

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 14 May 2026

N. Geng
Z. Li
Huawei Technologies
G. Mishra
Verizon Inc.
10 November 2025

APN Scope and Gap Analysis
draft-geng-apn-scope-gap-analysis-00

Abstract

The APN work in IETF is focused on developing a framework and set of mechanisms to derive, convey and use an attribute allowing the implementation of fine-grain user group-level and application group-level requirements in the network layer. APN aims to apply various policies in different nodes along a network path onto a traffic flow altogether, for example, 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. Currently there is still no way to efficiently realize this composite network service provisioning along the path. This document further clarifies the scope of the APN work and describes the solution gap analysis.

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 <https://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 14 May 2026.

Copyright Notice

Copyright (c) 2025 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
2. Requirements Language	3
3. Terminologies	3
4. APN Framework and Scope	3
5. Example Use Case and Existing Issues	5
6. Basic Solution and Benefits	6
7. Solution Gap Analysis	7
7.1. IPv6/MPLS Flow Label	7
7.2. Subscriber Identifier and Performance Policy Identifier in NSH	8
7.3. IOAM Flow ID	9
7.4. Binding SID	9
7.5. FlowSpec Label	9
7.6. Group Policy ID	9
7.7. Detnet Flow Identification	9
7.8. Network Slicing Resource ID	10
7.9. Service Path ID in NSH	10
7.10. DSCP	10
7.11. Summary	10
8. IANA Considerations	11
9. Acknowledgements	11
10. Informative References	11
Authors' Addresses	15

1. Introduction

Application-aware Networking (APN) is introduced in [I-D.li-apn-framework] and [I-D.li-apn-problem-statement-usecases]. APN conveys an attribute along with data packets into network and makes the network aware about data flow requirements at different granularity levels.

Such an attribute is acquired, constructed in a structured value, and then encapsulated in the packet. Such structured value is treated as an opaque object in the network to which the network operator applies policies in various nodes/service functions along the path and provides corresponding services.

This structured attribute can be encapsulated in various data planes adopted within a Network Operator controlled limited domain, e.g. MPLS, VXLAN, SR/SRv6 and other tunnel technologies, which waits to be further specified.

With APN, it becomes possible to apply various policies in different nodes along a network path onto a traffic flow altogether in a more efficient way, e.g., 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. Currently there is still no way to realize this composite network service provisioning along the path very efficiently. It may be possible to stack those various policies in a list of TLVs at the headend. However, this approach would introduce great complexities and impose big challenges on the hardware processing and forwarding.

The example use-case presented in this draft further expands on the rationale for such an attribute and how it can be derived and used in that specific context.

This document further clarifies the scope of the APN work and describes the solution gap analysis.

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

CPE: Customer Premises Equipment

DPI: Deep Packet Inspection

OS: Operating System

4. APN Framework and Scope

The APN framework is introduced in [I-D.li-apn-framework], as shown in the Figure 1.

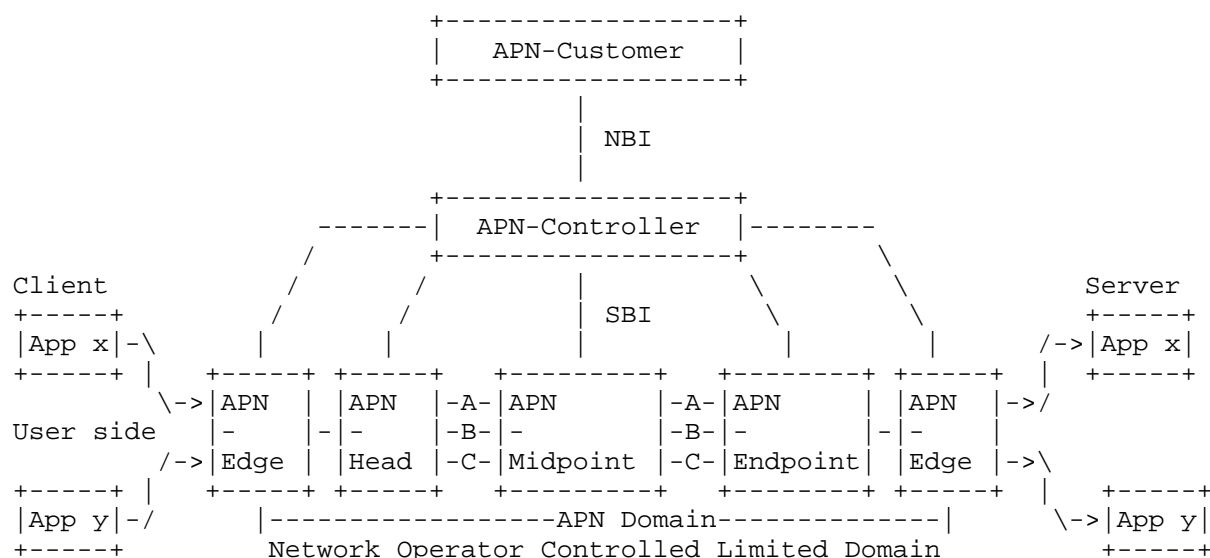


Figure 1. APN Framework and Scope

APN is only applied to an edge-to-edge tunnel encapsulation within a limited trusted domain. It means that the source and destination addresses of the packet are the endpoints of the tunnel (i.e. the domain edges), and nothing about the payload source and destination can be deduced, which substantially reduces the privacy concerns. Typically, an APN domain is defined as a Network Operator controlled limited domain (see Figure 1), in which MPLS, VXLAN, SR/Srv6 and other tunnel technologies are adopted to provide network services.

With APN, the attribute is acquired based on the existing information in the packet header (i.e. source and destination addresses, incoming L2 (or) MPLS encapsulation, incoming physical/virtual port information, the other fields of the 5-tuple if they are not encrypted) at the edge devices of the APN domain, added to the data packets along with the tunnel encapsulation, and delivered to the network, wherein, according to this attribute, corresponding network services are provisioned. When the packets leave the APN domain, the attribute is removed together with the tunnel encapsulation header.

5. Example Use Case and Existing Issues

To be more specific and more concrete, here we use SD-WAN as an example use case to further expand on the rationale for such attribute and how it can be derived and used in that specific context.

In the case of SD-WAN, an enterprise obtains WAN services from an SD-WAN provider so that its employees have access to the applications in the Cloud, and then the SD-WAN provider may buy WAN lines from a Network Operator. The enterprise may know what applications will use the SD-WAN services, but it will only provide the 5 tuples (i.e. source IP address, source port, destination IP address, destination port, transport protocol) of those applications to the SD-WAN provider. So, the SD-WAN provider does not know what applications it is serving, and will only provide 5 tuples to the Network Operator and the service performance requirements for steering their customer's traffic. In this way, the Network Operator does not know anything else about the traffic except the 5 tuples and requirements. Nowadays, SD-WAN is usually using 5-tuple to steer the traffic into corresponding WAN lines across the Network Operator's network [SD-WAN].

However, there are two main issues in the current SD-WAN deployments.

1) It is complicated to resolve the 5 tuples. Even worse, as the traffic is encrypted, it becomes impossible to obtain any transport layer information. Moreover, in the IPv6 data plane, with the extension headers being added before the upper layer, in some implementations it becomes very difficult and even impossible to obtain transport layer information because that information is located deep in the packet. So, there is no 5 tuples anymore, and maybe only 2 tuples are available.

2) Currently there is still no way to apply various policies in different nodes along the network path onto a traffic flow altogether, 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. It may be possible to stack those various policies in a list of TLVs at the headend. However, this approach would introduce great complexities and impose big challenges on the hardware processing and forwarding.

6. Basic Solution and Benefits

With APN, at the edge node, i.e. CPE, of the SD-WAN (see Figure 2), the 5-tuple, plus information related to user or application group-level requirements is constructed into a structured value, called APN attribute. This attribute is only meaningful for the network operators to apply various policies in different nodes/service functions, which can be enforced from the Controllers.

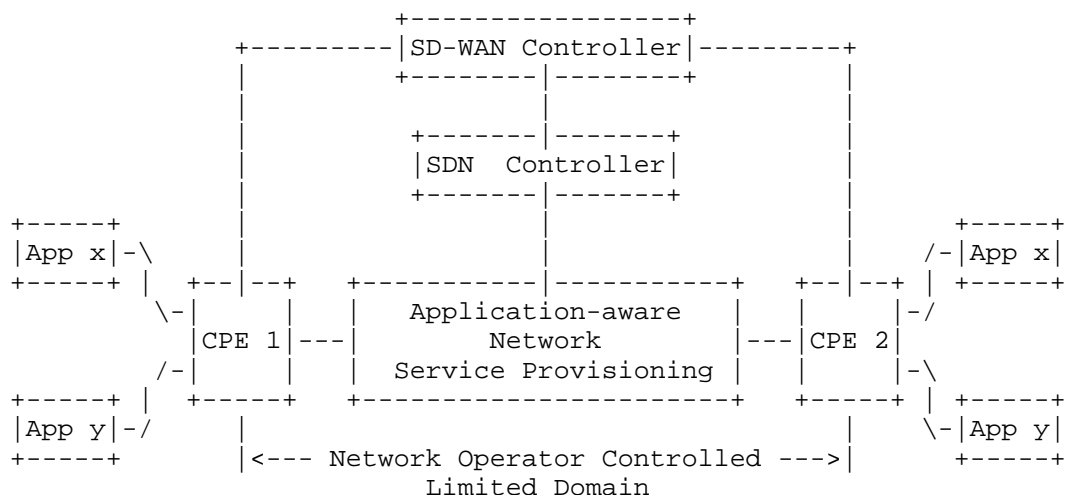


Figure 2. SD-WAN using the APN Framework

With such an attribute in the network, we can easily solve the two issues above-mentioned. For example, when the packet is sent from the CPE1 and the attribute is added along with the tunnel encapsulation, then it is not necessary to resolve the 5-tuple and perform the deep inspection in every node along the path. This attribute is encapsulated in the network layer and can be easily read by the routers and service functions. If the tunnel is based on the IPv6 data plane, for example, such an attribute can be encapsulated in an option of IPv6 hop-by-hop options header.

Since this attribute is taken as an object to the network, the network operators will simply place the policies in the nodes/service functions where this indicated traffic will go through, and the corresponding node/service function will just apply policies for this object. This can be easily done by utilizing this attribute, which is not possible with any current existing mechanism.

Such attribute will also bring other benefits, for example,

- * Improve the forwarding performance since it will only use 1 field in the IP layer instead of resolving 5 tuples, which will also improve the scalability.
- * Very flexible policy enforcement in various nodes and service functions along the network path.

Furthermore, with such attribute, more new services could be enabled, for example,

- * Even more fine-granularity performance measurement could be achieved and the granularity to be monitored and visualized can be controllable, which is able to relieve the processing pressure on the controller when it is facing the massive monitoring data.
- * The policy execution on the service function can be based only on this value and not based on 5-tuple, which can eliminate the need of deep packet inspection.
- * The underlay performance guarantee could be achieved for SD-WAN overlay services, such as explicit traffic engineering path satisfying SLA and selective visualized accurate performance measurement.

7. Solution Gap Analysis

There are already some solutions specified in IETF, which use identifier to perform traffic steering and service provisioning. However, the existing solutions are specific to a particular scenario or data plane. None of them is the same as APN and able to achieve the same effects.

7.1. IPv6/MPLS Flow Label

[RFC6437] specifies the IPv6 flow label which enables the IPv6 flow classification. However, the IPv6 flow label is mainly used for Equal Cost Multipath Routing (ECMP) and Link Aggregation [RFC6438].

Similarly, [RFC6391] describes a method of adding an additional Label Stack Entry (LSE) at the bottom of the stack in order to facilitate the load balancing of the flows within a pseudowire (PW) over the available ECMPs. A similar design for general MPLS use has also been proposed in [RFC6790] using the concept of Entropy Label.

7.2. Subscriber Identifier and Performance Policy Identifier in NSH

Subscriber Identifier and Performance Policy Identifier are specified in [RFC8979]. These identifiers are carried only in the Network Service Header (NSH) [RFC8300] Context Header, as shown in Figure 3, while the APN attribute can be carried in various data plane encapsulations.




Figure 3. Subscriber Identifier and Performance Policy Identifier

In this draft [RFC8979], the Subscriber Identifier carries an opaque local identifier that is assigned to a subscriber by a network operator, and the Performance Policy Identifier represents an opaque value pointing to specific performance policy to be enforced. In this way, in order to apply various policies in different nodes along the network path onto a traffic flow altogether, e.g., 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, those various policies would have to be stacked in a list of TLVs at the headend, introducing great complexities and big challenges on the hardware processing and forwarding.

The APN attribute is treated as an opaque object in the network, to which the network operator applies policies in various nodes/service functions along the path and provide corresponding services.

7.3. IOAM Flow ID

A 32-bit Flow ID is specified in [I-D.ietf-ippm-ioam-direct-export], which is used to correlate the exported data of the same flow from multiple nodes and from multiple packets, while the APN attribute can serve more various purposes.

7.4. Binding SID

The Binding SID (BSID) [RFC8402] is bound to an SR Policy, instantiation of which may involve a list of SIDs. Any packets received with an active segment equal to BSID are steered onto the bound SR Policy. A BSID may be either a local or a global SID. While the APN attribute is not bound to SR only, and it can be carried in various data plane encapsulations.

7.5. FlowSpec Label

The flow specification (FlowSpec) [RFC5575] is actually an n-tuple consisting of several matching criteria that can be applied to IP traffic, which include elements such as source and destination address prefixes, IP protocol, and transport protocol port numbers. In BGP VPN/MPLS networks, BGP FlowSpec can be extended to identify and change (push/swap/pop) the label(s) for traffic that matches a particular FlowSpec rule in [I-D.ietf-idr-flowspec-mpls-match] and [I-D.ietf-idr-bgp-flowspec-label]. In [I-D.liang-idr-bgp-flowspec-route], BGP is used to distribute the FlowSpec rule bound with label(s). While the APN attribute is not bound to MPLS only, and it can be carried in various data plane encapsulations.

7.6. Group Policy ID

The capabilities of the VXLAN-GPE protocol can be extended by defining next protocol "shim" headers that are used to implement new data plane functions. For example, Group Policy ID is carried in the Group-Based Policy (GBP) Shim header [I-D.lemon-vxlan-lisp-gpe-gbp]. GENEVE has similar ability as VXLAN-GPE to carry metadata.

7.7. Detnet Flow Identification

Identification and Specification of DetNet Flows is specified in [RFC9016]. DetNet MPLS flows can be identified and specified by the SLabel and the FLabelStack. The IP 6-tuple is used for DetNet IP flow identification, which consists of SourceIpAddress, DestinationIpAddress, Dscp, Protocol, SourcePort, and DestinationPort. IPv6FlowLabel and IPSecSpi are additional attributes that can be used for DetNet flow identification in

addition to the 6-tuple. Therefore, the Detnet IP Flow ID is logical and there is no such Flow ID carried for Detnet, but only the 6-tuple is directly used to identify the Detnet flows.

Only one exceptional case, in [I-D.ietf-spring-sr-redundancy-protection], the 32-bit flow identification (FID) identifies one specific Detnet flow of redundancy protection. This FID is usually allocated from centralized controller to the SR ingress node or redundancy node in SR network.

7.8. Network Slicing Resource ID

In [I-D.ietf-6man-enhanced-vpn-vtn-id], VTN Resource ID is a 4-octet identifier which uniquely identifies the set of network resources allocated to a VTN. For network slicing, the ID is used to indicate the network resources to be allocated to the network slices and it is not bound to any traffic flow.

APN is for traffic steering, while network slicing is about resource partition [I-D.ietf-teas-rfc3272bis].

7.9. Service Path ID in NSH

In [RFC8300], the Service Path Identifier (SPI) in the NSH uniquely identifies a Service Function Path (SFP). Participating nodes MUST use this identifier for SFP selection. The initial Classifier MUST set the appropriate SPI for a given classification result. For SFC, the ID is used to indicate a SF path and it is not bound to any traffic flow.

7.10. DSCP

DSCP [RFC4594] is mainly used for per hop QoS handling inside the network. It has 6 bits to set.

7.11. Summary

The comparison of the identifiers for the typical network services (incl. iOAM, Detnet, Network Slicing (NS), and Service Function Chaining (SFC)) is shown in the following Table from different aspects (incl. ID, Identification Object, Source (for generating the ID), Configuration (Conf.) node, and Size).

ID	Identification Object	Source	Conf. node	Size
APN	APN ID	The flow that needs fine-granular services	5-tuple Layer 2	Controller 32bits 128b
iOAM	Flow ID	The flow that needs performance monitoring	-	Controller Ingress 32bits
Detnet	Flow ID (6-tuple)	The flow that needs Detnet services	-	Controller -
Detnet	Flow ID	The redundant protection flow	-	Detnet Controller 32bits
NS	Resource ID	The network resources that are allocated to network slices	-	Controller 32bits
NSH	SPI	The SF Path	-	Controller 24bits
SFC	Performance Policy ID	The performance policy	-	Controller -
DSCP		Per-hop Behavior	-	6bits

Table 1. Comparison of the Identifiers

As driven by ever-emerging new 5G services, fine-granularity service provisioning becomes urgent. The existing solutions are either specific to a particular scenario or data plane. While APN aims to define a generalized attribute used for fine-granularity service provisioning, and can be carried in various data plane encapsulations.

8. IANA Considerations

There are no IANA considerations in this document.

9. Acknowledgements

The authors would like to acknowledge Martin Vigoureux, Alvaro Retana, Barry Leiba, Stefano Previdi, Adrian Farrel, and Daniel King for their valuable review and comments.

10. Informative References

[I-D.brockners-ippm-ioam-vxlan-gpe]

Brockners, F., Bhandari, S., Govindan, V. P., Pignataro, C., Gredler, H., Leddy, J., Youell, S., Mizrahi, T., Kfir, A., Gafni, B., Lapukhov, P., and M. Spiegel, "VXLAN-GPE Encapsulation for In Situ OAM (IOAM) Data", Work in Progress, Internet-Draft, draft-brockners-ippm-ioam-vxlan-gpe-05, 1 March 2024, <<https://datatracker.ietf.org/doc/html/draft-brockners-ippm-ioam-vxlan-gpe-05>>.

[I-D.ietf-6man-enhanced-vpn-vtn-id]

Dong, J., Li, Z., Xie, C., Ma, C., and G. S. Mishra, "Carrying Network Resource (NR) related Information in IPv6 Extension Header", Work in Progress, Internet-Draft, draft-ietf-6man-enhanced-vpn-vtn-id-13, 20 October 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-6man-enhanced-vpn-vtn-id-13>>.

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

liangqiandeng, Hares, S., You, J., Raszuk, R., and D. Ma, "Carrying Label Information for BGP FlowSpec", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-flowspec-label-02, 20 October 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-flowspec-label-02>>.

[I-D.ietf-idr-flowspec-mpls-match]

Yong, L., Hares, S., liangqiandeng, and J. You, "BGP Flow Specification Filter for MPLS Label", Work in Progress, Internet-Draft, draft-ietf-idr-flowspec-mpls-match-02, 20 October 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-flowspec-mpls-match-02>>.

[I-D.ietf-ippm-ioam-direct-export]

Song, H., Gafni, B., Brockners, F., Bhandari, S., and T. Mizrahi, "In Situ Operations, Administration, and Maintenance (IOAM) Direct Exporting", Work in Progress, Internet-Draft, draft-ietf-ippm-ioam-direct-export-11, 23 September 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-ippm-ioam-direct-export-11>>.

[I-D.ietf-sfc-serviceid-header]

Sarikaya, B., Von Hugo, D., and M. Boucadair, "Subscriber and Performance Policy Identifier Context Headers in the Network Service Header (NSH)", Work in Progress, Internet-Draft, draft-ietf-sfc-serviceid-header-14, 11 December 2020, <<https://datatracker.ietf.org/doc/html/draft-ietf-sfc-serviceid-header-14>>.

[I-D.ietf-spring-sr-redundancy-protection]

Geng, X., Chen, M., Camarillo, P., Mishra, G. S., Varga, B., and F. Fejes, "SRv6 for Redundancy Protection", Work in Progress, Internet-Draft, draft-ietf-spring-sr-redundancy-protection-05, 13 October 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-spring-sr-redundancy-protection-05>>.

[I-D.ietf-teas-rfc3272bis]

Farrel, A., "Overview and Principles of Internet Traffic Engineering", Work in Progress, Internet-Draft, draft-ietf-teas-rfc3272bis-27, 12 August 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-rfc3272bis-27>>.

[I-D.lemon-vxlan-lisp-gpe-gbp]

Lemon, J., Maino, F., Smith, M., and A. Isaac, "Group Policy Encoding with VXLAN-GPE and LISP-GPE", Work in Progress, Internet-Draft, draft-lemon-vxlan-lisp-gpe-gbp-02, 30 April 2019, <<https://datatracker.ietf.org/doc/html/draft-lemon-vxlan-lisp-gpe-gbp-02>>.

[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-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>>.

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

Liang, Q. and J. You, "BGP FlowSpec based Multi-dimensional Route Distribution", Work in Progress, Internet-Draft, draft-liang-idr-bgp-flowspec-route-00, 20 October 2014, <<https://datatracker.ietf.org/doc/html/draft-liang-idr-bgp-flowspec-route-00>>.

- [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>>.
- [RFC4594] Babiarz, J., Chan, K., and F. Baker, "Configuration Guidelines for DiffServ Service Classes", RFC 4594, DOI 10.17487/RFC4594, August 2006, <<https://www.rfc-editor.org/info/rfc4594>>.
- [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, <<https://www.rfc-editor.org/info/rfc5575>>.
- [RFC6391] Bryant, S., Ed., Filsfils, C., Drafz, U., Kompella, V., Regan, J., and S. Amante, "Flow-Aware Transport of Pseudowires over an MPLS Packet Switched Network", RFC 6391, DOI 10.17487/RFC6391, November 2011, <<https://www.rfc-editor.org/info/rfc6391>>.
- [RFC6437] Amante, S., Carpenter, B., Jiang, S., and J. Rajahalme, "IPv6 Flow Label Specification", RFC 6437, DOI 10.17487/RFC6437, November 2011, <<https://www.rfc-editor.org/info/rfc6437>>.
- [RFC6438] Carpenter, B. and S. Amante, "Using the IPv6 Flow Label for Equal Cost Multipath Routing and Link Aggregation in Tunnels", RFC 6438, DOI 10.17487/RFC6438, November 2011, <<https://www.rfc-editor.org/info/rfc6438>>.
- [RFC6790] Kompella, K., Drake, J., Amante, S., Henderickx, W., and L. Yong, "The Use of Entropy Labels in MPLS Forwarding", RFC 6790, DOI 10.17487/RFC6790, November 2012, <<https://www.rfc-editor.org/info/rfc6790>>.
- [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>>.
- [RFC8300] Quinn, P., Ed., Elzur, U., Ed., and C. Pignataro, Ed., "Network Service Header (NSH)", RFC 8300, DOI 10.17487/RFC8300, January 2018, <<https://www.rfc-editor.org/info/rfc8300>>.

- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.
- [RFC8979] Sarikaya, B., von Hugo, D., and M. Boucadair, "Subscriber and Performance Policy Identifier Context Headers in the Network Service Header (NSH)", RFC 8979, DOI 10.17487/RFC8979, February 2021, <<https://www.rfc-editor.org/info/rfc8979>>.
- [RFC9016] Varga, B., Farkas, J., Cummings, R., Jiang, Y., and D. Fedyk, "Flow and Service Information Model for Deterministic Networking (DetNet)", RFC 9016, DOI 10.17487/RFC9016, March 2021, <<https://www.rfc-editor.org/info/rfc9016>>.
- [SD-WAN] MEF 70.1 Draft (R1), available at <https://www.mef.net/wp-content/uploads/2020/08/MEF-70-1-Draft-R1.pdf>, "SD-WAN Service Attributes and Service Framework", August 2020.

Authors' Addresses

Nan Geng
Huawei Technologies
China
Email: gengnan@huawei.com

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

Gyan Mishra
Verizon Inc.
United States of America
Email: gyan.s.mishra@verizon.com