

IPPM
Internet-Draft
Intended status: Informational
Expires: April 28, 2022

M. Cociglio
M. Nilo
F. Bulgarella
Telecom Italia - TIM
G. Fioccola
Huawei Technologies
October 25, 2021

User Devices Explicit Monitoring
draft-cnbf-ippm-user-devices-explicit-monitoring-03

Abstract

This document describes a methodology to monitor network performance exploiting user devices. This can be achieved using the Explicit Flow Measurement Techniques, protocol independent methods that employ few marking bits, inside the header of each packet, for loss and delay measurement. User devices and servers, marking the traffic, signal these metrics to intermediate network observers allowing them to measure connection performance, and to locate the network segment where impairments happen. In addition or in alternative to network observers, a probe can be installed on the user device with remarkable benefits in terms of hardware deployment and measurement scalability.

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 April 28, 2022.

Copyright Notice

Copyright (c) 2021 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	Notational Conventions	3
3.	Explicit Performance Open Issues	3
4.	Explicit Performance Probes on User Devices	3
5.	Device Owner Activates Explicit Performance Measurements . .	4
6.	Who Will Handle the Performance Data?	4
7.	The Explicit Performance App	5
8.	Improvements of Explicit Flow Measurement Techniques Using Probes on User Devices	5
8.1.	Hidden Delay Bit Considerations	5
9.	Security Considerations	6
10.	Privacy Considerations	6
11.	IANA Considerations	6
12.	Change Log	6
13.	Contributors	6
14.	Acknowledgements	6
15.	References	6
15.1.	Normative References	6
15.2.	Informative References	7
	Authors' Addresses	7

[1.](#) Introduction

Explicit Performance Monitoring enables a passive observer (a probe) to measure delay and loss just watching the marking (a few header bits) of live traffic packets. It works on client-server protocols: e.g. QUIC [[QUIC-TRANSPORT](#)], TCP [[TCP](#)]. The different methods are described in [[EXPLICIT-FLOW-MEASUREMENTS](#)] and are inspired by [[AltMark](#)].

This document explains how to employ the methods described in [[EXPLICIT-FLOW-MEASUREMENTS](#)] by proposing the user device as a convenient place for the Explicit Performance Observer.

2. Notational Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

3. Explicit Performance Open Issues

There are some open issues to consider for the deployment of [[EXPLICIT-FLOW-MEASUREMENTS](#)]:

- Who decides whether to mark traffic? Explicit measures only work if both the server and the client mark the production traffic.
- What about scalability? Could network probes monitor all the connections? If they cannot, which ones to choose?
- Which connections to monitor within the network? Network probes need an effective way to identify which connections really need to be monitored.
- How to monitor both traffic directions? Not always possible for network probes (asymmetric connections).

4. Explicit Performance Probes on User Devices

This document proposes the user device (e.g. mobile phones, PCs) as a convenient place where to put the Explicit Performance Observer.

The placement of the observer on the user device helps to mitigate the issues reported in the previous section, in particular:

- The device should decide whether to mark the traffic or not.
- Regarding the scalability issue, on the user device there are few connections to monitor so it becomes less relevant.
- Connections eligible for monitoring should be the impaired ones. User devices and network probes can cooperate to achieve this goal. It is possible to set alarm thresholds on the user device and to signal to the network probes only the sessions with impairments. This allows to segment the performance measurements and to locate the faults. In this way network probes, that could also be embedded into network nodes, have to monitor a limited number of connections.
- Monitoring both directions is always possible on the user device.

5. Device Owner Activates Explicit Performance Measurements

The decision whether to activate the marking (e.g. [[SPIN-BIT](#)], [[ANRW19-PM-QUIC](#)], [[EXPLICIT-FLOW-MEASUREMENTS](#)]) or not should be made by the device owner by properly configuring the applications (e.g. browsers) based on connection-oriented protocols that support explicit measurements (e.g. QUIC).

All applications should provide the activation or deactivation of packet marking, for example by providing an user interface or exposing API.

So, during the client-server handshake, the client will decide whether the marking is active or not within a session and notify its decision to the server.

An example of a simple explicit marking agreement of a protocol is the following. This works if the usage of each performance bit is unique and predefined. An endpoint set to 0 all the explicit performance measurement bits to indicate its intention not to mark. Then:

- the client set at least one of its marking bits to 1 notifying the server of its intention to use that/those marking bits; the server adapts according to the client's will;
- the server set at least one of its marking bits to 1; if the client does not start marking the same bit/bits, then the marking for that/those bits is aborted.

The best would be if both client and server started using the same marking bits from the beginning of the connection. In this case no alignment between endpoints would be required. This mechanism works best if, where possible, measurements start using 1 as the first marking value.

6. Who Will Handle the Performance Data?

Performance data are stored only on the user device or also sent to "external bodies" according to the will of the device owner.

The main recipient would be the Internet Service Provider. Indeed, as explained in the previous section, this enables user device and network probes coordination that permits an improved performance measurement approach.

Moreover these data could also be of interest for the national regulatory authorities or others authorized subjects.

7. The Explicit Performance App

This methodology could be implemented with an "Explicit Performance App" installed on the user device.

The App should perform the following tasks:

- collect user preferences;
- activate/deactivate marking on device Apps (e.g. browsers);
- implement the observer;
- show performances to the user;
- send data to the "Explicit Performance Management Center";
- set performance thresholds.

8. Improvements of Explicit Flow Measurement Techniques Using Probes on User Devices

- Spin bit and Delay bit: the observer-server RTT component measured on the user device is equivalent to the RTT, but without including the client-side application delay and therefore more precise.
- sSquare bit: would measure the End-to-End loss rate in the download direction instead of upstream loss rate.
- Loss event bit: would measure, as before, the End-to-End loss rate in both directions. Moreover, in the upload direction, the signal would be "clean" since it is captured at the origin and therefore not affected by losses.
- Reflection square bit: would measure the RT loss rate instead of three-quarters connection loss rate.

8.1. Hidden Delay Bit Considerations

The Explicit Flow Measurements draft introduces a new Delay Bit feature capable of masking the RTT of the connection to the observers on the network. To use this feature, the client must select an Additional Delay used to delay the client-side reflection of marked samples. Clearly, the introduction by the client of a reflection delay makes the client-observer component of the RTT inaccurate.

Using this feature on a user device probe has several advantages:

- A system-wide Additional Delay can be selected and periodically updated making it common to all applications installed on the device.
- The hidden Delay Bit produces the same metrics of the Delay Bit since the observer-server RTT, measured on the client, is equal to the end-to-end RTT of the connection.
- The user device can easily communicate the Additional Delay to network probes whenever an alarm threshold is triggered. In this way, the observer can compute the e2e RTT of the connection.

9. Security Considerations

TBD

10. Privacy Considerations

TBD

11. IANA Considerations

This document makes no request of IANA.

12. Change Log

TBD

13. Contributors

TBD

14. Acknowledgements

TBD

15. References

15.1. 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>>.
- [TCP] Postel, J., "Transmission Control Protocol", STD 7, [RFC 793](#), DOI 10.17487/RFC0793, September 1981, <<https://www.rfc-editor.org/info/rfc793>>.

15.2. Informative References

- [AltMark] Fioccola, G., Ed., Capello, A., Cociglio, M., Castaldelli, L., Chen, M., Zheng, L., Mirsky, G., and T. Mizrahi, "Alternate-Marking Method for Passive and Hybrid Performance Monitoring", [RFC 8321](#), DOI 10.17487/RFC8321, January 2018, <<https://www.rfc-editor.org/info/rfc8321>>.
- [ANRW19-PM-QUIC] Bulgarella, F., Cociglio, M., Fioccola, G., Marchetto, G., and R. Sisto, "Performance measurements of QUIC communications", Proceedings of the Applied Networking Research Workshop, DOI 10.1145/3340301.3341127, July 2019.
- [EXPLICIT-FLOW-MEASUREMENTS] Cociglio, M., Ferrieux, A., Fioccola, G., Lubashev, I., Bulgarella, F., Hamchaoui, I., Nilo, M., Sisto, R., and D. Tikhonov, "Explicit Flow Measurements Techniques", [draft-ietf-ippm-explicit-flow-measurements-00](#) (work in progress), October 2021.
- [QUIC-TRANSPORT] Iyengar, J., Ed. and M. Thomson, Ed., "QUIC: A UDP-Based Multiplexed and Secure Transport", [RFC 9000](#), DOI 10.17487/RFC9000, May 2021, <<https://www.rfc-editor.org/info/rfc9000>>.
- [SPIN-BIT] Trammell, B., Vaere, P. D., Even, R., Fioccola, G., Fossati, T., Ihlar, M., Morton, A., and E. Stephan, "Adding Explicit Passive Measurability of Two-Way Latency to the QUIC Transport Protocol", [draft-trammell-quick-spin-03](#) (work in progress), May 2018.

Authors' Addresses

Mauro Cociglio
Telecom Italia - TIM
Via Reiss Romoli, 274
Torino 10148
Italy

E-Mail: mauro.cociglio@telecomitalia.it

Massimo Nilo
Telecom Italia - TIM
Via Reiss Romoli, 274
Torino 10148
Italy

EMail: massimo.nilo@telecomitalia.it

Fabio Bulgarella
Telecom Italia - TIM
Via Reiss Romoli, 274
Torino 10148
Italy

EMail: fabio.bulgarella@guest.telecomitalia.it

Giuseppe Fioccola
Huawei Technologies
Riesstrasse, 25
Munich 80992
Germany

EMail: giuseppe.fioccola@huawei.com

