

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

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

Application-aware Networking (APN) Framework
draft-li-rtgwg-apn-framework-01

Abstract

A multitude of applications are carried over the network, which have varying needs for network bandwidth, latency, jitter, and packet loss, etc. Some new emerging applications have very demanding performance requirements. However, in current networks, the network and applications are decoupled, that is, the network is not aware of the applications' requirements in a fine granularity. Therefore, it is difficult to provide truly fine-granularity traffic operations for the applications and guarantee their SLA requirements.

This document proposes a new framework, named Application-aware Networking (APN), where application-aware information (i.e. APN attribute) including APN identification (ID) and/or APN parameters (e.g. network performance requirements) is encapsulated at network edge devices and carried in packets traversing an APN domain in order to facilitate service provisioning, perform fine-granularity traffic steering and network resource adjustment.

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	4
3. Terminology	4
4. APN Framework and Key Components	5
5. APN Requirements	7
5.1. APN Attribute Conveying Requirements	8
5.1.1. Protocol Extensions Requirements	9
5.2. APN attribute Handling Requirements	11
5.2.1. Fine granular SLA Guarantee	11
5.2.2. Fine granular network slicing	11
5.2.3. Fine granular deterministic networking	12
5.2.4. Fine granular service function chaining	12
5.2.5. Fine granular network measurement	13
6. Illustration	13
6.1. Example use case description	13
6.2. User Group and Application Group Design	14
6.3. Derive the User Group and User Group at APN Edge	16
6.4. Access Right Check at the edge of the backbone network	16
6.5. SLA Guarantee in the backbone network	17
6.5.1. Network Measurement	17
6.5.2. Traffic Steering	18
7. Benefits	18
8. IANA Considerations	19
9. Security Considerations	19
10. Acknowledgements	20

11. Co-authors	20
12. Contributors	21
13. References	22
13.1. Normative References	22
13.2. Informative References	22
Authors' Addresses	22

1. Introduction

A multitude of applications are carried over the network, which have varying needs for network bandwidth, latency, jitter, and packet loss, etc. Some applications such as online gaming and live video streaming have very demanding network requirements and therefore require special treatment in the network. However, in current networks, the network and applications are decoupled, that is, the network is not aware of the applications' requirements in a fine granularity. Therefore, it is difficult to provide truly fine-granularity traffic operations for the applications and guarantee their SLA requirements accordingly.

[I-D.li-apn-problem-statement-usecases] describes the challenges of traditional differentiated service provisioning methods, such as five tuples used for ACL/PBR causing coarse granularity as well as orchestration and SDN-based solution causing long control loops.

This document proposes a new framework, named Application-aware Networking (APN), where application-aware information (APN attribute) including application-aware identification (APN ID) and application-aware parameters (APN Parameters), is encapsulated at network edge devices and carried along with the encapsulation of the tunnel used by the packet when traversing the APN domain. By APN domain we intend the operator infrastructure where APN is used from edge to edge (ingress to egress) and where the packet is encapsulated using an outer header incorporating the APN information. The APN attribute will facilitate service provisioning and provide fine-granularity services in the APN domain.

The APN attribute is acquired at the ingress of the APN domain based on the existing information in the incoming 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) in the edge devices. The APN information is then added to the data packets along with the tunnel encapsulation. The packet traverses the domain and, when exiting the domain, the outer header along with the APN information is removed.

APN aims to leverage the ability to apply policies to traffic flows entering into the infrastructure (APN domain). For example, at the headend (ingress) traffic is steered into a given path/policy, at a midpoint node the corresponding performance measurement data is collected, and at a service node a given function is executed.

APN assumes traffic is tunnel encapsulated 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.

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

LB: Load Balancing

MPLS: Multiprotocol Label Switching

PBR: Policy Based Routing

QoE: Quality of Experience

SDN: Software Defined Networking

SLA: Service Level Agreement

SR: Segment Routing

SR-MPLS: Segment Routing over MPLS dataplane

SRv6: Segment Routing over IPv6 dataplane

4. APN Framework and Key Components

The APN framework is shown in Figure 1. The key components include App-aware Edge Device (APN-Edge), App-aware-process Head-End (APN-Head), App-aware-process Mid-Point (APN-Midpoint), App-aware-process End-Point (APN-Endpoint), and APN-Controller. The key interfaces include both the Northbound Interface (NBI) and the Southbound Interface (SBI) of the APN-Controller.

Packets carry application characteristic information (i.e. APN attribute) which includes the following information:

- * Application-aware identification (APN ID): identifies the set of attributes, indicating that all packets belonging to the same flow will be given the same treatment by the network. APN ID is mandatory.
- * Application-aware parameters (APN parameters): The typical application-aware parameters are the network performance requirement parameters including bandwidth, delay, delay variation, packet loss ratio, etc. APN parameters are optional.

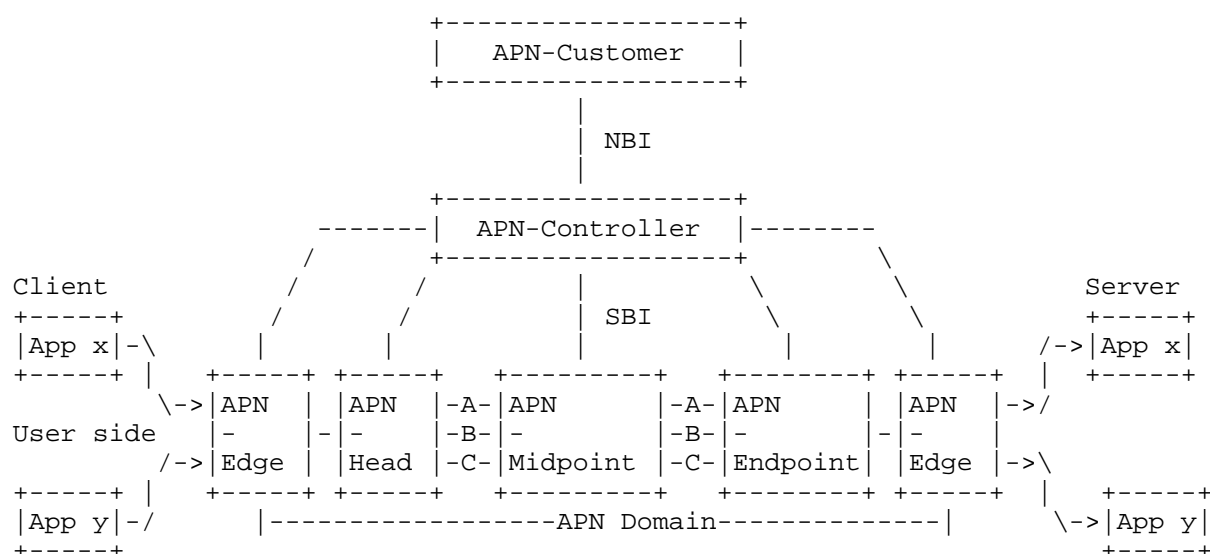


Figure 1: Framework and Key Components

The key components are introduced as follows.

- * App-aware Edge Device (APN-Edge): this network device receives packets from applications and obtains the APN attribute based on the configuration on this device according to the existing information in the packet header, such as 5-tuple, VLAN or double VLAN tagging (C-VLAN and S-VLAN). The APN-Edge device adds the APN attribute in the tunnel encapsulation. The packets carrying the APN attribute will be sent to the APN-Head, and the APN attribute will be used to apply various policies in different nodes along the network path onto the traffic flow, e.g., at the headend to steer into corresponding path satisfying SLAs, at the midpoint to collect corresponding performance measurement data, at the service function to execute particular policies. When the packets leave the APN domain, the APN attribute will be removed together with the tunnel encapsulation.
- * App-aware-process Head-End (APN-Head): This network device receives packets from the APN-Edge, obtains the APN attribute, and initiates the corresponding process. Generally, in order to satisfy different SLA requirements, a set of paths, tunnels or SR policies, are set up between the APN-Head and the APN-Endpoint. These multiple parallel paths have different SLA guarantees. The APN-Head maintains the matching relationship between the APN attribute and the paths between the APN-Head and the APN-Endpoint. The APN-Head determines the path between the APN-Head and the APN-Endpoint according to the APN attribute carried in the packets and the matching relationship with it, which satisfies the service requirements of the applications. The APN-Head forwards the packets along the path. The APN attribute conveyed by the packet received from the APN-Edge can also be copied or be mapped to the outgoing packet header.
- * App-aware-process Mid-Point (APN-Midpoint): the APN-Midpoint provides the path service and enforces various policies according to the APN attribute carried in the packets. The APN-Midpoint may also adjust the resource locally to guarantee the service requirements depending on a specific policy and the APN attribute conveyed by the packet. Policy definitions and mechanisms are out of the scope of this document.
- * App-aware-process End-Point (APN-Endpoint): the process of the specific service path will end at the APN-Endpoint. If the outer tunnel header for the path between the APN-Head and the APN-Endpoint exists, it will be removed by the APN-Endpoint. If the APN attribute is copied or mapped to the outer tunnel header by the APN-Head, it will also be removed along with the outer tunnel header.

Note that, the APN-Edge can co-exist with the APN-Head or APN-Endpoint, that is, one network device can implement the functionalities of both APN-Edge and the APN-Head/APN-Endpoint.

- * **APN-Controller:** the APN-Controller is a functional block responsible for planning and execution of how the APN attribute are allocated and maintained. It collects the service requirements using a NBI (northbound interface) and signals the APN related policies to the network devices using SBI (southbound interfaces). To be more specific, the APN related policies include the APN-marking policy (i.e., the matching state in order for the node to classify and encapsulate the incoming packet with the desired APN Information) being pushed on the APN-Edge, and the APN-servicing policy (i.e., the APN state in order to fulfill the service requirements using e.g. traffic steering and performance measurement) on the APN-Head, APN-Midpoint, APN-Endpoint.
- * **APN-Customer:** the APN-Customer manages the APN network services [I-D.li-apn-problem-statement-usecases] using the APN customer service model, which describes the requirements on the APN network services from the point of view of the APN customer.

The key interfaces are introduced as follows.

- * **The Northbound Interface (NBI) of the APN-Controller:** the requirements described in the APN service model by the APN-Customer is enforced via this interface to the APN-Controller, which will be translated into the APN policies, i.e., the APN-marking policy and the APN-servicing policy, in the APN-Controller.
- * **The Southbound Interface (SBI) of the APN-Controller:** the APN-marking policy and the APN-servicing policy are enforced via this interface from the APN-Controller to the relevant network devices. The candidate protocols for this interface are PCEP, BGP, YANG-based protocols (NETCONF/RESTCONF), etc.

These interfaces also include the operational status and monitoring information.

5. APN Requirements

This section specifies the requirements for supporting the APN framework, including the requirements for conveying and handling the APN attribute.

5.1. APN Attribute Conveying Requirements

The APN attribute consists of APN ID and APN parameters.

APN ID includes the following identifiers (IDs),

- * Application Group ID: identifies an application group of the traffic.
- * User Group ID: identifies the user group of the traffic.

APN ID can be acquired through different ways. In the APN framework it MUST be acquired according to the existing available information in the packet header without inspection of the payload.

The different combinations of the IDs can be used to provide different granularity of the service provisioning and SLA guarantee for the traffic.

The APN parameters are the network performance requirement parameters. The network service requirement can include the following parameters:

- * Bandwidth: the bandwidth requirement
- * Latency: the latency requirement
- * Packet loss ratio: the packet loss ratio requirement
- * Jitter: the jitter requirement

The different combinations of the parameters are for further expressing the more detailed service requirements, conveyed together with the APN ID, which can be used to match to appropriate tunnels/SR Policies and queues that can satisfy these service requirements.

APN attribute MUST be encapsulated within tunnels in the network layer. The tunnels include but not limit to MPLS, VxLAN, SR-MPLS, and SRv6. It can be extended according to requirements in the future.

[REQ 1a]. APN ID SHOULD include Application Group ID to indicate the application group that the packet belongs to.

[REQ 1b]. APN ID SHOULD include User Group ID to indicate the user group that the packet belongs to.

[REQ 1c]. APN ID MUST include either Application Group ID or User Group ID.

[REQ 1d]. APN ID MUST be acquired from the existing available information of the packet header without interference into the payload.

[REQ 1e]. APN parameters is OPTIONAL.

[REQ 1f]. APN attribute MUST be carried by the outer tunnel encapsulation.

[REQ 1g]. All the nodes along the path SHOULD be able to process the APN attribute if needed.

[REQ 1h]. The APN attribute is generated by the APN-Edge though local policy.

[REQ 1i]. The APN attribute SHOULD be kept intact when directly copied at the APN-Head and carried in the tunnel encapsulation.

[REQ 1j]. The APN attribute MUST be removed along with the tunnel encapsulation by the APN-Edge when the packets leave the APN domain.

[REQ 1k]. The APN attribute MUST NOT be encrypted when the APN packet is itself encrypted (e.g., the APN tunnel across the APN domain uses IPsec).

5.1.1.1. Protocol Extensions Requirements

The APN attribute is conveyed with the tunnel encapsulation. There are two typical types of tunnels:

- * MPLS-based tunnel: LDP tunnel, RSVP-TE tunnel, SR-MPLS tunnel or policy, etc.
- * IPv6-based tunnel: IPv6-based VxLAN tunnel, IPv6-based UDP tunnel, IPv6-based GRE tunnel, SRv6 tunnel or policy, etc.

In order to support encapsulation of APN attribute, the MPLS data plane and IPv6 data plane need to be extended.

In order to support acquiring the APN attribute according to the existing available information in the packet header, YANG models should be defined to configure the mapping between the application/user group ID and the existing information in the packet header and configure the corresponding APN attribute for the application/user group. It can also be implemented with protocol extensions such as BGP/BGP-LS and PCEP which can advertise the information from the central controller to the APN-Edge.

In addition, in the APN domain, the above-mentioned mapping and applying APN parameters may also be advertised from the APN-Edge/APN-Head to other devices or from the network devices to the central controller in the APN domain. IGP extensions or BGP-LS extensions should be introduced to achieve the purposes.

[REQ 1-1a] MPLS encapsulation SHOULD be extended to be able to carry the APN attribute for MPLS-based tunnels.

[REQ 1-1b] IPv6 encapsulation SHOULD be extended to be able to carry the APN attribute for IPv6-based tunnels.

[REQ 1-1c] YANG models SHOULD be defined to implement the mapping between the application/user group ID and the existing available information in the packet header and configure the corresponding APN parameters.

[REQ 1-1d] BGP extensions SHOULD be defined to advertise the mapping between the application/user group ID and the existing available information in the packet header and the corresponding APN parameters from the central controller to the APN-Edge in the APN domain.

[REQ 1-1e] PCEP extensions SHOULD be defined to advertise the mapping between the application/user group ID and the existing available information in the packet header and the corresponding APN parameters from the central controller to the APN-Edge in the APN domain.

[REQ 1-1f] IGP extensions SHOULD be defined to advertise the mapping between the application/user group ID and the existing available information in the packet header and the corresponding APN parameters from the APN-Edge to the network devices in the APN domain.

[REQ 1-1g] BGP-LS extensions SHOULD be defined to advertise the mapping between the application/user group ID and the existing available information in the packet header and the corresponding APN parameters from the network devices to the central controller in the APN domain.

5.2. APN attribute Handling Requirements

The APN Head and APN-Midpoint perform matching operation against the APN attribute, that is, to match IDs and/or service requirements to the corresponding network resources such as tunnels/SR policies and queues.

5.2.1. Fine granular SLA Guarantee

In order to achieve better Quality of Experience (QoE) of end users and engage customers, the network needs to be able to provide fine-granularity SLA guarantee [I-D.li-apn-problem-statement-usecases].

[REQ 2-1a]. With the APN attribute, the APN-Head SHOULD be able to steer the traffic to the tunnel/SR policy that satisfies the matching operation.

[REQ 2-1b]. With the APN attribute, the APN-Head SHOULD be able to trigger the setup of the tunnel/SR policy that satisfies the matching operation.

[REQ 2-1c]. With the APN attribute, the APN-Head and APN-Midpoint SHOULD be able to steer the traffic to the queue that satisfies the matching operation.

[REQ 2-1d]. With the APN attribute, the APN-Head and APN-Midpoint SHOULD be able to trigger the configuration of the queue that satisfies the matching operation.

[REQ 2-1e]. If the tunnels are used to satisfy the performance requirements, the APN-Head SHOULD be able to copy or map the APN attribute conveyed by the packet received from the APN-Edge to the outer tunnel header.

[REQ 2-1f]. If the tunnels are used to satisfy the performance requirements and the APN attribute are conveyed along with the outer tunnel, the APN-Endpoint MUST remove the APN attribute along with the outer tunnel.

5.2.2. Fine granular network slicing

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

APN is to help the operator of a network to steer some of the traffic tagged with an APN attribute to a certain network slice based on the SLA agreement with its customer.

[REQ 2-2a]. With the APN attribute, the APN-Head SHOULD be able to steer the traffic to the slice that satisfies the matching operation.

[REQ 2-2b]. With the APN attribute, the APN-Midpoint SHOULD be able to associate the traffic to the resources in the slice that satisfies the matching operation.

5.2.3. Fine granular deterministic networking

Along the path each node needs to provide guaranteed bandwidth, bounded latency, and other properties relevant to the transport of time-sensitive data for the Detnet flows that coexist with the best-effort traffic.

APN is to help the operator of a network to steer some of the traffic tagged with an APN attribute to a certain deterministic path based on the SLA agreement with its customer.

[REQ 2-3a]. With the APN attribute, the APN-Head MUST be able to steer the traffic to the appropriate path that satisfies the matching operation.

[REQ 2-3b]. With the APN attribute, the APN-Head MUST be able to trigger the setup of the appropriate path that satisfies the matching operation for the Detnet flows.

[REQ 2-3c]. With the APN attribute, the APN-Midpoint MUST be able to associate the traffic to the resources along the path that satisfies the performance guarantee.

[REQ 2-3d]. With the APN attribute, the APN-Midpoint MUST be able to reserve the resources for the Detnet flows along the path that satisfies the performance guarantee.

5.2.4. Fine granular service function chaining

The end-to-end service delivery often needs to go through various service functions, including traditional network service functions such as firewalls, LB as well as new application-specific functions, both physical and virtual. SFC is applicable to both fixed and mobile networks as well as data center networks.

APN is to help the operator of a network to steer some of the traffic tagged with an APN attribute to a certain service function chain based on the SLA agreement with its customer. On each service function along the service function chain, the policy can be enforced based on the APN attribute in the outer header.

[REQ 2-4a]. With the APN attribute, the App-aware-process devices SHOULD be able to steer the traffic to the appropriate service function.

[REQ 2-4b]. The App-aware-process devices including VAS SHOULD be able to process the APN attribute carried in the packets.

5.2.5. Fine granular network measurement

Network measurement can be used for verifying whether the network performance requirements have been satisfied, as well as locating silent failure and predicting QoE satisfaction, which enables real-time SLA awareness/proactive OAM and potential resource adjustments.

APN is to help the operator of a network to trigger performance measurement for the traffic tagged with an APN attribute based on its customer' consent.

[REQ 2-5a]. The App-aware-process devices SHOULD be able to perform IOAM based on the APN attribute.

[REQ 2-5b]. The network measurement results can be reported based on the APN attribute and verify whether the performance requirements are satisfied.

6. Illustration

In order to better clarify what APN can enable with the introduced APN attribute compared to the existing network without APN, we illustrate how APN works through an example use case, which is also a typical network service being provisioned nowadays, i.e. the Cloud Leased Line service. In order to make the tunnel description much easier to understand, we use the recent technology in IETF, i.e. SRv6.

6.1. Example use case description

We take the "SRv6-based Cloud Leased Line Service" as an illustrative example to show how APN is needed and can be beneficial.

Enterprises usually buy Cloud Leased Line Service to interconnect their local sites to Cloud. Generally, the Cloud Leased Line Service needs to go across multiple domains which are owned by the same operator and can be controlled by multiple controllers and an orchestrator/super-controller.

Due to management reasons, the network information in the intermediate domain cannot be advertised to other domains, so the ingress node cannot set up an appropriate E2E path. In that case, the intermediate domain is treated as a black box, and no fine grain traffic steering and other services can be provisioned.

The example of the network to provide the cloud leased lined service reference diagram is shown as the following figure. The network is composed by three network domains including the two metro networks in the City A and City B and the backbone network which connects the two metro networks. The cloud leased line services is provided to the specific enterprise whose branches located in different cities need to access the cloud-based service located in the City B.

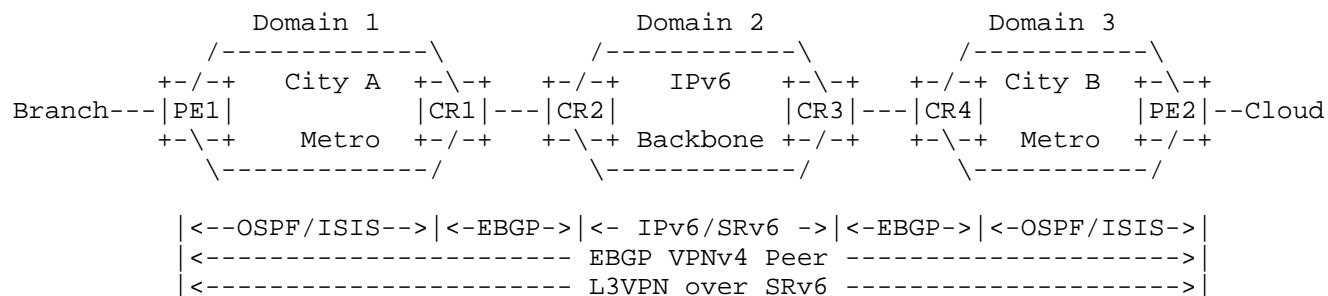


Figure 2: Reference diagram for the example use case illustration

6.2. User Group and Application Group Design

The user groups can be designed as follows:

	User Group
Enterprise A/Branch 1/Office Users	001001001
Enterprise A/Branch 1/R&D Users	001001002
Enterprise A/Branch 1/IT Users	001001003
Enterprise A/Branch 1/VIP Users	001001004
Enterprise A/Branch 2/Office Users	001002001
Enterprise A/Branch 2/R&D Users	001002002
Enterprise A/Branch 2/IT Users	001002003
Enterprise A/Branch 2/VIP Users	001002004
Enterprise A/Branch 3/Office Users	001003001
Enterprise A/Branch 3/R&D Users	001003002
Enterprise A/Branch 3/IT Users	001003003
Enterprise A/Branch 3/VIP Users	001003004

In the IP address design, the IPv6 address blocks allocated to the branches are as follows :

	IPv6 Address
Enterprise A/Branch 1/Office Users	2001:DB8:A:11::/56
Enterprise A/Branch 1/R&D Users	2001:DB8:A:12::/56
Enterprise A/Branch 1/IT Users	2001:DB8:A:13::/56
Enterprise A/Branch 1/VIP Users	2001:DB8:A:1D::/56
Enterprise A/Branch 2/Office Users	2001:DB8:A:21::/56
Enterprise A/Branch 2/R&D Users	2001:DB8:A:22::/56
Enterprise A/Branch 2/IT Users	2001:DB8:A:23::/56
Enterprise A/Branch 2/VIP Users	2001:DB8:A:2D::/56
Enterprise A/Branch 3/Office Users	2001:DB8:A:31::/ 56
Enterprise A/Branch 3/R&D Users	2001:DB8:A:32::/56
Enterprise A/Branch 3/IT Users	2001:DB8:A:33::/56
Enterprise A/Branch 3/VIP Users	2001:DB8:A:3D::/56

The application groups provided by the cloud can be designed as follows:

	Application Group
Enterprise A/Office Audio Applications	101001001
Enterprise A/Office Video Applications	101001002
Enterprise A/Office Data Applications	101001003
Enterprise A/R&D Audio Applications	101002001
Enterprise A/R&D Video Applications	101002002
Enterprise A/R&D Data Applications	101002003
Enterprise A/IT Audio Applications	101003001
Enterprise A/IT Video Applications	101003002
Enterprise A/IT Data Applications	101003003

In the address design, the IPv6 address blocks allocated to the applications of Enterprise A in the cloud is 2001:DB8:A1::/48A1::A:/16. The port number can be used to identify different applications.

	IPv6 Address	Port Number
Enterprise A/Office Audio Applications	2001:DB8:A1:A1::/64	1718, 1719
Enterprise A/Office Video Applications	2001:DB8:A1:A1::/64	5060, 5061
Enterprise A/Office Data Applications	2001:DB8:A1:A1::/64	21, 80
Enterprise A/R&D Audio Applications	2001:DB8:A1:A2::/64	1718, 1719
Enterprise A/R&D Video Applications	2001:DB8:A1:A2::/64	5060, 5061
Enterprise A/R&D Data Applications	2001:DB8:A1:A2::/64	21, 80
Enterprise A/IT Audio Applications	2001:DB8:A1:A3::/64	1718, 1719
Enterprise A/IT Video Applications	2001:DB8:A1:A3::/64	5060, 5061
Enterprise A/IT Data Applications	2001:DB8:A1:A3::/64	21, 80

6.3. Derive the User Group and User Group at APN Edge

The cloud leased line service adopts the SRv6-based L3VPN to traverse the network. The following policy can be applied at the APN edges of the City A1:

```
Match:
  VPN1
  Source Address 2001:DB8:A:11::/56
Action
  Set user-group 001001001

Match:
  VPN1
  destination Address 2001:DB8:A1:A1::/64
  destination port 1718,1719
Action
  Set app-group 101001001
```

6.4. Access Right Check at the edge of the backbone network

The following check can be applied at the edge of the IP backbone network:

```
Match:
  user-group 001001001
  app-group 101002001, 101002002, 101002003, 101003001, 101003002, 101003003
Action
  Deny

Match:
  user-group 001001001
  app-group 101001001, 101001002, 101001003
Action
  Permit
```

The policy means that the office users of the branch 1 can only access the office applications.

If the address allocation is changed. For example, one office user of the branch1's IPv6 address is changed to 2001:DB8:A:15::/56 because of the mobile office.

We only need to add the following policy at the APN edge:

```
Match:
  VPN1
  Source Address 2001:DB8:A:15::/56
Action
  Set user-group 001001001
```

The policy in the backbone network which is based on the user group and the application group is not necessary to change.

6.5. SLA Guarantee in the backbone network

Due to management reasons, the network information in the intermediate domain cannot be advertised to other domains, so the ingress node cannot set up an appropriate TE path, the intermediate domain is treated as a black box and no fine grain traffic steering can be performed.

In this case, we consider fine grain traffic steering in Domain 2 on top of the SRv6-based Cloud Leased Line Service for the purpose of SLA Guarantee.

6.5.1. Network Measurement

In order to guarantee SLA for the VIP users, the following network measurement policy can be applied in the backbone network:

```
Match:
  User-group 001001004 application group 101001002
  User-group 001002004 application group 101002002
  User-group 001003004 application group 101003002
Action
  Apply IOAM
```

The policy is to apply the IOAM as the network measurement for the VIP users of the branches to access the video applications. From the above illustration, there is the following observation:

When there is no APN deployed, at CR2, the 5 tuples of the original packets will need to be resolved since they have been encapsulated, and then IOAM can be triggered based on the 5 tuples. This resolution process is costly and consumes a lot of hardware resources. If Domain 3 needs to trigger IOAM, the same resolution process will have to be done at CR4.

When there is APN deployed, at CR1, the APN attribute is tagged. When these packets arrive at CR2, only the APN attribute in the outer header will be read out, based on which the IOAM can be triggered in Domain 2. That is, no 5-tuple resolution process is needed at CR2 but only checking the APN attribute in the outer header.

6.5.2. Traffic Steering

If the SLA guarantee of the VIP users accessing the video applications does not satisfy the requirements through the network measurement based on the IOAM, the SRv6 policy can be setup. For example, the SRv6 policy 1 which can satisfy the SLA requirement is set up. Then the following policy can be downloaded to the edge:

Match:

User-group 001001004 application group 101001002

User-group 001002004 application group 101002002

User-group 001003004 application group 101003002

Action

Redirect SRv6 Policy 1

The policy is to steer the traffic of the VIP users to the SRv6 policy in the backbone to satisfy the requirement .

From the above illustration, there is the similar observation as the network measurement:

When there is no APN deployed, at CR2, the 5 tuples of the original packets will need to be resolved since they have been encapsulated, and then the traffic can be steered into SRv6 policy 2 based on the 5 tuples. This resolution process is costly and consumes a lot of hardware resources.

When there is APN deployed, at CR1, the APN attribute is tagged. When these packets arrive at CR2, only the APN attribute in the outer header will be read out, based on which the traffic can be steered into SRv6 policy 2 in Domain 2.

7. Benefits

The APN attribute allows the network devices to only look at one easily-accessible field in the outer header, without having to resolve the 5 tuples of the original packets that are deeply encapsulated in the tunnel encapsulation.

The APN attribute allows to simplify the policy control at every policy enforcement point within the network. The APN attribute allows to reducing each matching entry of policy filter since it is

only one field and hardware resources are saved. Since APN attribute is relatively stable, it introduces the possibilities of eliminating the "stale" policy filter entries. In most cases, the APN attribute is centralized configured and distributed to all the policy enforcement points, which saves the policy filter configurations per node and simplifies the OM.

The structured APN attribute allows to express fine granular service requirements, e.g. MKT-user-group/app-group, RD-user-group/app-group, latency.

The structured APN attribute allows to match to the evolving fine granular differentiated network capabilities, e.g. SR policy with low latency and high reliability guaranteed.

In a tunnel across multiple domains of the same operator using the APN attribute in the outer header the operator can easily support multiple services not just a single one in a particular domain as illustrated in the use case illustration section.

When there is no APN, to achieve the same, now the operator may have two options: 1. Add all the policy identifiers at the tunnel headend with various further encapsulations and enforce the policies based on them at the intermediate policy enforcement nodes along the tunnel, 2. Resolve the original 5 tuples being encapsulated inside the tunnel which will be very costly and sometimes impossible.

Moreover, the policy enforcement table in the intermediate policy enforcement nodes is significantly reduced. Because before operator needs to resolve the 5 tuple but now with APN, operator only needs to read the APN attribute in one field of the outer header.

Since the 5 tuples of the traffic are changing frequently due to service deployment or management issues the policy enforcement table in the policy enforcement nodes is not stable and there is always a lot of stale entries in the table. But now since the APN attribute is a mapping of the 5 tuples operator will have a relatively stable policy enforcement table on their nodes.

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 following specifications are defined:

[S1]. The APN attribute MUST be conveyed along with the tunnel information in the APN domain. The APN attribute is encapsulated and removed at the APN-Edge.

[S2]. The APN ID (including the Application Group ID and the User Group 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:

[P1]. If the APN attribute is conveyed without the tunnel information, the packet MUST be dropped.

[P2]. If the APN attribute is 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.

[P3]. If the network service requirements exceed the specification for the specific Application Group ID and/or User Group ID, it should trigger the alarm information. The packet should be discarded to prevent abusing of the resources.

[P4]. There should be rate-limiting policy at the APN-Edge to prevent the traffic belonging to a specific Application Group ID and/or User Group ID from exceeding the preset limit.

10. Acknowledgements

The authors would like to acknowledge Robert Raszuk (Bloomberg LP), Yukito Ueno (NTT Communications Corporation), and Dhruv Dhody for their valuable reviews 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

Email: daniel.bernier@bell.ca

Chongfeng Xie
China Telecom

Email: xiechf@chinatelecom.cn

Feng Yang
China Mobile

Email: yangfeng@chinamobile.com

Zhuangzhuang Qin
China Unicom

Email: qinzhuangzhuang@chinaunicom.cn

Chang Liu
China Unicom

Email: liuc131@chinaunicom.cn

Gyan Mishra
Verizon

Email: hayabusagsm@gmail.com

Luis M. Contreras
Telefonica

Email: contreras.ietf@gmail.com

Luc-Fabrice Ndifor Ngwa
MTN

Email: Luc-Fabrice.Ndifor@mtn.com

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>>.

13.2. Informative References

- [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>>.

Authors' Addresses

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

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

Cong Li
China Telecom
China
Email: licong@chinatelecom.cn

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

Chang Cao
China Unicom
China
Email: caoc15@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