CORE Working Group
Internet-Draft
Intended status: Standards Track
Expires: 23 April 2026

G. Fioccola
T. Zhou
Huawei
M. Cociglio

F. Bulgarella
Telecom Italia
Y. Zhu
China Telecom

20 October 2025

Constrained Application Protocol (CoAP) Performance Measurement Option
draft-ietf-core-coap-pm-05

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 hop-by-hop. The endpoints cooperate by marking and, possibly, mirroring information on the round-trip connection.

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 23 April 2026.

Copyright Notice

Copyright (c) 2025 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

In the CoAP protocol [RFC7252], reliability is provided by marking a message as Confirmable (CON), as to be retransmitted if not acknowledged by an ACK message. A message that does not require reliable transmission can be sent as a Non-confirmable message (NON).

In case of CoAP reliable mode, Message IDs and ACKs could potentially 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 a CoAP environment including resource-constrained nodes. And it is in any case limited to RTT and end-to-end losses.

A mechanism to measure both end-to-end and hop-by-hop performance in CoAP can be useful to verify that the operational requirements are met, 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.

[RFC9506] 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 probe. These methodologies could potentially be used and extended in CoAP.

[RFC9506] defines different combinations of bits. Such flexibility is convenient when using a protocol with a limited number of eligible bits to exploit, e.g., QUIC. Different alternatives have been proposed, 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) allow performing both end-to-end and hop-by-hop measurements and can be useful for an operator or an enterprise that is managing a constrained, low-power and lossy network.

This document ultimately is intended to be published as a standards-track RFC. Its current stage of development, it is complete and stable enough to be used as a basis for experiments, the evaluation and continuation of which will lead to further evolution towards the intended standards-track document.

1.1. Requirements Language

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. Performance Measurement methods for CoAP

The approach proposed in this document relies on a new Performance Measurement (PM) Option for CoAP [RFC7252]. This new option is defined in Section 3 and carries PM bits.

The PM bits that are included in the Option are:

* sQuare bit (Q bit), based on [RFC9341] and further described in [RFC9506];

* 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 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 sQuare wave, and it is not necessary to store timestamps, since the RTT is the the time interval between a request and its response, if any.

2. Enabling on-path probe (proxy, gateway) metrics.

2.1. sQuare bit

The sQuare bit algorithm consists of creating square waves of a known length (e.g. 64 packets). Each communicating endpoint can set the Q bit and toggle its value each time a fixed number of messages have been sent. The number of packets can be easily recognized and packet loss can be measured.

An on-path probe can read the Q bit signal and perform the measurements. All the possible measurements (end-to-end, hop-by-hop) that are enabled by the Q bit are detailed in [RFC9341] and [RFC9506].

3. CoAP Performance Measurement Option

Table 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].

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

Table 1: CoAP PM Option Properties (C=Critical, U=Unsafe, N=NoCacheKey, R=Repeatable)

The CoAP PM Option is Elective and Proxy Unsafe. But as discussed in Section 4.3, it MAY also be Safe-to-Forward in some implementations with non-caching proxies.

As detailed in Section 4.5, the option can be of class U, I and E in terms of OSCORE processing.

3.1. Structure of the PM Option

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

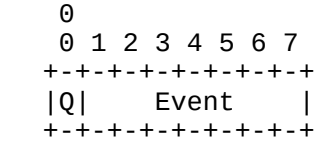

Figure 1: Value of the CoAP Performance Measurement Option

Where:

* Q bit is described in [RFC9506];

* Event bits MAY be used to encode additional Loss and Delay information based on well-defined encoding; they can also be used by on-path probes. If these Event bits are all zero, they MUST be ignored on receipt.

The Event bits can be divided into two parts, for instance: loss event bits and delay event bits. Based on the average RTT, an endpoint 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 probe becomes aware of the network conditions by simply reading these Event bits and without applying any algorithm.

The on-path probe 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 Option 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 exclude the option in the outgoing message. In this way the other CoAP endpoints become aware that the measurement cannot be executed in that case.

It is RECOMMENDED to insert the option immediately before the transmission in order to avoid unexpected behavior in case of retransmissions. Further details about the the accuracy of the measurements can be found in [RFC9506].

The fixed number of packets to create the Q bit signal is predefined: its value is configured from the beginning for all the CoAP endpoints, as also mentioned in Section 5.

It is worth mentioning that in some specific circumstances, e.g. CoAP clients that "observe" resources [RFC7641] or empty-ACKs, the measurements can be done only for one direction. For bidirectional measurements, it is required to have traffic in both directions.

4. Application Scenarios

The main usage of the CoAP PM Option is to do end-to-end measurement between the client and the server, but it can also allow on-path 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 nodes could be:

* Probes that must be able to see deep into application.

* Proxies that, as specified in [RFC7252], are CoAP endpoints tasked by CoAP clients to perform requests on their behalf.

4.1. 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.

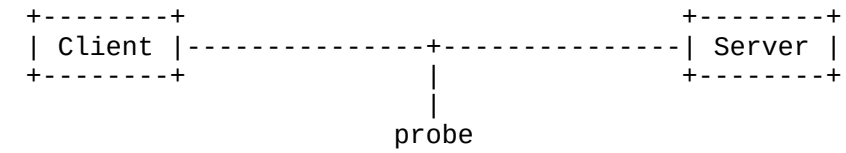

Figure 2: Scenario with non-proxying endpoints

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 [RFC9506]) if there is a probe (e.g. network functions).

* on-path intra-domain loss and delay portion if there are more than one probe as a result of the difference between the computed upstream or downstream components (as explained in [RFC9506]).

The on-path network probes can read Q 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 vice versa.

If the CoAP PM Option is applied between client and server, with the Q bit and by applying [RFC9341], a probe can do hop-by-hop measurements for loss and delay and segment where possible between the probes, according to the methodologies described in [RFC9506].

4.2. Collaborating proxies

The proxies can be "collaborating": this means that they understand and are configured to handle the CoAP PM Option. The CoAP PM Option can be handled on the client, on the server and on each Proxies.

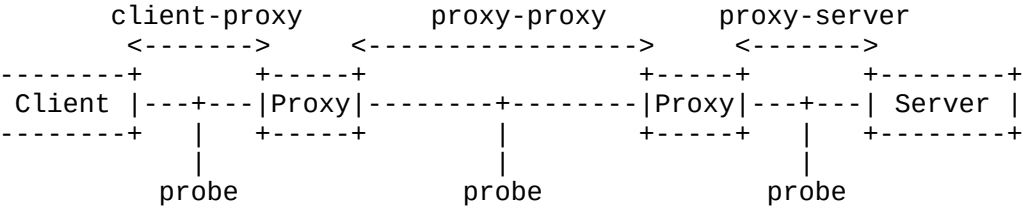

Figure 3: Scenario with collaborating proxies

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

It can be possible to apply the CoAP PM Option between Client and Proxy, between the Proxies and between Proxy and Server. The sessions are separated end-to-end between Client and Server and the enabled measurements can be done on the separated sessions:

* loss and delay measurements between Client and Proxy, between Proxies and between Proxy and Server,

* on-path upstream and downstream loss and delay components (as explained in [RFC9506]) on each Proxy,

* on-path upstream and downstream loss and delay components (as explained in [RFC9506]) on each Probe,

* end-to-end loss and delay measurements as a result of the addition of the loss and delay contributions of the separated sessions.

* on-path intra-domain loss and delay portion as a result of the difference between the computed upstream or downstream components (as explained in [RFC9506]).

So, 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 by applying [RFC9341] 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 [RFC9506].

4.3. Non-collaborating proxies

The proxies can be non-collaborating and this means that they do not handle the CoAP PM Option. The CoAP PM Option can be applied end-to-end between client and server.

There are some issues that may occur in case of non-collaborating proxies. In general, since CoAP proxies hide the identity of the client, the data would appear mixed after the proxy in the presence of more than one client doing the measurements. Similarly, since CoAP proxies could also apply caching, it can happen to receive mixed signals in the presence of cache entries.

In case of collaborating proxies, these issues are solved because the measurements can be segmented and done between the Client and a Proxy or between the Proxies or between a Proxy and the Server. In this case, it could be possible for the proxy to still use the PM Option for the bundle of clients for a specific server.

While, in case of non-collaborating proxies, it is RECOMMENDED to use the Option only for a single client and a single server at once in order to avoid that traffic from different clients would be mixed. But, if the proxy has also cached data, the data can be reordered and mixed, so that they cannot be used for measurement. For this reason, the PM Option is defined as Proxy Unsafe: it is intended to be unsafe for forwarding by a proxy that does not understand it. In conclusion, if there are non-collaborating and caching proxies, the measurements would not be possible.

4.3.1. Non-caching proxies

An implementation MAY consider the PM Option as Safe-to-Forward if the proxies are non-caching in general or in the only case the PM Option is included in the message.

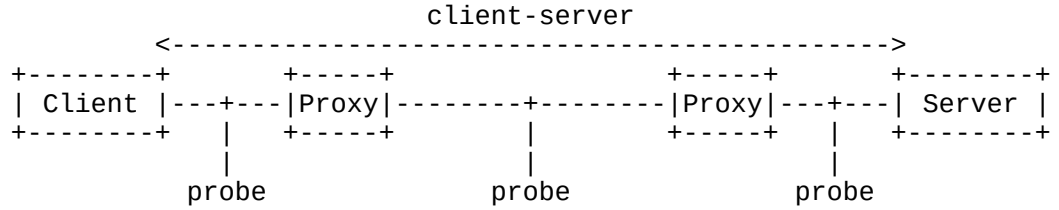

Figure 4: Scenario with non-collaborating and non-caching proxies

In case of non-collaborating and non-caching proxies, proxies MAY be configured to handle the PM Option as Safe-to-Forward, and it means that not recognized option MUST be forwarded. Therefore, the enabled measurements for a single client and a single server at once 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 [RFC9506]) on each Probe,

* on-path intra-domain loss and delay portion as a result of the difference between the computed upstream or downstream components (as explained in [RFC9506]).

4.4. DTLS

CoAP can be secured using Datagram TLS (DTLS) [RFC6347] over UDP and it can prevent on-path measures in case of non-proxying endpoints. When a client uses a collaborating proxy the sessions client-proxy, proxy-proxy, proxy-server are secured using DTLS, but the separated sessions can still be measured. An on-path probe cannot perform the measurements in any case.

4.5. OSCORE

The 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 network administrators 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 options (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 in all the cases. In case of collaborating proxies the separated sessions client-proxy, proxy-proxy, proxy-server cannot be measured.

Outer options (Class U or I) are intended to be 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 enabling hop-by-hop and segmented loss and delay measurements on the proxies.

If an OSCORE endpoint sends both outer and inner option, the inner is for measuring the connection to the end-to-end peer, and the outer can be used for measuring the connection to next proxy.

If the PM option is used as an Outer Option, it may also be integrity-protected, to be reliably processed and this would require using also DTLS or an OSCORE association with a proxy [I-D.ietf-core-oscore-capable-proxies].

5. 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; this is out of scope of this document.

6. Congestion Control

As specified in Section 4.7 of [RFC7252], clients (including proxies) have to 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.

7. Security Considerations

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

A CoAP endpoint can use the CoAP PM Option to affect the measures of a network into which it is making requests by maliciously specifying a wrong option value. 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 used with DTLS [RFC6347] and it can prevent on-path measures by on-path probes while it is still possible to do measurements on collaborating proxies, as explained above.

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 with the option value, but it doesn't mean that the endpoints are not lying to the probe. 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 probe is the attacker.

8. 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	[This document]

Table 2: CoAP PM Option Number

9. References

9.1. 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/rfc/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/rfc/rfc7252>.

[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/rfc/rfc8174>.

[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/rfc/rfc8613>.

[RFC9341] Fioccola, G., Ed., Cociglio, M., Mirsky, G., Mizrahi, T., and T. Zhou, "Alternate-Marking Method", RFC 9341, DOI 10.17487/RFC9341, December 2022, <https://www.rfc-editor.org/rfc/rfc9341>.

9.2. Informative References

[I-D.ietf-core-oscore-capable-proxies] Tiloca, M. and R. H旦glund, "OSCORE-capable Proxies", Work in Progress, Internet-Draft, draft-ietf-core-oscore-capable-proxies-05, 3 September 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-core-oscore-capable-proxies-05>.

[RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, DOI 10.17487/RFC6347, January 2012, <https://www.rfc-editor.org/rfc/rfc6347>.

[RFC7641] Hartke, K., "Observing Resources in the Constrained Application Protocol (CoAP)", RFC 7641, DOI 10.17487/RFC7641, September 2015, <https://www.rfc-editor.org/rfc/rfc7641>.

[RFC9000] 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/rfc/rfc9000>.

[RFC9312] K端hlewind, M. and B. Trammell, "Manageability of the QUIC Transport Protocol", RFC 9312, DOI 10.17487/RFC9312, September 2022, <https://www.rfc-editor.org/rfc/rfc9312>.

[RFC9506] Cociglio, M., Ferrieux, A., Fioccola, G., Lubashev, I., Bulgarella, F., Nilo, M., Hamchaoui, I., and R. Sisto, "Explicit Host-to-Network Flow Measurements Techniques", RFC 9506, DOI 10.17487/RFC9506, October 2023, <https://www.rfc-editor.org/rfc/rfc9506>.

Acknowledgements

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

Contributors

Massimo Nilo
Telecom Italia

Email: massimo.nilo@telecomitalia.it

Fabrizio Milan
Telecom Italia
Email: fabrizio.milan@telecomitalia.it

Authors' Addresses

Giuseppe Fioccola
Huawei
Viale Martesana, 12
20055 Vimodrone (Milan)
Italy
Email: giuseppe.fioccola@huawei.com

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

Mauro Cociglio
Italy
Email: mauro.cociglio@outlook.com

Fabio Bulgarella
Telecom Italia
Italy
Email: fabio.bulgarella@guest.telecomitalia.it

Yongqing Zhu
China Telecom
China
Email: zhuyq8@chinatelecom.cn