

CoRE Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: November 4, 2017

G. Selander  
J. Mattsson  
F. Palombini  
Ericsson AB  
L. Seitz  
SICS Swedish ICT  
May 03, 2017

**Object Security of CoAP (OSCOAP)**  
**draft-ietf-core-object-security-03**

**Abstract**

This document defines Object Security of CoAP (OSCOAP), a method for application layer protection of the Constrained Application Protocol (CoAP), using the CBOR Object Signing and Encryption (COSE). OSCOAP provides end-to-end encryption, integrity and replay protection to CoAP payload, options, and header fields, as well as a secure message binding. OSCOAP is designed for constrained nodes and networks and can be used across intermediaries and over any layer. The use of OSCOAP is signaled with the CoAP option Object-Security, also defined in this document.

**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 <http://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 November 4, 2017.

**Copyright Notice**

Copyright (c) 2017 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

(<http://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

<a href="#">1.</a>	<a href="#">Introduction</a>	<a href="#">3</a>
<a href="#">1.1.</a>	<a href="#">Terminology</a>	<a href="#">4</a>
<a href="#">2.</a>	<a href="#">The Object-Security Option</a>	<a href="#">5</a>
<a href="#">3.</a>	<a href="#">The Security Context</a>	<a href="#">6</a>
<a href="#">3.1.</a>	<a href="#">Security Context Definition</a>	<a href="#">6</a>
<a href="#">3.2.</a>	<a href="#">Derivation of Security Context Parameters</a>	<a href="#">9</a>
<a href="#">3.3.</a>	<a href="#">Requirements on the Security Context Parameters</a>	<a href="#">10</a>
<a href="#">4.</a>	<a href="#">Protected CoAP Message Fields</a>	<a href="#">11</a>
<a href="#">4.1.</a>	<a href="#">CoAP Payload</a>	<a href="#">12</a>
<a href="#">4.2.</a>	<a href="#">CoAP Header</a>	<a href="#">12</a>
<a href="#">4.3.</a>	<a href="#">CoAP Options</a>	<a href="#">12</a>
<a href="#">5.</a>	<a href="#">The COSE Object</a>	<a href="#">18</a>
<a href="#">5.1.</a>	<a href="#">Plaintext</a>	<a href="#">19</a>
<a href="#">5.2.</a>	<a href="#">Additional Authenticated Data</a>	<a href="#">19</a>
<a href="#">6.</a>	<a href="#">Sequence Numbers, Replay, Message Binding, and Freshness</a>	<a href="#">20</a>
<a href="#">6.1.</a>	<a href="#">AEAD Nonce Uniqueness</a>	<a href="#">20</a>
<a href="#">6.2.</a>	<a href="#">Replay Protection</a>	<a href="#">20</a>
<a href="#">6.3.</a>	<a href="#">Sequence Number and Replay Window State</a>	<a href="#">21</a>
<a href="#">6.4.</a>	<a href="#">Freshness</a>	<a href="#">22</a>
<a href="#">6.5.</a>	<a href="#">Delay and Mismatch Attacks</a>	<a href="#">23</a>
<a href="#">7.</a>	<a href="#">Processing</a>	<a href="#">23</a>
<a href="#">7.1.</a>	<a href="#">Protecting the Request</a>	<a href="#">23</a>
<a href="#">7.2.</a>	<a href="#">Verifying the Request</a>	<a href="#">23</a>
<a href="#">7.3.</a>	<a href="#">Protecting the Response</a>	<a href="#">25</a>
<a href="#">7.4.</a>	<a href="#">Verifying the Response</a>	<a href="#">25</a>
<a href="#">8.</a>	<a href="#">OSCOAP Compression</a>	<a href="#">26</a>
<a href="#">8.1.</a>	<a href="#">Encoding of the Object-Security Option</a>	<a href="#">27</a>
<a href="#">8.2.</a>	<a href="#">Examples</a>	<a href="#">28</a>
<a href="#">9.</a>	<a href="#">Web Linking</a>	<a href="#">29</a>
<a href="#">10.</a>	<a href="#">Security Considerations</a>	<a href="#">29</a>
<a href="#">11.</a>	<a href="#">Privacy Considerations</a>	<a href="#">31</a>
<a href="#">12.</a>	<a href="#">IANA Considerations</a>	<a href="#">32</a>
<a href="#">12.1.</a>	<a href="#">CoAP Option Numbers Registry</a>	<a href="#">32</a>
<a href="#">12.2.</a>	<a href="#">Media Type Registrations</a>	<a href="#">32</a>
<a href="#">12.3.</a>	<a href="#">CoAP Content Format Registration</a>	<a href="#">33</a>
<a href="#">13.</a>	<a href="#">Acknowledgments</a>	<a href="#">34</a>
<a href="#">14.</a>	<a href="#">References</a>	<a href="#">34</a>
<a href="#">14.1.</a>	<a href="#">Normative References</a>	<a href="#">34</a>



<a href="#">14.2.</a>	Informative References . . . . .	<a href="#">35</a>
<a href="#">Appendix A.</a>	Test Vectors . . . . .	<a href="#">36</a>
<a href="#">Appendix B.</a>	Examples . . . . .	<a href="#">36</a>
<a href="#">B.1.</a>	Secure Access to Sensor . . . . .	<a href="#">36</a>
<a href="#">B.2.</a>	Secure Subscribe to Sensor . . . . .	<a href="#">37</a>
<a href="#">Appendix C.</a>	Object Security of Content (OSCON) . . . . .	<a href="#">39</a>
<a href="#">C.1.</a>	Overhead OSCON . . . . .	<a href="#">40</a>
<a href="#">C.2.</a>	MAC Only . . . . .	<a href="#">41</a>
<a href="#">C.3.</a>	Signature Only . . . . .	<a href="#">41</a>
<a href="#">C.4.</a>	Authenticated Encryption with Additional Data (AEAD) . . . . .	<a href="#">42</a>
<a href="#">C.5.</a>	Symmetric Encryption with Asymmetric Signature (SEAS) . . . . .	<a href="#">43</a>
	Authors' Addresses . . . . .	<a href="#">43</a>

## 1. Introduction

The Constrained Application Protocol (CoAP) is a web application protocol, designed for constrained nodes and networks [[RFC7228](#)]. CoAP specifies the use of proxies for scalability and efficiency. At the same time CoAP [[RFC7252](#)] references DTLS [[RFC6347](#)] for security. Proxy operations on CoAP messages require DTLS to be terminated at the proxy. The proxy therefore not only has access to the data required for performing the intended proxy functionality, but is also able to eavesdrop on, or manipulate any part of the CoAP payload and metadata, in transit between client and server. The proxy can also inject, delete, or reorder packages since they are no longer protected by DTLS.

This document defines Object Security of CoAP (OSCOAP), a data object based security protocol, protecting CoAP message exchanges end-to-end, across intermediary nodes. An analysis of end-to-end security for CoAP messages through intermediary nodes is performed in [[I-D.hartke-core-e2e-security-reqs](#)], this specification addresses the forwarding case. In addition to the core features defined in [[RFC7252](#)], OSCOAP supports Observe [[RFC7641](#)] and Blockwise [[RFC7959](#)].

OSCOAP is designed for constrained nodes and networks and provides an in-layer security protocol for CoAP which does not depend on underlying layers. OSCOAP can be used anywhere that CoAP can be used, including unreliable transport [[RFC7228](#)], reliable transport [[I-D.ietf-core-coap-tcp-tls](#)], and non-IP transport [[I-D.bormann-6lo-coap-802-15-4e](#)]. OSCOAP may also be used to protect group communication for CoAP [[I-D.tiloca-core-multicast-oscoap](#)]. The use of OSCOAP does not affect the URI scheme and OSCOAP can therefore be used with any URI scheme defined for CoAP. The application decides the conditions for which OSCOAP is required.

OSCOAP builds on CBOR Object Signing and Encryption (COSE) [[I-D.ietf-cose-msg](#)], providing end-to-end encryption, integrity,



replay protection, and secure message binding. A compressed version of COSE is used, see [Section 8](#). The use of OSCOAP is signaled with the CoAP option Object-Security, defined in [Section 2](#). OSCOAP provides protection of CoAP payload, certain options, and header fields. The solution transforms an unprotected CoAP message into a protected CoAP message in the following way: the unprotected CoAP message is protected by including payload (if present), certain options, and header fields in a COSE object. The message fields that have been encrypted are removed from the message whereas the Object-Security option and the compressed COSE object are added, see Figure 1.



Figure 1: Sketch of OSCOAP

OSCOAP may be used in extremely constrained settings, where CoAP over DTLS may be prohibitive e.g. due to large code size. Alternatively, OSCOAP can be combined with DTLS, thereby enabling end-to-end security of e.g. CoAP payload and options, in combination with hop-by-hop protection of the entire CoAP message, during transport between end-point and intermediary node. Examples of the use of OSCOAP are given in [Appendix B](#).

The message protection provided by OSCOAP can alternatively be applied only to the payload of individual messages. We call this object security of content (OSCON), which is defined in [Appendix C](#).

### 1.1. Terminology

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 [\[RFC2119\]](#). These words may also appear in this document in lowercase, absent their normative meanings.



Readers are expected to be familiar with the terms and concepts described in CoAP [[RFC7252](#)], Observe [[RFC7641](#)], Blockwise [[RFC7959](#)], COSE [[I-D.ietf-cose-msg](#)], CBOR [[RFC7049](#)], CDDL [[I-D.greevenbosch-appsawg-cbor-cddl](#)], and constrained environments [[RFC7228](#)].

The terms Common/Sender/Recipient Context, Master Secret/Salt, Sender ID/Key/IV, Receptient ID/Key/IV and Context IV are defined in [Section 3.1](#).

## 2. The Object-Security Option

The Object-Security option (see Figure 2) indicates that OSCOAP is used to protect the CoAP message exchange. The Object-Security option is critical, safe to forward, part of the cache key, not repeatable, and opaque.

No.	C	U	N	R	Name	Format	Length
TBD	x				Object-Security	opaque	0-

C=Critical, U=Unsafe, N=NoCacheKey, R=Repeatable

Figure 2: The Object-Security Option

A successful response to a request with the Object-Security option SHALL contain the Object-Security option. A CoAP endpoint SHOULD NOT cache a response to a request with an Object-Security option, since the response is only applicable to the original client's request. The Object-Security option is included in the cache key for backward compatibility with proxies not recognizing the Object-Security option. The effect is that messages with the Object-Security option will never generate cache hits. For Max-Age processing, see [Section 4.3.1.1](#).

The protection is achieved by means of a COSE object (see [Section 5](#)), which is compressed and then included in the protected CoAP message. The placement of the COSE object depends on whether the method/response code allows payload (see [[RFC7252](#)]):

- o If the method/response code allows payload, then the compressed COSE object [Section 8](#) is the payload of the protected message, and the Object-Security option has length zero. An endpoint receiving a CoAP message with payload, that also contains a non-empty Object-Security option SHALL treat it as malformed and reject it.





- o If the method/response code does not allow payload, then the compressed COSE object [Section 8](#) is the value of the Object-Security option and the length of the Object-Security option is equal to the size of the compressed COSE object. An endpoint receiving a CoAP message without payload, that also contains an empty Object-Security option SHALL treat it as malformed and reject it.

The size of the COSE object depends on whether the method/response code allows payload, if the message is a request or response, on the set of options that are included in the unprotected message, the AEAD algorithm, the length of the information identifying the security context, and the length of the sequence number.

### **[3.](#) The Security Context**

OSCOAP uses COSE with an Authenticated Encryption with Additional Data (AEAD) algorithm between a CoAP client and a CoAP server. An implementation supporting this specification MAY only implement the client part or MAY only implement the server part.

This specification requires that client and server establish a security context to apply to the COSE objects protecting the CoAP messages. In this section we define the security context, and also specify how to derive the initial security contexts in client and server based on common shared secret and a key derivation function (KDF).

#### **[3.1.](#) Security Context Definition**

The security context is the set of information elements necessary to carry out the cryptographic operations in OSCOAP. For each endpoint, the security context is composed of a "Common Context", a "Sender Context", and a "Recipient Context".

The endpoints protect messages to send using the Sender Context and verify messages received using the Recipient Context, both contexts being derived from the Common Context and other data. Clients need to be able to retrieve the correct security context to use.

An endpoint uses its Sender ID (SID) to derive its Sender Context, and the other endpoint uses the same ID, now called Recipient ID (RID), to derive its Recipient Context. In communication between two endpoints, the Sender Context of one endpoint matches the Recipient Context of the other endpoint, and vice versa. Thus the two security contexts identified by the same IDs in the two endpoints are not the same, but they are partly mirrored. Retrieval and use of the security context are shown in Figure 3.



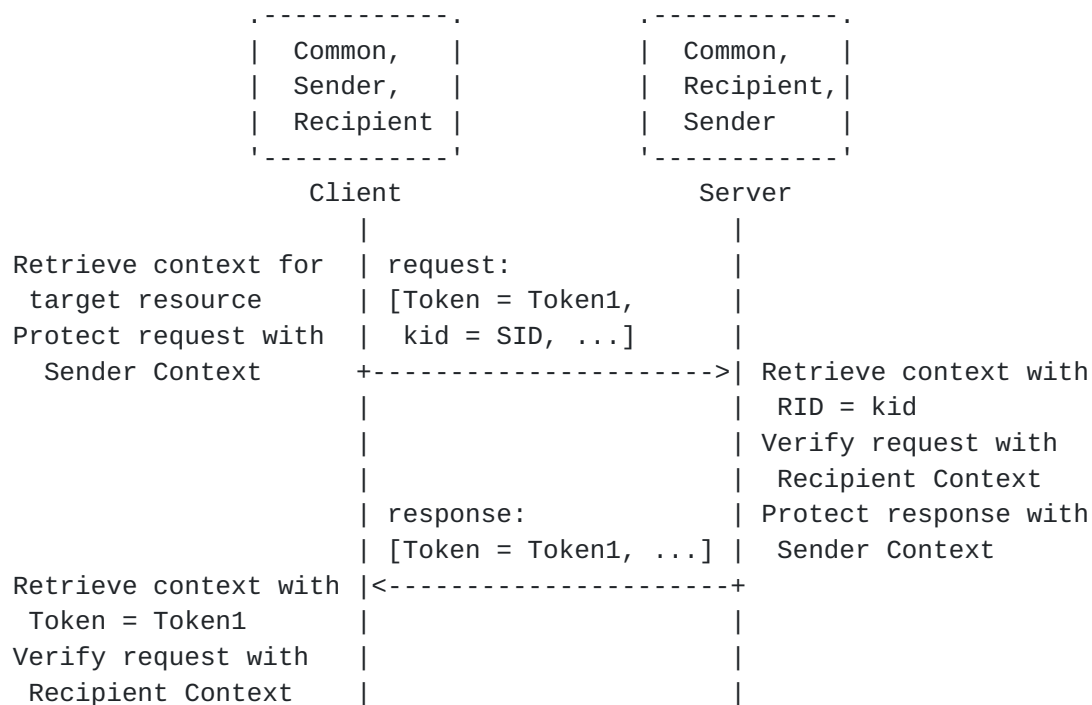


Figure 3: Retrieval and use of the Security Context

The Common Context contains the following parameters:

- o Algorithm (Alg). Value that identifies the COSE AEAD algorithm to use for encryption. Its value is immutable once the security context is established.
- o Master Secret. Variable length, uniformly random byte string containing the key used to derive traffic keys and IVs. Its value is immutable once the security context is established.
- o Master Salt (OPTIONAL). Variable length byte string containing the salt used to derive traffic keys and IVs. Its value is immutable once the security context is established.

The Sender Context contains the following parameters:

- o Sender ID. Variable length byte string identifying the Sender Context. Its value is immutable once the security context is established.
- o Sender Key. Byte string containing the symmetric key to protect messages to send. Derived from Common Context and Sender ID. Length is determined by Algorithm. Its value is immutable once the security context is established.



- o Sender IV. Byte string containing the IV to protect messages to send. Derived from Common Context and Sender ID. Length is determined by Algorithm. Its value is immutable once the security context is established.
- o Sequence Number. Non-negative integer used to protect requests and observe responses to send. Used as partial IV [[I-D.ietf-cose-msg](#)] to generate unique nonces for the AEAD. Maximum value is determined by Algorithm.

The Recipient Context contains the following parameters:

- o Recipient ID. Variable length byte string identifying the Recipient Context. Its value is immutable once the security context is established.
- o Recipient Key. Byte string containing the symmetric key to verify messages received. Derived from Common Context and Recipient ID. Length is determined by the Algorithm. Its value is immutable once the security context is established.
- o Recipient IV. Byte string containing the IV to verify messages received. Derived from Common Context and Recipient ID. Length is determined by Algorithm. Its value is immutable once the security context is established.
- o Replay Window. The replay window to verify requests and observe responses received.

When it is understood which context is referred to (Sender Context or Recipient Context), the term "Context IV" is used to denote the IV currently used with this context.

An endpoint may free up memory by not storing the Sender Key, Sender IV, Recipient Key, and Recipient IV, deriving them from the Common Context when needed. Alternatively, an endpoint may free up memory by not storing the Master Secret and Master Salt after the other parameters have been derived.

The endpoints MAY interchange the client and server roles while maintaining the same security context. When this happens, the former server still protects messages to send using its Sender Context, and verifies messages received using its Recipient Context. The same is also true for the former client. The endpoints MUST NOT change the Sender/Recipient ID. In other words, changing the roles does not change the set of keys to be used.



### **3.2. Derivation of Security Context Parameters**

The parameters in the security context are derived from a small set of input parameters. The following input parameters SHALL be pre-established:

- o Master Secret
- o Sender ID
- o Recipient ID

The following input parameters MAY be pre-established. In case any of these parameters is not pre-established, the default value indicated below is used:

- o AEAD Algorithm (Alg)
  - \* Default is AES-CCM-64-64-128 (COSE abbreviation: 12)
- o Master Salt
  - \* Default is the empty string
- o Key Derivation Function (KDF)
  - \* Default is HKDF SHA-256
- o Replay Window Type and Size
  - \* Default is DTLS-type replay protection with a window size of 32

How the input parameters are pre-established, is application specific. The EDHOC protocol [[I-D.selander-ace-cose-ecdhe](#)] enables the establishment of input parameters with the property of forward secrecy and negotiation of KDF and AEAD, it thus provides all necessary pre-requisite steps for using OSCOAP as defined here.

#### **3.2.1. Derivation of Sender Key/IV, Recipient Key/IV**

The KDF MUST be one of the HMAC based HKDF [[RFC5869](#)] algorithms defined in COSE. HKDF SHA-256 is mandatory to implement. The security context parameters Sender Key/IV and Recipient Key/IV SHALL be derived from the input parameters using the HKDF, which consists of the composition of the HKDF-Extract and HKDF-Expand steps ([RFC5869](#)):

output parameter = HKDF(salt, IKM, info, L)





where:

- o salt is the Master Salt as defined above
- o IKM is the Master Secret as defined above
- o info is a CBOR array consisting of:

```
info = [  
    id : bstr,  
    alg : int,  
    type : tstr,  
    L : int  
]
```

\* id is the Sender ID or Recipient ID

\* type is "Key" or "IV"

- o L is the size of the key/IV for the AEAD algorithm used, in octets.

For example, if the algorithm AES-CCM-64-64-128 (see Section 10.2 in [\[I-D.ietf-cose-msg\]](#)) is used, the value for L is 16 for keys and 7 for IVs.

### **3.2.2. Initial Sequence Numbers and Replay Window**

The Sequence Number is initialized to 0. The supported types of replay protection and replay window length is application specific and depends on the lower layers. Default is DTLS-type replay protection with a window size of 32 initiated as described in [Section 4.1.2.6 of \[RFC6347\]](#).

### **3.3. Requirements on the Security Context Parameters**

As collisions may lead to the loss of both confidentiality and integrity, Sender ID SHALL be unique in the set of all security contexts using the same Master Secret. Normally (e.g. when using EDHOC [\[I-D.selander-ace-cose-ecdhe\]](#)) Sender IDs can be very short. Note that Sender IDs of different lengths can be used with the same Master Secret. E.g. the SID with value 0x00 is different from the SID with the value 0x0000. If Sender ID uniqueness cannot be guaranteed, random Sender IDs MUST be used. Random Sender IDs MUST be long enough so that the probability of collisions is negligible.



To enable retrieval of the right Recipient Context, the Recipient ID SHOULD be unique in the sets of all Recipient Contexts used by an endpoint.

The same Master Salt MAY be used with several Master Secrets.

#### **4. Protected CoAP Message Fields**

OSCOAP transforms an unprotected CoAP message into a protected CoAP message, and vice versa. This section defines how the CoAP message fields are protected. Note that OSCOAP protects messages from the CoAP Requests/Responses layer only, and not from the Messaging layer ([Section 2 of \[RFC7252\]](#)): this means that RST and ACK empty messages are not protected, while ACK with piggybacked responses are protected using the process defined in this document. All the messages mentioned in this document refer to CON, NON and non-empty ACK messages.

OSCOAP protects as much of the unprotected CoAP message as possible, while still allowing forward proxy operations [[I-D.hartke-core-e2e-security-reqs](#)]. Message fields may either be

- o Class E: encrypted and integrity protected,
- o Class I: integrity protected only, or
- o Class U: unprotected.

This section also outlines how the message fields are transferred, a detailed description of the processing is provided in [Section 7](#). Message fields of the unprotected CoAP message are either transferred in the header/options part of the protected CoAP message, or in the plaintext of the COSE object. Depending on which, the location of the message field in the protected CoAP message is called "inner" or "outer":

- o Inner message field: message field included in the plaintext of the COSE object of the protected CoAP message (see [Section 5.1](#)). The inner message fields are by definition encrypted and integrity protected by the COSE object (Class E).
- o Outer message field: message field included in the header or options part of the protected CoAP message. The outer message fields are not encrypted and thus visible to an intermediary, but may be integrity protected by including the message field values in the Additional Authenticated Data (AAD) of the COSE object (see [Section 5.2](#)). I.e. outer message fields may be Class I or Class U.



Note that, even though the message formats are slightly different, OSCOAP complies with CoAP over unreliable transport [[RFC7252](#)] as well as CoAP over reliable transport [[I-D.ietf-core-coap-tcp-tls](#)].

#### **4.1. CoAP Payload**

The CoAP Payload SHALL be encrypted and integrity protected (Class E), and thus is an inner message field.

The sending endpoint writes the payload of the unprotected CoAP message into the plaintext of the COSE object.

The receiving endpoint verifies and decrypts the COSE object, and recreates the payload of the unprotected CoAP message.

#### **4.2. CoAP Header**

Many CoAP header fields are required to be read and changed during a normal message exchange or when traversing a proxy and thus cannot in general be protected between the endpoints, e.g. CoAP message layer fields such as Message ID.

The CoAP header field Code MUST be sent in plaintext to support RESTful processing, but MUST be integrity protected to prevent an intermediary from changing, e.g. from GET to DELETE (Class I). The CoAP version number MUST be integrity protected to prevent potential future version-based attacks (Class I). Note that while the version number is not sent in each CoAP message over reliable transport [[I-D.ietf-core-coap-tcp-tls](#)], its value is known to client and server.

The other CoAP header fields SHALL neither be integrity protected nor encrypted (Class U). All CoAP header fields are thus outer message fields.

The sending endpoint SHALL copy the header fields from the unprotected CoAP message to the header of the protected CoAP message. The receiving endpoint SHALL copy the header fields from the protected CoAP message to the header of the unprotected CoAP message. Both sender and receiver include the CoAP version number and header field Code in the AAD of the COSE object (see [Section 5.2](#)).

#### **4.3. CoAP Options**

Most options are encrypted and integrity protected (Class E), and thus inner message fields. But to allow certain proxy operations, some options have outer values, i.e. are present as options in the protected CoAP message. Certain options may have both an inner value



and a potentially different outer value, where the inner value is intended for the destination endpoint and the outer value is intended for the proxy.

A summary of how options are protected and processed is shown in Figure 4. Options within each class are protected and processed in a similar way, but certain options which require special processing are indicated by a \* in Figure 4 and described in the subsections below.

No.	Name	E	I	U
1	If-Match	x		
3	Uri-Host			x
4	ETag	x		
5	If-None-Match	x		
6	Observe		*	
7	Uri-Port			x
8	Location-Path	x		
11	Uri-Path	x		
12	Content-Format	x		
14	Max-Age	*		
15	Uri-Query	x		
17	Accept	x		
20	Location-Query	x		
23	Block2	*		
27	Block1	*		
28	Size2	*		
35	Proxy-Uri			*
39	Proxy-Scheme			x
60	Size1	*		

E=Encrypt and Integrity Protect, I=Integrity Protect only,  
U=Unprotected, \*=Special

Figure 4: Protection of CoAP Options

Unless specified otherwise, CoAP options not listed in Figure 4 SHALL be encrypted and integrity protected and processed as class E options.

Specifications of new CoAP options SHOULD define how they are processed with OSCOAP. New COAP options SHOULD be of class E and SHOULD NOT have outer values unless a forwarding proxy needs to read that option value. If a certain option has both inner and outer values, the two values SHOULD NOT be the same.





#### **4.3.1. Class E Options**

For options in class E (see Figure 4) the option value in the unprotected CoAP message, if present, SHALL be encrypted and integrity protected between the endpoints. Hence the actions resulting from the use of such options is analogous to communicating in a protected manner directly with the endpoint. For example, a client using an If-Match option will not be served by a proxy.

The sending endpoint SHALL write the class E option from the unprotected CoAP message into the plaintext of the COSE object.

Except for the special options described in the subsections, the sending endpoint SHALL NOT use the outer options of class E. However, note that an intermediary may, legitimately or not, add, change or remove the value of an outer option.

Except for the Block options [Section 4.3.1.2](#), the receiving endpoint SHALL discard any outer options of class E from the protected CoAP message and SHALL write the Class E options present in the plaintext of the COSE object into the unprotected CoAP message.

##### **4.3.1.1. Max-Age**

An inner Max-Age option, like other class E options, is used as defined in [\[RFC7252\]](#) taking into account that it is not accessible to proxies.

Since OSCOAP binds CoAP responses to requests, a cached response would not be possible to use for any other request. To avoid unnecessary caching, a server MAY add an outer Max-Age option with value zero to protected CoAP responses (see [Section 5.6.1 of \[RFC7252\]](#)). The outer Max-Age option is not integrity protected.

##### **4.3.1.2. The Block Options**

Blockwise [\[RFC7959\]](#) is an optional feature. An implementation MAY comply with [\[RFC7252\]](#) and the Object-Security option without implementing [\[RFC7959\]](#).

The Block options (Block1, Block2, Size1 and Size2) MAY be either only inner options, only outer options or both inner and outer options. The inner and outer options are processed independently.

The inner block options are used for endpoint-to-endpoint secure fragmentation of payload into blocks and protection of information about the fragmentation (block number, block size, last block). In this case, the CoAP client fragments the CoAP message as defined in



[RFC7959] before the message is processed by OSCOAP. The CoAP server first processes the OSCOAP message before processing blockwise as defined in [RFC7959].

There SHALL be a security policy defining a maximum unfragmented message size for inner Block options such that messages exceeding this size SHALL be fragmented by the sending endpoint.

Additionally, a proxy may arbitrarily do block fragmentation on any CoAP message, in particular an OSCOAP message, as defined in [RFC7959] and thereby add outer Block options to a block and send on the next hop. The outer block options are thus neither encrypted nor integrity protected.

An endpoint receiving a message with an outer Block option SHALL first process this option according to [RFC7959], until all blocks of the protected CoAP message has been received, or the cumulated message size of the exceeds the maximum unfragmented message size. In the latter case the message SHALL be discarded. In the former case, the processing of the protected CoAP message continues as defined in this document.

If the unprotected CoAP message in turn contains Block options, the receiving endpoint processes this according to [RFC7959].

TODO: Update processing to support multiple concurrently proceeding requests

#### **4.3.2. Class I Options**

A Class I option is an outer option and hence visible in the options part of the protected CoAP message. Except for special options described in the subsections, for options in Class I (see Figure 4) the option value SHALL be integrity protected between the endpoints, see (Section 5.2). Unless otherwise specified, the sending endpoint SHALL encode the Class I options in the protected CoAP message as described in Section 4.3.4.

##### **4.3.2.1. Observe**

Observe [RFC7641] is an optional feature. An implementation MAY support [RFC7252] and the Object-Security option without supporting [RFC7641]. The Observe option as used here targets the requirements on forwarding of [I-D.hartke-core-e2e-security-reqs] (Section 2.2.1.2).



In order for a proxy to support forwarding of Observe messages, there must be an Observe option present in options part of the protected CoAP message ([[RFC7641](#)]), so Observe must have an outer value:

- o The Observe option of the unprotected CoAP request SHALL be encoded in the protected CoAP request as described in [Section 4.3.4](#).

To secure the order of the notifications, responses with the Observe option SHALL be integrity protected in the following way:

- o The Observe option SHALL be included in the external\_aad of the response (see [Section 5.2](#)), with value set to the 3 least significant bytes of the Sequence Number of the response.

The Observe option in the CoAP request SHALL NOT be integrity protected, since it may be legitimately removed by a proxy.

If the Observe option is removed from a CoAP request by a proxy, then the server can still verify the request (as a non-Observe request), and produce a non-Observe response. If the OSCOAP client receives a response to an Observe request without an outer Observe value, then it MUST verify the response as a non-Observe response, i.e. not include the Sequence Number of the response in the external\_aad.

#### **[4.3.3](#). Class U Options**

Options in Class U have outer values and are used to support forward proxy operations. Unless otherwise specified, the sending endpoint SHALL encode the Class U options in the options part of the protected CoAP message as described in [Section 4.3.4](#).

##### **[4.3.3.1](#). Uri-Host, Uri-Port, and Proxy-Scheme**

The sending endpoint SHALL copy Uri-Host, Uri-Port, and Proxy-Scheme from the unprotected CoAP message to the options part of the protected CoAP message. When Uri-Host, Uri-Port, or Proxy-Scheme options are present, Proxy-Uri is not used [[RFC7252](#)].

##### **[4.3.3.2](#). Proxy-Uri**

Proxy-Uri, when present, is split by OSCOAP into class U options and class E options, which are processed accordingly. When Proxy-Uri is used in the unprotected CoAP message, Uri-\* are not present [[RFC7252](#)].

The sending endpoint SHALL first decompose the Proxy-Uri value of the unprotected CoAP message into the Proxy-Scheme, Uri-Host, Uri-Port,



Uri-Path and Uri-Query options (if present) according to [section 6.4 of \[RFC7252\]](#).

Uri-Path and Uri-Query are class E options and MUST be protected and processed as if obtained from the unprotected CoAP message, see [Section 4.3.1](#).

The value of the Proxy-Uri option of the protected CoAP message MUST be replaced with Proxy-Scheme, Uri-Host and Uri-Port options (if present) composed according to [section 6.5 of \[RFC7252\]](#) and MUST be processed as a class U option, see [Section 4.3.3](#).

An example of how Proxy-Uri is processed is given here. Assume that the unprotected CoAP message contains:

- o Proxy-Uri = "coap://example.com/resource?q=1"

During OSCOAP processing, Proxy-Uri is split into:

- o Proxy-Scheme = "coap"
- o Uri-Host = "example.com"
- o Uri-Port = "5863"
- o Uri-Path = "resource"
- o Uri-Query = "q=1"

Uri-Path and Uri-Query follow the processing defined in [Section 4.3.1](#), and are thus encrypted and transported in the COSE object. The remaining options are composed into the Proxy-Uri included in the options part of the protected CoAP message, which has value:

- o Proxy-Uri = "coap://example.com"

#### **[4.3.4](#). Outer Options in the Protected CoAP Message**

All options with outer values present in the protected CoAP message, including the Object-Security option, SHALL be encoded as described in [Section 3.1 of \[RFC7252\]](#), where the delta is the difference to the previously included outer option value.





## 5. The COSE Object

This section defines how to use COSE [[I-D.ietf-cose-msg](#)] to wrap and protect data in the unprotected CoAP message. OSCOAP uses the untagged COSE\_Encrypt0 structure with an Authenticated Encryption with Additional Data (AEAD) algorithm. The key lengths, IV lengths, and maximum sequence number are algorithm dependent.

The AEAD algorithm AES-CCM-64-64-128 defined in Section 10.2 of [[I-D.ietf-cose-msg](#)] is mandatory to implement. For AES-CCM-64-64-128 the length of Sender Key and Recipient Key is 128 bits, the length of nonce, Sender IV, and Recipient IV is 7 bytes. The maximum Sequence Number is specified in [Section 10](#).

The nonce is constructed as described in Section 3.1 of [[I-D.ietf-cose-msg](#)], i.e. by padding the partial IV (Sequence Number in network byte order) with zeroes and XORing it with the Context IV (Sender IV or Recipient IV), with the following addition: The most significant bit in the first byte of the Context IV SHALL be flipped for responses, in case there is a unique response (not Observe). In this way, the same sequence number can be reused for requests and corresponding responses, which reduces the size of the responses in the most common case. For detailed processing instructions, see [Section 7](#).

We denote by Plaintext the data that is encrypted and integrity protected, and by Additional Authenticated Data (AAD) the data that is integrity protected only.

The COSE Object SHALL be a COSE\_Encrypt0 object with fields defined as follows

- o The "protected" field is empty.
- o The "unprotected" field includes:
  - \* The "Partial IV" parameter. The value is set to the Sequence Number. The Partial IV SHALL be of minimum length needed to encode the sequence number. This parameter SHALL be present in requests. In case of Observe ([Section 4.3.2.1](#)) the Partial IV SHALL be present in the response, and otherwise the Partial IV SHALL NOT be present in the response.
  - \* The "kid" parameter. The value is set to the Sender ID (see [Section 3](#)). This parameter SHALL be present in requests and SHALL NOT be present in responses.



- o The "ciphertext" field is computed from the Plaintext (see [Section 5.1](#)) and the Additional Authenticated Data (AAD) (see [Section 5.2](#)) following Section 5.2 of [\[I-D.ietf-cose-msg\]](#).

The encryption process is described in Section 5.3 of [\[I-D.ietf-cose-msg\]](#).

### 5.1. Plaintext

The Plaintext is formatted as a CoAP message without Header (see Figure 5) consisting of:

- o all Class E option values [Section 4.3.1](#) present in the unprotected CoAP message (see [Section 4.3](#)). The options are encoded as described in [Section 3.1 of \[RFC7252\]](#), where the delta is the difference to the previously included Class E option; and
- o the Payload of unprotected CoAP message, if present, and in that case prefixed by the one-byte Payload Marker (0xFF).

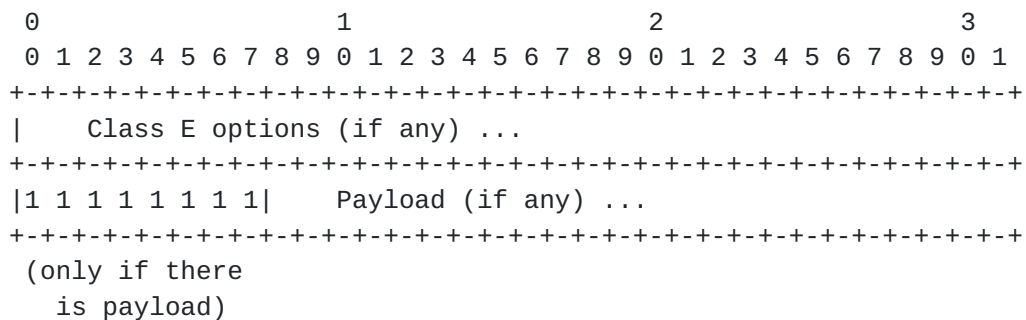


Figure 5: Plaintext

### 5.2. Additional Authenticated Data

The external\_aad SHALL be a CBOR array as defined below:

```

external_aad = [
  ver : uint,
  code : uint,
  options : bstr,
  alg : int,
  request_kid : bstr,
  request_seq : bstr
]

```

where:



- o ver: contains the CoAP version number, as defined in [Section 3 of \[RFC7252\]](#).
- o code: contains is the CoAP Code of the unprotected CoAP message, as defined in [Section 3 of \[RFC7252\]](#).
- o options: contains the Class I options [Section 4.3.2](#) present in the unprotected CoAP message encoded as described in [Section 3.1 of \[RFC7252\]](#), where the delta is the difference to the previously included class I option
- o alg: contains the Algorithm from the security context used for the exchange (see [Section 3.1](#)).
- o request\_kid: contains the value of the 'kid' in the COSE object of the request (see [Section 5](#)).
- o request\_seq: contains the value of the 'Partial IV' in the COSE object of the request (see [Section 5](#)).

## **6. Sequence Numbers, Replay, Message Binding, and Freshness**

Sequence numbers and replay window are initialized as defined in [Section 3.2.2](#).

### **6.1. AEAD Nonce Uniqueness**

An AEAD nonce MUST NOT be used more than once per AEAD key. In order to assure unique nonces, each Sender Context contains a Sequence Number used to protect requests, and - in case of Observe - responses. The maximum sequence number is algorithm dependent, see [Section 10](#). If the Sequence Number exceeds the maximum sequence number, the endpoint MUST NOT process any more messages with the given Sender Context. The endpoint SHOULD acquire a new security context (and consequently inform the other endpoint) before this happens. The latter is out of scope of this document.

### **6.2. Replay Protection**

In order to protect from replay of messages, each Recipient Context contains a Replay Window used to verify request, and - in case of Observe - responses. A receiving endpoint SHALL verify that a Sequence Number (Partial IV) received in the COSE object has not been received before in the Recipient Context. For requests, if this verification fails and the message received is a CON message, the server SHALL respond with a 4.00 Bad Request error message. The diagnostic payload MAY contain the "Replay protection failed" string. For responses, if this verification fails and the message received is



a CON message, the client SHALL respond with an empty ACK and stop processing the response.

The size and type of the Replay Window depends on the use case and lower protocol layers. In case of reliable and ordered transport from endpoint to endpoint, the recipient MAY just store the last received sequence number and require that newly received Sequence Numbers equals the last received Sequence Number + 1.

### **6.3. Sequence Number and Replay Window State**

To prevent reuse of the Nonce/Sequence Number with the same key, or from accepting replayed messages, a node needs to handle the situation of suddenly losing sequence number and replay window state in RAM, e.g. as a result of a reboot.

After boot, a node MAY reject to use existing security contexts from before it booted and MAY establish a new security context with each party it communicates, e.g. using EDHOC [[I-D.selander-ace-cose-ecdhe](#)]. However, establishing a fresh security context may have a non-negligible cost in terms of e.g. power consumption.

If a stored security context is to be used after reboot, then the node MUST NOT reuse a previous Sequence Number and MUST NOT accept previously accepted messages.

#### **6.3.1. The Basic Case**

To prevent reuse of Sequence Number, the node MAY perform the following procedure during normal operations:

- o Before sending a message, the client stores in persistent memory a sequence number associated to the stored security context higher than any sequence number which has been or are being sent using this security context. After boot, the client does not use any lower sequence number in a request than what was persistently stored with that security context.
- \* Storing to persistent memory can be costly. Instead of storing a sequence number for each request, the client may store Seq + K to persistent memory every K requests, where Seq is the current sequence number and  $K > 1$ . This is a trade-off between the number of storage operations and efficient use of sequence numbers.

To prevent accepting replay of previously received messages, the node MAY perform the following procedure:





- o After boot, before verifying a message using a security context stored before boot, the server synchronizes the replay window so that no old messages are being accepted. The server uses the Repeat option [[I-D.mattsson-core-coap-actuators](#)] for synchronizing the replay window: For each stored security context, the first time after boot the server receives an OSCOAP request, it generates a pseudo-random nonce and responds with the Repeat option set to the nonce as described in [[I-D.mattsson-core-coap-actuators](#)]. If the server receives a repeated OSCOAP request containing the Repeat option and the same nonce, and if the server can verify the request, then the sequence number obtained in the repeated message is set as the lower limit of the replay window.

#### **[6.3.2.](#) The Observe Case**

To prevent reuse of Sequence Number in case of Observe, the node MAY perform the following procedure during normal operations:

- o Before sending a notification, the server stores in persistent memory a sequence number associated to the stored security context higher than any sequence number for which a notification has been or are being sent using this security context. After boot, the server does not use any lower sequence number in an Observe response than what was persistently stored with that security context.
- \* Storing to persistent memory can be costly. Instead of storing a sequence number for each notification, the server may store  $\text{Seq} + K$  to persistent memory every  $K$  requests, where  $\text{Seq}$  is the current sequence number and  $K > 1$ . This is a trade-off between the number of storage operations and efficient use of sequence numbers.

Note that a client MAY continue an ongoing observation after reboot using a stored security context. With Observe, the client can only verify the order of the notifications, as they may be delayed. If the client wants to synchronize with a server resource it MAY restart an observation.

#### **[6.4.](#) Freshness**

For responses without Observe, OSCOAP provides absolute freshness. For requests, and responses with Observe, OSCOAP provides relative freshness in the sense that the sequence numbers allows a recipient to determine the relative order of messages.



For applications having stronger demands on freshness (e.g. control of actuators), OSCOAP needs to be augmented with mechanisms providing absolute freshness [[I-D.mattsson-core-coap-actuators](#)].

### **6.5. Delay and Mismatch Attacks**

In order to prevent response delay and mismatch attacks [[I-D.mattsson-core-coap-actuators](#)] from on-path attackers and compromised proxies, OSCOAP binds responses to the request by including the request's ID (Sender ID or Recipient ID) and sequence number in the AAD of the response. The server therefore needs to store the request's ID (Sender ID or Recipient ID) and sequence number until all responses have been sent.

## **7. Processing**

### **7.1. Protecting the Request**

Given an unprotected request, the client SHALL perform the following steps to create a protected request:

1. Retrieve the Sender Context associated with the target resource.
2. Compose the Additional Authenticated Data, as described in [Section 5](#).
3. Compose the AEAD nonce by XORing the Context IV (Sender IV) with the partial IV (Sequence Number in network byte order).
4. Encrypt the COSE object using the Sender Key. Compress the COSE Object as specified in [Section 8](#).
5. Format the protected CoAP message according to [Section 4](#). The Object-Security option is added, see [Section 4.3.4](#).
6. Store the association Token - Security Context. The client SHALL be able to find the Recipient Context from the Token in the response.
7. Increment the Sequence Number by one.

### **7.2. Verifying the Request**

A server receiving a request containing the Object-Security option SHALL perform the following steps:

1. Process outer Block options according to [[RFC7959](#)], until all blocks of the request have been received, see [Section 4.3.1.2](#).



2. Decompress the COSE Object ([Section 8](#)) and retrieve the Recipient Context associated with the Recipient ID in the 'kid' parameter. If the request is a CON message, and:

- \* either the decompression or the COSE message fails to decode, the server SHALL respond with a 4.02 Bad Option error message. The diagnostic payload SHOULD contain the string "Failed to decode COSE".
- \* the server fails to retrieve a Recipient Context with Recipient ID corresponding to the 'kid' parameter received, the server SHALL respond with a 4.01 Unauthorized error message. The diagnostic payload MAY contain the string "Security context not found".

If the request is a NON message and either the decompression or the COSE message fails to decode, or the server fails to retrieve a Recipient Context with Recipient ID corresponding to the 'kid' parameter received, then the server SHALL stop processing the request.

1. Verify the Sequence Number in the 'Partial IV' parameter, as described in [Section 6](#).
2. Compose the Additional Authenticated Data, as described in [Section 5](#).
3. Compose the AEAD nonce by XORing the Context IV (Recipient IV) with the padded 'Partial IV' parameter, received in the COSE Object.
4. Decrypt the COSE object using the Recipient Key.
  - \* If decryption fails, the server MUST stop processing the request and, if the request is a CON message, the server MUST respond with a 4.00 Bad Request error message. The diagnostic payload MAY contain the "Decryption failed" string.
  - \* If decryption succeeds, update the Recipient Replay Window, as described in [Section 6](#).
5. Add decrypted options and payload to the unprotected request, processing the E options as described in ([Section 4](#)). The Object-Security option is removed.
6. The unprotected CoAP request is processed according to [[RFC7252](#)]



### **7.3. Protecting the Response**

Given an unprotected response, the server SHALL perform the following steps to create a protected response:

1. Retrieve the Sender Context in the Security Context used to verify the request.
2. Compose the Additional Authenticated Data, as described in [Section 5](#).
3. Compose the AEAD nonce
  - \* If Observe is not used, compose the AEAD nonce by XORing the Context IV (Sender IV with the most significant bit in the first byte flipped) with the padded Partial IV parameter from the request.
  - \* If Observe is used, compose the AEAD nonce by XORing the Context IV (Sender IV) with the Partial IV of the response (Sequence Number in network byte order).
4. Encrypt the COSE object using the Sender Key. Compress the COSE Object as specified in [Section 8](#).
5. Format the protected CoAP message according to [Section 4](#). The Object-Security option is added, see [Section 4.3.4](#).
6. If Observe is used, increment the Sequence Number by one.

### **7.4. Verifying the Response**

A client receiving a response containing the Object-Security option SHALL perform the following steps:

1. Process outer Block options according to [\[RFC7959\]](#), until all blocks of the protected CoAP message have been received, see [Section 4.3.1.2](#).
2. Retrieve the Recipient Context associated with the Token. Decompress the COSE Object ([Section 8](#)). If the response is a CON message and either the decompression or the COSE message fails to decode, then the client SHALL send an empty ACK back and stop processing the response. If the response is a NON message and any of the previous conditions appear, then the client SHALL simply stop processing the response.





1. For Observe notifications, verify the Sequence Number in the 'Partial IV' parameter as described in [Section 6](#).
2. Compose the Additional Authenticated Data, as described in [Section 5](#).
3. Compose the AEAD nonce
  - \* If the Observe option is not present in the response, compose the AEAD nonce by XORing the Context IV (Recipient IV with the the most significant bit in the first byte flipped) with the padded Partial IV parameter from the request.
  - \* If the Observe option is present in the response, compose the AEAD nonce by XORing the Context IV (Recipient IV) with the padded Partial IV parameter from the response.
4. Decrypt the COSE object using the Recipient Key.
  - \* If decryption fails, the client MUST stop processing the response and, if the request is a CON message, the client MUST respond with an empty ACK back.
  - \* If decryption succeeds and Observe is used, update the Recipient Replay Window, as described in [Section 6](#).
5. Add decrypted options or payload to the unprotected response overwriting any outer E options (see [Section 4](#)). The Object-Security option is removed.
  - \* If Observe is used, replace the Observe value with the 3 least significant bytes in the sequence number.
6. The unprotected CoAP response is processed according to [\[RFC7252\]](#)

## **8. OSCOAP Compression**

The Concise Binary Object Representation (CBOR) [\[RFC7049\]](#) combines very small message sizes with extensibility. The CBOR Object Signing and Encryption (COSE) [\[I-D.ietf-cose-msg\]](#) uses CBOR to create compact encoding of signed and encrypted data. COSE is however constructed to support a large number of different stateless use cases, and is not fully optimized for use as a stateful security protocol, leading to a larger than necessary message expansion. In this section we define a simple stateless compression mechanism for OSCOAP, which significantly reduces the per-packet overhead.



### 8.1. Encoding of the Object-Security Option

The value of the Object-Security option SHALL be encoded as follows:

- o The first byte MUST encode a set of flags and the length of the Partial IV parameter.
  - \* The three least significant bits encode the Partial IV size. If their value is 0, the Partial IV is not present in the compressed message.
  - \* The fourth least significant bit is set to 1 if the kid is present in the compressed message.
  - \* The fifth-eighth least significant bits (= most significant half-byte) are reserved and SHALL be set to zero when not in use.
- o The following n bytes (n being the value of the Partial IV size in the first byte) encode the value of the Partial IV, if the Partial IV is present (size not 0).
- o The following byte encodes the size of the kid parameter, if the kid is present (flag bit set to 1)
- o The following m bytes (m given by the previous byte) encode the value of the kid, if the kid is present (flag bit set to 1)
- o The remainining bytes encode the ciphertext.

The presence of Partial IV and kid in requests and responses is specified in [Section 5](#), and summarized in Figure 6.

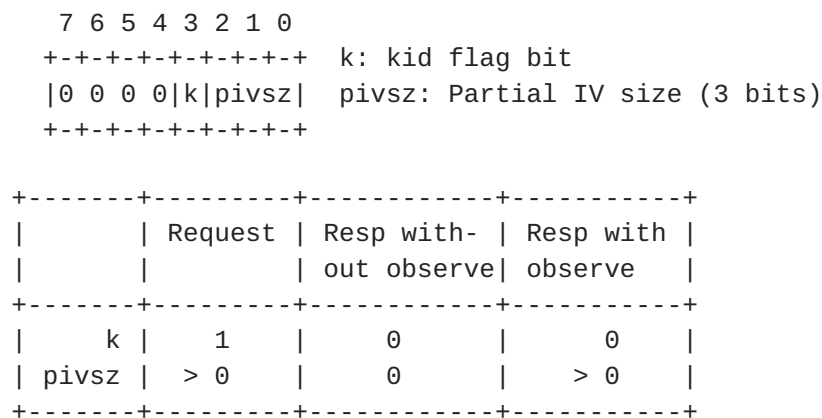


Figure 6: Flag byte for OSCOAP compression



## **8.2. Examples**

This section provides examples of COSE Objects before and after OSCOAP compression.

### **8.2.1. Example: Request**

Before compression:

```
[
  h'',
  { 4:h'25', 6:h'05' },
  h'aea0155667924dff8a24e4cb35b9'
]
```

0x83 40 a2 04 41 25 06 41 05 4e ae a0 15 56 67 92  
4d ff 8a 24 e4 cb 35 b9 (24 bytes)

After compression:

First byte: 0b00001001 = 0x09

0x09 05 01 25 ae a0 15 56 67 92 4d ff 8a 24 e4 cb  
35 b9 (18 bytes)

### **8.2.2. Example: Response (without Observe)**

Before compression:

```
[
  h'',
  {},
  h'aea0155667924dff8a24e4cb35b9'
]
```

0x83 40 a0 4e ae a0 15 56 67 92 4d ff 8a 24 e4 cb  
35 b9 (18 bytes)

After compression:

First byte: 0b00000000 = 0x00

0x00 ae a0 15 56 67 92 4d ff 8a 24 e4 cb 35 b9  
(15 bytes)



### **8.2.3. Example: Response (with Observe)**

Before compression:

```
[
  h'',
  { 6:h'07' },
  h'aea0155667924dff8a24e4cb35b9'
]
```

```
0x83 40 a1 06 41 07 4e ae a0 15 56 67 92 4d ff
8a 24 e4 cb 35 b9 (21 bytes)
```

After compression:

First byte: 0b00000001 = 0x01

```
0x01 07 ae a0 15 56 67 92 4d ff 8a 24 e4 cb 35 b9
(16 bytes)
```

## **9. Web Linking**

The use of OSCOAP MAY be indicated by a target attribute "osc" in a web link [[RFC5988](#)] to a CoAP resource. This attribute is a hint indicating that the destination of that link is to be accessed using OSCOAP. Note that this is simply a hint, it does not include any security context material or any other information required to run OSCOAP.

A value MUST NOT be given for the "osc" attribute; any present value MUST be ignored by parsers. The "osc" attribute MUST NOT appear more than once in a given link-value; occurrences after the first MUST be ignored by parsers.

## **10. Security Considerations**

In scenarios with intermediary nodes such as proxies or brokers, transport layer security such as DTLS only protects data hop-by-hop. As a consequence the intermediary nodes can read and modify information. The trust model where all intermediate nodes are considered trustworthy is problematic, not only from a privacy perspective, but also from a security perspective, as the intermediaries are free to delete resources on sensors and falsify commands to actuators (such as "unlock door", "start fire alarm", "raise bridge"). Even in the rare cases, where all the owners of the intermediary nodes are fully trusted, attacks and data breaches make such an architecture brittle.





DTLS protects hop-by-hop the entire CoAP message, including header, options, and payload. OSCOAP protects end-to-end the payload, and all information in the options and header, that is not required for forwarding (see [Section 4](#)). DTLS and OSCOAP can be combined, thereby enabling end-to-end security of CoAP payload, in combination with hop-by-hop protection of the entire CoAP message, during transport between end-point and intermediary node.

The CoAP message layer, however, cannot be protected end-to-end through intermediary devices since the parameters Type and Message ID, as well as Token and Token Length may be changed by a proxy. Moreover, messages that are not possible to verify should for security reasons not always be acknowledged but in some cases be silently dropped. This would not comply with CoAP message layer, but does not have an impact on the application layer security solution, since message layer is excluded from that.

The use of COSE to protect CoAP messages as specified in this document requires an established security context. The method to establish the security context described in [Section 3.2](#) is based on a common shared secret material in client and server, which may be obtained e.g. by using EDHOC [[I-D.selander-ace-cose-ecdhe](#)] or the ACE framework [[I-D.ietf-ace-oauth-authz](#)]. An OSCOAP profile of ACE is described in [[I-D.seitz-ace-oscoap-profile](#)].

The mandatory-to-implement AEAD algorithm AES-CCM-64-64-128 is selected for broad applicability in terms of message size ( $2^{64}$  blocks) and maximum number of messages ( $2^{56}$ ). Compatibility with CCM\* is achieved by using the algorithm AES-CCM-16-64-128 [[I-D.ietf-cose-msg](#)].

Most AEAD algorithms require a unique nonce for each message, for which the sequence numbers in the COSE message field "Partial IV" is used. If the recipient accepts any sequence number larger than the one previously received, then the problem of sequence number synchronization is avoided. With reliable transport it may be defined that only messages with sequence number which are equal to previous sequence number + 1 are accepted. The alternatives to sequence numbers have their issues: very constrained devices may not be able to support accurate time, or to generate and store large numbers of random nonces. The requirement to change key at counter wrap is a complication, but it also forces the user of this specification to think about implementing key renewal.

The maximum sequence number to guarantee nonce uniqueness ([Section 6.1](#)) is algorithm dependent. Using AES\_CCM, with the maximum sequence number SHALL be  $2^{(\min(\text{nonce length in bits}, 56) - 1)} - 1$ . The "-1" in the exponent stems from the same partial IV and



flipped bit of IV ([Section 5](#)) is used in request and response. The compression algorithm ([Section 8](#)) assumes that the partial IV is 56 bits or less (which is the reason for  $\min(,)$  in the exponent).

The inner block options enable the sender to split large messages into protected blocks such that the receiving node can verify blocks before having received the complete message. The outer block options allow for arbitrary proxy fragmentation operations that cannot be verified by the endpoints, but can by policy be restricted in size since the encrypted options allow for secure fragmentation of very large messages. A maximum message size (above which the sending endpoint fragments the message and the receiving endpoint discards the message, if complying to the policy) may be obtained as part of normal resource discovery.

Applications need to use a padding scheme if the content of a message can be determined solely from the length of the payload. As an example, the strings "YES" and "NO" even if encrypted can be distinguished from each other as there is no padding supplied by the current set of encryption algorithms. Some information can be determined even from looking at boundary conditions. An example of this would be returning an integer between 0 and 100 where lengths of 1, 2 and 3 will provide information about where in the range things are. Three different methods to deal with this are: 1) ensure that all messages are the same length. For example using 0 and 1 instead of 'yes' and 'no'. 2) Use a character which is not part of the responses to pad to a fixed length. For example, pad with a space to three characters. 3) Use the PKCS #7 style padding scheme where  $m$  bytes are appended each having the value of  $m$ . For example, appending a 0 to "YES" and two 1's to "NO". This style of padding means that all values need to be padded.

## **11. Privacy Considerations**

Privacy threats executed through intermediate nodes are considerably reduced by means of OSCOAP. End-to-end integrity protection and encryption of CoAP payload and all options that are not used for forwarding, provide mitigation against attacks on sensor and actuator communication, which may have a direct impact on the personal sphere.

The unprotected options (Figure 4) may reveal privacy sensitive information. In particular Uri-Host SHOULD NOT contain privacy sensitive information.

CoAP headers sent in plaintext allow for example matching of CON and ACK (CoAP Message Identifier), matching of request and responses (Token) and traffic analysis.



## **12. IANA Considerations**

Note to RFC Editor: Please replace all occurrences of "[[this document]]" with the RFC number of this specification.

### **12.1. CoAP Option Numbers Registry**

The Object-Security option is added to the CoAP Option Numbers registry:

+-----+	-----+	-----+	-----+
Number	Name	Reference	
+-----+	-----+	-----+	-----+
TBD	Object-Security	[[this document]]	
+-----+	-----+	-----+	-----+

### **12.2. Media Type Registrations**

The "application/oscon" media type is added to the Media Types registry:



Type name: application

Subtype name: oscon

Required parameters: N/A

Optional parameters: N/A

Encoding considerations: binary

Security considerations: See [Appendix C](#) of this document.

Interoperability considerations: N/A

Published specification: [[this document]] (this document)

Applications that use this media type: To be identified

Fragment identifier considerations: N/A

Additional information:

\* Magic number(s): N/A

\* File extension(s): N/A

\* Macintosh file type code(s): N/A

Person & email address to contact for further information:

Goeran Selander <goran.selander@ericsson.com>

Intended usage: COMMON

Restrictions on usage: N/A

Author: Goeran Selander, [goran.selander@ericsson.com](mailto:goran.selander@ericsson.com)

### **[12.3.](#) CoAP Content Format Registration**

The "application/oscon" content format is added to the CoAP Content Format registry:

Media type	Encoding	ID	Reference
application/oscon	-	70	[[this document]]





### **13. Acknowledgments**

The following individuals provided input to this document: Christian Amsuess, Carsten Bormann, Joakim Brorsson, Martin Gunnarsson, Klaus Hartke, Jim Schaad, Marco Tiloca, and Malisa Vuksinić.

Ludwig Seitz and Goeran Selander worked on this document as part of the CelticPlus project CyberWI, with funding from Vinnova.

### **14. References**

#### **14.1. Normative References**

- [I-D.ietf-cose-msg] Schaad, J., "CBOR Object Signing and Encryption (COSE)", [draft-ietf-cose-msg-24](#) (work in progress), November 2016.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC5869] Krawczyk, H. and P. Eronen, "HMAC-based Extract-and-Expand Key Derivation Function (HKDF)", [RFC 5869](#), DOI 10.17487/RFC5869, May 2010, <<http://www.rfc-editor.org/info/rfc5869>>.
- [RFC5988] Nottingham, M., "Web Linking", [RFC 5988](#), DOI 10.17487/RFC5988, October 2010, <<http://www.rfc-editor.org/info/rfc5988>>.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", [RFC 6347](#), DOI 10.17487/RFC6347, January 2012, <<http://www.rfc-editor.org/info/rfc6347>>.
- [RFC7049] Bormann, C. and P. Hoffman, "Concise Binary Object Representation (CBOR)", [RFC 7049](#), DOI 10.17487/RFC7049, October 2013, <<http://www.rfc-editor.org/info/rfc7049>>.
- [RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", [RFC 7252](#), DOI 10.17487/RFC7252, June 2014, <<http://www.rfc-editor.org/info/rfc7252>>.
- [RFC7641] Hartke, K., "Observing Resources in the Constrained Application Protocol (CoAP)", [RFC 7641](#), DOI 10.17487/RFC7641, September 2015, <<http://www.rfc-editor.org/info/rfc7641>>.



- [RFC7959] Bormann, C. and Z. Shelby, Ed., "Block-Wise Transfers in the Constrained Application Protocol (CoAP)", [RFC 7959](#), DOI 10.17487/RFC7959, August 2016, <<http://www.rfc-editor.org/info/rfc7959>>.

#### **14.2. Informative References**

- [I-D.bormann-6lo-coap-802-15-ie]  
Bormann, C., "Constrained Application Protocol (CoAP) over IEEE 802.15.4 Information Element for IETF", [draft-bormann-6lo-coap-802-15-ie-00](#) (work in progress), April 2016.
- [I-D.greevenbosch-appsawg-cbor-cddl]  
Birkholz, H., Vigano, C., and C. Bormann, "CBOR data definition language (CDDL): a notational convention to express CBOR data structures", [draft-greevenbosch-appsawg-cbor-cddl-10](#) (work in progress), March 2017.
- [I-D.hartke-core-e2e-security-reqs]  
Selander, G., Palombini, F., and K. Hartke, "Requirements for CoAP End-To-End Security", [draft-hartke-core-e2e-security-reqs-02](#) (work in progress), January 2017.
- [I-D.ietf-ace-oauth-authz]  
Seitz, L., Selander, G., Wahlstroem, E., Erdtman, S., and H. Tschofenig, "Authentication and Authorization for Constrained Environments (ACE)", [draft-ietf-ace-oauth-authz-06](#) (work in progress), March 2017.
- [I-D.ietf-core-coap-tcp-tls]  
Bormann, C., Lemay, S., Tschofenig, H., Hartke, K., Silverajan, B., and B. Raymor, "CoAP (Constrained Application Protocol) over TCP, TLS, and WebSockets", [draft-ietf-core-coap-tcp-tls-08](#) (work in progress), April 2017.
- [I-D.mattsson-core-coap-actuators]  
Mattsson, J., Fornehed, J., Selander, G., and F. Palombini, "Controlling Actuators with CoAP", [draft-mattsson-core-coap-actuators-02](#) (work in progress), November 2016.
- [I-D.seitz-ace-oscoap-profile]  
Seitz, L. and F. Palombini, "OSCOAP profile of ACE", [draft-seitz-ace-oscoap-profile-01](#) (work in progress), October 2016.



[I-D.selander-ace-cose-ecdhe]

Selander, G., Mattsson, J., and F. Palombini, "Ephemeral Diffie-Hellman Over COSE (EDHOC)", [draft-selander-ace-cose-ecdhe-06](#) (work in progress), April 2017.

[I-D.tiloca-core-multicast-oscoap]

Tiloca, M., Selander, G., and F. Palombini, "Secure group communication for CoAP", [draft-tiloca-core-multicast-oscoap-01](#) (work in progress), March 2017.

[RFC7228] Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", [RFC 7228](#), DOI 10.17487/RFC7228, May 2014, <<http://www.rfc-editor.org/info/rfc7228>>.

## [Appendix A](#). Test Vectors

TODO: This section needs to be updated.

## [Appendix B](#). Examples

This section gives examples of OSCOAP. The message exchanges are made, based on the assumption that there is a security context established between client and server. For simplicity, these examples only indicate the content of the messages without going into detail of the COSE message format.

### [B.1](#). Secure Access to Sensor

This example targets the scenario in Section 3.1 of [\[I-D.hartke-core-e2e-security-reqs\]](#) and illustrates a client requesting the alarm status from a server.



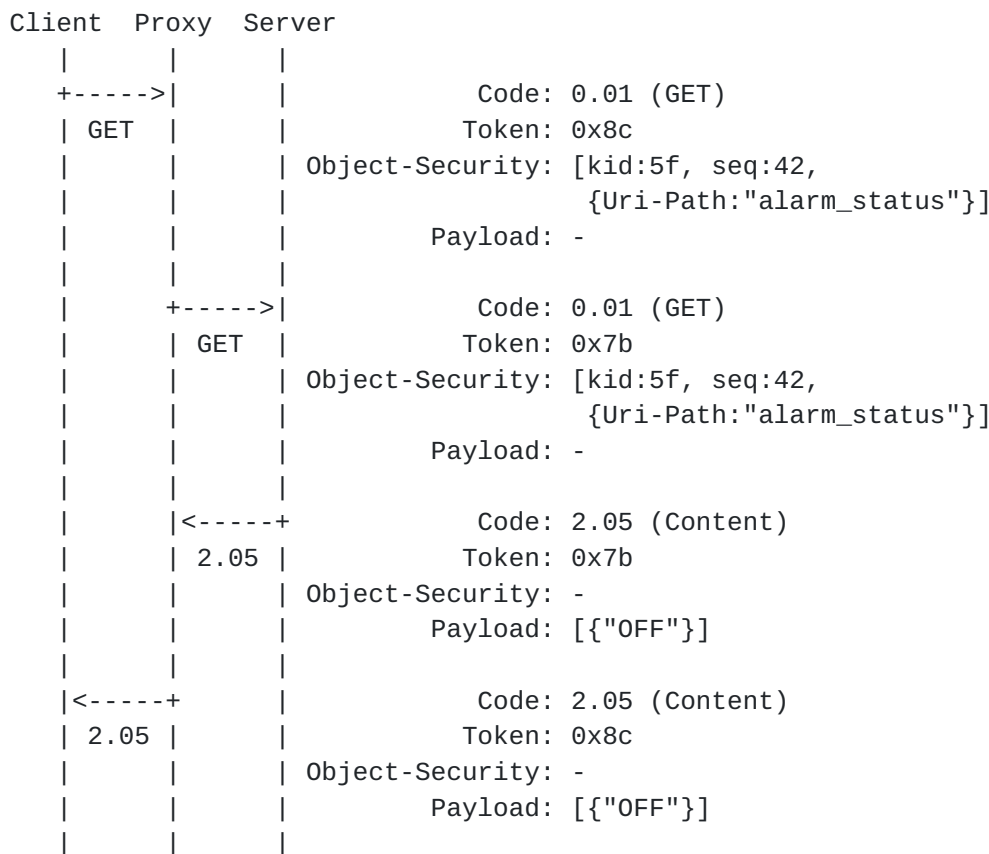


Figure 7: Secure Access to Sensor. Square brackets [ ... ] indicate a COSE object. Curly brackets { ... } indicate encrypted data.

Since the method (GET) doesn't allow payload, the Object-Security option carries the COSE object as its value. Since the response code (Content) allows payload, the COSE object is carried as the CoAP payload.

The COSE header of the request contains an identifier (5f), indicating which security context was used to protect the message and a sequence number (42). The option Uri-Path ("alarm\_status") and payload ("OFF") are encrypted.

The server verifies that the sequence number has not been received before. The client verifies that the response is bound to the request.

## B.2. Secure Subscribe to Sensor

This example targets the scenario in Section 3.2 of [\[I-D.hartke-core-e2e-security-reqs\]](#) and illustrates a client requesting subscription to a blood sugar measurement resource (GET





/glucose), first receiving the value 220 mg/dl and then a second value 180 mg/dl.

Client	Proxy	Server
+----->		Code: 0.01 (GET)
GET		Token: 0x83
		Observe: 0
		Object-Security: [kid:ca, seq:15,
		{Uri-Path:"glucose"}]
		Payload: -
	+----->	Code: 0.01 (GET)
	GET	Token: 0xbe
		Observe: 0
		Object-Security: [kid:ca, seq:15,
		{Uri-Path:"glucose"}]
		Payload: -
	<-----+	Code: 2.05 (Content)
	2.05	Token: 0xbe
		Observe: 000032
		Object-Security: -
		Payload: [seq:32, {Content-Format:0, "220"}]
<-----+		Code: 2.05 (Content)
2.05		Token: 0x83
		Observe: 000032
		Object-Security: -
		Payload: [seq:32, {Content-Format:0, "220"}]
...	...	...
	<-----+	Code: 2.05 (Content)
	2.05	Token: 0xbe
		Observe: 000036
		Object-Security: -
		Payload: [seq:36, {Content-Format:0, "180"}]
<-----+		Code: 2.05 (Content)
2.05		Token: 0x83
		Observe: 000036
		Object-Security: -
		Payload: [seq:36, {Content-Format:0, "180"}]

Figure 8: Secure Subscribe to Sensor. Square brackets [ ... ] indicate a COSE object. Curly brackets { ... } indicate encrypted data.



Since the method (GET) doesn't allow payload, the Object-Security option carries the COSE object as its value. Since the response code (Content) allows payload, the COSE object is carried as the CoAP payload.

The COSE header of the request contains an identifier (ca), indicating the security context used to protect the message and a Sequence Number (15). The COSE header of the responses contains sequence numbers (32 and 36). The options Content-Format (0) and the payload ("220" and "180"), are encrypted. The Observe option is integrity protected. The shown Observe values (000032 and 000036) are the ones that the client will see after OSCOAP processing.

The server verifies that the sequence number has not been received before. The client verifies that the sequence number has not been received before and that the responses are bound to the request.

#### [Appendix C](#). Object Security of Content (OSCON)

TODO: This section needs to be updated.

OSCOAP protects message exchanges end-to-end between a certain client and a certain server, targeting the security requirements for forward proxy of [[I-D.hartke-core-e2e-security-reqs](#)]. In contrast, many use cases require one and the same message to be protected for, and verified by, multiple endpoints, see caching proxy section of [[I-D.hartke-core-e2e-security-reqs](#)]. Those security requirements can be addressed by protecting essentially the payload/content of individual messages using the COSE format ([[I-D.ietf-cose-msg](#)]), rather than the entire request/response message exchange. This is referred to as Object Security of Content (OSCON).

OSCON transforms an unprotected CoAP message into a protected CoAP message in the following way: the payload of the unprotected CoAP message is wrapped by a COSE object, which replaces the payload of the unprotected CoAP message. We call the result the "protected" CoAP message.

The unprotected payload shall be the plaintext/payload of the COSE object. The 'protected' field of the COSE object 'Headers' shall include the context identifier, both for requests and responses. If the unprotected CoAP message includes a Content-Format option, then the COSE object shall include a protected 'content type' field, whose value is set to the unprotected message Content-Format value. The Content-Format option of the protected CoAP message shall be replaced with "application/oscon" ([Section 12](#))



The COSE object shall be protected (encrypted) and verified (decrypted) as described in ([[I-D.ietf-cose-msg](#)]).

Most AEAD algorithms require a unique nonce for each message. Sequence numbers for partial IV as specified for OSCOAP may be used for replay protection as described in [Section 6](#). The use of time stamps in the COSE header parameter 'operation time' [[I-D.ietf-cose-msg](#)] for freshness may be used.

OSCON shall not be used in cases where CoAP header fields (such as Code or Version) or CoAP options need to be integrity protected or encrypted. OSCON shall not be used in cases which require a secure binding between request and response.

The scenarios in Sections [3.3](#) - [3.5](#) of [[I-D.hartke-core-e2e-security-reqs](#)] assume multiple recipients for a particular content. In this case the use of symmetric keys does not provide data origin authentication. Therefore the COSE object should in general be protected with a digital signature.

### **[C.1](#). Overhead OSCON**

In general there are four different kinds of modes that need to be supported: message authentication code, digital signature, authenticated encryption, and symmetric encryption + digital signature. The use of digital signature is necessary for applications with many legitimate recipients of a given message, and where data origin authentication is required.

To distinguish between these different cases, the tagged structures of COSE are used (see Section 2 of [[I-D.ietf-cose-msg](#)]).

The sizes of COSE messages for selected algorithms are detailed in this section.

The size of the header is shown separately from the size of the MAC/signature. A 4-byte Context Identifier and a 1-byte Sequence Number are used throughout all examples, with these values:

- o Cid: 0xa1534e3c
- o Seq: 0xa3

For each scheme, we indicate the fixed length of these two parameters ("Cid+Seq" column) and of the Tag ("MAC"/"SIG"/"TAG"). The "Message OH" column shows the total expansions of the CoAP message size, while the "COSE OH" column is calculated from the previous columns.



Overhead incurring from CBOR encoding is also included in the COSE overhead count.

To make it easier to read, COSE objects are represented using CBOR's diagnostic notation rather than a binary dump.

### **C.2. MAC Only**

This example is based on HMAC-SHA256, with truncation to 8 bytes (HMAC 256/64).

Since the key is implicitly known by the recipient, the COSE\_Mac0\_Tagged structure is used (Section 6.2 of [\[I-D.ietf-cose-msg\]](#)).

The object in COSE encoding gives:

```

996(                                     # COSE_Mac0_Tagged
  [
    h'a20444a1534e3c0641a3', # protected:
                                {04:h'a1534e3c',
                                06:h'a3'}
    {},                        # unprotected
    h'',                      # payload
    MAC                       # truncated 8-byte MAC
  ]
)
```

This COSE object encodes to a total size of 26 bytes.

Figure 9 summarizes these results.

Structure	Tid	MAC	COSE OH	Message OH
COSE_Mac0_Tagged	5 B	8 B	13 B	26 B

Figure 9: Message overhead for a 5-byte Tid using HMAC 256/64

### **C.3. Signature Only**

This example is based on ECDSA, with a signature of 64 bytes.

Since only one signature is used, the COSE\_Sign1\_Tagged structure is used (Section 4.2 of [\[I-D.ietf-cose-msg\]](#)).

The object in COSE encoding gives:





```

997(                                     # COSE_Sign1_Tagged
  [
    h'a20444a1534e3c0641a3', # protected:
                                {04:h'a1534e3c',
                                06:h'a3'}
    {},                         # unprotected
    h'',                        # payload
    SIG                         # 64-byte signature
  ]
)

```

This COSE object encodes to a total size of 83 bytes.

Figure 10 summarizes these results.

Structure	Tid	SIG	COSE OH	Message OH
COSE_Sign1_Tagged	5 B	64 B	14 B	83 bytes

Figure 10: Message overhead for a 5-byte Tid using 64 byte ECDSA signature.

#### C.4. Authenticated Encryption with Additional Data (AEAD)

This example is based on AES-CCM with the Tag truncated to 8 bytes.

Since the key is implicitly known by the recipient, the COSE\_Encrypt0\_Tagged structure is used (Section 5.2 of [\[I-D.ietf-cose-msg\]](#)).

The object in COSE encoding gives:

```

993(                                     # COSE_Encrypt0_Tagged
  [
    h'a20444a1534e3c0641a3', # protected:
                                {04:h'a1534e3c',
                                06:h'a3'}
    {},                         # unprotected
    ciphertext                  # ciphertext including truncated 8-byte TAG
  ]
)

```

This COSE object encodes to a total size of 25 bytes.

Figure 11 summarizes these results.



Structure	Tid	TAG	COSE OH	Message OH
COSE_Encrypt0_Tagged	5 B	8 B	12 B	25 bytes

Figure 11: Message overhead for a 5-byte Tid using AES\_128\_CCM\_8.

### C.5. Symmetric Encryption with Asymmetric Signature (SEAS)

This example is based on AES-CCM and ECDSA with 64 bytes signature. The same assumption on the security context as in [Appendix C.4](#). COSE defines the field 'counter signature w/o headers' that is used here to sign a COSE\_Encrypt0\_Tagged message (see Section 3 of [\[I-D.ietf-cose-msg\]](#)).

The object in COSE encoding gives:

```

993(                                     # COSE_Encrypt0_Tagged
[
  h'a20444a1534e3c0641a3', # protected:
                           {04:h'a1534e3c',
                           06:h'a3'}
  {9:SIG},                 # unprotected:
                           09: 64 bytes signature
  ciphertext                # ciphertext including truncated 8-byte TAG
]
)

```

This COSE object encodes to a total size of 92 bytes.

Figure 12 summarizes these results.

Structure	Tid	TAG	SIG	COSE OH	Message OH
COSE_Encrypt0_Tagged	5 B	8 B	64 B	15 B	92 B

Figure 12: Message overhead for a 5-byte Tid using AES-CCM countersigned with ECDSA.

#### Authors' Addresses

Goeran Selander  
Ericsson AB

Email: [goran.selander@ericsson.com](mailto:goran.selander@ericsson.com)



John Mattsson  
Ericsson AB

Email: [john.mattsson@ericsson.com](mailto:john.mattsson@ericsson.com)

Francesca Palombini  
Ericsson AB

Email: [francesca.palombini@ericsson.com](mailto:francesca.palombini@ericsson.com)

Ludwig Seitz  
SICS Swedish ICT

Email: [ludwig@sics.se](mailto:ludwig@sics.se)