```
Workgroup: CoRE Working Group
Internet-Draft:
draft-ietf-core-oscore-key-update-00
Updates: <u>8613</u> (if approved)
Published: 6 December 2021
Intended Status: Standards Track
Expires: 9 June 2022
Authors: R. Höglund M. Tiloca
RISE AB RISE AB
Key Update for OSCORE (KUDOS)
```

# Abstract

Object Security for Constrained RESTful Environments (OSCORE) uses AEAD algorithms to ensure confidentiality and integrity of exchanged messages. Due to known issues allowing forgery attacks against AEAD algorithms, limits should be followed on the number of times a specific key is used for encryption or decryption. This document defines how two OSCORE peers must follow these limits and what steps they must take to preserve the security of their communications. Therefore, this document updates RFC8613. Furthermore, this document specifies Key Update for OSCORE (KUDOS), a lightweight procedure that two peers can use to update their keying material and establish a new OSCORE Security Context.

#### 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 <u>https://datatracker.ietf.org/drafts/current/</u>.

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 9 June 2022.

### Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents

(<u>https://trustee.ietf.org/license-info</u>) 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

- <u>1</u>. <u>Introduction</u>
- <u>1.1</u>. <u>Terminology</u>
- 2. AEAD Key Usage Limits in OSCORE
  - 2.1. Problem Overview
    - 2.1.1. Limits for 'q' and 'v'
  - 2.2. Additional Information in the Security Context
    - 2.2.1. Common Context
    - 2.2.2. <u>Sender Context</u>
    - 2.2.3. Recipient Context
  - 2.3. OSCORE Messages Processing
    - 2.3.1. Protecting a Request or a Response
    - <u>2.3.2</u>. <u>Verifying a Request or a Response</u>
- 3. Current methods for Rekeying OSCORE
- 4. Key Update for OSCORE (KUDOS)
  - 4.1. Extensions to the OSCORE Option
  - <u>4.2</u>. <u>Function for Security Context Update</u>
  - <u>4.3</u>. Establishment of the New OSCORE Security Context
    - <u>4.3.1</u>. <u>Client-Initiated Key Update</u>
    - <u>4.3.2</u>. <u>Server-Initiated Key Update</u>
  - <u>4.4</u>. <u>Retention Policies</u>
  - <u>4.5</u>. <u>Discussion</u>
- 5. <u>Security Considerations</u>
- <u>6</u>. <u>IANA Considerations</u>
  - 6.1. OSCORE Flag Bits Registry
- <u>7</u>. <u>References</u>
  - 7.1. Normative References
  - 7.2. Informative References

<u>Acknowledgments</u>

<u>Authors' Addresses</u>

# 1. Introduction

Object Security for Constrained RESTful Environments (OSCORE) [<u>RFC8613</u>] provides end-to-end protection of CoAP [<u>RFC7252</u>] messages at the application-layer, ensuring message confidentiality and integrity, replay protection, as well as binding of response to request between a sender and a recipient. In particular, OSCORE uses AEAD algorithms to provide confidentiality and integrity of messages exchanged between two peers. Due to known issues allowing forgery attacks against AEAD algorithms, limits should be followed on the number of times a specific key is used to perform encryption or decryption [I-D.irtfcfrg-aead-limits].

Should these limits be exceeded, an adversary may break the security properties of the AEAD algorithm, such as message confidentiality and integrity, e.g. by performing a message forgery attack. The original OSCORE specification [RFC8613] does not consider such limits.

This document updates [RFC8613] as follows.

\*It defines when a peer must stop using an OSCORE Security Context shared with another peer, due to the reached key usage limits. When this happens, the two peers have to establish a new Security Context with new keying material, in order to continue their secure communication with OSCORE.

\*It specifies KUDOS, a lightweight key update procedure that the two peers can use in order to update their current keying material and establish a new OSCORE Security Context. This deprecates and replaces the procedure specified in <u>Appendix B.2</u> of [<u>RFC8613</u>].

# **1.1. Terminology**

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 [<u>RFC2119</u>] [<u>RFC8174</u>] when, and only when, they appear in all capitals, as shown here.

Readers are expected to be familiar with the terms and concepts related to the CoAP [RFC7252] and OSCORE [RFC8613] protocols.

# 2. AEAD Key Usage Limits in OSCORE

The following sections details how key usage limits for AEAD algorithms must be considered when using OSCORE. It covers specific limits for common AEAD algorithms used with OSCORE; necessary additions to the OSCORE Security Context, updates to the OSCORE message processing, and existing methods for rekeying OSCORE.

#### 2.1. Problem Overview

The OSCORE security protocol [<u>RFC8613</u>] uses AEAD algorithms to provide integrity and confidentiality of messages, as exchanged between two peers sharing an OSCORE Security Context.

When processing messages with OSCORE, each peer should follow specific limits as to the number of times it uses a specific key. This applies separately to the Sender Key used to encrypt outgoing messages, and to the Recipient Key used to decrypt and verify incoming protected messages.

Exceeding these limits may allow an adversary to break the security properties of the AEAD algorithm, such as message confidentiality and integrity, e.g. by performing a message forgery attack.

The following refers to the two parameters 'q' and 'v' introduced in  $[\underline{I-D.irtf-cfrg-aead-limits}]$ , to use when deploying an AEAD algorithm.

\*'q': this parameter has as value the number of messages protected with a specific key, i.e. the number of times the AEAD algorithm has been invoked to encrypt data with that key.

\*'v': this parameter has as value the number of alleged forgery attempts that have been made against a specific key, i.e. the amount of failed decryptions that has been done with the AEAD algorithm for that key.

When a peer uses OSCORE:

\*The key used to protect outgoing messages is its Sender Key, in its Sender Context.

\*The key used to decrypt and verify incoming messages is its Recipient Key, in its Recipient Context.

Both keys are derived as part of the establishment of the OSCORE Security Context, as defined in <u>Section 3.2</u> of [<u>RFC8613</u>].

As mentioned above, exceeding specific limits for the 'q' or 'v' value can weaken the security properties of the AEAD algorithm used, thus compromising secure communication requirements.

Therefore, in order to preserve the security of the used AEAD algorithm, OSCORE has to observe limits for the 'q' and 'v' values, throughout the lifetime of the used AEAD keys.

# 2.1.1. Limits for 'q' and 'v'

Formulas for calculating the security levels as Integrity Advantage (IA) and Confidentiality Advantage (CA) probabilities, are presented in [<u>I-D.irtf-cfrg-aead-limits</u>]. These formulas take as input specific values for 'q' and 'v' (see section <u>Section 2.1</u>) and for 'l', i.e., the maximum length of each message (in cipher blocks).

For the algorithms that can be used as AEAD Algorithm for OSCORE shows in Figure 1, the key property to achieve is having IA and CA values which are no larger than  $p = 2^{-64}$ , which will ensure a safe security level for the AEAD Algorithm. This can be entailed by using the values  $q = 2^{20}$ ,  $v = 2^{20}$ , and  $l = 2^{10}$ , that this document recommends to use for these algorithms.

Figure 1 shows the resulting IA and CA probabilities enjoyed by the considered algorithms, when taking the value of 'q', 'v' and 'l' above as input to the formulas defined in [<u>I-D.irtf-cfrg-aead-limits</u>].

+	.+	+
Algorithm name	IA probability	CA probability
AEAD_AES_128_CCM	2^-64	2^-66
AEAD_AES_128_GCM	2^-97     2^-97	2^-89   2^-89
AEAD_CHACHA20_POLY1305	2^-73	-

Figure 1: Probabilities for algorithms based on chosen q, v and l values.

For the AEAD\_AES\_128\_CCM\_8 algorithm when used as AEAD Algorithm for OSCORE, larger IA and CA values are achieved, depending on the value of 'q', 'v' and 'l'. <u>Figure 2</u> shows the resulting IA and CA probabilities enjoyed by AEAD\_AES\_128\_CCM\_8, when taking different values of 'q', 'v' and 'l' as input to the formulas defined in [<u>1</u>\_<u>D.irtf-cfrg-aead-limits</u>].

As shown in Figure 2, it is especially possible to achieve the lowest IA =  $2^{-54}$  and a good CA =  $2^{-70}$  by considering the largest possible value of the (q, v, l) triplet equal to ( $2^{20}$ ,  $2^{10}$ ,  $2^{8}$ ), while still keeping a good security level. Note that the value of 'l' does not impact on IA, while CA displays good values for every considered value of 'l'.

When AEAD\_AES\_128\_CCM\_8 is used as AEAD Algorithm for OSCORE, this document recommends to use the triplet  $(q, v, 1) = (2^{20}, 2^{10}, 2^{8})$ 

and to never use a triplet (q, v, 1) such that the resulting IA and CA probabilities are higher than  $2^{-54}$ .

++	+	+
'q', 'v' and 'l'	IA probability	CA probability
+	+	·
q=2^20, v=2^20, l=2^8	2^-44	2^-70
q=2^15, v=2^20, l=2^8	2^-44	2^-80
q=2^10, v=2^20, l=2^8	2^-44	2^-90
q=2^20, v=2^15, l=2^8	2^-49	2^-70
q=2^15, v=2^15, l=2^8	2^-49	2^-80
q=2^10, v=2^15, l=2^8	2^-49	2^-90
q=2^20, v=2^14, l=2^8	2^-50	2^-70
q=2^15, v=2^14, l=2^8	2^-50	2^-80
q=2^10, v=2^14, l=2^8	2^-50	2^-90
q=2^20, v=2^10, l=2^8	2^-54	2^-70
q=2^15, v=2^10, l=2^8	2^-54	2^-80
q=2^10, v=2^10, l=2^8	2^-54	2^-90
+	+	••••••
q=2^20, v=2^20, l=2^6	2^-44	2^-74
q=2^15, v=2^20, l=2^6	2^-44	2^-84
q=2^10, v=2^20, l=2^6	2^-44	2^-94
q=2^20, v=2^15, l=2^6	2^-49	2^-74
q=2^15, v=2^15, l=2^6	2^-49	2^-84
q=2^10, v=2^15, l=2^6	2^-49	2^-94
q=2^20, v=2^14, l=2^6	2^-50	2^-74
q=2^15, v=2^14, l=2^6	2^-50	2^-84
q=2^10, v=2^14, l=2^6	2^-50	2^-94
q=2^20, v=2^10, l=2^6	2^-54	2^-74
q=2^15, v=2^10, l=2^6	2^-54	2^-84
q=2^10, v=2^10, l=2^6	2^-54	2^-94
++		+

Figure 2: Probabilities for AEAD\_AES\_128\_CCM\_8 based on chosen q, v and l values.

The algorithms using AES presented in this draft all use a block size of 16 bytes (128 bits), while AEAD\_CHACHA20\_POLY1305 uses a block size of 64 bytes (512 bits). As 'l' is defined as the maximum size of each message in blocks, different block sizes will result in different maximum messages sizes for the same value of 'l'. Figure 3 presents the resulting maximum message size in bytes for the different algorithms and values of 'l' presented in this document.

+	_ +		+
Algorithm name 	l=2^6 in   bytes	l=2^8 in     bytes	l=2^10 in   bytes
	-+	+	
AEAD_AES_128_CCM	1024	4096	16384
AEAD_AES_128_GCM	1024	4096	16384
AEAD_AES_256_GCM	1024	4096	16384
AEAD_AES_128_CCM_8	1024	4096	16384
AEAD_CHACHA20_POLY1305	4096	16384	65536
+	-+	++	+

Figure 3: Maximum length of each message (in bytes)

# 2.2. Additional Information in the Security Context

In addition to what defined in <u>Section 3.1</u> of [<u>RFC8613</u>], the OSCORE Security Context MUST also include the following information.

# 2.2.1. Common Context

The Common Context is extended to include the following parameter.

\*'exp': with value the expiration time of the OSCORE Security Context, as a non-negative integer. The parameter contains a numeric value representing the number of seconds from 1970-01-01T00:00:00Z UTC until the specified UTC date/time, ignoring leap seconds, analogous to what specified for NumericDate in Section 2 of [RFC7519].

At the time indicated in this field, a peer MUST stop using this Security Context to process any incoming or outgoing message, and is required to establish a new Security Context to continue OSCORE-protected communications with the other peer.

# 2.2.2. Sender Context

The Sender Context is extended to include the following parameters.

\*'count\_q': a non-negative integer counter, keeping track of the current 'q' value for the Sender Key. At any time, 'count\_q' has as value the number of messages that have been encrypted using the Sender Key. The value of 'count\_q' is set to 0 when establishing the Sender Context.

\*'limit\_q': a non-negative integer, which specifies the highest value that 'count\_q' is allowed to reach, before stopping using the Sender Key to process outgoing messages. The value of 'limit\_q' depends on the AEAD algorithm specified in the Common Context, considering the properties of that algorithm. The value of 'limit\_q' is determined according to <u>Section 2.1.1</u>.

# 2.2.3. Recipient Context

The Recipient Context is extended to include the following parameters.

\*'count\_v': a non-negative integer counter, keeping track of the current 'v' value for the Recipient Key. At any time, 'count\_v' has as value the number of failed decryptions occurred on incoming messages using the Recipient Key. The value of 'count\_v' is set to 0 when establishing the Recipient Context.

\*'limit\_v': a non-negative integer, which specifies the highest value that 'count\_v' is allowed to reach, before stopping using the Recipient Key to process incoming messages.

The value of 'limit\_v' depends on the AEAD algorithm specified in the Common Context, considering the properties of that algorithm. The value of 'limit\_v' is determined according to Section 2.1.1.

#### 2.3. OSCORE Messages Processing

In order to keep track of the 'q' and 'v' values and ensure that AEAD keys are not used beyond reaching their limits, the processing of OSCORE messages is extended as defined in this section. A limitation that is introduced is that, in order to not exceed the selected value for 'l', the total size of the COSE plaintext, authentication Tag, and possible cipher padding for a message may not exceed the block size for the selected algorithm multiplied with 'l'.

In particular, the processing of OSCORE messages follows the steps outlined in <u>Section 8</u> of [<u>RFC8613</u>], with the additions defined below.

# 2.3.1. Protecting a Request or a Response

Before encrypting the COSE object using the Sender Key, the 'count\_q' counter MUST be incremented.

If 'count\_q' exceeds the 'limit\_q' limit, the message processing MUST be aborted. From then on, the Sender Key MUST NOT be used to encrypt further messages.

### 2.3.2. Verifying a Request or a Response

If an incoming message is detected to be a replay (see <u>Section 7.4</u> of [<u>RFC8613</u>]), the 'count\_v' counter MUST NOT be incremented.

If the decryption and verification of the COSE object using the Recipient Key fails, the 'count\_v' counter MUST be incremented.

After 'count\_v' has exceeded the 'limit\_v' limit, incoming messages MUST NOT be decrypted and verified using the Recipient Key, and their processing MUST be aborted.

# 3. Current methods for Rekeying OSCORE

Before the limit of 'q' or 'v' defined in <u>Section 2.1.1</u> has been reached for an OSCORE Security Context, the two peers have to establish a new OSCORE Security Context, in order to continue using OSCORE for secure communication.

In practice, the two peers have to establish new Sender and Recipient Keys, as the keys actually used by the AEAD algorithm. When this happens, both peers reset their 'count\_q' and 'count\_v' values to 0 (see <u>Section 2.2</u>).

Other specifications define a number of ways to accomplish this, as summarized below.

\*The two peers can run the procedure defined in <u>Appendix B.2</u> of [<u>RFC8613</u>]. That is, the two peers exchange three or four messages, protected with temporary Security Contexts adding randomness to the ID Context.

As a result, the two peers establish a new OSCORE Security Context with new ID Context, Sender Key and Recipient Key, while keeping the same OSCORE Master Secret and OSCORE Master Salt from the old OSCORE Security Context.

This procedure does not require any additional components to what OSCORE already provides, and it does not provide perfect forward secrecy.

The procedure defined in <u>Appendix B.2</u> of [<u>RFC8613</u>] is used in 6TiSCH networks [<u>RFC7554</u>][<u>RFC8180</u>] when handling failure events. That is, a node acting as Join Registrar/Coordinator (JRC) assists new devices, namely "pledges", to securely join the network as per the Constrained Join Protocol [<u>RFC9031</u>]. In particular, a pledge exchanges OSCORE-protected messages with the JRC, from which it obtains a short identifier, link-layer keying material and other configuration parameters. As per <u>Section 8.3.3</u> of [<u>RFC9031</u>], a JRC that experiences a failure event may likely lose information about joined nodes, including their assigned identifiers. Then, the reinitialized JRC can establish a new OSCORE Security Context with each pledge, through the procedure defined in <u>Appendix B.2</u> of [<u>RFC8613</u>].

\*The two peers can run the OSCORE profile [<u>I-D.ietf-ace-oscore-</u> <u>profile</u>] of the Authentication and Authorization for Constrained Environments (ACE) Framework [<u>I-D.ietf-ace-oauth-authz</u>].

When a CoAP client uploads an Access Token to a CoAP server as an access credential, the two peers also exchange two nonces. Then, the two peers use the two nonces together with information provided by the ACE Authorization Server that issued the Access Token, in order to derive an OSCORE Security Context.

This procedure does not provide perfect forward secrecy.

\*The two peers can run the EDHOC key exchange protocol based on Diffie-Hellman and defined in [<u>I-D.ietf-lake-edhoc</u>], in order to establish a pseudo-random key in a mutually authenticated way.

Then, the two peers can use the established pseudo-random key to derive external application keys. This allows the two peers to securely derive especially an OSCORE Master Secret and an OSCORE Master Salt, from which an OSCORE Security Context can be established.

This procedure additionally provides perfect forward secrecy.

\*If one peer is acting as LwM2M Client and the other peer as LwM2M Server, according to the OMA Lightweight Machine to Machine Core specification [LwM2M], then the LwM2M Client peer may take the initiative to bootstrap again with the LwM2M Bootstrap Server, and receive again an OSCORE Security Context. Alternatively, the LwM2M Server can instruct the LwM2M Client to initiate this procedure.

If the OSCORE Security Context information on the LwM2M Bootstrap Server has been updated, the LwM2M Client will thus receive a fresh OSCORE Security Context to use with the LwM2M Server.

In addition to that, the LwM2M Client, the LwM2M Server as well as the LwM2M Bootstrap server are required to use the procedure defined in <u>Appendix B.2</u> of [<u>RFC8613</u>] and overviewed above, when they use a certain OSCORE Security Context for the first time [<u>LwM2M-Transport</u>].

Manually updating the OSCORE Security Context at the two peers should be a last resort option, and it might often be not practical or feasible. Even when any of the alternatives mentioned above is available, it is RECOMMENDED that two OSCORE peers update their Security Context by using the KUDOS procedure as defined in <u>Section 4</u> of this document.

It is RECOMMENDED that the peer initiating the key update procedure starts it before reaching the 'q' or 'v' limits. Otherwise, the AEAD keys possibly to be used during the key update procedure itself may already be or become invalid before the rekeying is completed, which may prevent a successful establishment of the new OSCORE Security Context altogether.

# 4. Key Update for OSCORE (KUDOS)

This section defines KUDOS, a lightweight procedure that two OSCORE peers can use to update their keying material and establish a new OSCORE Security Context.

KUDOS relies on the support function updateCtx() defined in <u>Section</u> <u>4.2</u> and the message exchange defined in <u>Section 4.3</u>. The following properties are fulfilled.

\*KUDOS can be initiated by either peer. In particular, the client or the server may start KUDOS by sending the first rekeying message.

\*The new OSCORE Security Context enjoys Perfect Forward Secrecy.

\*The same ID Context value used in the old OSCORE Security Context is preserved in the new Security Context. Furthermore, the ID Context value never changes throughout the KUDOS execution.

\*KUDOS is robust against a peer rebooting, and it especially avoids the reuse of AEAD (nonce, key) pairs.

\*KUDOS completes in one round trip. The two peers achieve mutual proof-of-possession in the following exchange, which is protected with the newly established OSCORE Security Context.

# 4.1. Extensions to the OSCORE Option

In order to support the message exchange for establishing a new OSCORE Security Context as defined in <u>Section 4.3</u>, this document extends the use of the OSCORE option originally defined in [<u>RFC8613</u>] as follows.

\*This document defines the usage of the seventh least significant bit, called "Extension-1 Flag", in the first byte of the OSCORE option containing the OSCORE flag bits. This flag bit is specified in <u>Section 6.1</u>. When the Extension-1 Flag is set to 1, the second byte of the OSCORE option MUST include the set of OSCORE flag bits 8-15.

\*This document defines the usage of the first least significant bit "ID Detail Flag", 'd', in the second byte of the OSCORE option containing the OSCORE flag bits. This flag bit is specified in <u>Section 6.1</u>.

When it is set to 1, the compressed COSE object contains an 'id detail', to be used for the steps defined in <u>Section 4.3</u>. In particular, the 1 byte following 'kid context' (if any) encodes the length x of 'id detail', and the following x bytes encode 'id detail'.

\*The second-to-eighth least significant bits in the second byte of the OSCORE option containing the OSCORE flag bits are reserved for future use. These bits SHALL be set to zero when not in use. According to this specification, if any of these bits are set to 1, the message is considered to be malformed and decompression fails as specified in item 2 of <u>Section 8.2</u> of [<u>RFC8613</u>].

Figure 4 shows the OSCORE option value including also 'id detail'.

| kid (if any) ... | +-----+

Figure 4: The OSCORE option value, including 'id detail'

#### 4.2. Function for Security Context Update

The updateCtx() function shown in <u>Figure 5</u> takes as input a nonce N as well as an OSCORE Security Context CTX\_IN, and returns as output a new OSCORE Security Context CTX\_OUT.

As a first step, the updateCtx() function derives the new values of the Master Secret and Master Salt for CTX\_OUT, according to one of the two following methods. The used method depends on how the two peers established their original Security Context, i.e., the Security Context that they shared before performing KUDOS with one another for the first time.

\*If the original Security Context was established by running the EDHOC protocol [<u>I-D.ietf-lake-edhoc</u>], the following applies.

First, the EDHOC key PRK\_4x3m shared by the two peers is updated using the EDHOC-KeyUpdate() function defined in <u>Section 4.4</u> of [<u>I-D.ietf-lake-edhoc</u>], which takes the nonce N as input.

After that, the EDHOC-Exporter() function defined in <u>Section 4.3</u> of [<u>I-D.ietf-lake-edhoc</u>] is used to derive the new values for the Master Secret and Master Salt, consistently with what is defined in <u>Appendix A.2</u> of [<u>I-D.ietf-lake-edhoc</u>]. In particular, the context parameter provided as second argument to the EDHOC-Exporter() function is the empty CBOR byte string (0x40) [<u>RFC8949</u>], which is denoted as h''.

Note that, compared to the compliance requirements in <u>Section 7</u> of [<u>I-D.ietf-lake-edhoc</u>], a peer MUST support the EDHOC-KeyUpdate() function, in case it establishes an original Security Context through the EDHOC protocol and intends to perform KUDOS.

\*If the original Security Context was established through other means than the EDHOC protocol, the new Master Secret is derived through an HKDF-Expand() step, which takes as input N as well as the Master Secret value from the Security Context CTX\_IN. Instead, the new Master Salt takes N as value.

In either case, the derivation of new values follows the same approach used in TLS 1.3, which is also based on HKDF-Expand (see <u>Section 7.1</u> of [<u>RFC8446</u>]) and used for computing new keying material in case of key update (see <u>Section 4.6.3</u> of [<u>RFC8446</u>]).

After that, the new Master Secret and Master Salt parameters are used to derive a new Security Context CTX\_OUT as per <u>Section 3.2</u> of [<u>RFC8613</u>]. Any other parameter required for the derivation takes the same value as in the Security Context CTX\_IN. Finally, the function returns the newly derived Security Context CTX\_OUT.

```
updateCtx(N, CTX_IN) {
  CTX_OUT
            // The new Security Context
  MSECRET_NEW // The new Master Secret
              // The new Master Salt
  MSALT_NEW
  if <the original Security Context was established through EDHOC> {
    EDHOC-KeyUpdate(N)
    // This results in updating the key PRK_4x3m of the
    // EDHOC session, i.e., PRK_4x3m = Extract(N, PRK_4x3m)
    MSECRET_NEW = EDHOC-Exporter("OSCORE_Master_Secret",
                                 h'', key_length)
      = EDHOC-KDF(PRK_4x3m, TH_4,
                  "OSCORE_Master_Secret", h'', key_length)
    MSALT_NEW = EDHOC-Exporter("OSCORE_Master_Salt",
                               h'', salt_length)
      = EDHOC-KDF(PRK_4x3m, TH_4,
                  "OSCORE_Master_Salt", h'', salt_length)
  }
  else {
    Master Secret Length = < Size of CTX_IN.MasterSecret in bytes >
    MSECRET_NEW = HKDF-Expand-Label(CTX_IN.MasterSecret, Label,
                                    N, Master Secret Length)
                = HKDF-Expand(CTX_IN.MasterSecret, HkdfLabel,
                              Master Secret Length)
    MSALT_NEW = N;
  }
  < Derive CTX_OUT using MSECRET_NEW and MSALT_NEW,
    together with other parameters from CTX_IN >
  Return CTX_OUT;
}
Where HkdfLabel is defined as
struct {
    uint16 length = Length;
    opaque label<7..255> = "oscore " + Label;
    opaque context<0..255> = Context;
} HkdfLabel;
   Figure 5: Function for deriving a new OSCORE Security Context
```

#### 4.3. Establishment of the New OSCORE Security Context

This section defines the actual KUDOS procedure performed by two peers to update their OSCORE keying material. Before starting KUDOS, the two peers share the OSCORE Security Context CTX\_OLD. Once completed the KUDOS execution, the two peers agree on a newly established OSCORE Security Context CTX\_NEW.

In particular, each peer contributes by generating a fresh value R1 or R2, and providing it to the other peer. The byte string concatenation of the two values, hereafter denoted as R1 | R2, is used as input N by the updateCtx() function, in order to derive the new OSCORE Security Context CTX\_NEW. As for any new OSCORE Security Context, the Sender Sequence Number and the replay window are re-initialized accordingly (see Section 3.2.2 of [RFC8613]).

Once a peer has successfully derived the new OSCORE Security Context CTX\_NEW, that peer MUST terminate all the ongoing observations it has with the other peer as protected with the old Security Context CTX\_OLD.

Once a peer has successfully decrypted and verified an incoming message protected with CTX\_NEW, that peer MUST discard the old Security Context CTX\_OLD.

KUDOS can be started by the client or the server, as defined in <u>Section 4.3.1</u> and <u>Section 4.3.2</u>, respectively. The following properties hold for both the client- and server-initiated version of KUDOS.

\*The initiator always offers the fresh value R1.

\*The responder always offers the fresh value R2.

\*The responder is always the first one deriving the new OSCORE Security Context CTX\_NEW.

\*The initiator is always the first one achieving key confirmation, hence able to safely discard the old OSCORE Security Context CTX\_OLD.

\*Both the initiator and the responder use the same respective OSCORE Sender ID and Recipient ID. Also, they both preserve and use the same OSCORE ID Context from CTX\_OLD.

The length of the nonces R1, and R2 is application specific. The application needs to set the length of each nonce such that the probability of its value being repeated is negligible; typically, at least 8 bytes long.

# 4.3.1. Client-Initiated Key Update

Figure 6 shows the KUDOS workflow with the client acting as initiator.



#### Figure 6: Client-Initiated KUDOS Workflow

First, the client generates a random value R1, and uses the nonce N
= R1 together with the old Security Context CTX\_OLD, in order to
derive a temporary Security Context CTX\_1. Then, the client sends an
OSCORE request to the server, protected with the Security Context
CTX\_1. In particular, the request has the 'd' flag bit set to 1 and
specifies R1 as 'id detail' (see Section 4.1).

Upon receiving the OSCORE request, the server retrieves the value R1 from the 'id detail' of the request, and uses the nonce N = R1 together with the old Security Context CTX\_OLD, in order to derive the temporary Security Context CTX\_1. Then, the server verifies the request by using the Security Context CTX\_1.

After that, the server generates a random value R2, and uses the nonce N = R1 | R2 together with the old Security Context CTX\_OLD, in order to derive the new Security Context CTX\_NEW. Then, the server sends an OSCORE response to the client, protected with the new Security Context CTX\_NEW. In particular, the response has the 'd' flag bit set to 1 and specifies R2 as 'id detail'.

Upon receiving the OSCORE response, the client retrieves the value R2 from the 'id detail' of the response. Since the client has received a response to an OSCORE request it made with the 'd' flag bit set to 1, the client uses the nonce N = R1 | R2 together with the old Security Context CTX\_OLD, in order to derive the new Security Context CTX\_NEW. Finally, the client verifies the response by using the Security Context CTX\_NEW and deletes the old Security Context CTX\_OLD.

After that, the client can send a new OSCORE request protected with the new Security Context CTX\_NEW. When successfully verifying the request using the Security Context CTX\_NEW, the server deletes the old Security Context CTX\_OLD and can reply with an OSCORE response protected with the new Security Context CTX\_NEW.

From then on, the two peers can protect their message exchanges by using the new Security Context CTX\_NEW.

### 4.3.2. Server-Initiated Key Update

Figure 7 shows the KUDOS workflow with the server acting as initiator.



First, the client sends a normal OSCORE request to the server, protected with the old Security Context CTX\_OLD and with the 'd' flag bit set to 0.

Upon receiving the OSCORE request and after having verified it with the old Security Context CTX\_OLD as usual, the server generates a random value R1 and uses the nonce N = R1 together with the old Security Context CTX\_OLD, in order to derive a temporary Security Context CTX\_1. Then, the server sends an OSCORE response to the client, protected with the Security Context CTX\_1. In particular, the response has the 'd' flag bit set to 1 and specifies R1 as 'id detail' (see Section 4.1).

Upon receiving the OSCORE response, the client retrieves the value R1 from the 'id detail' of the response, and uses the nonce N = R1 together with the old Security Context CTX\_OLD, in order to derive the temporary Security Context CTX\_1. Then, the client verifies the response by using the Security Context CTX\_1.

After that, the client generates a random value R2, and uses the nonce N = R1 | R2 together with the old Security Context CTX\_OLD, in order to derive the new Security Context CTX\_NEW. Then, the client sends an OSCORE request to the server, protected with the new Security Context CTX\_NEW. In particular, the request has the 'd' flag bit set to 1 and specifies R1 | R2 as 'id detail'.

Upon receiving the OSCORE request, the server retrieves the value R1 | R2 from the request. Then, the server verifies that: i) the value R1 is identical to the value R1 specified in a previous OSCORE response with the 'd' flag bit set to 1; and ii) the value R1 | R2 has not been received before in an OSCORE request with the 'd' flag bit set to 1. If the verification succeeds, the server uses the nonce N = R1 | R2 together with the old Security Context CTX\_OLD, in order to derive the new Security Context CTX\_NEW. Finally, the server verifies the request by using the Security Context CTX\_NEW and deletes the old Security Context CTX\_OLD.

After that, the server can send an OSCORE response protected with the new Security Context CTX\_NEW. When successfully verifying the response using the Security Context CTX\_NEW, the client deletes the old Security Context CTX\_OLD.

From then on, the two peers can protect their message exchanges by using the new Security Context CTX\_NEW.

### 4.4. Retention Policies

Applications MAY define policies that allows a peer to also temporarily keep the old Security Context CTX\_OLD, rather than simply overwriting it to become CTX\_NEW. This allows the peer to decrypt late, still on-the-fly incoming messages protected with CTX\_OLD.

When enforcing such policies, the following applies.

\*Outgoing messages MUST be protected by using only CTX\_NEW.

\*Incoming messages MUST first be attempted to decrypt by using CTX\_NEW. If decryption fails, a second attempt can use CTX\_OLD.

\*When an amount of time defined by the policy has elapsed since the establishment of CTX\_NEW, the peer deletes CTX\_OLD.

# 4.5. Discussion

KUDOS is intended to deprecate and replace the procedure defined in <u>Appendix B.2</u> of [<u>RFC8613</u>], as fundamentally achieving the same goal, while displaying a number of improvements and advantages.

In particular, it is especially convenient for the handling of failure events concerning the JRC node in 6TiSCH networks (see <u>Section 3</u>). That is, among its intrinsic advantages compared to the procedure defined in <u>Appendix B.2</u> of [<u>RFC8613</u>], KUDOS preserves the same ID Context value, when establishing a new OSCORE Security Context.

Since the JRC uses ID Context values as identifiers of network nodes, namely "pledge identifiers", the above implies that the JRC does not have anymore to perform a mapping between a new, different ID Context value and a certain pledge identifier (see <u>Section 8.3.3</u> of [<u>RFC9031</u>]). It follows that pledge identifiers can remain constant once assigned, and thus ID Context values used as pledge identifiers can be employed in the long-term as originally intended.

# 5. Security Considerations

This document mainly covers security considerations about using AEAD keys in OSCORE and their usage limits, in addition to the security considerations of [RFC8613].

Depending on the specific key update procedure used to establish a new OSCORE Security Context, the related security considerations also apply.

TODO: Add more considerations.

### 6. IANA Considerations

This document has the following actions for IANA.

# 6.1. OSCORE Flag Bits Registry

IANA is asked to add the following entries to the "OSCORE Flag Bits" registry within the "Constrained RESTful Environments (CoRE) Parameters" registry group.

+	+	+	++
Bit	Name	Description	Reference
Position			
	Extension-1 Flag     	Set to 1 if the OSCORE   Option specifies a   second byte of OSCORE   flag bits	[This   Document]   
15	ID Detail Flag	Set to 1 if the	[This
		compressed COSE object	Document]
		contains 'id detail'	

# 7. References

# 7.1. Normative References

- [I-D.ietf-lake-edhoc] Selander, G., Mattsson, J. P., and F. Palombini, "Ephemeral Diffie-Hellman Over COSE (EDHOC)", Work in Progress, Internet-Draft, draft-ietf-lakeedhoc-12, 20 October 2021, <<u>https://www.ietf.org/archive/</u> id/draft-ietf-lake-edhoc-12.txt>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/ RFC2119, March 1997, <<u>https://www.rfc-editor.org/info/</u> rfc2119>.
- [RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", RFC 7252, DOI 10.17487/ RFC7252, June 2014, <<u>https://www.rfc-editor.org/info/</u> rfc7252>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<u>https://www.rfc-editor.org/info/rfc8174</u>>.

(OSCORE)", RFC 8613, DOI 10.17487/RFC8613, July 2019, <<u>https://www.rfc-editor.org/info/rfc8613</u>>.

[RFC8949] Bormann, C. and P. Hoffman, "Concise Binary Object Representation (CBOR)", STD 94, RFC 8949, DOI 10.17487/ RFC8949, December 2020, <<u>https://www.rfc-editor.org/info/</u> rfc8949>.

### 7.2. Informative References

- [I-D.ietf-ace-oauth-authz] Seitz, L., Selander, G., Wahlstroem, E., Erdtman, S., and H. Tschofenig, "Authentication and Authorization for Constrained Environments (ACE) using the OAuth 2.0 Framework (ACE-OAuth)", Work in Progress, Internet-Draft, draft-ietf-ace-oauth-authz-46, 8 November 2021, <<u>https://www.ietf.org/archive/id/draft-ietf-ace-oauth-authz-46.txt</u>>.
- [I-D.ietf-ace-oscore-profile] Palombini, F., Seitz, L., Selander, G., and M. Gunnarsson, "OSCORE Profile of the Authentication and Authorization for Constrained Environments Framework", Work in Progress, Internet-Draft, draft-ietf-ace-oscore-profile-19, 6 May 2021, <<u>https://www.ietf.org/archive/id/draft-ietf-ace-oscoreprofile-19.txt</u>>.
- [I-D.irtf-cfrg-aead-limits] Günther, F., Thomson, M., and C. A. Wood, "Usage Limits on AEAD Algorithms", Work in Progress, Internet-Draft, draft-irtf-cfrg-aead-limits-03, 12 July 2021, <<u>https://www.ietf.org/archive/id/draft-</u> irtf-cfrg-aead-limits-03.txt>.
- [LwM2M] Open Mobile Alliance, "Lightweight Machine to Machine Technical Specification - Core, Approved Version 1.2, OMA-TS-LightweightM2M\_Core-V1\_2-20201110-A", November 2020, <<u>http://www.openmobilealliance.org/release/</u> LightweightM2M/V1\_2-20201110-A/OMA-TS-LightweightM2M\_Core-V1\_2-20201110-A.pdf>.
- [LwM2M-Transport] Open Mobile Alliance, "Lightweight Machine to Machine Technical Specification - Transport Bindings, Approved Version 1.2, OMA-TS-LightweightM2M\_Transport-V1\_2-20201110-A", November 2020, <<u>http://</u> www.openmobilealliance.org/release/LightweightM2M/ V1\_2-20201110-A/OMA-TS-LightweightM2M\_Transport-V1\_2-20201110-A.pdf>.
- [RFC7519] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Token (JWT)", RFC 7519, DOI 10.17487/RFC7519, May 2015, <<u>https://www.rfc-editor.org/info/rfc7519</u>>.

# [RFC7554]

Watteyne, T., Ed., Palattella, M., and L. Grieco, "Using IEEE 802.15.4e Time-Slotted Channel Hopping (TSCH) in the Internet of Things (IoT): Problem Statement", RFC 7554, DOI 10.17487/RFC7554, May 2015, <<u>https://www.rfc-</u> editor.org/info/rfc7554>.

- [RFC8180] Vilajosana, X., Ed., Pister, K., and T. Watteyne, "Minimal IPv6 over the TSCH Mode of IEEE 802.15.4e (6TiSCH) Configuration", BCP 210, RFC 8180, DOI 10.17487/ RFC8180, May 2017, <<u>https://www.rfc-editor.org/info/</u> rfc8180>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS)
  Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446,
  August 2018, <<u>https://www.rfc-editor.org/info/rfc8446</u>>.
- [RFC9031] Vučinić, M., Ed., Simon, J., Pister, K., and M. Richardson, "Constrained Join Protocol (CoJP) for 6TiSCH", RFC 9031, DOI 10.17487/RFC9031, May 2021, <<u>https://www.rfc-editor.org/info/rfc9031</u>>.

# Acknowledgments

The authors sincerely thank Christian Amsuess, John Mattsson and Goeran Selander for their feedback and comments.

The work on this document has been partly supported by VINNOVA and the Celtic-Next project CRITISEC; and by the H2020 project SIFIS-Home (Grant agreement 952652).

# Authors' Addresses

Rikard Höglund RISE AB Isafjordsgatan 22 SE-16440 Stockholm Kista Sweden

Email: rikard.hoglund@ri.se

Marco Tiloca RISE AB Isafjordsgatan 22 SE-16440 Stockholm Kista Sweden

Email: marco.tiloca@ri.se