Network Working Group T. Fossati
Internet-Draft Nokia
Intended status: Standards Track G. Fairhurst
Expires: June 19, 2019 University of Aberdeen
P. Aranda Gutierrez
Universidad Carlos III de Madrid
M. Kuehlewind
ETH Zurich
December 16, 2018

A Loss-Latency Trade-off Signal for the Mobile Network
draft-fossati-tsvwg-lola-00

Abstract

This document proposes a marking scheme for tagging low-latency flows (for example: interactive voice and video, gaming, machine to machine applications) that is safe to use by the mobile network for matching such flows to suitable per-hop behaviors (EPS bearers defined by 3GPP) in its core and radio segments. The suggested scheme re-uses NQB, a DiffServ-based signalling scheme with compatible rate-delay trade-off semantics that has been recently introduced in the context of fixed access to allow differential treatment of non-queue building vs queue building flows.

Table of Contents

1. Introduction

Today's mobile networks are configured to bundle all flows to and from the Internet into a single "default" EPS bearer whose buffering characteristics are not compatible with low-latency traffic. The established behaviour is partly rooted in the desire to prioritise operators' voice services over competing over-the-top services. Of late, said business consideration seems to have lost momentum and the incentives might now be aligned towards allowing a more suitable treatment of Internet real-time flows. However, a couple of preconditions need to be satisfied before we can move on from the status quo. First, the real-time flows must be efficiently identified so that they can be quickly assigned to the "right" EPS bearer. This is especially important with the rising popularity of encrypted and multiplexed transports, which has the potential of increasing the cost/accuracy ratio of Multi-Field (MF) based classification over the acceptable threshold. Second, the signal must be such that malicious or badly configured nodes can't abuse it. Today's mobile networks take a rather extreme posture in this respect by actively discarding (remarking or bleaching [Custura]) DiffServ signalling coming from an interconnect. Therefore, the signal must

be modelled in a way that the mobile network can either trust it or, even better, avoid the need of trusting it in the first place. The Rate-Delay trade-off [Podlesny] satisfies the above requirements and has been shown [Fossati] to integrate well with a simplified LTE QoS setup that uses one dedicated low-latency bearer in addition to the default.

This document suggests reusing the Non Queue Building (NQB) signalling protocol described in [I-D.white-tsvwg-nqb] as the method employed by endpoints to mark their real-time flows and by the LTE network to classify and route these flows via a suitable (low-latency) bearer through the LTE core network and E-UTRAN.

2. Terminology

- DPI: Deep Packet Inspection

- EPS bearer: Evolved Packet System bearer, a virtual circuit with a given set of QoS attributes which spans the entire mobile network including the LTE core and E-UTRAN segments;

- GBR: Guaranteed Bit Rate. EPS bearers can be GBR, in which case they are guaranteed to not drop packets under congestion, or non-GBR, in which case no guarantee of delivery is made by the mobile network;

- LTE: 3GPP Long Term Evolution, aka 4G;

- E-UTRAN: LTE Radio Access Network;

- QCI: QoS Class Identifier. In LTE networks, EPS bearers are partitioned into equivalency classes modulo the QoS treatment they receive. QCI is an integer that labels a specific QoS class. Its semantics is consistently understood by all network elements involved in packet forwarding;

- UE: User Equipment, any device (e.g., smartphone, laptop, tablet) attached to an LTE network.

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.

3. DiffServ Code

Given the semantic equivalence of a Loss-Latency trade-off with the Non Queue Building (NQB) behaviour, this document reuses the NQB DSCP (Section 4 of [I-D.white-tsvwg-nqb]) as-is.

4. Marking

Endpoints SHOULD mark packets that belong to the Best Effort class and are latency sensitive by assigning the NQB DSCP value to the DS field.

The marking could also be added to other BE traffic if the default class could be reclassified by the network to use the NQB DSCP before the packet enters the mobile domain - for example by a classifier in the SGi-LAN or in an LTE router.

5. Relationship to a Mobile DiffServ Domain

The Mobile network is configured to give UEs a dedicated, low-latency, non-GBR, EPS bearer with QCI 7 in addition to the default EPS bearer.

A packet carrying the NQB DSCP shall be routed through the dedicated low-latency EPS bearer. A packet that has no associated NQB marking shall be routed through the default EPS bearer.

6. On Remarking and Bleaching

NQB markings SHOULD be preserved when forwarding via an interconnect.

The NQB DSCP can have end-to-end semantics and this might benefit any NQB-marked traffic if utilised by other path elements (e.g. allowing DS treatment if a bottleneck link happens to be in the part of the path outisde the mobile access network segment).

7. IANA Considerations

This document makes no request to IANA.

8. Security Considerations

Internet applications cannot benefit from wrongly indicating low-latency as they have to pay the expense of high loss as a trade-off. Hence there is no incentive for Internet applications to set the wrong code-point.

The NQB signal is not integrity protected and could be flipped by an on-path attacker. This might negatively affect the QoS of the tampered flow.

9. Privacy Considerations

As described in [Shbair] state of art encrypted traffic analysis based machine learning can successfully identify the type of transported application (e.g., HTTPS, SMTP, P2P, VoIP, SSH, Skype) with good accuracy and without any need to access the clear-text. In this context, despite it being coarse grained, a 1-bit signal such as the one described in this document might be used to improve the precision of the classifier.

10. Acknowledgments

We would like to thank the authors of the "Latency Loss Tradeoff PHB Group" draft: Jianjie You, Michael Welzl, Brian Trammell and Kevin Smith. Big thanks to Chris Seal, Dan Druta, Diego Lopez, Shamit Bhat, Georg Mayer, Florin Baboescu, James Gruessing for the help.

This work is partially supported by the European Commission under Horizon 2020 grant agreement no. 688421 Measurement and Architecture for a Middleboxed Internet (MAMI), and by the Swiss State Secretariat for Education, Research, and Innovation under contract no. 15.0268. This support does not imply endorsement.

11. References

11.1. Normative References

[I-D.white-tsvwg-nqb]
White, G., "Identifying and Handling Non Queue Building Flows in a Bottleneck Link", draft-white-tsvwg-nqb-00 (work in progress), October 2018.

[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>.

11.2. Informative References

[Custura] Custura, A., Venne, A., and G. Fairhurst, "Exploring DSCP modification pathologies in mobile edge networks", TMA , 2017.

[Fossati] Fossati, T., Aranda Gutierrez, P., Kuehlewind, M., and D. Lopez, "Loss Latency Tradeoff and the Mobile Network", 2018, <https://github.com/mami-project/roadshows/blob/master/ietf-103-lola/hotrfc/build/main.pdf>.

[Podlesny] Podlesny, M. and S. Gorinsky, "Rd Network Services: Differentiation Through Performance Incentives", SIGCOMM , 2008.

[Shbair] Shbair, W., Cholez, T., Francois, J., and I. Chrisment, "A Multi-Level Framework to Identify HTTPS Services", NOMS , 2016.

Authors' Addresses

Thomas Fossati
Nokia

Email: thomas.fossati@nokia.com

Gorry Fairhurst
University of Aberdeen

Email: gorry@erg.abdn.ac.uk

Pedro A. Aranda Gutierrez
Universidad Carlos III de Madrid

Email: paranda@it.uc3m.es

Mirja Kuehlewind
ETH Zurich

Email: mirja.kuehlewind@tik.ee.ethz.ch