

Internet Engineering Task Force
Internet-Draft
Updates: 4210 (if approved)
Intended status: Standards Track
Expires: July 31, 2020

H. Brockhaus
Siemens
January 28, 2020

CMP Updates
draft-brockhaus-lamps-cmp-updates-03

Abstract

This document contains a set of updates to the base syntax of Certificate Management Protocol (CMP) version 2. This document updates RFC 4210.

Specifically, the CMP services updated in this document comprise the enabling of using EnvelopedData instead of EncryptedValue and the definition of extended key usages to identify certificates of CMP endpoints on certification and registration authorities.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on July 31, 2020.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect

to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. History of changes	2
2. Introduction	3
2.1. Convention and Terminology	3
3. Updates to RFC 4210 – Certificate Management Protocol (CMP)	4
3.1. New Section 1.1. – Changes since RFC 4210	4
3.2. New Section 4.5 – Extended Key Usage	5
3.3. Replace Section 5.1.3.4 – Multiple Protection	6
3.4. Replace Section 5.2.2. – Encrypted Values	7
3.5. Update Section 5.3.4. – Certification Response	9
3.6. Replace Section 5.3.19.9. – Revocation Passphrase	9
3.7. Update Section 5.3.22 – Polling Request and Response	10
3.8. Update Appendix B – The Use of Revocation Passphrase	11
3.9. Update Appendix C – Request Message Behavioral Clarifications	12
3.10. Update Appendix D.4. – Initial Registration/Certification (Basic Authenticated Scheme)	12
4. IANA Considerations	12
5. Security Considerations	13
6. Acknowledgements	13
7. References	13
7.1. Normative References	13
7.2. Informative References	14
Appendix A. ASN.1 Modules	14
Author's Address	15

1. History of changes

From version 02 -> 03:

- o Added some clarification in Section 3.1.

From version 01 -> 02:

- o Added clarification to section on multiple protection
- o Added clarification on new EKUs after some exchange with Tomas Gustavsson
- o Reused OIDs from RFC 6402 [RFC6402] as suggested by Sean Turner at IETF 106

- o Added clarification on the field containing the key identifier for a revocation passphrase
- o Minor changes in wording

From version 00 -> 01:

- o Added a section describing the new extended key usages
- o Completed the section on changes to the specification of encrypted values
- o Added a section on clarification to Appendix D.4
- o Minor generalization in RFC 4210 [RFC4210] Sections 5.1.3.4 and 5.3.22
- o Minor changes in wording

2. Introduction

While using CMP [RFC4210] in industrial and IoT environments and developing the Lightweight CMP Profile [I-D.brockhaus-lamps-lightweight-cmp-profile] some limitations were identified in the original CMP specification. This document updates RFC 4210 [RFC4210] to overcome these limitations.

In general, this document aims to improve the crypto agility of CMP to be flexible to react on future advances in cryptography.

This document also introduces new extended key usages to identify CMP endpoints on registration and certification authorities.

2.1. Convention and 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 RFC 2119 [RFC2119].

In this document, these words will appear with that interpretation only when in ALL CAPS. Lower case uses of these words are not to be interpreted as carrying significance described in RFC 2119.

Technical terminology is used in conformance with RFC 4210 [RFC4210], RFC 4211 [RFC4211], and RFC 5280 [RFC5280]. The following key words are used:

CA: Certification authority, which issues certificates.

RA: Registration authority, an optional system component to which a CA delegates certificate management functions such as authorization checks.

KGA: Key generation authority, which generates key pairs on behalf of an EE. The KGA could be co-located with an RA or a CA.

EE: End entity, a user, device, or service that holds a PKI certificate. An identifier for the EE is given as its subject of the certificate.

3. Updates to RFC 4210 - Certificate Management Protocol (CMP)

3.1. New Section 1.1. - Changes since RFC 4210

The following subsections describe feature updates to RFC 4210 [RFC4210]. They are always related to the base specification. Hence references to the original sections in RFC 4210 [RFC4210] are used whenever possible.

Insert this section at the end of the current Section 1.

The following updates are made in draft-brockhaus-lamps-cmp-updates:

- o Add new extended key usages for different CMP server types, e.g. registration authority and certification authority, to express the authorization of the entity identified in the certificate containing the respective extended key usage extension to act as the indicated PKI management component.
- o Extend the description of multiple protection to cover additional use cases, e.g., batch processing of messages.
- o Offering EnvelopedData as another choice next to EncryptedValue to extend crypto agility in CMP. Note that according to RFC 4211 [RFC4211] section 2.1.9 the use of the EncryptedValue structure has been deprecated in favor of the EnvelopedData structure. For reasons of completeness and consistency the exchange of EncryptedValue with EncryptedKey is performed not only where required for the needed crypto agility for protection of centrally generated private keys, but also for other purposes like encryption of certificates and revocation passphrases.
- o Extend the usage of polling also to p10cr messages.

3.2. New Section 4.5 - Extended Key Usage

Insert this section.

The Extended Key Usage (EKU) extension indicates the purposes for which the certified public key may be used. It therefore restricts the use of a certificate to specific applications. Certificates used for CMP message protection or signed data for central key generation SHOULD use one of the following EKUs to express its authorization for acting as the PKI management entities described below. The ASN.1 to define these EKUs is:

```
id-kp-cmpCA OBJECT IDENTIFIER ::= { id-kp 27 }
id-kp-cmpRA OBJECT IDENTIFIER ::= { id-kp 28 }
id-kp-cmpKGA OBJECT IDENTIFIER ::= { id-kp ... }
```

< TBD: id-kp-cmpKGA to be defined. >

Note: RFC 6402 section 2.10 [RFC6402] specifies OIDs for a CMC CA and a CMC RA. As the functionality of a CA and RA is not specific to whether use CMC or CMP as certificate management protocol, the same OIDs SHALL be used for a cmpCA and a cmpRA.

< TBD: It needs to be clarified, if the Name and Description of the OIDs can be adapted or extended to avoid confusion as they currently only refer to CMC endpoints. >

The description of the PKI entity for each of the EKUs is as follows:

CMP Certification Authorities are CMP endpoints on CA equipment as described in section 3.1.1.2. The key used in the context of CMP management operations, especially CMP message protection, need not be the same key that signs the certificates. It is necessary, however, to ensure that the entity acting as cmpCA is authorized to do so. Therefore, the cmpCA MUST do one of the following,

- o use the CA private key on the CMP endpoint, or
- o explicitly designate this authority to another entity.

CMP message protection delegation on the CA SHALL be designated by the inclusion of id-kp-cmpCA in an extended key usage certificate extension included in the CMP response signer's certificate. This certificate MUST be issued directly by the CA that is identified in the request.

Note: Using a separate key pair for protecting CMP management operations at the CA decreases the number of operations of the private key used to sign certificates.

CMP Registration Authorities are CMP endpoints on RA equipment as described in section 3.1.1.3. A cmpRA is identified by the id-kp-cmpRA extended key usage. This extended key usage is placed into RA certificates.

CMP Key Generation Authorities are identified by the id-kp-cmpKGA extended key usage. Though the cmpKGA knows the private key it generated on behalf of the end entity, this is a very sensible service and needs specific authorization. This authorization is either with the CA certificate itself, or indicated by placing the id-kp-cmpKGA extended key usage into the cmpRA or cmpCA certificate used to authenticate the origin of the private key to express the authorization to offer this service.

Note: In device PKIs, especially those issuing IDevID certificates, CA may have very long validity (including the GeneralizedTime value 99991231235959Z to indicate an indefinite expiration date as specified in IEEE 802.1AR section 8.5 [IEEE802.1AR] and RFC 5280 Section 4.1.2.5 [RFC5280]). Such validity periods SHOULD NOT be used for protection of CMP messages. Certificates for delegated CMP message protection (cmpCA, cmpRA, cmpKGA) MUST NOT use indefinite expiration date.

< TBD: In bigger PKI installations the CA equipment may host, and an RA equipment may serve several CAs. These CAs, especially those issuing IDevID certificates may have very long validities and use specific algorithms not suitable for protection of day-to-day PKI management operations on a CMC, CMP or TLS level. Therefore, it may be an advantage to utilize a specific 'Infrastructure' CA for issuing CMC, CMP and TLS certificates to protect PKI management operations for other CAs hosted on that PKI. A mechanism would be needed to securely delegate authorization to act as a cmpCA, cmpRA, or cmpKGA for a specific CA without directly issuing the cmpCA, cmpRA, and cmpKGA certificates. I am happy for any suitable suggestions to address this issue. >

3.3. Replace Section 5.1.3.4 - Multiple Protection

Section 5.1.3.4 of RFC 4210 [RFC4210] describes the nested message. This document opens the usage of nested messages also for batch transport of PKI messages between different PKI management entities.

Replace the text of the section with the following text.

In cases where an end entity sends a protected PKI message to an RA, the RA MAY forward that message to a CA, adding its own protection (which MAY be a MAC or a signature, depending on the information and certificates shared between the RA and the CA). There are different use cases for such multi protected messages.

- o The RA confirms the validation and authorization of a message and forwards the original message unchanged.
- o The RA collects several messages and forwards them in a batch. This can for instance be used to bridge an off-line connection between two PKI management entities. In an up-stream connection request messages and in a down-stream connection response or announcement messages will be collected in the batch.
- o The RA modifies the message(s) in some way (e.g., add or modify particular field values or add new extensions) before forwarding them, then it MAY create its own desired PKIBody. In case the changes made by the RA to PKIMessage breaks the POP, the RA MUST either set the POP RAVerified or include the original PKIMessage from the EE in the generalInfo field of PKIHeader of the nested message (to force the CA to check POP on the original message). The infoType to be used in this situation is {id-it 15} (see Section 5.3.19 for the value of id-it) and the infoValue is PKIMessages (contents MUST be in the same order as the requests in PKIBody). For simplicity reasons, if batching is used in combination with inclusion of the original PKIMessage in the generalInfo field, all messages in the batch MUST be of the same type (e.g., ir).

These use cases are accomplished by nesting the messages sent by the end entity within a new PKI message. The structure used is as follows.

NestedMessageContent ::= PKIMessages

(The use of PKIMessages, a SEQUENCE OF PKIMessage, lets the RA batch the requests of several EEs in a single new message.)

3.4. Replace Section 5.2.2. - Encrypted Values

Section 5.2.2 of RFC 4210 [RFC4210] describes the usage of EncryptedValue to transport encrypted data. This document extends the encryption of data to also use EnvelopedData.

Replace the text of the section with the following text.

Where encrypted data (restricted, in this specification, to be either private keys, certificates or passwords) are sent in PKI messages, the EncryptedKey data structure is used.

```
EncryptedKey ::= CHOICE {  
    encryptedValue      EncryptedValue, -- deprecated  
    envelopedData      [0] EnvelopedData }
```

See CRMF [RFC4211] for EncryptedKey and EncryptedValue syntax and for EnvelopedData syntax see CMS [RFC5652]. Using the EncryptedKey data structure, the choice to either use EncryptedValue (for backward compatibility only) or EnvelopedData is offered. The use of the EncryptedValue structure has been deprecated in favor of the EnvelopedData structure. Therefore, it is recommended to use EnvelopedData.

The EncryptedKey data structure is used in CMP to either transport a private key, certificate or revocation passphrase in encrypted form.

EnvelopedData is used as follows:

- o Contains only one recipientInfo structure because the content is encrypted only for one recipient.
- o Contains private key in a SignedData structure as specified in CMS section 5 [RFC5652] signed by the Key Generation Authority.
- o Contains certificate or revocation passphrase directly in the encryptedContent field.

Note: When transferring a centrally generated private key in a certificate response message to the EE, the algorithm identifier and the associated public key will anyhow be transported in this response message. Therefore, the private key will not be delivered in a key package structure as specified in [RFC5958] and [RFC6032]. But the wrapping of the private key in a SignedData structure that is wrapped in the EnvelopedData structure as specified in [RFC6032] is applied.

The content of the EnvelopedData structure, as specified in CMS section 3 [RFC5652], MUST be encrypted using a newly generated symmetric content-encryption key. This content-encryption key MUST be securely provided to the recipient using one of three key management techniques.

The choice of the key management technique to be used by the sender depends on the credential available for the recipient:

- o Jointly shared secret: The content-encryption key will be protected using the symmetric key-encryption key management technique, as specified in CMS section 5.2.3 [RFC5652].
- o Recipient's certificate that contains a key usage extension asserting keyAgreement: The content-encryption key will be protected using the key agreement key management technique, as specified in CMS section 5.2.2 [RFC5652].
- o Recipient's certificate that contains a key usage extension asserting keyEncipherment: The content-encryption key will be protected using the key transport key management technique, as specified in CMS section 5.2.1 [RFC5652].

3.5. Update Section 5.3.4. - Certification Response

Section 5.3.4 of RFC 4210 [RFC4210] describes the Certification Response. This document updates the syntax by using EncryptedKey instead of EncryptedValue as described in Section 3.1 above.

Replace the ASN.1 syntax of CertifiedKeyPair and CertOrEncCert with the following text.

```
CertifiedKeyPair ::= SEQUENCE {  
    certOrEncCert      CertOrEncCert,  
    privateKey          [0] EncryptedKey      OPTIONAL,  
    -- see [CRMF] for comment on encoding  
    publicationInfo [1] PKIPublicationInfo  OPTIONAL  
}
```

```
CertOrEncCert ::= CHOICE {  
    certificate      [0] Certificate,  
    encryptedCert    [1] EncryptedKey  
}
```

Add the following paragraphs to the end of the section.

The use of EncryptedKey is described in section 5.2.2.

3.6. Replace Section 5.3.19.9. - Revocation Passphrase

Section 5.3.19.9 of RFC 4210 [RFC4210] describes the provisioning of a revocation passphrase for authenticating a later revocation request. This document updates the handling by using EncryptedKey instead of EncryptedValue to transport this information as described in Section 3.1 above.

Replace the text of the section with the following text.

This MAY be used by the EE to send a passphrase to a CA/RA for the purpose of authenticating a later revocation request (in the case that the appropriate signing private key is no longer available to authenticate the request). See Appendix B for further details on the use of this mechanism.

GenMsg: {id-it 12}, EncryptedKey
GenRep: {id-it 12}, < absent >

The use of EncryptedKey is described in section 5.2.2.

3.7. Update Section 5.3.22 - Polling Request and Response

Section 5.3.22 of RFC 4210 [RFC4210] describes when and how polling messages are used. This document adds the polling mechanism also to outstanding p10cr transactions.

Replace all paragraphs in front of the state machine diagram with the following text.

This pair of messages is intended to handle scenarios in which the client needs to poll the server in order to determine the status of an outstanding ir, cr, p10cr, or kur transaction (i.e., when the "waiting" PKIStatus has been received).

```
PollReqContent ::= SEQUENCE OF SEQUENCE {  
    certReqId    INTEGER }  
  
PollRepContent ::= SEQUENCE OF SEQUENCE {  
    certReqId    INTEGER,  
    checkAfter   INTEGER, -- time in seconds  
    reason       PKIFreeText OPTIONAL }
```

The following clauses describe when polling messages are used, and how they are used. It is assumed that multiple certConf messages can be sent during transactions. There will be one sent in response to each ip, cp, or kup that contains a CertStatus for an issued certificate.

- 1 In response to an ip, cp, or kup message, an EE will send a certConf for all issued certificates and, following the ack, a pollReq for all pending certificates.
- 2 In response to a pollReq, a CA/RA will return an ip, cp, or kup if one or more of the pending certificates is ready; otherwise, it will return a pollRep.

- 3 If the EE receives a pollRep, it will wait for at least as long as the checkAfter value before sending another pollReq.
- 4 If an ip, cp, or kup is received in response to a pollReq, then it will be treated in the same way as the initial response.

Note: A pl0cr message contains exactly one CertificationRequestInfo data structure as specified in PKCS#10 [RFC2986] but not certificate request number. Therefore, the certReqId MUST be set to 0 in all following messages of this transaction.

3.8. Update Appendix B - The Use of Revocation Passphrase

Appendix B of RFC 4210 [RFC4210] describes the usage of the revocation passphrase. As this document updates RFC 4210 [RFC4210] to utilize EncryptedKey instead of EncryptedValue as described in Section 3.1 above, the description is updated accordingly.

Replace the first bullet point of this section with the following text.

- o The OID and value specified in Section 5.3.19.9 of RFC 4210 [RFC4210] MAY be sent in a GenMsg message at any time, or MAY be sent in the generalInfo field of the PKIHeader of any PKIMessage at any time. (In particular, the EncryptedKey as described in section 5.2.2 may be sent in the header of the certConf message that confirms acceptance of certificates requested in an initialization request or certificate request message.) This conveys a revocation passphrase chosen by the entity (i.e., for use of EnvelopedData this is in the decrypted bytes of encryptedContent of the EnvelopedData structure and for use of EncryptedValue this is in the decrypted bytes of the encValue field) to the relevant CA/RA; furthermore, the transfer is accomplished with appropriate confidentiality characteristics.

Replace the third bullet point of this section with the following text.

- o When using EnvelopedData the unprotectedAttrs and when using EncryptedValue the valueHint field MAY contain a key identifier (chosen by the entity, along with the passphrase itself) to assist in later retrieval of the correct passphrase (e.g., when the revocation request is constructed by the entity and received by the CA/RA).

< TBD: The attribute structure containing the key identifier in the unprotectedAttr field could either be pkcs-9-at-friendlyName or pkcs-

9-at-localKeyId as specified in RFC 2985 section 5.5 [RFC2985]. Are there preferences for either one? >

3.9. Update Appendix C - Request Message Behavioral Clarifications

Appendix C of RFC 4210 [RFC4210] provides clarifications to the request message behavior. As this document updates RFC 4210 [RFC4210] to utilize EncryptedKey instead of EncryptedValue as described in Section 3.1 above, the description is updated accordingly.

Replace the note coming after the ASN.1 syntax of POPOPrivKey of this section with the following text.

```
-- *****
-- * the type of "thisMessage" is given as BIT STRING in RFC 4211
-- * [RFC4211]; it should be "EncryptedKey" (in accordance with
-- * Section 5.2.2, "Encrypted Values", of this specification).
-- * Therefore, this document makes the behavioral clarification of
-- * specifying that the contents of "thisMessage" MUST be encoded
-- * either as EnvelopedData or EncryptedValue (only for backward
-- * compatibility) and then wrapped in a BIT STRING. This allows
-- * the necessary conveyance and protection of the private key
-- * while maintaining bits-on-the-wire compatibility with RFC 4211
-- * [RFC4211].
-- *****
```

3.10. Update Appendix D.4. - Initial Registration/Certification (Basic Authenticated Scheme)

Appendix D.4 of RFC 4210 [RFC4210] provides the initial registration/certification scheme. This scheme shall continue to use EncryptedValue for backward compatibility reasons.

Replace the comment after the privateKey field of crc[1].certifiedKeyPair in the syntax of the Initialization Response message with the following text.

```
-- see Appendix C, Request Message Behavioral Clarifications
-- for backward compatibility reasons, use EncryptedValue
```

4. IANA Considerations

< TBD: Add any IANA considerations >

5. Security Considerations

No changes are made to the existing security considerations of RFC 4210 [RFC4210].

6. Acknowledgements

Special thank goes to Jim Schaad for his guidance and the inspiration on structuring and writing this document I got from [RFC6402] that updates CMC.

I also like to thank all reviewers of this document for their valuable feedback.

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC2985] Nystrom, M. and B. Kaliski, "PKCS #9: Selected Object Classes and Attribute Types Version 2.0", RFC 2985, DOI 10.17487/RFC2985, November 2000, <<https://www.rfc-editor.org/info/rfc2985>>.
- [RFC2986] Nystrom, M. and B. Kaliski, "PKCS #10: Certification Request Syntax Specification Version 1.7", RFC 2986, DOI 10.17487/RFC2986, November 2000, <<https://www.rfc-editor.org/info/rfc2986>>.
- [RFC4210] Adams, C., Farrell, S., Kause, T., and T. Mononen, "Internet X.509 Public Key Infrastructure Certificate Management Protocol (CMP)", RFC 4210, DOI 10.17487/RFC4210, September 2005, <<https://www.rfc-editor.org/info/rfc4210>>.
- [RFC4211] Schaad, J., "Internet X.509 Public Key Infrastructure Certificate Request Message Format (CRMF)", RFC 4211, DOI 10.17487/RFC4211, September 2005, <<https://www.rfc-editor.org/info/rfc4211>>.

- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <<https://www.rfc-editor.org/info/rfc5280>>.
- [RFC5652] Housley, R., "Cryptographic Message Syntax (CMS)", STD 70, RFC 5652, DOI 10.17487/RFC5652, September 2009, <<https://www.rfc-editor.org/info/rfc5652>>.
- [RFC6402] Schaad, J., "Certificate Management over CMS (CMC) Updates", RFC 6402, DOI 10.17487/RFC6402, November 2011, <<https://www.rfc-editor.org/info/rfc6402>>.

7.2. Informative References

- [I-D.brockhaus-lamps-lightweight-cmp-profile] Brockhaus, H., Fries, S., and D. Oheimb, "Lightweight CMP Profile", draft-brockhaus-lamps-lightweight-cmp-profile-02 (work in progress), December 2019.
- [RFC5958] Turner, S., "Asymmetric Key Packages", RFC 5958, DOI 10.17487/RFC5958, August 2010, <<https://www.rfc-editor.org/info/rfc5958>>.
- [RFC6032] Turner, S. and R. Housley, "Cryptographic Message Syntax (CMS) Encrypted Key Package Content Type", RFC 6032, DOI 10.17487/RFC6032, December 2010, <<https://www.rfc-editor.org/info/rfc6032>>.

Appendix A. ASN.1 Modules

Changes to the following parts are needed

- o Import from PKCS-9

friendlyName, localKeyId

```
FROM PKCS-9 {iso(1) member-body(2) us(840) rsadsi(113549) pkcs(1)
pkcs-9(9) modules(0) pkcs-9(1)}
```

< TBD: Either friendlyName or localKeyId need to be imported here. >

- o Import from PKIXCRMF-2005

```
CertTemplate, PKIPublicationInfo, EncryptedKey, CertId,
CertReqMessages
  FROM PKIXCRMF-2005 {iso(1) identified-organization(3)
    dod(6) internet(1) security(5) mechanisms(5) pkix(7)
    id-mod(0) id-mod-crmf2005(36)}
```

o In CertifiedKeyPair, CertOrEncCert and id-it-revPassphrase

```
CertifiedKeyPair ::= SEQUENCE {
  certOrEncCert      CertOrEncCert,
  privateKey          [0] EncryptedKey OPTIONAL,
  -- see [CRMF] for comment on encoding
  publicationInfo [1] PKIPublicationInfo OPTIONAL
}

CertOrEncCert ::= CHOICE {
  certificate      [0] CMPCertificate,
  encryptedCert    [1] EncryptedKey
}

-- id-it-revPassphrase OBJECT IDENTIFIER ::= {id-it 12}
-- RevPassphraseValue ::= EncryptedKey

--
-- Extended Key Usage extension for PKI entities used in
-- CMP operations
--

id-kp-cmpCA OBJECT IDENTIFIER ::= { id-kp 27 }
id-kp-cmpRA OBJECT IDENTIFIER ::= { id-kp 28 }
id-kp-cmpKGA OBJECT IDENTIFIER ::= { id-kp ... }

< TBD: id-kp-cmpKGA to be defined. >

< TBD: If needed the complete ASN.1 Module from RFC 4210 section
needs to be copied here. >
```

Author's Address

Hendrik Brockhaus
Siemens AG
Otto-Hahn-Ring 6
Munich 81739
Germany

Email: hendrik.brockhaus@siemens.com
URI: <http://www.siemens.com/>

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: July 30, 2020

H. Brockhaus
S. Fries
D. von Oheimb
Siemens
January 27, 2020

Lightweight CMP Profile
draft-brockhaus-lamps-lightweight-cmp-profile-03

Abstract

The goal of this document is to facilitate interoperability and automation by profiling the Certificate Management Protocol (CMP) version 2 and the related Certificate Request Message Format (CRMF) version 2 and the HTTP Transfer for the Certificate Management Protocol. It specifies a subset of CMP and CRMF focusing on typical uses cases relevant for managing certificates of devices in many industrial and IoT scenarios. To limit the overhead of certificate management for more constrained devices only the most crucial types of transactions are specified as mandatory. To foster interoperability also in more complex scenarios, other types of transactions are specified as recommended or optional.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on July 30, 2020.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

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

(<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. History of changes	3
2. Introduction	5
2.1. Motivation for profiling CMP	5
2.2. Motivation for a lightweight profile for CMP	6
2.3. Existing CMP profiles	7
2.4. Compatibility with existing CMP profiles	9
2.5. Scope of this document	10
2.6. Structure of this document	11
2.7. Convention and Terminology	11
3. Architecture and use cases	12
3.1. Solution architecture	12
3.2. Basic generic CMP message content	13
3.3. Supported use cases	14
3.3.1. Mandatory use cases	14
3.3.2. Recommended Use Cases	14
3.3.3. Optional use cases	15
3.4. CMP message transport	15
4. Generic parts of the PKI message	16
4.1. General description of the CMP message header	17
4.2. General description of the CMP message protection	18
4.3. General description of CMP message extraCerts	19
5. End Entity focused certificate management use cases	19
5.1. Requesting a new certificate from a PKI	20
5.1.1. A certificate from a new PKI with signature protection	21
5.1.2. A certificate from a trusted PKI with signature protection	27
5.1.3. Update an existing certificate with signature protection	27
5.1.4. A certificate from a PKI with MAC protection	28
5.1.5. A certificate from a legacy PKI using PKCS#10 request	30
5.1.6. Generate the key pair centrally at the (L)RA/CA	32
5.1.6.1. Using symmetric key-encryption key management technique	37
5.1.6.2. Using key agreement key management technique	38
5.1.6.3. Using key transport key management technique	39
5.1.7. Delayed enrollment	40
5.2. Revoking a certificate	45

5.3.	Error reporting	47
5.4.	Support messages	49
5.4.1.	General message and response	49
5.4.2.	Get CA certificates	51
5.4.3.	Get root CA certificate update	51
5.4.4.	Get certificate request parameters	53
5.4.5.	Get certificate management configuration	54
5.4.6.	Get enrollment voucher	56
6.	LRA and RA focused certificate management use cases	57
6.1.	Forwarding of messages	57
6.1.1.	Not changing protection	59
6.1.2.	Replacing protection	60
6.1.2.1.	Keeping proof-of-possession	60
6.1.2.2.	Breaking proof-of-possession	61
6.1.3.	Adding Protection	61
6.1.4.	Initiating delayed enrollment	61
6.2.	Revoking certificates on behalf of another's entities	61
6.3.	Error reporting	62
7.	CMP message transport variants	63
7.1.	HTTP transport	63
7.2.	HTTPS transport using certificates	65
7.3.	HTTPS transport using shared secrets	65
7.4.	File-based transport	66
7.5.	CoAP transport	66
7.6.	Piggybacking on other reliable transport	66
8.	IANA Considerations	66
9.	Security Considerations	66
10.	Acknowledgements	66
11.	References	67
11.1.	Normative References	67
11.2.	Informative References	68
Appendix A.	Additional Stuff	70
Authors' Addresses	70

1. History of changes

Note: This section will be deleted in the final version of the document.

From version 02 -> 03:

- o Added a short summary of [RFC4210] Appendix D and E in Section 2.3.
- o Clarified some references to different sections and added some clarification in response to feedback from Michael Richardson and Tomas Gustavsson.

- o Added an additional label to the operational path to address multiple CAs or certificate profiles in Section 7.1.

From version 01 -> 02:

- o Added some clarification on the key management techniques for protection of centrally generated keys in Section 5.1.6.
- o Added some clarifications on the certificates for root CA certificate update in Section 5.4.3.
- o Added a section to specify the usage of nested messages for RAs to add an additional protection for further discussion, see Section 6.1.3.
- o Added a table containing endpoints for HTTP transport in Section 7.1 to simplify addressing PKI management entities.
- o Added some Todos resulting from discussion with Tomas Gustavsson.
- o Minor clarifications and changes in wording.

From version 00 -> 01:

- o Added a section to specify the enrollment with a already trusted PKI for further discussion, see Section 5.1.2.
- o Complete specification of requesting a certificate from a legacy PKI using a PKCS#10 [RFC2986] request in Section 5.1.5.
- o Complete specification of adding central generation of a key pair on behalf of an end entity in Section 5.1.6.
- o Complete specification of handling delayed enrollment due to asynchronous message delivery in Section 5.1.7.
- o Complete specification of additional support messages, e.g., to update a Root CA certificate or to request an RFC 8366 [RFC8366] voucher, in Section 5.4.
- o Minor changes in wording.

From version draft-brockhaus-lamps-industrial-cmp-profile-00 -> brockhaus-lamps-lightweight-cmp-profile-00:

- o Change focus from industrial to more multi-purpose use cases and lightweight CMP profile.

- o Incorporate the omitted confirmation into the header specified in Section 4.1 and described in the standard enrollment use case in Section 5.1.1 due to discussion with Tomas Gustavsson.
- o Change from OPTIONAL to RECOMMENDED for use case 'Revoke another's entities certificate' in Section 6.2, because it is regarded as important functionality in many environments to enable the management station to revoke EE certificates.
- o Complete the specification of the revocation message flow in Section 5.2 and Section 6.2.
- o The CoAP based transport mechanism and piggybacking of CMP messages on top of other reliable transport protocols is out of scope of this document and would need to be specified in another document.
- o Further minor changes in wording.

2. Introduction

This document specifies PKI management operations supporting machine-to-machine and IoT use cases. The focus lies on maximum automation and interoperable implementation of all involved PKI entities from end entities (EE) through an optional Local Registration Authority (LRA) and the RA up to the CA. The profile makes use of the concepts and syntax specified in CMP [RFC4210], CRMF [RFC4211], HTTP transfer for CMP [RFC6712], and CMP Updates [I-D.brockhaus-lamps-cmp-updates]. Especially CMP and CRMF are very feature-rich standards, while only a limited subset of the specified functionality is needed in many environments. Additionally, the standards are not always precise enough on how to interpret and implement the described concepts. Therefore, we aim at tailoring and specifying in more detail how to use these concepts to implement lightweight automated certificate management.

2.1. Motivation for profiling CMP

CMP was standardized in 1999 and is implemented in several CA products. In 2005 a completely reworked and enhanced version 2 of CMP [RFC4210] and CRMF [RFC4211] has been published followed by a document specifying a transfer mechanism for CMP messages using http [RFC6712] in 2012.

Though CMP is a very solid and capable protocol it could be used more widely. The most important reason for not more intense application of CMP appears to be that the protocol is offering a large set of features and options but being not always precise enough and leaving

room for interpretation. On the one hand, this makes CMP applicable to a very wide range of scenarios, but on the other hand a full implementation of all options is unrealistic because this would take enormous effort.

Moreover, many details of the CMP protocol have been left open or have not been specified in full preciseness. The profiles specified in Appendix D and E of [RFC4210] offer some more detailed certificate use cases. But the specific needs of highly automated scenarios for a machine-to-machine communication are not covered sufficiently.

As also 3GPP and UNISG already put across, profiling is a way of coping with the challenges mentioned above. To profile means to take advantage of the strengths of the given protocol, while explicitly narrowing down the options it provides to exactly those needed for the purpose(s) at hand and eliminating all identified ambiguities. In this way all the general and applicable aspects of the protocol can be taken over and only the peculiarities of the target scenario need to be dealt with specifically.

Doing such a profiling for a new target environment can be a high effort because the range of available options needs to be well understood and the selected options need to be consistent with each other and with the intended usage scenario. Since most industrial use cases typically have much in common it is worth sharing this effort, which is the aim of this document. Other standardization bodies can then reference the profile from this document and do not need to come up with individual profiles.

2.2. Motivation for a lightweight profile for CMP

The profiles specified in Appendix D and E of CMP have been developed in particular to manage certificates of human end entities. With the evolution of distributed systems and client-server architectures, certificates for machines and applications on them have become widely used. This trend has strengthened even more in emerging industrial and IoT scenarios. CMP is sufficiently flexible to support these very well.

Today's IT security architectures for industrial solutions typically use certificates for endpoint authentication within protocols like IPSec, TLS, or SSH. Therefore, the security of these architectures highly relies upon the security and availability of the implemented certificate management procedures.

Due to increasing security in operational networks as well as availability requirements, especially on critical infrastructures and systems with a high volume of certificates, a state-of-the-art

certificate management must be constantly available and cost-efficient, which calls for high automation and reliability. The NIST Cyber Security Framework [NIST-CSFW] also refers to proper processes for issuance, management, verification, revocation, and audit for authorized devices, users and processes involving identity and credential management. Such PKI operation according to commonly accepted best practices is also required in IEC 62443-3-3 [IEC62443-3-3] for security level 2 up to security level 4.

Further challenges in many industrial systems are network segmentation and asynchronous communication, where PKI operation is often not deployed on-site but in a more protected environment of a data center or trust center. Certificate management must be able to cope with such network architectures. CMP offers the required flexibility and functionality, namely self-contained messages, efficient polling, and support for asynchronous message transfer with end-to-end security.

2.3. Existing CMP profiles

As already stated, CMP contains profiles with mandatory and optional transactions in the Appendixes D and E of [RFC4210]. Those profiles focus on management of human user certificates and do only partly address the specific needs for certificate management automation for unattended machine or application-oriented end entities.

[RFC4210] specifies in Appendix D the following mandatory PKI management operations (all require support of, in the meantime outdated, algorithms, e.g., SHA-1 and 3-DES; all operations may enroll up to two certificates, one for a locally generated and another optional one for a centrally generated key pair; all require use of certConf/PKIConf messages for confirmation):

- o Initial registration/certification; an (uninitialized) end entity requests a (first) certificate from a CA using shared secret based message authentication. The content is similar to PKI management operation specified in Section 5.1.4 of this document.
- o Certificate request; an (initialized) end entity requests a certificate from a CA (for any reason) using signature or shared secret based message authentication. The content is similar to PKI management operation specified in Section 5.1.2 of this document.
- o Key update; an (initialized) end entity requests a certificate from a CA (to update the key pair and/or corresponding certificate that it already possesses) using signature or shared secret based

message authentication. The content is similar to PKI management operation specified in Section 5.1.3 of this document.

Due to the two certificates that may be enrolled and the shared secret based authentication, these PKI management operations focuss more on the enrollment of a human users at a PKI.

[RFC4210] specifies in Appendix E the following optional transactions (all require support of, in the meantime outdated, algorithms, e.g., SHA-1 and 3-DES):

- o Root CA key update; a root CA updates its key pair and produces a CA key update announcement message that can be made available (via some transport mechanism) to the relevant end entities. This operation only supports a push and no pull model. The content is similar to PKI management operation specified in Section 5.4.3 of this document.
- o Information request/response; an end entity sends a general message to the PKI requesting details that will be required for later PKI management operations. The content is similar to PKI management operation specified in Section 5.4.4 and Section 5.4.5 of this document.
- o Cross-certification request/response (1-way); creation of a single cross-certificate (i.e., not two at once). The requesting CA MAY choose who is responsible for publication of the cross-certificate created by the responding CA through use of the PKIPublicationInfo control.
- o In-band initialization using external identity certificate (this PKI management operation may also enroll up to two certificates and requires use of certConf/PKIConf messages for confirmation as specified in Appendix D of [RFC4210]). An (uninitialized) end entity wishes to initialize into the PKI with a CA, CA-1. It uses, for authentication purposes, a pre-existing identity certificate issued by another (external) CA, CA-X. A trust relationship must already have been established between CA-1 and CA-X so that CA-1 can validate the EE identity certificate signed by CA-X. Furthermore, some mechanism must already have been established within the Personal Security Environment (PSE) of the EE that would allow it to authenticate and verify PKIMessages signed by CA-1. The content is similar to PKI management operation specified in Section 5.1.1 of this document. The trust establishment of the EE in CA-1 and of the CA/RA in CA-X can be automatized using, e.g., the exchange of a certificate management configuration as specified in Section 5.4.5 or an enrollment voucher as specified in Section 5.4.6 of this document.

Both Appendixes focus on EE to CA/RA PKI management operations and do not address further profiling of RA to CA communication as typically used for full backend automation.

3GPP makes use of CMP [RFC4210] in its Technical Specification 133 310 [ETSI-3GPP] for automatic management of IPsec certificates in UMTS, LTE, and 5G backbone networks. Since 2010 a dedicated CMP profile for initial certificate enrollment and update transactions between end entities and the RA/CA is specified in the document.

UNISIG has included a CMP profile for certificate enrollment in the subset 137 specifying the ETRAM/ECTS on-line key management for train control systems [UNISIG] in 2015.

Both standardization bodies use CMP [RFC4210], CRMF [RFC4211], and HTTP transfer for CMP [RFC6712] to add tailored means for automated certificate management for unattended machine or application-oriented end entities.

2.4. Compatibility with existing CMP profiles

The profile specified in this document is compatible with CMP [RFC4210] Appendixes D and E (PKI Management Message Profiles), with the following exceptions:

- o signature-based protection is the default protection; initial transactions may also use HMAC,
- o certification of a second key pair within the same transaction is not supported,
- o proof-of-possession (POPO) with self-signature of the certTemplate according to [RFC4211] section 4.1 clause 3 is the recommended default POPO method (deviations are possible by EEs when requesting central key generation and by (L)RAs when using raVerified),
- o confirmation of newly enrolled certificates may be omitted, and
- o all transactions consist of request-response message pairs originating at the EE, i.e., announcement messages are omitted.

The profile specified in this document is compatible with the CMP profile for UMTS, LTE, and 5G network domain security and authentication framework [ETSI-3GPP], except that:

- o protection of initial transactions may be HMAC-based,

- o the subject name is mandatory in certificate templates, and
- o confirmation of newly enrolled certificates may be omitted.

The profile specified in this document is compatible with the CMP profile for on-line key management in rail networks as specified in UNISIG subset-137 [UNISIG], except that:

- o as of RFC 4210 [RFC4210] the messageTime is required to be Greenwich Mean Time coded as generalizedTime (Note: While UNISIG explicitly states that the messageTime is required to be 'UTC time', it is not clear if this means a coding as UTCTime or generalizedTime and if other time zones than Greenwich Mean Time shall be allowed. Therefore UNISIG may be in conflict with RFC 4210 [RFC4210]. Both time formats are described in RFC 5280 [RFC5280] section 4.1.2.5.), and
- o in case the request message is MAC protected, also the response, certConf, and PKIconf messages have a MAC-based protection (Note: if changing to signature protection of the response the caPubs field cannot be used securely anymore.).

2.5. Scope of this document

This document specifies requirements on generating messages on the sender side. It does not specify strictness of verification on the receiving side and how in detail to handle error cases.

Especially on the EE side this profile aims at a lightweight protocol that can be implemented on more constrained devices. On the side of the central PKI management entities the profile accepts higher resource needed.

For the sake of robustness and preservation of security properties implementations should, as far as security is not affected, adhere to Postel's law: "Be conservative in what you do, be liberal in what you accept from others" (often reworded as: "Be conservative in what you send, be liberal in what you accept").

When in Section 4, Section 5, and Section 6 a field of the ASN.1 syntax as defined in RFC 4210 [RFC4210] and RFC 4211 [RFC4211] is not explicitly specified, it SHOULD not be used by the sending entity. The receiving entity MUST NOT require its absence and if present MUST gracefully handle its presence.

2.6. Structure of this document

Section 3 introduces the general PKI architecture and approach to certificate management using CMP that is assumed in this document. Then it enlists the PKI management operations specified in this document and describes them in general words. The list of supported certificate management use cases is divided into mandatory, recommended, and optional ones.

Section 4 profiles the CMP message header, protection, and extraCerts section as they are general elements of CMP messages.

Section 5 profiles the exchange of CMP messages between an EE and the first PKI management entities. There are various flavors of certificate enrollment requests optionally with polling, revocation, error handling, and general support transactions.

Section 6 profiles the exchange between PKI management entities. These are in the first place the forwarding of messages coming from or going to an EE. This includes also initiating delayed delivery of messages, which involves polling. Additionally, it specifies transactions where the PKI component manages certificates on behalf of an EE or for itself.

Section 7 outlines different mechanisms for CMP message transfer, namely http-based transfer as already specified in [RFC6712], using an additional TLS layer, or offline file-based transport. CoAP [RFC7252] and piggybacking CMP messages on other protocols is out of scope and left for further documents.

2.7. Convention and 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 RFC 2119 [RFC2119].

In this document, these words will appear with that interpretation only when in ALL CAPS. Lower case uses of these words are not to be interpreted as carrying significance described in RFC 2119.

Technical terminology is used in conformance with RFC 4210 [RFC4210], RFC 4211 [RFC4211], RFC 5280 [RFC5280], and IEEE 802.1AR [IEEE802.1AR]. The following key words are used:

CA: Certification authority, which issues certificates.

RA: Registration authority, an optional system component to which a CA delegates certificate management functions such as authorization checks.

LRA: Local registration authority, an optional RA system component with proximity to the end entities.

KGA: Key generation authority, an optional system component, typically co-located with an LRA, RA, or CA, that offers key generation services to end entities.

EE: End entity, a user, device, or service that holds a PKI certificate. An identifier for the EE is given as the subject of its certificate.

3. Architecture and use cases

3.1. Solution architecture

Typically, a machine EE will be equipped with a manufacturer issued certificate during production. Such a manufacturer issued certificate is installed during production to identify the device throughout its lifetime. This manufacturer certificate can be used to protect the initial enrollment of operational certificates after installation of the EE in a plant or industrial network. An operational certificate is issued by the owner or operator of the device to identify the device during operation, e.g., within a security protocol like IPsec, TLS, or SSH. In IEEE 802.1AR [IEEE802.1AR] a manufacturer certificate is called IDevID certificate and an operational certificate is called LDevID certificate.

All certificate management transactions are initiated by the EE. The EE creates a CMP request message, protects it using its manufacturer or operational certificate, if available, and sends it to its locally reachable PKI component. This PKI component may be an LRA, RA, or the CA, which checks the request, responds to it itself, or forwards the request upstream to the next PKI component. In case an (L)RA changes the CMP request message header or body or wants to prove a successful verification or authorization, it can apply a protection of its own. Especially the communication between an LRA and RA can be performed synchronously or asynchronously. Synchronous communication describes a timely uninterrupted communication between two communication partners, as asynchronous communication is not performed in a timely consistent manner, e.g., because of a delayed message delivery.




Figure 1: Certificate management on site

In operation environments a layered LRA-RA-CA architecture can be deployed, e.g., with LRAs bundling requests from multiple EEs at dedicated locations and one (or more than one) central RA aggregating the requests from multiple LRAs. Every (L)RA in this scenario will have its own dedicated certificate containing an extended key usage as specified in CMP Updates [I-D.brockhaus-lamps-cmp-updates] and private key allowing it to protect CMP messages it processes (CMP signing key/certificate). The figure above shows an architecture using one LRA and one RA. It is also possible to have only an RA or multiple LRAs and/or RAs. Depending on the network infrastructure, the communication between different PKI components may be synchronous online-communication, delayed asynchronous communication, or even offline file transfer.

As this profile focusses on specifying the pull model, where the EE always requests a specific PKI management operation. CMP response messages, especially in case of central key generation, as described in Section 5.1.6, can also be used to deliver proactively to the EE to implement the push model.

Third-party CAs typically implement different variants of CMP or even use proprietary interfaces for certificate management. Therefore, the LRA or the RA may need to adapt the exchanged CMP messages to the flavor of communication required by the CA.

3.2. Basic generic CMP message content

Section 4 specifies the generic parts of the CMP messages as used later in Section 5 and Section 6.

- o Header of a CMP message; see Section 4.1.
- o Protection of a CMP message; see Section 4.2.

- o ExtraCerts field of a CMP message; see Section 4.3.

3.3. Supported use cases

Following the outlined scope from Section 2.5, this section gives a brief overview of the certificate management use cases specified in Section 5 and Section 6 and points out, if an implementation by compliant EE or PKI component is mandatory, recommended or optional.

3.3.1. Mandatory use cases

The mandatory uses case in this document shall limit the overhead of certificate management for more constrained devices to the most crucial types of transactions.

Section 5 – End Entity focused certificate management use cases

- o Request a certificate from a new PKI with signature protection; see Section 5.1.1.
- o Request to update an existing certificate with signature protection; see Section 5.1.3.
- o Error reporting; see Section 5.3.

Section 6 – LRA and RA focused certificate management use cases

- o Forward messages without changes; see Section 6.1.1.
- o Forward messages with replaced protection and raVerified as proof-of-possession; see Section 6.1.2.2.
- o Error reporting; see Section 6.3.

3.3.2. Recommended Use Cases

Additional recommended use cases shall support some more complex scenarios, that are considered as beneficial for environments with more specific boundary conditions.

Section 5 – End Entity focused certificate management use cases

- o Request a certificate from a PKI with MAC protection; see Section 5.1.4.
- o Handle delayed enrollment due to asynchronous message delivery; see Section 5.1.7.

< TBD: There still some discussion ongoing if this should be recommended or optional. >

- o Revoke an own certificate.

Section 6 – LRA and RA focused certificate management use cases

- o Revoke another's entities certificate.

3.3.3. Optional use cases

The optional use cases support specific requirements seen only in a subset of environments.

Section 5 – End Entity focused certificate management use cases

- o Request a certificate from a legacy PKI using a PKCS#10 [RFC2986] request; see Section 5.1.5.
- o Add central generation of a key pair to a certificate request; see Section 5.1.6. If central key generation is supported, the key agreement key management technique is REQUIRED to be supported, and the key transport and symmetric key-encryption key management techniques are OPTIONAL.
- o Additional support messages, e.g., to update a Root CA certificate or to request an RFC 8366 [RFC8366] voucher; see Section 5.4.

Section 6 – LRA and RA focused certificate management use cases

- o Initiate delayed enrollment due to asynchronous message delivery; see Section 6.1.4.

3.4. CMP message transport

On different links between PKI entities, e.g., EE<->RA and RA<->CA, different transport MAY be used. As CMP has only very limited requirement regarding the mechanisms used for message transport and in different environments different transport mechanisms are supported, e.g. HTTP, CoAP, or even offline files based, this document requires no specific transport protocol to be supported by all conforming implementations.

HTTP transfer is RESOMMENDED to use for all PKI entities, but there is no transport specified as mandatory to be flexible for devices with special constrains to choose whatever transport is suitable.

Recommended transport

- o Transfer CMP messages using HTTP; see Section 7.1.

Optional transport

- o Transfer CMP messages using HTTPS with certificate-based authentication; see Section 7.2.
- o Transfer CMP messages using HTTPS with shared-secret based protection; see Section 7.3.
- o File-based CMP message transport.

< TBD: Motivation see Section 7.4 >

< TBD: Michael Richardson proposed to also specify a CoAP based message transport profile. If there is further support for this profile and someone volunteering to provide the necessary input for this section, I would add it to the document. >

4. Generic parts of the PKI message

To reduce redundancy in the text and to ease implementation, the contents of the header, protection, and extraCerts fields of the CMP messages used in the transactions specified in Section 5 and Section 6 are standardized to the maximum extent possible. Therefore, the generic parts of a CMP message are described centrally in this section.

As described in section 5.1 of [RFC4210], all CMP messages have the following general structure:

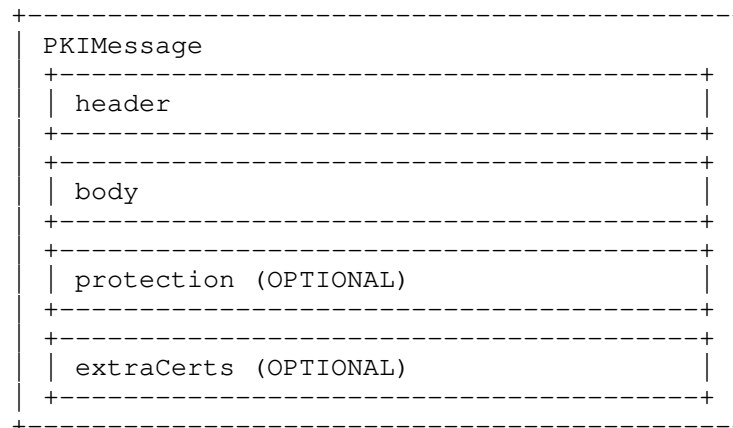


Figure 2: CMP message structure

The general contents of the message header, protection, and extraCerts fields are specified in the Section 4.1 to Section 4.3.

In case a specific CMP message needs different contents in the header, protection, or extraCerts fields, the differences are described in the respective message.

The CMP message body contains the message-specific information. It is described in the context of Section 5 and Section 6.

The behavior in case an error occurs while handling a CMP message is described in Section 6.3.

4.1. General description of the CMP message header

This section describes the generic header field of all CMP messages with signature-based protection. The only variations described here are in the fields recipient, transactionID, and recipNonce of the first message of a transaction.

In case a message has MAC-based protection the changes are described in the respective section. The variations will affect the fields sender, protectionAlg, and senderKID.

For requirements about proper random number generation please refer to [RFC4086]. Any message-specific fields or variations are described in the respective sections of this chapter.

header

```

  pvno                                REQUIRED
    -- MUST be set to 2 to indicate CMP V2
  sender                              REQUIRED
    -- MUST be the subject of the signing certificate used for
    -- protection of this message
  recipient                            REQUIRED
    -- MUST be the name of the intended recipient
    -- If this is the first message of a transaction: SHOULD be the
    -- subject of the issuing CA certificate
    -- In all other messages: SHOULD be the same name as in the
    -- sender field of the previous message in this transaction
  messageTime                          RECOMMENDED
    -- MUST be the time at which the message was produced, if
    -- present
  protectionAlg                        REQUIRED
    -- MUST be the algorithm identifier of the signature or algorithm
    -- id-PasswordBasedMac algorithm used for calculation of the
    -- protection bits
    -- The signature algorithm MUST be consistent with the
```



```

-- SubjectPublicKeyInfo field of the signer's certificate
-- The hash algorithm used SHOULD be SHA-256
algorithm                REQUIRED
-- MUST be the OID of the signature algorithm, like
-- sha256WithRSAEncryption or ecdsa-with-SHA256, or
-- id-PasswordBasedMac
senderKID                RECOMMENDED
-- MUST be the SubjectKeyIdentifier, if available, of the
-- certificate used for protecting this message
transactionID           REQUIRED
-- If this is the first message of a transaction:
-- MUST be 128 bits of random data for the start of a
-- transaction to reduce the probability of having the
-- transactionID already in use at the server
-- In all other messages:
-- MUST be the value from the previous message in the same
-- transaction
senderNonce              REQUIRED
-- MUST be fresh 128 random bits
recipNonce               RECOMMENDED
-- If this is the first message of a transaction: SHOULD be
-- absent
-- In all other messages: MUST be present and contain the value
-- from senderNonce of the previous message in the same
-- transaction
generalInfo              OPTIONAL
  implicitConfirm         OPTIONAL
    ImplicitConfirmValue  REQUIRED
-- The field is optional though it only applies to
-- ir/cr/kur/pl0cr requests and ip/cp/kup responses
-- ImplicitConfirmValue of the request message MUST be NULL if
-- the EE wants to request not to send a confirmation message
-- ImplicitConfirmValue MUST be set to NULL if the (L)RA/CA wants
-- to grant not sending a confirmation message

```

4.2. General description of the CMP message protection

This section describes the generic protection field of all CMP messages with signature-based protection.

```

protection                REQUIRED
-- MUST contain the signature calculated using the signature
-- algorithm specified in protectionAlg

```

Only for MAC-based protection major differences apply as described in the respective message.

The CMP message protection provides, if available, message origin authentication and integrity protection for the CMP message header and body. The CMP message extraCerts is not covered by this protection.

NOTE: The requirements for checking certificates given in [RFC5280] MUST be followed for the CMP message protection. In case the CMP signer certificate is not the CA certificate that signed the newly issued certificate, certificate status checking SHOULD be used for the CMP signer certificates of communication partners.

4.3. General description of CMP message extraCerts

This section describes the generic extraCerts field of all CMP messages with signature-based protection.

```
extraCerts                                RECOMMENDED
-- SHOULD contain the signing certificate together with its
-- chain, if needed
-- If present, the first certificate in this field MUST
-- be the certificate used for signing this message
-- Self-signed certificates SHOULD NOT be included in
-- extraCerts and MUST NOT be trusted based on the listing in
-- extraCerts in any case
```

5. End Entity focused certificate management use cases

This chapter focuses on the communication of the EE and the first PKI component it talks to. Depending on the network and PKI solution, this will either be the LRA, the RA or the CA.

Profiles of the Certificate Management Protocol (CMP) [RFC4210] handled in this chapter cover the following certificate management use cases:

- o Requesting a certificate from a PKI with variations like initial requests and updating, central key generation and different protection means
- o Revocation of a certificate
- o General messages for further support functions

The use cases mainly specify the message body of the CMP messages and utilize the specification of the message header, protection and extraCerts as specified in Section 5.

The behavior in case an error occurs is described in Section 5.3.

This chapter is aligned to Appendix D and Appendix E of [RFC4210]. The general rules for interpretation stated in Appendix D.1 in [RFC4210] need to be applied here, too.

This document does not mandate any specific supported algorithms like Appendix D.2 of [RFC4210], [ETSI-3GPP], and [UNISIG] do. Using the message sequences described here require agreement upon the algorithms to support and thus the algorithm identifiers for the specific target environment.

5.1. Requesting a new certificate from a PKI

There are different approaches to request a certificate from a PKI.

These approaches differ on the one hand in the way the EE can authenticate itself to the PKI it wishes to get a new certificate from and on the other hand in its capabilities to generate a proper new key pair. The authentication means may be as follows:

- o Using a certificate from a trusted PKI and the corresponding private key, e.g., a manufacturer certificate
- o Using the certificate to be updated and the corresponding private key
- o Using a shared secret known to the EE and the PKI

Typically, such EE requests a certificate from a CA. When the (L)RA/CA responds with a message containing a certificate, the EE MUST reply with a confirmation message. The (L)RA/CA then MUST send confirmation back, closing the transaction.

The message sequences in this section allow the EE to request certification of a locally generated public-private key pair. For requirements about proper random number and key generation please refer to [RFC4086]. The EE MUST provide a signature-based proof-of-possession of the private key associated with the public key contained in the certificate request as defined by [RFC4211] section 4.1 case 3. To this end it is assumed that the private key can technically be used as signing key. The most commonly used algorithms are RSA and ECDSA, which can technically be used for signature calculation regardless of potentially intended restrictions of the key usage.

The requesting EE provides the binding of the proof-of-possession to its identity by signature-based or MAC-based protection of the CMP request message containing that POPO. The (L)RA/CA needs to verify whether this EE is authorized to obtain a certificate with the

requested subject and other attributes and extensions. Especially when removing the protection provided by the EE and applying a new protection the (L)RA MUST verify in particular the included proof-of-possession self-signature of the certTemplate using the public key of the requested certificate and MUST check that the EE, as authenticated by the message protection, is authorized to request a certificate with the subject as specified in the certTemplate (see Section 6.1.2).

There are several ways to install the Root CA certificate of a new PKI on an EE. The installation can be performed in an out-of-band manner, using general messages, a voucher [RFC8366], or other formats for enrolment, or in-band of CMP by the caPubs field in the certificate response message. In case the installation of the new Root CA certificate is performed using the caPubs field, the certificate response message MUST be properly authenticated, and the sender of this message MUST be authorized to install new Root CA certificates on the EE. This authorization MUST be indicated by the extended key usage in the (L)RA/CA certificate as specified in CMP Updates [I-D.brockhaus-lamps-cmp-updates].

5.1.1.1. A certificate from a new PKI with signature protection

This message sequence should be used by an EE to request a certificate of a new PKI using an existing certificate from an external PKI, e.g., a manufacturer certificate, to prove its identity to the new PKI. The EE already has established trust in this new PKI it is about to enroll to, e.g., by configuration means. The initialization request message is signature-protected using the existing certificate.

Preconditions:

- 1 The EE MUST have a certificate enrolled by an external PKI in advance to this transaction to authenticate itself to the (L)RA/CA using signature-based protection, e.g., using a manufacturer certificate.
- 2 The EE SHOULD know the subject name of the new CA it requests a certificate from; this name MAY be established using an enrollment voucher or other configuration means. If the EE does not know the name of the CA, the (L)RA/CA MUST know where to route this request to.
- 3 The EE MUST authenticate responses from the (L)RA/CA; trust MAY be established using an enrollment voucher or other configuration means

- 4 The (L)RA/CA MUST trust the external PKI the EE uses to authenticate itself; trust MAY be established using some configuration means

This message sequence is like that given in [RFC4210] Appendix E.7.

Message flow:

Step#	EE		(L)RA/CA
1	format ir		
2		-> ir	->
3			handle, re-protect or forward ir
4			format or receive ip
5			possibly grant implicit confirm
6		<- ip	<-
7	handle ip		
8			In case of status "rejection" in the ip message, no certConf and pkiConf are sent
9	format certConf (optional)		
10		-> certConf	->
11			handle, re-protect or forward certConf
12			format or receive PKIConf
13		<- pkiConf	<-
14	handle pkiConf (optional)		

For this message sequence the EE MUST include exactly one single CertReqMsg in the ir. If more certificates are required, further requests MUST be sent using separate CMP Messages. If the EE wants to omit sending a certificate confirmation message after receiving the ip to reduce the number of protocol messages exchanged in a transaction, it MUST request this by setting the implicitControlValue in the ir to NULL.

If the CA accepts the request it MUST return the new certificate in the certifiedKeyPair field of the ip message. If the EE requested to omit sending a certConf message after receiving the ip, the (L)RA/CA MAY confirm this by also setting the implicitControlValue in the ip to NULL.

If the EE did not request implicit confirmation or the request was not granted by the (L)RA/CA the confirmation as follows MUST be performed. If the EE successfully receives the certificate and accepts it, the EE MUST send a certConf message, which MUST be

answered by the (L)RA/CA with a pkiConf message. If the (L)RA/CA does not receive the expected certConf message in time it MUST handle this like a rejection by the EE.

If the certificate request was refused by the CA, the (L)RA/CA must return an ip message containing the status code "rejection" and no certifiedKeyPair field. Such an ip message MUST NOT be followed by the certConf and pkiConf messages.

Detailed message description:

Certification Request -- ir

Field	Value
-------	-------

header

-- As described in section 4.1

body

-- The request of the EE for a new certificate

ir	REQUIRED
----	----------

-- MUST be exactly one CertReqMsg

-- If more certificates are required, further requests MUST be

-- packaged in separate PKI Messages

certReq	REQUIRED
---------	----------

certReqId	REQUIRED
-----------	----------

-- MUST be set to 0

certTemplate	REQUIRED
--------------	----------

version	OPTIONAL
---------	----------

-- MUST be 2 if supplied.

subject	REQUIRED
---------	----------

-- MUST contain the suggested subject name of the EE

-- certificate

publicKey	REQUIRED
-----------	----------

algorithm	REQUIRED
-----------	----------

-- MUST include the subject public key algorithm ID and value

-- In case a central key generation is requested, this field

-- contains the algorithm and parameter preferences of the

-- requesting entity regarding the to-be-generated key pair

subjectPublicKey	REQUIRED
------------------	----------

-- MUST contain the public key to be included into the requested

-- certificate in case of local key-generation

-- MUST contain a zero-length BIT STRING in case a central key

-- generation is requested

-- MUST include the subject public key algorithm ID and value

extensions	OPTIONAL
------------	----------

-- MAY include end-entity-specific X.509 extensions of the

-- requested certificate like subject alternative name,

```

-- key usage, and extended key usage
Popo                                REQUIRED
  POPOSigningKey                    OPTIONAL
-- MUST be used in case subjectPublicKey contains a public key
-- MUST be absent in case subjectPublicKey contains a
-- zero-length BIT STRING
  POPOSigningKey                    REQUIRED
  poposkInput                        PROHIBITED
-- MUST NOT be used because subject and publicKey are both
-- present in the certTemplate
  algorithmIdentifier               REQUIRED
-- The signature algorithm MUST be consistent with the
-- publicKey field of the certTemplate
-- The hash algorithm used SHOULD be SHA-256
  signature                         REQUIRED
-- MUST be the signature computed over the DER-encoded
-- certTemplate

protection                          REQUIRED
  -- As described in section 4.2

extraCerts                          REQUIRED
  -- As described in section 4.3

Certification Response -- ip

Field                               Value

header
  -- As described in section 4.1

body
  -- The response of the CA to the request as appropriate
  ip                                REQUIRED
  caPubs                            OPTIONAL
  -- MAY be used
  -- If used it MUST contain only the root certificate of the
  -- certificate contained in certOrEncCert
  response                          REQUIRED
  -- MUST be exactly one CertResponse
  certReqId                         REQUIRED
  -- MUST be set to 0
  status                            REQUIRED
  -- PKIStatusInfo structure MUST be present
  status                            REQUIRED
  -- positive values allowed: "accepted", "grantedWithMods"
  -- negative values allowed: "rejection"

```

```

-- In case of rejection no certConf and pkiConf messages will
-- be sent
    statusString          OPTIONAL
-- MAY be any human-readable text for debugging, logging or to
-- display in a GUI
    failInfo              OPTIONAL
-- MUST be present if status is "rejection" and in this case
-- the transaction MUST be terminated
-- MUST be absent if the status is "accepted" or
-- "grantedWithMods"
    certifiedKeyPair      OPTIONAL
-- MUST be present if status is "accepted" or "grantedWithMods"
-- MUST be absent if status is "rejection"
    certOrEncCert         REQUIRED
-- MUST be present when certifiedKeyPair is present
    certificate            REQUIRED
-- MUST be present when certifiedKeyPair is present
-- MUST contain the newly enrolled X.509 certificate
    privateKey            OPTIONAL
-- MUST be absent in case of local key-generation
-- MUST contain the encrypted private key in an EnvelopedData
-- structure as specified in section 5.1.5 in case the private
-- key was generated centrally

protection                REQUIRED
-- As described in section 4.2

extraCerts                REQUIRED
-- As described in section 4.3
-- MUST contain the chain of the issued certificate
-- Duplicate certificates MAY be omitted

Certificate Confirmation -- certConf

Field                    Value

header
-- As described in section 4.1

body
-- The message of the EE sends confirmation to the (L)RA/CA
-- to accept or reject the issued certificates
certConf                REQUIRED
-- MUST be exactly one CertStatus
CertStatus               REQUIRED
    certHash              REQUIRED
-- MUST be the hash of the certificate, using the same hash

```



```

-- algorithm as used to create the certificate signature
certReqId          REQUIRED
-- MUST be set to 0
status             RECOMMENDED
-- PKIStatusInfo structure SHOULD be present
-- Omission indicates acceptance of the indicated certificate
status             REQUIRED
-- positive values allowed: "accepted"
-- negative values allowed: "rejection"
statusString       OPTIONAL
-- MAY be any human-readable text for debugging or logging
failInfo           OPTIONAL
-- MUST be present if status is "rejection"
-- MUST be absent if the status is "accepted"

protection          REQUIRED
-- As described in section 4.2
-- MUST use the same certificate as for protection of the ir

extraCerts          RECOMMENDED
-- SHOULD contain the protection certificate together with its
-- chain
-- If present, the first certificate in this field MUST be the
-- certificate used for signing this message
-- Self-signed certificates SHOULD NOT be included in
-- extraCerts and
-- MUST NOT be trusted based on the listing in extraCerts in
-- any case

PKI Confirmation -- pkiConf

Field                Value

header
-- As described in section 4.1

body
  pkiConf             REQUIRED
-- The content of this field MUST be NULL

protection          REQUIRED
-- As described in section 4.2
-- SHOULD use the same certificate as for protection of the ip

extraCerts          RECOMMENDED
-- SHOULD contain the protection certificate together with its
-- chain

```

- If present, the first certificate in this field MUST be the
- certificate used for signing this message
- Self-signed certificates SHOULD NOT be included in extraCerts
- and
- MUST NOT be trusted based on the listing in extraCerts in
- any case

5.1.2. A certificate from a trusted PKI with signature protection

< TBD: In case the PKI is already trusted the cr/cp messages could be used instead of ir/ip. It needs to be decided, whether an additional section should be added here, or the previous section should be extended to also cover this use case. >

5.1.3. Update an existing certificate with signature protection

This message sequence should be used by an EE to request an update of one of the certificates it already has and that is still valid. The EE uses the certificate it wishes to update to prove its identity and possession of the private key for the certificate to be updated to the PKI. Therefore, the key update request message is signed using the certificate that is to be updated.

The general message flow for this message sequence is the same as given in Section 5.1.1.

Preconditions:

- 1 The certificate the EE wishes to update MUST NOT be expired or revoked.
- 2 A new public-private key pair SHOULD be used.

The message sequence for this exchange is like that given in [RFC4210] Appendix D.6.

The message sequence for this exchange is identical to that given in Section 5.1.1, with the following changes:

- 1 The body of the first request and response MUST be kur and kup, respectively.
- 2 Protection of the kur MUST be performed using the certificate to be updated.
- 3 The subject field of the CertTemplate MUST contain the subject name of the existing certificate to be updated, without modifications.

- 4 The CertTemplate MUST contain the subject, issuer and publicKey fields only.
- 5 The regCtrl OldCertId SHOULD be used to make clear, even in case an (L)RA changes the message protection, which certificate is to be.
- 6 The caPubs field in the kup message MUST be absent.

As part of the certReq structure of the kur the control is added right after the certTemplate.

```
controls
  type                                RECOMMENDED
  -- MUST be the value id-regCtrl-oldCertID, if present
  value
    issuer                            REQUIRED
    serialNumber                      REQUIRED
  -- MUST contain the issuer and serialNumber of the certificate
  -- to be updated
```

5.1.4. A certificate from a PKI with MAC protection

This message sequence should be used by an EE to request a certificate of a new PKI without having a certificate to prove its identity to the target PKI, but there is a shared secret established between the EE and the PKI. Therefore, the initialization request is MAC-protected using this shared secret. The (L)RA checking the MAC-protection SHOULD replace this protection according to Section 6.1.2 in case the next hop does not know the shared secret.

For requirements with regard to proper random number and key generation please refer to [RFC4086].

The general message flow for this message sequence is the same as given in Section 5.1.1.

Preconditions:

- 1 The EE and the (L)RA/CA MUST share a symmetric key, this MAY be established by a service technician during initial local configuration.
- 2 The EE SHOULD know the subject name of the new CA it requests a certificate from; this name MAY be established using an enrollment voucher or other configuration means. If the EE does not know the name of the CA, the (L)RA/CA MUST know where to route this request to.

- 3 The EE MUST authenticate responses from the (L)RA/CA; trust MAY be established using the shared symmetric key.

The message sequence for this exchange is like that given in [RFC4210] Appendix D.4.

The message sequence for this exchange is identical to that given in Section 5.1.1, with the following changes:

- 1 The protection of all messages MUST be calculated using Message Authentication Code (MAC); the protectionAlg field MUST be id-PasswordBasedMac as described in section 5.1.3.1 of [RFC4210].
- 2 The sender MUST contain a name representing the originator of the message. The senderKID MUST contain a reference all participating entities can use to identify the symmetric key used for the protection.
- 3 The extraCerts of the ir, certConf, and PKIConf messages MUST be absent.
- 4 The extraCerts of the ip message MUST contain the chain of the issued certificate and root certificates SHOULD not be included and MUST NOT be trusted in any case.

Part of the protectionAlg structure, where the algorithm identifier MUST be id-PasswordBasedMac, is a PBMPParameter sequence. The fields of PBMPParameter SHOULD remain constant for message protection throughout this certificate management transaction to reduce the computational overhead.

```

PBMPParameter          REQUIRED
  salt                  REQUIRED
  -- MUST be the random value to salt the secret key
  owf                   REQUIRED
  -- MUST be the algorithm identifier for the one-way function
  -- used
  -- The one-way function SHA-1 MUST be supported due to
  -- [RFC4211] requirements, but SHOULD NOT be used any more
  -- SHA-256 SHOULD be used instead
  iterationCount        REQUIRED
  -- MUST be a limited number of times the OWF is applied
  -- To prevent brute force and dictionary attacks a reasonable
  -- high number SHOULD be used
  mac                   REQUIRED
  -- MUST be the algorithm identifier of the MAC algorithm used
  -- The MAC function HMAC-SHA1 MUST be supported due to
  -- [RFC4211] requirements, but SHOULD NOT be used any more
  -- HMAC-SHA-256 SHOULD be used instead

```

< TBD: SHA-1 is no collision resistant hash algorithm. Due to this fact the usage of SHA-1 has significantly decreased. Currently HMAC-SHA-1 seems relatively secure, it is currently recommended by cryptographers to also depreciate the uses of SHA-1 in the context of HMAC calculation. Should we depreciate the support of SHA-1 here completely? >

5.1.5. A certificate from a legacy PKI using PKCS#10 request

This message sequence should be used by an EE to request a certificate of a legacy PKI only capable to process PKCS#10 [RFC2986] certification requests. The EE can prove its identity to the target PKI by using various protection means as described in Section 5.1.1 or Section 5.1.4.

In contrast to the other transactions described in Section 5.1, this transaction uses PKCS#10 [RFC2986] instead of CRMF [RFC4211] for the certificate request for compatibility reasons with legacy CA systems that require a PKCS#10 certificate request and cannot process CMP [RFC4210] or CRMF [RFC4211] messages. In such case the (L)RA must extract the PKCS#10 certificate request from the p10cr and provides it separately to the CA.

The general message flow for this message sequence is the same as given in Section 5.1.1, but the public key is contained in the subjectPKInfo of the PKCS#10 certificate request.

Preconditions:

- 1 The EE MUST either have a certificate enrolled from this or any other accepted PKI, or a shared secret known to the PKI and the EE to authenticate itself to the (L)RA/CA.
- 2 The EE SHOULD know the subject name of the CA it requests a certificate from; this name MAY be established using an enrollment voucher or other configuration means. If the EE does not know the name of the CA, the (L)RA/CA MUST know where to route this request to.
- 3 The EE MUST authenticate responses from the (L)RA/CA; trust MAY be established by an available root certificate, using an enrollment voucher, or other configuration means.
- 4 The (L)RA/CA MUST trust the current or the PKI the EE uses to authenticate itself; trust MAY be established by a corresponding available root certificate or using some configuration means.

The profile for this exchange is identical to that given in Section 5.1.1, with the following changes:

- 1 The body of the first request and response MUST be p10cr and cp, respectively.
- 2 The subject name of the CA MUST be in the recipient field of the p10cr message header.
- 3 The certReqId in the cp message MUST be 0.
- 4 The caPubs field in the cp message SHOULD be absent.

Detailed description of the p10cr message:

Certification Request -- p10cr

Field	Value
header	-- As described in section 4.1
body	-- The request of the EE for a new certificate using a PKCS#10 -- certificate request
p10cr	REQUIRED
CertificationRequestInfo	REQUIRED
version	REQUIRED
	-- MUST be set to 0 to indicate PKCS#10 V1.7
subject	REQUIRED
	-- MUST contain the suggested subject name of the EE
subjectPKInfo	REQUIRED
	-- MUST include the subject public key algorithm ID and value
attributes	OPTIONAL
	-- MAY contain a set of end-entity-specific attributes or X.509
	-- extensions to be included in the requested certificate or used
	-- otherwise
signatureAlgorithm	REQUIRED
	-- The signature algorithm MUST be consistent with the
	-- subjectPKInfo field. The hash algorithm used SHOULD be SHA-256
signature	REQUIRED
	-- MUST containing the self-signature for proof-of-possession
protection	REQUIRED
	-- As described in section 4.2
extraCerts	REQUIRED
	-- As described in section 4.3

5.1.6. Generate the key pair centrally at the (L)RA/CA

This functional extension can be applied in combination with certificate enrollment as described in Section 5.1.1 and Section 5.1.4. The functional extension can be used in case an EE is not able or is not willing to generate its own new public-private key pair itself. It is a matter of the local implementation which central PKI components will perform the key generation. This component must have a proper (L)RA/CA certificate containing the additional extended key usage id-kp-cmcKGA to be identified by the EE as a legitimate key-generation instance. In case the (L)RA generated

the new key pair for the EE, it can use Section 5.1.1 to Section 5.1.4 to request the certificate for this key pair as usual.

Generally speaking, in a machine-to-machine scenario it is strongly preferable to generate public-private key pairs locally at the EE. Together with proof-of-possession of the private key in the certification request, this is to make sure that only the entity identified in the newly issued certificate is the only entity who ever hold the private key.

There are some cases where an EE is not able or not willing to locally generate the new key pair. Reasons for this may be the following:

- o Lack of sufficient initial entropy.

Note: Good random numbers are not only needed for key generation, but also for session keys and nonces in any security protocol. Therefore, we believe that a decent security architecture should anyways support good random number generation on the EE side or provide enough entropy for the RNG seed during manufacturing to guarantee good initial pseudo-random number generation.

- o Due to lack of computational resources, e.g., in case of RSA keys.

Note: As key generation can be performed in advance to the certificate enrollment communication, it is typical not time critical.

Note: Besides the initial enrollment right after the very first bootup of the device, where entropy available on the device may be insufficient, we do not see any good reason for central key generation.

Note: As mentioned in Section 3.1 central key generation may be required in a push model, where the certificate response message is transferred by the (L)RA/CA to the EE without receiving a previous request message.

If the EE wishes to request central key generation, it MUST fill the subjectPublicKey field in the certTemplate structure of the request message with a zero-length BIT STRING. This indicates to the (L)RA/CA that a new key pair shall be generated centrally on behalf of the EE.

Note: As the protection of centrally generated keys in the response message is being extended from EncryptedValue to EncryptedKey by CMP Updates [I-D.brockhaus-lamps-cmp-updates] also the alternative

EnvelopedData can be used. In CRMF Section 2.1.9 [RFC4211] the use of EncryptedValue has been deprecated in favor of the EnvelopedData structure. Therefore, this profile specifies using EnvelopedData as specified in CMS Section 6 [RFC5652] to offer more crypto agility.




Figure 3: Encrypted private key container

The (L)RA/CA delivers the private key in the privateKey field in the certifiedKeyPair structure of the response message also containing the newly issued certificate.

The private key MUST be wrapped in a SignedData structure, as specified in CMS Section 5 [RFC5652], signed by the KGA generating the key pair. The signature MUST be performed using a CMP signer certificate asserting the extended key usage kp-id-cmpKGA as described in CMP Updates [I-D.brockhaus-lamps-cmp-updates] to show the authorization to generate key pairs on behalf of an EE.

This SignedData structure MUST be wrapped in an EnvelopedData structure, as specified in CMS Section 6 [RFC5652], encrypting it using a newly generated symmetric content-encryption key.

Note: Instead of the specification in CMP Appendix D 4.4 [RFC4210] this content-encryption key is not generated on the EE side. As we just mentioned, central key generation should only be used in this profile in case of lack of randomness on the EE.

As part of the EnvelopedData structure this content-encryption key MUST be securely provided to the EE using one of three key management techniques. The choice of the key management technique to be used by the (L)RA/CA depends on the authentication mechanism the EE choose to protect the request message, see CMP Updates section 3.4 [I-D.brockhaus-lamps-cmp-updates] for more details on which key management technique to use.

- o MAC protected request message: The content-encryption key SHALL be protected using the symmetric key-encryption key management technique, see Section 5.1.6.1, only if the EE used MAC protection for the respected request message.
- o Signature protected request message using a certificate that contains a key usage extension asserting keyAgreement: The content-encryption key SHALL be protected using the key agreement key management technique, see Section 5.1.6.2, if the certificate used by the EE for signing the respective request message contains the key usage keyAgreement. If the certificate also contains the key usage keyEncipherment, the key transport key management technique SHALL NOT be used.
- o Signature protected request message using a certificate that contains a key usage extension asserting keyEncipherment: The content-encryption key SHALL be protected using the key transport key management technique, see Section 5.1.6.3, if the certificate used by the EE for signing the respective request message contains the key usage keyEncipherment and not keyAgreement.

The key agreement key management technique can be supported by most signature algorithms, as key transport key management technique can only be supported by a very limited number of algorithms. The symmetric key-encryption key management technique shall only be used in combination with MAC protection, wich is a side-line in this profile. Therefore, this profile REQUIRES support of the key agreement key management technique and the key transport and symmetric key-encryption key management techniques are OPTIONAL.

For encrypting the SignedData structure containing the private key a fresh content-encryption key MUST be generated with enough entropy with regard to the used symmetric encryption algorithm.

Note: Depending on the lifetime of the certificate and the criticality of the generated private key, it is advisable to use the strongest possible symmetric encryption algorithm. Therefore, this specification recommends using at least AES-256.

The detailed description of the privateKey field looks like this:

```
privateKey          OPTIONAL
-- MUST be an envelopedData structure as specified in
-- CMS [RFC5652] section 6
    version          REQUIRED
-- MUST be set to 2
    recipientInfos    REQUIRED
-- MUST be exactly one RecipientInfo
```

```
    recipientInfo      REQUIRED
-- MUST be either KEKRecipientInfo (see section 5.1.5.1),
-- KeyAgreeRecipientInfo (see section 5.1.5.2), or
-- KeyTransRecipientInfo (see section 5.1.5.3) is used
    encryptedContentInfo
                                REQUIRED
    contentType        REQUIRED
-- MUST be id-signedData
    contentEncryptionAlgorithm
                                REQUIRED
-- MUST be the algorithm identifier of the symmetric
-- content-encryption algorithm used
-- As private keys need long-term protection, the use of AES-256
-- or a stronger symmetric algorithm is RECOMMENDED
    encryptedContent    REQUIRED
-- MUST be the encrypted signedData structure as specified in
-- CMS [RFC5652] section 5
    version            REQUIRED
-- MUST be set to 3
    digestAlgorithms
                                REQUIRED
-- MUST be exactly one digestAlgorithm identifier
    digestAlgorithmIdentifier
                                REQUIRED
-- MUST be the OID of the digest algorithm used for generating
-- the signature
-- The hash algorithm used SHOULD be SHA-256
    encapContentInfo
                                REQUIRED
-- MUST be the content that is to be signed
    contentType        REQUIRED
-- MUST be id-data
    content            REQUIRED
-- MUST be the privateKey as OCTET STRING
    certificates        REQUIRED
-- SHOULD contain the signing certificate together with its chain
-- If present, the first certificate in this field MUST
-- be the certificate used for signing this content
-- Self-signed certificates SHOULD NOT be included
-- and MUST NOT be trusted based on the listing in any case
    crls                OPTIONAL
-- MAY be present to provide status information on the signer or
-- its CA certificates
    signerInfos        REQUIRED
-- MUST be exactly one signerInfo
    version            REQUIRED
-- MUST be set to 3
    sid                REQUIRED
```

```
        subjectKeyIdentifier
            REQUIRED
-- MUST be the subjectKeyIdentifier of the signer's certificate
        digest algorithm
            REQUIRED
-- MUST be the same OID as in digest algorithm
        signatureAlgorithm
            REQUIRED
-- MUST be the algorithm identifier of the signature algorithm
-- used for calculation of the signature bits,
-- like sha256WithRSAEncryption or ecdsa-with-SHA256
-- The signature algorithm MUST be consistent with the
-- SubjectPublicKeyInfo field of the signer's certificate
        signature      REQUIRED
-- MUST be the result of the digital signature generation
```

5.1.6.1. Using symmetric key-encryption key management technique

This key management technique can be applied in combination with the message flow specified in Section 5.1.4 using MAC protected CMP messages. The shared secret used for the MAC protection MUST also be used for the encryption of the content-encryption key but with a different seed in the PBMPParameter sequence. To use this key management technique the KEKRecipientInfo structure MUST be used in the contentInfo field.

The KEKRecipientInfo structure included into the envelopedData structure is specified in CMS Section 6.2.3 [RFC5652].

The detailed description of the KEKRecipientInfo structure looks like this:

```

        recipientInfo      REQUIRED
-- MUST be KEKRecipientInfo as specified in
-- CMS section 6.2.3 [RFC5652]
        version            REQUIRED
-- MUST be set to 4
        kekid              REQUIRED
        keyIdentifier      REQUIRED
-- MUST contain the same value as the senderKID in the respective
-- request messages
        keyEncryptionAlgorithm
                                REQUIRED
-- MUST be id-PasswordBasedMac
        PBMPParameter      REQUIRED
        salt                REQUIRED
-- MUST be the random value to salt the secret key
-- MUST be a different value than used in the PBMPParameter
-- data structure of the CMP message protection in the
-- header of this message
        owf                 REQUIRED
-- MUST be the same value than used in the PBMPParameter
-- data structure in the header of this message
        iterationCount
                                REQUIRED
-- MUST be a limited number of times the OWF is applied
-- To prevent brute force and dictionary attacks a reasonable
-- high number SHOULD be used
        mac                 REQUIRED
-- MUST be the same as in the contentEncryptionAlgorithm field
        encryptedKey       REQUIRED
-- MUST be the encrypted content-encryption key

```

< TBD: To make use of a different symmetric keys for encrypting the private key and for MAC-protection of the CMP message, we derive another key using the same PBMPParameter structure from CMP, even though from the perspective of field names, it is not intended to be used for deriving encryption keys. Does anyone sees a better solution here? >

5.1.6.2. Using key agreement key management technique

This key management technique can be applied in combination with the message flow specified in Section 5.1.1 using signature-based protected CMP messages. The public key of the EE certificate used for the signature-based protection of the request message MUST also be used for the Ephemeral-Static Diffie-Hellmann key establishment of

the content-encryption key. To use this key management technique the KeyAgreeRecipientInfo structure MUST be used in the contentInfo field.

The KeyAgreeRecipientInfo structure included into the envelopedData structure is specified in CMS Section 6.2.2 [RFC5652].

The detailed description of the KeyAgreeRecipientInfo structure looks like this:

```

    recipientInfo      REQUIRED
-- MUST be KeyAgreeRecipientInfo as specified in
    version            REQUIRED
-- MUST be set to 3
    originator         REQUIRED
-- MUST contain the originatorKey sequence
    algorithm          REQUIRED
-- MUST be the algorithm identifier of the
-- static-ephemeral Diffie-Hellmann algorithm
    publicKey          REQUIRED
-- MUST be the ephemeral public key of the sending party
    ukm                OPTIONAL
-- MUST be used when 1-pass ECMQV is used
    keyEncryptionAlgorithm
                        REQUIRED
-- MUST be the same as in the contentEncryptionAlgorithm field
    recipientEncryptedKeys
                        REQUIRED
-- MUST be exactly one recipientEncryptedKey sequence
    recipientEncryptedKey
                        REQUIRED
    rid                REQUIRED
    rKeyId              REQUIRED
    subjectKeyID        REQUIRED
-- MUST contain the same value as the senderKID in the respective
-- request messages
    encryptedKey        REQUIRED
-- MUST be the encrypted content-encryption key

```

5.1.6.3. Using key transport key management technique

This key management technique can be applied in combination with the message flow specified in Section 5.1.1 using signature-based protected CMP messages. The public key of the EE certificate used for the signature-based protection of the request message MUST also be used for key encipherment of the content-encryption key. To use

this key management technique the KeyTransRecipientInfo structure MUST be used in the contentInfo field.

The KeyTransRecipientInfo structure included into the envelopedData structure is specified in CMS Section 6.2.1 [RFC5652].

The detailed description of the KeyTransRecipientInfo structure looks like this:

```
        recipientInfo      REQUIRED
-- MUST be KeyTransRecipientInfo as specified in
-- CMS section 6.2.1 [RFC5652]
        version            REQUIRED
-- MUST be set to 2
        rid                REQUIRED
        subjectKeyIdentifier
                           REQUIRED
-- MUST contain the same value as the senderKID in the respective
-- request messages
        keyEncryptionAlgorithm
                           REQUIRED
-- MUST contain the key encryption algorithm identifier used for
-- public key encryption
        encryptedKey       REQUIRED
-- MUST be the encrypted content-encryption key
```

5.1.7. Delayed enrollment

This functional extension can be applied in combination with certificate enrollment as described in Section 5.1.1 to Section 5.1.5. The functional extension can be used in case a (L)RA/CA cannot respond to the certificate request in a timely manner, e.g., due to offline upstream communication or required registration officer interaction. Depending on the PKI architecture, it is not necessary that the PKI component directly communicating with the EE initiates the delayed enrollment.

The PKI component initiating the delayed enrollment MUST include the status "waiting" in the response and this response MUST not contain the newly issued certificate. When receiving a response with status "waiting" the EE MUST send a poll request to the (L)RA/CA. The PKI component that initiated the delayed enrollment MUST answer with a poll response containing a checkAfter time. This value indicates the minimum number of seconds that must elapse before the EE sends another poll request. As soon as the (L)RA/CA can provide the final response message for the initial request of the EE, it MUST provide this in response to a poll request. After receiving this response,

the EE can continue the original message sequence as described in the respective section of this document, e.g., send a certConf message.

Typically, intermediate PKI entities SHOULD NOT change the sender and recipient nonce even in case an intermediate (L)RA modifies a request or a response message. In the special case of polling between EE and LRA with offline transport between an LRA and RA, see Section 6.1.4, an exception occurs. The EE and LRA exchange pollReq and pollRep messages handle the nonce words as described. When, after pollRep, the final response from the CA arrives at the LRA, the next response will contain the recipientNonce set to the value of the senderNonce in the original request message (copied by the CA). The LRA needs to replace the recipientNonce in this case with the senderNonce of the last pollReq because the EE will validate it in this way.

Message flow:

```

Step# EE                                     (L) RA/CA
1  format ir/cr/p10cr/kur
   As described in the
   respective section
   in this document
2  ->ir/cr/p10cr/kur->
3                                     handle request as described
                                     in the respective section
                                     in this document
4                                     in case no immediate final
                                     response is possible,
                                     receive or format ip, cp
                                     or kup message containing
                                     status "waiting"
5                                     <- ip/cp/kup <-
6  handle ip/cp/kup
7  format pollReq
8  -> pollReq ->
9                                     handle, re-protect or
                                     forward pollReq
10                                    in case the requested
                                     certificate or a
                                     corresponding response
                                     message is available,
                                     receive or format ip, cp,
                                     or kup containing the
                                     issued certificate, or
                                     format or receive pollRep
                                     with appropriate
                                     checkAfter value
11                                    <- pollRep <-
12  handle pollRep
13  let checkAfter
   time elapse
14  continue with line 7

```

Detailed description of the first ip/cp/kup:

Response with status 'waiting' -- ip/cp/kup

Field	Value
-------	-------

header

```

-- MUST contain a header as described for the first response
-- message of the respective scheme

```

```

body
  -- The response of the (L)RA/CA to the request in case no
  -- immediate appropriate response can be sent
  ip/cp/kup                REQUIRED
  response                 REQUIRED
  -- MUST be exactly one CertResponse
  certReqId               REQUIRED
  -- MUST be set to 0
  status                  REQUIRED
  -- PKIStatusInfo structure MUST be present
  status                 REQUIRED
  -- MUST be set to "waiting"
  statusString            OPTIONAL
  -- MAY be any human-readable text for debugging, logging or to
  -- display in a GUI
  failInfo               PROHIBITED
  certifiedKeyPair       PROHIBITED

protection                REQUIRED
  -- MUST contain protection as described for the first response
  -- message of the respective profile, but
  -- MUST use the protection key of the (L)RA/CA initiating the
  -- delayed enrollment and creating this response message

extraCerts                REQUIRED
  -- MUST contain certificates as described for the first response
  -- message of the respective profile.
  -- As no new certificate is issued yet, no respective certificate
  -- chain is included.

Polling Request -- pollReq

Field                    Value

header
  -- MUST contain a header as described for the certConf message
  -- of the respective scheme

body
  -- The message of the EE asks for the final response or for a
  -- time to check again
  pollReq                REQUIRED
  certReqId             REQUIRED
  -- MUST be exactly one value
  -- MUST be set to 0

protection                REQUIRED

```

-- MUST contain protection as described for the certConf message
 -- of the respective profile

extraCerts OPTIONAL
 -- If present, it MUST contain certificates as described for the
 -- certConf message of the respective profile.

Polling Response -- pollRep

Field	Value
-------	-------

header	
	-- MUST contain a header as described for the pkiConf message
	-- of the respective scheme

body	pollRep
	-- The message indicated the time to after which the EE may
	-- send another pollReq messaged for this transaction
pollRep	REQUIRED
	-- MUST be exactly one set of the following values
certReqId	REQUIRED
	-- MUST be set to 0
checkAfter	REQUIRED
	-- time in seconds to elapse before a new pollReq may be sent by
	-- the EE

protection	REQUIRED
	-- MUST contain protection as described for the pkiConf message
	-- of the respective profile, but
	-- MUST use the protection key of the (L)RA/CA that initiated the
	-- delayed enrollment and is creating this response message

extraCerts	OPTIONAL
	-- If present, it MUST contain certificates as described for the
	-- pkiConf message of the respective profile.

Final response -- ip/cp/kup

Field	Value
-------	-------

header	
	-- MUST contain a header as described for the first
	-- response message of the respective scheme
	-- but the recipientNonce MUST be the senderNonce of the last
	-- pollReq message

```

body
  -- The response of the (L)RA/CA to the initial request as
  -- described in the respective profile

protection                                REQUIRED
  -- MUST contain protection as described for the first response
  -- message of the respective profile, but
  -- MUST use the protection key of the (L)RA/CA that initiated the
  -- delayed enrollment and forwarding the response message

extraCerts                                REQUIRED
  -- MUST contain certificates as described for the first
  -- response message of the respective profile

```

5.2. Revoking a certificate

This message sequence should be used by an entity to request the revocation of a certificate. Here the revocation request is used by an EE to revoke one of its own certificates. A (L)RA could also act as an EE to revoke one of its own certificates.

The revocation request message MUST be signed using the certificate that is to be revoked to prove the authorization to revoke to the PKI. The revocation request message is signature-protected using this certificate.

An EE requests the revocation of an own certificate at the CA that issued this certificate. The (L)RA/CA responds with a message that contains the status of the revocation from the CA.

Preconditions:

- 1 The certificate the EE wishes to revoke is not yet expired or revoked.

Message flow:

Step#	EE		(L)RA/CA
1	format rr		
2		-> rr	->
3			handle, re-protect or forward rr
4			receive rp
5		<- rp	<-
6	handle rp		

For this profile, the EE MUST include exactly one RevDetails structure in the rr. In case no error occurred the response to the rr MUST be an rp message. The (L)RA/CA MUST produce a rp containing a status field with a single set of values.

Detailed message description:

Revocation Request -- rr

Field	Value
-------	-------

header

-- As described in section 4.1

body

-- The request of the EE to revoke its certificate

rr REQUIRED

-- MUST contain exactly one element of type RevDetails

-- If more revocations are desired, further requests MUST be

-- packaged in separate PKI Messages

certDetails REQUIRED

-- MUST be present and is of type CertTemplate

serialNumber REQUIRED

-- MUST contain the certificate serialNumber attribute of the

-- X.509 certificate to be revoked

issuer REQUIRED

-- MUST contain the issuer attribute of the X.509 certificate to

-- be revoked

crlEntryDetails REQUIRED

-- MUST contain exactly one reasonCode of type CRLReason (see

-- [RFC5280] section 5.3.1)

-- If the reason for this revocation is not known or shall not be

-- published the reasonCode MUST be 0 = unspecified

protection REQUIRED

-- As described in section 4.2 and the private key related to the

-- certificate to be revoked

extraCerts REQUIRED

-- As described in section 4.3

Revocation Response -- rp

Field	Value
-------	-------

header

-- As described in section 4.1

```
body
  -- The responds of the (L)RA/CA to the request as appropriate
  rp                                REQUIRED
  status                            REQUIRED
  -- MUST contain exactly one element of type PKIStatusInfo
  status                            REQUIRED
  -- positive value allowed: "accepted"
  -- negative value allowed: "rejection"
  statusString                      OPTIONAL
  -- MAY be any human-readable text for debugging, logging or to
  -- display in a GUI
  failInfo                          OPTIONAL
  -- MAY be present if and only if status is "rejection"

protection                          REQUIRED
  -- As described in section 4.2

extraCerts                          REQUIRED
```

5.3. Error reporting

This functionality should be used by an EE to report any error conditions upstream to the (L)RA/CA. Error reporting by the (L)RA downstream to the EE is described in Section 6.3.

In case the error condition is related to specific details of an ip, cp, or kup response message and a confirmation is expected the error condition MUST be reported in the respective certConf message with negative contents.

General error conditions, e.g., problems with the message header, protection, or extraCerts, and negative feedback on rp, pollRep, or pkiConf messages MAY be reported in the form of an error message.

In both situations the error is reported in the PKIStatusInfo structure of the respective message.

The (L)RA/CA MUST respond to an error message with a pkiConf message, or with another error message if any part of the header is not valid. Both sides MUST treat this message as the end of the current transaction.

The PKIStatusInfo structure is used to report errors. The PKIStatusInfo structure SHOULD consist of the following fields:

- o status: Here the PKIStatus value rejection is the only one allowed.

- o `statusString`: Here any human-readable valid value for logging or to display in a GUI SHOULD be added.
- o `failInfo`: Here the `PKIFailureInfo` values MAY be used in the following way. For explanation of the reason behind a specific value, please refer to [RFC4210] Appendix F.
- * `transactionIdInUse`: This is sent in case the received request contains a transaction ID that is already in use for another transaction. An EE receiving such error message SHOULD resend the request in a new transaction using a different transaction ID.
- * `systemUnavail` or `systemFailure`: This is sent in case a back-end system is not available or currently not functioning correctly. An EE receiving such error message SHOULD resend the request in a new transaction after some time.

Detailed error message description:

Error Message -- error

Field	Value
-------	-------

header

-- As described in section 4.1

body

-- The message sent by the EE or the (L)RA/CA to indicate an error that occurred

error	REQUIRED
-------	----------

pKIStatusInfo	REQUIRED
---------------	----------

status	REQUIRED
--------	----------

-- MUST have the value "rejection"

statusString	RECOMMENDED
--------------	-------------

-- SHOULD be any human-readable text for debugging, logging

-- or to display in a GUI

failInfo	OPTIONAL
----------	----------

-- MAY be present

protection	REQUIRED
------------	----------

-- As described in section 4.2

extraCerts	OPTIONAL
------------	----------

-- As described in section 4.3

5.4. Support messages

The following support messages offer on demand in-band transport of content that may be provided by the (L)RA/CA and relevant to the EE. The general messages and general response are used for this purpose. Depending on the environment, these requests are answered by the LRA, RA, or CA.

The general message and general response transport InfoTypeAndValue structures. In addition to those infoType values defined in CMP [RFC4210] further OIDs MAY be defined to define new PKI management operations, or general-purpose messages as needed in a specific environment.

Possible content described here address:

- o Request of CA certificates
- o Update of Root CA certificates
- o Parameters needed for a planned certificate request message
- o Voucher request and enrollment voucher exchange

5.4.1. General message and response

The general message transaction is similar to that given in CMP Appendix E.5 [RFC4210]. In this section the general message (genm) and general response (genp) are described. The specific InfoTypeAndValue structures are described in the following sections.

The behavior in case an error occurs is described in Section 5.3.

Message flow:

Step#	EE		(L) RA/CA
1	format genm		
2		-> genm	->
3			handle, re-protect or forward genm
4			format or receive genp
5		<- genp	<-
6	handle genp		

Detailed message description:

General Message -- genm

Field	Value
header	
	-- As described in section 4.1
body	
	-- The request of the EE to receive information
genm	REQUIRED
	-- MUST contain exactly one element of type
	-- InfoTypeAndValue
infoType	REQUIRED
	-- MUST be the OID identifying the specific scheme
	-- described below
infoValue	OPTIONAL
	-- MUST be as described in the specific scheme described
	-- below
protection	REQUIRED
	-- As described in section 4.2
extraCerts	REQUIRED
	-- As described in section 4.3

General Response -- genp

Field	Value
header	
	-- As described in section 4.1
body	
	-- The response of the (L)RA/CA to the information request
genp	REQUIRED
	-- MUST contain exactly one element of type
	-- InfoTypeAndValue
infoType	REQUIRED
	-- MUST be the OID identifying the specific scheme
	-- described below
infoValue	OPTIONAL
	-- MUST be as described in the specific scheme described
	-- below
protection	REQUIRED
	-- As described in section 4.2
extraCerts	REQUIRED
	-- As described in section 4.3

5.4.2. Get CA certificates

This scheme can be used by an EE to request CA certificates from the (L)RA/CA.

An EE requests CA certificates from the (L)RA/CA by sending a general message with OID `id-it-getCaCerts`. The (L)RA/CA responds with a general response with the same OID that either contains a SEQUENCE of certificates populated with the available CA intermediate and issuing CA certificates or with no content in case no CA certificate is available.

< NOTE: The OID `id-it-getCaCerts` is not yet defined. It should be registered in the tree 1.3.6.1.5.5.7.4 (`id-it`) like other `infoType` OIDs, see CMP Appendix F [RFC4210] on page 92. >

The profile for this exchange is as given in Section 5.4.1, with the following specific content:

- 1 the body MUST contain as `infoType` the OID `id-it-getCaCerts`
- 2 the `infoValue` of the request MUST be absent
- 3 if present, the `infoValue` of the response MUST be `caCerts` field

The `infoValue` field of the general response containing the `id-it-getCaCerts` OID looks like this:

<code>infoValue</code>	OPTIONAL
-- MUST be absent if no CA certificate is available	
-- MUST be present if CA certificates are available	
<code>caCerts</code>	REQUIRED
-- MUST be present if <code>infoValue</code> is present	
-- MUST be a sequence of <code>CMPCertificate</code>	

5.4.3. Get root CA certificate update

This scheme can be used by an EE to request an update of an existing root CA Certificate by the EE. It utilizes the `CAKeyUpdAnnContent` structure as described in CMP Appendix E.4 [RFC4210] as response to a respective general message.

An EE requests a root CA certificate update from the (L)RA/CA by sending a general message with OID `id-it-caKeyUpdateInfo` as `infoType` and no `infoValue`. The (L)RA/CA responds with a general response with the same OID that either contains the update of the root CA certificate consisting of up to three certificates, or with no content in case no update is available.

These three certificates are described in more detail in section 4.4.1, section 6.2, and Appendix E.3 of [RFC4210]. The newWithNew certificate is the new root CA certificates and is REQUIRED to be present in the response message. The newWithOld certificate RECOMMENDED to be present in the response message though it is required for those cases where the receiving entity trusts the old root CA certificate and wishes to gain trust in the new root CA certificate. The oldWithNew certificate is OPTIONAL though it is only needed in a scenario where the requesting entity already trusts the new root CA certificate and wants to gain trust in the old root certificate.

The profile for this exchange is as given in Section 5.4.1, with the following specific content:

- 1 the body MUST contain as infoType the OID id-it-caKeyUpdateInfo
- 2 the infoValue of the request MUST be absent
- 3 if present, the infoValue of the response MUST be a CAKeyUpdAnnContent structure

The infoValue field of the general response containing the id-it-caKeyUpdateInfo extension looks like this:

```

    infoValue                OPTIONAL
-- MUST be absent if no update of the root CA certificate is
    available
-- MUST be present if an update of the root CA certificate
-- is available
    caKeyUpdateInfo          REQUIRED
-- MUST be present and be of type CAKeyUpdAnnContent
    oldWithNew                OPTIONAL
-- MUST be present if infoValue is present
-- MUST contain an X.509 certificate containing the old public
-- root CA key signed with the new private root CA key
    newWithOld                RECOMMENDED
-- MUST be present if infoValue is present
-- MUST contain an X.509 certificate containing the new public
-- root CA key signed with the old private root CA key
    newWithNew                REQUIRED
-- MUST be present if infoValue is present
-- MUST contain the new root CA certificate

```

5.4.4. Get certificate request parameters

This scheme can be used by an EE to request configuration parameters for a planned certificate request transaction.

An EE requests certificate request parameters from the (L)RA/CA by sending a general message with OID id-it-getCSRParam. The (L)RA/CA responds with a general response with the same OID that either contains the required fields, e.g., algorithm identifier for key pair generation or other attributes and extensions or with no content in case no specific requirements are made by the (L)RA/CA.

< NOTE: The OID id-it-getCSRParam is not yet defined. It should be registered in the tree 1.3.6.1.5.5.7.4 (id-it) like other infoType OIDs, see CMP Appendix F [RFC4210] on page 92. >

The EE SHOULD follow the requirements from the recieved CertTemplate and the optional RSA key length. In case a field is present but the value is absent, it means that this field is required but its content has to be provided by the EE.

< TBD: There is some more explanation needed to explain how to prefill the certTemplate structure. Possibly an example will help to clarify this. >

The profile for this exchange is as given in Section 5.4.1, with the following specific content:

- 1 the body MUST contain as infoType the OID id-it-getCSRParam
- 2 the infoValue of the request MUST be absent
- 3 if present, the infoValue of the response MUST be a SEQUENCE of a certTemplate structure and an rsaKeyLen field of type INTEGER

The infoValue field of the general response containing the id-it-getCSRParam OID looks like this:

```
    infoValue                                OPTIONAL
-- MUST be absent if no requirements are available
-- MUST be present if the (L)RA/CA has any requirements on the
-- content of the certificates to be requested.
    certTemplate                            REQUIRED
-- MUST be present if infoValue is present
-- MUST contain the prefilled certTemplate structure
    rsaKeyLen                               OPTIONAL
-- This field is of type INTEGER. Any reasonable RSA key length
-- SHOULD be specified if the algorithm in the
-- subjectPublicKeyInfo field of the certTemplate is of type
-- rsaEncryption.
```

5.4.5. Get certificate management configuration

This scheme can be used by an EE to request the current certificate management configuration information by the EE in advance to a planned certificate management transaction, e.g., in case no out-of-band transport is available. Such certificate management configuration can consist of all information the EE needs to know to generate and deliver a proper certificate request, such as

- o algorithm, curve, and key length for key generation
- o various certificate attributes and extensions to be used for the certificate request
- o specific host name, port and path on the RA/LRA to send this CMP request to
- o Infrastructure Root CA Certificate, e.g., the root of the (L)RA TLS and CMP signer certificates.

There is an overlap with Section 5.4.2 with regard to transport of CA certificates and with Section 5.4.4 with regard to key generation parameter and certificate request attributes and extensions. This profile offers to request a proprietary configuration file containing all information needed in one exchange.

< TBD: Especially with section 5.4.4 there is some overlap regarding algorithms, attributes and, extensions of the certificate that will be requested. It needs to be decided if both variants have a right to exist next to the other or if one option should be removed from this document. >

An EE requests certificate management configuration from the (L)RA/CA by sending a general message with the OID `id-it-getCertMgtConfig`. The (L)RA/CA responds with a general response with the same OID that either contains a `certMgtConfig` field containing the configuration file encoded as OCTET STRING or with no content in case no certificate management configuration is available.

< NOTE: The OID `id-it-getCertMgtConfig` is not yet defined. It should be registered in the tree 1.3.6.1.5.5.7.4 (`id-it`) like other `infoType` OIDs, see CMP Appendix F [RFC4210] on page 92. >

The EE SHOULD use the contents of this `certMgtConfig` to format and deliver the certificate request. The certificate management configuration may contain contact details, e.g., like an URI and issuing CA distinguished name, where to address the request messages to and may also contain certificate request parameters as described in Section 5.4.4.

The `certMgtConfig` field may be of any format suitable for the EE, e.g., CMS [RFC5652], JWT [RFC7519] or, XML [W3C_XML]. The `certMgtConfig` contents MAY be signed, e.g., like CMS SignedData [RFC5652], JWS [RFC7515] or, XML-DSig [W3C_XML-Dsig]. For interoperability the format of the `certMgtConfig` field should be specified in detail if needed.

The profile for this exchange is as given in Section 5.4.1, with the following specific content:

- 1 the body MUST contain as `infoType` the OID `id-it-getCertMgtConfig`
- 2 the `infoValue` of the request MUST be absent
- 3 if present, the `infoValue` of the response MUST be a `certMgtConfig` structure

The `infoValue` field of the general response containing the `id-it-getCertMgtConfig` extension looks like this:

```
infoValue          OPTIONAL
-- MUST be absent if no certificate management configuration
-- is available
-- MUST be present if the (L)RA/CA provides any certificate
-- management configuration
  certMgtConfig    REQUIRED
-- MUST be present if infoValue is present
-- MUST contain the certificate management configuration as OCTET
-- OCTET STRING
```

5.4.6. Get enrollment voucher

This scheme can be used by an EE to request an enrollment voucher containing the root certificate of a new, additional, or alternative PKI to establish trust in this PKI, e.g., in case no out-of-band transport is available. Such an enrollment voucher can be used in advance to an enrollment to this new environment. It may contain further information depending on the use case.

An EE requests an enrollment voucher from the (L)RA/CA by sending a general message. The (L)RA/CA responds with a general response with the same OID that either contains the voucher or with no content in case no voucher is available.

The (L)RA MAY use the content of the voucherRequest to get an enrollment voucher from other backend components, e.g., as described in BRSKI [I-D.ietf-anima-bootstrapping-keyinfra]. The EE SHOULD use the contents of the received enrollmentVoucher to authenticate the (L)RA/CA it is about to enroll to. The enrollment voucher may for example contain the Root CA certificate of the new PKI or the CMP signer certificate of the (L)RA. The general response message MUST be properly authenticated and the sender of this message MUST be authorized to install new root certificates. One example for an enrollment voucher is specified in RFC8366 [RFC8366].

The voucherRequest and enrollmentVoucher fields may be of any format suitable for the EE, e.g., CMS [RFC5652], JWT [RFC7519] or, XML [W3C_XML]. The voucherRequest and enrollmentVoucher contents MAY contain a signature, e.g., CMS SignedData [RFC5652], JWS [RFC7515] or, XML-DSig [W3C_XML-DSig]. For interoperability the format of the voucherRequest and enrollmentVoucher field should be specified in detail if needed, e.g., as defined in BRSKI [I-D.ietf-anima-bootstrapping-keyinfra] and RFC8366 [RFC8366].

< TBD: The content of the voucherRequest and enrollmentVoucher fields can also be limited to the specifications in BRSKI [I-D.ietf-anima-bootstrapping-keyinfra] and RFC8366 [RFC8366]. >

The profile for this exchange is as given in Section 5.4.1, with the following specific content:

- 1 the body MUST contain as infoType the OID id-it-getEnrollmentVoucher
- 2 if present, the infoValue of the request MUST be a voucherRequest structure

- 3 if present, the infoValue of the response MUST be an enrollmentVoucher structure

The infoValue field of the general message containing the id-it-getEnrollmentVoucher extension looks like this:

```
infoValue          OPTIONAL
-- MUST be absent if no voucher request is available
-- MUST be present if the EE provides the voucher request
  voucherRequest    REQUIRED
-- MUST be present if infoValue is present
-- MUST contain the voucher request as OCTET STRING
```

The infoValue field of the general response containing the id-it-getEnrollmentVoucher extension looks like this:

```
infoValue          OPTIONAL
-- MUST be absent if no enrollment voucher is available
-- MUST be present if the (L)RA/CA provides the enrollment
  voucher
    enrollmentVoucher  REQUIRED
-- MUST be present if infoValue is present
-- MUST contain the enrollment voucher as OCTET STRING
```

6. LRA and RA focused certificate management use cases

This chapter focuses on the communication of PKI backend components with each other. Depending on the network and PKI solution design, these will either be an LRA, RA or CA.

Typically, an (L)RA forwards messages from downstream, but it may also reply to them itself. Besides forwarding of received messages an (L)RA could also need to revoke certificates of EEs, report errors, or may need to manage its own certificates.

< TBD: In CMP Updates [I-D.brockhaus-lamps-cmp-updates] additional extended key usages like id-kp-cmpRA will be defined to indicate that a key pair is entitled to be used for signature-based protection of a CMP message by an (L)RA/CA. >

6.1. Forwarding of messages

Each CMP request message (i.e., ir, cr, p10cr, kur, pollReq, or certConf) or error message coming from an EE or the previous (downstream) PKI component MUST be sent to the next (upstream) PKI component. This PKI component MUST forward response messages to the next (downstream) PKI component or EE.

The (L)RA SHOULD verify the protection, the syntax, the required message fields, the message type, and if applicable the authorization and the proof-of-possession of the message. Additional checks or actions MAY be applied depending on the PKI solution requirements and concept. If one of these verification procedures fails, the (L)RA SHOULD respond with a negative response message and SHOULD not forward the message further upstream. General error conditions should be handled as described in Section 5.3 and Section 6.3.

An (L)RA SHOULD not change the received message if not necessary. The (L)RA SHOULD only update the message protection if it is technically necessary. Concrete PKI system specifications may define in more detail if and when to do so.

This is particularly relevant in the upstream communication of a request message.

Each hop in a chain of PKI components has one or more functionalities, e.g.,

- o An (L)RA may need to verify the identities of EEs or base authorization decisions for certification request processing on specific knowledge of the local setup, e.g., by consulting an inventory or asset management system.
- o An (L)RA may need to add fields to certificate request messages.
- o An (L)RA may need to store data from a message in a database for later usage or documentation purposes.
- o An (L)RA may provide traversal of a network boundary.
- o An (L)RA may need to double-check if the messages transferred back and forth are properly protected and well formed.
- o An (L)RA may provide a proof that it has performed all required checks.
- o An (L)RA may initiate a delayed enrollment due to offline upstream communication or registration officer interaction.
- o An (L)RA may grant the request of an EE to omit sending a confirmation message.
- o An RA can collect messages from different LRAs and forward them to the CA.

Therefore, the decision if a message should be forwarded

- o unchanged with the original protection,
- o unchanged with a new protection, or
- o changed with a new protection

depends on the PKI solution design and the associated security policy (CP/CPS [RFC3647]).

< TBD: In [CMP Updates] different circumstances that require adding of an additional protection by an (L)RA or batching CMP messages at an (L)RA by using the nested messages is described. It needs to be decided which of these variants should be specified here. Finally, I guess they will all be OPTIONAL. >

This section specifies the different options an (L)RA may implement and use.

An (L)RA MAY update the protection of a message

- o if the (L)RA performs changes to the header or the body of the message,
- o if the (L)RA needs to prove checks or validations performed on the message to one of the next (upstream) PKI components,
- o if the (L)RA needs to protect the message using a key and certificate from a different PKI, or
- o if the (L)RA needs to replace a MAC based-protection.

This is particularly relevant in the upstream communication of certificate request messages.

The message protection covers only the header and the body and not the extraCerts. The (L)RA MAY change the extraCerts in any of the following message adaptations, e.g., to sort or add needed or to delete needless certificates to support the next hop. This may be particularly helpful to extend upstream messages with additional certificates or to reduce the number of certificates in downstream messages when forwarding to constrained devices.

6.1.1. Not changing protection

This message adaptation can be used by any (L)RA to forward an original CMP message without changing the header, body or protection. In any of these cases the (L)RA acts more like a proxy, e.g., on a

network boundary, implementing no specific RA-like security functionality to the PKI.

This message adaptation MUST be used for forwarding kur messages that must not be approved by the respective (L)RA.

6.1.2. Replacing protection

The following two message adaptations can be used by any (L)RA to forward a CMP message with or without changes, but providing its own protection using its CMP signer key providing approval of this message. In this case the (L)RA acts as an actual Registration Authority (RA), which implements important security functionality of the PKI.

Before replacing the existing protection by a new protection, the (L)RA MUST verify the protection provided by the EE or by the previous PKI component and approve its content including any own modifications. For certificate requests the (L)RA MUST verify in particular the included proof-of-possession self-signature of the certTemplate using the public key of the requested certificate and MUST check that the EE, as authenticated by the message protection, is authorized to request a certificate with the subject as specified in the certTemplate.

In case the received message has been protected by a CA or another (L)RA, the current (L)RA MUST verify its protection and approve its content including any own modifications. For certificate requests the (L)RA MUST check that the other (L)RA, as authenticated by the message protection, is authorized to issue or forward the request.

These message adaptations MUST NOT be applied to kur request messages as described in Section 5.1.3 since their original protection using the key and certificate to be updated needs to be preserved, unless the regCtrl OldCertId is used to clearly identify the certificate to be updated.

6.1.2.1. Keeping proof-of-possession

This message adaptation can be used by any (L)RA to forward a CMP message with or without modifying the message header or body while preserving any included proof-of-possession.

By replacing the existing protection using its own CMP signer key the (L)RA provides a proof of verifying and approving of the message as described above.

In case the (L)RA modifies the certTemplate of an ir or cr message, the message adaptation in Section 6.1.2.2 needs to be applied instead.

6.1.2.2. Breaking proof-of-possession

This message adaptation can be used by any (L)RA to forward an ir or cr message with modifications of the certTemplate i.e., modification, addition, or removal of fields. Such changes will break the proof-of-possession provided by the EE in the original message.

By replacing the existing or applying an initial protection using its own CMP signer key the (L)RA provides a proof of verifying and approving the new message as described above.

In addition to the above the (L)RA MUST verify in particular the proof-of-possession contained in the original message as described above. If these checks were successfully performed the (L)RA MUST change the popo to raVerified.

The popo field MUST contain the raVerified choice in the certReq structure of the modified message as follows:

```
popo
  raVerified          REQUIRED
-- MUST have the value NULL and indicates that the (L)RA
-- verified the popo of the original message.
```

6.1.3. Adding Protection

< TBD: In [CMP Updates] different circumstances that require adding of an additional protection by an (L)RA or batching CMP messages at an (L)RA by using the nested messages is described. It needs to be decided which of these variants should be specified here. Finally, I guess they will all be OPTIONAL. >

6.1.4. Initiating delayed enrollment

This message adaptation can be used by an (L)RA to initiate delayed enrollment. In this case a (L)RA/CA MUST add the status waiting in the response message. The (L)RA/CA MUST then reply to the pollReq messages as described in Section 5.1.7.

6.2. Revoking certificates on behalf of another's entities

This message sequence can be used by an (L)RA to revoke a certificate of any other entity. This revocation request message MUST be signed

by the (L)RA using its own CMP signer key to prove to the PKI authorization to revoke the certificate on behalf of the EE.

The general message flow for this profile is the same as given in section Section 5.2.

Preconditions:

- 1 the certificate to be revoked MUST be known to the (L)RA
- 2 the (L)RA MUST have the authorization to revoke the certificates of other entities issued by the corresponding CA

The profile for this exchange is identical to that given in section Section 5.2, with the following changes:

- 1 it is not required that the certificate to be revoked is not yet expired or revoked
- 2 the (L)RA acts as EE for this message exchange
- 3 the rr messages MUST be signed using the CMP signer key of the (L)RA.

6.3. Error reporting

This functionality should be used by the (L)RA to report any error conditions downstream to the EE. Potential error reporting by the EE upstream to the (L)RA/CA is described in Section 5.3.

In case the error condition is related to specific details of an ir, cr, pl0cr, or kur request message it MUST be reported in the specific response message, i.e., an ip, cp, or kup with negative contents.

General error conditions, e.g., problems with the message header, protection, or extraCerts, and negative feedback on rr, pollReq, certConf, or error messages MUST be reported in the form of an error message.

In both situations the (L)RA reports the errors in the PKIStatusInfo structure of the respective message as described in Section 5.3.

An EE receiving any such negative feedback SHOULD log the error appropriately and MUST terminate the current transaction.

7. CMP message transport variants

The CMP messages are designed to be self-contained, such that in principle any transport can be used. HTTP SHOULD be used for online transport while file-based transport MAY be used in case offline transport is required. In case HTTP transport is not desired or possible, CMP messages MAY also be piggybacked on any other reliable transport protocol, e.g., CoAP [RFC7252].

Independently of the means of transport it could happen that messages are lost, or a communication partner does not respond. In order to prevent waiting indefinitely, each CMP client component SHOULD use a configurable per-request timeout, and each CMP server component SHOULD use a configurable per-response timeout in case a further message is to be expected from the client side. In this way a hanging transaction can be closed cleanly with an error and related resources (for instance, any cached extraCerts) can be freed.

7.1. HTTP transport

This transport mechanism can be used by an EE and (L)RA/CA to transfer CMP messages over HTTP. If HTTP transport is used the specifications as described in [RFC6712] MUST be followed.

Each PKI management entity supporting HTTP(S) transport MUST support the use of the path-prefix of '/.well-known/' as defined in [RFC5785] and the registered name of 'cmp' to ease interworking in a multi-vendor environment.

The CMP client MUST be configured with sufficient information to form the CMP server URI. This MUST be at least the authority portion of the URI, e.g., 'www.example.com:80', or the full operational path of the CA/RA. An additional arbitrary label, e.g., 'arbitraryLabel1', MAY be configured as a separate component or as part of the full operational path to provide further information to address multiple CAs or certificate profiles. A valid full operational path can look like this:

- 1 http://www.example.com/.well-known/cmp/keyupdate
- 2 http://www.example.com/.well-known/cmp/arbitraryLabel1/keyupdate

PKI management operations MUST use the following URI path:

PKI management operation	Path	Details
Enroll client to new PKI (REQUIRED)	/initialization	Section 5.1.1
Enroll client to existing PKI (OPTIONAL)	/certification	Section 5.1.2
Update client certificate (REQUIRED)	/keyupdate	Section 5.1.3
Enroll client using PKCS#10 (OPTIONAL)	/p10	Section 5.1.5
Enroll client using central key generation (OPTIONAL)	/serverkeygen	Section 5.1.6
Revoke client certificate (RECOMMENDED)	/revocation	Section 5.2
Get CA certificates (OPTIONAL)	/getCAcert	Section 5.4.2
Get root CA certificate update (OPTIONAL)	/getRootCAcertUpdate	Section 5.4.3
Get certificate request parameters (OPTIONAL)	/getCSRparam	Section 5.4.4
Get certificate management configuration (OPTIONAL)	/getCertMgtConfig	Section 5.4.5
Get enrollment voucher (OPTIONAL)	/getVoucher	Section 5.4.6

Table 1: HTTP endpoints

Subsequent certConf, error, and pollReq messages are sent to the URI of the respective PKI management operation.

< TBD: It needs to be defined if specific path values for communication between PKI management entities as specified in section 6 are needed, e.g., 'forward' or 'nested'.>

7.2. HTTPS transport using certificates

This transport mechanism can be used by an EE and (L)RA/CA to further protect the HTTP transport as described in Section 7.1 using TLS 1.2 [RFC5246] or TLS 1.3 [RFC8446] as described in [RFC2818] with certificate-based authentication. Using this transport mechanism, the CMP transport via HTTPS MUST use TLS server authentication and SHOULD use TLS client authentication.

EE:

- o The EE SHOULD use a TLS client certificate as far as available. If no dedicated TLS certificate is available the EE SHOULD use an already existing certificate identifying the EE (e.g., a manufacturer certificate).
- o If no TLS certificate is available at the EE, server-only authenticated TLS SHOULD be used.
- o The EE MUST validate the TLS server certificate of its communication partner.

(L)RA:

- o Each (L)RA SHOULD use a TLS client certificate on its upstream (client) interface.
- o Each (L)RA SHOULD use a TLS server certificate on its downstream (server) interface.
- o Each (L)RA MUST validate the TLS certificate of its communication partner.

NOTE: The requirements for checking certificates given in [RFC5280], [RFC5246] and [RFC8446] MUST be followed for the TLS layer. Certificate status checking SHOULD be used for the TLS certificates of communication partners.

7.3. HTTPS transport using shared secrets

This transport mechanism can be used by an EE and (L)RA/CA to further protect the HTTP transport as described in Section 7.1 using TLS 1.2 [RFC5246] or TLS 1.3 [RFC8446] as described in [RFC2818] with mutual authentication based on shared secrets as described in [RFC5054].

EE:

- o The EE MUST use the shared symmetric key for authentication.

(L)RA:

- o The (L)RA MUST use the shared symmetric key for authentication.

7.4. File-based transport

For offline transfer file-based transport MAY be used. Offline transport is typically used between LRA and RA nodes.

Connection and error handling mechanisms like those specified for HTTP in [RFC6712] need to be implemented.

< TBD: Details need to be defined later >

7.5. CoAP transport

In constrained environments where no HTTP transport is desired or possible, CoAP [RFC7252] MAY be used instead. Connection and error handling mechanisms like those specified for HTTP in [RFC6712] may need to be implemented.

Such specification is out of scope of this document and would need to be specifies in a separate document.

7.6. Piggybacking on other reliable transport

For online transfer where no HTTP transport is desired or possible CMP messages MAY also be transported on some other reliable protocol. Connection and error handling mechanisms like those specified for HTTP in [RFC6712] need to be implemented.

Such specification is out of scope of this document and would need to be specifies in a separate document, e.g., in the scope of the respective transport protocol used.

8. IANA Considerations

<Add any IANA considerations>

9. Security Considerations

<Add any security considerations>

10. Acknowledgements

We would like to thank the various reviewers of this CMP profile.

11. References

11.1. Normative References

- [I-D.brockhaus-lamps-cmp-updates]
Brockhaus, H., "CMP Updates", draft-brockhaus-lamps-cmp-updates-02 (work in progress), December 2019.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC2986] Nystrom, M. and B. Kaliski, "PKCS #10: Certification Request Syntax Specification Version 1.7", RFC 2986, DOI 10.17487/RFC2986, November 2000, <<https://www.rfc-editor.org/info/rfc2986>>.
- [RFC4086] Eastlake 3rd, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security", BCP 106, RFC 4086, DOI 10.17487/RFC4086, June 2005, <<https://www.rfc-editor.org/info/rfc4086>>.
- [RFC4210] Adams, C., Farrell, S., Kause, T., and T. Mononen, "Internet X.509 Public Key Infrastructure Certificate Management Protocol (CMP)", RFC 4210, DOI 10.17487/RFC4210, September 2005, <<https://www.rfc-editor.org/info/rfc4210>>.
- [RFC4211] Schaad, J., "Internet X.509 Public Key Infrastructure Certificate Request Message Format (CRMF)", RFC 4211, DOI 10.17487/RFC4211, September 2005, <<https://www.rfc-editor.org/info/rfc4211>>.
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <<https://www.rfc-editor.org/info/rfc5280>>.
- [RFC5652] Housley, R., "Cryptographic Message Syntax (CMS)", STD 70, RFC 5652, DOI 10.17487/RFC5652, September 2009, <<https://www.rfc-editor.org/info/rfc5652>>.
- [RFC5785] Nottingham, M. and E. Hammer-Lahav, "Defining Well-Known Uniform Resource Identifiers (URIs)", RFC 5785, DOI 10.17487/RFC5785, April 2010, <<https://www.rfc-editor.org/info/rfc5785>>.

- [RFC6712] Kaese, T. and M. Peylo, "Internet X.509 Public Key Infrastructure -- HTTP Transfer for the Certificate Management Protocol (CMP)", RFC 6712, DOI 10.17487/RFC6712, September 2012, <<https://www.rfc-editor.org/info/rfc6712>>.

11.2. Informative References

- [ETSI-3GPP] 3GPP, "TS33.310; Network Domain Security (NDS); Authentication Framework (AF); Release 16; V16.1.0", December 2018, <http://www.3gpp.org/ftp/Specs/archive/33_series/33.310/>.
- [I-D.ietf-anima-bootstrapping-keyinfra] Pritikin, M., Richardson, M., Eckert, T., Behringer, M., and K. Watsen, "Bootstrapping Remote Secure Key Infrastructures (BRSKI)", draft-ietf-anima-bootstrapping-keyinfra-34 (work in progress), January 2020.
- [IEC62443-3-3] IEC, "Industrial communication networks - Network and system security - Part 3-3: System security requirements and security levels", IEC 62443-3-3, August 2013, <<https://webstore.iec.ch/publication/7033>>.
- [IEEE802.1AR] IEEE, "802.1AR Secure Device Identifier", June 2018, <<http://standards.ieee.org/findstds/standard/802.1AR-2009.html>>.
- [NIST-CSFW] NIST, "Framework for Improving Critical Infrastructure Cybersecurity Version 1.1", April 2018, <<https://www.nist.gov/publications/framework-improving-critical-infrastructure-cybersecurity-version-11>>.
- [RFC2818] Rescorla, E., "HTTP Over TLS", RFC 2818, DOI 10.17487/RFC2818, May 2000, <<https://www.rfc-editor.org/info/rfc2818>>.
- [RFC3647] Chokhani, S., Ford, W., Sabett, R., Merrill, C., and S. Wu, "Internet X.509 Public Key Infrastructure Certificate Policy and Certification Practices Framework", RFC 3647, DOI 10.17487/RFC3647, November 2003, <<https://www.rfc-editor.org/info/rfc3647>>.

- [RFC5054] Taylor, D., Wu, T., Mavrogiannopoulos, N., and T. Perrin, "Using the Secure Remote Password (SRP) Protocol for TLS Authentication", RFC 5054, DOI 10.17487/RFC5054, November 2007, <<https://www.rfc-editor.org/info/rfc5054>>.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/RFC5246, August 2008, <<https://www.rfc-editor.org/info/rfc5246>>.
- [RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", RFC 7252, DOI 10.17487/RFC7252, June 2014, <<https://www.rfc-editor.org/info/rfc7252>>.
- [RFC7515] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Signature (JWS)", RFC 7515, DOI 10.17487/RFC7515, May 2015, <<https://www.rfc-editor.org/info/rfc7515>>.
- [RFC7519] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Token (JWT)", RFC 7519, DOI 10.17487/RFC7519, May 2015, <<https://www.rfc-editor.org/info/rfc7519>>.
- [RFC8366] Watsen, K., Richardson, M., Pritikin, M., and T. Eckert, "A Voucher Artifact for Bootstrapping Protocols", RFC 8366, DOI 10.17487/RFC8366, May 2018, <<https://www.rfc-editor.org/info/rfc8366>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.
- [UNISIG] UNISIG, "Subset-137; ERTMS/ETCS On-line Key Management FFFIS; V1.0.0", December 2015, <https://www.era.europa.eu/filebrowser/download/542_en>.
- [W3C_XML] W3C, "Extensible Markup Language (XML) 1.0", W3C XML, November 2008, <<https://www.w3.org/TR/xml/>>.
- [W3C_XML-Dsig] W3C, "XML Signature Syntax and Processing Version 2.0", W3C XML-DSIG, July 2015, <<https://www.w3.org/TR/xmlsig-core2/>>.

Appendix A. Additional Stuff

This becomes an Appendix.

Authors' Addresses

Hendrik Brockhaus
Siemens AG
Otto-Hahn-Ring 6
Munich 81739
Germany

Email: hendrik.brockhaus@siemens.com
URI: <http://www.siemens.com/>

Steffen Fries
Siemens AG
Otto-Hahn-Ring 6
Munich 81739
Germany

Email: steffen.fries@siemens.com
URI: <http://www.siemens.com/>

David von Oheimb
Siemens AG
Otto-Hahn-Ring 6
Munich 81739
Germany

Email: david.von.oheimb@siemens.com
URI: <http://www.siemens.com/>

INTERNET-DRAFT
Internet Engineering Task Force (IETF)
Intended Status: Proposed Standard
Expires: 18 March 2020

R. Housley
Vigil Security

18 September 2019

Use of the HSS/LMS Hash-based Signature Algorithm
in the Cryptographic Message Syntax (CMS)
<draft-ietf-lamps-cms-hash-sig-10>

Abstract

This document specifies the conventions for using the Hierarchical Signature System (HSS) / Leighton-Micali Signature (LMS) hash-based signature algorithm with the Cryptographic Message Syntax (CMS). In addition, the algorithm identifier and public key syntax are provided. The HSS/LMS algorithm is one form of hash-based digital signature; it is described in RFC 8554.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

Copyright and License Notice

Copyright (c) 2019 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

1. Introduction	3
1.1. ASN.1	3
1.2. Terminology	3
1.3. Motivation	3
2. HSS/LMS Hash-based Signature Algorithm Overview	4
2.1. Hierarchical Signature System (HSS)	4
2.2. Leighton-Micali Signature (LMS)	5
2.3. Leighton-Micali One-time Signature Algorithm (LM-OTS)	6
3. Algorithm Identifiers and Parameters	7
4. HSS/LMS Public Key Identifier	8
5. Signed-data Conventions	8
6. Security Considerations	9
7. IANA Considerations	10
8. References	10
8.1. Normative References	10
8.2. Informative References	11
Appendix: ASN.1 Module	13
Acknowledgements	14
Author's Address	14

1. Introduction

This document specifies the conventions for using the Hierarchical Signature System (HSS) / Leighton-Micali Signature (LMS) hash-based signature algorithm with the Cryptographic Message Syntax (CMS) [CMS] signed-data content type. The LMS system provides a one-time digital signature that is a variant of Merkle Tree Signatures (MTS). The HSS is built on top of the LMS system to efficiently scale for a larger numbers of signatures. The HSS/LMS algorithm is one form of hash-based digital signature, and it is described in [HASHSIG]. The HSS/LMS signature algorithm can only be used for a fixed number of signing operations with a given private key, and the number of signing operations depends upon the size of the tree. The HSS/LMS signature algorithm uses small public keys, and it has low computational cost; however, the signatures are quite large. The HSS/LMS private key can be very small when the signer is willing to perform additional computation at signing time; alternatively, the private key can consume additional memory and provide a faster signing time. The HSS/LMS signatures [HASHSIG] are currently defined to use exclusively SHA-256 [SHS].

1.1. ASN.1

CMS values are generated using ASN.1 [ASN1-B], using the Basic Encoding Rules (BER) and the Distinguished Encoding Rules (DER) [ASN1-E].

1.2. 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 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

1.3. Motivation

Recent advances in cryptanalysis [BH2013] and progress in the development of quantum computers [NAS2019] pose a threat to widely deployed digital signature algorithms. As a result, there is a need to prepare for a day that cryptosystems such as RSA and DSA that depend on discrete logarithm and factoring cannot be depended upon.

If large-scale quantum computers are ever built, these computers will be able to break many of the public-key cryptosystems currently in use. A post-quantum cryptosystem [PQC] is a system that is secure against quantum computers that have more than a trivial number of quantum bits (qubits). It is open to conjecture when it will be

feasible to build such computers; however, RSA, DSA, ECDSA, and EdDSA are all vulnerable if large-scale quantum computers come to pass.

Since the HSS/LMS signature algorithm does not depend on the difficulty of discrete logarithm or factoring, the HSS/LMS signature algorithm is considered to be post-quantum secure. One use of post-quantum secure signatures is the protection of software updates, perhaps using the format described in [FWPROT], to enable deployment of software that implements new cryptosystems.

2. HSS/LMS Hash-based Signature Algorithm Overview

Merkle Tree Signatures (MTS) are a method for signing a large but fixed number of messages. An MTS system depends on a one-time signature method and a collision-resistant hash function.

This specification makes use of the hash-based algorithm specified in [HASHSIG], which is the Leighton and Micali adaptation [LM] of the original Lamport-Diffie-Winternitz-Merkle one-time signature system [M1979] [M1987] [M1989a] [M1989b].

As implied by the name, the hash-based signature algorithm depends on a collision-resistant hash function. The hash-based signature algorithm specified in [HASHSIG] uses only the SHA-256 one-way hash function [SHS], but it establishes an IANA registry [IANA-LMS] to permit the registration of additional one-way hash functions in the future.

2.1. Hierarchical Signature System (HSS)

The MTS system specified in [HASHSIG] uses a hierarchy of trees. The Hierarchical N-time Signature System (HSS) allows subordinate trees to be generated when needed by the signer. Otherwise, generation of the entire tree might take weeks or longer.

An HSS signature as specified in [HASHSIG] carries the number of signed public keys (Nspk), followed by that number of signed public keys, followed by the LMS signature as described in Section 2.2. The public key for the top-most LMS tree is the public key of the HSS system. The LMS private key in the parent tree signs the LMS public key in the child tree, and the LMS private key in the bottom-most tree signs the actual message. The signature over the public key and the signature over the actual message are LMS signatures as described in Section 2.2.

The elements of the HSS signature value for a stand-alone tree (a top tree with no children) can be summarized as:

```
u32str(0) ||
lms_signature /* signature of message */
```

where, `u32str()` and `||` are used as defined in [HASHSIG].

The elements of the HSS signature value for a tree with `Nspk` signed public keys can be summarized as:

```
u32str(Nspk) ||
signed_public_key[0] ||
signed_public_key[1] ||
...
signed_public_key[Nspk-2] ||
signed_public_key[Nspk-1] ||
lms_signature /* signature of message */
```

where, as defined in Section 3.3 of [HASHSIG], the `signed_public_key` structure contains the `lms_signature` over the public key followed by the public key itself. Note that `Nspk` is the number of levels in the hierarchy of trees minus 1.

2.2. Leighton-Micali Signature (LMS)

Each tree in the system specified in [HASHSIG] uses the Leighton-Micali Signature (LMS) system. LMS systems have two parameters. The first parameter is the height of the tree, `h`, which is the number of levels in the tree minus one. The [HASHSIG] specification supports five values for this parameter: `h=5`; `h=10`; `h=15`; `h=20`; and `h=25`. Note that there are 2^h leaves in the tree. The second parameter, `m`, is the number of bytes output by the hash function, and it is the amount of data associated with each node in the tree. The [HASHSIG] specification supports only the SHA-256 hash function [SHS], with `m=32`. As a result, the [HASHSIG] specification supports five tree sizes; they are identified as:

```
LMS_SHA256_M32_H5;
LMS_SHA256_M32_H10;
LMS_SHA256_M32_H15;
LMS_SHA256_M32_H20; and
LMS_SHA256_M32_H25.
```

The [HASHSIG] specification establishes an IANA registry [IANA-LMS] to permit the registration of additional hash functions and additional tree sizes in the future.

As specified in [HASHSIG], the LMS public key consists of four elements: the `lms_algorithm_type` from the list above, the `otstype` to identify the LM-OTS type as discussed in Section 2.3, the private key identifier (I) as described in Section 5.3 of [HASHSIG], and the m-byte string associated with the root node of the tree (`T[1]`).

The LMS public key can be summarized as:

```
u32str(lms_algorithm_type) || u32str(otstype) || I || T[1]
```

As specified in [HASHSIG], an LMS signature consists of four elements: the number of the leaf (`q`) associated with the LM-OTS signature, an LM-OTS signature as described in Section 2.3, a typecode indicating the particular LMS algorithm, and an array of values that is associated with the path through the tree from the leaf associated with the LM-OTS signature to the root. The array of values contains the siblings of the nodes on the path from the leaf to the root but does not contain the nodes on the path itself. The array for a tree with height `h` will have `h` values. The first value is the sibling of the leaf, the next value is the sibling of the parent of the leaf, and so on up the path to the root.

The four elements of the LMS signature value can be summarized as:

```
u32str(q) ||
ots_signature ||
u32str(type) ||
path[0] || path[1] || ... || path[h-1]
```

2.3. Leighton-Micali One-time Signature Algorithm (LM-OTS)

Merkle Tree Signatures (MTS) depend on a one-time signature method, and [HASHSIG] specifies the use of the LM-OTS, which has five parameters:

- `n` - The length in bytes of the hash function output. [HASHSIG] supports only SHA-256 [SHS], with `n=32`.
- `H` - A preimage-resistant hash function that accepts byte strings of any length, and returns an `n`-byte string.
- `w` - The width in bits of the Winternitz coefficients. [HASHSIG] supports four values for this parameter: `w=1`; `w=2`; `w=4`; and `w=8`.
- `p` - The number of `n`-byte string elements that make up the LM-OTS signature.

ls - The number of bits that are left-shifted in the final step of the checksum function, which is defined in Section 4.4 of [HASHSIG].

The values of p and ls are dependent on the choices of the parameters n and w, as described in Appendix B of [HASHSIG].

The [HASHSIG] specification supports four LM-OTS variants:

```
LMOTS_SHA256_N32_W1;
LMOTS_SHA256_N32_W2;
LMOTS_SHA256_N32_W4; and
LMOTS_SHA256_N32_W8.
```

The [HASHSIG] specification establishes an IANA registry [IANA-LMS] to permit the registration of additional variants in the future.

Signing involves the generation of C, an n-byte random value.

The LM-OTS signature value can be summarized as the identifier of the LM-OTS variant, the random value, and a sequence of hash values (y[0] through y[p-1]) that correspond to the elements of the public key as described in Section 4.5 of [HASHSIG]:

```
u32str(otstype) || C || y[0] || ... || y[p-1]
```

3. Algorithm Identifiers and Parameters

The algorithm identifier for an HSS/LMS hash-based signatures is:

```
id-alg-hss-lms-hashsig OBJECT IDENTIFIER ::= { iso(1)
  member-body(2) us(840) rsadsi(113549) pkcs(1) pkcs9(9)
  smime(16) alg(3) 17 }
```

When this object identifier is used for an HSS/LMS signature, the AlgorithmIdentifier parameters field MUST be absent (that is, the parameters are not present; the parameters are not set to NULL).

The signature value is a large OCTET STRING as described in Section 2 of this document. The signature format is designed for easy parsing. The HSS, LMS, and LMOTS component of the signature value each format include a counter and a type code that indirectly provide all of the information that is needed to parse the value during signature validation.

The signature value identifies the hash function used in the HSS/LMS tree. In [HASHSIG] uses only the SHA-256 hash function [SHS], but it establishes an IANA registry [IANA-LMS] to permit the registration of

additional hash functions in the future.

4. HSS/LMS Public Key Identifier

The AlgorithmIdentifier for an HSS/LMS public key uses the id-alg-hss-lms-hashsig object identifier, and the parameters field MUST be absent.

When this AlgorithmIdentifier appears in the SubjectPublicKeyInfo field of an X.509 certificate [RFC5280], the certificate key usage extension MAY contain digitalSignature, nonRepudiation, keyCertSign, and cRLSign; however, it MUST NOT contain other values.

```
pk-HSS-LMS-HashSig PUBLIC-KEY ::= {  
    IDENTIFIER id-alg-hss-lms-hashsig  
    KEY HSS-LMS-HashSig-PublicKey  
    PARAMS ARE absent  
    CERT-KEY-USAGE  
        { digitalSignature, nonRepudiation, keyCertSign, cRLSign } }
```

```
HSS-LMS-HashSig-PublicKey ::= OCTET STRING
```

Note that the id-alg-hss-lms-hashsig algorithm identifier is also referred to as id-alg-mts-hashsig. This synonym is based on the terminology used in an early draft of the document that became [HASHSIG].

The public key value is an OCTET STRING. Like the signature format, it is designed for easy parsing. The value is the number of levels in the public key, L, followed by the LMS public key.

The HSS/LMS public key value can be described as:

```
u32str(L) || lms_public_key
```

Note that the public key for the top-most LMS tree is the public key of the HSS system. When L=1, the HSS system is a single tree.

5. Signed-data Conventions

As specified in [CMS], the digital signature is produced from the message digest and the signer's private key. The signature is computed over different values depending on whether signed attributes are absent or present.

When signed attributes are absent, the HSS/LMS signature is computed over the content. When signed attributes are present, a hash is computed over the content using the same hash function that is used

in the HSS/LMS tree, and then a message-digest attribute is constructed with the hash of the content, and then the HSS/LMS signature is computed over the DER-encoded set of signed attributes (which MUST include a content-type attribute and a message-digest attribute). In summary:

```
IF (signed attributes are absent)
THEN HSS_LMS_Sign(content)
ELSE message-digest attribute = Hash(content);
     HSS_LMS_Sign(DER(SignedAttributes))
```

When using [HASHSIG], the fields in the SignerInfo are used as follows:

digestAlgorithm MUST contain the one-way hash function used in the HSS/LMS tree. In [HASHSIG], SHA-256 is the only supported hash function, but other hash functions might be registered in the future. For convenience, the AlgorithmIdentifier for SHA-256 from [PKIXASN1] is repeated here:

```
mda-sha256 DIGEST-ALGORITHM ::= {
    IDENTIFIER id-sha256
    PARAMS TYPE NULL ARE preferredAbsent }
```

```
id-sha256 OBJECT IDENTIFIER ::= { joint-iso-itu-t(2)
    country(16) us(840) organization(1) gov(101) csor(3)
    nistAlgorithms(4) hashalgs(2) 1 }
```

signatureAlgorithm MUST contain id-alg-hss-lms-hashsig, and the algorithm parameters field MUST be absent.

signature contains the single HSS signature value resulting from the signing operation as specified in [HASHSIG].

6. Security Considerations

Implementations MUST protect the private keys. Compromise of the private keys may result in the ability to forge signatures. Along with the private key, the implementation MUST keep track of which leaf nodes in the tree have been used. Loss of integrity of this tracking data can cause a one-time key to be used more than once. As a result, when a private key and the tracking data are stored on non-volatile media or stored in a virtual machine environment, failed writes, virtual machine snapshotting or cloning, and other operational concerns must be considered to ensure confidentiality and integrity.

When generating an LMS key pair, an implementation MUST generate each

key pair independently of all other key pairs in the HSS tree.

An implementation MUST ensure that a LM-OTS private key is used to generate a signature only one time, and ensure that it cannot be used for any other purpose.

The generation of private keys relies on random numbers. The use of inadequate pseudo-random number generators (PRNGs) to generate these values can result in little or no security. An attacker may find it much easier to reproduce the PRNG environment that produced the keys, searching the resulting small set of possibilities, rather than brute force searching the whole key space. The generation of quality random numbers is difficult, and [RFC4086] offers important guidance in this area.

The generation of hash-based signatures also depends on random numbers. While the consequences of an inadequate pseudo-random number generator (PRNG) to generate these values is much less severe than in the generation of private keys, the guidance in [RFC4086] remains important.

When computing signatures, the same hash function SHOULD be used to compute the message digest of the content and the signed attributes, if they are present.

7. IANA Considerations

SMI Security for S/MIME Module Identifier (1.2.840.113549.1.9.16.0) registry, change the reference for value 64 to point to this document.

In the SMI Security for S/MIME Algorithms (1.2.840.113549.1.9.16.3) registry, change the description for value 17 to "id-alg-hss-lms-hashsig" and change the reference to point to this document.

Also, add the following note to the registry:

Value 17, "id-alg-hss-lms-hashsig", is also referred to as "id-alg-mts-hashsig".

8. References

8.1. Normative References

- [ASN1-B] ITU-T, "Information technology -- Abstract Syntax Notation One (ASN.1): Specification of basic notation", ITU-T Recommendation X.680, 2015.

- [ASN1-E] ITU-T, "Information technology -- ASN.1 encoding rules: Specification of Basic Encoding Rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER)", ITU-T Recommendation X.690, 2015.
- [CMS] Housley, R., "Cryptographic Message Syntax (CMS)", STD 70, RFC 5652, DOI 10.17487/RFC5652, September 2009, <<http://www.rfc-editor.org/info/rfc5652>>.
- [HASHSIG] McGrew, D., Curcio, M., and S. Fluhrer, "Leighton-Micali Hash-Based Signatures", RFC 8554, April 2019, <<https://rfc-editor.org/rfc/rfc8554.txt>>.
- [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>>.
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <<https://www.rfc-editor.org/info/rfc5280>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [SHS] National Institute of Standards and Technology (NIST), FIPS Publication 180-3: Secure Hash Standard, October 2008.

8.2. Informative References

- [BH2013] Ptacek, T., T. Ritter, J. Samuel, and A. Stamos, "The Factoring Dead: Preparing for the Cryptopocalypse", August 2013. <<https://media.blackhat.com/us-13/us-13-Stamos-The-Factoring-Dead.pdf>>
- [CMSASN1] Hoffman, P. and J. Schaad, "New ASN.1 Modules for Cryptographic Message Syntax (CMS) and S/MIME", RFC 5911, DOI 10.17487/RFC5911, June 2010, <<http://www.rfc-editor.org/info/rfc5911>>.

- [CMSASN1U] Schaad, J. and S. Turner, "Additional New ASN.1 Modules for the Cryptographic Message Syntax (CMS) and the Public Key Infrastructure Using X.509 (PKIX)", RFC 6268, DOI 10.17487/RFC6268, July 2011, <<http://www.rfc-editor.org/info/rfc6268>>.
- [FWPROT] Housley, R., "Using Cryptographic Message Syntax (CMS) to Protect Firmware Packages", RFC 4108, DOI 10.17487/RFC4108, August 2005, <<http://www.rfc-editor.org/info/rfc4108>>.
- [IANA-LMS] IANA Registry for Leighton-Micali Signatures (LMS). <<https://www.iana.org/assignments/leighton-micali-signatures/leighton-micali-signatures.xhtml>>.
- [LM] Leighton, T. and S. Micali, "Large provably fast and secure digital signature schemes from secure hash functions", U.S. Patent 5,432,852, July 1995.
- [M1979] Merkle, R., "Secrecy, Authentication, and Public Key Systems", Stanford University Information Systems Laboratory Technical Report 1979-1, 1979.
- [M1987] Merkle, R., "A Digital Signature Based on a Conventional Encryption Function", Lecture Notes in Computer Science crypto87, 1988.
- [M1989a] Merkle, R., "A Certified Digital Signature", Lecture Notes in Computer Science crypto89, 1990.
- [M1989b] Merkle, R., "One Way Hash Functions and DES", Lecture Notes in Computer Science crypto89, 1990.
- [NAS2019] National Academies of Sciences, Engineering, and Medicine, "Quantum Computing: Progress and Prospects", The National Academies Press, DOI 10.17226/25196, 2019.
- [PKIXASN1] Hoffman, P. and J. Schaad, "New ASN.1 Modules for the Public Key Infrastructure Using X.509 (PKIX)", RFC 5912, DOI 10.17487/RFC5912, June 2010, <<http://www.rfc-editor.org/info/rfc5912>>.
- [PQC] Bernstein, D., "Introduction to post-quantum cryptography", 2009. <http://www.pqcrypto.org/www.springer.com/cda/content/document/cda_downloaddocument/9783540887010-c1.pdf>

[RFC4086] Eastlake 3rd, D., Schiller, J., and S. Crocker,
 "Randomness Requirements for Security", BCP 106, RFC 4086,
 DOI 10.17487/RFC4086, June 2005, <<http://www.rfc-editor.org/info/rfc4086>>.

Appendix: ASN.1 Module

<CODE STARTS>

MTS-HashSig-2013

```
{ iso(1) member-body(2) us(840) rsadsi(113549) pkcs(1) pkcs9(9)
  id-smime(16) id-mod(0) id-mod-mts-hashsig-2013(64) }
```

DEFINITIONS IMPLICIT TAGS ::= BEGIN

EXPORTS ALL;

IMPORTS

```
PUBLIC-KEY, SIGNATURE-ALGORITHM, SMIME-CAPS
FROM AlgorithmInformation-2009 -- RFC 5911 [CMSASN1]
{ iso(1) identified-organization(3) dod(6) internet(1)
  security(5) mechanisms(5) pkix(7) id-mod(0)
  id-mod-algorithmInformation-02(58) } ;
```

--

-- Object Identifiers

--

```
id-alg-hss-lms-hashsig OBJECT IDENTIFIER ::= { iso(1)
  member-body(2) us(840) rsadsi(113549) pkcs(1) pkcs9(9)
  smime(16) alg(3) 17 }
```

```
id-alg-mts-hashsig OBJECT IDENTIFIER ::= id-alg-hss-lms-hashsig
```

--

-- Signature Algorithm and Public Key

--

```
sa-HSS-LMS-HashSig SIGNATURE-ALGORITHM ::= {
  IDENTIFIER id-alg-hss-lms-hashsig
  PARAMS ARE absent
  PUBLIC-KEYS { pk-HSS-LMS-HashSig }
  SMIME-CAPS { IDENTIFIED BY id-alg-hss-lms-hashsig } }
```

```
pk-HSS-LMS-HashSig PUBLIC-KEY ::= {
    IDENTIFIER id-alg-hss-lms-hashsig
    KEY HSS-LMS-HashSig-PublicKey
    PARAMS ARE absent
    CERT-KEY-USAGE
        { digitalSignature, nonRepudiation, keyCertSign, cRLSign } }

HSS-LMS-HashSig-PublicKey ::= OCTET STRING

--
-- Expand the signature algorithm set used by CMS [CMSASN1U]
--

SignatureAlgorithmSet SIGNATURE-ALGORITHM ::=
    { sa-HSS-LMS-HashSig, ... }

--
-- Expand the S/MIME capabilities set used by CMS [CMSASN1]
--

SMimeCaps SMIME-CAPS ::=
    { sa-HSS-LMS-HashSig.&smimeCaps, ... }

END

<CODE ENDS>
```

Acknowledgements

Many thanks to Scott Fluhrer, Jonathan Hammell, Ben Kaduk, Panos Kampanakis, Barry Leiba, John Mattsson, Jim Schaad, Sean Turner, Daniel Van Geest, Roman Danyliw, Dale Worley, and Joe Clarke for their careful review and comments.

Author's Address

Russ Housley
Vigil Security, LLC
516 Dranesville Road
Herndon, VA 20170
USA

EMail: housley@vigilsec.com

INTERNET-DRAFT
Internet Engineering Task Force (IETF)
Intended Status: Proposed Standard
Expires: 23 February 2020

R. Housley
Vigil Security
23 August 2019

Using Pre-Shared Key (PSK) in the Cryptographic Message Syntax (CMS)
<draft-ietf-lamps-cms-mix-with-psk-07.txt>

Abstract

The invention of a large-scale quantum computer would pose a serious challenge for the cryptographic algorithms that are widely deployed today. The Cryptographic Message Syntax (CMS) supports key transport and key agreement algorithms that could be broken by the invention of such a quantum computer. By storing communications that are protected with the CMS today, someone could decrypt them in the future when a large-scale quantum computer becomes available. Once quantum-secure key management algorithms are available, the CMS will be extended to support the new algorithms, if the existing syntax does not accommodate them. In the near-term, this document describes a mechanism to protect today's communication from the future invention of a large-scale quantum computer by mixing the output of key transport and key agreement algorithms with a pre-shared key.

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."

Copyright Notice

Copyright (c) 2019 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

1. Introduction	2
1.1. Terminology	4
1.2. ASN.1	4
1.3. Version Numbers	4
2. Overview	4
3. KeyTransPSKRecipientInfo	6
4. KeyAgreePSKRecipientInfo	7
5. Key Derivation	9
6. ASN.1 Module	10
7. Security Considerations	13
8. Privacy Considerations	15
9. IANA Considerations	15
10. References	16
10.1. Normative References	16
10.2. Informative References	16
Appendix A: Key Transport with PSK Example	17
A.1. Originator Processing Example	18
A.2. ContentInfo and AuthEnvelopedData	20
A.3. Recipient Processing Example	22
Appendix B: Key Agreement with PSK Example	23
B.1. Originator Processing Example	23
B.2. ContentInfo and AuthEnvelopedData	26
B.3. Recipient Processing Example	27
Acknowledgements	29
Author's Address	29

1. Introduction

The invention of a large-scale quantum computer would pose a serious challenge for the cryptographic algorithms that are widely deployed today [S1994]. It is an open question whether or not it is feasible to build a large-scale quantum computer, and if so, when that might happen [NAS2019]. However, if such a quantum computer is invented, many of the cryptographic algorithms and the security protocols that use them would become vulnerable.

The Cryptographic Message Syntax (CMS) [RFC5652][RFC5083] supports key transport and key agreement algorithms that could be broken by the invention of a large-scale quantum computer [C2PQ]. These algorithms include RSA [RFC8017], Diffie-Hellman [RFC2631], and Elliptic Curve Diffie-Hellman [RFC5753]. As a result, an adversary that stores CMS-protected communications today, could decrypt those communications in the future when a large-scale quantum computer becomes available.

Once quantum-secure key management algorithms are available, the CMS will be extended to support them, if the existing syntax does not already accommodate the new algorithms.

In the near-term, this document describes a mechanism to protect today's communication from the future invention of a large-scale quantum computer by mixing the output of existing key transport and key agreement algorithms with a pre-shared key (PSK). Secure communication can be achieved today by mixing a strong PSK with the output of an existing key transport algorithm, like RSA [RFC8017], or an existing key agreement algorithm, like Diffie-Hellman [RFC2631] or Elliptic Curve Diffie-Hellman [RFC5753]. A security solution that is believed to be quantum resistant can be achieved by using a PSK with sufficient entropy along with a quantum resistant key derivation function (KDF), like HKDF [RFC5869], and a quantum resistant encryption algorithm, like 256-bit AES [AES]. In this way, today's CMS-protected communication can be resistant to an attacker with a large-scale quantum computer.

In addition, there may be other reasons for including a strong PSK besides protection against the future invention of a large-scale quantum computer. For example, there is always the possibility of a cryptanalytic breakthrough on one or more of the classic public-key algorithm, and there are longstanding concerns about undisclosed trapdoors in Diffie-Hellman parameters [FGHT2016]. Inclusion of a strong PSK as part of the overall key management offer additional protection against these concerns.

Note that the CMS also supports key management techniques based on symmetric key-encryption keys and passwords, but they are not discussed in this document because they are already quantum resistant. The symmetric key-encryption key technique is quantum resistant when used with an adequate key size. The password technique is quantum resistant when used with a quantum-resistant key derivation function and a sufficiently large password.

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

1.2. ASN.1

CMS values are generated using ASN.1 [X680], which uses the Basic Encoding Rules (BER) and the Distinguished Encoding Rules (DER) [X690].

1.3. Version Numbers

The major data structures include a version number as the first item in the data structure. The version number is intended to avoid ASN.1 decode errors. Some implementations do not check the version number prior to attempting a decode, and then if a decode error occurs, the version number is checked as part of the error handling routine. This is a reasonable approach; it places error processing outside of the fast path. This approach is also forgiving when an incorrect version number is used by the sender.

Whenever the structure is updated, a higher version number will be assigned. However, to ensure maximum interoperability, the higher version number is only used when the new syntax feature is employed. That is, the lowest version number that supports the generated syntax is used.

2. Overview

The CMS enveloped-data content type [RFC5652] and the CMS authenticated-enveloped-data content type [RFC5083] support both key transport and key agreement public-key algorithms to establish the key used to encrypt the content. No restrictions are imposed on the key transport or key agreement public-key algorithms, which means that any key transport or key agreement algorithm can be used, including algorithms that are specified in the future. In both cases, the sender randomly generates the content-encryption key, and then all recipients obtain that key. All recipients use the sender-generated symmetric content-encryption key for decryption.

This specification defines two quantum-resistant ways to establish a symmetric key-encryption key, which is used to encrypt the sender-generated content-encryption key. In both cases, the PSK is used as one of the inputs to a key-derivation function to create a quantum-

resistant key-encryption key. The PSK MUST be distributed to the sender and all of the recipients by some out-of-band means that does not make it vulnerable to the future invention of a large-scale quantum computer, and an identifier MUST be assigned to the PSK. It is best if each PSK has a unique identifier; however, if a recipient has more than one PSK with the same identifier, the recipient can try each of them in turn. A PSK is expected to be used with many messages, with a lifetime of weeks or months.

The content-encryption key or content-authenticated-encryption key is quantum-resistant, and the sender establishes it using these steps:

When using a key transport algorithm:

1. The content-encryption key or the content-authenticated-encryption key, called CEK, is generated at random.
2. The key-derivation key, called KDK, is generated at random.
3. For each recipient, the KDK is encrypted in the recipient's public key, then the key derivation function (KDF) is used to mix the pre-shared key (PSK) and the KDK to produce the key-encryption key, called KEK.
4. The KEK is used to encrypt the CEK.

When using a key agreement algorithm:

1. The content-encryption key or the content-authenticated-encryption key, called CEK, is generated at random.
2. For each recipient, a pairwise key-encryption key, called KEK1, is established using the recipient's public key and the sender's private key. Note that KEK1 will be used as a key-derivation key.
3. For each recipient, the key derivation function (KDF) is used to mix the pre-shared key (PSK) and the pairwise KEK1, and the result is called KEK2.
4. For each recipient, the pairwise KEK2 is used to encrypt the CEK.

As specified in Section 6.2.5 of [RFC5652], recipient information for additional key management techniques are represented in the OtherRecipientInfo type. Two key management techniques are specified in this document, and they are each identified by a unique ASN.1 object identifier.

The first key management technique, called `keyTransPSK`, see Section 3, uses a key transport algorithm to transfer the key-derivation key from the sender to the recipient, and then the key-derivation key is mixed with the PSK using a KDF. The output of the KDF is the key-encryption key, which is used for the encryption of the content-encryption key or content-authenticated-encryption key.

The second key management technique, called `keyAgreePSK`, see Section 4, uses a key agreement algorithm to establish a pairwise key-encryption key, which is then mixed with the PSK using a KDF to produce a second pairwise key-encryption key, which is then used to encrypt the content-encryption key or content-authenticated-encryption key.

3. `keyTransPSK`

Per-recipient information using `keyTransPSK` is represented in the `KeyTransPSKRecipientInfo` type, which is indicated by the `id-ori-keyTransPSK` object identifier. Each instance of `KeyTransPSKRecipientInfo` establishes the content-encryption key or content-authenticated-encryption key for one or more recipients that have access to the same PSK.

The `id-ori-keyTransPSK` object identifier is:

```
id-ori OBJECT IDENTIFIER ::= { iso(1) member-body(2) us(840)
    rsadsi(113549) pkcs(1) pkcs-9(9) smime(16) TBD1 }

id-ori-keyTransPSK OBJECT IDENTIFIER ::= { id-ori 1 }
```

The `KeyTransPSKRecipientInfo` type is:

```
KeyTransPSKRecipientInfo ::= SEQUENCE {
    version CMSVersion, -- always set to 0
    pskid PreSharedKeyIdentifier,
    kdfAlgorithm KeyDerivationAlgorithmIdentifier,
    keyEncryptionAlgorithm KeyEncryptionAlgorithmIdentifier,
    ktris KeyTransRecipientInfos,
    encryptedKey EncryptedKey }

PreSharedKeyIdentifier ::= OCTET STRING

KeyTransRecipientInfos ::= SEQUENCE OF KeyTransRecipientInfo
```

The fields of the KeyTransPSKRecipientInfo type have the following meanings:

version is the syntax version number. The version MUST be 0. The CMSVersion type is described in Section 10.2.5 of [RFC5652].

pskid is the identifier of the PSK used by the sender. The identifier is an OCTET STRING, and it need not be human readable.

kdfAlgorithm identifies the key-derivation algorithm, and any associated parameters, used by the sender to mix the key-derivation key and the PSK to generate the key-encryption key. The KeyDerivationAlgorithmIdentifier is described in Section 10.1.6 of [RFC5652].

keyEncryptionAlgorithm identifies a key-encryption algorithm used to encrypt the content-encryption key. The KeyEncryptionAlgorithmIdentifier is described in Section 10.1.3 of [RFC5652].

ktris contains one KeyTransRecipientInfo type for each recipient; it uses a key transport algorithm to establish the key-derivation key. That is, the encryptedKey field of KeyTransRecipientInfo contains the key-derivation key instead of the content-encryption key. KeyTransRecipientInfo is described in Section 6.2.1 of [RFC5652].

encryptedKey is the result of encrypting the content-encryption key or the content-authenticated-encryption key with the key-encryption key. EncryptedKey is an OCTET STRING.

4. keyAgreePSK

Per-recipient information using keyAgreePSK is represented in the KeyAgreePSKRecipientInfo type, which is indicated by the id-ori-keyAgreePSK object identifier. Each instance of KeyAgreePSKRecipientInfo establishes the content-encryption key or content-authenticated-encryption key for one or more recipients that have access to the same PSK.

The id-ori-keyAgreePSK object identifier is:

```
id-ori-keyAgreePSK OBJECT IDENTIFIER ::= { id-ori 2 }
```

The KeyAgreePSKRecipientInfo type is:

```
KeyAgreePSKRecipientInfo ::= SEQUENCE {  
    version CMSVersion, -- always set to 0  
    pskid PreSharedKeyIdentifier,  
    originator [0] EXPLICIT OriginatorIdentifierOrKey,  
    ukm [1] EXPLICIT UserKeyingMaterial OPTIONAL,  
    kdfAlgorithm KeyDerivationAlgorithmIdentifier,  
    keyEncryptionAlgorithm KeyEncryptionAlgorithmIdentifier,  
    recipientEncryptedKeys RecipientEncryptedKeys }
```

The fields of the KeyAgreePSKRecipientInfo type have the following meanings:

version is the syntax version number. The version MUST be 0. The CMSVersion type is described in Section 10.2.5 of [RFC5652].

pskid is the identifier of the PSK used by the sender. The identifier is an OCTET STRING, and it need not be human readable.

originator is a CHOICE with three alternatives specifying the sender's key agreement public key. Implementations MUST support all three alternatives for specifying the sender's public key. The sender uses their own private key and the recipient's public key to generate a pairwise key-encryption key. A key derivation function (KDF) is used to mix the PSK and the pairwise key-encryption key to produce a second key-encryption key. The OriginatorIdentifierOrKey type is described in Section 6.2.2 of [RFC5652].

ukm is optional. With some key agreement algorithms, the sender provides a User Keying Material (UKM) to ensure that a different key is generated each time the same two parties generate a pairwise key. Implementations MUST accept a KeyAgreePSKRecipientInfo SEQUENCE that includes a ukm field. Implementations that do not support key agreement algorithms that make use of UKMs MUST gracefully handle the presence of UKMs. The UserKeyingMaterial type is described in Section 10.2.6 of [RFC5652].

kdfAlgorithm identifies the key-derivation algorithm, and any associated parameters, used by the sender to mix the pairwise key-encryption key and the PSK to produce a second key-encryption key of the same length as the first one. The KeyDerivationAlgorithmIdentifier is described in Section 10.1.6 of [RFC5652].

keyEncryptionAlgorithm identifies a key-encryption algorithm used to encrypt the content-encryption key or the content-authenticated-encryption key. The KeyEncryptionAlgorithmIdentifier type is described in Section 10.1.3 of [RFC5652].

recipientEncryptedKeys includes a recipient identifier and encrypted key for one or more recipients. The KeyAgreeRecipientIdentifier is a CHOICE with two alternatives specifying the recipient's certificate, and thereby the recipient's public key, that was used by the sender to generate a pairwise key-encryption key. The encryptedKey is the result of encrypting the content-encryption key or the content-authenticated-encryption key with the second pairwise key-encryption key. EncryptedKey is an OCTET STRING. The RecipientEncryptedKeys type is defined in Section 6.2.2 of [RFC5652].

5. Key Derivation

Many key derivation functions (KDFs) internally employ a one-way hash function. When this is the case, the hash function that is used is indirectly indicated by the KeyDerivationAlgorithmIdentifier. HKDF [RFC5869] is one example of a KDF that makes use of a hash function.

Other KDFs internally employ an encryption algorithm. When this is the case, the encryption that is used is indirectly indicated by the KeyDerivationAlgorithmIdentifier. For example, AES-128-CMAC can be used for randomness extraction in a KDF as described in [NIST2018].

A KDF has several input values. This section describes the conventions for using the KDF to compute the key-encryption key for KeyTransPSKRecipientInfo and KeyAgreePSKRecipientInfo. For simplicity, the terminology used in the HKDF [RFC5869] specification is used here.

The KDF inputs are:

IKM is the input keying material; it is the symmetric secret input to the KDF. For KeyTransPSKRecipientInfo, it is the key-derivation key. For KeyAgreePSKRecipientInfo, it is the pairwise key-encryption key produced by the key agreement algorithm.

salt is an optional non-secret random value. Many KDFs do not require a salt, and the KeyDerivationAlgorithmIdentifier assignments for HKDF [RFC8619] do not offer a parameter for a salt. If a particular KDF requires a salt, then the salt value is provided as a parameter of the KeyDerivationAlgorithmIdentifier.

L is the length of output keying material in octets; the value depends on the key-encryption algorithm that will be used. The algorithm is identified by the KeyEncryptionAlgorithmIdentifier. In addition, the OBJECT IDENTIFIER portion of the KeyEncryptionAlgorithmIdentifier is included in the next input value, called info.

info is optional context and application specific information. The DER-encoding of CMSORIforPSKOtherInfo is used as the info value, and the PSK is included in this structure. Note that EXPLICIT tagging is used in the ASN.1 module that defines this structure. For KeyTransPSKRecipientInfo, the ENUMERATED value of 5 is used. For KeyAgreePSKRecipientInfo, the ENUMERATED value of 10 is used. CMSORIforPSKOtherInfo is defined by the following ASN.1 structure:

```
CMSORIforPSKOtherInfo ::= SEQUENCE {
    psk                OCTET STRING,
    keyMgmtAlgType      ENUMERATED {
        keyTrans        (5),
        keyAgree         (10) },
    keyEncryptionAlgorithm KeyEncryptionAlgorithmIdentifier,
    pskLength           INTEGER (1..MAX),
    kdkLength           INTEGER (1..MAX) }
```

The fields of type CMSORIforPSKOtherInfo have the following meanings:

psk is an OCTET STRING; it contains the PSK.

keyMgmtAlgType is either set to 5 or 10. For KeyTransPSKRecipientInfo, the ENUMERATED value of 5 is used. For KeyAgreePSKRecipientInfo, the ENUMERATED value of 10 is used.

keyEncryptionAlgorithm is the KeyEncryptionAlgorithmIdentifier, which identifies the algorithm and provides algorithm parameters, if any.

pskLength is a positive integer; it contains the length of the PSK in octets.

kdkLength is a positive integer; it contains the length of the key-derivation key in octets. For KeyTransPSKRecipientInfo, the key-derivation key is generated by the sender. For KeyAgreePSKRecipientInfo, the key-derivation key is the pairwise key-encryption key produced by the key agreement algorithm.

The KDF output is:

OKM is the output keying material, which is exactly L octets. The OKM is the key-encryption key that is used to encrypt the content-encryption key or the content-authenticated-encryption key.

An acceptable KDF MUST accept IKM, L, and info inputs; and acceptable KDF MAY also accept salt and other inputs. All of these inputs MUST influence the output of the KDF. If the KDF requires a salt or other inputs, then those inputs MUST be provided as parameters of the KeyDerivationAlgorithmIdentifier.

6. ASN.1 Module

This section contains the ASN.1 module for the two key management techniques defined in this document. This module imports types from other ASN.1 modules that are defined in [RFC5912] and [RFC6268].

<CODE BEGINS>

CMSORIforPSK-2019

```
{ iso(1) member-body(2) us(840) rsadsi(113549) pkcs(1) pkcs-9(9)
  smime(16) modules(0) id-mod-cms-ori-psk-2019(TBD0) }
```

```
DEFINITIONS EXPLICIT TAGS ::=
BEGIN
```

```
-- EXPORTS All
```

```
IMPORTS
```

```
AlgorithmIdentifier{}, KEY-DERIVATION
```

```
FROM AlgorithmInformation-2009 -- [RFC5912]
{ iso(1) identified-organization(3) dod(6) internet(1)
  security(5) mechanisms(5) pkix(7) id-mod(0)
  id-mod-algorithmInformation-02(58) }
```

```
OTHER-RECIPIENT, OtherRecipientInfo, CMSVersion,
KeyTransRecipientInfo, OriginatorIdentifierOrKey,
UserKeyingMaterial, RecipientEncryptedKeys, EncryptedKey,
KeyDerivationAlgorithmIdentifier, KeyEncryptionAlgorithmIdentifier
FROM CryptographicMessageSyntax-2010 -- [RFC6268]
{ iso(1) member-body(2) us(840) rsadsi(113549)
  pkcs(1) pkcs-9(9) smime(16) modules(0)
  id-mod-cms-2009(58) } ;
```

```
--
-- OtherRecipientInfo Types (ori-)
--

SupportedOtherRecipInfo OTHER-RECIPIENT ::= {
    ori-keyTransPSK |
    ori-keyAgreePSK,
    ... }

--
-- Key Transport with Pre-Shared Key
--

ori-keyTransPSK OTHER-RECIPIENT ::= {
    KeyTransPSKRecipientInfo IDENTIFIED BY id-ori-keyTransPSK }

id-ori OBJECT IDENTIFIER ::= { iso(1) member-body(2) us(840)
    rsadsi(113549) pkcs(1) pkcs-9(9) smime(16) TBD1 }

id-ori-keyTransPSK OBJECT IDENTIFIER ::= { id-ori 1 }

KeyTransPSKRecipientInfo ::= SEQUENCE {
    version CMSVersion, -- always set to 0
    pskid PreSharedKeyIdentifier,
    kdfAlgorithm KeyDerivationAlgorithmIdentifier,
    keyEncryptionAlgorithm KeyEncryptionAlgorithmIdentifier,
    ktris KeyTransRecipientInfos,
    encryptedKey EncryptedKey }

PreSharedKeyIdentifier ::= OCTET STRING

KeyTransRecipientInfos ::= SEQUENCE OF KeyTransRecipientInfo

--
-- Key Agreement with Pre-Shared Key
--

ori-keyAgreePSK OTHER-RECIPIENT ::= {
    KeyAgreePSKRecipientInfo IDENTIFIED BY id-ori-keyAgreePSK }

id-ori-keyAgreePSK OBJECT IDENTIFIER ::= { id-ori 2 }
```

```

KeyAgreePSKRecipientInfo ::= SEQUENCE {
    version CMSVersion, -- always set to 0
    pskid PreSharedKeyIdentifier,
    originator [0] EXPLICIT OriginatorIdentifierOrKey,
    ukm [1] EXPLICIT UserKeyingMaterial OPTIONAL,
    kdfAlgorithm KeyDerivationAlgorithmIdentifier,
    keyEncryptionAlgorithm KeyEncryptionAlgorithmIdentifier,
    recipientEncryptedKeys RecipientEncryptedKeys }

--
-- Structure to provide 'info' input to the KDF,
-- including the Pre-Shared Key
--

CMSORIforPSKOtherInfo ::= SEQUENCE {
    psk OCTET STRING,
    keyMgmtAlgType ENUMERATED {
        keyTrans (5),
        keyAgree (10) },
    keyEncryptionAlgorithm KeyEncryptionAlgorithmIdentifier,
    pskLength INTEGER (1..MAX),
    kdkLength INTEGER (1..MAX) }

END

<CODE ENDS>

```

7. Security Considerations

The security considerations in related to the CMS enveloped-data content type in [RFC5652] and the security considerations related to the CMS authenticated-enveloped-data content type in [RFC5083] continue to apply.

Implementations of the key derivation function must compute the entire result, which in this specification is a key-encryption key, before outputting any portion of the result. The resulting key-encryption key must be protected. Compromise of the key-encryption key may result in the disclosure of all content-encryption keys or content-authenticated-encryption keys that were protected with that keying material, which in turn may result in the disclosure of the content. Note that there are two key-encryption keys when a PSK with a key agreement algorithm is used, with similar consequence for the compromise of either one of these keys.

Implementations must protect the pre-shared key (PSK), key transport private key, the agreement private key, and the key-derivation key. Compromise of the PSK will make the encrypted content vulnerable to

the future invention of a large-scale quantum computer. Compromise of the PSK and either the key transport private key or the agreement private key may result in the disclosure of all contents protected with that combination of keying material. Compromise of the PSK and the key-derivation key may result in disclosure of all contents protected with that combination of keying material.

A large-scale quantum computer will essentially negate the security provided by the key transport algorithm or the key agreement algorithm, which means that the attacker with a large-scale quantum computer can discover the key-derivation key. In addition a large-scale quantum computer effectively cuts the security provided by a symmetric key algorithm in half. Therefore, the PSK needs at least 256 bits of entropy to provide 128 bits of security. To match that same level of security, the key derivation function needs to be quantum-resistant and produce a key-encryption key that is at least 256 bits in length. Similarly, the content-encryption key or content-authenticated-encryption key needs to be at least 256 bits in length.

When using a PSK with a key transport or a key agreement algorithm, a key-encryption key is produced to encrypt the content-encryption key or content-authenticated-encryption key. If the key-encryption algorithm is different than the algorithm used to protect the content, then the effective security is determined by the weaker of the two algorithms. If, for example, content is encrypted with 256-bit AES, and the key is wrapped with 128-bit AES, then at most 128 bits of protection is provided. Implementers must ensure that the key-encryption algorithm is as strong or stronger than the content-encryption algorithm or content-authenticated-encryption algorithm.

The selection of the key-derivation function imposes an upper bound on the strength of the resulting key-encryption key. The strength of the selected key-derivation function should be at least as strong as the key-encryption algorithm that is selected. NIST SP 800-56C Revision 1 [NIST2018] offers advice on the security strength of several popular key-derivation functions.

Implementers should not mix quantum-resistant key management algorithms with their non-quantum-resistant counterparts. For example, the same content should not be protected with KeyTransRecipientInfo and KeyTransPSKRecipientInfo. Likewise, the same content should not be protected with KeyAgreeRecipientInfo and KeyAgreePSKRecipientInfo. Doing so would make the content vulnerable to the future invention of a large-scale quantum computer.

Implementers should not send the same content in different messages,

one using a quantum-resistant key management algorithm and the other using a non-quantum-resistant key management algorithm, even if the content-encryption key is generated independently. Doing so may allow an eavesdropper to correlate the messages, making the content vulnerable to the future invention of a large-scale quantum computer.

This specification does not require that PSK is known only by the sender and recipients. The PSK may be known to a group. Since confidentiality depends on the key transport or key agreement algorithm, knowledge of the PSK by other parties does not enable inherently eavesdropping. However, group members can record the traffic of other members, and then decrypt it if they ever gain access to a large-scale quantum computer. Also, when many parties know the PSK, there are many opportunities for theft of the PSK by an attacker. Once an attacker has the PSK, they can decrypt stored traffic if they ever gain access to a large-scale quantum computer in the same manner as a legitimate group member.

Sound cryptographic key hygiene is to use a key for one and only one purpose. Use of the recipient's public key for both the traditional CMS and the PSK-mixing variation specified in this document would be a violation of this principle; however, there is no known way for an attacker to take advantage of this situation. That said, an application should enforce separation whenever possible. For example, a purpose identifier for use in the X.509 extended key usage certificate extension [RFC5280] could be identified in the future to indicate that a public key should only be used in conjunction with a PSK, or only without.

Implementations must randomly generate key-derivation keys as well as the content-encryption keys or content-authenticated-encryption keys. Also, the generation of public/private key pairs for the key transport and key agreement algorithms rely on a random numbers. The use of inadequate pseudo-random number generators (PRNGs) to generate cryptographic keys can result in little or no security. An attacker may find it much easier to reproduce the PRNG environment that produced the keys, searching the resulting small set of possibilities, rather than brute force searching the whole key space. The generation of quality random numbers is difficult. [RFC4086] offers important guidance in this area.

Implementers should be aware that cryptographic algorithms become weaker with time. As new cryptanalysis techniques are developed and computing performance improves, the work factor to break a particular cryptographic algorithm will be reduced. Therefore, cryptographic algorithm implementations should be modular, allowing new algorithms to be readily inserted. That is, implementers should be prepared for the set of supported algorithms to change over time.

The security properties provided by the mechanisms specified in this document can be validated using formal methods. A ProVerif proof in [H2019] shows that an attacker with a large-scale quantum computer that is capable of breaking the Diffie-Hellman key agreement algorithm cannot disrupt the delivery of the content-encryption key to the recipient and the attacker cannot learn the content-encryption key from the protocol exchange.

8. Privacy Considerations

An observer can see which parties are using each PSK simply by watching the PSK key identifiers. However, the addition of these key identifiers is not really making privacy worse. When key transport is used, the RecipientIdentifier is always present, and it clearly identifies each recipient to an observer. When key agreement is used, either the IssuerAndSerialNumber or the RecipientKeyIdentifier is always present, and these clearly identify each recipient.

9. IANA Considerations

One object identifier for the ASN.1 module in Section 6 was assigned in the SMI Security for S/MIME Module Identifiers (1.2.840.113549.1.9.16.0) [IANA-MOD] registry:

```
id-mod-cms-ori-psk-2019 OBJECT IDENTIFIER ::= {
    iso(1) member-body(2) us(840) rsadsi(113549) pkcs(1)
    pkcs-9(9) smime(16) mod(0) TBD0 }
```

One new registry was created for Other Recipient Info Identifiers within the SMI Security for S/MIME Mail Security (1.2.840.113549.1.9.16) [IANA-SMIME] registry:

```
id-ori OBJECT IDENTIFIER ::= { iso(1) member-body(2) us(840)
    rsadsi(113549) pkcs(1) pkcs-9(9) smime(16) TBD1 }
```

Updates to the new registry are to be made according to the Specification Required policy as defined in [RFC8126]. The expert is expected to ensure that any new values identify additions RecipientInfo structures for use with the CMS. Object identifiers for other purposes should not be assigned in this arc.

Two assignments were made in the new SMI Security for Other Recipient Info Identifiers (1.2.840.113549.1.9.16.TBD1) [IANA-ORI] registry with references to this document:

```
id-ori-keyTransPSK OBJECT IDENTIFIER ::= {  
    iso(1) member-body(2) us(840) rsadsi(113549) pkcs(1)  
    pkcs-9(9) smime(16) id-ori(TBD1) 1 }
```

```
id-ori-keyAgreePSK OBJECT IDENTIFIER ::= {  
    iso(1) member-body(2) us(840) rsadsi(113549) pkcs(1)  
    pkcs-9(9) smime(16) id-ori(TBD1) 2 }
```

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5083] Housley, R., "Cryptographic Message Syntax (CMS) Authenticated-Enveloped-Data Content Type", RFC 5083, November 2007.
- [RFC5652] Housley, R., "Cryptographic Message Syntax (CMS)", RFC 5652, September 2009.
- [RFC5912] Hoffman, P., and J. Schaad, "New ASN.1 Modules for the Public Key Infrastructure Using X.509 (PKIX)", RFC 5912, June 2010.
- [RFC6268] Schaad, J., S. Turner, "Additional New ASN.1 Modules for the Cryptographic Message Syntax (CMS) and the Public Key Infrastructure Using X.509 (PKIX)", RFC 6268, July 2011.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, June 2017.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, May 2017.
- [X680] ITU-T, "Information technology -- Abstract Syntax Notation One (ASN.1): Specification of basic notation", ITU-T Recommendation X.680, 2015.

- [X690] ITU-T, "Information technology -- ASN.1 encoding rules: Specification of Basic Encoding Rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER)", ITU-T Recommendation X.690, 2015.

10.2. Informative References

- [AES] National Institute of Standards and Technology, FIPS Pub 197: Advanced Encryption Standard (AES), 26 November 2001.
- [C2PQ] Hoffman, P., "The Transition from Classical to Post-Quantum Cryptography", work-in-progress, draft-hoffman-c2pq-05, August 2018.
- [FGHT2016] Fried, J., Gaudry, P., Heninger, N., and E. Thome, "A kilobit hidden SNFS discrete logarithm computation", Cryptology ePrint Archive, Report 2016/961, 2016. <https://eprint.iacr.org/2016/961.pdf>.
- [H2019] Hammell, J., "Re: [lamps] WG Last Call for draft-ietf-lamps-cms-mix-with-psk", <https://mailarchive.ietf.org/arch/msg/spasm/_6d_4jp3sOprAnbU2fp_yp_-6-k>, 27 May 2019.
- [IANA-MOD] <https://www.iana.org/assignments/smi-numbers/smi-numbers.xhtml#security-smime-0>.
- [IANA-SMIME] <https://www.iana.org/assignments/smi-numbers/smi-numbers.xhtml#security-smime>.
- [IANA-ORI] <https://www.iana.org/assignments/smi-numbers/smi-numbers.xhtml#security-smime-TBD1>.
- [NAS2019] National Academies of Sciences, Engineering, and Medicine, "Quantum Computing: Progress and Prospects", The National Academies Press, DOI 10.17226/25196, 2019.
- [NIST2018] Barker, E., Chen, L., and R. Davis, "Recommendation for Key-Derivation Methods in Key-Establishment Schemes", NIST Special Publication 800-56C Rev. 1, April 2018, <<https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-56Cr1.pdf>>.
- [S1994] Shor, P., "Algorithms for Quantum Computation: Discrete Logarithms and Factoring", Proceedings of the 35th Annual Symposium on Foundations of Computer Science, 1994, pp. 124-134.

- [RFC2631] Rescorla, E., "Diffie-Hellman Key Agreement Method", RFC 2631, June 1999.
- [RFC4086] D. Eastlake 3rd, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security", RFC 4086, June 2005.
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, May 2008.
- [RFC5753] Turner, S., and D. Brown, "Use of Elliptic Curve Cryptography (ECC) Algorithms in Cryptographic Message Syntax (CMS)", RFC 5753, January 2010.
- [RFC5869] Krawczyk, H., and P. Eronen, "HMAC-based Extract-and-Expand Key Derivation Function (HKDF)", RFC 5869, May 2010.
- [RFC8017] Moriarty, K., Ed., Kaliski, B., Jonsson, J., and A. Rusch, "PKCS #1: RSA Cryptography Specifications Version 2.2", RFC 8017, November 2016.
- [RFC8619] Housley, R., "Algorithm Identifiers for the HMAC-based Extract-and-Expand Key Derivation Function (HKDF)", June 2019.

Appendix A: Key Transport with PSK Example

This example shows the establishment of an AES-256 content-encryption key using:

- a pre-shared key of 256 bits;
- key transport using RSA PKCS#1 v1.5 with a 3072-bit key;
- key derivation using HKDF with SHA-384; and
- key wrap using AES-256-KEYWRAP.

In real-world use, the originator would encrypt the key-derivation key in their own RSA public key as well as the recipient's public key. This is omitted in an attempt to simplify the example.

A.1. Originator Processing Example

The pre-shared key known to Alice and Bob, in hexadecimal:

c244cdd11a0d1f39d9b61282770244fb0f6befb91ab7f96cb05213365cf95b15

The identifier assigned to the pre-shared key is:
ptf-kmc:13614122112

Alice obtains Bob's public key:

-----BEGIN PUBLIC KEY-----

```
MIIBOjANBgkqhkiG9w0BAQEFAAOCAQY8AMIIBigKCAYEA3ocW14cxncPJ47fnEjBZ
AyfC2lqapL3ET4jvV6C7gGeVrRQxWPDwl+cFYBBR2ej3j3/0ecDmu+XuVi2+s5JH
Keeza+itfuhsz3yifgeEpeK8T+SusHhn20/NBLhYKbh3kiAcCgQ56dpDrDvDcLqq
vS3jg/VO+OPnZbofoH0Oevt8Q/roahJe1PlIyQ4udWB8zZezJ4mLLfbOA9YVaYXx
2AHHZJevo3nmRnlGJXo6mE00E/6qkhjDHKSMDl2WG6mO9TCDZc9qY3cAJDU6Ir0v
SH7qU18/vN13y4UOFkn8hM4kmZ6bJqbZt5NbJHtY4uQ0VMW3RyESzhrO02mrp39a
uLNhH3EXdXaVltk75H3qC7zJaeGWMJyQfOE3YfEGRKn8fxubji716D8UecAxAzFy
FL6m1JiOyV5acAiOpXN14qRYZdHnXOM9DqGIGpoeY1UuD4Mo05osOqOUpBJHA9fS
whSZG7VNf+vgNWTLNYSYLI04KiMdulnvU6ds+QPz+KKtAgMBAAE=
```

-----END PUBLIC KEY-----

Bob's RSA public key has the following key identifier:

9eeb67c9b95a74d44d2f16396680e801b5cba49c

Alice randomly generates a content-encryption key:

c8adc30f4a3e20ac420caa76a68f5787c02ab42afea20d19672fd963a5338e83

Alice randomly generates a key-derivation key:

df85af9e3cebffde6e9b9d24263db31114d0a8e33a0d50e05eb64578ccde81eb

Alice encrypts the key-derivation key in Bob's public key:

```
4e6200431ed95e0e28f7288dba56d4b90e75959e068884664c43368f3d978f3d
8179e5837e3c27bf8dc1f6e2827b9ede969be77417516de07d90e37c560add01
48deb0c9178088ccb72c068d8a9076b6a5e7ecc9093e30fdeaec9e138d80626
74fcf685f3082b910839551cd8741beedeee6e87c08ff84f03ba87118730cdf7
667002316f1a29a6cc596c7ddf95a51e398927d1916bf27929945de080fc7c80
6af6281aed6492acffa4ef1b4f53e67fca9a417db2350a2277d586ee3cabefd3
b4a44f04d3c6803d54fe9a7159210dabedda9a94f310d303331da51c0218d92a
2efb003792259195a9fd4cc403af613fdf1a6265ea70bf702fd1c6f734264c9a
59196e8e8fd657fa028e272ef741eb7711fd5b3f4ea7da9c33df66bf487da710
1c9bbfddaf1c073900a3ea99da513d8aa32605db07dc1c47504cab30c9304a85
d87377f603ec3df4056ddcf3d756fb7ed98254421a4ae151f17ad4e28c5ea077
63358dfb1ef5f73435f337b21a38c1a3fa697a530dd97e462f6b5fb2052a2d53
```

Alice produces a 256-bit key-encryption key with HKDF using SHA-384; the secret value is the key-derivation key; the 'info' is the DER-encoded CMSORIforPSKOtherInfo structure with the following values:

```
0  56: SEQUENCE {
2  32:  OCTET STRING
      :    C2 44 CD D1 1A 0D 1F 39 D9 B6 12 82 77 02 44 FB
      :    0F 6B EF B9 1A B7 F9 6C B0 52 13 36 5C F9 5B 15
36  1:  ENUMERATED 5
39 11:  SEQUENCE {
41  9:  OBJECT IDENTIFIER aes256-wrap
      :    { 2 16 840 1 101 3 4 1 45 }
      :    }
52  1:  INTEGER 32
55  1:  INTEGER 32
      :    }
```

The DER encoding of CMSORIforPSKOtherInfo produces 58 octets:

```
30380420c244cdd11a0d1f39d9b61282770244fb0f6befb91ab7f96cb0521336
5cf95b150a0105300b060960864801650304012d020120020120
```

The HKDF output is 256 bits:

```
a14d87451dfd11d83cd54ffe2bd38c49a2adfed3ac49f1d3e62bbdc64ae43b32
```

Alice uses AES-KEY-WRAP to encrypt the 256-bit content-encryption key with the key-encryption key:

```
ae4ea1d99e78fcdceal2d9f10d991ac71502939ee0c30ebdcc97dd1fc5ba3566
c83d0dd5d1b4faa5
```

Alice encrypts the content using AES-256-GCM with the content-encryption key. The 12-octet nonce used is:

```
cafebabefacedbaddecaf888
```

The content plaintext is:

```
48656c6c6f2c20776f726c6421
```

The resulting ciphertext is:

```
9af2d16f21547fcefed9b3ef2d
```

The resulting 12-octet authentication tag is:

```
a0e5925cc184e0172463c44c
```


A.2. ContentInfo and AuthEnvelopedData

Alice encodes the AuthEnvelopedData and the ContentInfo, and sends the result to Bob. The resulting structure is:

```

0 650: SEQUENCE {
4 11: OBJECT IDENTIFIER authEnvelopedData
: { 1 2 840 113549 1 9 16 1 23 }
17 633: [0] {
21 629: SEQUENCE {
25 1: INTEGER 0
28 551: SET {
32 547: [4] {
36 11: OBJECT IDENTIFIER ** Placeholder **
: { 1 2 840 113549 1 9 16 TBD 1 }
49 530: SEQUENCE {
53 1: INTEGER 0
56 19: OCTET STRING 'ptf-kmc:13614122112'
77 13: SEQUENCE {
79 11: OBJECT IDENTIFIER ** Placeholder **
: { 1 2 840 113549 1 9 16 3 TBD }
: }
92 11: SEQUENCE {
94 9: OBJECT IDENTIFIER aes256-wrap
: { 2 16 840 1 101 3 4 1 45 }
: }
105 432: SEQUENCE {
109 428: SEQUENCE {
113 1: INTEGER 2
116 20: [0]
: 9E EB 67 C9 B9 5A 74 D4 4D 2F 16 39 66 80 E8 01
: B5 CB A4 9C
138 13: SEQUENCE {
140 9: OBJECT IDENTIFIER rsaEncryption
: { 1 2 840 113549 1 1 1 }
151 0: NULL
: }
153 384: OCTET STRING
: 18 09 D6 23 17 DF 2D 09 55 57 3B FE 75 95 EB 6A
: 3D 57 84 6C 69 C1 49 0B F1 11 1A BB 40 0C D8 B5
: 26 5F D3 62 4B E2 D8 E4 CA EC 6A 12 36 CA 38 E3
: A0 7D AA E0 5F A1 E3 BC 59 F3 AD A8 8D 95 A1 6B
: 06 85 20 93 C7 C5 C0 05 62 ED DF 02 1D FE 68 7C
: 18 A1 3A AB AA 59 92 30 6A 1B 92 73 D5 01 C6 5B
: FD 1E BB A9 B9 D2 7F 48 49 7F 3C 4F 3C 13 E3 2B
: 2A 19 F1 7A CD BC 56 28 EF 7F CA 4F 69 6B 7E 92
: 66 22 0D 13 B7 23 AD 41 9E 5E 98 2A 80 B7 6C 77
: FF 9B 76 B1 04 BA 30 6D 4B 4D F9 25 57 E0 7F 0E

```

```

:      95 9A 43 6D 14 D5 72 3F AA 8F 66 35 40 D0 E3 71
:      4B 7F 20 9D ED 67 EA 33 79 CD AB 84 16 72 07 D2
:      AC 8D 3A DA 12 43 B7 2F 3A CF 91 3E F1 D9 58 20
:      6D F2 9C 09 E1 EC D2 0B 82 BE 5D 69 77 6F FE F7
:      EB F6 31 C0 D9 B7 15 BF D0 24 F3 05 1F FF 48 76
:      1D 73 17 19 2C 38 C6 D5 86 BD 67 82 2D B2 61 AA
:      08 C7 E4 37 34 D1 2D E0 51 32 15 4A AC 6B 2B 28
:      5B CD FA 7C 65 89 2F A2 63 DB AB 64 88 43 CC 66
:      27 84 29 AC 15 5F 3B 9E 5B DF 99 AE 4F 1B B2 BC
:      19 6C 17 A1 99 A5 CF F7 80 32 11 88 F1 9D B3 6F
:      4B 16 5F 3F 03 F7 D2 04 3D DE 5F 30 CD 8B BB 3A
:      38 DA 9D EC 16 6C 36 4F 8B 7E 99 AA 99 FB 42 D6
:      1A FF 3C 85 D7 A2 30 74 2C D3 AA F7 18 2A 25 3C
:      B5 02 C4 17 62 21 97 F1 E9 81 83 D0 4E BF 5B 5D
:      }
:      }
541 40:      OCTET STRING
:      AE 4E A1 D9 9E 78 FC DC EA 12 D9 F1 0D 99 1A C7
:      15 02 93 9E E0 C3 0E BD CC 97 DD 1F C5 BA 35 66
:      C8 3D 0D D5 D1 B4 FA A5
:      }
:      }
:      }
583 55:      SEQUENCE {
585 9:      OBJECT IDENTIFIER data { 1 2 840 113549 1 7 1 }
596 27:      SEQUENCE {
598 9:      OBJECT IDENTIFIER aes256-GCM
:      { 2 16 840 1 101 3 4 1 46 }
609 14:      SEQUENCE {
611 12:      OCTET STRING CA FE BA BE FA CE DB AD DE CA F8 88
:      }
:      }
625 13:      [0] 9A F2 D1 6F 21 54 7F CE FE D9 B3 EF 2D
:      }
640 12:      OCTET STRING A0 E5 92 5C C1 84 E0 17 24 63 C4 4C
:      }
:      }
:      }

```

A.3. Recipient Processing Example

Bob's private key:

```
-----BEGIN RSA PRIVATE KEY-----
MIIG5AIBAACKAYEA3ocW14cxncPJ47fnEjBZAyfc2lqapL3ET4jvV6C7gGeVrRQx
WPDwl+cFYBBR2ej3j3/0ecDmu+XuVi2+s5JHKEeza+itfuhsz3yifgeEpeK8T+Su
sHhn20/NBLhYKbh3kiAcCgQ56dpDrDvDcLqqvS3jg/VO+OPnZbofoHO0evt8Q/ro
ahJe1Pl1IyQ4udWB8zZezJ4mLLfbOA9YVaYXx2AAHZJevo3nmRnlGJXo6mE00E/6q
khjDhKSMdl2WG6mO9TCDZc9qY3cAJDU6Ir0vSH7qU18/vN13y4UOfkn8m4kmZ6b
JqbZt5NbJhtY4uQ0VMW3RyESzhrO02mrp39auLNNH3EXdXaV1tk75H3qC7zJaeGW
MJyQfOE3YfEGRKn8fxubji716D8UecAxAzFyFL6m1JiOyV5acAiOpxN14qRYZdHn
XOM9DqGIGpoeY1UuD4Mo05osOqOUpBJHA9fSwShSG7VNf+vgNWTNLNYSYLI04KiMd
ulnvU6ds+QPz+KKtAgMBAAECggGATffkSkUjjJCjLvDk4aScpSx6+Rakf2hrdS3x
jwghyUfAXgTTeUQQBs1HVtHCgxQd+qlXyn3/qu8TeZVwG4NPztyi/Z5yB1wOGJEV
3k8N/ytul6pJFFn6p48VM01bUdTrkMJbXERe6g/rr6dBQeeItCaOK7N5SIJH3Oqh
9xYuB5tH4rquCdYLmt17Tx8CaVqU9qPY3vOdQEOWIjjMV8uQUR8rHSO9Kksj8AGs
Lq9kcUPpvgJc2oqMRcNePS2WVh8xPFktRLLRazgLP8STHAtjT6S1J2UzkUqfDHGK
q/BoXxBdu6L1VDwdnIS5HXL54ElcXWsoOyKF8/ilmhRUIUWRZFmlS1ok8IC5IgX
UdL9rJVZFTRLyAwmcCEvRMlasbBrhyEyshSouN5nHJi2WVJ+wShiJeKl1qeLlpMk
HrdIYBq4Nz7/zXmiQphpAy+yQeanhP8O406C8e7RwKdpXe44su4Z8fEgA5yQx0u7
8yR1EhGKydx5bhBLR5CmlVM7rT2BAoHBAP/+e5gZLNf/ECTEBZjeiJ0Vshsz0oUq
haUQPA+9Bx9pytsoKm5oQhB7QDaxAvrn8/FUW2aAkaXsaj9F+/q30AYSQtExai9J
fdKKook3oimN8/yNRsKmhfjGOj8hd4+GjX0qoMSBCEVdT+bAjjry8wgQrqReuZnu
oXU85dmb3jvv0uIczIKvTIEyjXE5afjQIJLmZFXsBm09BG87Ia5EFUKly96BOMJh
/QWEzuYYXDqOFFfzQtkAefXNFW21Kz4Hw2QKBwQDeiGh41xCGTjECvG7fauMGlu+q
DSdYyMHif6t6mx57eS16EjvOrlXKItYhIyZW8Kw0rf/CSB2j8ig1GkMLTogrGIJ1
0322o50F0r5oOmZPueer4pOyAP0fgQ8DD1L3JBpY68/8MhYbsizVrR+Ar4jM0f96
W2bF5Xj3h+fQTDmKx6VrCCQ6miRmBUZH+ZPs5n/lyOzAYrqiKOanaiHy4mjRvlsy
mjZ6z5CG8sISqCLQ/k3Qli5pOY/v0rdBjgwAW/UCgcEAqGVYGjKdXCzuDvf9EpV4
mpTWB6yIV2ckaPon/tZi5BgsmEPwvZYzt0vMbu28Px7sSpkqUuBKbzJ4pcy8uC3I
SuYiTAhMiHS4rxIBX3BYXSudd2RD4vG1+XM0h6jVRHXHh0nOXdvfgnmigPGz3jVJ
B8oph/jD8O2Yck4YCTDOXPEi8Rjusxzro+whvRR+kG0gsGGcKSVNCPj1fNISEte4
gJId701mUAzeDjn/VaS/PXQovEMolssPPKn9NocbKbpAoHBAJnFHJun122W/lrr
ppmPnIzji30YVcYOA5vlqLKYGaAsnfYqP1WUNgfVhq2jRsrHx9cnHQI9Hu442PvI
x+c5H30YFJ4ipE3eRRRmAUi4ghY5WgD+1hw8fqyUW7E715LbSbGEUVXtrkU5G64T
UR91LEyMF8OPATdiV/KD4PWykgaqRm3tVEuCVACDTQkqNsOOi3YPQcm270w6gxfQ
SOEy/kdhCFexJFA8uZvmh6Cp2crczxyBilR/yCxCqKOONqlFdOQKBwFbJk5eHPjJz
AYueKMQESPGYCrwIqxgZGCxqaeVArHvKsEDx5whI6JWoFYVkfA8F0MyhukoEb/2x
2qB5T88Dg3EbqjTiLg3qxrWJ2OxtUo8pBP2I2wbl2N0wzcbrlYhzEZ8bJyxZu5i1
sYILC8PJ4Qzw6jS4Qpm4y1WHz8e/ElW6VyfmljZYA7f9WMntdfeQVqCVzNTvKn6f
hg6GSpJTzp4LV3ougi9nQuWXZF2wInsXkLYpsiMbL6Fz34RwohJtYA==
-----END RSA PRIVATE KEY-----
```

Bob decrypts the key-derivation key with his RSA private key:

```
df85af9e3cebffde6e9b9d24263db31114d0a8e33a0d50e05eb64578ccde81eb
```

Bob produces a 256-bit key-encryption key with HKDF using SHA-384; the secret value is the key-derivation key; the 'info' is the DER-encoded CMSORIforPSKOtherInfo structure with the same values as shown in A.1. The HKDF output is 256 bits:

a14d87451dfd11d83cd54ffe2bd38c49a2adfed3ac49f1d3e62bbdc64ae43b32

Bob uses AES-KEY-WRAP to decrypt the content-encryption key with the key-encryption key; the content-encryption key is:

c8adc30f4a3e20ac420caa76a68f5787c02ab42afea20d19672fd963a5338e83

Bob decrypts the content using AES-256-GCM with the content-encryption key, and checks the received authentication tag. The 12-octet nonce used is:

cafebabefacedbaddecaf888

The 12-octet authentication tag is:

a0e5925cc184e0172463c44c

The received ciphertext content is:

9af2d16f21547fcefed9b3ef2d

The resulting plaintext content is:

48656c6c6f2c20776f726c6421

Appendix B: Key Agreement with PSK Example

This example shows the establishment of an AES-256 content-encryption key using:

- a pre-shared key of 256 bits;
- key agreement using ECDH on curve P-384 and X9.63 KDF with SHA-384;
- key derivation using HKDF with SHA-384; and
- key wrap using AES-256-KEYWRAP.

In real-world use, the originator would treat themselves as an additional recipient by performing key agreement with their own static public key and the ephemeral private key generated for this message. This is omitted in an attempt to simplify the example.

B.1. Originator Processing Example

The pre-shared key known to Alice and Bob, in hexadecimal:

4aa53cbf500850dd583a5d9821605c6fa228fb5917f87c1c078660214e2d83e4

The identifier assigned to the pre-shared key is:

ptf-kmc:216840110121

Alice randomly generates a content-encryption key:

```
937b1219a64d57ad81c05cc86075e86017848c824d4e85800c731c5b7b091033
```

Alice obtains Bob's static ECDH public key:

```
-----BEGIN PUBLIC KEY-----
```

```
MHYwEAYHKOZIzj0CAQYFK4EEACIDYgAESCgPBO9nmUwGrgbGEOFY9HR/bCo0WyeY
/dePQVrwZmwN2yMJmO2dlkWCvLTz8U7atinxyIRe9CV54yau1KWU/wbkhPDnzuSM
YkcpXMG032z3JetEloW5aFOja13vv/W5
```

```
-----END PUBLIC KEY-----
```

It has a key identifier of:

```
e8218b98b8b7d86b5e9ebdc8aeb8c4ecdc05c529
```

Alice generates an ephemeral ECDH key pair on the same curve:

```
-----BEGIN EC PRIVATE KEY-----
```

```
MIGkAgEBBDCMiWLG44ik+L8cYVvJrQdLcFA+PwlgRF+Wt1Ab25qUh8OB7OePWjxp
/b8P6IOuI6GgBwYFK4EEACKhZANiAAQ5G0EmJk/2ks8sXY1kzbuG3Uu3ttWwQRXA
LFDJICjvYfr+yTpOQVkcHm88FAh9MEkw4NKctokKNgpsiXyrT3DtOg76oIYENpPb
GE5lJdjPx9sBsZQdABwlsU0Zb7P/7i8=
```

```
-----END EC PRIVATE KEY-----
```

Alice computes a shared secret, called Z, using the Bob's static ECDH public key and her ephemeral ECDH private key; Z is:

```
3f015ed0ff4b99523a95157bbe77e9cc0ee52fcffeb7e41eac79d1c11b6cc556
19cf8807e6d800c2de40240fe0e26adc
```

Alice computes the pairwise key-encryption key, called KEK1, from Z using the X9.63 KDF with the ECC-CMS-SharedInfo structure with the following values:

```
0  21: SEQUENCE {
2  11: SEQUENCE {
4   9: OBJECT IDENTIFIER aes256-wrap
      : { 2 16 840 1 101 3 4 1 45 }
      : }
15  6: [2] {
17  4: OCTET STRING 00 00 00 20
      : }
      : }
```

The DER encoding of ECC-CMS-SharedInfo produces 23 octets:

```
3015300b060960864801650304012da206040400000020
```

The X9.63 KDF output is the 256-bit KEK1:

```
27dc25ddb0b425f7a968ceada80a8f73c6ccaab115baafcce4a22a45d6b8f3da
```

Alice produces the 256-bit KEK2 with HKDF using SHA-384; the secret value is KEK1; the 'info' is the DER-encoded CMSORIforPSKOtherInfo structure with the following values:

```

0   56: SEQUENCE {
2   32:   OCTET STRING
      :     4A A5 3C BF 50 08 50 DD 58 3A 5D 98 21 60 5C 6F
      :     A2 28 FB 59 17 F8 7C 1C 07 86 60 21 4E 2D 83 E4
36  1:   ENUMERATED 10
39  11:  SEQUENCE {
41   9:   OBJECT IDENTIFIER aes256-wrap
      :     { 2 16 840 1 101 3 4 1 45 }
      :   }
52  1:   INTEGER 32
55  1:   INTEGER 32
      :   }

```

The DER encoding of CMSORIforPSKOtherInfo produces 58 octets:

```

303804204aa53cbf500850dd583a5d9821605c6fa228fb5917f87c1c07866021
4e2d83e40a010a300b060960864801650304012d020120020120

```

The HKDF output is the 256-bit KEK2:

```

7de693ee30ae22b5f8f6cd026c2164103f4e1430f1ab135dc1fb98954f9830bb

```

Alice uses AES-KEY-WRAP to encrypt the content-encryption key with the KEK2; the wrapped key is:

```

229fe0b45e40003e7d8244ec1b7e7ffb2c8dca16c36f5737222553a71263a92b
de08866a602d63f4

```

Alice encrypts the content using AES-256-GCM with the content-encryption key. The 12-octet nonce used is:

```

dbaddecabf888cafebabeface

```

The plaintext is:

```

48656c6c6f2c20776f726c6421

```

The resulting ciphertext is:

```

fc6d6f823e3ed2d209d0c6ffcf

```

The resulting 12-octet authentication tag is:

```

550260c42e5b29719426c1ff

```

B.2. ContentInfo and AuthEnvelopedData

Alice encodes the AuthEnvelopedData and the ContentInfo, and sends the result to Bob. The resulting structure is:

```
0 327: SEQUENCE {
4 11: OBJECT IDENTIFIER authEnvelopedData
: { 1 2 840 113549 1 9 16 1 23 }
17 310: [0] {
21 306: SEQUENCE {
25 1: INTEGER 0
28 229: SET {
31 226: [4] {
34 11: OBJECT IDENTIFIER ** Placeholder **
: { 1 2 840 113549 1 9 16 TBD 2 }
47 210: SEQUENCE {
50 1: INTEGER 0
53 20: OCTET STRING 'ptf-kmc:216840110121'
75 85: [0] {
77 83: [1] {
79 19: SEQUENCE {
81 6: OBJECT IDENTIFIER
: dhSinglePass-stdDH-sha256kdf-scheme
: { 1 3 132 1 11 1 }
89 9: OBJECT IDENTIFIER aes256-wrap
: { 2 16 840 1 101 3 4 1 45 }
: }
100 60: BIT STRING, encapsulates {
103 57: OCTET STRING
: 1B 41 26 26 4F F6 92 CF 2C 5D 8D 64 CD BB 86 DD
: 4B B7 B6 D5 B0 41 15 C0 2C 50 C9 20 28 EF 61 FA
: FE C9 3A 4E 41 59 1C 86 6F 3C 14 08 7D 30 49 30
: E0 D2 9C B6 89 0A 36 0A 6C
: }
: }
: }
162 13: SEQUENCE {
164 11: OBJECT IDENTIFIER ** Placeholder **
: { 1 2 840 113549 1 9 16 3 TBD }
: }
177 11: SEQUENCE {
179 9: OBJECT IDENTIFIER aes256-wrap
: { 2 16 840 1 101 3 4 1 45 }
: }
190 68: SEQUENCE {
192 66: SEQUENCE {
```

```

194 22:      [0] {
196 20:      OCTET STRING
      :      E8 21 8B 98 B8 B7 D8 6B 5E 9E BD C8 AE B8 C4 EC
      :      DC 05 C5 29
      :      }
218 40:      OCTET STRING
      :      22 9F E0 B4 5E 40 00 3E 7D 82 44 EC 1B 7E 7F FB
      :      2C 8D CA 16 C3 6F 57 37 22 25 53 A7 12 63 A9 2B
      :      DE 08 86 6A 60 2D 63 F4
      :      }
      :      }
      :      }
      :      }
      :      }
260 55:      SEQUENCE {
262 9:        OBJECT IDENTIFIER data { 1 2 840 113549 1 7 1 }
273 27:      SEQUENCE {
275 9:        OBJECT IDENTIFIER aes256-GCM
      :        { 2 16 840 1 101 3 4 1 46 }
286 14:      SEQUENCE {
288 12:        OCTET STRING DB AD DE CA F8 88 CA FE BA BE FA CE
      :        }
      :      }
302 13:      [0] FC 6D 6F 82 3E 3E D2 D2 09 D0 C6 FF CF
      :      }
317 12:      OCTET STRING 55 02 60 C4 2E 5B 29 71 94 26 C1 FF
      :      }
      :      }
      :      }

```

B.3. Recipient Processing Example

Bob obtains Alice's ephemeral ECDH public key from the message:

```

-----BEGIN PUBLIC KEY-----
MHYwEAYHkoZiZj0CAQYFK4EEACIDYgAEORTBJiZP9pLPLF2NZM27ht1Lt7bVseEV
wCxQySAo72H6/sk6TkFZHIzVPBQIfTBJMODSnLaJCjYKbKl8q09w7ToO+qCGBDaT
2xhOZSXYz8fbAbGUHQAcJbFNGW+z/+4v
-----END PUBLIC KEY-----

```

Bob's static ECDH private key:

```

-----BEGIN EC PRIVATE KEY-----
MIGkAgEBBDAnJ4hB+tTUN9X03/W0RsrYy+qcptlRSYkhaDIIsQYPXfTU0ugjJEmRk
NTPj4y1IRjegBwYFK4EEACKhZANiAARJwY8E72eZTAauBsYSgVj0dH9sKjRbJ5j9
149BWvBmbA3bIwmY7Z3WRYK8tPPxTtq2KfHIhF70JXnjJq7UpZT/BuSE8OfO5Ixi
RynEwajfbPcl60SWhbloU6NrXe+/9bk=
-----END EC PRIVATE KEY-----

```


Bob computes a shared secret, called Z, using the Alice's ephemeral ECDH public key and his static ECDH private key; Z is:

```
3f015ed0ff4b99523a95157bbe77e9cc0ee52fcffeb7e41eac79d1c11b6cc556
19cf8807e6d800c2de40240fe0e26adc
```

Bob computes the pairwise key-encryption key, called KEK1, from Z using the X9.63 KDF with the ECC-CMS-SharedInfo structure with the values shown in B.1. The X9.63 KDF output is the 256-bit KEK1:

```
27dc25ddb0b425f7a968ceada80a8f73c6ccaab115baafcce4a22a45d6b8f3da
```

Bob produces the 256-bit KEK2 with HKDF using SHA-384; the secret value is KEK1; the 'info' is the DER-encoded CMSORIforPSKOtherInfo structure with the values shown in B.1. The HKDF output is the 256-bit KEK2:

```
7de693ee30ae22b5f8f6cd026c2164103f4e1430f1ab135dc1fb98954f9830bb
```

Bob uses AES-KEY-WRAP to decrypt the content-encryption key with the KEK2; the content-encryption key is:

```
937b1219a64d57ad81c05cc86075e86017848c824d4e85800c731c5b7b091033
```

Bob decrypts the content using AES-256-GCM with the content-encryption key, and checks the received authentication tag. The 12-octet nonce used is:

```
dbaddecaf888cafebabeface
```

The 12-octet authentication tag is:

```
550260c42e5b29719426c1ff
```

The received ciphertext content is:

```
fc6d6f823e3ed2d209d0c6ffcf
```

The resulting plaintext content is:

```
48656c6c6f2c20776f726c6421
```

Acknowledgements

Many thanks to Roman Danyliw, Ben Kaduk, Burt Kaliski, Panos Kampanakis, Jim Schaad, Robert Sparks, Sean Turner, and Daniel Van Geest for their review and insightful comments. They have greatly improved the design, clarity, and implementation guidance.

Author's Address

Russ Housley
Vigil Security, LLC
516 Dranesville Road
Herndon, VA 20170
USA
EMail: housley@vigilsec.com

LAMPS WG
Internet-Draft
Updates: 3370 (if approved)
Intended status: Standards Track
Expires: March 19, 2020

P. Kampanakis
Cisco Systems
Q. Dang
NIST
September 16, 2019

Use of the SHAKE One-way Hash Functions in the Cryptographic Message
Syntax (CMS)
draft-ietf-lamps-cms-shakes-18

Abstract

This document updates the "Cryptographic Message Syntax Algorithms" (RFC3370) and describes the conventions for using the SHAKE family of hash functions in the Cryptographic Message Syntax as one-way hash functions with the RSA Probabilistic signature and ECDSA signature algorithms. The conventions for the associated signer public keys in CMS are also described.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on March 19, 2020.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must

include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Change Log	2
2. Introduction	5
2.1. Terminology	6
3. Identifiers	6
4. Use in CMS	7
4.1. Message Digests	7
4.2. Signatures	8
4.2.1. RSASSA-PSS Signatures	8
4.2.2. ECDSA Signatures	9
4.3. Public Keys	9
4.4. Message Authentication Codes	10
5. IANA Considerations	10
6. Security Considerations	11
7. Acknowledgements	11
8. References	11
8.1. Normative References	11
8.2. Informative References	12
Appendix A. ASN.1 Module	14
Authors' Addresses	18

1. Change Log

[EDNOTE: Remove this section before publication.]

- o draft-ietf-lamps-cms-shake-18:
 - * Minor ASN.1 changes.
- o draft-ietf-lamps-cms-shake-17:
 - * Minor updates for EDNOTE accuracy.
- o draft-ietf-lamps-cms-shake-16:
 - * Minor nits.
 - * Using bytes instead of bits for consistency.
- o draft-ietf-lamps-cms-shake-15:
 - * Minor editorial nits.

- o draft-ietf-lamps-cms-shake-14:
 - * Fixing error with incorrect preimage resistance bits for SHA128 and SHA256.
- o draft-ietf-lamps-cms-shake-13:
 - * Addressing comments from Dan M.'s secdir review.
 - * Addressing comment from Scott B.'s opsdireview about references in the abstract.
- o draft-ietf-lamps-cms-shake-12:
 - * Nits identified by Roman, Barry L. in ballot position review.
- o draft-ietf-lamps-cms-shake-11:
 - * Minor nits.
 - * Nits identified by Roman in AD Review.
- o draft-ietf-lamps-cms-shake-10:
 - * Updated IANA considerations section to request for OID assignments.
- o draft-ietf-lamps-cms-shake-09:
 - * Fixed minor text nit.
 - * Updates in Sec Considerations section.
- o draft-ietf-lamps-cms-shake-08:
 - * id-shake128-len and id-shake256-len were replaced with id-shal28 with 32 bytes output length and id-shake256 with 64 bytes output length.
 - * Fixed a discrepancy between section 3 and 4.4 about the KMAC OIDs that have parameters as optional.
- o draft-ietf-lamps-cms-shake-07:
 - * Small nit from Russ while in WGLC.
- o draft-ietf-lamps-cms-shake-06:

- * Incorporated Eric's suggestion from WGLC.
- o draft-ietf-lamps-cms-shake-05:
 - * Added informative references.
 - * Updated ASN.1 so it compiles.
 - * Updated IANA considerations.
- o draft-ietf-lamps-cms-shake-04:
 - * Added RFC8174 reference and text.
 - * Explicitly explained why RSASSA-PSS-params are omitted in section 4.2.1.
 - * Simplified Public Keys section by removing redundant info from RFCs.
- o draft-ietf-lamps-cms-shake-03:
 - * Removed paragraph suggesting KMAC to be used in generating k in Deterministic ECDSA. That should be RFC6979-bis.
 - * Removed paragraph from Security Considerations that talks about randomness of k because we are using deterministic ECDSA.
 - * Completed ASN.1 module and fixed KMAC ASN.1 based on Jim's feedback.
 - * Text fixes.
- o draft-ietf-lamps-cms-shake-02:
 - * Updates based on suggestions and clarifications by Jim.
 - * Started ASN.1 module.
- o draft-ietf-lamps-cms-shake-01:
 - * Significant reorganization of the sections to simplify the introduction, the new OIDs and their use in CMS.
 - * Added new OIDs for RSASSA-PSS that hardcodes hash, salt and MGF, according the WG consensus.

- * Updated Public Key section to use the new RSASSA-PSS OIDs and clarify the algorithm identifier usage.
- * Removed the no longer used SHAKE OIDs from section 3.1.
- o draft-ietf-lamps-cms-shake-00:
 - * Various updates to title and section names.
 - * Content changes filling in text and references.
- o draft-dang-lamps-cms-shakes-hash-00:
 - * Initial version

2. Introduction

The "Cryptographic Message Syntax (CMS)" [RFC5652] is used to digitally sign, digest, authenticate, or encrypt arbitrary message contents. "Cryptographic Message Syntax (CMS) Algorithms" [RFC3370] defines the use of common cryptographic algorithms with CMS. This specification updates RFC3370 and describes the use of the SHAKE128 and SHAKE256 specified in [SHA3] as new hash functions in CMS. In addition, it describes the use of these functions with the RSASSA-PSS signature algorithm [RFC8017] and the Elliptic Curve Digital Signature Algorithm (ECDSA) [X9.62] with the CMS signed-data content type.

In the SHA-3 family, two extendable-output functions (SHAKEs), SHAKE128 and SHAKE256, are defined. Four other hash function instances, SHA3-224, SHA3-256, SHA3-384, and SHA3-512, are also defined but are out of scope for this document. A SHAKE is a variable length hash function defined as $\text{SHAKE}(M, d)$ where the output is a d -bits-long digest of message M . The corresponding collision and second-preimage-resistance strengths for SHAKE128 are $\min(d/2, 128)$ and $\min(d, 128)$ bits, respectively (Appendix A.1 [SHA3]). And the corresponding collision and second-preimage-resistance strengths for SHAKE256 are $\min(d/2, 256)$ and $\min(d, 256)$ bits, respectively. In this specification we use $d=256$ (for SHAKE128) and $d=512$ (for SHAKE256).

A SHAKE can be used in CMS as the message digest function (to hash the message to be signed) in RSASSA-PSS and ECDSA, message authentication code and as the mask generation function (MGF) in RSASSA-PSS. This specification describes the identifiers for SHAKEs to be used in CMS and their meaning.

2.1. 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 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Identifiers

This section identifies eight new object identifiers (OIDs) for using SHAKE128 and SHAKE256 in CMS.

Two object identifiers for SHAKE128 and SHAKE256 hash functions are defined in [shake-nist-oids] and we include them here for convenience.

```
id-shake128 OBJECT IDENTIFIER ::= { joint-iso-itu-t(2)
    country(16) us(840) organization(1) gov(101) csor(3)
    nistAlgorithm(4) 2 11 }
```

```
id-shake256 OBJECT IDENTIFIER ::= { joint-iso-itu-t(2)
    country(16) us(840) organization(1) gov(101) csor(3)
    nistAlgorithm(4) 2 12 }
```

In this specification, when using the id-shake128 or id-shake256 algorithm identifiers, the parameters MUST be absent. That is, the identifier SHALL be a SEQUENCE of one component, the OID.

[I-D.ietf-lamps-pkix-shake] [EDNOTE: Update reference with the RFC when it is published.] defines two identifiers for RSASSA-PSS signatures using SHAKEs which we include here for convenience.

```
id-RSASSA-PSS-SHAKE128 OBJECT IDENTIFIER ::= { iso(1)
    identified-organization(3) dod(6) internet(1)
    security(5) mechanisms(5) pkix(7) algorithms(6) 30 }
```

```
id-RSASSA-PSS-SHAKE256 OBJECT IDENTIFIER ::= { iso(1)
    identified-organization(3) dod(6) internet(1)
    security(5) mechanisms(5) pkix(7) algorithms(6) 31 }
```

The same RSASSA-PSS algorithm identifiers can be used for identifying public keys and signatures.

[I-D.ietf-lamps-pkix-shake] [EDNOTE: Update reference with the RFC when it is published.] also defines two algorithm identifiers of ECDSA signatures using SHAKEs which we include here for convenience.


```
id-ecdsa-with-shake128 OBJECT IDENTIFIER ::= { iso(1)
    identified-organization(3) dod(6) internet(1)
    security(5) mechanisms(5) pkix(7) algorithms(6) 32 }
```

```
id-ecdsa-with-shake256 OBJECT IDENTIFIER ::= { iso(1)
    identified-organization(3) dod(6) internet(1)
    security(5) mechanisms(5) pkix(7) algorithms(6) 33 }
```

The parameters for the four RSASSA-PSS and ECDSA identifiers MUST be absent. That is, each identifier SHALL be a SEQUENCE of one component, the OID.

Two object identifiers for KMACs using SHAKE128 and SHAKE256 as defined in by the National Institute of Standards and Technology (NIST) in [shake-nist-oids] and we include them here for convenience.

```
id-KmacWithSHAKE128 OBJECT IDENTIFIER ::= { joint-iso-itu-t(2)
    country(16) us(840) organization(1) gov(101) csor(3)
    nistAlgorithm(4) 2 19 }
```

```
id-KmacWithSHAKE256 OBJECT IDENTIFIER ::= { joint-iso-itu-t(2)
    country(16) us(840) organization(1) gov(101) csor(3)
    nistAlgorithm(4) 2 20 }
```

The parameters for id-KmacWithSHAKE128 and id-KmacWithSHAKE256 are OPTIONAL.

Section 4.1, Section 4.2.1, Section 4.2.2 and Section 4.4 specify the required output length for each use of SHAKE128 or SHAKE256 in message digests, RSASSA-PSS, ECDSA and KMAC.

4. Use in CMS

4.1. Message Digests

The id-shake128 and id-shake256 OIDs (Section 3) can be used as the digest algorithm identifiers located in the SignedData, SignerInfo, DigestedData, and the AuthenticatedData digestAlgorithm fields in CMS [RFC5652]. The OID encoding MUST omit the parameters field and the output length of SHAKE128 or SHAKE256 as the message digest MUST be 32 or 64 bytes, respectively.

The digest values are located in the DigestedData field and the Message Digest authenticated attribute included in the signedAttributes of the SignedData signerInfo. In addition, digest values are input to signature algorithms. The digest algorithm MUST be the same as the message hash algorithms used in signatures.

4.2. Signatures

In CMS, signature algorithm identifiers are located in the `SignerInfo` `signatureAlgorithm` field of `SignedData` content type and countersignature attribute. Signature values are located in the `SignerInfo` signature field of `SignedData` content type and countersignature attribute.

Conforming implementations that process RSASSA-PSS and ECDSA with SHAKE signatures when processing CMS data MUST recognize the corresponding OIDs specified in Section 3.

When using RSASSA-PSS or ECDSA with SHAKEs, the RSA modulus or ECDSA curve order SHOULD be chosen in line with the SHAKE output length. Refer to Section 6 for more details.

4.2.1. RSASSA-PSS Signatures

The RSASSA-PSS algorithm is defined in [RFC8017]. When `id-RSASSA-PSS-SHAKE128` or `id-RSASSA-PSS-SHAKE256` specified in Section 3 is used, the encoding MUST omit the parameters field. That is, the `AlgorithmIdentifier` SHALL be a SEQUENCE of one component, `id-RSASSA-PSS-SHAKE128` or `id-RSASSA-PSS-SHAKE256`. [RFC4055] defines RSASSA-PSS-params that are used to define the algorithms and inputs to the algorithm. This specification does not use parameters because the hash, mask generation algorithm, trailer and salt are embedded in the OID definition.

The hash algorithm to hash a message being signed and the hash algorithm as the mask generation function used in RSASSA-PSS MUST be the same: both SHAKE128 or both SHAKE256. The output length of the hash algorithm which hashes the message SHALL be 32 (for SHAKE128) or 64 bytes (for SHAKE256).

The mask generation function takes an octet string of variable length and a desired output length as input, and outputs an octet string of the desired length. In RSASSA-PSS with SHAKEs, the SHAKEs MUST be used natively as the MGF function, instead of the MGF1 algorithm that uses the hash function in multiple iterations as specified in Section B.2.1 of [RFC8017]. In other words, the MGF is defined as the SHAKE128 or SHAKE256 with input being the `mgfSeed` for `id-RSASSA-PSS-SHAKE128` and `id-RSASSA-PSS-SHAKE256`, respectively. The `mgfSeed` is the seed from which mask is generated, an octet string [RFC8017]. As explained in Step 9 of section 9.1.1 of [RFC8017], the output length of the MGF is $\text{emLen} - \text{hLen} - 1$ bytes. `emLen` is the maximum message length $\text{ceil}((n-1)/8)$, where `n` is the RSA modulus in bits. `hLen` is 32 and 64-bytes for `id-RSASSA-PSS-SHAKE128` and `id-RSASSA-PSS-SHAKE256`, respectively. Thus when SHAKE is used as the MGF, the

SHAKE output length maskLen is $(8 * \text{emLen} - 264)$ or $(8 * \text{emLen} - 520)$ bits, respectively. For example, when RSA modulus n is 2048, the output length of SHAKE128 or SHAKE256 as the MGF will be 1784 or 1528-bits when id-RSASSA-PSS-SHAKE128 or id-RSASSA-PSS-SHAKE256 is used, respectively.

The RSASSA-PSS saltLength MUST be 32 bytes for id-RSASSA-PSS-SHAKE128 or 64 bytes for id-RSASSA-PSS-SHAKE256. Finally, the trailerField MUST be 1, which represents the trailer field with hexadecimal value 0xBC [RFC8017].

4.2.2. ECDSA Signatures

The Elliptic Curve Digital Signature Algorithm (ECDSA) is defined in [X9.62]. When the id-ecdsa-with-shake128 or id-ecdsa-with-shake256 (specified in Section 3) algorithm identifier appears, the respective SHAKE function is used as the hash. The encoding MUST omit the parameters field. That is, the AlgorithmIdentifier SHALL be a SEQUENCE of one component, the OID id-ecdsa-with-shake128 or id-ecdsa-with-shake256.

For simplicity and compliance with the ECDSA standard specification, the output length of the hash function must be explicitly determined. The output length for SHAKE128 or SHAKE256 used in ECDSA MUST be 32 or 64 bytes, respectively.

Conforming CA implementations that generate ECDSA with SHAKE signatures in certificates or CRLs SHOULD generate such signatures with a deterministically generated, non-random k in accordance with all the requirements specified in [RFC6979]. They MAY also generate such signatures in accordance with all other recommendations in [X9.62] or [SEC1] if they have a stated policy that requires conformance to those standards. Those standards have not specified SHAKE128 and SHAKE256 as hash algorithm options. However, SHAKE128 and SHAKE256 with output length being 32 and 64 octets, respectively can be used instead of 256 and 512-bit output hash algorithms such as SHA256 and SHA512.

4.3. Public Keys

In CMS, the signer's public key algorithm identifiers are located in the OriginatorPublicKey's algorithm attribute. The conventions and encoding for RSASSA-PSS and ECDSA public keys algorithm identifiers are as specified in Section 2.3 of [RFC3279], Section 3.1 of [RFC4055] and Section 2.1 of [RFC5480].

Traditionally, the rsaEncryption object identifier is used to identify RSA public keys. The rsaEncryption object identifier

continues to identify the public key when the RSA private key owner does not wish to limit the use of the public key exclusively to RSASSA-PSS with SHAKEs. When the RSA private key owner wishes to limit the use of the public key exclusively to RSASSA-PSS, the AlgorithmIdentifier for RSASSA-PSS defined in Section 3 SHOULD be used as the algorithm attribute in the OriginatorPublicKey sequence. Conforming client implementations that process RSASSA-PSS with SHAKE public keys in CMS message MUST recognize the corresponding OIDs in Section 3.

Conforming implementations MUST specify and process the algorithms explicitly by using the OIDs specified in Section 3 when encoding ECDSA with SHAKE public keys in CMS messages.

The identifier parameters, as explained in Section 3, MUST be absent.

4.4. Message Authentication Codes

KMAC message authentication code (KMAC) is specified in [SP800-185]. In CMS, KMAC algorithm identifiers are located in the AuthenticatedData macAlgorithm field. The KMAC values are located in the AuthenticatedData mac field.

When the id-KmacWithSHAKE128 or id-KmacWithSHAKE256 OID is used as the MAC algorithm identifier, the parameters field is optional (absent or present). If absent, the SHAKE256 output length used in KMAC is 32 or 64 bytes, respectively, and the customization string is an empty string by default.

Conforming implementations that process KMACs with the SHAKEs when processing CMS data MUST recognize these identifiers.

When calculating the KMAC output, the variable N is 0xD2B282C2, S is an empty string, and L, the integer representing the requested output length in bits, is 256 or 512 for KmacWithSHAKE128 or KmacWithSHAKE256, respectively, in this specification.

5. IANA Considerations

One object identifier for the ASN.1 module in Appendix A was requested for the SMI Security for S/MIME Module Identifiers (1.2.840.113549.1.9.16.0) registry:

Decimal	Description	References
70	CMSAlgsForSHAKE-2019	[EDNOTE: THIS RFC]

6. Security Considerations

This document updates [RFC3370]. The security considerations section of that document applies to this specification as well.

NIST has defined appropriate use of the hash functions in terms of the algorithm strengths and expected time frames for secure use in Special Publications (SPs) [SP800-78-4] and [SP800-107]. These documents can be used as guides to choose appropriate key sizes for various security scenarios.

SHAKE128 with output length of 32 bytes offers 128-bits of collision and preimage resistance. Thus, SHAKE128 OIDs in this specification are RECOMMENDED with 2048 (112-bit security) or 3072-bit (128-bit security) RSA modulus or curves with group order of 256-bits (128-bit security). SHAKE256 with 64 bytes output length offers 256-bits of collision and preimage resistance. Thus, the SHAKE256 OIDs in this specification are RECOMMENDED with 4096-bit RSA modulus or higher or curves with group order of at least 512 bits such as NIST Curve P-521 (256-bit security). Note that we recommended 4096-bit RSA because we would need 15360-bit modulus for 256-bits of security which is impractical for today's technology.

When more than two parties share the same message-authentication key, data origin authentication is not provided. Any party that knows the message-authentication key can compute a valid MAC, therefore the content could originate from any one of the parties.

7. Acknowledgements

This document is based on Russ Housley's draft [I-D.housley-lamps-cms-sha3-hash]. It replaces SHA3 hash functions by SHAKE128 and SHAKE256 as the LAMPS WG agreed.

The authors would like to thank Russ Housley for his guidance and very valuable contributions with the ASN.1 module. Valuable feedback was also provided by Eric Rescorla.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

- [RFC3370] Housley, R., "Cryptographic Message Syntax (CMS) Algorithms", RFC 3370, DOI 10.17487/RFC3370, August 2002, <<https://www.rfc-editor.org/info/rfc3370>>.
- [RFC4055] Schaad, J., Kaliski, B., and R. Housley, "Additional Algorithms and Identifiers for RSA Cryptography for use in the Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 4055, DOI 10.17487/RFC4055, June 2005, <<https://www.rfc-editor.org/info/rfc4055>>.
- [RFC5480] Turner, S., Brown, D., Yiu, K., Housley, R., and T. Polk, "Elliptic Curve Cryptography Subject Public Key Information", RFC 5480, DOI 10.17487/RFC5480, March 2009, <<https://www.rfc-editor.org/info/rfc5480>>.
- [RFC5652] Housley, R., "Cryptographic Message Syntax (CMS)", STD 70, RFC 5652, DOI 10.17487/RFC5652, September 2009, <<https://www.rfc-editor.org/info/rfc5652>>.
- [RFC8017] Moriarty, K., Ed., Kaliski, B., Jonsson, J., and A. Rusch, "PKCS #1: RSA Cryptography Specifications Version 2.2", RFC 8017, DOI 10.17487/RFC8017, November 2016, <<https://www.rfc-editor.org/info/rfc8017>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [SHA3] National Institute of Standards and Technology, U.S. Department of Commerce, "SHA-3 Standard - Permutation-Based Hash and Extendable-Output Functions", FIPS PUB 202, August 2015.
- [SP800-185] National Institute of Standards and Technology, "SHA-3 Derived Functions: cSHAKE, KMAC, TupleHash and ParallelHash. NIST SP 800-185", December 2016, <<http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-185.pdf>>.

8.2. Informative References

- [I-D.housley-lamps-cms-sha3-hash]
Housley, R., "Use of the SHA3 One-way Hash Functions in the Cryptographic Message Syntax (CMS)", draft-housley-lamps-cms-sha3-hash-00 (work in progress), March 2017.

- [I-D.ietf-lamps-pkix-shake]
Kampanakis, P. and Q. Dang, "Internet X.509 Public Key Infrastructure: Additional Algorithm Identifiers for RSASSA-PSS and ECDSA using SHAKEs", draft-ietf-lamps-pkix-shake-15 (work in progress), July 2019.
- [RFC3279] Bassham, L., Polk, W., and R. Housley, "Algorithms and Identifiers for the Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 3279, DOI 10.17487/RFC3279, April 2002, <<https://www.rfc-editor.org/info/rfc3279>>.
- [RFC5753] Turner, S. and D. Brown, "Use of Elliptic Curve Cryptography (ECC) Algorithms in Cryptographic Message Syntax (CMS)", RFC 5753, DOI 10.17487/RFC5753, January 2010, <<https://www.rfc-editor.org/info/rfc5753>>.
- [RFC5911] Hoffman, P. and J. Schaad, "New ASN.1 Modules for Cryptographic Message Syntax (CMS) and S/MIME", RFC 5911, DOI 10.17487/RFC5911, June 2010, <<https://www.rfc-editor.org/info/rfc5911>>.
- [RFC6268] Schaad, J. and S. Turner, "Additional New ASN.1 Modules for the Cryptographic Message Syntax (CMS) and the Public Key Infrastructure Using X.509 (PKIX)", RFC 6268, DOI 10.17487/RFC6268, July 2011, <<https://www.rfc-editor.org/info/rfc6268>>.
- [RFC6979] Pornin, T., "Deterministic Usage of the Digital Signature Algorithm (DSA) and Elliptic Curve Digital Signature Algorithm (ECDSA)", RFC 6979, DOI 10.17487/RFC6979, August 2013, <<https://www.rfc-editor.org/info/rfc6979>>.
- [SEC1] Standards for Efficient Cryptography Group, "SEC 1: Elliptic Curve Cryptography", May 2009, <<http://www.secg.org/sec1-v2.pdf>>.
- [shake-nist-oids]
National Institute of Standards and Technology, "Computer Security Objects Register", October 2017, <<https://csrc.nist.gov/Projects/Computer-Security-Objects-Register/Algorithm-Registration>>.

- [SP800-107] National Institute of Standards and Technology (NIST), "SP800-107: Recommendation for Applications Using Approved Hash Algorithms", May 2014, <https://csrc.nist.gov/csrc/media/publications/sp/800-107/rev-1/final/documents/draft_revised_sp800-107.pdf>.
- [SP800-78-4] National Institute of Standards and Technology (NIST), "SP800-78-4: Cryptographic Algorithms and Key Sizes for Personal Identity Verification", May 2014, <https://csrc.nist.gov/csrc/media/publications/sp/800-78/4/final/documents/sp800_78-4_revised_draft.pdf>.
- [X9.62] American National Standard for Financial Services (ANSI), "X9.62-2005 Public Key Cryptography for the Financial Services Industry: The Elliptic Curve Digital Signature Standard (ECDSA)", November 2005.

Appendix A. ASN.1 Module

This appendix includes the ASN.1 modules for SHAKEs in CMS. This module includes some ASN.1 from other standards for reference.

```
CMSAlgsForSHAKE-2019 { iso(1) member-body(2) us(840)
    rsadsi(113549) pkcs(1) pkcs-9(9) smime(16) modules(0)
    id-mod-cms-shakes-2019(70) }

DEFINITIONS EXPLICIT TAGS ::=

BEGIN

-- EXPORTS ALL;

IMPORTS

DIGEST-ALGORITHM, MAC-ALGORITHM, SMIME-CAPS
FROM AlgorithmInformation-2009
    { iso(1) identified-organization(3) dod(6) internet(1) security(5)
      mechanisms(5) pkix(7) id-mod(0)
      id-mod-algorithmInformation-02(58) }

RSAPublicKey, rsaEncryption, id-ecPublicKey
FROM PKIXAlgs-2009 { iso(1) identified-organization(3) dod(6)
    internet(1) security(5) mechanisms(5) pkix(7) id-mod(0)
    id-mod-pkix1-algorithms2008-02(56) }

sa-rsaspssWithSHAKE128, sa-rsaspssWithSHAKE256,
```



```
sa-ecdsaWithSHAKE128, sa-ecdsaWithSHAKE256
FROM PKIXAlgsForSHAKE-2019 {
    iso(1) identified-organization(3) dod(6)
    internet(1) security(5) mechanisms(5) pkix(7) id-mod(0)
    id-mod-pkix1-shakes-2019(94) } ;

-- Message Digest Algorithms (mda-)
-- used in SignedData, SignerInfo, DigestedData,
-- and the AuthenticatedData digestAlgorithm
-- fields in CMS
--
-- This expands MessageAuthAlgs from [RFC5652] and
-- MessageDigestAlgs in [RFC5753]
--
-- MessageDigestAlgs DIGEST-ALGORITHM ::= {
--     mda-shake128      |
--     mda-shake256,
--     ...
-- }

--
-- One-Way Hash Functions
-- SHAKE128
mda-shake128 DIGEST-ALGORITHM ::= {
    IDENTIFIER id-shake128 -- with output length 32 bytes.
}
id-shake128 OBJECT IDENTIFIER ::= { joint-iso-itu-t(2) country(16)
                                     us(840) organization(1) gov(101)
                                     csor(3) nistAlgorithm(4)
                                     hashAlgs(2) 11 }

-- SHAKE256
mda-shake256 DIGEST-ALGORITHM ::= {
    IDENTIFIER id-shake256 -- with output length 64 bytes.
}
id-shake256 OBJECT IDENTIFIER ::= { joint-iso-itu-t(2) country(16)
                                     us(840) organization(1) gov(101)
                                     csor(3) nistAlgorithm(4)
                                     hashAlgs(2) 12 }

--
-- Public key algorithm identifiers located in the
-- OriginatorPublicKey's algorithm attribute in CMS.
-- And Signature identifiers used in SignerInfo
-- signatureAlgorithm field of SignedData content
-- type and countersignature attribute in CMS.
--
-- From RFC5280, for reference.
```

```
-- rsaEncryption OBJECT IDENTIFIER ::= { pkcs-1 1 }
-- When the rsaEncryption algorithm identifier is used
-- for a public key, the AlgorithmIdentifier parameters
-- field MUST contain NULL.
--
id-RSASSA-PSS-SHAKE128 OBJECT IDENTIFIER ::= { iso(1)
    identified-organization(3) dod(6) internet(1)
    security(5) mechanisms(5) pkix(7) algorithms(6) 30 }

id-RSASSA-PSS-SHAKE256 OBJECT IDENTIFIER ::= { iso(1)
    identified-organization(3) dod(6) internet(1)
    security(5) mechanisms(5) pkix(7) algorithms(6) 31 }

-- When the id-RSASSA-PSS-* algorithm identifiers are used
-- for a public key or signature in CMS, the AlgorithmIdentifier
-- parameters field MUST be absent. The message digest algorithm
-- used in RSASSA-PSS MUST be SHAKE128 or SHAKE256 with a 32 or
-- 64 byte outout length, respectively. The mask generation
-- function MUST be SHAKE128 or SHAKE256 with an output length
-- of  $(8 * \text{ceil}((n-1)/8) - 264)$  or  $(8 * \text{ceil}((n-1)/8) - 520)$  bits,
-- respectively, where n is the RSA modulus in bits.
-- The RSASSA-PSS saltLength MUST be 32 or 64 bytes, respectively.
-- The trailerField MUST be 1, which represents the trailer
-- field with hexadecimal value 0xBC. Regardless of
-- id-RSASSA-PSS-* or rsaEncryption being used as the
-- AlgorithmIdentifier of the OriginatorPublicKey, the RSA
-- public key MUST be encoded using the RSAPublicKey type.

-- From RFC4055, for reference.
-- RSAPublicKey ::= SEQUENCE {
--     modulus INTEGER, -- -- n
--     publicExponent INTEGER } -- -- e

id-ecdsa-with-shake128 OBJECT IDENTIFIER ::= { iso(1)
    identified-organization(3) dod(6) internet(1)
    security(5) mechanisms(5) pkix(7) algorithms(6) 32 }

id-ecdsa-with-shake256 OBJECT IDENTIFIER ::= { iso(1)
    identified-organization(3) dod(6) internet(1)
    security(5) mechanisms(5) pkix(7) algorithms(6) 33 }

-- When the id-ecdsa-with-shake* algorithm identifiers are
-- used in CMS, the AlgorithmIdentifier parameters field
-- MUST be absent and the signature algorithm should be
-- deterministic ECDSA [RFC6979]. The message digest MUST
-- be SHAKE128 or SHAKE256 with a 32 or 64 byte outout
-- length, respectively. In both cases, the ECDSA public key,
-- MUST be encoded using the id-ecPublicKey type.
```

```

-- From RFC5480, for reference.
-- id-ecPublicKey OBJECT IDENTIFIER ::= {
--   iso(1) member-body(2) us(840) ansi-X9-62(10045) keyType(2) 1 }
--   -- The id-ecPublicKey parameters must be absent or present
--   -- and are defined as
-- ECPParameters ::= CHOICE {
--   namedCurve          OBJECT IDENTIFIER
--   -- -- implicitCurve   NULL
--   -- -- specifiedCurve  SpecifiedECDomain
-- }

-- This expands SignatureAlgorithms from [RFC5912]
--
-- SignatureAlgs SIGNATURE-ALGORITHM ::= {
--   sa-rsassaPssWithSHAKE128 |
--   sa-rsassaPssWithSHAKE256 |
--   sa-ecdsaWithSHAKE128 |
--   sa-ecdsaWithSHAKE256,
--   ...
-- }

-- This expands MessageAuthAlgs from [RFC5652] and [RFC6268]
--
-- Message Authentication (maca-) Algorithms
-- used in AuthenticatedData macAlgorithm in CMS
--
MessageAuthAlgs MAC-ALGORITHM ::= {
  maca-KMACwithSHAKE128 |
  maca-KMACwithSHAKE256,
  ...
}

-- This expands SMimeCaps from [RFC5911]
--
SMimeCaps SMIME-CAPS ::= {
  -- sa-rsassaPssWithSHAKE128.&smimeCaps |
  -- sa-rsassaPssWithSHAKE256.&smimeCaps |
  -- sa-ecdsaWithSHAKE128.&smimeCaps |
  -- sa-ecdsaWithSHAKE256.&smimeCaps,
  maca-KMACwithSHAKE128.&smimeCaps |
  maca-KMACwithSHAKE256.&smimeCaps,
  ...
}

--
-- KMAC with SHAKE128
maca-KMACwithSHAKE128 MAC-ALGORITHM ::= {
  IDENTIFIER id-KMACWithSHAKE128

```

```

    PARAMS TYPE KMACwithSHAKE128-params ARE optional
    -- If KMACwithSHAKE128-params parameters are absent
    -- the SHAKE128 output length used in KMAC is 256 bits
    -- and the customization string is an empty string.
    IS-KEYED-MAC TRUE
    SMIME-CAPS {IDENTIFIED BY id-KMACWithSHAKE128}
}
id-KMACWithSHAKE128 OBJECT IDENTIFIER ::= { joint-iso-itu-t(2)
    country(16) us(840) organization(1)
    gov(101) csor(3) nistAlgorithm(4)
    hashAlgs(2) 19 }
KMACwithSHAKE128-params ::= SEQUENCE {
    kMACOutputLength    INTEGER DEFAULT 256, -- Output length in bits
    customizationString OCTET STRING DEFAULT ''H
}

-- KMAC with SHAKE256
maca-KMACwithSHAKE256 MAC-ALGORITHM ::= {
    IDENTIFIER id-KMACWithSHAKE256
    PARAMS TYPE KMACwithSHAKE256-params ARE optional
    -- If KMACwithSHAKE256-params parameters are absent
    -- the SHAKE256 output length used in KMAC is 512 bits
    -- and the customization string is an empty string.
    IS-KEYED-MAC TRUE
    SMIME-CAPS {IDENTIFIED BY id-KMACWithSHAKE256}
}
id-KMACWithSHAKE256 OBJECT IDENTIFIER ::= { joint-iso-itu-t(2)
    country(16) us(840) organization(1)
    gov(101) csor(3) nistAlgorithm(4)
    hashAlgs(2) 20 }
KMACwithSHAKE256-params ::= SEQUENCE {
    kMACOutputLength    INTEGER DEFAULT 512, -- Output length in bits
    customizationString OCTET STRING DEFAULT ''H
}

END

```

Authors' Addresses

Panos Kampanakis
Cisco Systems

Email: pkampana@cisco.com

Quynh Dang
NIST
100 Bureau Drive
Gaithersburg, MD 20899

Email: quynh.Dang@nist.gov

Network Working Group
Internet-Draft
Intended status: Informational
Expires: December 31, 2019

R. Housley
Vigil Security
June 29, 2019

Hash Of Root Key Certificate Extension
draft-ietf-lamps-hash-of-root-key-cert-extn-07

Abstract

This document specifies the Hash Of Root Key certificate extension. This certificate extension is carried in the self-signed certificate for a trust anchor, which is often called a Root Certification Authority (CA) certificate. This certificate extension unambiguously identifies the next public key that will be used at some point in the future as the next Root CA certificate, eventually replacing the current one.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on December 31, 2019.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
1.1. Terminology	2
1.2. ASN.1	3
2. Overview	3
3. Hash Of Root Key Certificate Extension	4
4. IANA Considerations	4
5. Operational Considerations	4
6. Security Considerations	6
7. Acknowledgements	7
8. References	7
8.1. Normative References	8
8.2. Informative References	9
Appendix A. ASN.1 Module	9
Author's Address	11

1. Introduction

This document specifies the Hash Of Root Key X.509 version 3 certificate extension. The extension is an optional addition to the Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile [RFC5280]. The certificate extension facilitates the orderly transition from one Root Certification Authority (CA) public key to the next. It does so by publishing the hash value of the next generation public key in the current self-signed certificate. This hash value is a commitment to a particular public key in the next generation self-signed certificate. This commitment allows a relying party to unambiguously recognize the next generation self-signed certificate when it becomes available, install the new self-signed certificate in the trust anchor store, and eventually remove the previous one from the trust anchor store.

A Root CA Certificate MAY include the Hashed Root Key certificate extension to provide the hash value of the next public key that will be used by the Root CA.

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

1.2. ASN.1

Certificates [RFC5280] use ASN.1 [X680]; Distinguished Encoding Rules (DER) [X690] are REQUIRED for certificate signing and validation.

2. Overview

Before the initial deployment of the Root CA, the following are generated:

- R1 = The initial Root key pair
- R2 = The second generation Root key pair
- H2 = Thumbprint (hash) of the public key of R2
- C1 = Self-signed certificate for R1, which also contains H2

C1 is a self-signed certificate, and it contains H2 within the HashOfRootKey extension. C1 is distributed as part of the initial the system deployment. The HashOfRootKey certificate extension is described in Section 3.

When the time comes to replace the initial Root CA certificate, R1, the following are generated:

- R3 = The third generation Root key pair
- H3 = Thumbprint (hash) the public key of R3
- C2 = Self-signed certificate for R2, which contains H3

This is an iterative process. That is, R4 and H4 are generated when it is time for C3 to replace C2. And so on.

The successor to the Root CA self-signed certificate can be delivered by any means. Whenever a new Root CA self-signed certificate is received, the recipient is able to verify that the potential Root CA certificate links back to a previously authenticated Root CA certificate with the hashOfRootKey certificate extension. That is, the recipient verifies the signature on the self-signed certificate and verifies that the hash of the DER-encoded SubjectPublicKeyInfo from the potential Root CA certificate matches the value from the HashOfRootKey certificate extension of the current Root CA certificate. Checking the self-signed certificate signature ensures that the certificate contains the subject name, public key algorithm identifier, and public key algorithm parameters intended by the key owner; these are important inputs to certification path validation as defined in Section 6 of [RFC5280]. Checking the hash of the SubjectPublicKeyInfo ensures that the certificate contains the intended public key. If either check fails, then the potential Root CA certificate is not a valid replacement, and the recipient continues to use the current Root CA certificate. If both checks

succeed, then the recipient adds the potential Root CA certificate to the trust anchor store. As discussed in Section 5, the recipient can remove the current Root CA certificate immediately in some situations. In other situations, the recipient waits an appropriate amount of time to ensure that existing certification paths continue to validate.

3. Hash Of Root Key Certificate Extension

The HashOfRootKey certificate extension MUST NOT be critical.

The following ASN.1 [X680][X690] syntax defines the HashOfRootKey certificate extension:

```
ext-HashOfRootKey EXTENSION ::= {      -- Only in Root CA certificates
    SYNTAX          HashedRootKey
    IDENTIFIED BY   id-ce-hashOfRootKey
    CRITICALITY     {FALSE} }

HashedRootKey ::= SEQUENCE {
    hashAlg          HashAlgorithm,      -- Hash algorithm used
    hashValue        OCTET STRING }     -- Hash of DER-encoded
                                         -- SubjectPublicKeyInfo

id-ce-hashOfRootKey ::= OBJECT IDENTIFIER { 1 3 6 1 4 1 51483 2 1 }
```

The definitions of EXTENSION and HashAlgorithm can be found in [RFC5912].

The hashAlg indicates the one-way hash algorithm that was used to compute the hash value.

The hashValue contains the hash value computed from the next generation public key. The public key is DER-encoded SubjectPublicKeyInfo as defined in [RFC5280].

4. IANA Considerations

This document makes no requests of the IANA.

5. Operational Considerations

Guidance on the transition from one trust anchor to another is available in Section 4.4 of [RFC4210]. In particular, the oldWithNew and newWithOld advice ensures that relying parties are able to validate certificates issued under the current Root CA certificate and the next generation Root CA certificate throughout the transition. The notAfter field in the oldWithNew certificate MUST

cover the validity period of all unexpired certificates issued under the old Root CA private key. Further, this advice SHOULD be followed by Root CAs to avoid the need for all relying parties to make the transition at the same time.

After issuing the newWithOld certificate, the Root CA MUST stop using the old private key to sign certificates.

Some enterprise and application-specific environments offer a directory service or certificate repository to make certificate and CRLs available to relying parties. Section 3 in [RFC5280] describes a certificate repository. When a certificate repository is available, the oldWithNew and newWithOld certificates SHOULD be published before the successor to the current Root CA self-signed certificate is released. Recipients that are able to obtain the oldWithNew certificate SHOULD immediately remove the old Root CA self-signed certificate from the trust anchor store.

In environments without such a directory service or repository, like the Web PKI, recipients need a way to obtain the oldWithNew and newWithOld certificates. The Root CA SHOULD include the subject information access extension [RFC5280] with the accessMethod set to id-ad-caRepository and the assessLocation set to the HTTP URL that can be used to fetch a DER-encoded "certs-only" (simple PKI response) message as specified in [RFC5272] in all of their self-signed certificates. The Root CA SHOULD publish the "certs-only" message with the oldWithNew certificate and the newWithOld certificate before the subsequent Root CA self-signed certificate is released. The "certs-only" message format allows certificates to be added and removed from the bag of certificates over time, so the same HTTP URL can be used throughout the lifetime of the Root CA.

In environments without such a directory service or repository, recipients SHOULD keep both the old and replacement Root CA self-signed certificates in the trust anchor store for some amount of time to ensure that all end-entity certificates can be validated until they expire. The recipient MAY keep the old Root CA self-signed certificate until all of the certificates in the local cache that are subordinate to it have expired.

Certification path construction is more complex when the trust anchor store contains multiple self-signed certificates with the same distinguished name. For this reason, the replacement Root CA self-signed certificate SHOULD contain a different distinguished name than the one it is replacing. One approach is to include a number as part of the name that is incremented with each generation, such as "Example CA", "Example CA G2", "Example CA G3", and so on.

Changing names from one generation to another can lead to confusion when reviewing the history of a trust anchor store. To assist with such review, a recipient MAY create an audit entry to capture the old and replacement self-signed certificates.

The Root CA must securely back up the yet-to-be-deployed key pair. If the Root CA stores the key pair in a hardware security module, and that module fails, the Root CA remains committed to the key pair that is no longer available. This leaves the Root CA with no alternative but to deploy a new self-signed certificate that contains a newly-generated key pair in the same manner as the initial self-signed certificate, thus losing the benefits of the Hash Of Root Key certificate extension altogether.

6. Security Considerations

The security considerations from [RFC5280] apply, especially the discussion of self-issued certificates.

The Hash Of Root Key certificate extension facilitates the orderly transition from one Root CA public key to the next by publishing the hash value of the next generation public key in the current certificate. This allows a relying party to unambiguously recognize the next generation public key when it becomes available; however, the full public key is not disclosed until the Root CA releases the next generation certificate. In this way, attackers cannot begin to analyze the public key before the next generation Root CA self-signed certificate is released.

The Root CA needs to ensure that the public key in the next generation certificate is as strong or stronger than the key that it is replacing. Of course, a significant advance in cryptanalytic capability can break the yet-to-be-deployed key pair. Such advances are rare and difficult to predict. If such an advance occurs, the Root CA remains committed to the now broken key. This leaves the Root CA with no alternative but to deploy a new self-signed certificate that contains a newly-generated key pair, most likely using a different signature algorithm, in the same manner as the initial self-signed certificate, thus losing the benefits of the Hash Of Root Key certificate extension altogether.

The Root CA needs to employ a hash function that is resistant to preimage attacks [RFC4270]. A first-preimage attack against the hash function would allow an attacker to find another input that results published hash value. For the attack to be successful, the input would have to be a valid SubjectPublicKeyInfo that contains a public key that corresponds to a private key known to the attacker. A second-preimage attack becomes possible once the Root CA releases the

next generation public key, which makes the input to the hash function available to the attacker and everyone else. Again, the attacker needs to find a valid SubjectPublicKeyInfo that contains the public key that corresponds to a private key known to the attacker. If the employed hash function is broken after the Root CA publishes the self-signed certificate with the HashOfRootKey certificate extension, an attacker would be able to trick the recipient into installing the incorrect next generation certificate in the trust anchor store.

If an early release of the next generation public key occurs and the Root CA is concerned that attackers were given too much lead time to analyze that public key, then the Root CA can transition to a freshly generated key pair by rapidly performing two transitions. The first transition takes the Root CA to the key pair that suffered the early release, and it causes the Root CA to generate the subsequent Root key pair. The second transition occurs when the Root CA is confident that the population of relying parties have completed the first transition, and it takes the Root CA to the freshly generated key pair. Of course, the second transition also causes the Root CA to generate another key pair that is reserved for future use. Queries for the CRLs associated with certificates that are subordinate to the self-signed certificate can give some indication for the number of relying parties that are still actively using the self-signed certificates.

7. Acknowledgements

The Secure Electronic Transaction (SET) [SET] specification published by MasterCard and VISA in 1997 includes a very similar certificate extension. The SET certificate extension has essentially the same semantics, but the syntax fairly different.

CTIA - The Wireless Association - is developing a public key infrastructure that will make use of the certificate extension described in this document, and the object identifiers used in the ASN.1 module were assigned by CTIA.

Many thanks to Stefan Santesson, Jim Schaad, Daniel Kahn Gillmor, Joel Halpern, Paul Hoffman, Rich Salz, and Ben Kaduk. Their review and comments have greatly improved the document, especially the Operational Considerations and Security Considerations sections.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC4210] Adams, C., Farrell, S., Kause, T., and T. Mononen, "Internet X.509 Public Key Infrastructure Certificate Management Protocol (CMP)", RFC 4210, DOI 10.17487/RFC4210, September 2005, <<https://www.rfc-editor.org/info/rfc4210>>.
- [RFC4270] Hoffman, P. and B. Schneier, "Attacks on Cryptographic Hashes in Internet Protocols", RFC 4270, DOI 10.17487/RFC4270, November 2005, <<https://www.rfc-editor.org/info/rfc4270>>.
- [RFC5272] Schaad, J. and M. Myers, "Certificate Management over CMS (CMC)", RFC 5272, DOI 10.17487/RFC5272, June 2008, <<https://www.rfc-editor.org/info/rfc5272>>.
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <<https://www.rfc-editor.org/info/rfc5280>>.
- [RFC5912] Hoffman, P. and J. Schaad, "New ASN.1 Modules for the Public Key Infrastructure Using X.509 (PKIX)", RFC 5912, DOI 10.17487/RFC5912, June 2010, <<https://www.rfc-editor.org/info/rfc5912>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [X680] ITU-T, "Information technology -- Abstract Syntax Notation One (ASN.1): Specification of basic notation", ITU-T Recommendation X.680, 2015.
- [X690] ITU-T, "Information Technology -- ASN.1 encoding rules: Specification of Basic Encoding Rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER)", ITU-T Recommendation X.690, 2015.

8.2. Informative References

[SET] MasterCard and VISA, "SET Secure Electronic Transaction Specification -- Book 2: Programmer's Guide, Version 1.0", May 1997.

Appendix A. ASN.1 Module

The following ASN.1 module provides the complete definition of the HashOfRootKey certificate extension.

```

HashedRootKeyCertExtn { 1 3 6 1 4 1 51483 0 1 }

DEFINITIONS IMPLICIT TAGS ::=
BEGIN

-- EXPORTS All

IMPORTS

HashAlgorithm
  FROM PKIX1-PSS-OAEP-Algorithms-2009 -- [RFC5912]
    { iso(1) identified-organization(3) dod(6) internet(1)
      security(5) mechanisms(5) pkix(7) id-mod(0)
      id-mod-pkix1-rsa-pkalgs-02(54) }

EXTENSION
  FROM PKIX-CommonTypes-2009
    { iso(1) identified-organization(3) dod(6) internet(1)
      security(5) mechanisms(5) pkix(7) id-mod(0)
      id-mod-pkixCommon-02(57) } ;

--
-- Expand the certificate extensions list in [RFC5912]
--

CertExtensions EXTENSION ::= {
  ext-HashOfRootKey, ... }

--
-- HashOfRootKey Certificate Extension
--

ext-HashOfRootKey EXTENSION ::= { -- Only in Root CA certificates
  SYNTAX          HashedRootKey
  IDENTIFIED BY   id-ce-hashOfRootKey
  CRITICALITY     {FALSE} }

HashedRootKey ::= SEQUENCE {
  hashAlg      HashAlgorithm, -- Hash algorithm used
  hashValue    OCTET STRING } -- Hash of DER-encoded
                                -- SubjectPublicKeyInfo

id-ce-hashOfRootKey OBJECT IDENTIFIER ::= { 1 3 6 1 4 1 51483 2 1 }

END

```

Author's Address

Russ Housley
Vigil Security
516 Dranesville Road
Herndon, VA 20170
US

Email: housley@vigilsec.com

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 30, 2020

A. Melnikov
Isode Ltd
B. Hoeneisen
Ucom.ch
October 28, 2019

Problem Statement and Requirements for Header Protection
draft-ietf-lamps-header-protection-requirements-01

Abstract

Privacy and security issues with email header protection in S/MIME have been identified for some time. However, the desire to fix these issues has only recently been expressed in the IETF LAMPS Working Group. The existing S/MIME specification is likely to be updated regarding header protection.

This document describes the problem statement, generic use cases, and requirements of header protection.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 30, 2020.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect

to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
1.1. Requirements Language	4
1.2. Terms	4
2. Problem Statement	4
2.1. Privacy	4
2.2. Security	5
2.3. Usability	5
2.4. Interoperability	5
3. Use Cases	5
3.1. Interactions	5
3.2. Protection Levels	6
4. Requirements	7
4.1. General Requirements	7
4.1.1. Sending Side	7
4.1.2. Receiving Side	8
4.2. Additional Requirements for Backward-Compatibility With Legacy Clients Unaware of Header Protection	8
4.2.1. Sending side	8
4.2.2. Receiving side	9
4.3. Additional Requirements for Backward-Compatibility with Legacy Header Protection Systems (if supported)	9
4.3.1. Sending Side	9
4.3.2. Receiving Side	9
5. Security Considerations	9
6. Privacy Considerations	10
7. IANA Considerations	10
8. Acknowledgments	10
9. References	10
9.1. Normative References	10
9.2. Informative References	11
Appendix A. Implementation Considerations	12
A.1. Options to Achieve Header Protection	12
A.1.1. Option 1: Memory Hole	12
A.1.2. Option 2: Wrapping with message/rfc822 or message/global	12
A.1.3. Option 2.1: Progressive Header Disclosure	13
A.1.4. Examples	14
A.2. Sending Side Considerations	20
A.2.1. Candidate Header Fields for Header Protection	20
A.3. Receiving Side Considerations	21
A.3.1. Which Header Fields to Display to User	22

A.3.2. Mail User Agent Algorithm for deciding which version of a header field to display	22
Appendix B. Document Changelog	22
Appendix C. Open Issues	23
Authors' Addresses	23

1. Introduction

A range of protocols for the protection of electronic mail (email) exist, which allow to assess the authenticity and integrity of the email headers section or selected header fields (HF) from the domain-level perspective, specifically DomainKeys Identified Mail (DKIM) [RFC6376] and Sender Policy Framework (SPF) [RFC7208], and Domain-based Message Authentication, Reporting, and Conformance (DMARC) [RFC7489]. These protocols, while essential to responding to a range of attacks on email, do not offer full end-to-end protection to the header section and are not capable of providing privacy for the information contained therein.

The need for means of Data Minimization, which includes data sparseness and hiding all technically concealable information whenever possible, has grown in importance over the past several years.

A standard for end-to-end protection of the email header section exists for S/MIME version 3.1 and later. (cf. [RFC8551]):

In order to protect outer, non-content-related message header fields (for instance, the "Subject", "To", "From", and "Cc" fields), the sending client MAY wrap a full MIME message in a message/rfc822 wrapper in order to apply S/MIME security services to these header fields.

No mechanism for header protection (HP) has been standardized for PGP (Pretty Good Privacy) [RFC4880] yet.

Several varying implementations of end-to-end protections for email header sections exist, though the total number of such implementations appears to be rather low.

Some LAMPS WG participants expressed the opinion that whatever mechanism will be chosen, it should not be limited to S/MIME, but also applicable to PGP/MIME.

This document describes the problem statement (Section 2), generic use cases (Section 3) and requirements for Header Protection (Section 4). In Appendix A, possible solutions to address the challenge and some best practices are collected. In any case, the final solution is to be determined by the IETF LAMPS WG.

1.1. Requirements Language

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

1.2. Terms

The following terms are defined for the scope of this document:

- o Header Protection (HP): cryptographic protection of email Header Sections for signatures and encryption
- o Header Field (HF): cf. [RFC5322]
- o Header Section (HS): cf. [RFC5322]
- o Man-in-the-middle (MITM) attack: cf. [RFC4949], which states: "A form of active wiretapping attack in which the attacker intercepts and selectively modifies communicated data to masquerade as one or more of the entities involved in a communication association."
- o 'Signature and encryption', 'signature only' or 'encryption only' are further explained in Section 3.2.

2. Problem Statement

The LAMPS charter contains the following Work Item:

Update the specification for the cryptographic protection of email headers - both for signatures and encryption - to improve the implementation situation with respect to privacy, security, usability and interoperability in cryptographically-protected electronic mail. Most current implementations of cryptographically-protected electronic mail protect only the body of the message, which leaves significant room for attacks against otherwise-protected messages.

In the following a set of challenges to be addressed:

[[TODO: enhance this section, add more items to the following]]

2.1. Privacy

- o Data Minimization, which includes data sparseness and hiding all technically concealable information whenever possible

2.2. Security

- o MITM attacks (cf. [RFC4949])

2.3. Usability

- o User interaction / User experience

2.4. Interoperability

- o Interoperability with [RFC8551] implementations

3. Use Cases

In the following a list of the generic use cases that need to be addressed for messages with Header Protection (HP). These use cases apply independently of whether S/MIME, PGP/MIME or any other technology is used to achieve HP.

3.1. Interactions

The main interaction case for Header Protection (HP) is:

- 1) Both peers (sending and receiving side) fully support HP

For backward compatibility of legacy clients - unaware of any HP - the following intermediate interactions need to be considered as well:

- 2) The sending side fully supports HP, while the receiving side does not support any HP
- 3) The sending side does not support any HP, while the receiving side fully supports HP
- 4) Neither the sending side nor the receiving side supports any HP (trivial case)

The following intermediate use cases may need to be considered as well for backward compatibility with legacy HP systems, such as S/MIME version 3.1 and later (cf. [RFC8551]), in the following designated as legacy HP:

- 5) The sending side fully supports HP, while the receiving side supports legacy HP only
- 6) The sending side supports legacy HP only, while the receiving side fully supports HP
- 7) Both peers (sending and receiving side) support legacy HP only
- 8) The sending side supports legacy HP only, while the receiving side does not support any HP
- 9) The sending side does not support any HP, while the receiving side supports legacy HP only

Note: It is to be decided whether to ensure legacy HP systems do not conflict with any new solution for HP at all or whether (and to which degree) backward compatibility to legacy HP systems shall be maintained.

[[TODO: Decide in which form legacy HP requirements should remain in this document.]]

3.2. Protection Levels

The following protection levels need to be considered:

a) Signature and encryption

Messages containing a cryptographic signature which are also encrypted.

Sending and receiving side SHOULD implement 'signature and encryption', which is the default to use on the sending side.

b) Signature only

Messages containing a cryptographic signature, but which no encryption is applied to.

Certain implementations MAY decide to send 'signature only' messages, depending on the circumstances and customer requirements. Sending and Receiving sides SHOULD implement 'signature only'.

c) Encryption only

Messages that encryption is applied to which do not contain a cryptographic signature.

'Encryption only' is NOT RECOMMENDED on the sending side, however the receiving side needs certain guidelines on how to process received 'encrypted only' messages

4. Requirements

The following is a list of requirements that need to be addressed independently of whether S/MIME, PGP/MIME or any other technology is used to apply HP to.

4.1. General Requirements

Note: This subsection lists the requirements to address use case 1) (cf. Section 3.1).

- G1: Define the HP format for all protection levels (cf. above), which includes MIME structure, Content-Type (including all parameters, such as "charset" and "name"), Content-Disposition (including all parameters, such as "filename"), and Content-Transfer-Encoding.
- G2: To foster wide implementation of the new solution, it shall be easily implementable. Unless needed for maximizing protection and privacy, existing implementations shall not require substantial changes in the existing code base. In particular also MIME libraries widely used shall not need to be changed to comply with the new mechanism for HP.
- G3: There SHOULD be a single format that covers all protection levels (cf. above).

[[TODO: Should this one remain in the document?]]
- G4: Ensure that man-in-the-middle attack (MITM, cf. [RFC4949]), in particular downgrade attacks, are mitigated to the greatest extent possible.

4.1.1. Sending Side

- GS1: Determine which Header Fields (HFs) should or must be protected for 'signature only' emails at a minimum.
- GS2: Determine which HFs should or must be sent in clear text (i.e., included in the outer header) for emails with (signature and) encryption applied.
- GS3: Determine which HFs should not or must not be sent in clear text (i.e., not be included in the outer header) of an email with (signature and) encryption applied.
- GS4: Determine which HFs to not include to any HP part (e.g. Bcc).

4.1.2. Receiving Side

- GR1: Determine how HFs should be displayed to the user in case of conflicting information between the protected and unprotected HFs.
- GR2: Ensure that man-in-the-middle attacks (MITM, cf. [RFC4949]), in particular downgrade attacks, can be detected.
- GR3: Define how emails that 'encryption only' was applied to are to be treated.

4.2. Additional Requirements for Backward-Compatibility With Legacy Clients Unaware of Header Protection

Note: This sub-section addresses the use cases 2) - 4) (cf. Section 3.1)

- B1: Define a means to distinguish between forwarded emails and encapsulated emails using new HP mechanism.

4.2.1. Sending side

- BS1: Define how full HP support can be indicated to outgoing emails.
- BS2: Define how full HP support of the receiver can be detected or derived.
- BS3: Ensure a HP-unaware receiving side easily can display the "Subject" HF to the user.

4.2.2. Receiving side

BR1: Define how full HP support can be detected in incoming emails.

4.3. Additional Requirements for Backward-Compatibility with Legacy Header Protection Systems (if supported)

Note: This sub-section addresses the use cases 5) - 9) (cf. Section 3.1).

LS1: Depending on the solution, define a means to distinguish between forwarded emails, legacy encapsulated emails, and encapsulated emails using new HP mechanism.

LS2: The solution should be backward compatible to existing solutions and aim to minimize the implementation effort to include support for existing solutions.

4.3.1. Sending Side

LSS1: Determine how legacy HP support can be indicated to outgoing emails.

LSS2: Determine how legacy HP support of the receiver can be detected or derived.

4.3.2. Receiving Side

LSR1: Determine how legacy HP support can be detected in incoming emails.

5. Security Considerations

This document talks about UI considerations, including security considerations, when processing messages protecting Header Fields. One of the goals of this document is to specify UI for displaying such messages which is less confusing/misleading for the end-user and thus more secure.

The document does not define a new protocol, and thus does not create any new security concerns not already covered by S/MIME [RFC8551], MIME [RFC2045] and Email [RFC5322] in general.

6. Privacy Considerations

[[TODO]]

7. IANA Considerations

This document requests no action from IANA.

[[RFC Editor: This section may be removed before publication.]]

8. Acknowledgments

The authors would like to thank the following people who have provided helpful comments and suggestions for this document: David Wilson, Kelly Bristol, Robert Williams, Steve Kille, and Wei Chuang.

Essential parts of [I-D.luck-lamps-pep-header-protection] have been merged into this document. Special thanks to its author Claudio Luck. For further Acknowledgments, please refer to Acknowledgments section of [I-D.luck-lamps-pep-header-protection].

David Wilson came up with the idea of defining a new Content-Type header field parameter to distinguish forwarded messages from inner header field protection constructs.

9. References

9.1. Normative References

- [RFC2045] Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies", RFC 2045, DOI 10.17487/RFC2045, November 1996, <<https://www.rfc-editor.org/info/rfc2045>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC5322] Resnick, P., Ed., "Internet Message Format", RFC 5322, DOI 10.17487/RFC5322, October 2008, <<https://www.rfc-editor.org/info/rfc5322>>.
- [RFC8551] Schaad, J., Ramsdell, B., and S. Turner, "Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 4.0 Message Specification", RFC 8551, DOI 10.17487/RFC8551, April 2019, <<https://www.rfc-editor.org/info/rfc8551>>.

9.2. Informative References

- [I-D.luck-lamps-pep-header-protection]
Luck, C., "pretty Easy privacy (pEp): Progressive Header Disclosure", draft-luck-lamps-pep-header-protection-03 (work in progress), July 2019.
- [I-D.marques-pep-email]
Marques, H., "pretty Easy privacy (pEp): Email Formats and Protocols", draft-marques-pep-email-02 (work in progress), October 2018.
- [RFC3501] Crispin, M., "INTERNET MESSAGE ACCESS PROTOCOL - VERSION 4rev1", RFC 3501, DOI 10.17487/RFC3501, March 2003, <<https://www.rfc-editor.org/info/rfc3501>>.
- [RFC4880] Callas, J., Donnerhacke, L., Finney, H., Shaw, D., and R. Thayer, "OpenPGP Message Format", RFC 4880, DOI 10.17487/RFC4880, November 2007, <<https://www.rfc-editor.org/info/rfc4880>>.
- [RFC4949] Shirey, R., "Internet Security Glossary, Version 2", FYI 36, RFC 4949, DOI 10.17487/RFC4949, August 2007, <<https://www.rfc-editor.org/info/rfc4949>>.
- [RFC6376] Crocker, D., Ed., Hansen, T., Ed., and M. Kucherawy, Ed., "DomainKeys Identified Mail (DKIM) Signatures", STD 76, RFC 6376, DOI 10.17487/RFC6376, September 2011, <<https://www.rfc-editor.org/info/rfc6376>>.
- [RFC6532] Yang, A., Steele, S., and N. Freed, "Internationalized Email Headers", RFC 6532, DOI 10.17487/RFC6532, February 2012, <<https://www.rfc-editor.org/info/rfc6532>>.
- [RFC7208] Kitterman, S., "Sender Policy Framework (SPF) for Authorizing Use of Domains in Email, Version 1", RFC 7208, DOI 10.17487/RFC7208, April 2014, <<https://www.rfc-editor.org/info/rfc7208>>.
- [RFC7489] Kucherawy, M., Ed. and E. Zwicky, Ed., "Domain-based Message Authentication, Reporting, and Conformance (DMARC)", RFC 7489, DOI 10.17487/RFC7489, March 2015, <<https://www.rfc-editor.org/info/rfc7489>>.

Appendix A. Implementation Considerations

[[Note: Please be advised that this part of the document is early work-in-progress.]]

This Appendix A contains additional information and considerations regarding the implementation. Although not (strictly) part of the requirements, this is useful to better understand them. Parts of the text in this Appendix A will likely be moved to the upcoming implementation document.

A.1. Options to Achieve Header Protection

The following are current options for addressing Email Header Protection. The IETF LAMPS WG may choose from these options in order to update [RFC8551].

A.1.1. Option 1: Memory Hole

The Memory Hole approach works by copying the normal message header fields into the MIME header section of the top level protected body part. Since the MIME body part header section is itself covered by the protection mechanisms (signature and/or encryption) it shares the protections of the message body.

[[TODO: add more information on memory hole]]

A.1.2. Option 2: Wrapping with message/rfc822 or message/global

Wrapping with message/rfc822 (or message/global) works by copying the normal message header fields into the MIME header section of the top level protect body part

[[TODO: consider rephrasing, as not only the header fields is copied, but also the content.]]

and then prepending them with "Content-Type: message/rfc822; forwarded=no\r\n" or "Content-Type: message/global; forwarded=no\r\n", where \r\n is US-ASCII CR followed by US-ASCII LF (see also Appendix A.1.2.1). Since the MIME body part header section is itself covered by the protection mechanisms (signature and/or encryption) it shares the protections of the message body.

A.1.2.1. Content-Type Parameter "forwarded"

This section outlines how the new "forwarded" Content-Type header field parameter could be defined (probably in a separate document) and how header section wrapping works:

This document defines a new Content-Type header field parameter [RFC2045] with name "forwarded". The parameter value is case-insensitive and can be either "yes" or "no". (The default value being "yes"). The parameter is only meaningful with media type "message/rfc822" and "message/global" [RFC6532] when used within S/MIME or PGP/MIME signed or encrypted body parts. The value "yes" means that the message nested inside "message/rfc822" ("message/global") is a forwarded message and not a construct created solely to protect the inner header section.

Instructions in [RFC8551] describing how to protect the Email message header section [RFC5322], by wrapping the message inside a message/rfc822 container [RFC2045] are thus updated to read:

In order to protect outer, non-content-related message header fields (for instance, the "Subject", "To", "From", and "Cc" fields), the sending client MAY wrap a full MIME message in a message/rfc822 wrapper in order to apply S/MIME security services to these header fields. It is up to the receiving client to decide how to present this "inner" header section along with the unprotected "outer" header section.

When an S/MIME message is received, if the top-level protected MIME entity has a Content-Type of message/rfc822 or message/global without the "forwarded" parameter or with the "forwarded" parameter set to "no", it can be assumed that the intent was to provide header protection. This entity SHOULD be presented as the top-level message, taking into account header section merging issues as previously discussed.

A.1.3. Option 2.1: Progressive Header Disclosure

This option is similar to Option 2 (cf. Appendix A.1.2). It also makes use the Content-Type parameter "forwarded" (cf. Appendix A.1.2.1).

pEp for email [I-D.marques-pep-email] defines a fixed MIME structure for its innermost message structure. Security comes just next after privacy in pEp, for which reason the application of signatures without encryption to messages in transit is not considered purposeful. pEp for email, either expects to transfer messages in cleartext without signature or encryption, or transfer them encrypted and with enclosed signature and necessary public keys so that replies can be immediately upgraded to encrypted messages.

The pEp message format is equivalent to the S/MIME standard in ensuring header protection, in that the whole message is protected instead, by wrapping it and providing cryptographic services to the

whole original message. However, for the purpose of allowing the insertion of public keys, the root entity of the protected message is thus nested once more into an additional multipart/mixed MIME entity. The current pEp proposal is for PGP/MIME, while an extension to S/MIME is also on the roadmap.

pEp has also implemented the above (in Appendix A.1.2.1) described Content-Type parameter "forwarded" to distinguish between encapsulated and forwarded emails.

More information on progressive header disclosure can be found in [I-D.luck-lamps-pep-header-protection].

A.1.4. Examples

Examples in subsequent sections assume that an email client is trying to protect (sign) the following initial message:

Date: Mon, 25 Sep 2017 17:31:42 +0100 (GMT Daylight Time)
From: "Alexey Melnikov" <alexey.melnikov@example.net>
Message-ID: <e4a483cb-1dfb-481d-903b-298c92c21f5e@matt.example.net>
MIME-Version: 1.0
MMHS-Primary-Precedence: 3
Subject: Meeting at my place
To: somebody@example.net
X-Mailer: Isode Harrier Web Server
Content-Type: text/plain; charset=us-ascii

This is an important message that I don't want to be modified.

Without message header protection the corresponding signed message might look like this. (Lines prepended by "O: " are the outer header.)

O: Date: Mon, 25 Sep 2017 17:31:42 +0100 (GMT Daylight Time)
O: Message-ID: <e4a483cb-1dfb-481d-903b-298c92c21f5e@matt.example.net>
O: Subject: Meeting at my place
O: From: "Alexey Melnikov" <alexey.melnikov@example.net>
O: MIME-Version: 1.0
O: Content-Type: multipart/signed; charset=us-ascii; micalg=sha1;
O: protocol="application/pkcs7-signature";
O: boundary=.cbel6d2a-e1a3-4220-b821-38348fc97237

This is a multipart message in MIME format.
--.cbel6d2a-e1a3-4220-b821-38348fc97237
Content-Type: text/plain; charset=us-ascii

This is an important message that I don't want to be modified.

--.cbel6d2a-e1a3-4220-b821-38348fc97237
Content-Transfer-Encoding: base64
Content-Type: application/pkcs7-signature

[[base-64 encoded signature]]

--.cbel6d2a-e1a3-4220-b821-38348fc97237--

A.1.4.1. Option 1: Memory Hole

The following example demonstrates how header section and payload of a protect body part might look like. For example, this will be the first body part of a multipart/signed message or the signed and/or encrypted payload of the application/pkcs7-mime body part. Lines

prepended by "O: " are the outer header section. Lines prepended by "I: " are the inner header section.

```
O: Date: Mon, 25 Sep 2017 17:31:42 +0100 (GMT Daylight Time)
O: Message-ID: <e4a483cb-1dfb-481d-903b-298c92c21f5e@matt.example.net>
O: Subject: Meeting at my place
O: From: "Alexey Melnikov" <alexey.melnikov@example.net>
O: MIME-Version: 1.0
O: Content-Type: multipart/signed; charset=us-ascii; micalg=sha1;
O: protocol="application/pkcs7-signature";
O: boundary=.cbe16d2a-e1a3-4220-b821-38348fc97237
```

This is a multipart message in MIME format.

--.cbe16d2a-e1a3-4220-b821-38348fc97237

```
I: Date: Mon, 25 Sep 2017 17:31:42 +0100 (GMT Daylight Time)
I: From: "Alexey Melnikov" <alexey.melnikov@example.net>
I: Message-ID: <e4a483cb-1dfb-481d-903b-298c92c21f5e@matt.example.net>
I: MIME-Version: 1.0
I: MMHS-Primary-Precedence: 3
I: Subject: Meeting at my place
I: To: somebody@example.net
I: X-Mailer: Isode Harrier Web Server
I: Content-Type: text/plain; charset=us-ascii
```

This is an important message that I don't want to be modified.

--.cbe16d2a-e1a3-4220-b821-38348fc97237

Content-Transfer-Encoding: base64

Content-Type: application/pkcs7-signature

[[base-64 encoded signature]]

--.cbe16d2a-e1a3-4220-b821-38348fc97237--

[[TODO (AM): HB: Not sure whether the Outer Subject HF is replaced by "Encrypted Message" (or alike). Please verify.]]

A.1.4.2. Option 2: Wrapping with message/rfc822 or message/global

The following example demonstrates how header section and payload of a protect body part might look like. For example, this will be the first body part of a multipart/signed message or the signed and/or encrypted payload of the application/pkcs7-mime body part. Lines prepended by "O: " are the outer header section. Lines prepended by "I: " are the inner header section. Lines prepended by "W: " are the wrapper.


```
O: Date: Mon, 25 Sep 2017 17:31:42 +0100 (GMT Daylight Time)
O: Message-ID: <e4a483cb-1dfb-481d-903b-298c92c21f5e@matt.example.net>
O: Subject: Meeting at my place
O: From: "Alexey Melnikov" <alexey.melnikov@example.net>
O: MIME-Version: 1.0
O: Content-Type: multipart/signed; charset=us-ascii; micalg=sha1;
O:   protocol="application/pkcs7-signature";
O:   boundary=.cbel6d2a-ela3-4220-b821-38348fc97237
```

This is a multipart message in MIME format.

--.cbel6d2a-ela3-4220-b821-38348fc97237

W: Content-Type: message/rfc822; forwarded=no

W:

```
I: Date: Mon, 25 Sep 2017 17:31:42 +0100 (GMT Daylight Time)
I: From: "Alexey Melnikov" <alexey.melnikov@example.net>
I: Message-ID: <e4a483cb-1dfb-481d-903b-298c92c21f5e@matt.example.net>
I: MIME-Version: 1.0
I: MMHS-Primary-Precedence: 3
I: Subject: Meeting at my place
I: To: somebody@example.net
I: X-Mailer: Isode Harrier Web Server
I: Content-Type: text/plain; charset=us-ascii
```

This is an important message that I don't want to be modified.

--.cbel6d2a-ela3-4220-b821-38348fc97237

Content-Transfer-Encoding: base64

Content-Type: application/pkcs7-signature

[[base-64 encoded signature]]

--.cbel6d2a-ela3-4220-b821-38348fc97237--

A.1.4.3. Option 2.1 Progressive Header Disclosure

The following example demonstrates how header section and payload of a protect body part might look like. For example, this will be the first body part of a multipart/signed message or the signed and encrypted payload of the application/pkcs7-mime body part. Lines prepended by "O: " are the outer header section. Lines prepended by "I: " are the inner header section. Lines prepended by "W: " are the wrapper.

The main difference compared to Option 2 is an additional multipart/mixed Content-Type containing the original message (as a whole) and the public key (of the sender).

Note: This example is derived from the pEp's PGP/MIME implementation and adjusted to the above S/MIME examples. The pEp implementations do not support S/MIME yet; therefore the following can serve no more as for illustrative purpose. Specific examples can be found in [I-D.luck-lamps-pep-header-protection].

O: Date: Mon, 25 Sep 2017 17:31:42 +0100 (GMT Daylight Time)
O: Message-ID: <e4a483cb-1dfb-481d-903b-298c92c21f5e@matt.example.net>
O: Subject: Meeting at my place
O: From: "Alexey Melnikov" <alexey.melnikov@example.net>
W: MIME-Version: 1.0
W: Content-Type: multipart/mixed;
W: boundary="6b8b4567327b23c6643c986966334873"
W:
W: --6b8b4567327b23c6643c986966334873
W: Content-Type: message/rfc822; forwarded="no"
W:
I: Date: Mon, 25 Sep 2017 17:31:42 +0100 (GMT Daylight Time)
I: Message-ID: <e4a483cb-1dfb-481d-903b-298c92c21f5e@matt.example.net>
I: Subject: Meeting at my place
I: From: "Alexey Melnikov" <alexey.melnikov@example.net>
I: MIME-Version: 1.0
I: Content-Type: multipart/signed; charset=us-ascii; micalg=sha1;
I: protocol="application/pkcs7-signature";
I: boundary=.cbe16d2a-e1a3-4220-b821-38348fc97237

This is a multipart message in MIME format.
--.cbe16d2a-e1a3-4220-b821-38348fc97237
Content-Type: text/plain; charset=us-ascii

This is an important message that I don't want to be modified.

--.cbe16d2a-e1a3-4220-b821-38348fc97237
Content-Transfer-Encoding: base64
Content-Type: application/pkcs7-signature

[[base-64 encoded signature]]

--.cbe16d2a-e1a3-4220-b821-38348fc97237--
W: --6b8b4567327b23c6643c986966334873
W: Content-Type: application/pgp-keys
W: Content-Disposition: attachment; filename="pEpkey.asc"
W:
-----BEGIN PGP PUBLIC KEY BLOCK-----
...
-----END PGP PUBLIC KEY BLOCK-----
W:
W: --6b8b4567327b23c6643c986966334873--

A.2. Sending Side Considerations

A.2.1. Candidate Header Fields for Header Protection

[[TODO: This section is very early stage and needs more work.]]

For a 'signature only' (cf. Section 3.2) message, it is RECOMMENDED that all "outer" header fields are identical to the "inner" protected header fields. This would mean that all header fields are signed. In this case, the "outer" header fields simply match the protected header fields. And in the case that the "outer" header fields differ, they can simply be replaced with their protected versions when displayed to the user.

[[TODO: Decide whether "Bcc" header field should be excluded. Also verify whether this requirement applies generally or just for specific implementations.]]

When generating S/MIME messages with applied (signature and) encryption to protect header fields:

1. If a header field is being encrypted because it is sensitive, its true value MUST NOT be included in the outer header. If the header field is mandatory according to [RFC5322], a stub value (or a value indicating that the outer value is not to be used) is to be included in the outer header section.
2. The outer header section SHOULD be minimal in order to avoid disclosure of confidential information. It is recommended that the outer header section only contains "Date" (set to the same value as in the inner header field, or, if the Date value is also sensitive, to Monday 9am of the same week), possibly "Subject" and "To"/"Cc" header fields. ("From", "Date", and at least one destination header field is mandatory as per [RFC5322].) In particular, Keywords, In-Reply-To and References header fields SHOULD NOT be included in the outer header; "To" and "Cc" header fields should be omitted and replaced with "Bcc: undisclosed-recipients;".

But note that having key header fields duplicated in the outer header is convenient for many message stores (e.g. IMAP) and clients that can't decode S/MIME encrypted messages. In particular, Subject/To/Cc/Bcc/Date header field values are returned in IMAP ENVELOPE FETCH data item [RFC3501], which is frequently used by IMAP clients in order to avoid parsing message header.

3. The "Subject" header field value of the outer header section SHOULD either be identical to the inner "Subject" header field value, or contain a clear indication that the outer value is not to be used for display (the inner header field value would contain the true value).

Note that recommendations listed above typically only apply to non MIME header fields (header fields with names not starting with "Content-" prefix), but there are exceptions, e.g. Content-Language.

Note that the above recommendations can also negatively affect anti-spam processing.

Messages containing at least one recipient address in the Bcc header field may appear in up to three different variants:

1. The message for the recipient addresses listed in To or Cc header fields, which must not include the Bcc header field neither for signature calculation nor for encryption.
2. The message(s) sent to the recipient addresses in the Bcc header field, which depends on the implementation:
 - a) One message for each recipient in the Bcc header field separately with a Bcc header field containing only the address of the recipient it is sent to
 - b) The same message for each recipient in the Bcc header field with a Bcc header field containing an indication such as "Undisclosed recipients" (but no addressees)
 - c) The same message for each recipient in the Bcc header field which does not include a Bcc header field (this message is identical to 1. / cf. above)
3. The message stored in the 'Sent'-Folder of the sender, which usually contains the Bcc unchanged from the original message, i.e. with all recipient addresses.

Regarding the Bcc header field there should be no difference between the inner and the outer header section.

A.3. Receiving Side Considerations

A.3.1. Which Header Fields to Display to User

When displaying S/MIME messages which protect header fields (independent of which protection level 'signature and encryption', 'signature only' or 'encryption only' is applied to (cf. Section 3.2):

1. The outer header fields might be tampered with, so a receiving client SHOULD ignore them, unless they are protected in some other way(*). If a header field is present in the inner header, only the inner header field value MUST be displayed (and the corresponding outer value must be ignored). If a particular header field is only present in the outer header, it MAY be ignored (not displayed) or it MAY be displayed with a clear indicator that it is not trustworthy(*).

(*) - this only applies if the header field is not protected in some other way, for example with a DKIM signature that validates and is trusted.

A.3.2. Mail User Agent Algorithm for deciding which version of a header field to display

[[TODO: describe how to recurse to find the innermost protected root body part, extract header fields from it and propagate them to the top level. This should also work for triple-wrapped messages.]]

Appendix B. Document Changelog

[[RFC Editor: This section is to be removed before publication]]

- o draft-ietf-lamps-header-protection-requirements-00
 - * Initial version
- o draft-ietf-lamps-header-protection-requirements-01
 - * Moved Implementation Considerations to Appendix (HB)
 - * Shortened abstract (HB)
 - * Many editorial changes, e.g., replaced "content-type" with "Content-Type". (HB)
 - * Added example for Option 2.1 / pEp (HB)
 - * Added (short) definition of Header Protection (HB)

- * Added more information regarding Bcc (feedback IETF-105) (HB)
- * Simplified GS3 (HB)
- * Added GR3 (HB)

Appendix C. Open Issues

[[RFC Editor: This section should be empty and is to be removed before publication.]]

- o Enhance Introduction and Problem Statement sections
- o Decide in which form legacy HP requirements should remain in this document
- o Improve definitions in Section 3.2
- o Should requirement G3 remain? If you consider improve / rewrite it.
- o Add more text on Memory Hole
- o Rephrase Appendix A.1.2
- o Resolve question regarding Bcc in Appendix A.2.1
- o Rewrite Appendix A.2.1
- o Write Appendix A.3.2
- o Correct terminology for Header(s) and Header Fields throughout the document (editorial).
 - * Header: Whole Header Section of the message
 - * Header Field: Part / single Line inside a Header (Section)
- o Replace "email" by "email message" as needed

Authors' Addresses

Alexey Melnikov
Isode Ltd
14 Castle Mews
Hampton, Middlesex TW12 2NP
UK

Email: alexey.melnikov@isode.com

Bernie Hoeneisen
Ucom Standards Track Solutions GmbH
CH-8046 Zuerich
Switzerland

Phone: +41 44 500 52 40
Email: bernie@ietf.hoeneisen.ch (bernhard.hoeneisen AT ucom.ch)
URI: <https://ucom.ch/>

LAMPS WG
Internet-Draft
Updates: 3279 (if approved)
Intended status: Standards Track
Expires: January 22, 2020

P. Kampanakis
Cisco Systems
Q. Dang
NIST
July 21, 2019

Internet X.509 Public Key Infrastructure: Additional Algorithm
Identifiers for RSASSA-PSS and ECDSA using SHAKEs
draft-ietf-lamps-pkix-shake-15

Abstract

Digital signatures are used to sign messages, X.509 certificates and CRLs. This document updates the "Algorithms and Identifiers for the Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List Profile" (RFC3279) and describes the conventions for using the SHAKE function family in Internet X.509 certificates and revocation lists as one-way hash functions with the RSA Probabilistic signature and ECDSA signature algorithms. The conventions for the associated subject public keys are also described.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 22, 2020.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents

carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Change Log	2
2. Introduction	5
3. Terminology	5
4. Identifiers	5
5. Use in PKIX	6
5.1. Signatures	6
5.1.1. RSASSA-PSS Signatures	7
5.1.2. ECDSA Signatures	8
5.2. Public Keys	9
6. IANA Considerations	9
7. Security Considerations	10
8. Acknowledgements	10
9. References	11
9.1. Normative References	11
9.2. Informative References	12
Appendix A. ASN.1 module	13
Authors' Addresses	17

1. Change Log

[EDNOTE: Remove this section before publication.]

- o draft-ietf-lamps-pkix-shake-15:
 - * Minor editorial nits.
- o draft-ietf-lamps-pkix-shake-14:
 - * Fixing error with incorrect preimage resistance bits for SHA128 and SHA256.
- o draft-ietf-lamps-pkix-shake-13:
 - * Addressing one applicable comment from Dan M. about sec levels while in secdir review of draft-ietf-lamps-cms-shakes.
 - * Addressing comment from Scott B.'s opsdirev review about references in the abstract.
- o draft-ietf-lamps-pkix-shake-12:

- * Nits identified by Roman, Eric V. Ben K., Barry L. in ballot position review.
- o draft-ietf-lamps-pkix-shake-11:
 - * Nits identified by Roman in AD Review.
- o draft-ietf-lamps-pkix-shake-10:
 - * Updated IANA considerations section to request for OID assignments.
- o draft-ietf-lamps-pkix-shake-09:
 - * Fixed minor text nits.
 - * Added text name allocation for SHAKEs in IANA considerations.
 - * Updates in Sec Considerations section.
- o draft-ietf-lamps-pkix-shake-08:
 - * Small nits from Russ while in WGLC.
- o draft-ietf-lamps-pkix-shake-07:
 - * Incorporated Eric's suggestion from WGLC.
- o draft-ietf-lamps-pkix-shake-06:
 - * Added informative references.
 - * Updated ASN.1 so it compiles.
 - * Updated IANA considerations.
- o draft-ietf-lamps-pkix-shake-05:
 - * Added RFC8174 reference and text.
 - * Explicitly explained why RSASSA-PSS-params are omitted in section 5.1.1.
 - * Simplified Public Keys section by removing redundant info from RFCs.
- o draft-ietf-lamps-pkix-shake-04:

- * Removed paragraph suggesting KMAC to be used in generating k in Deterministic ECDSA. That should be RFC6979-bis.
- * Removed paragraph from Security Considerations that talks about randomness of k because we are using deterministic ECDSA.
- * Various ASN.1 fixes.
- * Text fixes.
- o draft-ietf-lamps-pkix-shake-03:
 - * Updates based on suggestions and clarifications by Jim.
 - * Added ASN.1.
- o draft-ietf-lamps-pkix-shake-02:
 - * Significant reorganization of the sections to simplify the introduction, the new OIDs and their use in PKIX.
 - * Added new OIDs for RSASSA-PSS that hardcode hash, salt and MGF, according the WG consensus.
 - * Updated Public Key section to use the new RSASSA-PSS OIDs and clarify the algorithm identifier usage.
 - * Removed the no longer used SHAKE OIDs from section 3.1.
 - * Consolidated subsection for message digest algorithms.
 - * Text fixes.
- o draft-ietf-lamps-pkix-shake-01:
 - * Changed titles and section names.
 - * Removed DSA after WG discussions.
 - * Updated shake OID names and parameters, added MGF1 section.
 - * Updated RSASSA-PSS section.
 - * Added Public key algorithm OIDs.
 - * Populated Introduction and IANA sections.
- o draft-ietf-lamps-pkix-shake-00:

* Initial version

2. Introduction

[RFC3279] defines cryptographic algorithm identifiers for the Internet X.509 Certificate and Certificate Revocation Lists (CRL) profile [RFC5280]. This document updates RFC3279 and defines identifiers for several cryptographic algorithms that use variable length output SHAKE functions introduced in [SHA3] which can be used with .

In the SHA-3 family, two extendable-output functions (SHAKEs), SHAKE128 and SHAKE256, are defined. Four other hash function instances, SHA3-224, SHA3-256, SHA3-384, and SHA3-512, are also defined but are out of scope for this document. A SHAKE is a variable length hash function defined as $\text{SHAKE}(M, d)$ where the output is a d -bits-long digest of message M . The corresponding collision and second-preimage-resistance strengths for SHAKE128 are $\min(d/2, 128)$ and $\min(d, 128)$ bits, respectively (Appendix A.1 [SHA3]). And the corresponding collision and second-preimage-resistance strengths for SHAKE256 are $\min(d/2, 256)$ and $\min(d, 256)$ bits, respectively.

A SHAKE can be used as the message digest function (to hash the message to be signed) in RSASSA-PSS [RFC8017] and ECDSA [X9.62] and as the hash in the mask generation function (MGF) in RSASSA-PSS.

3. 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 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

4. Identifiers

This section defines four new object identifiers (OIDs), for RSASSA-PSS and ECDSA with each of SHAKE128 and SHAKE256. The same algorithm identifiers can be used for identifying a public key in RSASSA-PSS.

The new identifiers for RSASSA-PSS signatures using SHAKEs are below.

```
id-RSASSA-PSS-SHAKE128 OBJECT IDENTIFIER ::= { iso(1)
    identified-organization(3) dod(6) internet(1)
    security(5) mechanisms(5) pkix(7) algorithms(6)
    TBD1 }

id-RSASSA-PSS-SHAKE256 OBJECT IDENTIFIER ::= { iso(1)
    identified-organization(3) dod(6) internet(1)
    security(5) mechanisms(5) pkix(7) algorithms(6)
    TBD2 }
```

The new algorithm identifiers of ECDSA signatures using SHAKEs are below.

```
id-ecdsa-with-shake128 OBJECT IDENTIFIER ::= { iso(1)
    identified-organization(3) dod(6) internet(1)
    security(5) mechanisms(5) pkix(7) algorithms(6)
    TBD3 }

id-ecdsa-with-shake256 OBJECT IDENTIFIER ::= { iso(1)
    identified-organization(3) dod(6) internet(1)
    security(5) mechanisms(5) pkix(7) algorithms(6)
    TBD4 }
```

The parameters for the four identifiers above MUST be absent. That is, the identifier SHALL be a SEQUENCE of one component, the OID.

Section 5.1.1 and Section 5.1.2 specify the required output length for each use of SHAKE128 or SHAKE256 in RSASSA-PSS and ECDSA. In summary, when hashing messages to be signed, output lengths of SHAKE128 and SHAKE256 are 256 and 512 bits respectively. When the SHAKEs are used as mask generation functions RSASSA-PSS, their output length is $(8 * \text{ceil}((n-1)/8) - 264)$ or $(8 * \text{ceil}((n-1)/8) - 520)$ bits, respectively, where n is the RSA modulus size in bits.

5. Use in PKIX

5.1. Signatures

Signatures are used in a number of different ASN.1 structures. As shown in the ASN.1 representation from [RFC5280] below, in an X.509 certificate, a signature is encoded with an algorithm identifier in the signatureAlgorithm attribute and a signatureValue attribute that contains the actual signature.

```
Certificate ::= SEQUENCE {  
    tbsCertificate      TBSCertificate,  
    signatureAlgorithm  AlgorithmIdentifier,  
    signatureValue      BIT STRING }
```

The identifiers defined in Section 4 can be used as the AlgorithmIdentifier in the signatureAlgorithm field in the sequence Certificate and the signature field in the sequence TBSCertificate in X.509 [RFC5280]. The parameters of these signature algorithms are absent as explained in Section 4.

Conforming CA implementations MUST specify the algorithms explicitly by using the OIDs specified in Section 4 when encoding RSASSA-PSS or ECDSA with SHAKE signatures in certificates and CRLs. Conforming client implementations that process certificates and CRLs using RSASSA-PSS or ECDSA with SHAKE MUST recognize the corresponding OIDs. Encoding rules for RSASSA-PSS and ECDSA signature values are specified in [RFC4055] and [RFC5480], respectively.

When using RSASSA-PSS or ECDSA with SHAKEs, the RSA modulus and ECDSA curve order SHOULD be chosen in line with the SHAKE output length. Refer to Section 7 for more details.

5.1.1.1. RSASSA-PSS Signatures

The RSASSA-PSS algorithm is defined in [RFC8017]. When id-RSASSA-PSS-SHAKE128 or id-RSASSA-PSS-SHAKE256 specified in Section 4 is used, the encoding MUST omit the parameters field. That is, the AlgorithmIdentifier SHALL be a SEQUENCE of one component, id-RSASSA-PSS-SHAKE128 or id-RSASSA-PSS-SHAKE256. [RFC4055] defines RSASSA-PSS-params that are used to define the algorithms and inputs to the algorithm. This specification does not use parameters because the hash, mask generation algorithm, trailer and salt are embedded in the OID definition.

The hash algorithm to hash a message being signed and the hash algorithm used as the mask generation function in RSASSA-PSS MUST be the same: both SHAKE128 or both SHAKE256. The output length of the hash algorithm which hashes the message SHALL be 32 (for SHAKE128) or 64 bytes (for SHAKE256).

The mask generation function takes an octet string of variable length and a desired output length as input, and outputs an octet string of the desired length. In RSASSA-PSS with SHAKEs, the SHAKEs MUST be used natively as the MGF function, instead of the MGF1 algorithm that uses the hash function in multiple iterations as specified in Section B.2.1 of [RFC8017]. In other words, the MGF is defined as the SHAKE128 or SHAKE256 output of the mgfSeed for id-RSASSA-PSS-

SHAKE128 and id-RSASSA-PSS-SHAKE256, respectively. The mgfSeed is the seed from which mask is generated, an octet string [RFC8017]. As explained in Step 9 of section 9.1.1 of [RFC8017], the output length of the MGF is $\text{emLen} - \text{hLen} - 1$ bytes. emLen is the maximum message length $\text{ceil}((n-1)/8)$, where n is the RSA modulus in bits. hLen is 32 and 64-bytes for id-RSASSA-PSS-SHAKE128 and id-RSASSA-PSS-SHAKE256, respectively. Thus when SHAKE is used as the MGF, the SHAKE output length maskLen is $(8 * \text{emLen} - 264)$ or $(8 * \text{emLen} - 520)$ bits, respectively. For example, when RSA modulus n is 2048, the output length of SHAKE128 or SHAKE256 as the MGF will be 1784 or 1528-bits when id-RSASSA-PSS-SHAKE128 or id-RSASSA-PSS-SHAKE256 is used, respectively.

The RSASSA-PSS saltLength MUST be 32 bytes for id-RSASSA-PSS-SHAKE128 or 64 bytes for id-RSASSA-PSS-SHAKE256. Finally, the trailerField MUST be 1, which represents the trailer field with hexadecimal value 0xBC [RFC8017].

5.1.2. ECDSA Signatures

The Elliptic Curve Digital Signature Algorithm (ECDSA) is defined in [X9.62]. When the id-ecdsa-with-shake128 or id-ecdsa-with-shake256 (specified in Section 4) algorithm identifier appears, the respective SHAKE function (SHAKE128 or SHAKE256) is used as the hash. The encoding MUST omit the parameters field. That is, the AlgorithmIdentifier SHALL be a SEQUENCE of one component, the OID id-ecdsa-with-shake128 or id-ecdsa-with-shake256.

For simplicity and compliance with the ECDSA standard specification, the output length of the hash function must be explicitly determined. The output length, d , for SHAKE128 or SHAKE256 used in ECDSA MUST be 256 or 512 bits, respectively.

Conforming CA implementations that generate ECDSA with SHAKE signatures in certificates or CRLs SHOULD generate such signatures with a deterministically generated, non-random k in accordance with all the requirements specified in [RFC6979]. They MAY also generate such signatures in accordance with all other recommendations in [X9.62] or [SEC1] if they have a stated policy that requires conformance to those standards. Those standards have not specified SHAKE128 and SHAKE256 as hash algorithm options. However, SHAKE128 and SHAKE256 with output length being 32 and 64 octets, respectively, can be used instead of 256 and 512-bit output hash algorithms such as SHA256 and SHA512.

5.2. Public Keys

Certificates conforming to [RFC5280] can convey a public key for any public key algorithm. The certificate indicates the public key algorithm through an algorithm identifier. This algorithm identifier is an OID and optionally associated parameters. The conventions and encoding for RSASSA-PSS and ECDSA public keys algorithm identifiers are as specified in Section 2.3.1 and 2.3.5 of [RFC3279], Section 3.1 of [RFC4055] and Section 2.1 of [RFC5480].

Traditionally, the `rsaEncryption` object identifier is used to identify RSA public keys. The `rsaEncryption` object identifier continues to identify the subject public key when the RSA private key owner does not wish to limit the use of the public key exclusively to RSASSA-PSS with SHAKes. When the RSA private key owner wishes to limit the use of the public key exclusively to RSASSA-PSS with SHAKes, the `AlgorithmIdentifiers` for RSASSA-PSS defined in Section 4 SHOULD be used as the algorithm field in the `SubjectPublicKeyInfo` sequence [RFC5280]. Conforming client implementations that process RSASSA-PSS with SHAKE public keys when processing certificates and CRLs MUST recognize the corresponding OIDs.

Conforming CA implementations MUST specify the X.509 public key algorithm explicitly by using the OIDs specified in Section 4 when encoding ECDSA with SHAKE public keys in certificates and CRLs. Conforming client implementations that process ECDSA with SHAKE public keys when processing certificates and CRLs MUST recognize the corresponding OIDs.

The identifier parameters, as explained in Section 4, MUST be absent.

6. IANA Considerations

One object identifier for the ASN.1 module in Appendix A is requested for the SMI Security for PKIX Module Identifiers (1.3.6.1.5.5.7.0) registry:

Decimal	Description	References
TBD	id-mod-pkix1-shakes-2019	[EDNOTE: THIS RFC]

IANA is requested to update the SMI Security for PKIX Algorithms [SMI-PKIX] (1.3.6.1.5.5.7.6) registry with four additional entries:

Decimal	Description	References
TBD1	id-RSASSA-PSS-SHAKE128	[EDNOTE: THIS RFC]
TBD2	id-RSASSA-PSS-SHAKE256	[EDNOTE: THIS RFC]
TBD3	id-ecdsa-with-shake128	[EDNOTE: THIS RFC]
TBD4	id-ecdsa-with-shake256	[EDNOTE: THIS RFC]

IANA is also requested to update the Hash Function Textual Names Registry [Hash-Texts] with two additional entries for SHAKE128 and SHAKE256:

Hash Function Name	OID	Reference
shake128	2.16.840.1.101.3.4.2.11	[EDNOTE: THIS RFC]
shake256	2.16.840.1.101.3.4.2.12	[EDNOTE: THIS RFC]

7. Security Considerations

This document updates [RFC3279]. The security considerations section of that document applies to this specification as well.

NIST has defined appropriate use of the hash functions in terms of the algorithm strengths and expected time frames for secure use in Special Publications (SPs) [SP800-78-4] and [SP800-107]. These documents can be used as guides to choose appropriate key sizes for various security scenarios.

SHAKE128 with output length of 256-bits offers 128-bits of collision and preimage resistance. Thus, SHAKE128 OIDs in this specification are RECOMMENDED with 2048 (112-bit security) or 3072-bit (128-bit security) RSA modulus or curves with group order of 256-bits (128-bit security). SHAKE256 with 512-bits output length offers 256-bits of collision and preimage resistance. Thus, the SHAKE256 OIDs in this specification are RECOMMENDED with 4096-bit RSA modulus or higher or curves with group order of at least 521-bits (256-bit security). Note that we recommended 4096-bit RSA because we would need 15360-bit modulus for 256-bits of security which is impractical for today's technology.

8. Acknowledgements

We would like to thank Sean Turner, Jim Schaad and Eric Rescorla for their valuable contributions to this document.

The authors would like to thank Russ Housley for his guidance and very valuable contributions with the ASN.1 module.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC3279] Bassham, L., Polk, W., and R. Housley, "Algorithms and Identifiers for the Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 3279, DOI 10.17487/RFC3279, April 2002, <<https://www.rfc-editor.org/info/rfc3279>>.
- [RFC4055] Schaad, J., Kaliski, B., and R. Housley, "Additional Algorithms and Identifiers for RSA Cryptography for use in the Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 4055, DOI 10.17487/RFC4055, June 2005, <<https://www.rfc-editor.org/info/rfc4055>>.
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <<https://www.rfc-editor.org/info/rfc5280>>.
- [RFC5480] Turner, S., Brown, D., Yiu, K., Housley, R., and T. Polk, "Elliptic Curve Cryptography Subject Public Key Information", RFC 5480, DOI 10.17487/RFC5480, March 2009, <<https://www.rfc-editor.org/info/rfc5480>>.
- [RFC8017] Moriarty, K., Ed., Kaliski, B., Jonsson, J., and A. Rusch, "PKCS #1: RSA Cryptography Specifications Version 2.2", RFC 8017, DOI 10.17487/RFC8017, November 2016, <<https://www.rfc-editor.org/info/rfc8017>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

- [SHA3] National Institute of Standards and Technology (NIST), "SHA-3 Standard – Permutation-Based Hash and Extendable-Output Functions FIPS PUB 202", August 2015, <<https://www.nist.gov/publications/sha-3-standard-permutation-based-hash-and-extendable-output-functions>>.

9.2. Informative References

- [Hash-Texts] IANA, "Hash Function Textual Names", July 2017, <<https://www.iana.org/assignments/hash-function-text-names/hash-function-text-names.xhtml>>.
- [RFC5912] Hoffman, P. and J. Schaad, "New ASN.1 Modules for the Public Key Infrastructure Using X.509 (PKIX)", RFC 5912, DOI 10.17487/RFC5912, June 2010, <<https://www.rfc-editor.org/info/rfc5912>>.
- [RFC6979] Pornin, T., "Deterministic Usage of the Digital Signature Algorithm (DSA) and Elliptic Curve Digital Signature Algorithm (ECDSA)", RFC 6979, DOI 10.17487/RFC6979, August 2013, <<https://www.rfc-editor.org/info/rfc6979>>.
- [SEC1] Standards for Efficient Cryptography Group, "SEC 1: Elliptic Curve Cryptography", May 2009, <<http://www.secg.org/sec1-v2.pdf>>.
- [SMI-PKIX] IANA, "SMI Security for PKIX Algorithms", March 2019, <<https://www.iana.org/assignments/smi-numbers/smi-numbers.xhtml#smi-numbers-1.3.6.1.5.5.7.6>>.
- [SP800-107] National Institute of Standards and Technology (NIST), "SP800-107: Recommendation for Applications Using Approved Hash Algorithms", May 2014, <https://csrc.nist.gov/csrc/media/publications/sp/800-107/rev-1/final/documents/draft_revised_sp800-107.pdf>.
- [SP800-78-4] National Institute of Standards and Technology (NIST), "SP800-78-4: Cryptographic Algorithms and Key Sizes for Personal Identity Verification", May 2014, <https://csrc.nist.gov/csrc/media/publications/sp/800-78/4/final/documents/sp800_78-4_revised_draft.pdf>.

[X9.62] American National Standard for Financial Services (ANSI), "X9.62-2005: Public Key Cryptography for the Financial Services Industry: The Elliptic Curve Digital Signature Standard (ECDSA)", November 2005.

Appendix A. ASN.1 module

This appendix includes the ASN.1 module for SHAKEs in X.509. This module does not come from any existing RFC.

```
PKIXAlgsForSHAKE-2019 { iso(1) identified-organization(3) dod(6)
    internet(1) security(5) mechanisms(5) pkix(7) id-mod(0)
    id-mod-pkix1-shakes-2019(TBD) }

DEFINITIONS EXPLICIT TAGS ::=

BEGIN

-- EXPORTS ALL;

IMPORTS

-- FROM [RFC5912]

PUBLIC-KEY, SIGNATURE-ALGORITHM, DIGEST-ALGORITHM, SMIME-CAPS
FROM AlgorithmInformation-2009
    { iso(1) identified-organization(3) dod(6) internet(1) security(5)
      mechanisms(5) pkix(7) id-mod(0)
      id-mod-algorithmInformation-02(58) }

-- FROM [RFC5912]

RSAPublicKey, rsaEncryption, pk-rsa, pk-ec,
CURVE, id-ecPublicKey, ECPPoint, ECPParameters, ECDSA-Sig-Value
FROM PKIXAlgs-2009 { iso(1) identified-organization(3) dod(6)
    internet(1) security(5) mechanisms(5) pkix(7) id-mod(0)
    id-mod-pkix1-algorithms2008-02(56) }
;

--
-- Message Digest Algorithms (mda-)
--
DigestAlgorithms DIGEST-ALGORITHM ::= {
    -- This expands DigestAlgorithms from [RFC5912]
    mda-shake128      |
    mda-shake256,
    ...
}
```

```
--
-- One-Way Hash Functions
--

-- SHAKE128
mda-shake128 DIGEST-ALGORITHM ::= {
  IDENTIFIER id-shake128 -- with output length 32 bytes.
}
id-shake128 OBJECT IDENTIFIER ::= { joint-iso-itu-t(2) country(16)
                                     us(840) organization(1) gov(101)
                                     csor(3) nistAlgorithm(4)
                                     hashAlgs(2) 11 }

-- SHAKE256
mda-shake256 DIGEST-ALGORITHM ::= {
  IDENTIFIER id-shake256 -- with output length 64 bytes.
}
id-shake256 OBJECT IDENTIFIER ::= { joint-iso-itu-t(2) country(16)
                                     us(840) organization(1) gov(101)
                                     csor(3) nistAlgorithm(4)
                                     hashAlgs(2) 12 }

--
-- Public Key (pk-) Algorithms
--
PublicKeys PUBLIC-KEY ::= {
  -- This expands PublicKeys from [RFC5912]
  pk-rsaSSA-PSS-SHAKE128 |
  pk-rsaSSA-PSS-SHAKE256,
  ...
}

-- The hashAlgorithm is mda-shake128
-- The maskGenAlgorithm is id-shake128
-- Mask Gen Algorithm is SHAKE128 with output length
--  $(8 * \text{ceil}((n-1)/8) - 264)$  bits, where n is the RSA
-- modulus in bits.
-- The saltLength is 32. The trailerField is 1.
pk-rsaSSA-PSS-SHAKE128 PUBLIC-KEY ::= {
  IDENTIFIER id-RSASSA-PSS-SHAKE128
  KEY RSAPublicKey
  PARAMS ARE absent
  -- Private key format not in this module --
  CERT-KEY-USAGE { nonRepudiation, digitalSignature,
                  keyCertSign, cRLSign }
}

-- The hashAlgorithm is mda-shake256
```

```

-- The maskGenAlgorithm is id-shake256
-- Mask Gen Algorithm is SHAKE256 with output length
-- (8*ceil((n-1)/8) - 520)-bits, where n is the RSA
-- modulus in bits.
-- The saltLength is 64. The trailerField is 1.
pk-rsaSSA-PSS-SHAKE256 PUBLIC-KEY ::= {
  IDENTIFIER id-RSASSA-PSS-SHAKE256
  KEY RSAPublicKey
  PARAMS ARE absent
  -- Private key format not in this module --
  CERT-KEY-USAGE { nonRepudiation, digitalSignature,
                    keyCertSign, cRLSign }
}

--
-- Signature Algorithms (sa-)
--
SignatureAlgs SIGNATURE-ALGORITHM ::= {
  -- This expands SignatureAlgorithms from [RFC5912]
  sa-rsaspssWithSHAKE128 |
  sa-rsaspssWithSHAKE256 |
  sa-ecdsaWithSHAKE128 |
  sa-ecdsaWithSHAKE256,
  ...
}

--
-- SMIME Capabilities (sa-)
--
SMimeCaps SMIME-CAPS ::= {
  -- The expands SMimeCaps from [RFC5912]
  sa-rsaspssWithSHAKE128.&smimeCaps |
  sa-rsaspssWithSHAKE256.&smimeCaps |
  sa-ecdsaWithSHAKE128.&smimeCaps |
  sa-ecdsaWithSHAKE256.&smimeCaps,
  ...
}

-- RSASSA-PSS with SHAKE128
sa-rsaspssWithSHAKE128 SIGNATURE-ALGORITHM ::= {
  IDENTIFIER id-RSASSA-PSS-SHAKE128
  PARAMS ARE absent
  -- The hashAlgorithm is mda-shake128
  -- The maskGenAlgorithm is id-shake128
  -- Mask Gen Algorithm is SHAKE128 with output length
  -- (8*ceil((n-1)/8) - 264) bits, where n is the RSA
  -- modulus in bits.
  -- The saltLength is 32. The trailerField is 1

```

```
HASHES { mda-shake128 }
PUBLIC-KEYS { pk-rsa | pk-rsaSSA-PSS-SHAKE128 }
SMIME-CAPS { IDENTIFIED BY id-RSASSA-PSS-SHAKE128 }
}
id-RSASSA-PSS-SHAKE128 OBJECT IDENTIFIER ::= { iso(1)
    identified-organization(3) dod(6) internet(1)
    security(5) mechanisms(5) pkix(7) algorithms(6)
    TBD1 }

-- RSASSA-PSS with SHAKE256
sa-rsassaPssWithSHAKE256 SIGNATURE-ALGORITHM ::= {
    IDENTIFIER id-RSASSA-PSS-SHAKE256
    PARAMS ARE absent
    -- The hashAlgorithm is mda-shake256
    -- The maskGenAlgorithm is id-shake256
    -- Mask Gen Algorithm is SHAKE256 with output length
    -- (8*ceil((n-1)/8) - 520)-bits, where n is the
    -- RSA modulus in bits.
    -- The saltLength is 64. The trailerField is 1.
    HASHES { mda-shake256 }
    PUBLIC-KEYS { pk-rsa | pk-rsaSSA-PSS-SHAKE256 }
    SMIME-CAPS { IDENTIFIED BY id-RSASSA-PSS-SHAKE256 }
}
id-RSASSA-PSS-SHAKE256 OBJECT IDENTIFIER ::= { iso(1)
    identified-organization(3) dod(6) internet(1)
    security(5) mechanisms(5) pkix(7) algorithms(6)
    TBD2 }

-- ECDSA with SHAKE128
sa-ecdsaWithSHAKE128 SIGNATURE-ALGORITHM ::= {
    IDENTIFIER id-ecdsa-with-shake128
    VALUE ECDSA-Sig-Value
    PARAMS ARE absent
    HASHES { mda-shake128 }
    PUBLIC-KEYS { pk-ec }
    SMIME-CAPS { IDENTIFIED BY id-ecdsa-with-shake128 }
}
id-ecdsa-with-shake128 OBJECT IDENTIFIER ::= { iso(1)
    identified-organization(3) dod(6) internet(1)
    security(5) mechanisms(5) pkix(7) algorithms(6)
    TBD3 }

-- ECDSA with SHAKE256
sa-ecdsaWithSHAKE256 SIGNATURE-ALGORITHM ::= {
    IDENTIFIER id-ecdsa-with-shake256
    VALUE ECDSA-Sig-Value
    PARAMS ARE absent
    HASHES { mda-shake256 }
```



```
PUBLIC-KEYS { pk-ec }
  SMIME-CAPS { IDENTIFIED BY id-ecdsa-with-shake256 }
}
id-ecdsa-with-shake256 OBJECT IDENTIFIER ::= { iso(1)
  identified-organization(3) dod(6) internet(1)
  security(5) mechanisms(5) pkix(7) algorithms(6)
  TBD4 }
```

END

Authors' Addresses

Panos Kampanakis
Cisco Systems

Email: pkampana@cisco.com

Quynh Dang
NIST
100 Bureau Drive, Stop 8930
Gaithersburg, MD 20899-8930
USA

Email: quynh.dang@nist.gov

Network Working Group
Internet-Draft
Obsoletes: 6844 (if approved)
Intended status: Standards Track
Expires: December 1, 2019

P. Hallam-Baker
R. Stradling
Sectigo
J. Hoffman-Andrews
Let's Encrypt
May 30, 2019

DNS Certification Authority Authorization (CAA) Resource Record
draft-ietf-lamps-rfc6844bis-07

Abstract

The Certification Authority Authorization (CAA) DNS Resource Record allows a DNS domain name holder to specify one or more Certification Authorities (CAs) authorized to issue certificates for that domain name. CAA Resource Records allow a public Certification Authority to implement additional controls to reduce the risk of unintended certificate mis-issue. This document defines the syntax of the CAA record and rules for processing CAA records by certificate issuers.

This document obsoletes RFC 6844.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on December 1, 2019.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

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

(<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Definitions	3
2.1. Requirements Language	3
2.2. Defined Terms	4
3. Relevant Resource Record Set	5
4. Mechanism	6
4.1. Syntax	6
4.1.1. Canonical Presentation Format	7
4.2. CAA issue Property	8
4.3. CAA issuewild Property	9
4.4. CAA iodef Property	10
4.5. Critical Flag	11
5. Security Considerations	11
5.1. Use of DNS Security	12
5.2. Non-Compliance by Certification Authority	12
5.3. Mis-Issue by Authorized Certification Authority	12
5.4. Suppression or Spoofing of CAA Records	12
5.5. Denial of Service	13
5.6. Abuse of the Critical Flag	13
6. Deployment Considerations	13
6.1. Blocked Queries or Responses	14
6.2. Rejected Queries and Malformed Responses	14
6.3. Delegation to Private Nameservers	14
6.4. Bogus DNSSEC Responses	14
7. Differences versus RFC6844	15
8. IANA Considerations	15
9. Acknowledgements	16
10. References	16
10.1. Normative References	16
10.2. Informative References	17
Authors' Addresses	17

1. Introduction

The Certification Authority Authorization (CAA) DNS Resource Record allows a DNS domain name holder to specify the Certification Authorities (CAs) authorized to issue certificates for that domain name. Publication of CAA Resource Records allows a public

Certification Authority to implement additional controls to reduce the risk of unintended certificate mis-issue.

Like the TLSA record defined in DNS-Based Authentication of Named Entities (DANE) [RFC6698], CAA records are used as a part of a mechanism for checking PKIX [RFC6698] certificate data. The distinction between the two specifications is that CAA records specify an authorization control to be performed by a certificate issuer before issue of a certificate and TLSA records specify a verification control to be performed by a relying party after the certificate is issued.

Conformance with a published CAA record is a necessary but not sufficient condition for issuance of a certificate.

Criteria for inclusion of embedded trust anchor certificates in applications are outside the scope of this document. Typically, such criteria require the CA to publish a Certification Practices Statement (CPS) that specifies how the requirements of the Certificate Policy (CP) are achieved. It is also common for a CA to engage an independent third-party auditor to prepare an annual audit statement of its performance against its CPS.

A set of CAA records describes only current grants of authority to issue certificates for the corresponding DNS domain name. Since certificates are valid for a period of time, it is possible that a certificate that is not conformant with the CAA records currently published was conformant with the CAA records published at the time that the certificate was issued. Relying parties **MUST NOT** use CAA records as part of certificate validation.

CAA records **MAY** be used by Certificate Evaluators as a possible indicator of a security policy violation. Such use **SHOULD** take account of the possibility that published CAA records changed between the time a certificate was issued and the time at which the certificate was observed by the Certificate Evaluator.

2. Definitions

2.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2.2. Defined Terms

The following terms are used in this document:

Certificate: An X.509 Certificate, as specified in [RFC5280].

Certificate Evaluator: A party other than a Relying Party that evaluates the trustworthiness of certificates issued by Certification Authorities.

Certification Authority (CA): An Issuer that issues certificates in accordance with a specified Certificate Policy.

Certificate Policy (CP): Specifies the criteria that a Certification Authority undertakes to meet in its issue of certificates. See [RFC3647].

Certification Practices Statement (CPS): Specifies the means by which the criteria of the Certificate Policy are met. In most cases, this will be the document against which the operations of the Certification Authority are audited. See [RFC3647].

Domain Name: The label assigned to a node in the Domain Name System.

Domain Name System (DNS): The Internet naming system specified in [RFC1034] and [RFC1035].

DNS Security (DNSSEC): Extensions to the DNS that provide authentication services as specified in [RFC4033], [RFC4034], [RFC4035], [RFC5155], and revisions.

Fully-Qualified Domain Name (FQDN): A Domain Name that includes the labels of all superior nodes in the Domain Name System.

Issuer: An entity that issues certificates. See [RFC5280].

Property: The tag-value portion of a CAA Resource Record.

Property Tag: The tag portion of a CAA Resource Record.

Property Value: The value portion of a CAA Resource Record.

Resource Record (RR): A particular entry in the DNS including the owner name, class, type, time to live, and data, as defined in [RFC1034] and [RFC2181].

Resource Record Set (RRSet): A set of Resource Records of a particular owner name, class, and type. The time to live on all RRs

within an RRSet is always the same, but the data may be different among RRs in the RRSet.

Relevant Resource Record Set (Relevant RRSet): A set of CAA Resource Records resulting from applying the algorithm in Section 3 to a specific Fully-Qualified Domain Name or Wildcard Domain Name.

Relying Party: A party that makes use of an application whose operation depends on use of a certificate for making a security decision. See [RFC5280].

Wildcard Domain Name: A Domain Name consisting of a single asterisk character followed by a single full stop character ("*.") followed by a Fully-Qualified Domain Name.

3. Relevant Resource Record Set

Before issuing a certificate, a compliant CA MUST check for publication of a Relevant RRSet. If such an RRSet exists, a CA MUST NOT issue a certificate unless the CA determines that either (1) the certificate request is consistent with the applicable CAA Resource Record set or (2) an exception specified in the relevant Certificate Policy or Certification Practices Statement applies. If the Relevant RRSet for a Fully-Qualified Domain Name or Wildcard Domain Name contains no Property Tags that restrict issuance (for instance, if it contains only iodef Property Tags, or only Property Tags unrecognized by the CA), CAA does not restrict issuance.

A certificate request MAY specify more than one Fully-Qualified Domain Name and MAY specify Wildcard Domain Names. Issuers MUST verify authorization for all the Fully-Qualified Domain Names and Wildcard Domain Names specified in the request.

The search for a CAA RRSet climbs the DNS name tree from the specified label up to but not including the DNS root '.' until a CAA RRSet is found.

Given a request for a specific Fully-Qualified Domain Name X, or a request for a Wildcard Domain Name *.X, the Relevant Resource Record Set RelevantCAASet(X) is determined as follows (in pseudocode):

Let CAA(X) be the RRSet returned by performing a CAA record query for the Fully-Qualified Domain Name X, according to the lookup algorithm specified in RFC 1034 section 4.3.2 (in particular chasing aliases). Let Parent(X) be the Fully-Qualified Domain Name produced by removing the leftmost label of X.

```

RelevantCAASet(domain):
  while domain is not ".":
    if CAA(domain) is not Empty:
      return CAA(domain)
    domain = Parent(domain)
  return Empty

```

For example, processing CAA for the Fully-Qualified Domain Name "X.Y.Z" where there are no CAA records at any level in the tree RelevantCAASet would have the following steps:

```

CAA("X.Y.Z.") = Empty; domain = Parent("X.Y.Z.") = "Y.Z."
CAA("Y.Z.")   = Empty; domain = Parent("Y.Z.")   = "Z."
CAA("Z.")     = Empty; domain = Parent("Z.")     = "."
return Empty

```

Processing CAA for the Fully-Qualified Domain Name "A.B.C" where there is a CAA record "issue example.com" at "B.C" would terminate early upon finding the CAA record:

```

CAA("A.B.C.") = Empty; domain = Parent("A.B.C.") = "B.C."
CAA("B.C.")   = "issue example.com"
return "issue example.com"

```

4. Mechanism

4.1. Syntax

A CAA Resource Record contains a single Property consisting of a tag-value pair. A Fully-Qualified Domain Name MAY have multiple CAA RRs associated with it and a given Property Tag MAY be specified more than once across those RRs.

The RDATA section for a CAA Resource Record contains one Property. A Property consists of the following:

+0-1-2-3-4-5-6-7-	0-1-2-3-4-5-6-7-	
Flags	Tag Length = n	
+-----+	+...+-----+	
Tag char 0	Tag char 1	... Tag char n-1
+-----+	+...+-----+	
+-----+	+.....+-----+	
Value byte 0	Value byte 1 Value byte m-1
+-----+	+.....+-----+	

Where n is the length specified in the Tag length field and m is the remaining octets in the Value field. They are related by $(m = d - n - 2)$ where d is the length of the RDATA section.

The fields are defined as follows:

Flags: One octet containing the following field:

Bit 0, Issuer Critical Flag: If the value is set to '1', the Property is critical. A Certification Authority MUST NOT issue certificates for any FQDN the Relevant RRSset for that FQDN contains a CAA critical Property for an unknown or unsupported Property Tag.

Note that according to the conventions set out in [RFC1035], bit 0 is the Most Significant Bit and bit 7 is the Least Significant Bit. Thus, the Flags value 1 means that bit 7 is set while a value of 128 means that bit 0 is set according to this convention.

All other bit positions are reserved for future use.

To ensure compatibility with future extensions to CAA, DNS records compliant with this version of the CAA specification MUST clear (set to "0") all reserved flags bits. Applications that interpret CAA records MUST ignore the value of all reserved flag bits.

Tag Length: A single octet containing an unsigned integer specifying the tag length in octets. The tag length MUST be at least 1.

Tag: The Property identifier, a sequence of US-ASCII characters.

Tags MAY contain US-ASCII characters 'a' through 'z', 'A' through 'Z', and the numbers 0 through 9. Tags MUST NOT contain any other characters. Matching of tags is case insensitive.

Tags submitted for registration by IANA MUST NOT contain any characters other than the (lowercase) US-ASCII characters 'a' through 'z' and the numbers 0 through 9.

Value: A sequence of octets representing the Property Value. Property Values are encoded as binary values and MAY employ sub-formats.

The length of the value field is specified implicitly as the remaining length of the enclosing RDATA section.

4.1.1. Canonical Presentation Format

The canonical presentation format of the CAA record is:

CAA <flags> <tag> <value>

Where:

Flags: Is an unsigned integer between 0 and 255.

Tag: Is a non-zero-length sequence of US-ASCII letters and numbers in lower case.

Value: The value field, expressed as a contiguous set of characters without interior spaces, or as a quoted string. See the <character-string> format specified in [RFC1035], Section 5.1, but note that the value field contains no length byte and is not limited to 255 characters.

4.2. CAA issue Property

If the issue Property Tag is present in the Relevant RRSset for a Fully-Qualified Domain Name, it is a request that Issuers

1. Perform CAA issue restriction processing for the FQDN, and
2. Grant authorization to issue certificates containing that FQDN to the holder of the issuer-domain-name or a party acting under the explicit authority of the holder of the issuer-domain-name.

The CAA issue Property Value has the following sub-syntax (specified in ABNF as per [RFC5234]).

```
issue-value = *WSP [issuer-domain-name *WSP] [";" *WSP [parameters *WSP]]
```

```
issuer-domain-name = label *("." label)
```

```
label = (ALPHA / DIGIT) *( *("-") (ALPHA / DIGIT))
```

```
parameters = (parameter *WSP ";" *WSP parameters) / parameter
```

```
parameter = tag *WSP "=" *WSP value
```

```
tag = (ALPHA / DIGIT) *( *("-") (ALPHA / DIGIT))
```

```
value = *("%x21-3A / %x3C-7E")
```

For consistency with other aspects of DNS administration, FQDN values are specified in letter-digit-hyphen Label (LDH-Label) form.

The following CAA record set requests that no certificates be issued for the FQDN 'certs.example.com' by any Issuer other than cal.example.net or ca2.example.org.

```
certs.example.com      CAA 0 issue "cal.example.net"
certs.example.com      CAA 0 issue "ca2.example.org"
```

Because the presence of an issue Property Tag in the Relevant RRSset for an FQDN restricts issuance, FQDN owners can use an issue Property Tag with no issuer-domain-name to request no issuance.

For example, the following RRSSet requests that no certificates be issued for the FQDN 'nocerts.example.com' by any Issuer.

```
nocerts.example.com      CAA 0 issue ";"
```

An issue Property Tag where the issue-value does not match the ABNF grammar MUST be treated the same as one specifying an empty issuer-domain-name. For example, the following malformed CAA RRsSet forbids issuance:

```
malformed.example.com    CAA 0 issue "%%%"
```

CAA authorizations are additive; thus, the result of specifying both an empty issuer-domain-name and a non-empty issuer-domain-name is the same as specifying just the non-empty issuer-domain-name.

An Issuer MAY choose to specify parameters that further constrain the issue of certificates by that Issuer, for example, specifying that certificates are to be subject to specific validation policies, billed to certain accounts, or issued under specific trust anchors.

For example, if cal.example.net has requested its customer accountable.example.com to specify their account number "230123" in each of the customer's CAA records using the (CA-defined) "account" parameter, it would look like this:

```
accountable.example.com  CAA 0 issue "cal.example.net; account=230123"
```

The semantics of parameters to the issue Property Tag are determined by the Issuer alone.

4.3. CAA issuewild Property

The issuewild Property Tag has the same syntax and semantics as the issue Property Tag except that it only grants authorization to issue certificates that specify a Wildcard Domain Name and issuewild properties take precedence over issue properties when specified. Specifically:

issuewild properties MUST be ignored when processing a request for a Fully-Qualified Domain Name that is not a Wildcard Domain Name.

If at least one issuewild Property is specified in the Relevant RRsSet for a Wildcard Domain Name, all issue properties MUST be ignored when processing a request for that Wildcard Domain Name.

For example, the following RRsSet requests that only cal.example.net issue certificates for "wild.example.com" or "sub.wild.example.com",

and that `_only_ ca2.example.org` issue certificates for `"*.wild.example.com"` or `"*.sub.wild.example.com"`). Note that this presumes there are no CAA RRs for `sub.wild.example.com`.

```
wild.example.com      CAA 0 issue "ca1.example.net"
wild.example.com      CAA 0 issuewild "ca2.example.org"
```

The following RRSset requests that `_only_ ca1.example.net` issue certificates for `"wild2.example.com"`, `"*.wild2.example.com"` or `"*.sub.wild2.example.com"`.

```
wild2.example.com      CAA 0 issue "ca1.example.net"
```

The following RRSset requests that `_only_ ca2.example.org` issue certificates for `"*.wild3.example.com"` or `"*.sub.wild3.example.com"`. It does not permit any Issuer to issue for `"wild3.example.com"` or `"sub.wild3.example.com"`.

```
wild3.example.com      CAA 0 issuewild "ca2.example.org"
wild3.example.com      CAA 0 issue ";"
```

The following RRSset requests that `_only_ ca2.example.org` issue certificates for `"*.wild3.example.com"` or `"*.sub.wild3.example.com"`. It permits any Issuer to issue for `"wild3.example.com"` or `"sub.wild3.example.com"`.

```
wild3.example.com      CAA 0 issuewild "ca2.example.org"
```

4.4. CAA iodef Property

The iodef Property specifies a means of reporting certificate issue requests or cases of certificate issue for domains for which the Property appears in the Relevant RRSset, when those requests or issuances violate the security policy of the Issuer or the FQDN holder.

The Incident Object Description Exchange Format (IODEF) [RFC7970] is used to present the incident report in machine-readable form.

The iodef Property Tag takes a URL as its Property Value. The URL scheme type determines the method used for reporting:

`mailto:` The IODEF incident report is reported as a MIME email attachment to an SMTP email that is submitted to the mail address specified. The mail message sent SHOULD contain a brief text message to alert the recipient to the nature of the attachment.

http or https: The IODEF report is submitted as a Web service request to the HTTP address specified using the protocol specified in [RFC6546].

These are the only supported URL schemes.

The following RRSset specifies that reports may be made by means of email with the IODEF data as an attachment, a Web service [RFC6546], or both:

```
report.example.com      CAA 0 issue "cal.example.net"
report.example.com      CAA 0 iodef "mailto:security@example.com"
report.example.com      CAA 0 iodef "http://iodef.example.com/"
```

4.5. Critical Flag

The critical flag is intended to permit future versions of CAA to introduce new semantics that MUST be understood for correct processing of the record, preventing conforming CAs that do not recognize the new semantics from issuing certificates for the indicated FQDNs.

In the following example, the Property with a Property Tag of 'tbs' is flagged as critical. Neither the cal.example.net CA nor any other Issuer is authorized to issue for "new.example.com" (or any other domains for which this is the Relevant RRSset) unless the Issuer has implemented the processing rules for the 'tbs' Property Tag.

```
new.example.com      CAA 0 issue "cal.example.net"
new.example.com      CAA 128 tbs "Unknown"
```

5. Security Considerations

CAA records assert a security policy that the holder of an FQDN wishes to be observed by Issuers. The effectiveness of CAA records as an access control mechanism is thus dependent on observance of CAA constraints by Issuers.

The objective of the CAA record properties described in this document is to reduce the risk of certificate mis-issue rather than avoid reliance on a certificate that has been mis-issued. DANE [RFC6698] describes a mechanism for avoiding reliance on mis-issued certificates.

5.1. Use of DNS Security

Use of DNSSEC to authenticate CAA RRs is strongly RECOMMENDED but not required. An Issuer MUST NOT issue certificates if doing so would conflict with the Relevant RRSets, irrespective of whether the corresponding DNS records are signed.

DNSSEC provides a proof of non-existence for both DNS Fully-Qualified Domain Names and RRSets within FQDNs. DNSSEC verification thus enables an Issuer to determine if the answer to a CAA record query is empty because the RRSSet is empty or if it is non-empty but the response has been suppressed.

Use of DNSSEC allows an Issuer to acquire and archive a proof that they were authorized to issue certificates for the FQDN. Verification of such archives may be an audit requirement to verify CAA record processing compliance. Publication of such archives may be a transparency requirement to verify CAA record processing compliance.

5.2. Non-Compliance by Certification Authority

CAA records offer CAs a cost-effective means of mitigating the risk of certificate mis-issue: the cost of implementing CAA checks is very small and the potential costs of a mis-issue event include the removal of an embedded trust anchor.

5.3. Mis-Issue by Authorized Certification Authority

Use of CAA records does not prevent mis-issue by an authorized Certification Authority, i.e., a CA that is authorized to issue certificates for the FQDN in question by CAA records.

FQDN holders SHOULD verify that the CAs they authorize to issue certificates for their FQDNs employ appropriate controls to ensure that certificates are issued only to authorized parties within their organization.

Such controls are most appropriately determined by the FQDN holder and the authorized CA(s) directly and are thus out of scope of this document.

5.4. Suppression or Spoofing of CAA Records

Suppression of the CAA record or insertion of a bogus CAA record could enable an attacker to obtain a certificate from an Issuer that was not authorized to issue for an affected FQDN.

Where possible, Issuers SHOULD perform DNSSEC validation to detect missing or modified CAA record sets.

In cases where DNSSEC is not deployed for a corresponding FQDN, an Issuer SHOULD attempt to mitigate this risk by employing appropriate DNS security controls. For example, all portions of the DNS lookup process SHOULD be performed against the authoritative name server. Data cached by third parties MUST NOT be relied on as the sole source of DNS CAA information but MAY be used to support additional anti-spoofing or anti-suppression controls.

5.5. Denial of Service

Introduction of a malformed or malicious CAA RR could in theory enable a Denial-of-Service (DoS) attack. This could happen by modification of authoritative DNS records or by spoofing inflight DNS responses.

This specific threat is not considered to add significantly to the risk of running an insecure DNS service.

An attacker could, in principle, perform a DoS attack against an Issuer by requesting a certificate with a maliciously long DNS name. In practice, the DNS protocol imposes a maximum name length and CAA processing does not exacerbate the existing need to mitigate DoS attacks to any meaningful degree.

5.6. Abuse of the Critical Flag

A Certification Authority could make use of the critical flag to trick customers into publishing records that prevent competing Certification Authorities from issuing certificates even though the customer intends to authorize multiple providers. This could happen if the customers were setting CAA records based on data provided by the CA rather than generating those records themselves.

In practice, such an attack would be of minimal effect since any competent competitor that found itself unable to issue certificates due to lack of support for a Property marked critical should investigate the cause and report the reason to the customer. The customer will thus discover that they had been deceived.

6. Deployment Considerations

A CA implementing CAA may find that they receive errors looking up CAA records. The following are some common causes of such errors, so that CAs may provide guidance to their subscribers on fixing the underlying problems.

6.1. Blocked Queries or Responses

Some middleboxes, in particular anti-DDoS appliances, may be configured to drop DNS packets of unknown types, or may start dropping such packets when they consider themselves under attack. This generally manifests as a timed-out DNS query, or a SERVFAIL at a local recursive resolver.

6.2. Rejected Queries and Malformed Responses

Some authoritative nameservers respond with REJECTED or NOTIMP when queried for a Resource Record type they do not recognize. At least one authoritative resolver produces a malformed response (with the QR bit set to 0) when queried for unknown Resource Record types. Per RFC 1034, the correct response for unknown Resource Record types is NOERROR.

6.3. Delegation to Private Nameservers

Some FQDN administrators make the contents of a subdomain unresolvable on the public Internet by delegating that subdomain to a nameserver whose IP address is private. A CA processing CAA records for such subdomains will receive SERVFAIL from its recursive resolver. The CA MAY interpret that as preventing issuance. FQDN administrators wishing to issue certificates for private FQDNs SHOULD use split-horizon DNS with a publicly available nameserver, so that CAs can receive a valid, empty CAA response for those FQDNs.

6.4. Bogus DNSSEC Responses

Queries for CAA Resource Records are different from most DNS RR types, because a signed, empty response to a query for CAA RRs is meaningfully different from a bogus response. A signed, empty response indicates that there is definitely no CAA policy set at a given label. A bogus response may mean either a misconfigured zone, or an attacker tampering with records. DNSSEC implementations may have bugs with signatures on empty responses that go unnoticed, because for more common Resource Record types like A and AAAA, the difference to an end user between empty and bogus is irrelevant; they both mean a site is unavailable.

In particular, at least two authoritative resolvers that implement live signing had bugs when returning empty Resource Record sets for DNSSEC-signed zones, in combination with mixed-case queries. Mixed-case queries, also known as DNS 0x20, are used by some recursive resolvers to increase resilience against DNS poisoning attacks. DNSSEC-signing authoritative resolvers are expected to copy the same capitalization from the query into their ANSWER section, but sign the

response as if they had used all lowercase. In particular, PowerDNS versions prior to 4.0.4 had this bug.

7. Differences versus RFC6844

This document obsoletes RFC6844. The most important change is to the Certification Authority Processing section. RFC6844 specified an algorithm that performed DNS tree-climbing not only on the FQDN being processed, but also on all CNAMEs and DNAMEs encountered along the way. This made the processing algorithm very inefficient when used on FQDNs that utilize many CNAMEs, and would have made it difficult for hosting providers to set CAA policies on their own FQDNs without setting potentially unwanted CAA policies on their customers' FQDNs. This document specifies a simplified processing algorithm that only performs tree climbing on the FQDN being processed, and leaves processing of CNAMEs and DNAMEs up to the CA's recursive resolver.

This document also includes a "Deployment Considerations" section detailing experience gained with practical deployment of CAA enforcement among CAs in the WebPKI.

This document clarifies the ABNF grammar for the issue and issuewild tags and resolves some inconsistencies with the document text. In particular, it specifies that parameters are separated with semicolons. It also allows hyphens in Property Tags.

This document also clarifies processing of a CAA RRset that is not empty, but contains no issue or issuewild tags.

This document removes the section titled "The CAA RR Type," merging it with "Mechanism" because the definitions were mainly duplicates. It moves the "Use of DNS Security" section into Security Considerations. It renames "Certification Authority Processing" to "Relevant Resource Record Set," and emphasizes the use of that term to more clearly define which domains are affected by a given RRset.

8. IANA Considerations

IANA is requested to add [[[RFC Editor: Please replace with this RFC]]] as a reference for the Certification Authority Restriction Flags and Certification Authority Restriction Properties registries, and update references to [RFC6844] within those registries to refer to [[[RFC Editor: Please replace with this RFC]]]. IANA is also requested to update the CAA TYPE in the DNS Parameters registry with a reference to [[[RFC Editor: Please replace with this RFC]]].

9. Acknowledgements

The authors would like to thank the following people who contributed to the design and documentation of this work item: Corey Bonnell, Chris Evans, Stephen Farrell, Jeff Hodges, Paul Hoffman, Tim Hollebeek, Stephen Kent, Adam Langley, Ben Laurie, James Manger, Chris Palmer, Scott Schmit, Sean Turner, and Ben Wilson.

10. References

10.1. Normative References

- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <<https://www.rfc-editor.org/info/rfc1034>>.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <<https://www.rfc-editor.org/info/rfc1035>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC2181] Elz, R. and R. Bush, "Clarifications to the DNS Specification", RFC 2181, DOI 10.17487/RFC2181, July 1997, <<https://www.rfc-editor.org/info/rfc2181>>.
- [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, DOI 10.17487/RFC4033, March 2005, <<https://www.rfc-editor.org/info/rfc4033>>.
- [RFC4034] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", RFC 4034, DOI 10.17487/RFC4034, March 2005, <<https://www.rfc-editor.org/info/rfc4034>>.
- [RFC4035] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", RFC 4035, DOI 10.17487/RFC4035, March 2005, <<https://www.rfc-editor.org/info/rfc4035>>.
- [RFC5155] Laurie, B., Sisson, G., Arends, R., and D. Blacka, "DNS Security (DNSSEC) Hashed Authenticated Denial of Existence", RFC 5155, DOI 10.17487/RFC5155, March 2008, <<https://www.rfc-editor.org/info/rfc5155>>.

- [RFC5234] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, DOI 10.17487/RFC5234, January 2008, <<https://www.rfc-editor.org/info/rfc5234>>.
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <<https://www.rfc-editor.org/info/rfc5280>>.
- [RFC6546] Trammell, B., "Transport of Real-time Inter-network Defense (RID) Messages over HTTP/TLS", RFC 6546, DOI 10.17487/RFC6546, April 2012, <<https://www.rfc-editor.org/info/rfc6546>>.
- [RFC6698] Hoffman, P. and J. Schlyter, "The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA", RFC 6698, DOI 10.17487/RFC6698, August 2012, <<https://www.rfc-editor.org/info/rfc6698>>.
- [RFC6844] Hallam-Baker, P. and R. Stradling, "DNS Certification Authority Authorization (CAA) Resource Record", RFC 6844, DOI 10.17487/RFC6844, January 2013, <<https://www.rfc-editor.org/info/rfc6844>>.
- [RFC7970] Danyliw, R., "The Incident Object Description Exchange Format Version 2", RFC 7970, DOI 10.17487/RFC7970, November 2016, <<https://www.rfc-editor.org/info/rfc7970>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

10.2. Informative References

- [RFC3647] Chokhani, S., Ford, W., Sabett, R., Merrill, C., and S. Wu, "Internet X.509 Public Key Infrastructure Certificate Policy and Certification Practices Framework", RFC 3647, DOI 10.17487/RFC3647, November 2003, <<https://www.rfc-editor.org/info/rfc3647>>.

Authors' Addresses

Phillip Hallam-Baker

Email: phill@hallambaker.com

Rob Stradling
Sectigo Ltd.

Email: rob@sectigo.com

Jacob Hoffman-Andrews
Let's Encrypt

Email: jsha@letsencrypt.org

LAMPS
Internet-Draft
Intended status: Standards Track
Expires: 12 August 2022

M. Ounsworth
Entrust
M. Pala
CableLabs
8 February 2022

Composite Signatures For Use In Internet PKI
draft-ounsworth-pq-composite-sigs-06

Abstract

With the widespread adoption of post-quantum cryptography will come the need for an entity to possess multiple public keys on different cryptographic algorithms. Since the trustworthiness of individual post-quantum algorithms is at question, a multi-key cryptographic operation will need to be performed in such a way that breaking it requires breaking each of the component algorithms individually. This requires defining new structures for holding composite signature data.

This document defines the structures `CompositeSignatureValue`, and `CompositeParams`, which are sequences of the respective structure for each component algorithm. This document also defines processes for generating and verifying composite signatures. This document makes no assumptions about what the component algorithms are, provided that their algorithm identifiers and signature generation and verification processes are defined.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 12 August 2022.

Copyright Notice

Copyright (c) 2022 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
1.1. Terminology	3
2. Composite Identifiers and Structures	4
2.1. Algorithm Identifier	5
2.2. Composite Keys	5
2.2.1. Key Usage Bits	5
2.3. Composite Signature	6
2.4. Encoding Rules	6
3. Composite Signature Processes	7
3.1. Composite Signature Generation Process	7
3.2. Composite-OR Signature Generation Process	8
3.3. Composite Signature Verification Process	9
3.4. Composite-OR Signature Verification	11
3.4.1. Composite-OR Legacy Mode	11
4. In Practice	12
4.1. Cryptographic protocols	13
5. IANA Considerations	14
6. Security Considerations	14
6.1. Policy for Deprecated and Acceptable Algorithms	14
7. Appendices	14
7.1. ASN.1 Module	14
7.2. Intellectual Property Considerations	16
8. Contributors and Acknowledgements	16
8.1. Making contributions	17
9. Normative References	17
Authors' Addresses	18

1. Introduction

During the transition to post-quantum cryptography, there will be uncertainty as to the strength of cryptographic algorithms; we will no longer fully trust traditional cryptography such as RSA, Diffie-Hellman, DSA and their elliptic curve variants, but we will also not fully trust their post-quantum replacements until they have had sufficient scrutiny. Unlike previous cryptographic algorithm migrations, the choice of when to migrate and which algorithms to migrate to, is not so clear. Even after the migration period, it may be advantageous for an entity's cryptographic identity to be composed of multiple public-key algorithms.

The deployment of composite signatures using post-quantum algorithms will face two challenges

- * Algorithm strength uncertainty: During the transition period, some post-quantum signature and encryption algorithms will not be fully trusted, while also the trust in legacy public key algorithms will start to erode. A relying party may learn some time after deployment that a public key algorithm has become untrustworthy, but in the interim, they may not know which algorithm an adversary has compromised.
- * Backwards compatibility: During the transition period, post-quantum algorithms will not be supported by all clients.

This document provides a mechanism to address algorithm strength uncertainty by building on `~~ reference draft-ounsworth-pq-composite-pubkeys ~~` by providing formats for encoding multiple signature values into existing public signature fields, as well as the process for validating a composite signature. Backwards compatibility is addressed via the Composite-OR mechanism described herein.

This document is intended for general applicability anywhere that digital signatures are used within PKIX and CMS structures.

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

The following terms are used in this document:

ALGORITHM: An information object class for identifying the type of cryptographic operation to be performed. This document is primarily concerned with algorithms for producing digital signatures.

BER: Basic Encoding Rules (BER) as defined in [X.690].

COMPONENT ALGORITHM: A single basic algorithm which is contained within a composite algorithm.

COMPOSITE ALGORITHM: An algorithm which is a sequence of two or more component algorithms, as defined in Section 2.

DER: Distinguished Encoding Rules as defined in [X.690].

LEGACY: For the purposes of this document, a legacy key or signature is a non-composite key or signature.

PUBLIC / PRIVATE KEY: The public and private portion of an asymmetric cryptographic key, making no assumptions about which algorithm.

SIGNATURE: A digital cryptographic signature, making no assumptions about which algorithm.

2. Composite Identifiers and Structures

In order for signatures to be composed of multiple algorithms, we define encodings consisting of a sequence of signature primitives (aka "component algorithms") such that these structures can be used as a drop-in replacement for existing signature fields such as those found in PKCS#10 [RFC2986], CMP [RFC4210], X.509 [RFC5280], CMS [RFC5652].

This section defines the following structures:

- * The id-alg-composite is an AlgorithmIdentifier identifying a composite signature object.

The sa-CompositeSignature AlgorithmIdentifier and the corresponding CompositeParams identify the algorithm(s) used in a composite signature.

- * The CompositeSignatureValue, carries a sequence of signatures that are generated by a CompositePrivateKey, and can be verified with the corresponding CompositePublicKey.

EDNOTE 2: the choice to define composite algorithm parameters as a sequence inside the existing fields avoids the exponential proliferation of OIDs that are needed for each combination of

signature algorithms in other schemes for achieving multi-key certificates. This scheme also naturally extends from 2-keypair to n-keypair keys and certificates.

EDNOTE 2a: We have heard community feedback that the ASN.1 structures presented here are too flexible in that allow arbitrary combinations of an arbitrary number of signature algorithms. The feedback is that this is too much of a "footgun" for implementors and sysadmins. We are working on an alternative formulation using ASN.1 information object classes that allow for compiling explicit pairs of algorithmIDs. We would love community feedback on which approach is preferred. See slide 30 of this presentation: <https://datatracker.ietf.org/meeting/interim-2021-lamps-01/materials/slides-interim-2021-lamps-01-sessa-position-presentation-by-mike-ounsworth-00.pdf>

2.1. Algorithm Identifier

The following object identifier is used for identifying a composite signature. Additional encoding information is provided below for each of these objects.

```
id-alg-composite OBJECT IDENTIFIER ::= {
    iso(1) identified-organization(3) dod(6) internet(1) private(4)
    enterprise(1) OpenCA(18227) Algorithms(2) id-alg-composite(1) }
```

EDNOTE 3: this is a temporary OID for the purposes of prototyping. We are requesting IANA to assign a permanent OID, see Section 5.

2.2. Composite Keys

A Composite signature MUST be associated with a Composite public key as defined in `reference draft-ounsworth-pq-composite-pubkey`.

2.2.1. Key Usage Bits

For protocols such as X.509 [RFC5280] that specify key usage along with the public key, then the composite public key associated with a composite signature MUST have a signing-type key usage.

If the keyUsage extension is present in a Certification Authority (CA) certificate that indicates id-composite-key, then any combination of the following values MAY be present:

```
digitalSignature;
nonRepudiation;
keyCertSign; and
cRLSign.
```


If the keyUsage extension is present in an End Entity (EE) certificate that indicates id-composite-key, then any combination of the following values MAY be present:

digitalSignature; and
nonRepudiation;

2.3. Composite Signature

The ASN.1 algorithm object for a composite signature is:

```
sa-CompositeSignature SIGNATURE-ALGORITHM ::= {  
  IDENTIFIER id-alg-composite  
  VALUE CompositeSignatureValue  
  PARAMS TYPE CompositeParams ARE required  
  PUBLIC-KEYS { pk-Composite }  
  SMIME-CAPS { IDENTIFIED BY id-alg-composite } }  
}
```

The following algorithm parameters MUST be included:

CompositeParams ::= SEQUENCE SIZE (2..MAX) OF AlgorithmIdentifier

The signature's CompositeParams sequence MUST contain the same component algorithms listed in the same order as in the associated CompositePrivateKey and CompositePublicKey.

The output of the composite signature algorithm is the DER encoding of the following structure:

CompositeSignatureValue ::= SEQUENCE SIZE (2..MAX) OF BIT STRING

Where each BIT STRING within the SEQUENCE is a signature value produced by one of the component keys. It MUST contain one signature value produced by each component algorithm, and in the same order as in the associated CompositeParams object.

The choice of SEQUENCE OF BIT STRING, rather than for example a single BIT STRING containing the concatenated signature values, is to gracefully handle variable-length signature values by taking advantage of ASN.1's built-in length fields.

2.4. Encoding Rules

Many protocol specifications will require that composite signature data structures be represented by an octet string or bit string.

When an octet string is required, the DER encoding of the composite data structure SHALL be used directly.

When a bit string is required, the octets of the DER encoded composite data structure SHALL be used as the bits of the bit string, with the most significant bit of the first octet becoming the first bit, and so on, ending with the least significant bit of the last octet becoming the last bit of the bit string.

In the interests of simplicity and avoiding compatibility issues, implementations that parse these structures MAY accept both BER and DER.

3. Composite Signature Processes

This section specifies the processes for generating and verifying composite signatures.

This process addresses algorithm strength uncertainty by providing the verifier with parallel signatures from all the component signature algorithms; thus breaking the composite signature would require breaking all of the component signatures.

3.1. Composite Signature Generation Process

Generation of a composite signature involves applying each component algorithm's signature process to the input message according to its specification, and then placing each component signature value into the CompositeSignatureValue structure defined in Section 2.3.

The following process is used to generate composite signature values.

Input:

K1, K2, ..., Kn Private keys for the n component signature algorithms, a CompositePrivateKey
M Message to be signed, an octet string

Output:

S The signatures, a CompositeSignatureValue

Signature Generation Process:

1. Generate the n component signatures independently, according to their algorithm specifications.

 for i := 1 to n
 Si := Sign(Ki, M)
2. Encode each component signature S1, S2, ..., Sn into a BIT STRING according to its algorithm specification.

 S ::= Sequence { S1, S2, ..., Sn }
3. Output S

Since recursive composite public keys are disallowed in ~~ Reference draft-ounsworth-pq-composite-pubkeys sec-composite-pub-keys ~~, no component signature may itself be a composite; ie the signature generation process MUST fail if one of the private keys K1, K2, ..., Kn is a composite with the OID id-alg-composite.

A composite signature MUST produce and include in the output a signature value for every component key in the corresponding CompositePrivateKey. For this mode, please see Composite-OR in section Section 3.2.

3.2. Composite-OR Signature Generation Process

EDNOTE: This section was written with the intention of keeping the primary Composite OID reserved for the simple and strict mode; if you want to do either a simple OR, or a custom policy then we have given a different OID. We are still debating whether this is useful to specify at issuing time, or whether this is adding needless complexity to the draft.

If the algorithm ID of the public key associated with this signature is id-composite-or-key then the signer MAY use only a subset of the component keys and therefore produce fewer signatures than the number of component keys.

Composite-OR signature generation uses the same structures and algorithms as Composite, with the difference that the signature generation process may emit a null instead of a signature value in step 1 for one or more component algorithms. A Composite-OR signature MUST NOT be entirely null; it must contain at least one valid signature.

The design intent of this mode is to support migration scenarios where an end entity has been issued keys on algorithms that either itself or the peer with which it is communicating do not (yet) support. This design allows for both the mode where the signer omits signatures that it knows its peer cannot process in order to save bandwidth and performance, and the mode where it includes all component signatures and allows the verifier to choose how many to verify. The latter is RECOMMENDED for signatures that need both short-term backwards compatibility as well as long-term security.

EDNOTE: Do we want to allow a Composite-OR with only a single signature to produce non-composite signatureAlgorithm and signatureValue as per [RFC5280]? Advantages: bandwidth savings of an extra OID and some sequences with one element. Disadvantages: ambiguous whether a signature is traditional or composite until you look at the corresponding public key.

3.3. Composite Signature Verification Process

Verification of a composite signature involves applying each component algorithm's verification process according to its specification.

In the absence of an application profile specifying otherwise, compliant applications MUST output "Valid signature" (true) if and only if all component signatures were successfully validated, and "Invalid signature" (false) otherwise.

The following process is used to perform this verification.

Input:

P Signer's composite public key
M Message whose signature is to be verified, an octet string
S Composite Signature to be verified
A Composite Algorithm identifier

Output:

Validity "Valid signature" (true) if the composite signature
 is valid, "Invalid signature" (false) otherwise.

Signature Verification Procedure::

1. Parse P, S, A into the component public keys, signatures,
and algorithm identifiers

P1, P2, ..., Pn := Desequence(P)
S1, S2, ..., Sn := Desequence(S)
A1, A2, ..., An := Desequence(A)

If Error during Desequencing, or the three sequences have
different numbers of elements, or any of the public keys P1, P2, ..., Pn or
algorithm identifiers A1, A2, ..., An are composite with the OID
id-alg-composite then output "Invalid signature" and stop.

2. Check each component signature individually, according to its
algorithm specification.
If any fail, then the entire signature validation fails.

for i := 1 to n
 if not verify(Pi, M, Si), then
 output "Invalid signature"

if all succeeded, then
 output "Valid signature"

Since recursive composite public keys are disallowed in ~~ Reference
draft-ounsworth-pq-composite-keys sec-composite-pub-keys ~~, no
component signature may be composite; ie the signature verification
procedure MUST fail if any of the public keys P1, P2, ..., Pn or
algorithm identifiers A1, A2, ..., An are composite with the OID id-
alg-composite.

3.4. Composite-OR Signature Verification

EDNOTE: This section was written with the intention of keeping the primary Composite OID reserved for the simple and strict mode; if you want to do either a simple OR, or a custom policy then we have given a different OID. We are still debating whether this is useful to specify at issuing time, or whether this is adding needless complexity to the draft.

When the public key associated with the signature being verified has algorithm id-composite-or-key, then an alternate verification processes MAY be used, at the discretion of the implementor. In this section we provide some examples of alternate verification processes.

If the signature is a traditional (non-composite) algorithm and value or a composite signature with a single component, then it MAY be considered valid if it verifies under one of the component keys.

If the signature is composite, then the implementor MAY implement policy for which combinations are acceptable.

EDNOTE: Does this mean Composite-OR end entity certificates need to be issued by a PKI that is marked as Composite-OR all the way to the top so that verifiers that do not support all the algorithms don't fail? Need to think more about the security implications of allowing a Composite-or in an end entity cert implicitly turning all Composite algIDs into Composite-or algIDs in its cert chain.

EDNOTE: Do we need to specify the semantics of verifying an "n of m" subset signature? I suspect that specifying this in general will be a rat's nest of edge cases, so I propose to "leave this to the implementor".

3.4.1. Composite-OR Legacy Mode

The Composite-OR Legacy Mode is provided to facilitate migration by allowing existing PKI entities (including root CAs, intermediate CAs, and end entities) to have their existing keys re-certified inside a Composite-OR structure along with Post-Quantum keys, and for signatures made by that key prior to the migration to remain valid. Note that Composite-OR Legacy Mode is only provided for signature verification, and not for signature generation; legacy signatures SHOULD NOT be produced from a Composite key.

EDNOTE: to further solidify this, we could add a clause that Legacy Mode signatures are to fail if the signature was produced after notBefore date of the Composite-OR certificate?

In Composite-OR Legacy Mode, a legacy signature algorithm and legacy signature value MAY be validated against a Composite-OR public key. The legacy signature algorithm is to be interpreted by the verifier as a sa-CompositeSignature with CompositeParams in the following way:

```
CompositeParams {legacyAlgorithmIdentifier, null, ..., null}
```

with the correct number of nulls to match the Composite-OR public key that the signature is being verified against. For the purposes of a signature validation under Composite-OR Legacy Mode, a null AlgorithmIdentifier is considered to be a match for the corresponding algorithm in the Composite-OR public key.

The legacy signature value is to be interpreted by the verifier as a sa-CompositeSignature with CompositeParams in the following way:

```
CompositeSignatureValue {legacySignatureValue, null, ..., null}
```

with the correct number of nulls to match the Composite-OR public key that the signature is being verified against. The verification algorithm in section Section 3.4 applies.

Security consideration: when implementing Composite-OR Legacy Mode, it is important to catch the edge case of {null, null, ..., null} for both AlgorithmIdentifier and SignatureValue and return Invalid Signature.

It is RECOMMENDED that Composite-OR Legacy Mode be implemented as an optional mode in the verifier that can be enabled or disabled by runtime configuration or policy.

EDNOTE: the signing public key is often identified in the signed document by issuer+serialNumber or by an SKI containing a hash of the public key value. Might need X.509 extensions identifying the SKI of the legacy cert it's replacing?

4. In Practice

This section addresses practical issues of how this draft affects other protocols and standards.

~~~ BEGIN EDNOTE 10~~~

EDNOTE 10: Possible topics to address:

- \* The size of these certs and cert chains.

- \* In particular, implications for (large) composite keys / signatures / certs on the handshake stages of TLS and IKEv2.
- \* If a cert in the chain is a composite cert then does the whole chain need to be of composite Certs?
- \* We could also explain that the root CA cert does not have to be of the same algorithms. The root cert SHOULD NOT be transferred in the authentication exchange to save transport overhead and thus it can be different than the intermediate and leaf certs.
- \* We could talk about overhead (size and processing).
- \* We could also discuss backwards compatibility.
- \* We could include a subsection about implementation considerations.

~~~ END EDNOTE 10~~~

4.1. Cryptographic protocols

This section talks about how protocols like (D)TLS and IKEv2 are affected by this specifications. It will not attempt to solve all these problems, but it will explain the rationale, how things will work and what open problems need to be solved. Obvious issues that need to be discussed.

- * How does the protocol declare support for composite signatures? TLS has hooks for declaring support for specific signature algorithms, however it would need to be extended, because the client would need to declare support for both the composite infrastructure, as well as for the various component signature algorithms.
- * How does the protocol use the multiple keys. The obvious way would be to have the server sign using its composite public key; is this sufficient.
- * Overhead; including certificate size, signature processing time, and size of the signature.
- * How to deal with crypto protocols that use public key encryption algorithms; this document only lists how to work with signature algorithms. Encoding composite public keys is straightforward; encoding composite ciphertexts is less so - we decided to put that off to another draft.

5. IANA Considerations

The ASN.1 module OID is TBD. The id-alg-composite OID is to be assigned by IANA. The authors suggest that IANA assign an OID on the id-pkix arc:

```
id-alg-composite OBJECT IDENTIFIER ::= {  
    iso(1) identified-organization(3) dod(6) internet(1) security(5)  
    mechanisms(5) pkix(7) algorithms(6) composite(??) }
```

6. Security Considerations

6.1. Policy for Deprecated and Acceptable Algorithms

Traditionally, a public key, certificate, or signature contains a single cryptographic algorithm. If and when an algorithm becomes deprecated (for example, RSA-512, or SHA1), it is obvious that structures using that algorithm are implicitly revoked.

In the composite model this is less obvious since a single public key, certificate, or signature may contain a mixture of deprecated and non-deprecated algorithms. Moreover, implementers may decide that certain cryptographic algorithms have complementary security properties and are acceptable in combination even though neither algorithm is acceptable by itself.

Specifying a modified verification algorithm to handle these situations is beyond the scope of this draft, but could be desirable as the subject of an application profile document, or to be up to the discretion of implementers.

2. Check policy to see whether A1, A2, ..., An constitutes a valid combination of algorithms.

```
if not checkPolicy(A1, A2, ..., An), then  
    output "Invalid signature"
```

While intentionally not specified in this document, implementors should put careful thought into implementing a meaningful policy mechanism within the context of their signature verification engines, for example only algorithms that provide similar security levels should be combined together.

7. Appendices

7.1. ASN.1 Module

<CODE STARTS>

Composite-Signatures-2019
{ TBD }

DEFINITIONS IMPLICIT TAGS ::= BEGIN

EXPORTS ALL;

IMPORTS

PUBLIC-KEY, SIGNATURE-ALGORITHM
FROM AlgorithmInformation-2009 -- RFC 5912 [X509ASN1]
{ iso(1) identified-organization(3) dod(6) internet(1)
security(5) mechanisms(5) pkix(7) id-mod(0)
id-mod-algorithmInformation-02(58) }

SubjectPublicKeyInfo
FROM PKIX1Explicit-2009
{ iso(1) identified-organization(3) dod(6) internet(1)
security(5) mechanisms(5) pkix(7) id-mod(0)
id-mod-pkix1-explicit-02(51) }

OneAsymmetricKey
FROM AsymmetricKeyPackageModuleV1
{ iso(1) member-body(2) us(840) rsadsi(113549) pkcs(1)
pkcs-9(9) smime(16) modules(0)
id-mod-asymmetricKeyPkgV1(50) } ;

--
-- Object Identifiers
--

id-alg-composite OBJECT IDENTIFIER ::= { TBD }

--
-- Public Key
--

pk-Composite PUBLIC-KEY ::= {
IDENTIFIER id-alg-composite
KEY CompositePublicKey
PARAMS ARE absent
PRIVATE-KEY CompositePrivateKey
}

CompositePublicKey ::= SEQUENCE SIZE (2..MAX) OF SubjectPublicKeyInfo

CompositePrivateKey ::= SEQUENCE SIZE (2..MAX) OF OneAsymmetricKey

```
--
-- Signature Algorithm
--

sa-CompositeSignature SIGNATURE-ALGORITHM ::= {
    IDENTIFIER id-alg-composite
    VALUE CompositeSignatureValue
    PARAMS TYPE CompositeParams ARE required
    PUBLIC-KEYS { pk-Composite }
    SMIME-CAPS { IDENTIFIED BY id-alg-composite } }

CompositeParams ::= SEQUENCE SIZE (2..MAX) OF AlgorithmIdentifier

CompositeSignatureValue ::= SEQUENCE SIZE (2..MAX) OF BIT STRING

END

<CODE ENDS>
```

7.2. Intellectual Property Considerations

The following IPR Disclosure relates to this draft:

<https://datatracker.ietf.org/ipr/3588/>

8. Contributors and Acknowledgements

This document incorporates contributions and comments from a large group of experts. The Editors would especially like to acknowledge the expertise and tireless dedication of the following people, who attended many long meetings and generated millions of bytes of electronic mail and VOIP traffic over the past year in pursuit of this document:

John Gray (Entrust), Serge Mister (Entrust), Scott Fluhrer (Cisco Systems), Panos Kampanakis (Cisco Systems), Daniel Van Geest (ISARA), Tim Hollebeek (Digicert), and Francois Rousseau.

We are grateful to all, including any contributors who may have been inadvertently omitted from this list.

This document borrows text from similar documents, including those referenced below. Thanks go to the authors of those documents.

"Copying always makes things easier and less error prone" -
[RFC8411].

8.1. Making contributions

Additional contributions to this draft are welcome. Please see the working copy of this draft at, as well as open issues at:

<https://github.com/EntrustCorporation/draft-ounsworth-composite-sigs>

9. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC2986] Nystrom, M. and B. Kaliski, "PKCS #10: Certification Request Syntax Specification Version 1.7", RFC 2986, DOI 10.17487/RFC2986, November 2000, <<https://www.rfc-editor.org/info/rfc2986>>.
- [RFC4210] Adams, C., Farrell, S., Kause, T., and T. Mononen, "Internet X.509 Public Key Infrastructure Certificate Management Protocol (CMP)", RFC 4210, DOI 10.17487/RFC4210, September 2005, <<https://www.rfc-editor.org/info/rfc4210>>.
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <<https://www.rfc-editor.org/info/rfc5280>>.
- [RFC5652] Housley, R., "Cryptographic Message Syntax (CMS)", STD 70, RFC 5652, DOI 10.17487/RFC5652, September 2009, <<https://www.rfc-editor.org/info/rfc5652>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8411] Schaad, J. and R. Andrews, "IANA Registration for the Cryptographic Algorithm Object Identifier Range", RFC 8411, DOI 10.17487/RFC8411, August 2018, <<https://www.rfc-editor.org/info/rfc8411>>.
- [X.690] ITU-T, "Information technology - ASN.1 encoding Rules: Specification of Basic Encoding Rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER)", ISO/IEC 8825-1:2015, November 2015.

Authors' Addresses

Mike Ounsworth
Entrust Limited
2500 Solandt Road -- Suite 100
Ottawa, Ontario K2K 3G5
Canada

Email: mike.ounsworth@entrust.com

Massimiliano Pala
CableLabs

Email: director@openca.org

TBD
Internet-Draft
Updates: 5480 (if approved)
Intended status: Standards Track
Expires: February 14, 2020

T. Ito
SECOM CO., LTD.
S. Turner
sn3rd
August 13, 2019

Clarifications for Elliptic Curve Cryptography Subject Public Key
Information
draft-turner-5480-ku-clarifications-01

Abstract

This document updates RFC 5480 to specify semantics for the keyEncipherment and dataEncipherment key usage bits when used in certificates that support Elliptic Curve Cryptography.

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 February 14, 2020.

Copyright Notice

Copyright (c) 2019 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

| | |
|--------------------------------------|---|
| 1. Introduction | 2 |
| 2. Terminology | 2 |
| 3. Updates to Section 3 | 2 |
| 4. Security Considerations | 3 |
| 5. IANA Considerations | 3 |
| 6. Normative References | 3 |
| Authors' Addresses | 3 |

1. Introduction

[RFC5480] specifies the syntax and semantics for the Subject Public Key Information field in certificates that support Elliptic Curve Cryptography. As part of these semantics, it defines what combinations are permissible for the values of the key usage extensions [RFC5280]. [RFC5480] specifies 7 of the 9 values; it makes no mention of keyEncipherment and dataEncipherment key usage bits. This document corrects this omission, by updating Section 3 of [RFC5480] to make it clear that neither keyEncipherment nor the dataEncipherment key usage bits are set for key agreement algorithms.

2. 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 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Updates to Section 3

If the keyUsage extension is present in a certificate that indicates id-ecPublicKey as algorithm of AlgorithmIdentifier [RFC2986] in SubjectPublicKeyInfo, then following values MUST NOT be present:

keyEncipherment; and
dataEncipherment.

If the keyUsage extension is present in a certificate that indicates id-ecDH or id-ecMQV in SubjectPublicKeyInfo, then the following values also MUST NOT be present:

keyEncipherment; and
dataEncipherment.

4. Security Considerations

This document introduces no new security considerations beyond those found in [RFC5480].

5. IANA Considerations

This document makes no request of IANA.

6. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC2986] Nystrom, M. and B. Kaliski, "PKCS #10: Certification Request Syntax Specification Version 1.7", RFC 2986, DOI 10.17487/RFC2986, November 2000, <<https://www.rfc-editor.org/info/rfc2986>>.
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <<https://www.rfc-editor.org/info/rfc5280>>.
- [RFC5480] Turner, S., Brown, D., Yiu, K., Housley, R., and T. Polk, "Elliptic Curve Cryptography Subject Public Key Information", RFC 5480, DOI 10.17487/RFC5480, March 2009, <<https://www.rfc-editor.org/info/rfc5480>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

Authors' Addresses

Tadahiko Ito
SECOM CO., LTD.

Email: tadahiko.ito.public@gmail.com

Sean Turner
sn3rd

Email: sean@sn3rd.com