

Internet Draft
Intended status: Informational
Expires: August 2009

A. Santoni
Actalis S.p.A.
February 28, 2009

Syntax for binding documents with time stamps

[draft-santoni-timestampeddata-05.txt](#)

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/ietf/1id-abstracts.txt>

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

Copyright Notice

Copyright (c) 2009 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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may

not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Abstract

This document describes an envelope which can be used to bind a file (not necessarily protected by means of cryptographic techniques) with one or more time-stamp tokens obtained for that file, where "time-stamp token" has the meaning defined in [RFC 3161](#) or its successors. Additional types of temporal evidence are also allowed.

The proposed envelope is based on the Cryptographic Message Syntax as defined in [RFC 3852](#).

Conventions used in this document

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](#) [[KEYWORDS](#)].

The terms "document" and "file" are used interchangeably.

Table of Contents

1.	Introduction.....	2
2.	Syntax for TimeStampedData.....	3
3.	Compliance requirements.....	6
4.	Recommended processing.....	6
4.1.	Generating a new TimeStampedData structure.....	6
4.2.	Verifying an existing TimeStampedData structure.....	7
4.3.	Extending the validity of an existing TimeStampedData structure.....	8
5.	Security Considerations.....	9
6.	IANA Considerations.....	9
7.	Normative references.....	10
8.	Informative references.....	10
	APPENDIX A: ASN.1 Module.....	11
	Acknowledgments.....	13
	Author's Address.....	13

[1. Introduction](#)

Time-stamping has become the standard technique for proving the existence of a document before a certain point in time. Several legislations around the world embrace the concept and provide for time-stamping services, mainly for the purpose of extending the validity of signed documents. However, while time-stamping enhances digital signatures, its value does not depend on them. It can clearly be useful to time-stamp a document even if this is not signed. And it can also be useful, or even mandatory in some cases, to time-stamp a signed document in its entirety, regardless of how many signatures it

contains.

Santoni

Expires August 28, 2009

[Page 2]

When a time-stamp is related to a digital signature, there already exists a way to keep the two pieces together: [RFC 3161](#) describes how one or more TimeStampTokens can be included in a SignerInfo structure as unsigned attributes. On the other hand, there is no standard way to keep together a time-stamped document, whether signed or not, and the related time-stamps.

In such cases two approaches are typically being adopted:

- o time-stamps are kept as separate files (keeping track of what time-stamps belong to what documents is up to the user);
- o an ad hoc solution is adopted for specific applications, like e.g. a ZIP archive or a proprietary "envelope" of some kind.

Both solutions impede interoperability, the objective of this memo.

This document describes a simple syntax for binding one document (actually, any kind of file) to the corresponding temporal evidence, this latter being typically represented by one or more [RFC 3161](#) TimeStampTokens. Additional types of temporal evidence, like e.g. an [RFC 4998](#) EvidenceRecord, are also supported via an "open" syntax. However, for the sake of interoperability, the emphasis is given to TimeStampTokens.

The proposed syntax is broadly based on the Cryptographic Message Syntax (CMS) defined in [RFC 3852](#) [[CMS](#)].

2. Syntax for TimeStampedData

The proposed data structure is called TimeStampedData and it is based on the ContentInfo envelope defined in [[CMS](#)]:

```
ContentInfo ::= SEQUENCE {  
    contentType ContentType,  
    content [0] EXPLICIT ANY DEFINED BY contentType }
```

```
ContentType ::= OBJECT IDENTIFIER
```

While CMS defines six content types (data, signed-data, enveloped-data, digested-data, encrypted-data, and authenticated-data), this memo defines an additional content type, timestamped-data, identified by the following OID:

```
id-ct-timestampedData OBJECT IDENTIFIER ::= {
```

```
iso(1) member-body(2) us(840) rsadsi(113549) pkcs(1)
pkcs9(9) id-smime(16) id-ct(1) 31 }
```

Santoni

Expires August 28, 2009

[Page 3]

This particular OID signals that the content field of the ContentInfo has the following syntax:

```
TimeStampedData ::= SEQUENCE {
    version                INTEGER { v1(1) },
    dataUri                IA5String OPTIONAL,
    metaData               MetaData OPTIONAL,
    content                OCTET STRING OPTIONAL,
    temporalEvidence       Evidence
}

MetaData ::= SEQUENCE {
    hashProtected          BOOLEAN,
    fileName               UTF8String OPTIONAL,
    mediaType              IA5String OPTIONAL,
    otherMetaData          Attributes OPTIONAL
}

Attributes ::=
    SET SIZE(1..MAX) OF Attribute -- according to RFC 3852

Evidence ::= CHOICE {
    tstEvidence            [0] SEQUENCE SIZE(1..MAX) OF
                               TimeStampTokenEvidence, -- see RFC 3161
    ersEvidence            [1] EvidenceRecord,          -- see RFC 4998
    otherEvidence          [2] OtherEvidence
}

OtherEvidence ::= SEQUENCE {
    oeType                 OBJECT IDENTIFIER,
    oeValue                 ANY DEFINED BY oeType }

TimeStampTokenEvidence ::=
    SEQUENCE SIZE(1..MAX) OF TimeStampAndCRL

TimeStampAndCRL ::= SEQUENCE {
    timeStamp              TimeStampToken, -- according to RFC 3161
    crl                    CertificateList OPTIONAL -- according to RFC 5280
}
```

The version field contains the version number of the TimeStampedData syntax. It SHALL be 1 for this version of the document.

The dataUri field contains a URI reference conforming to [[URI](#)]. When the content field is absent, dataUri MUST be present and contain an URI allowing to retrieve the document which was time-stamped (unless the document is later moved). When the content field is present, this

field MAY also be present.

The metaData field contains metadata related to the document which was time-stamped, if applicable. In particular:

The `hashProtected` field indicates whether the subsequent fields have been included in the computation of the digest within the first `TimeStampToken` (see further on). This makes it possible to detect a subsequent alteration of the metadata.

The `fileName` field contains the original filename of the document which was time-stamped.

The `mediaType` field contains a media type/subtype and possible parameters for the time-stamped document, according to [\[MIME\]](#). This information may help decide how to "open" or deal with the time-stamped document.

The `otherMetaData` field contains further attributes of the time-stamped document (like e.g. a description, claimed author, etc.), where each attribute is specified by an object identifier and a corresponding set of values, as described in [\[CMS\]](#). When this field is present, it MUST contain at least one Attribute.

Within the `metaData` field, if present, at least one of the `fileName`, `mediaType` and `otherMetaData` subfields MUST be present.

The `content` field, when present, carries the entire contents, in its original format and encoding, of the document which was time-stamped. This can actually be any kind of data like e.g. a text document, an executable, a movie, a message, etc. The omission of the `content` field makes it possible to bind the temporal evidence to external data. In such a case, the temporal evidence is computed as though the `content` field was present.

The `temporalEvidence` field carries the evidence that the time-stamped document did exist before a certain point in time. Several types of evidence are allowed, but compliant applications are only required to support the [RFC 3161](#) type, namely the `tstEvidence` choice.

The `tstEvidence` sequence MUST contain at least one element of type `TimeStampTokenEvidence`.

The `TimeStampTokenEvidence` sequence MUST contain at least one element of type `TimeStampAndCRL`.

The elements of the `TimeStampTokenEvidence` sequence MUST conform to the following rule:

- o if the `metaData` field is absent or the value of its `hashProtected` field is `FALSE`, then the `TimeStampToken` within the first element SHALL be computed over the value octets of the `content` field (if this field is absent, use the octets retrieved via the `dataUri` field);

Santoni

Expires August 28, 2009

[Page 5]

- o otherwise (the metaData field is present and the value of its hashProtected field is TRUE), the TimeStampToken within the first element SHALL be computed over the concatenation of the following fields:
 - the value octets of the fileName field, when present;
 - the value octets of the mediaType field, when present;
 - the value octets of the otherMetaData field, when present;
 - the value octets of the content field (if this field is absent, use the octets retrieved via the dataUri field);
- o the TimeStampToken within the second element SHALL be computed over the first element;
- o the TimeStampToken within each subsequent element SHALL be computed over its preceding element in the sequence.

Within the TimeStampAndCRL construct, the optional crl field carries a suitable CRL (Certificate Revocation List) demonstrating that the certificate of the TSA that issued the TimeStampToken was not revoked at the time when the subsequent element in the TimeStampTokenEvidence sequence was added. See the Security Considerations section for further discussion on this topic.

3. Compliance requirements

Compliant applications MUST support at least the [RFC 3161](#)-based type of evidence (i.e. the tstEvidence CHOICE).

4. Recommended processing

This section is focused on the [RFC 3161](#)-based type of evidence.

4.1. Generating a new TimeStampedData structure

In this case, applications are supposed to behave like follows:

- o populate the version field with the integer value v1(1);
- o if a self-contained envelope is to be generated, always populate the content field with the content of the file in its original format and encoding; depending on the application, the dataUri field may also be added;
- o otherwise (a detached envelope is to be generated), always populate the dataUri field with the URI of the time-stamped

document (e.g. <http://foo.example.com/Contract12345.pdf>);
using an absolute URI or a relative reference depends on the
application;

- o if the metaData field is being added, decide on the value of its hashProtected field; set its value to TRUE if the application needs the remaining fields of the metaData construct to be hash-protected as described in [Section 2](#); otherwise set it to FALSE;
- o if the metaData field is being added, optionally populate the fileName field (e.g. "Contract12345.pdf"), the mediaType field with a suitable media type/subtype and possible parameters according to [\[MIME\]](#), and the otherMetaData field, depending on the application;
- o select a suitable one-way hash function and compute a hash value using that function over the content, or the concatenation of the metaData and the content, as described in [Section 2](#); this hash value will then be used for requesting the first TimeStampToken;
- o obtain the first temporal evidence from a TSA and add it to the temporalEvidence field;
- o insert the TimeStampedData into a ContentInfo structure, with the id-ct-timestampedData OID in the contentType field;
- o DER-encode the ContentInfo structure and save it with a reasonable file name (e.g. derived from the name of the time-stamped file).

[4.2](#). Verifying an existing TimeStampedData structure

In this case, applications are supposed to behave like follows:

- o check that the contentType field of the ContentInfo structure has the expected value (id-ct-timestampedData) in its contentType field; then, extract the inner TimeStampedData structure and continue processing;
- o check the version field (it should be v1);
- o check that the temporalEvidence field is not empty;
- o check if the content is present; if it is not, use the dataUri field to retrieve the file;
- o open the first element of the TimeStampTokenEvidence sequence, open the time-stamp token within it and use the hash function which was used to obtain it to re-compute the hash of the fields indicated in [Section 2](#); if the re-computed hash value matches the one within the time-stamp token, continue processing; otherwise, the TimeStampedData structure has been modified;
- o validate the temporalEvidence by checking that:

- each TimestampToken in the chain under examination does contain the correct digest value (according to the rule described in [Section 2](#)) and it was signed by a trusted TSA,

- the corresponding TSA signing certificate was not revoked at the time when the subsequent TimeStampToken has been issued, based on the associated CRL;
- o depending on the application, use the temporal evidence for whatever purpose the application was designed for;
- o depending on the application, show the dataUri, the fileName, the mediaType, the otherMetaData, and the temporal evidence to the user;
- o depending on the application, save the content to a separate file;
- o depending on the application, store at a different place the content which has been retrieved using the dataUri field, then update the dataUri field accordingly;
- o depending on the application, show the time-stamped file to the user, possibly by activating a suitable "viewer".

4.3. Extending the validity of an existing TimeStampedData structure

In this case, applications are supposed to behave like follows:

- o validate the TimeStampedData structure as described above;
- o select the time-stamp token from the last TimeStampAndCRL element in the chain under examination and obtain the latest available CRL for the corresponding TSA certificate (if this CRL is not fresh enough, wait until the next one is available), then store it in the TimeStampAndCRL element;
- o instantiate a new TimeStampAndCRL element and obtain a new time-stamp token computed over the previous element, according to the rule described in [Section 2](#); insert the new time-stamp token into the new TimeStampAndCRL element, then append this latter to the end of the current chain.

See the "Security Considerations" section for further discussion on extending the validity of an existing TimeStampedData structure.

Santoni

Expires August 28, 2009

[Page 8]

5. Security Considerations

When the `metaData` field is present and the `hashProtected` sub-field is set to `TRUE`, the metadata are also included in the computation of the digest within the first time-stamp token, so that any subsequent alteration of the metadata will be easily detected. However, the integrity of hash-protected metadata does not imply that the metadata were correct at the time when the `TimeStampedData` object was created. That can only be inferred by other means (e.g. from context). For instance, when `TimeStampedData` objects are created by an archival service provider, it may be reasonable to assume that the metadata are correct at creation time. Instead, when a `TimeStampedData` object is received from an unknown party, the recipient cannot safely assume that the metadata are correct, lacking further information.

In general, a time-stamp token should not be considered valid after the certificate of the issuing TSA is expired (also depending on the legislation and the policy under which the TSA operates). However, a time-stamp token can itself be time-stamped to extend the validity of the TSA's signature. By repeatedly applying this technique, a whole chain of time-stamp tokens can be grown to extend the validity of the first one ad libitum. This approach can thus be adopted to extend the validity of a `TimeStampedData` structure beyond the expiry date of the first `TimeStampToken` within it, by adding further elements to the `TimeStampTokenEvidence` sequence according to the rule described in [Section 2](#). Of course, each additional `TimeStampToken` must be added in a timely manner (before the previous one is expired or has been revoked).

The validity extension technique described above requires that the TSA signing certificates can still be verified long after they have expired, typically by checking a CRL. This latter must be captured at the suitable time, because expired certificates are typically removed from the CRL regardless of their being revoked. The `TimeStampAndCRL` construct allows adding a CRL next to the related `TimeStampToken`, so that the TSA certificate will still be verifiable at any later time. The CRL must be captured at the time when another element is about to be added to the `TimeStampTokenEvidence` sequence, or even later -- to allow for a last-minute revocation request to be processed by the CA (see the discussion about "grace periods" in [\[CADES\]](#)).

6. IANA Considerations

None.

Santoni

Expires August 28, 2009

[Page 9]

7. Normative references

- [KEYWORDS] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
<http://www.ietf.org/rfc/rfc2119.txt>
- [TSP] Adams, C., Cain, P., Pinkas, D. and R. Zuccherato, "Internet X.509 Public Key Infrastructure Time-Stamp Protocol (TSP)", [RFC 3161](#), August 2001.
<http://www.ietf.org/rfc/rfc3161.txt>
- [CMS] Housley, R., "Cryptographic Message Syntax (CMS)", [RFC 3852](#), July 2004.
<http://www.ietf.org/rfc/rfc3852.txt>
- [MIME] Borenstein, N., and N. Freed, "Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies", [RFC 2045](#), November 1996.
<http://www.ietf.org/rfc/rfc2045.txt>
- [ERS] Gondrom, T., Brandner, R., and Pordesch, U., "Evidence Record Syntax (ERS)", [RFC 4998](#), August 2007.
<http://www.ietf.org/rfc/rfc4998.txt>
- [PKIX1] 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.
<http://www.ietf.org/rfc/rfc5280.txt>
- [URI] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, [RFC 3986](#), January 2005.
<http://www.ietf.org/rfc/rfc3986.txt>

8. Informative references

- [CADES] Pinkas, D., Pope, N., and J. Ross, "CMS Advanced Electronic Signatures (CADES)", [RFC 5126](#), February 2008.
<http://www.ietf.org/rfc/rfc5126.txt>

Santoni

Expires August 28, 2009

[Page 10]

APPENDIX A: ASN.1 Module

The ASN.1 module contained in this appendix defines the structures that are needed to implement this specification. It is expected to be used in conjunction with the ASN.1 modules in [\[CMS\]](#), [\[TSP\]](#), [\[PKIX1\]](#), and [\[ERS\]](#).

TimeStampedDataModule

```
{ iso(1) member-body(2) us(840) rsadsi(113549) pkcs(1)
  pkcs-9(9) smime(16) modules(0) 35 }
```

DEFINITIONS IMPLICIT TAGS ::=

BEGIN

IMPORTS

-- Imports from [RFC 3852](#) [\[CMS\]](#)

Attribute

FROM CryptographicMessageSyntax2004

```
{ iso(1) member-body(2) us(840) rsadsi(113549)
  pkcs(1) pkcs-9(9) smime(16) modules(0) cms-2004(24) }
```

-- Imports from [RFC 3161](#) [\[TSP\]](#)

TimeStampToken

FROM PKIXTSP

```
{ iso(1) identified-organization(3) dod(6) internet(1)
  security(5) mechanisms(5) pkix(7) id-mod(0)
  id-mod-tsp(13)}
```

-- Imports from [RFC 5280](#) [\[PKIX1\]](#)

CertificateList

FROM PKIX1Explicit88

```
{ iso(1) identified-organization(3) dod(6) internet(1)
  security(5) mechanisms(5) pkix(7) id-mod(0)
  id-pkix1-explicit-88(18)}
```

-- Imports from [RFC 4998](#) [\[ERS\]](#)

EvidenceRecord

FROM ERS

```
{ iso(1) identified-organization(3) dod(6) internet(1)
  security(5) mechanisms(5) ltans(11) id-mod(0)
  id-mod-ers88(2) id-mod-ers88-v1(1) };
```

-- TimeStampedData Content Type and Object Identifier

```
id-ct-timestampedData OBJECT IDENTIFIER ::= {
  iso(1) member-body(2) us(840) rsadsi(113549) pkcs(1) pkcs9(9)
  id-smime(16) id-ct(1) 31 }
```

Santoni

Expires August 28, 2009

[Page 11]

```
TimeStampedData ::= SEQUENCE {
    version            INTEGER { v1(1) },
    dataUri            IA5String OPTIONAL,
    metaData           MetaData OPTIONAL,
    content            OCTET STRING OPTIONAL,
    temporalEvidence   Evidence
}

MetaData ::= SEQUENCE {
    hashProtected      BOOLEAN,
    fileName           UTF8String OPTIONAL,
    mediaType          IA5String OPTIONAL,
    otherMetaData      Attributes OPTIONAL
}

Attributes ::=
    SET SIZE(1..MAX) OF Attribute -- according to RFC 3852

Evidence ::= CHOICE {
    tstEvidence        [0] SEQUENCE SIZE(1..MAX) OF
                        TimeStampTokenEvidence, -- see RFC 3161
    ersEvidence        [1] EvidenceRecord,      -- see RFC 4998
    otherEvidence      [2] OtherEvidence
}

OtherEvidence ::= SEQUENCE {
    oeType             OBJECT IDENTIFIER,
    oeValue            ANY DEFINED BY oeType }

TimeStampTokenEvidence ::=
    SEQUENCE SIZE(1..MAX) OF TimeStampAndCRL

TimeStampAndCRL ::= SEQUENCE {
    timeStamp          TimeStampToken, -- according to RFC 3161
    crl               CertificateList OPTIONAL -- according to RFC 5280
}

END
```

Santoni

Expires August 28, 2009

[Page 12]

Acknowledgments

Thanks to Stephen Kent for encouraging the author in the early stages of this work.

Thanks to Russ Housley for reviewing this memo, suggesting useful amendments and assigning a value to the OIDs herein defined.

Thanks are also due to other people who reviewed this memo and helped improving it, but prefer not to be mentioned.

Author's Address

Adriano Santoni
Actalis S.p.A.
Via Taramelli 26
I-20124 Milano

E-mail: adriano.santoni@actalis.it

