

CORE Working Group
Internet-Draft
Intended status: Standards Track
Expires: 27 April 2023

G. Fioccola
T. Zhou
Huawei
M. Cociglio
F. Bulgarella
M. Nilo
Telecom Italia
24 October 2022

Constrained Application Protocol (CoAP) Performance Measurement Option draft-fz-core-coap-pm-03

Abstract

This document specifies a method for the Performance Measurement of the Constrained Application Protocol (CoAP). A new CoAP option is defined in order to enable network telemetry both end-to-end and on-path.

Requirements Language

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 [RFC 2119](#) [[RFC2119](#)].

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 27 April 2023.

Copyright Notice

Copyright (c) 2022 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 [2](#)
- [2.](#) Performance Measurement methods for CoAP [3](#)
 - [2.1.](#) sSquare bit and Spin bit [4](#)
 - [2.2.](#) Combined sSquare bit [4](#)
- [3.](#) CoAP Performance Measurement Option [5](#)
- [4.](#) Structure of the PM Option [6](#)
- [5.](#) Application Scenarios [8](#)
 - [5.1.](#) Non-proxying endpoints [8](#)
 - [5.2.](#) Collaborating or Non-collaborating proxies [9](#)
 - [5.3.](#) OSCORE [11](#)
- [6.](#) Management and Configuration [12](#)
- [7.](#) Congestion Control [12](#)
- [8.](#) Security Considerations [12](#)
- [9.](#) IANA Considerations [13](#)
- [10.](#) Acknowledgements [13](#)
- [11.](#) References [13](#)
 - [11.1.](#) Normative References [13](#)
 - [11.2.](#) Informative References [13](#)
- Authors' Addresses [14](#)

[1.](#) Introduction

[RFC7252] define the CoAP Protocol. In CoAP, reliability is provided by marking a message as Confirmable (CON) with ACKs. A message that does not require reliable transmission can be sent as a Non-confirmable message (NON).

In case of CoAP reliable mode there are Message IDs and ACKs, that could eventually be used to measure Round-Trip Time (RTT) and losses. But it can be resource-consuming for constrained nodes since they have to look at Message IDs and take timestamps. These operations are expensive in terms of resources. In case of CoAP unreliable mode, there is no ACK and, consequently, it is not possible to measure RTT and losses.

Thus, there is no easy way to measure the performance metrics in COAP environment to satisfy the low resources of constrained nodes. And it is in any case limited to RTT and end-to-end losses.

A mechanism to measure the performance in CoAP can be useful to verify and meet the operational requirements, but it should be a simple mechanism for network diagnostic to be developed on constrained nodes requiring just a minimal amount of collaboration from the endpoints.

[I-D.ietf-ippm-explicit-flow-measurements] describes the methodologies for Explicit Flow Measurement (EFM). The EFM techniques employ few marking bits, inside the header of each packet, for loss and delay measurement. These are relevant for encrypted protocols, e.g. QUIC [[RFC9000](#)], where there are only few bits available in the non-encrypted header in order to allow passive performance metrics from an on-path observer. These methodologies could potentially be used and extended in CoAP.

[I-D.ietf-ippm-explicit-flow-measurements] defines different combinations of bits because the number of bits in QUIC is limited and different experiments have been done. But all these methods together imply complex algorithms that do not apply well to the CoAP environment.

This document aims to create an easy way to allow performance measurement for CoAP, by defining a new option, called Performance Measurement (PM) CoAP Option. The CoAP performance metrics (e.g. RTT and losses) can be useful for an operator or an enterprise that is managing a constrained, low-power and lossy network.

2. Performance Measurement methods for CoAP

CoAP [[RFC7252](#)] defines a number of options that can be included in a message. For this reason, a new option for CoAP, carrying Performance Measurement (PM) bits is the approach followed by this document.

The PM bits that are included in the Option are:

- * sQuare bit (Q bit), based on [[I-D.ietf-ippm-rfc8321bis](#)] and further described in [[I-D.ietf-ippm-explicit-flow-measurements](#)];
- * Spin bit (S bit), described in [[RFC9312](#)] and included as optional bit in [[RFC9000](#)];
- * Loss and Delay event information for further usage.

A requirement to enable PM methods in COAP environment is that the methodologies and the algorithm needs to be kept simple. For this reason, the idea is to re-apply only the S bit and Q bit.

Thus, the advantages of using the CoAP PM Option are:

- 1) Simplification because it is not needed to read Message IDs, indeed there is a well-defined sSquare wave, and it is not necessary to store timestamps, since the duration of the Spin Bit period is equal to RTT.
- 2) Enabling easy on-path observer (proxy, gateway) metrics.

2.1. sSquare bit and Spin bit

The sSquare bit algorithm consists of creating square waves of a known length (e.g. 64 packets). Each side of the connection can set the Q bit and toggle its value every fixed number of packets. The number of packets can be easily recognized and packet loss can be measured.

The Spin bit algorithm consists of creating a square wave signal on the data flow, using a bit, whose length is equal to RTT. The Spin bit causes one bit to 'spin', generating one edge (a transition from 0 to 1 or from 1 to 0) once per end-to-end RTT. The Spin bit is set by both sides to the same value for as long as one round trip lasts and then it toggles the value.

All the possible measurements (end-to-end, hop-by-hop) that are enabled by Q bit and S bit are detailed in [\[I-D.ietf-ippm-explicit-flow-measurements\]](#).

2.2. Combined sSquare bit

The synergy between S bit and Q bit is also possible. As described above, the length of the Q bit square waves is fixed (e.g. a predefined number of packets) in this way each endpoint can detect a packet loss if it receives less packets than expected. It is possible to potentiate the Q bit signal by incorporating RTT information as well. This implies a little modification to the algorithm of the Q bit that could also be used alone:

A single packet in a period of the square wave can be selected and set to the opposite value of that period. After one RTT it comes back and another packet is selected and set again to the opposite value of that period. And the process can start again. By measuring the distance between these special packets, it is possible to measure the RTT in addition to packet loss. The periods with the special packets have one packet less than expected but this is easy to recognize and to take into account by both endpoints.

This mechanism uses a single bit that serves two purposes: a loss indicator and a delay indicator. It is worth highlighting that the mechanism is similar to the Delay bit (D bit), described in [\[I-D.ietf-ippm-explicit-flow-measurements\]](#). Indeed, the Delay bit is set only once per RTT and a single packet with the marked Delay bit bounces between a client and a server.

The Q bit and D bit signals use two single bit values and the new signal is a Combined square bit (C bit) signal. The C bit value is given by an Exclusive OR operation (XOR) between the two Q bit and D bit values: $C = Q \text{ XOR } D$.

Since C bit incorporates both Q bit and D bit information, the same considerations for the two separate signals in [\[I-D.ietf-ippm-explicit-flow-measurements\]](#) can also be extended in the case of C bit. Therefore, all the possible measurements (end-to-end, hop-by-hop) that are enabled by using only C bit can be found in [\[I-D.ietf-ippm-explicit-flow-measurements\]](#) by merging Q bit and D bit derived measurements.

3. CoAP Performance Measurement Option

Figure 1 shows the property of the CoAP Performance Measurement (PM) Option. The formatting of this table is reported in [\[RFC7252\]](#). The C, U, N, and R columns indicate the properties Critical, Unsafe, NoCacheKey, and Repeatable as defined in [\[RFC7252\]](#). None of these properties is marked for the PM options.

Number	C	U	N	R	Name	Format	Length	Default
TBD			x		PM	uint	1	0

Figure 1: CoAP PM Option Properties

The CoAP PM Option is Elective, Safe-to-Forward and it is not to be included in the Cache-Key (NoCacheKey is set).

Note that it could be possible to make use of one bit in the option to identify the mode. In this way two patterns can be defined.

4. Structure of the PM Option

The value of the PM option is a 1 byte unsigned integer. This integer value encodes the following fields:




Figure 2: CoAP Performance Measurement Option

Where:

- * The Mode bit (M bit) can be set to 1 or 0 and it is used to identify whether the Option follows pattern 1 (M bit = 0) or pattern 2 (M bit = 1).
- * Pattern bits can be of two kinds as reported below.

The PM Option can employ two patterns based on the value of the M bit:




Figure 3: CoAP Performance Measurement Option pattern 1




Figure 4: CoAP Performance Measurement Option pattern 2

The COAP Option could be defined with 2 PM bits (S and Q) or defined with a single PM bit (C bit).

Where:

- * C bit is used in pattern 1. It is based on the enhancement of the Q bit signal with the S bit information. The two methods are described in [[I-D.ietf-ippm-explicit-flow-measurements](#)] and coupled as detailed in [Section 2.2](#);
- * Q bit is used in pattern 2. It is described in [[I-D.ietf-ippm-explicit-flow-measurements](#)];
- * S bit is used in pattern 2. It is described in [[RFC9312](#)] and also embedded in the QUIC Protocol [[RFC9000](#)];
- * Event bits MAY encode additional Loss and Delay information based on well-defined encoding and they can also be used by on-path observers. If these Event bits are all zero, they MUST be ignored on receipt.

It is worth noting that the only differences between the two patterns are related to the accuracy of the measurements. Further details can be found in [[I-D.ietf-ippm-explicit-flow-measurements](#)].

The Event bits can be divided into two parts, for instance: loss event bits and delay event bits. Based on the average RTT, an end point can define different levels of thresholds and set the delay event bits accordingly. The same applies to loss event bits. In this way an on-path observer becomes aware of the network conditions by simply reading these Event bits.

The on-path observer can read the event signaling bits and could be the Proxy or the Gateway which interconnects disjointed CoAP networks. It MAY communicate with Client and Server to set some parameters such as timeout based on the network performance.

The CoAP PM Options described in this document can be used in both requests and responses. If a CoAP endpoint does not implement the measurement methodologies, it can simply leave the default value (all bits are zero). In this way the other CoAP endpoints become aware that the measurement cannot be executed in that case.

The fixed number of packets to create the Q bit (or C bit) signal is predefined and its value is configured from the beginning for all the CoAP endpoints.

5. Application Scenarios

The main usage of the CoAP PM Options is to do end-to-end measurement between the client and the server but it can also allow split measurements. The on-path measurement is the additional feature. This information can be used to monitor the network in order to check the operational performance and to employ further network optimization.

The intermediaries or on-path observers could be:

Network Functions or Probes that must be able to see deep into application.

Gateway or Proxies that, as specified in [\[RFC7252\]](#), are CoAP endpoints tasked by CoAP clients to perform requests on their behalf.

5.1. Non-proxying endpoints

The CoAP PM Option can be applied end-to-end between client and server and, since it is Elective, it can be ignored by an endpoint that does not understand it.




Figure 5: Scenario with non-proxying endpoints

The enabled measurements are:

- * end-to-end loss and delay measurements between Client and Server,
- * on-path upstream and downstream loss and delay components (as explained in [\[I-D.ietf-ippm-explicit-flow-measurements\]](#)) if there is an observer (e.g. network functions or probes).
- * on-path intra-domain loss and delay portion as a result of the difference between the computed upstream or downstream components (as explained in [\[I-D.ietf-ippm-explicit-flow-measurements\]](#)).

The on-path network probes can read Q bit and S bit (or C bit) and implement the relevant algorithms to measure losses and RTT. Otherwise they can simply read the Event bits and be informed about the performance without implementing any algorithm. The event signaling bits can be sent from the Server (that can do the performance measurement calculation) to the Client, or viceversa.

if the CoAP PM Option is applied between client and the server, an Observer can measure the total RTT by using the S bit, indeed it allows RTT measurement for all the intermediate points. Additionally, with the Q bit and by applying [[I-D.ietf-ippm-rfc8321bis](#)], it is also possible to do hop-by-hop measurements for loss and delay and segment where possible, according to the methodologies described in [[I-D.ietf-ippm-explicit-flow-measurements](#)]. Alternatively, it is possible to use the C bit to get the same information for loss and delay as explained in [[I-D.ietf-ippm-explicit-flow-measurements](#)].

5.2. Collaborating or Non-collaborating proxies

The CoAP PM Option can be applied end-to-end between client and server (or between Proxies), and since it is Safe-to-Forward, it is intended to be safe for forwarding by a proxy that does not understand it.




Figure 6: Scenario with proxies

The proxies can be collaborating and it means that they understand and are configured to handle the CoAP PM Option. Otherwise, the proxies can be non-collaborating and this means that they do not handle the CoAP PM Option.

In case of collaborating proxies, the enabled measurements are different depending on where it is applied the CoAP PM Option:

- * It can be possible to apply the CoAP PM Option between Client and Server and the enabled measurements can be:
 - end-to-end loss and delay measurements between Client and Server,

- on-path upstream and downstream loss and delay components (as explained in [[I-D.ietf-ippm-explicit-flow-measurements](#)]) on each Proxy,
 - on-path upstream and downstream loss and delay components (as explained in [[I-D.ietf-ippm-explicit-flow-measurements](#)]) on each Observer,
 - on-path intra-domain loss and delay portion as a result of the difference between the computed upstream or downstream components (as explained in [[I-D.ietf-ippm-explicit-flow-measurements](#)]).
- * It can also be possible to apply the CoAP PM Option between the collaborating Proxies (instead of Client and Server) and the enabled measurements can be:
- end-to-end loss and delay measurements between Proxies,
 - on-path upstream and downstream loss and delay components (as explained in [[I-D.ietf-ippm-explicit-flow-measurements](#)]) on each Observer,
 - on-path intra-domain loss and delay portion as a result of the difference between the computed upstream or downstream components (as explained in [[I-D.ietf-ippm-explicit-flow-measurements](#)]).

In case of non-collaborating proxies, since Safe-to-Forward options that are not recognized MUST be forwarded, the enabled measurements can be:

- * end-to-end loss and delay measurements between Client and Server,
- * on-path upstream and downstream loss and delay components (as explained in [[I-D.ietf-ippm-explicit-flow-measurements](#)]) on each Observer,
- * on-path intra-domain loss and delay portion as a result of the difference between the computed upstream or downstream components (as explained in [[I-D.ietf-ippm-explicit-flow-measurements](#)]).

If there are CoAP proxies, the measurement can be done between the Proxies or between a Proxy and the Client or between a Proxy and the Server. It can be done through Spin bit or by applying [[I-D.ietf-ippm-rfc8321bis](#)] on the sQuare Bit signal. Therefore, it is also possible to do hop-by-hop measurements for loss and delay and segment where possible according to the methodologies described in [[I-D.ietf-ippm-explicit-flow-measurements](#)].

Since CoAP proxies hide the identity of the client and could also apply caching, on the server side the data would appear mixed in presence of more than one client doing the measurements. Similarly, clients would receive mixed signals in presence of cache entries. But, as previously mentioned, the measurements can be segmented and done between the Proxies or between a Proxy and the Client or between a Proxy and the Server. The Server can distinguish the source client by using additional flow information such as the IP addresses. It could also be possible to bundle different clients if they are mixed. So, it is worth highlighting that an on-path observer can find useful information both on the client-proxy link and on proxy-server link:

On the link from a proxy to the server, traffic from different clients would be mixed. In this case, the proxy can still use the PM Option to set S bit and Q bit (or C bit) for the bundle of clients for a specific server. The measurement can be done but it is an information related to a bundle of clients. An alternative can be to use the Option only for a single client at once in order to avoid to do a grouped measurement.

Conversely, on the link from the client to the proxy, communication may happen with different servers, and in this case it is necessary to check the other fields to understand the server.

5.3. OSCORE

CoAP PM Option can be used with OSCORE [[RFC8613](#)]. Since an OSCORE message may contain both an Inner and an Outer instance of a certain CoAP message field, the CoAP PM Option can be an Inner option or an Outer option based on the specific applications and required security and privacy. Then the operators can put their measurement probes in one or more places to break down the different RTT and loss contributions where it is relevant (e.g. at the ingress/egress of their respective network segments).

Inner option message fields (Class E) are used to communicate directly with the other endpoint and are encrypted and integrity protected. If the CoAP PM Option is sent as Inner Option, it only enables end-to-end measurements.

Outer option message fields (Class U or I) are used to support proxy operations and are unprotected or integrity protected only. If the CoAP PM Option is sent as Outer Option, it allows both end-to-end and on-path measurements by also enabling hop-by-hop and segmented measurements for loss and delay.

6. Management and Configuration

The measurement points can perform RTT and packet loss calculation without the need of any Network Management System (NMS) to collect information. It may be possible that the measurement points inform the NMS if there are particular network conditions (e.g. high packet loss or high RTT). For some parameters (e.g. 64 packets sQuare Bit signal) It is assumed static configuration on the client. There are several alternatives for the implementation but this is out of scope of this document.

7. Congestion Control

As specified in [[RFC7252](#)], clients (including proxies) MUST strictly limit the number of simultaneous outstanding interactions that they maintain to a given server (including proxies) to NSTART. The default value of NSTART is 1 but a value for NSTART greater than one is also possible. The CoAP PM Option implementation must not affect CoAP congestion control mechanisms.

8. Security Considerations

Security considerations related to CoAP proxying are discussed in [[RFC7252](#)].

A CoAP endpoint can use the CoAP PM Options to affect the measures of a network into which it is making requests by maliciously modifying the value of the option. Also, the PM bits may reveal performance information outside the administrative domain. To prevent that, a CoAP proxy that is located at the boundary of an administrative domain MAY be instructed to strip the payload or part of it before forwarding the message.

It is worth highlighting what happens if devices, transport network and server are operated by different administrative domains. Security concerns need to be taken into account.

CoAP can be secured using Datagram TLS (DTLS) [[RFC6347](#)] over UDP and it can prevent on-path measures in some cases. When a client uses a proxy the session towards the proxy can be secured using DTLS and the same from there on. In this case the separated sessions can still be measured.

CoAP can also be used with OSCORE [[RFC8613](#)] and the CoAP PM options can be integrity protected end-to-end by OSCORE. In this case, as explained above and differently from DTLS, the CoAP PM can easily work with OSCORE. OSCORE ensures end-to-end integrity protection and would tell the endpoints if someone tampered, but it doesn't mean that the endpoints are not lying to the observer. However it is possible to assume that for the typical COAP applications it is less likely that the endpoints are attackers while it is more likely that an on-path observer is the attacker.

9. IANA Considerations

IANA is requested to add the following entry to the "CoAP Option Numbers" sub-registry available at <https://www.iana.org/assignments/core-parameters/core-parameters.xhtml#option-numbers>:

Number	Name	Reference
TBD	PM Option	[This draft]

Figure 7: CoAP PM Option Numbers

10. Acknowledgements

The authors would like to thank Christian Amsuess, Carsten Bormann, Marco Tiloca, Thomas Fossati for the precious comments and suggestions.

11. References

11.1. 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>>.
- [RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", [RFC 7252](#), DOI 10.17487/RFC7252, June 2014, <<https://www.rfc-editor.org/info/rfc7252>>.
- [RFC8613] Selander, G., Mattsson, J., Palombini, F., and L. Seitz, "Object Security for Constrained RESTful Environments (OSCORE)", [RFC 8613](#), DOI 10.17487/RFC8613, July 2019, <<https://www.rfc-editor.org/info/rfc8613>>.

11.2. Informative References

- [I-D.ietf-ippm-explicit-flow-measurements]
Cociglio, M., Ferrieux, A., Fioccola, G., Lubashev, I.,
Bulgarella, F., Nilo, M., Hamchaoui, I., and R. Sisto,
"Explicit Flow Measurements Techniques", Work in Progress,
Internet-Draft, [draft-ietf-ippm-explicit-flow-measurements-02](https://www.ietf.org/archive/id/draft-ietf-ippm-explicit-flow-measurements-02), 13 October 2022,
<<https://www.ietf.org/archive/id/draft-ietf-ippm-explicit-flow-measurements-02.txt>>.
- [I-D.ietf-ippm-rfc8321bis]
Fioccola, G., Cociglio, M., Mirsky, G., Mizrahi, T., and
T. Zhou, "Alternate-Marking Method", Work in Progress,
Internet-Draft, [draft-ietf-ippm-rfc8321bis-03](https://www.ietf.org/archive/id/draft-ietf-ippm-rfc8321bis-03), 25 July
2022, <<https://www.ietf.org/archive/id/draft-ietf-ippm-rfc8321bis-03.txt>>.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer
Security Version 1.2", [RFC 6347](https://www.rfc-editor.org/info/rfc6347), DOI 10.17487/RFC6347,
January 2012, <<https://www.rfc-editor.org/info/rfc6347>>.
- [RFC9000] Iyengar, J., Ed. and M. Thomson, Ed., "QUIC: A UDP-Based
Multiplexed and Secure Transport", [RFC 9000](https://www.rfc-editor.org/info/rfc9000),
DOI 10.17487/RFC9000, May 2021,
<<https://www.rfc-editor.org/info/rfc9000>>.
- [RFC9312] Kühlewind, M. and B. Trammell, "Manageability of the QUIC
Transport Protocol", [RFC 9312](https://www.rfc-editor.org/info/rfc9312), DOI 10.17487/RFC9312,
September 2022, <<https://www.rfc-editor.org/info/rfc9312>>.

Authors' Addresses

Giuseppe Fioccola
Huawei
Riesstrasse, 25
80992 Munich
Germany
Email: giuseppe.fioccola@huawei.com

Tianran Zhou
Huawei
156 Beiqing Rd.
Beijing
100095
China
Email: zhoutianran@huawei.com

Mauro Cociglio
Telecom Italia
Via Reiss Romoli, 274
10148 Torino
Italy
Email: mauro.cociglio@outlook.com

Fabio Bulgarella
Telecom Italia
Via Reiss Romoli, 274
10148 Torino
Italy
Email: fabio.bulgarella@guest.telecomitalia.it

Massimo Nilo
Telecom Italia
Via Reiss Romoli, 274
10148 Torino
Italy
Email: massimo.nilo@telecomitalia.it