Workgroup: Network Working Group
Internet-Draft:
draft-per-app-networking-considerations-00
Published: 15 November 2020
Intended Status: Informational
Expires: 19 May 2021
Authors: L. Colitti  T. Pauly
Google  Apple Inc.

# Per-Application Networking Considerations

## Abstract

This document describes considerations for and implications of using application identifiers as a method of differentiating traffic on networks. Specifically, it discusses privacy considerations, possible mitigations, and considerations for user experience and API design.

## Discussion Venues

This note is to be removed before publishing as an RFC.

Source for this draft and an issue tracker can be found at https://github.com/tfpauly/per-app-networking-considerations.

## 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 19 May 2021.

## Copyright Notice

**Table of Contents**



## 1. Introduction

There are a number of use cases where network operators, or applications, might desire for application traffic to be treated differently by the network. Some examples are:

- Network-specific services. Applications might want to access local resources on a network that does not otherwise provide Internet access (for example, an entertainment system on an airplane).

- Per-application private networks. Certain applications, such as enterprise applications, might want to connect directly to the enterprise network in a secure fashion without using a device-wide VPN.

- Mobile network services. In mobile networks, applications like voice over LTE, IMS and RCS often use a different virtual network than general Internet traffic.

- Applications with specific performance requirements. Certain applications would benefit from particular scheduling or QoS policies - for example applications requiring low latency such as voice might be scheduled and queued differently from latency-insensitive traffic.

* Local breakout. In a mobile networks, applications might want to access resources through a different network interface (e.g., one that uses IPv6 addresses that are local to a specific area, and do not have a wide mobility).

* Zero-rating traffic. As allowed by regulators, certain classes of traffic (e.g., messaging or streaming video) might be exempt from metering on networks that are otherwise metered.

In existing networks, this is sometimes implemented by the network using deep packet inspection (e.g., flow tracking coupled with inspection of the SNI handshake). This is complex, implicates public policy concerns, and generally conflicts with the recommendations in [RFC7258]. The move towards encrypted protocols such as [RFC8484] and [I-D.ietf-tls-esni] will make this more difficult for some operators. Thus, if an application is to receive different treatment, the host or the application itself should be involved in requesting specific network treatment. This document explores the implications.

In this document, the term "application" refers to an application as understood by the user of the device.

### 1.1. Conventions and Definitions

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

## 2. Requesting differential treatment

There are already mechanisms for applications to request and obtain particular treatment by the network, or to communicate application identity to the network in order to obtain particular treatment. These include:

* Diffserv

* APN6

* Network tokens

* Slicing in 3GPP 5G networks

* Explicit application selection of a given Provisioning Domain (PvD) [RFC8801]

## 3. Open Internet implications

In certain regulatory regions, networks that provide general Internet access may not be permitted to discriminate between traffic sent to or from different lawful applications or websites, or such discrimination may be prohibited if commercially based. In a situation where the network operator has influence on the implementation of the user host (e.g., mobile networks where the handset is sold by the carrier), the device may be able to implement network policies directly, and thus may be impacted by neutrality considerations.

Neutrality concerns can be addressed by providing user control over assignment of particular applications to the particular network resources available to that user. Further, network neutrality implications may be reduced or avoided in some jurisdictions if the differential treatment occurs between different classes of traffic with different network requirements (e.g., bandwidth-intensive traffic vs. low-latency traffic) as opposed to between different applications with similar network requirements, and thus, by ensuring that the mechanism used to communicate requests to the network only specifies traffic classes and not individual applications.

## 4. Privacy implications

IETF guidance to avoid pervasive monitoring [RFC7258] is for network protocols to expose as little information as possible. Some of the proposed technologies for application signalling rely on the application exposing its identity to the network so that the network can then implement appropriate policies. This may provide the network with much more information than is needed to implement the desired behaviour. Information about which users are using specific applications, or visiting certain destinations, and when, can be highly privacy-sensitive.

Note that application identity can be exposed to the network even in the absence of explicit signalling. For example, if the host were to implement a network-set policy that requires that traffic from application X be sent on a different network path than all other traffic, the identity of application X would be exposed to the network as soon as it sends traffic.

Privacy concerns may also be reduced or avoided if the mechanism to request a different class of service only specifies the class of service (e.g., "low latency" or "streaming video") instead of the application originating the traffic.

In a situation where the network operator has influence over the implementation of the user host, the operator can still impose policies on what requests are possible - for example, the operator might choose to limit access to specialized services such as carrier messaging only to carrier applications. It is possible for such policies to preserve privacy if the policies specify general categories of traffic as opposed to specifying applications.

## 5. Mitigating implications via traffic categories

Many of these implications can be mitigated if the mechanism does not request different treatment of a service for a particular application, but instead specifies a general category of traffic, especially one that is defined based on traffic properties rather than commercial agreements.

Categories of traffic need to be sufficiently broad to not identify individual applications, and should be general enough that details about a user cannot be inferred merely by use of the category.

Consider the example a network that wants to provide differentiated service for a role-playing game application that can take advantage of a low-latency path. Several levels of categories could be defined. The following list shows some examples, in order of decreasing specificity:

1. Role-playing game

2. Game

3. Real-time/low-latency

The first category would not be an appropriate choice due to the privacy implications of identifying what kind of game a user plays. The second category is preferable, but the third is best since it defines a way to manage the network traffic without identifying anything about the content of the application.

Some use cases for traffic differentiation might need other kinds of categories. For example, operators might wish to zero-rate applications using categories based on payment tiers and rate-limiting.

## 6. User experience considerations

Privacy and neutrality concerns can be mitigated if the host's user is informed that particular applications are seeking or designated for particular treatment and consents to it. In order for consent to be meaningful, the user should be presented with a message that they understand. It may be difficult to balance the goal of providing

complete and accurate information with the goal of ensuring that the user understands the implications.

## 7. API considerations

It is desirable to provide an API layer that is not tied to specific network technologies (e.g., URSP, VPN, etc.). Having applications select a specific Provisioning Domain (PvD) could provide a useful layer of abstraction, as described in [I-D.ietf-taps-interface].

Any API should not involve revealing an application or user identity to the network via metadata without network authentication. Instead, the API should allow a given setting to be conditional on the identity of the network. For example, an application should express "use the zero-rated service for my app when on a particular carrier network", instead of blindly saying "this is my application identifier".

## 8. References

### 8.1. 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>.


### 8.2. Informative References


**[I-D.ietf-taps-interface]** Trammell, B., Welzl, M., Enghardt, T., Fairhurst, G., Kuehlewind, M., Perkins, C., Tiesel, P., Wood, C., and T. Pauly, "An Abstract Application Layer Interface to Transport Services", Work in Progress, Internet-Draft, draft-ietf-taps-interface-10, 2 November 2020, <http://www.ietf.org/internet-drafts/draft-ietf-taps-interface-10.txt>.

**[I-D.ietf-tls-esni]** Rescorla, E., Oku, K., Sullivan, N., and C. Wood, "TLS Encrypted Client Hello", Work in Progress, Internet-Draft, draft-ietf-tls-esni-08, 16 October 2020,

<http://www.ietf.org/internet-drafts/draft-ietf-tls-esni-08.txt>.

**[RFC7258]** Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", BCP 188, RFC 7258, DOI 10.17487/RFC7258, May 2014, <https://www.rfc-editor.org/info/rfc7258>.

**[RFC8484]** Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <https://www.rfc-editor.org/info/rfc8484>.

**[RFC8801]** Pfister, P., Vyncke, É., Pauly, T., Schinazi, D., and W. Shao, "Discovering Provisioning Domain Names and Data", RFC 8801, DOI 10.17487/RFC8801, July 2020, <https://www.rfc-editor.org/info/rfc8801>.

## Acknowledgments

Thanks to Adi Masputra and Elliot Briggs for their inputs to this discussion.

## Authors' Addresses

Lorenzo Colitti
Google
Shibuya 3-21-3,
Japan

Email: lorenzo@google.com

Tommy Pauly
Apple Inc.
One Apple Park Way
Cupertino, California 95014,
United States of America

Email: tpauly@apple.com