

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 16 May 2026

Z. Li
Huawei Technologies
D. Voyer
Bell Canada
C. Xie
China Telecom
P. Liu
China Mobile
Z. Qin
China Unicom
G. Mishra
Verizon Inc.
N. Geng
Huawei Technologies
12 November 2025

Problem Statement and Use Cases of Application-aware Networking (APN)
draft-li-apn-problem-statement-usecases-09

Abstract

Network operators are facing the challenge of providing better network services for users. As the ever-developing 5G and industrial verticals evolve, more and more services that have diverse network requirements such as ultra-low latency and high reliability are emerging, and therefore differentiated service treatment is desired by users. On the other hand, as network technologies such as Hierarchical QoS (H-QoS), SR Policy, and Network Slicing keep evolving, the network has the capability to provide more fine-granularity differentiated services. However, network operators are typically unaware of the applications that are traversing their network infrastructure, which means that not very effective differentiated service treatment can be provided to the traffic flows. As network technologies evolve including deployments of IPv6, SRv6, Segment Routing over MPLS dataplane, the programmability provided by IPv6 and Segment Routing can be augmented by conveying application related information into the network satisfying the fine-granularity requirements.

This document analyzes the existing problems caused by lack of service awareness, and outlines various use cases that could benefit from an Application-aware Networking (APN) framework.

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 16 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	3
2. Requirements Language	3
3. Terminology	3
4. Problem Statement	4
4.1. Challenges of lack of fine-granularity service information	4
4.2. Challenges of Traditional Differentiated Service Provisioning	5
4.3. Challenges of Supporting New 5G and Edge Computing Technologies	7
5. Key Elements of Application-aware Networking (APN)	7
6. Scenarios of APN Domains	8
7. Use cases for Application-aware Networking (APN)	10
7.1. Application-aware SLA Guarantee	11
7.2. Application-aware network slicing	11
7.3. Application-aware Deterministic Networking	12
7.4. Application-aware Service Function Chaining	12
7.5. Application-aware Network Measurement	13
8. IANA Considerations	13
9. Security Considerations	13

10. Acknowledgements	14
11. Co-authors	14
12. Contributors	14
13. References	15
13.1. Normative References	15
13.2. Informative References	15
Authors' Addresses	16

1. Introduction

Due to the requirement for differentiated traffic treatment driven by diverse new services, the ability to convey the application-aware information and program the network infrastructure accordingly to provide fine-grained services is becoming increasingly necessary for network operators. The Application-aware Networking (APN) framework is being defined to address the requirements and some use cases are described in this document. APN takes advantage of network programmability by conveying application related information in the data plane so to facilitate network operators in providing fine-grained services.

One of the key objectives of APN is for network operators to provide fine-granularity SLA guarantees instead of coarse-grain traffic operations. This will allow to provide differentiated services for different applications and increase revenue accordingly. Among various applications being carried and running in the network, some revenue-producing applications such as online gaming, video streaming, and enterprise video conferencing have much more demanding performance requirements such as low network latency and high bandwidth. In order to achieve better Quality of Experience (QoE) for end users and engage customers, the network needs to be able to provide fine-granularity and even application group-level SLA guarantee.

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. Terminology

ACL: Access Control List

APN: Application-aware Networking

APN6: Application-aware Networking for IPv6/SRv6

DPI: Deep Packet Inspection

PBR: Policy Based Routing

QoE: Quality of Experience

SDN: Software Defined Networking

SLA: Service Level Agreement

MPLS: Multiprotocol Label Switching

SR: Segment Routing

SRv6: Segment Routing over IPv6 dataplane

SR-MPLS: Segment Routing over MPLS dataplane

VPN: Virtual Private Network

TE: Traffic Engineering

FRR: Fast Reroute

CAPEX: Capital expenditures

OPEX: Operating expenditures

4. Problem Statement

This section summarizes the challenges currently faced by network operators when attempting to provide fine-grained traffic operations to satisfy the various requirements demanded by new applications that require differentiated service treatment.

4.1. Challenges of lack of fine-granularity service information

In today's networks, the infrastructure through which the traffic is forwarded is not able to provide fine-granularity SLAs satisfying specific requirements. It is therefore difficult for network operators to provide fine-grained traffic operations for various performance-demanding applications. In order to satisfy the SLA requirements network operators continue to increase the network bandwidth but only carrying very light traffic load (in general, around 30%-40% of its capacity).

As network technologies keep evolving, the network capability has been greatly enhanced and is able to provide fine-granularity service provisioning. For example,

H-QoS: provides hierarchical fine-grained QoS services.

SR Policy: provides the ability to handle a large number of explicit and flexible SR paths in order for services to select the policy that satisfies their SLA requirements.

Network Slicing: provides the ability to define a number of isolated network slices having different set of requirements.

IOAM: provides more accurate performance measurement of the traffic flow.

In summary, driven by the ever-emerging diverse demanding services, the lack of the fine-granularity information about the services in the network will cause the following issues: 1) the service information is not clearly described and known by the network, 2) the fine-granularity service provisioning capability is not fully utilized, 3) a fine-granularity service scheduling and measurement cannot be achieved.

4.2. Challenges of Traditional Differentiated Service Provisioning

The traditional ways used to provide fine-grained service provisioning face some challenges. The network devices mainly rely on the 5-tuple of the packets or DPI. However, there are some challenges for these traditional methods in differentiated service provisioning:

1. Five Tuples used for ACL/PBR: five tuples are widely used for ACL/PBR matching of traffic. However, these features cannot provide enough information for the fine-grained service process, and can only provide indirect application-level information which needs to be translated. Generally, ACLs involve high overhead on the forwarding process. Moreover, in some cases such as tunnel encapsulation and IPv6 data plane with a list of extension headers, it becomes impossible to resolve the 5 tuples due to the transport layer information being pushed very deep in the packet.
2. Deep Packet Inspection (DPI): If more information is needed, it must be extracted using DPI which can inspect deep into the packets for application specific information. However, this will introduce more CAPEX and OPEX for the network operator and impose security and privacy challenges.

3. Orchestration and SDN-based Solution: In the era of SDN, typically, an SDN controller is used to manage and operate the network infrastructure and orchestrator elements introduce application requirements so that the network is programmed accordingly. The SDN controller can be aware of the service requirements of the applications on the network through the interface with the orchestrator, and the service requirement is used by the controller for traffic management over the network. However, this method raises the following problems:
- The whole loop is long and time-consuming which is not suitable for fast service provisioning for critical applications;
 - Too many interfaces are involved in the loop, as shown in Figure 1, which introduce challenges of standardization and inter-operability.

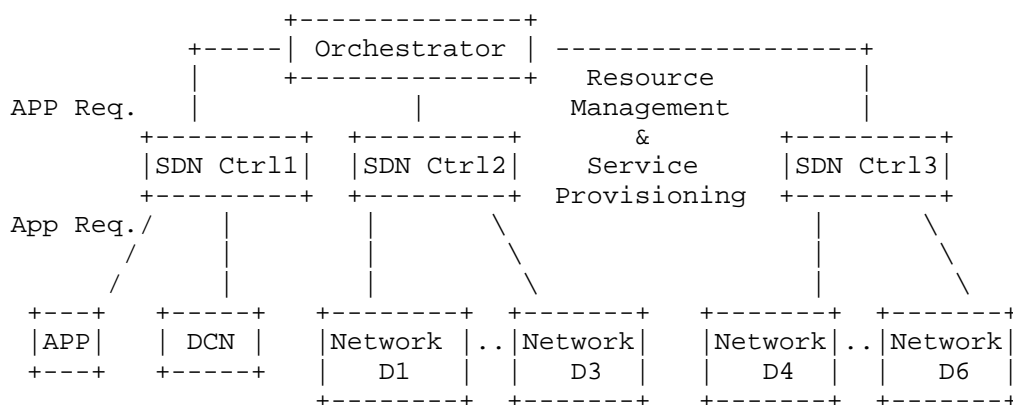


Figure 1: Multiple interfaces involved in the long service-provisioning loop

In [I-D.peng-apn-scope-gap-analysis], some mechanisms that have been specified in IETF and using attribute/identifier to perform traffic steering and service provisioning are analyzed. The existing solutions are specific to a particular scenario or data plane, and a generalized method used for fine-grained service provisioning is still missing.

4.3. Challenges of Supporting New 5G and Edge Computing Technologies

New technologies such as 5G, IoT, and edge computing, are continuously developing leading to more and more new types of services accessing the network. Large volumes of network traffic with diverse requirements such as low latency and high reliability are therefore rapidly increasing. If traditional methods for differentiation of traffic continue to be utilized, it will cause much higher CAPEX and OPEX to satisfy the ever-developing applications' diverse requirements.

5. Key Elements of Application-aware Networking (APN)

Application-aware Networking (APN) aims to address the problems mentioned in Section 4, associated with fine-grained traffic operations that are required in order to satisfy the various application-awareness requirements demanded by new services that need differentiated service treatment.

In APN, the application-aware information (APN Attribute) is derived according to the existing information in the packet header and encapsulated along with the encapsulation of the tunnel. With the APN attribute, fine-granularity network services can be provisioned within the APN domain accordingly. The APN attribute can include application-aware ID (APN ID) and application-aware parameters (APN Parameters). The APN ID can be derived, and applied to the packet, through a mapping mechanism leveraging the existing information in the packet header. The APN parameters can be applied for the corresponding APN ID according to the local policy. The typical APN parameters are the network performance requirements.

APN has the following key elements:

1. Application-aware information (APN attribute) is conveyed in the data plane through augmentation of existing encapsulations such as IPv6, SRv6 and MPLS. The conveyed APN attribute includes APN ID and/or APN parameters. This information is acquired at the edge of the APN domain according to the existing information in the packet header. When a data packet uses APN and conveys the application-aware information, it is referred in this document as an APN packet.

2. Application-aware information and network service provisioning matching providing fine-granularity network service provisioning (traffic operations) and SLA guarantee based on the APN attribute carried in APN packets. According to the APN attribute, appropriate network services are selected, provisioned, and provided to the demanding applications to satisfy their service requirements.
3. Measurement of the network performance so to maintain the match between the applications requirements and corresponding network services for a better fine-granularity SLA compliance. The network measurement methods include in-band and out-of-band, passive, active, per-packet, per-flow, per node, end-to-end, etc. These methods can also be integrated.

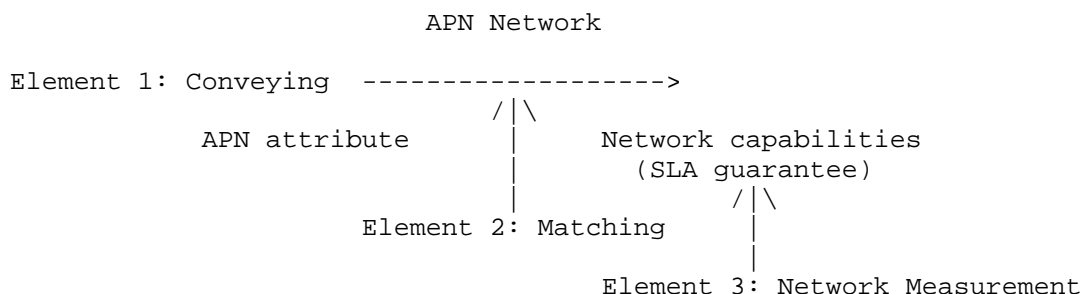


Figure 2: Illustration of the key elements of APN

6. Scenarios of APN Domains

1. SD-WAN scenario

The SD-WAN scenario is shown in the following figure. With APN, at the edge node, i.e. CPE, of the SD-WAN, the 5-tuple, plus information related to user or application group-level requirements is constructed into the APN attribute. When the packet is sent from the CPE, the attribute is added along with the tunnel encapsulation. 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.

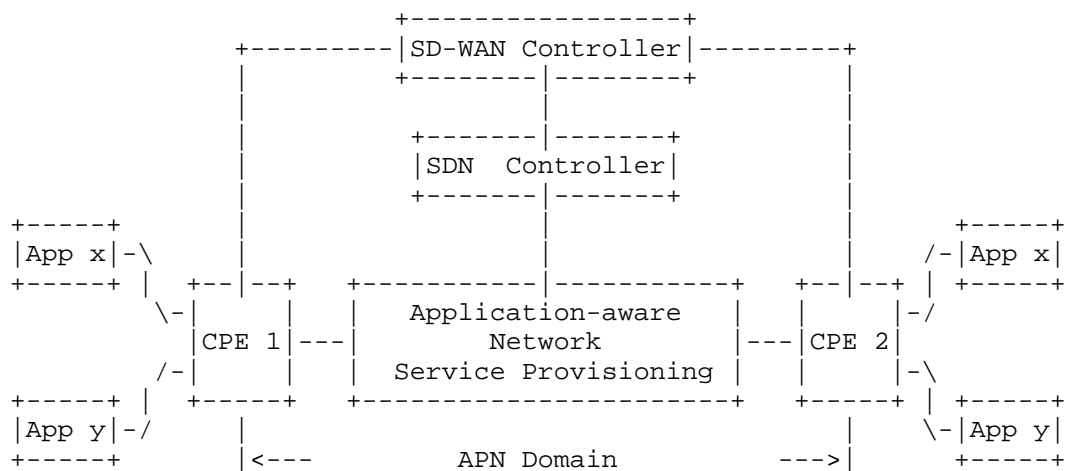


Figure 3: APN Domain in the Scenario of SD-WAN

2. Home broadband scenario

In the home broadband scenario, generally a home broadband user is authorized by the BNG. If the validation is passed and the access control is released, the user group can start enjoying the value-added service. With APN, when the traffic traverses the metro network, the traffic flow can be indicated by the APN attribute that is added/removed at the edge devices of the Metro Network (APN domain) based on the mapping from the existing information (e.g. the QinQ which is composed of C-VLAN and S-VLAN) in the packet header and then carried in the tunnel encapsulation header. The APN attribute will facilitate the fine-granular service in the APN domain. Once the packets leave the APN domain, the APN attribute will be removed together with the tunnel encapsulation header.

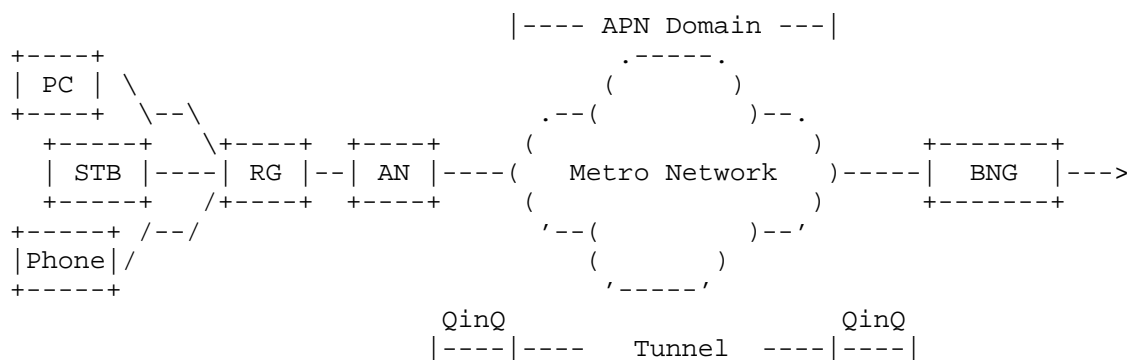


Figure 4: Home Broadband Scenario

3. Mobile broadband scenario

In the mobile broadband scenario, a UE is authorized by the 5GC function, and the traffic steering and QoS policy are enforced by the UPF (User Plane Function) node. If the validation is passed and the access control is released, the user can start enjoying the value-added service. With APN, when the traffic traverses the mobile transport network, the traffic flow can be indicated by the APN attribute that is added at the edge devices of the mobile transport network (APN domain) based on mapping from the existing information (e.g. GTP-u tunnel encapsulation information) in the packet header and then carried in the tunnel encapsulation header. The APN attribute will facilitate the fine-granular service in the APN domain. Once the packets leave the APN domain, the APN attribute will be removed together with the tunnel encapsulation header.

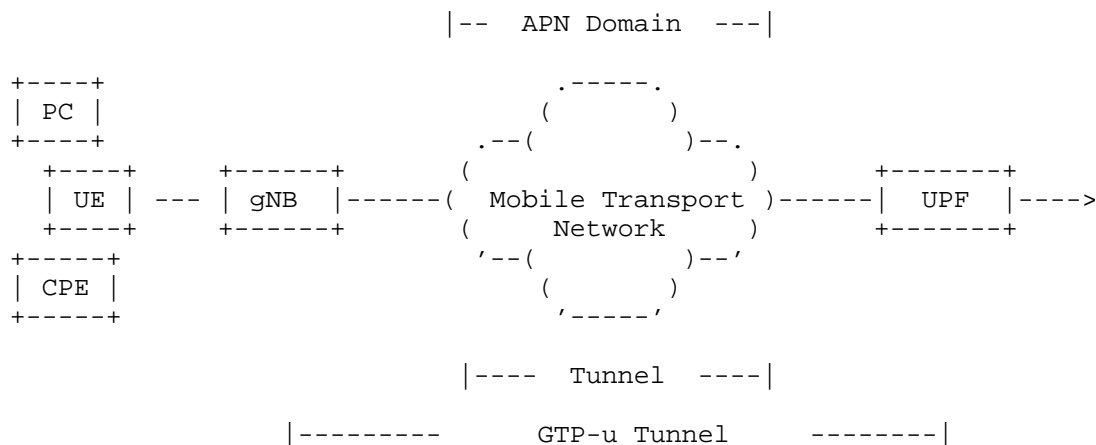


Figure 5: Mobile Broadband Scenario

7. Use cases for Application-aware Networking (APN)

This section illustrates some of the use cases that can benefit from APN. The corresponding requirements for APN are also outlined.

7.1. Application-aware SLA Guarantee

As stated in section 1, among various applications being carried and running in the network, some revenue-producing applications such as online gaming, video streaming, and enterprise video conferencing have much more demanding performance requirements such as low network latency and high bandwidth. In order to achieve better Quality of Experience (QoE) for end users and engage customers, the network needs to be able to provide fine-granularity and even application group-level SLA guarantee. Differentiated service provisioning is also desired.

The APN framework MUST address the following requirements:

- * APN needs to perform the three key elements as described in Section 5.
- * Support application group-level fine-granularity traffic operation that may include finer QoS scheduling.

7.2. Application-aware network slicing

More and more applications/services with diverse requirements are being carried over and sharing a common operators' network infrastructure. However, it is still desirable to have customized network transport that can support some applications' specific requirements, taking into consideration service and resource isolation, which drives the concept of network slicing.

Network slicing provides ways to partition the network infrastructure in either the control plane or data plane into multiple network slices that are running in parallel. These network slices can serve diverse services and fulfill their various requirements at the same time. For example, the mission critical application that requires ultra-low latency and high reliability can be provisioned over a separate network slice.

The APN framework MUST address the following requirements:

- * APN needs to perform the three key elements as described in Section 5 in the context of network slicing.
- * For the element 2, the APN framework MUST allow to assign a given traffic flow to specific network slice according to the APN attribute carried in the APN packet.
- * For the element 3, the APN framework MUST allow the network measurement of each network slice.

7.3. Application-aware Deterministic Networking

[RFC8578] documents use cases for diverse industry applications that require deterministic flows over multi-hop paths. Deterministic flows provide guaranteed bandwidth, bounded latency, and other properties relevant to the transport of time-sensitive data, and can coexist on an IP network with best-effort traffic. It also provides for highly reliable flows through provision for redundant paths.

The APN framework MUST address the following requirements:

- * APN needs to perform the three key elements as described in Section 5 in the context of deterministic networking.
- * For the element 2, the APN framework MUST allow to assign a given traffic flow to a specific deterministic path according to the APN attribute carried in the APN packet.
- * For the element 3, the APN framework MUST allow the network measurement of each application-aware deterministic path.

7.4. Application-aware Service Function Chaining

End-to-end service delivery often needs to go through various service functions including traditional network service functions such as firewalls, DPIs as well as new application-specific functions, both physical and virtual. The definition and instantiation of an ordered set of service functions and subsequent steering of the traffic through them is called Service Function Chaining (SFC) [RFC7665]. SFC is applicable to both fixed and mobile networks as well as data center networks.

Generally, in order to manipulate a specific traffic flow along the SFC, a DPI needs to be deployed as the first service function of the chain to detect the application, which will impose high CAPEX and consume long processing time. For encrypted traffic, it even becomes impossible to inspect the traffic flow.

The APN framework MUST address the following requirements:

- * APN needs to perform the three key elements as described in Section 5 in the context of service function chaining.
- * For the element 1, class information can be conveyed.

- * For the element 2, the APN framework MUST allow to assign a given traffic flow to a specific service function chain and MUST allow the subsequent steering according to the APN attribute carried in the APN packets.
- * For the element 3, the APN framework MUST allow the network measurement of each application-aware service function chain.

7.5. Application-aware Network Measurement

Network measurement can be used for locating silent failure and predicting QoE satisfaction, which enables real-time SLA awareness/proactive OAM. Operations, Administration, and Maintenance (OAM) refers to a toolset for fault detection and isolation, and network performance measurement. In-situ Operations, Administration, and Maintenance (IOAM) records operational and telemetry information in the packet while the packet traverses a path between two points in the network.

The APN framework MUST address the following requirements:

- * APN needs to perform the two key elements as described in Section 5 in the context of network measurement. The network measurement in the element 3 does not need to be considered here.

8. IANA Considerations

This document does not include an IANA request.

9. Security Considerations

In the APN work, in order to reduce the privacy and security issues, the APN attribute MUST be conveyed along with the tunnel information in the APN domain. The APN attribute is encapsulated and removed at the edge of the APN domain. The APN ID MUST be acquired from the existing available information in the packet header without interference into the payload.

According to the above specifications, the APN attribute is only produced and used locally within the APN domain without the involvement of the host/application side.

In order to prevent the malicious attack through the APN attribute, the following policies can be configured at the network devices of the APN domain. If the APN attribute is conveyed without the tunnel information, the packet MUST be dropped. If the APN attributes are not known to the APN domain, it should trigger the alarm information. The packet can be forwarded without being processed or dropped

depending on the local policy. If the network service requirements exceed the specification for the specific APN ID, it should trigger the alarm information. The packet should be discarded to prevent abusing of the resources. There should be rate-limiting policy at the edge of the APN domain to prevent the traffic belonging to a specific APN ID from exceeding the preset limit.

10. Acknowledgements

The authors would like to acknowledge Robert Raszuk (Bloomberg LP) and Yukito Ueno (NTT Communications Corporation) for their valuable review and comments.

11. Co-authors

Kentaro Ebisawa
Toyota Motor Corporation
Japan

Email: ebisawa@toyota-tokyo.tech

Stefano Previdi
Huawei Technologies
Italy

Email: stefano@previdi.net

James N Guichard
Futurewei Technologies Ltd.
USA

Email: jguichar@futurewei.com

12. Contributors

Daniel Bernier
Bell Canada
Canada

Email: daniel.bernier@bell.ca

Liang Geng
China Mobile
China

Email: gengliang@chinamobile.com

Chang Cao
China Unicom
China

Email: caocl5@chinaunicom.cn

Chang Liu
China Unicom
China

Email: liucl31@chinaunicom.cn

Cong Li
China Telecom
China

Email: licong@chinatelecom.cn

13. References

13.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [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>>.
- [RFC8578] Grossman, E., Ed., "Deterministic Networking Use Cases", RFC 8578, DOI 10.17487/RFC8578, May 2019, <<https://www.rfc-editor.org/info/rfc8578>>.
- [RFC7665] Halpern, J., Ed. and C. Pignataro, Ed., "Service Function Chaining (SFC) Architecture", RFC 7665, DOI 10.17487/RFC7665, October 2015, <<https://www.rfc-editor.org/info/rfc7665>>.

13.2. Informative References

- [I-D.peng-apn-scope-gap-analysis] Peng, S., Li, Z., and G. S. Mishra, "APN Scope and Gap Analysis", Work in Progress, Internet-Draft, draft-peng-apn-scope-gap-analysis-07, 14 September 2023, <<https://datatracker.ietf.org/doc/html/draft-peng-apn-scope-gap-analysis-07>>.

Authors' Addresses

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

Daniel Voyer
Bell Canada
Canada
Email: daniel.voyer@bell.ca

Chongfeng Xie
China Telecom
China
Email: xiechf@chinatelecom.cn

Peng Liu
China Mobile
China
Email: liupengyjy@chinamobile.com

Zhuangzhuang Qin
China Unicom
China
Email: qinzhuangzhuang@chinaunicom.cn

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

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