested to assign a value in the Flow Specification Component Types Registry as follows:

Value	Name	Reference
TBD1	APN ID	This document

8.2. Opaque Extended Community - Grouping Identifier

The Grouping Identifier Opaque Extended Community is defined in this document and it is requested that a Sub-Type = TBD2 be assigned as follows.

Value	Name	Reference
TBD2	Grouping Identifier	This document

8.3. Extended Community Flow Specification Actions

The Extended Community Flow Specification Actions are defined in this document and it is requested that corresponding Sub-Types as shown in the following table be assigned.

Sub-Type Value	Name	Reference
TBD3	traffic-marking-apn	This document
TBD4	traffic-marking-apn-partial	This document
TBD5	inherit-apn	This document
TBD6	stitch-apn	This document

9. Acknowledgements

The authors would like to thank the careful reviews and valuable comments from Haibo Wang, Shunwan Zhuang, Stefano Previdi, and Donald Eastlake.

10. Security Considerations

The security considerations are the same as [RFC8955], [RFC8956], and [I-D.li-apn-framework].

11. References

11.1. Normative References

[I-D.hares-idr-flowspec-v2]
Hares, S., Eastlake, D. E., Yadlapalli, C., and S. Maduschke, "BGP Flow Specification Version 2", Work in Progress, Internet-Draft, draft-hares-idr-flowspec-v2-05, 4 February 2022, <https://datatracker.ietf.org/doc/html/draft-hares-idr-flowspec-v2-05>.

[I-D.li-apn-framework]
Li, Z., Peng, S., Voyer, D., Li, C., Liu, P., Cao, C., and G. S. Mishra, "Application-aware Networking (APN) Framework", Work in Progress, Internet-Draft, draft-li-apn-framework-07, 3 April 2023, <https://datatracker.ietf.org/doc/html/draft-li-apn-framework-07>.

[I-D.li-apn-header]
Li, Z., Peng, S., and S. Zhang, "Application-aware Networking (APN) Header", Work in Progress, Internet-Draft, draft-li-apn-header-04, 12 April 2023, <https://datatracker.ietf.org/doc/html/draft-li-apn-header-04>.

[I-D.li-apn-ipv6-encap]
Li, Z., Peng, S., and C. Xie, "Application-aware IPv6 Networking (APN6) Encapsulation", Work in Progress, Internet-Draft, draft-li-apn-ipv6-encap-06, 9 December 2022, <https://datatracker.ietf.org/doc/html/draft-li-apn-ipv6-encap-06>.

[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>.

[RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", RFC 4360, DOI 10.17487/RFC4360, February 2006, <https://www.rfc-editor.org/info/rfc4360>.

[RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, DOI 10.17487/RFC4760, January 2007, <https://www.rfc-editor.org/info/rfc4760>.

[RFC5701] Rekhter, Y., "IPv6 Address Specific BGP Extended Community Attribute", RFC 5701, DOI 10.17487/RFC5701, November 2009, <https://www.rfc-editor.org/info/rfc5701>.

[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>.

11.2. Informative References

[I-D.ietf-idr-flowspec-l2vpn]
Weiguo, H., Eastlake, D. E., Litkowski, S., and S. Zhuang, "BGP Dissemination of L2 Flow Specification Rules", Work in Progress, Internet-Draft, draft-ietf-idr-flowspec-l2vpn-21, 24 April 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-flowspec-l2vpn-21>.

[I-D.li-apn-problem-statement-usecases]
Li, Z., Peng, S., Voyer, D., Xie, C., Liu, P., Qin, Z., and G. S. Mishra, "Problem Statement and Use Cases of Application-aware Networking (APN)", Work in Progress, Internet-Draft, draft-li-apn-problem-statement-usecases-08, 3 April 2023, <https://datatracker.ietf.org/doc/html/draft-li-apn-problem-statement-usecases-08>.

[RFC2780] Bradner, S. and V. Paxson, "IANA Allocation Guidelines For Values In the Internet Protocol and Related Headers", BCP 37, RFC 2780, DOI 10.17487/RFC2780, March 2000, <https://www.rfc-editor.org/info/rfc2780>.

[RFC5871] Arkko, J. and S. Bradner, "IANA Allocation Guidelines for the IPv6 Routing Header", RFC 5871, DOI 10.17487/RFC5871, May 2010, <https://www.rfc-editor.org/info/rfc5871>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <https://www.rfc-editor.org/info/rfc8200>.

Authors' Addresses

Shuping Peng
Huawei Technologies
Beijing
China
Email: pengshuping@huawei.com

Zhenbin Li
Huawei Technologies
Beijing
China
Email: lizhenbin@huawei.com

Sheng Fang
Huawei Technologies
Beijing
China
Email: fangsheng@huawei.com

Yong Cui
Tsinghua University
Beijing
China
Email: cuiyong@tsinghua.edu.cn