Workgroup:
Transport and Services Working Group
Internet-Draft:
draft-bwbr-tsvwg-signaling-use-cases-00
Published: 4 March 2024
Intended Status: Informational
Expires: 5 September 2024
Authors: S. Rajagopalan
Cloud Software Group
D. Wing
Cloud Software Group
M. Boucadair
Orange
T. Reddy
Nokia

Signaling Use Cases for Traffic Differentiation

Abstract

Host-to-network signaling can improve the user experience by informing the network which flows are more important and which packets within a flow are more important. The differentiated service may be provided at the network (e.g., packet prioritization), the sender (e.g., adaptive transmission), or through cooperation of both the sender and the network.

This document outlines use-cases that highlight the need for a new signaling protocol from the receiver to its network elements which enables different traffic treatment.

About This Document

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

The latest revision of this draft can be found at https://danwing.github.io/signaling-use-cases/draft-wing-tsvwg-signaling-use-cases.html. Status information for this document may be found at https://datatracker.ietf.org/doc/draft-bwbr-tsvwg-signaling-use-cases/.

Discussion of this document takes place on the Transport and Services Working Group Working Group mailing list (mailto:tsvwg@ietf.org), which is archived at https://mailarchive.ietf.org/arch/browse/tsvwg/. Subscribe at https://www.ietf.org/mailman/listinfo/tsvwg/.

Source for this draft and an issue tracker can be found at https://github.com/danwing/signaling-use-cases.

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 5 September 2024.

Table of Contents

1. Introduction

Bandwidth constraints exist most predominently at the access network. Users of those networks run various hosts which run various applications, each having different needs for the best user experience. These requirements are not fixed but change over time

depending on the application and even depending on how the application is used.

The simple network diagram below shows where such bandwidth and performance constraints usually exist with a "B".

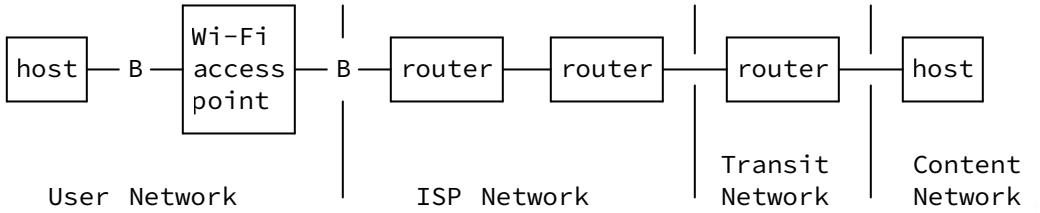

For traffic sent in either direction, the network network element(s) immediately prior to the bandwidth constraining link can be augmented with flow metadata. Such augmentation allows those network elements to make autonomous decisions to prioritize, delay, or drop packets especially to when performing Reactive Management.

A difficulty with this metadata augmentation is deciding which metadata to trust. Traffic arriving from a content provider cannot be differentiated from traffic arriving from other hosts on the Internet. The metadata signals from the content provider are more likely to be authentic but the metadata signals from other hosts may be wrong, undesired by the local host, or maliciously contain improper metadata. Attempts to automate identification of content providers have included HTTP "Host" header inspection and TLS SNI inspection which are expected to fail as encrypted SNI and privacy-enhancin MASQUE proxies become more prevalant. A remaining mechanism to authorize metadata signals from the content provider is to configure the ISP equipment with the content network's source IP addresses and provide only that traffic with differentiated service. However, such an arrangement benefits large players (large ISPs and large content network) and disadvantages small players (and new players). A more egalitarian approach would provide the same benefit to all parties -- large and small -- and also provide richer signaling to further improve user experience and metadata interoperability. This would allow all parties to become part of the "Internet fast lane".

The authorization problem exists with technologies as relatively simple as DiffServ and the problem persists with many other recently discussed metadata signaling mechanisms including embedding information in the UDP payload ([I-D.draft-trammell-plus-spec]), UDP options ([I-D.draft-kaippallimalil-tsvwg-media-hdr-wireless]), overloading the IPv6 Flow Label ([I-D.draft-cc-v6ops-wlcg-flow-label-marking], and Hop-by-Hop

Options. One mechanism suggested occasionally is to encrypt or integrity protect the metadata with a key; such a key could be established with a signaling protocol, see Section 3.3.

There is consensus that applications can benefit by signaling the network ([IAB], [ATIS]). This document provides use-cases to further detail the need of such signaling.

2. Conventions and Definitions

Intentional Management: network policy such as (monthly) bandwidth quota or bandwidth limit, or quality (delay and/or jitter)) assurances.

Reactive Management: network reactions to congestion events, with very short to very long durations (e.g., varying wireless and mobile air interface conditions).

3. Use Cases

3.1. Priority Between Flows (Inter-Flow)

Certain flows being received by an host or by an application are less or more important than other flows. For example, a host downloading a software update is generally considered less important than another host doing interactive audio/video or gaming. By signaling the relative importance of flows to a network element (e.g., router, MASQUE proxy), the network element can (de-)prioritize those flows to best accomodate the needs of the various applications (on a single host) and between hosts on a network.

3.1.1. Abuse and Constraints

It is important that not every flow be prioritized; otherwise, the network devolves into the best-effort network that existed prior to metadata signaling. It is a requirement that mechanisms exist to prevent this occurrence. The mechanism might be simple, for example a cellular network might allow one flow from a subscriber to declare itself as important; other flows with that subscriber are denied attempts to prioritize themselves. The mechanism might be more complex where authentication and authorization is performed by an enterprise network which, itself, decides which flows are important based on its policy and only the enterprise network communicates flow priorities to the ISP network. The enterprise might prioritize certain users (e.g., IT staff, CEO), certain equipment (audio/video equipment in a conference room), or whatever its policies it might want.

3.1.1.1. Interactive Media

Examples: VoIP, gaming, virtual desktop.

Requirement: Signal the flow needs low jitter and low delay. However, the network can only provide a limited amount of low jitter/low delay to each host, maybe as few as one. This requires signaling feedback indicating that low jitter and low delay flows are already subscribed to other hosts. In response, the user and the application will likely continue, occasionally re-attempting to get the desired quality of service from the network.

Todo: this section on cooperation needs editing.

3.1.1.2. Bulk Data Transfer

Examples: backup/restore, software update, RSS feed update, email

Requirement: Signal the flow as below best-effort.

3.2. Priority Within a Flow (Intra-Flow)

Interactive Audio/Video has long been using [RTP] which runs over UDP. As described in Section 2.3.7.2 of [RFC7478], there is value in differentiating between voice, video and data. Today's video streaming is exclusively over TCP but will migrate to QUIC and eventually is likely to support unreliable transport ([RFC9221], [I-D.draft-kpugin-rush]). With unreliable transport of video in RTP or QUIC, it is beneficial to differentiate the important video keyframes from other video frames. Other applications such as gaming and remote desktop also benefit from differentiating their packets to the network.

Many of these flows do not originate from a content provider's network. Thus, the flows originate from an IP address that is not known before connection establishment, so there needs to be a way for the client to authorize the network elements to honor the metadata of those packets.

3.3. Key Establishment

Various proposals have suggested establishing a key to validate per-packet metadata or to decrypt per-packet metadata. However, most proposals have not specified how this key would be established. A signaling protocol from the receiving host to its ISP could establish such a key. The host can then convey the key to the

sending host to use to integrity protect or encrypt the per-packet metadata.

> Note: The CPU overhead of validating or decrypting such per-packet metadata needs to be carefully considered by the signaling protocol proposing such keying.

3.4. Metadata Version/Capability Exchange

The sender has to convey metadata in a way that is understood by the various network elements on the path -- each of which might be operated by different entities and have different capabilities. For example, the Wi-Fi access point might be operated by an enterprise network, hotel, or home user, whereas the ISP's router is operated by the ISP. Each of those might support different versions of the same metadata, or might need the metadata expressed in different ways.

The signaling protocol would provide a way to learn the needs of those networks, and provide metadata signaling satisfying most or all of their needs.

4. Requirements Summary

TODO summary.

5. Security Considerations

TODO Security

6. IANA Considerations

This document has no IANA actions.

7. Informative References

[ATIS] "Content Classification for Traffic Optimization", 2023, <https://access.atis.org/higherlogic/ws/public/download/72240>.

[I-D.draft-cc-v6ops-wlcg-flow-label-marking] Carder, D. W., Chown, T., McKee, S., and M. Babik, "Use of the IPv6 Flow Label for WLCG Packet Marking", Work in Progress, Internet-Draft, draft-cc-v6ops-wlcg-flow-label-marking-02, 10 July 2023, <https://datatracker.ietf.org/doc/html/draft-cc-v6ops-wlcg-flow-label-marking-02>.

[I-D.draft-kaippallimalil-tsvwg-media-hdr-wireless] Kaippallimalil, J., Gundavelli, S., and S. Dawkins, "Media Handling Considerations for Wireless Networks",

Work in Progress, Internet-Draft, draft-kaippallimalil-tsvwg-media-hdr-wireless-04, 14 February 2024, <https://datatracker.ietf.org/doc/html/draft-kaippallimalil-tsvwg-media-hdr-wireless-04>.

[I-D.draft-kpugin-rush] Pugin, K., Frindell, A., Ferret, J. C., and J. Weissman, "RUSH - Reliable (unreliable) streaming protocol", Work in Progress, Internet-Draft, draft-kpugin-rush-02, 10 May 2023, <https://datatracker.ietf.org/doc/html/draft-kpugin-rush-02>.

[I-D.draft-trammell-plus-spec] Trammell, B. and M. Kühlewind, "Path Layer UDP Substrate Specification", Work in Progress, Internet-Draft, draft-trammell-plus-spec-01, 13 March 2017, <https://datatracker.ietf.org/doc/html/draft-trammell-plus-spec-01>.

[IAB] Arkko, J., Hardie, T., Pauly, T., and M. Kühlewind, "Considerations on Application - Network Collaboration Using Path Signals", RFC 9419, DOI 10.17487/RFC9419, July 2023, <https://www.rfc-editor.org/rfc/rfc9419>.

[RFC7478] Holmberg, C., Hakansson, S., and G. Eriksson, "Web Real-Time Communication Use Cases and Requirements", RFC 7478, DOI 10.17487/RFC7478, March 2015, <https://www.rfc-editor.org/rfc/rfc7478>.

[RFC9221] Pauly, T., Kinnear, E., and D. Schinazi, "An Unreliable Datagram Extension to QUIC", RFC 9221, DOI 10.17487/RFC9221, March 2022, <https://www.rfc-editor.org/rfc/rfc9221>.

[RTP] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, DOI 10.17487/RFC3550, July 2003, <https://www.rfc-editor.org/rfc/rfc3550>.

Acknowledgments

TODO acknowledge.

Authors' Addresses

Sridharan Rajagopalan
Cloud Software Group Holdings, Inc.
United States of America

Email: sridharan.girish@gmail.com

Dan Wing

Cloud Software Group Holdings, Inc.
United States of America

Email: danwing@gmail.com

Mohamed Boucadair
Orange
France

Email: mohamed.boucadair@orange.com

Tirumaleswar Reddy
Nokia
India

Email: kondtir@gmail.com