Authors: C. Bormann   B. Gamari   H. Birkholz
         Universität Bremen TZI   Well-Typed   Fraunhofer SIT

Concise Binary Object Representation (CBOR) Tags for Time, Duration, and Period

Abstract

The Concise Binary Object Representation (CBOR, RFC 8949) is a data
format whose design goals include the possibility of extremely small
code size, fairly small message size, and extensibility without the
need for version negotiation.

In CBOR, one point of extensibility is the definition of CBOR tags.
RFC 8949 defines two tags for time: CBOR tag 0 (RFC3339 time as a
string) and tag 1 (Posix time as int or float). Since then,
additional requirements have become known. The present document
defines a CBOR tag for time that allows a more elaborate
representation of time, as well as related CBOR tags for duration
and time period. It is intended as the reference document for the
IANA registration of the CBOR tags defined.

The present version (-05) adds CDDL definitions; this should now
address all WGLC comments.

About This Document

This note is to be removed before publishing as an RFC.

Status information for this document may be found at https://datatracker.ietf.org/doc/draft-ietf-cbor-time-tag/.

Discussion of this document takes place on the CBOR Working Group
mailing list (mailto:cbor@ietf.org), which is archived at https://mailarchive.ietf.org/arch/browse/cbor/. Subscribe at https://www.ietf.org/mailman/listinfo/cbor/.

Source for this draft and an issue tracker can be found at https://github.com/cbor-wg/time-tag.

Status of This Memo

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

Copyright Notice

Table of Contents

1. Introduction

The Concise Binary Object Representation (CBOR, [[RFC8949](#)]) provides
for the interchange of structured data without a requirement for a
pre-agreed schema. RFC 8949 defines a basic set of data types, as
well as a tagging mechanism that enables extending the set of data
types supported via an IANA registry.

In CBOR, one point of extensibility is the definition of CBOR tags.
RFC 8949 defines two tags for time: CBOR tag 0 (RFC3339 time as a
string) and tag 1 (Posix time as int or float). Since then,
additional requirements have become known. The present document
defines a CBOR tag for time that allows a more elaborate
representation of time, as well as related CBOR tags for duration
and time period. It is intended as the reference document for the
IANA registration of the CBOR tags defined.

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 term "byte" is used in its now customary sense as a synonym for
"octet". Where bit arithmetic is explained, this document uses the
notation familiar from the programming language C (including C++14's
0bnnn binary literals), except that the operator "**" stands for
exponentiation.

CBOR diagnostic notation is defined in [Section 8](#) of [[RFC8949](#)] and
[Appendix G](#) of [[RFC8610](#)]. A machine-processable model of the data
structures defined in this specification is provided throughout the
text using the Concise Data Definition Language, CDDL [[RFC8610](#)];
[Appendix A](#) provides the collected model information.

2. Objectives

For the time tag, the present specification addresses the following objectives that go beyond the original tags 0 and 1:

  *Additional resolution for epoch-based time (as in tag 1). CBOR
   tag 1 only provides for integer and up to binary64 floating point
   representation of times, limiting resolution to approximately
   microseconds at the time of writing (and progressively becoming
   worse over time).

  *Indication of timescale. Tags 0 and 1 are for UTC; however, some
   interchanges are better performed on TAI. Other timescales may be
   registered once they become relevant (e.g., one of the proposed
   successors to UTC that might no longer use leap seconds, or a
   scale based on smeared leap seconds).

By incorporating a way to transport [IXDTF] suffix information
(Section 3.6, Section 3.7), additional indications can be provided
of intents about the interpretation of the time given, in particular
for future times. Intents might include information about time
zones, daylight savings times, preferred calendar representations,
etc.

Semantics not covered by this document can be added by registering
additional map keys for the map inside the tag, the specification
for which is referenced by the registry entry (Section 7.3,
Section 3). For example, map keys could be registered for:

  *Direct representation of natural platform time formats. Some
   platforms use epoch-based time formats that require some
   computation to convert them into the representations allowed by
   tag 1; these computations can also lose precision and cause
   ambiguities. (The present specification does not take a position
   on whether tag 1 can be "fixed" to include, e.g., Decimal or
   BigFloat representations. It does define how to use these
   representations with the extended time format.)

Additional tags are defined for durations and periods.

3. Time Format

An extended time is indicated by CBOR tag 1001, the content of which
is a map data item (CBOR major type 5). The map may contain integer
(major types 0 and 1) or text string (major type 3) keys, with the
value type determined by each specific key. For negative integer
keys and text string values of the key, implementations **MUST** ignore
key/value pairs they do not understand. Conversely, for unsigned
integer keys, implementations **MUST** signal as an error key/value

pairs they do not understand or implement (these are either "base time" or "critical", see below).

The map must contain exactly one unsigned integer key that specifies the "base time", and may also contain one or more negative integer or text-string keys, which may encode supplementary information.

Supplementary information may also be provided by additional unsigned integer keys that are explicitly defined to provide supplementary information ("critical"; as these are required to be understood, there can be no confusion with Base Time keys).

Negative integer and text string keys always supply supplementary information ("elective", and this will not be explicitly stated below).

Supplementary information may include:

  *a higher precision time offset to be added to the Base Time,

  *a reference timescale and epoch different from the default UTC
   and 1970-01-01

  *information about clock quality parameters, such as source,
   accuracy, and uncertainty

Additional keys can be defined by registering them in the Map Key Registry ([Section 7.3](#)). Future keys may add:

  *intent information such as timezone and daylight savings time,
   and/or possibly positioning coordinates, to express information
   that would indicate a local time.

This document does not define supplementary text keys. A number of both unsigned and negative-integer keys are defined in the following subsections.

A formal definition of Tag 1001 in CDDL is:

```
Etime = #6.1001(etime-detailed)

etime-framework = {
  uint => any ; at least one base time
  * (nint/text) => any ; elective supplementary information
  * uint => any ; critical supplementary information
}

etime-detailed = ({
  $$ETIME-BASETIME
  ClockQuality-group
  * $$ETIME-ELECTIVE
  * $$ETIME-CRITICAL
  * ((nint/text) .feature "etime-elective-extension") => any
  * (uint .feature "etime-critical-extension") => any
}) .within etime-framework
```

3.1. Key 1

Key 1 indicates a Base Time value that is exactly like the data item
that would be tagged by CBOR tag 1 (Posix time [TIME_T] as int or
float). The time value indicated by the value under this key can be
further modified by other keys.

```
$$ETIME-BASETIME //= (1: ~time)
```

3.2. Keys 4 and 5

Keys 4 and 5 indicate a Base Time value and are like key 1, except
that the data item is an array as defined for CBOR tag 4 or 5,
respectively. This can be used to include a Decimal or Bigfloat
epoch-based float [TIME_T] in an extended time.

```
$$ETIME-BASETIME //= (4: ~decfrac)
$$ETIME-BASETIME //= (5: ~bigfloat)
```

3.3. Keys -3, -6, -9, -12, -15, -18

The keys -3, -6, -9, -12, -15 and -18 indicate additional decimal
fractions by giving an unsigned integer (major type 0) and scaling
this with the scale factor 1e-3, 1e-6, 1e-9, 1e-12, 1e-15, and
1e-18, respectively (see Table 1). More than one of these keys **MUST
NOT** be present in one extended time data item. These additional
fractions are added to a Base Time in seconds [SI-SECOND] indicated
by a Key 1, which then **MUST** also be present and **MUST** have an integer
value.

Key	meaning	example usage
-3	milliseconds	Java time
```

| Key | meaning | example usage |
|---|---|---|
| -6 | microseconds | (old) UNIX time |
| -9 | nanoseconds | (new) UNIX time |
| -12 | picoseconds | Haskell time |
| -15 | femtoseconds | (future) |
| -18 | attoseconds | (future) |

Table 1: Key for decimally scaled Fractions

```
$$ETIME-ELECTIVE //= (-3: uint)
$$ETIME-ELECTIVE //= (-6: uint)
$$ETIME-ELECTIVE //= (-9: uint)
$$ETIME-ELECTIVE //= (-12: uint)
$$ETIME-ELECTIVE //= (-15: uint)
$$ETIME-ELECTIVE //= (-18: uint)
```

3.4. Key -1: Timescale

Key -1 is used to indicate a timescale. The value 0 indicates UTC, with the POSIX epoch [TIME_T]; the value 1 indicates TAI, with the PTP (Precision Time Protocol) epoch [IEEE1588-2008].

```
$$ETIME-ELECTIVE //= (-1 => $ETIME-TIMESCALE)

$ETIME-TIMESCALE /= &(etime-utc: 0)
$ETIME-TIMESCALE /= &(etime-tai: 1)
```

If key -1 is not present, timescale value 0 is implied.

Additional values can be registered in the Timescale Registry (Section 7.2); values **MUST** be integers or text strings.

(Note that there should be no timescales "GPS" or "NTP" — instead, the time should be converted to TAI or UTC using a single addition or subtraction.)

$$t_{utc} = t_{ntp} - 2208988800$$
$$t_{tai} = t_{gps} + 315964819$$

Figure 1: Converting Common Offset Timescales

Editor's note: This initial set of timescales was deliberately chosen to be frugal, as the specification of the tag provides an extension point where additional timescales can be registered at any time. Registrations are clearly needed for earth-referenced timescales (such as UT1 and TT), as well as possibly for specific realizations of abstract time scales (such as TAI(USNO) which is more accurate as a constant offset basis for GPS times). While the registration process itself is trivial, these registrations need to

be made based on a solid specification of their actual definition.
Draft text for a specification of the UT1 time scale can be found at
https://github.com/cbor-wg/time-tag/pull/9.

3.5. Clock Quality

A number of keys are defined to indicate the quality of clock that
was used to determine the point in time.

The first three are analogous to clock-quality-grouping in
[RFC8575], which is in turn based on the definitions in
[IEEE1588-2008]; two more are specific to this document.

```
ClockQuality-group = (
  ? &(ClockClass: -2) => uint .size 1 ; PTP/RFC8575
  ? &(ClockAccuracy: -4) => uint .size 1 ; PTP/RFC8575
  ? &(OffsetScaledLogVariance: -5) => uint .size 2 ; PTP/RFC8575
  ? &(Uncertainty: -7) => ~time/~duration
  ? &(Guarantee: -8) => ~time/~duration
)
```

3.5.1. ClockClass (Key -2)

Key -2 (ClockClass) can be used to indicate the clock class as per
Table 5 of [IEEE1588-2008]. It is defined as a one-byte unsigned
integer as that is the range defined there.

3.5.2. ClockAccuracy (Key -4)

Key -4 (ClockAccuracy) can be used to indicate the clock accuracy as
per Table 6 of [IEEE1588-2008]. It is defined as a one-byte unsigned
integer as that is the range defined there. The range between 32 and
47 is a slightly distorted logarithmic scale from 25 ns to 1 s (see
Figure 2); the number 254 is the value to be used if an unknown
accuracy needs to be expressed.

$$enum_{acc} \approx 48 + \lfloor 2 \cdot log_{10} \frac{acc}{s} - \epsilon \rfloor$$

Figure 2: Approximate conversion from accuracy to accuracy enumeration
value

3.5.3. OffsetScaledLogVariance (Key -5)

Key -5 (OffsetScaledLogVariance) can be used to represent the
variance exhibited by the clock when it has lost its synchronization
with an external reference clock. The details for the computation of
this characteristic are defined in Section 7.6.3 of [IEEE1588-2008].

3.5.4. Uncertainty (Key -7)

Key -7 (Uncertainty) can be used to represent a known measurement
uncertainty for the clock, as a numeric value in seconds or as a
duration (Section 4).

For this document, uncertainty is defined as in Section 2.2.3 of
[GUM]: "parameter, associated with the result of a measurement, that
characterizes the dispersion of the values that could reasonably be
attributed to the measurand". More specifically, the value for this
key represents the extended uncertainty for k = 2, in seconds.

3.5.5. Guarantee (Key -8)

Key -8 (Guarantee) can be used to represent a stated guarantee for
the accuracy of the point in time, as a numeric value in seconds or
as a duration (Section 4) representing the maximum allowed deviation
from the true value.

While such a guarantee is unattainable in theory, existing standards
such as [RFC3161] stipulate the representation of such guarantees,
and therefore this format provides a way to represent them as well;
the time value given is nominally guaranteed to not deviate from the
actual time by more than the value of the guarantee, in seconds.

3.6. Keys -10, 10: Time Zone Hint

Keys -10 and 10 supply supplementary information, where key 10 is
critical.

They can be used to provide a hint about the time zone that would
best fit for displaying the time given to humans, using a text
string in the format defined for time-zone-name or time-numoffset in
[IXDTF]. Key -10 is equivalent to providing this information as an
elective hint, while key 10 provides this information as critical
(i.e., it **MUST** be used when interpreting the entry with this key).

Keys -10 and 10 **MUST NOT** both be present.

```
$$ETIME-ELECTIVE //= (-10: time-zone-info)
$$ETIME-CRITICAL //= (10: time-zone-info)

time-zone-info = tstr .abnf
                ("time-zone-name / time-numoffset" .det IXDTFtz)
IXDTFtz = '
   time-hour        = 2DIGIT  ; 00-23
   time-minute      = 2DIGIT  ; 00-59
   time-numoffset   = ("+" / "-") time-hour ":" time-minute



   time-zone-initial = ALPHA / "." / "_"
   time-zone-char    = time-zone-initial / DIGIT / "-" / "+"
   time-zone-part    = time-zone-initial *13(time-zone-char)
                        ; but not "." or ".."
   time-zone-name    = time-zone-part *("/" time-zone-part)
   ALPHA             =  %x41-5A / %x61-7A   ; A-Z / a-z
   DIGIT             =  %x30-39 ; 0-9
' ; extracted from [IXDTF] and [RFC3339]; update as needed
```

3.7. Keys -11, 11: IXDTF Suffix Information

Keys -11 and 11 supply supplementary information, where key 11 is
critical.

Similar to keys -10 and 10, keys -11 (elective) and 11 (critical)
can be used to provide additional information in the style of IXDTF
suffixes, such as the calendar that would best fit for displaying
the time given to humans. The key's value is a map that has IXDTF
suffix-key names as keys and corresponding suffix values as values,
specifically:

```
$$ETIME-ELECTIVE //= (-11: suffix-info-map)
$$ETIME-CRITICAL //= (11: suffix-info-map)

suffix-info-map = { * suffix-key => suffix-values }
suffix-key = tstr .abnf ("suffix-key" .det IXDTF)
suffix-values = one-or-more<suffix-value>
one-or-more<T> = T / [ 2* T ]
suffix-value = tstr .abnf ("suffix-value" .det IXDTF)

IXDTF = '
   key-initial        = lcalpha / "_"
   key-char           = key-initial / DIGIT / "-"
   suffix-key         = key-initial *key-char

   suffix-value       = 1*alphanum
   alphanum           = ALPHA / DIGIT
   lcalpha            =  %x61-7A
   ALPHA              =  %x41-5A / %x61-7A   ; A-Z / a-z
   DIGIT              =  %x30-39 ; 0-9
' ; extracted from [IXDTF]; update as needed!
```

When keys -11 and 11 both are present, the two maps **MUST NOT** have entries with the same map keys.

Figure 4 of [IXDTF] gives an example for an extended date-time with both time zone and suffix information:

```
1996-12-19T16:39:57-08:00[America/Los_Angeles][u-ca=hebrew]
```

A time tag that is approximating this example, in CBOR diagnostic notation, would be:

```
/ 1996-12-19T16:39:57-08:00[America//Los_Angeles][u-ca=hebrew] /
1001({ 1: 851042397,
     -10: "America/Los_Angeles",
     -11: { "u-ca": "hebrew" }
})
```

Note that both -10 and -11 are using negative keys and therefore provide elective information, as in the IXDTF form. Note also that in this example the time numeric offset (-08:00) is lost in translating from the [RFC3339] information in the IXDTF into a Posix time that can be included under Key 1 in a time tag.

4. Duration Format

A duration is the length of an interval of time. Durations in this format are given in SI seconds, possibly adjusted for conventional corrections of the timescale given (e.g., leap seconds).

Except for using Tag 1002 instead of 1001, durations are
structurally identical to time values.

```
Duration = #6.1001(etime-detailed)
```

Semantically, they do not measure the time elapsed from a given
epoch, but from the start to the end of (an otherwise unspecified)
interval of time.

In combination with an epoch identified in the context, a duration
can also be used to express an absolute time.

Without such context, durations are subject to some uncertainties
underlying the timescale used. E.g., for durations intended as a
determinant of future time periods, there is some uncertainty of
what irregularities (such as leap seconds, timescale corrections)
will be exhibited by the timescale in that period. For durations as
measurements of past periods, abstracting the period to a duration
loses some detail about timescale irregularities. For many
applications, these uncertainties are acceptable and thus the use of
durations is appropriate.

Note that [ISO8601:1988] durations are rather different from the
ones defined in the present specification; there is no intention to
support ISO 8601 durations here.

5. Period Format

A period is a specific interval of time, specified as either two
times giving the start and the end of that interval, or as one of
these two plus a duration.

They are given as an array of unwrapped time and duration elements,
tagged with Tag 1003:

```
Period = #6.1003([
  start: ~Time / null
  end: ~Time / null
  ? duration: ~Duration / null
])
```

If the third array element is not given, the duration element is
null. Exactly two out of the three elements must be non-null, this
can be clumsily expressed in CDDL as:

```
clumsy-Period = #6.1003([
  (start: ~Time,
   ((end: ~Time,
     ? duration: null) //
    (end: null,
     duration: ~Duration))) //
  (start: null,
   end: ~Time,
   duration: ~Duration)
])
```

6. CDDL typenames

When detailed validation is not needed, the type names defined in
[Figure 3](#) are recommended:

```
etime = #6.1001({* (int/tstr) => any})
duration = #6.1002({* (int/tstr) => any})
period = #6.1003([~etime/null, ~etime/null, ~duration/null])
```

Figure 3: Recommended type names for CDDL

7. IANA Considerations

7.1. CBOR tags

In the registry [IANA.cbor-tags], IANA has allocated the tags in
[Table 2](#) from what was at the time the FCFS space, with the present
document as the specification reference.

Tag	Data Item	Semantics
1001	map	[RFCthis] extended time
1002	map	[RFCthis] duration
1003	array	[RFCthis] period

Table 2: Values for Tags

IANA is requested to change the "Data Item" column for Tag 1003 from
"map" to "array".

7.2. Timescale Registry

This specification defines a new subregistry titled "Timescale
Registry" in the "CBOR Time Tag Parameters" registry [IANA.cbor-
time-tag-parameters], with a combination of "Expert Review" and "RFC
Required" as the Registration Procedure (Sections [4.5](#) and [4.7](#) of
[[BCP26](#)]).

Each entry needs to provide a timescale name (a sequence of
uppercase ASCII characters and digits, where a digit may not occur
at the start: [A-Z][A-Z0-9]*), a value (unsigned integer), and brief
description of the semantics, and a specification reference (RFC).
The initial contents are shown in Table 3.

Timescale	Value	Semantics	Reference
UTC	0	UTC with POSIX Epoch	[RFCthis]
TAI	1	TAI with PTP Epoch	[RFCthis]

Table 3: Initial Content of Timescale Registry

7.3. Map Key Registry

This specification defines a new subregistry titled "Map Key
Registry" in the "CBOR Time Tag Parameters" registry [IANA.cbor-
time-tag-parameters], with "Specification Required" as the
Registration Procedure (Section 4.6 of [BCP26]).

The designated expert is requested to assign the key values with the
shortest encodings (1+0 and 1+1 encoding) to registrations that are
likely to enjoy wide use and can benefit from short encodings.

Each entry needs to provide a map key value (integer), a brief
description of the semantics, and a specification reference (RFC).
The initial contents are shown in Table 3.

Value	Semantics	Reference
-18	attoseconds	[RFCthis]
-15	femtoseconds	[RFCthis]
-12	picoseconds	[RFCthis]
-11	IXDTF Suffix Information (elective)	[RFCthis], [IXDTF]
-10	IXDTF Time Zone Hint (elective)	[RFCthis], [IXDTF]
-9	nanoseconds	[RFCthis]
-8	Guarantee	[RFCthis]
-7	Uncertainty	[RFCthis]
-6	microseconds	[RFCthis]
-5	Offset-Scaled Log Variance	[RFCthis]
-4	Clock Accuracy	[RFCthis]
-3	milliseconds	[RFCthis]
-2	Clock Class	[RFCthis]
1	Base Time value (as in CBOR Tag 1)	[RFCthis]
4	Base Time value as in CBOR Tag 4	[RFCthis]
5	Base Time value as in CBOR Tag 5	[RFCthis]
10	IXDTF Time Zone Hint (critical)	[RFCthis], [IXDTF]
11	IXDTF Suffix Information (critical)	[RFCthis], [IXDTF]

Table 4

8. Security Considerations

The security considerations of RFC 8949 apply; the tags introduced here are not expected to raise security considerations beyond those.

Time, of course, has significant security considerations; these include the exploitation of ambiguities where time is security relevant (e.g., for freshness or in a validity span) or the disclosure of characteristics of the emitting system (e.g., time zone, or clock resolution and wall clock offset).

9. References

9.1. Normative References

[BCP26]   Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <https://www.rfc-editor.org/rfc/rfc8126>.

[GUM]   Joint Committee for Guides in Metrology, "Evaluation of measurement data — Guide to the expression of uncertainty in measurement", JCGM 100:2008, September 2008, <https://www.bipm.org/en/publications/guides/gum.html>.

[IANA.cbor-tags] IANA, "Concise Binary Object Representation (CBOR) Tags", <https://www.iana.org/assignments/cbor-tags>.

[IEEE1588-2008] IEEE, "1588-2008 - IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems", July 2008, <http://standards.ieee.org/findstds/standard/1588-2008.html>.

[IXDTF]   Sharma, U. and C. Bormann, "Date and Time on the Internet: Timestamps with additional information", Work in Progress, Internet-Draft, draft-ietf-sedate-datetime-extended-07, 19 January 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-sedate-datetime-extended-07>.

[RFC2119]   Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/

                  RFC2119, March 1997, <https://www.rfc-editor.org/rfc/
                  rfc2119>.

     [RFC8174]   Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC
                  2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174,
                  May 2017, <https://www.rfc-editor.org/rfc/rfc8174>.

     [RFC8610]   Birkholz, H., Vigano, C., and C. Bormann, "Concise Data
                  Definition Language (CDDL): A Notational Convention to
                  Express Concise Binary Object Representation (CBOR) and
                  JSON Data Structures", RFC 8610, DOI 10.17487/RFC8610,
                  June 2019, <https://www.rfc-editor.org/rfc/rfc8610>.

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

     [SI-SECOND] International Organization for Standardization (ISO),
                  "Quantities and units — Part 3: Space and time", ISO
                  80000-3, 1 March 2006.

     [TIME_T]   The Open Group Base Specifications, "Vol. 1: Base
                  Definitions, Issue 7", Section 4.15 'Seconds Since the
                  Epoch', IEEE Std 1003.1-2008, 2016 Edition, 2016,
                  <http://pubs.opengroup.org/onlinepubs/9699919799/
                  basedefs/V1_chap04.html#tag_04_16>.

9.2. Informative References

     [ISO8601:1988] ISO, "Data elements and interchange formats —
                  Information interchange — Representation of dates and
                  times", ISO 8601:1988, June 1988, <https://www.iso.org/
                  standard/15903.html>.

     [RFC3161]   Adams, C., Cain, P., Pinkas, D., and R. Zuccherato,
                  "Internet X.509 Public Key Infrastructure Time-Stamp
                  Protocol (TSP)", RFC 3161, DOI 10.17487/RFC3161, August
                  2001, <https://www.rfc-editor.org/rfc/rfc3161>.

     [RFC3339]   Klyne, G. and C. Newman, "Date and Time on the Internet:
                  Timestamps", RFC 3339, DOI 10.17487/RFC3339, July 2002,
                  <https://www.rfc-editor.org/rfc/rfc3339>.

     [RFC8575]   Jiang, Y., Ed., Liu, X., Xu, J., and R. Cummings, Ed.,
                  "YANG Data Model for the Precision Time Protocol (PTP)",
                  RFC 8575, DOI 10.17487/RFC8575, May 2019, <https://
                  www.rfc-editor.org/rfc/rfc8575>.

Appendix A. Collected CDDL

This appendix collects the CDDL rules spread over the document into
one convenient place.

```
Etime = #6.1001(etime-detailed)

etime-framework = {
  uint => any ; at least one base time
  * (nint/text) => any ; elective supplementary information
  * uint => any ; critical supplementary information
}

etime-detailed = ({
  $$ETIME-BASETIME
  ClockQuality-group
  * $$ETIME-ELECTIVE
  * $$ETIME-CRITICAL
  * ((nint/text) .feature "etime-elective-extension") => any
  * (uint .feature "etime-critical-extension") => any
}) .within etime-framework


$$ETIME-BASETIME //= (1: ~time)


$$ETIME-BASETIME //= (4: ~decfrac)
$$ETIME-BASETIME //= (5: ~bigfloat)


$$ETIME-ELECTIVE //= (-3: uint)
$$ETIME-ELECTIVE //= (-6: uint)
$$ETIME-ELECTIVE //= (-9: uint)
$$ETIME-ELECTIVE //= (-12: uint)
$$ETIME-ELECTIVE //= (-15: uint)
$$ETIME-ELECTIVE //= (-18: uint)


$$ETIME-ELECTIVE //= (-1 => $ETIME-TIMESCALE)

$ETIME-TIMESCALE /= &(etime-utc: 0)
$ETIME-TIMESCALE /= &(etime-tai: 1)


ClockQuality-group = (
  ? &(ClockClass: -2) => uint .size 1 ; PTP/RFC8575
  ? &(ClockAccuracy: -4) => uint .size 1 ; PTP/RFC8575
  ? &(OffsetScaledLogVariance: -5) => uint .size 2 ; PTP/RFC8575
  ? &(Uncertainty: -7) => ~time/~duration
  ? &(Guarantee: -8) => ~time/~duration
)


$$ETIME-ELECTIVE //= (-10: time-zone-info)
$$ETIME-CRITICAL //= (10: time-zone-info)
```

```
time-zone-info = tstr .abnf
                ("time-zone-name / time-numoffset" .det IXDTFtz)
IXDTFtz = '
   time-hour        = 2DIGIT  ; 00-23
   time-minute      = 2DIGIT  ; 00-59
   time-numoffset   = ("+" / "-") time-hour ":" time-minute


   time-zone-initial = ALPHA / "." / "_"
   time-zone-char    = time-zone-initial / DIGIT / "-" / "+"
   time-zone-part    = time-zone-initial *13(time-zone-char)
                       ; but not "." or ".."
   time-zone-name    = time-zone-part *("/" time-zone-part)
   ALPHA             =  %x41-5A / %x61-7A    ; A-Z / a-z
   DIGIT             =  %x30-39 ; 0-9
' ; extracted from [IXDTF] and [RFC3339]; update as needed


$$ETIME-ELECTIVE //= (-11: suffix-info-map)
$$ETIME-CRITICAL //= (11: suffix-info-map)

suffix-info-map = { * suffix-key => suffix-values }
suffix-key = tstr .abnf ("suffix-key" .det IXDTF)
suffix-values = one-or-more<suffix-value>
one-or-more<T> = T / [ 2* T ]
suffix-value = tstr .abnf ("suffix-value" .det IXDTF)

IXDTF = '
   key-initial      = lcalpha / "_"
   key-char         = key-initial / DIGIT / "-"
   suffix-key       = key-initial *key-char

   suffix-value     = 1*alphanum
   alphanum         = ALPHA / DIGIT
   lcalpha          =  %x61-7A
   ALPHA            =  %x41-5A / %x61-7A    ; A-Z / a-z
   DIGIT            =  %x30-39 ; 0-9
' ; extracted from [IXDTF]; update as needed!


Duration = #6.1001(etime-detailed)


Period = #6.1003([
  start: ~Time / null
  end: ~Time / null
  ? duration: ~Duration / null
])
```

```
clumsy-Period = #6.1003([
  (start: ~Time,
   ((end: ~Time,
     ? duration: null) //
    (end: null,
     duration: ~Duration))) //
  (start: null,
   end: ~Time,
   duration: ~Duration)
])


etime = #6.1001({* (int/tstr) => any})
duration = #6.1002({* (int/tstr) => any})
period = #6.1003([~etime/null, ~etime/null, ~duration/null])
```

Figure 4: Collected CDDL rules from this specification

Acknowledgements

Authors' Addresses

Carsten Bormann
Universität Bremen TZI
Postfach 330440
D-28359 Bremen
Germany

Phone: +49-421-218-63921
Email: cabo@tzi.org

Ben Gamari
Well-Typed
117 Middle Rd.
Portsmouth, NH 03801
United States

Email: ben@well-typed.com

Henk Birkholz
Fraunhofer Institute for Secure Information Technology
Rheinstrasse 75
64295 Darmstadt
Germany

Email: henk.birkholz@sit.fraunhofer.de