

Network Working Group  
Internet-Draft  
Intended status: Informational  
Expires: 9 April 2022

C. Bormann  
Universität Bremen TZI  
6 October 2021

Application-Oriented Literals in CBOR Extended Diagnostic Notation  
draft-bormann-cbor-edn-literals-00

Abstract

The Concise Binary Object Representation, CBOR (RFC 8949) defines a "diagnostic notation" in order to be able to converse about CBOR data items without having to resort to binary data.

This document specifies how to add application-oriented extensions to the diagnostic notation. It then defines two such extensions for the use of CBOR diagnostic notation with CoRAL and Constrained Resource Identifiers (draft-ietf-core-coral, draft-ietf-core-href).

Note

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

The content of this draft may preferably be distributed to a number of different documents. This is to be decided.

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 9 April 2022.

Copyright Notice

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

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<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. Application-Oriented Extension Literals . . . . . | 2 |
| 3. The "cri" Extension . . . . .                     | 3 |
| 4. The "dt" Extension . . . . .                      | 4 |
| 5. IANA Considerations . . . . .                     | 4 |
| 6. Security considerations . . . . .                 | 5 |
| 7. References . . . . .                              | 5 |
| 7.1. Normative References . . . . .                  | 5 |
| 7.2. Informative References . . . . .                | 6 |
| Acknowledgements . . . . .                           | 6 |
| Author's Address . . . . .                           | 6 |

## 1. Introduction

For the Concise Binary Object Representation, CBOR, Section 8 of [RFC8949] defines a "diagnostic notation" in order to be able to converse about CBOR data items without having to resort to binary data. Diagnostic notation is based on JSON, with extensions for representing CBOR constructs such as binary data and tags. (Standardizing this together with the actual interchange format does not serve to create another interchange format, but enables the use of a shared diagnostic notation in tools for and documents about CBOR.)

This document specifies how to add application-oriented extensions to the diagnostic notation. It then defines two such extensions for the use of CBOR diagnostic notation with CoRAL and Constrained Resource Identifiers [I-D.ietf-core-coral] [I-D.ietf-core-href].

## 2. Application-Oriented Extension Literals

This document extends the syntax used in diagnostic notation for byte string literals to also be available for application-oriented extensions.

As per Section 8 of [RFC8949], the diagnostic notation can notate byte strings in a number of [RFC4648] base encodings, where the encoded text is enclosed in single quotes, prefixed by an identifier (>h< for base16, >b32< for base32, >h32< for base32hex, >b64< for base64 or base64url).

This syntax can be thought to establish a name space, with the names "h", "b32", "h32", and "b64" taken, but other names being unallocated. The present specification defines additional names for this namespace, which we call `_application-extension identifiers_`. For the quoted string, the same rules apply as for byte strings. In particular, the escaping rules of JSON strings are applied equivalently for application-oriented extensions, e.g., `\\` stands for a single backslash and `\'` stands for a single quote.

An application-extension identifier is a name consisting of a lower-case ASCII letter (a-z) and zero or more additional ASCII characters that are either lower-case letters or digits (a-z0-9).

Application-extension identifiers are registered in a registry (Section 5). Prefixing a single-quoted string, an application-extension identifier is used to build an application-oriented extension literal, which stands for a CBOR data item the value of which is derived from the text given in the single-quoted string using a procedure defined in the specification for an application-extension identifier.

Examples for application-oriented extensions to CBOR diagnostic notation can be found in the following sections.

### 3. The "cri" Extension

The application-extension identifier "cri" is used to notate a Constrained Resource Identifier literal as per [I-D.ietf-core-href].

The text of the literal is a URI Reference as per [RFC3986] or an IRI Reference as per [RFC3987].

The value of the literal is a CRI that can be converted to the text of the literal using the procedure of Section 6.1 of [I-D.ietf-core-href]. Note that there may be more than one CRI that can be converted to the URI/IRI given; implementations are expected to favor the simplest variant available and make non-surprising choices otherwise.

As an example, the CBOR diagnostic notation

```
cri'https://example.com/bottarga/shaved'
```

is equivalent to

```
[-4, ["example", "com"], ["bottarga", "shaved"]]
```

#### 4. The "dt" Extension

The application-extension identifier "dt" is used to notate a date/time literal that can be used as an Epoch-Based Date/Time as per Section 3.4.2 of [RFC8949].

The text of the literal is a Standard Date/Time String as per Section 3.4.1 of [RFC8949].

The value of the literal is a number representing the result of a conversion of the given Standard Date/Time String to an Epoch-Based Date/Time. If fractional seconds are given in the text (production time-fraction in Appendix A of [RFC3339]), the value is a floating-point number; the value is an integer number otherwise.

As an example, the CBOR diagnostic notation

```
dt'1969-07-21T02:56:16Z'
```

is equivalent to

```
-14159024
```

#### 5. IANA Considerations

IANA is requested to create a registry [[where?]] for application-extension identifiers, with the initial content shown in Table 1.

| application-extension<br>identifier | description                        | reference |
|-------------------------------------|------------------------------------|-----------|
| h                                   | Reserved                           | RFC8949   |
| b32                                 | Reserved                           | RFC8949   |
| h32                                 | Reserved                           | RFC8949   |
| b64                                 | Reserved                           | RFC8949   |
| cri                                 | Constrained<br>Resource Identifier | RFCThis   |
| dt                                  | Date/Time                          | RFCThis   |

Table 1: Initial Content of application extension  
identifier registry

// (Define policy; detailed template)

## 6. Security considerations

The security considerations of [RFC8949] and [RFC8610] apply.

// Anything else meaningful to say here?

## 7. References

### 7.1. Normative References

[I-D.ietf-core-href]

Bormann, C. and H. Birkholz, "Constrained Resource Identifiers", Work in Progress, Internet-Draft, draft-ietf-core-href-06, 25 July 2021, <<https://www.ietf.org/archive/id/draft-ietf-core-href-06.txt>>.

[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/info/rfc3339>>.

- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <<https://www.rfc-editor.org/info/rfc3986>>.
- [RFC3987] Duerst, M. and M. Suignard, "Internationalized Resource Identifiers (IRIs)", RFC 3987, DOI 10.17487/RFC3987, January 2005, <<https://www.rfc-editor.org/info/rfc3987>>.
- [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/info/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/info/rfc8949>>.

## 7.2. Informative References

- [I-D.ietf-core-coral] Hartke, K., "The Constrained RESTful Application Language (CoRAL)", Work in Progress, Internet-Draft, draft-ietf-core-coral-03, 9 March 2020, <<https://www.ietf.org/archive/id/draft-ietf-core-coral-03.txt>>.
- [RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", RFC 4648, DOI 10.17487/RFC4648, October 2006, <<https://www.rfc-editor.org/info/rfc4648>>.

## Acknowledgements

The concept of application-oriented extensions to diagnostic notation, as well as the definition for the "dt" extension were inspired by the CoRAL work by Klaus Hartke.

## Author's Address

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

Phone: +49-421-218-63921

Internet-Draft

CBOR EDN Literals

October 2021

Email: [cabo@tzi.org](mailto: cabo@tzi.org)

Network Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: 25 April 2022

C. Bormann  
Universität Bremen TZI  
22 October 2021

Additional Control Operators for CDDL  
draft-ietf-cbor-cddl-control-07

## Abstract

The Concise Data Definition Language (CDDL), standardized in RFC 8610, provides "control operators" as its main language extension point.

The present document defines a number of control operators that were not yet ready at the time RFC 8610 was completed: `.plus`, `.cat` and `.det` for the construction of constants, `.abnf/.abnfb` for including ABNF (RFC 5234/RFC 7405) in CDDL specifications, and `.feature` for indicating the use of a non-basic feature in an instance.

## 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 25 April 2022.

## Copyright Notice

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

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<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 . . . . .                         | 3  |
| 2. Computed Literals . . . . .                     | 3  |
| 2.1. Numeric Addition . . . . .                    | 4  |
| 2.2. String Concatenation . . . . .                | 4  |
| 2.3. String Concatenation with Dedenting . . . . . | 5  |
| 3. Embedded ABNF . . . . .                         | 6  |
| 4. Features . . . . .                              | 8  |
| 5. IANA Considerations . . . . .                   | 10 |
| 6. Implementation Status . . . . .                 | 11 |
| 7. Security considerations . . . . .               | 11 |
| 8. References . . . . .                            | 12 |
| 8.1. Normative References . . . . .                | 12 |
| 8.2. Informative References . . . . .              | 12 |
| Acknowledgements . . . . .                         | 13 |
| Author's Address . . . . .                         | 13 |

## 1. Introduction

The Concise Data Definition Language (CDDL), standardized in [RFC8610], provides "control operators" as its main language extension point (Section 3.8 of [RFC8610]).

The present document defines a number of control operators that were not yet ready at the time RFC 8610 was completed:

| Name     | Purpose   |
|----------|---|
| .plus    | Numeric addition                                |
| .cat     | String Concatenation                            |
| .det     | String Concatenation, pre-dedenting             |
| .abnf    | ABNF in CDDL (text strings)                     |
| .abnfb   | ABNF in CDDL (byte strings)                     |
| .feature | Indicate name of feature used (extension point) |

Table 1: New control operators in this document

### 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.

This specification uses terminology from [RFC8610]. In particular, with respect to control operators, "target" refers to the left-hand side operand, and "controller" to the right-hand side operand. "Tool" refers to tools along the lines of that described in Appendix F of [RFC8610]. Note also that the data model underlying CDDL provides for text strings as well as byte strings as two separate types, which are then collectively referred to as "strings".

The term ABNF in this specification stands for the combination of [RFC5234] and [RFC7405], i.e., the ABNF control operators defined by this document allow use of the case-sensitive extensions defined in [RFC7405].

### 2. Computed Literals

CDDL as defined in [RFC8610] does not have any mechanisms to compute literals. To cover a large part of the use cases, this specification adds three control operators: .plus for numeric addition, .cat for string concatenation, and .det for string concatenation with dedenting of both sides (target and controller).

For these operators, as with all control operators, targets and controllers are types. The resulting type is therefore formally a function of the elements of the cross-product of the two types. Not all tools may be able to work with non-unique targets or controllers.

## 2.1. Numeric Addition

In many cases in a specification, numbers are needed relative to a base number. The `.plus` control identifies a number that is constructed by adding the numeric values of the target and of the controller.

Target and controller **MUST** be numeric. If the target is a floating point number and the controller an integer number, or vice versa, the sum is converted into the type of the target; converting from a floating point number to an integer selects its floor (the largest integer less than or equal to the floating point number, i.e., rounding towards negative infinity).

```
interval<BASE> = (  
    BASE => int           ; lower bound  
    (BASE .plus 1) => int ; upper bound  
    ? (BASE .plus 2) => int ; tolerance  
)  
  
X = 0  
Y = 3  
rect = {  
    interval<X>  
    interval<Y>  
}
```

Figure 1: Example: addition to a base value

The example in Figure 1 contains the generic definition of a CDDL group interval that gives a lower and an upper bound and optionally a tolerance. The parameter `BASE` allows the non-conflicting use of multiple of these interval groups in one map, by assigning different labels to the entries of the interval. `rect` combines two of these interval groups into a map, one group for the X dimension (using 0, 1, and 2 as labels) and one for Y dimension (using 3, 4, and 5 as labels).

## 2.2. String Concatenation

It is often useful to be able to compose string literals out of component literals defined in different places in the specification.

The `.cat` control identifies a string that is built from a concatenation of the target and the controller. Target and controller **MUST** be strings. The result of the operation has the type of the target. The concatenation is performed on the bytes in both strings. If the target is a text string, the result of that concatenation **MUST** be valid UTF-8.

```
a = "foo" .cat '
    bar
    baz
,
; on a system where the newline is \n, is the same string as:
b = "foo\n bar\n baz\n"
```

Figure 2: Example: concatenation of text and byte string

The example in Figure 2 builds a text string named `a` out of concatenating the target text string `"foo"` and the controller byte string entered in a text form byte string literal. (This particular idiom is useful when the text string contains newlines, which, as shown in the example for `b`, may be harder to read when entered in the format that the pure CDDL text string notation inherits from JSON.)

### 2.3. String Concatenation with Dedenting

Multi-line string literals for various applications, including embedded ABNF (Section 3), need to be set flush left, at least partially. Often, having some indentation in the source code for the literal can promote readability, as in Figure 3.

```
oid = bytes .abnfb ("oid" .det cbor-tags-oid)
roid = bytes .abnfb ("roid" .det cbor-tags-oid)

cbor-tags-oid = '
    oid = 1*arc
    roid = *arc
    arc = [nlsb] %x00-7f
    nlsb = %x81-ff *%x80-ff
,
```

Figure 3: Example: dedenting concatenation

The control operator `.det` works like `.cat`, except that both arguments (target and controller) are independently dedented before the concatenation takes place.

For the first rule in Figure 3, the result is equivalent to Figure 4.

```
oid = bytes .abnfb 'oid
oid = 1*arc
roid = *arc
arc = [nlsb] %x00-7f
nlsb = %x81-ff *%x80-ff
,
```

Figure 4: Dedenting example: result of first .det

For the purposes of this specification, we define dedenting as:

1. determining the smallest amount of left-most blank space (number of leading space characters) present in all the non-blank lines, and
2. removing exactly that number of leading space characters from each line. For blank (blank space only or empty) lines, there may be less (or no) leading space characters than this amount, in which case all leading space is removed.

(The name .det is a shortcut for "dedenting cat". The maybe more obvious name .dedcat has not been chosen as it is longer and may invoke unpleasant images.)

Occasionally, dedenting of only a single item is needed. This can be achieved by using this operator with an empty string, e.g., "" .det rhs or lhs .det "", which can in turn be combined with a .cat: in the construct lhs .cat (" " .det rhs), only rhs is dedented.

### 3. Embedded ABNF

Many IETF protocols define allowable values for their text strings in ABNF [RFC5234] [RFC7405]. It is often desirable to define a text string type in CDDL by employing existing ABNF embedded into the CDDL specification. Without specific ABNF support in CDDL, that ABNF would usually need to be translated into a regular expression (if that is even possible).

ABNF is added to CDDL in the same way that regular expressions were added: by defining a .abnf control operator. The target is usually text or some restriction on it, the controller is the text of an ABNF specification.

There are several small issues, with solutions given here:

- \* ABNF can be used to define byte sequences as well as UTF-8 text strings interpreted as Unicode scalar sequences. This means this specification defines two control operators: .abnfb for ABNF

denoting byte sequences and `.abnf` for denoting sequences of Unicode scalar values (codepoint) represented as UTF-8 text strings. Both control operators can be applied to targets of either string type; the ABNF is applied to sequence of bytes in the string interpreting that as a sequence of bytes (`.abnfb`) or as a sequence of code points represented as an UTF-8 text string (`.abnf`). The controller string **MUST** be a text string.

- \* ABNF defines a list of rules, not a single expression (called "elements" in [RFC5234]). This is resolved by requiring the controller string to be one valid "element", followed by zero or more valid "rule" separated from the element by a newline; so the controller string can be built by preceding a piece of valid ABNF by an "element" that selects from that ABNF and a newline.
- \* For the same reason, ABNF requires newlines; specifying newlines in CDDL text strings is tedious (and leads to essentially unreadable ABNF). The workaround employs the `.cat` operator introduced in Section 2.2 and the syntax for text in byte strings. As is customary for ABNF, the syntax of ABNF itself (NOT the syntax expressed in ABNF!) is relaxed to allow a single linefeed as a newline:

`CRLF = %x0A / %x0D.0A`

- \* One set of rules provided in an ABNF specification is often used in multiple positions, in particular staples such as `DIGIT` and `ALPHA`. (Note that all rules referenced need to be defined in each ABNF operator controller string -- there is no implicit import of [RFC5234] Core ABNF or other rules.) The composition this calls for can be provided by the `.cat` operator, and/or by `.det` if there is indentation to be disposed of.

These points are combined into an example in Figure 5, which uses ABNF from [RFC3339] to specify one each of the CBOR tags defined in [RFC8943] and [RFC8949].

```

; for RFC 8943
Tag1004 = #6.1004(text .abnf full-date)
; for RFC 8949
Tag0 = #6.0(text .abnf date-time)

full-date = "full-date" .cat rfc3339
date-time = "date-time" .cat rfc3339

; Note the trick of idiomatically starting with a newline, separating
; off the element in the concatenations above from the rule-list
rfc3339 = '
    date-fullyear    = 4DIGIT
    date-month       = 2DIGIT ; 01-12
    date-mday        = 2DIGIT ; 01-28, 01-29, 01-30, 01-31 based on
                        ; month/year
    time-hour        = 2DIGIT ; 00-23
    time-minute      = 2DIGIT ; 00-59
    time-second      = 2DIGIT ; 00-58, 00-59, 00-60 based on leap sec
                        ; rules
    time-secfrac     = "." 1*DIGIT
    time-numoffset   = ("+" / "-") time-hour ":" time-minute
    time-offset      = "Z" / time-numoffset

    partial-time     = time-hour ":" time-minute ":" time-second
                        [time-secfrac]
    full-date        = date-fullyear "-" date-month "-" date-mday
    full-time        = partial-time time-offset

    date-time        = full-date "T" full-time
' .det rfc5234-core

rfc5234-core = '
    DIGIT            = %x30-39 ; 0-9
    ; abbreviated here
,
```

Figure 5: Example: employing RFC 3339 ABNF for defining CBOR Tags

#### 4. Features

Commonly, the kind of validation enabled by languages such as CDDL provides a Boolean result: valid, or invalid.

In rapidly evolving environments, this is too simplistic. The data models described by a CDDL specification may continually be enhanced by additional features, and it would be useful even for a specification that does not yet describe a specific future feature to identify the extension point the feature can use, accepting such extensions while marking them as such.

The `.feature` control annotates the target as making use of the feature named by the controller. The latter will usually be a string. A tool that validates an instance against that specification may mark the instance as using a feature that is annotated by the specification.

More specifically, the tool's diagnostic output might contain the controller (right-hand side) as a feature name, and the target (left-hand side) as a feature detail. However, in some cases, the target has too much detail, and the specification might want to hint the tool that more limited detail is appropriate. In this case, the controller should be an array, with the first element being the feature name (that would otherwise be the entire controller), and the second element being the detail (usually another string), as illustrated in Figure 6.

```
foo = {  
  kind: bar / baz .feature (["foo-extensions", "bazify"])  
}  
bar = ...  
baz = ... ; complex stuff that doesn't all need to be in the detail
```

Figure 6: Providing explicit detail with `.feature`

Figure 7 shows what could be the definition of a person, with potential extensions beyond name and organization being marked further-person-extension. Extensions that are known at the time this definition is written can be collected into `$$person-extensions`. However, future extensions would be deemed invalid unless the wildcard at the end of the map is added. These extensions could then be specifically examined by a user or a tool that makes use of the validation result; the label (map key) actually used makes a fine feature detail for the tool's diagnostic output.

Leaving out the entire extension point would mean that instances that make use of an extension would be marked as whole-sale invalid, making the entire validation approach much less useful. Leaving the extension point in, but not marking its use as special, would render mistakes such as using the label "organisation" instead of "organization" invisible.



```

person = {
  ? name: text
  ? organization: text
  $$person-extensions
  * (text .feature "further-person-extension") => any
}

$$person-extensions // = (? bloodgroup: text)

```

Figure 7: Map extensibility with .feature

Figure 8 shows another example where .feature provides for type extensibility.

```

allowed-types = number / text / bool / null
               / [* number] / [* text] / [* bool]
               / (any .feature "allowed-type-extension")

```

Figure 8: Type extensibility with .feature

A CDDL tool may simply report the set of features being used; the control then only provides information to the process requesting the validation. One could also imagine a tool that takes arguments allowing the tool to accept certain features and reject others (enable/disable). The latter approach could for instance be used for a JSON/CBOR switch, as illustrated in Figure 9, using SenML [RFC8428] as the example data model used with both JSON and CBOR.

```

SenML-Record = {
; ...
  ? v => number
; ...
}
v = JC<"v", 2>
JC<J,C> = J .feature "json" / C .feature "cbor"

```

Figure 9: Describing variants with .feature

It remains to be seen if the enable/disable approach can lead to new idioms of using CDDL. The language currently has no way to enforce mutually exclusive use of features, as would be needed in this example.

## 5. IANA Considerations

This document requests IANA to register the contents of Table 2 into the registry "CDDL Control Operators" of [IANA.cddl]:

| Name     | Reference |
|----------|-----------|
| .plus    | [RFCthis] |
| .cat     | [RFCthis] |
| .det     | [RFCthis] |
| .abnf    | [RFCthis] |
| .abnfb   | [RFCthis] |
| .feature | [RFCthis] |

Table 2: New control operators to be registered

## 6. Implementation Status

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

An early implementation of the control operator `.feature` has been available in the CDDL tool described in Appendix F of [RFC8610] since version 0.8.11. The validator warns about each feature being used and provides the set of target values used with the feature. The other control operators defined in this specification are also implemented as of version 0.8.21 and 0.8.26 (double-handed `.det`).

Andrew Weiss' [CDDL-RS] has an ongoing implementation of this draft which is feature-complete except for the ABNF and dedenting support (<https://github.com/anweiss/cddl/pull/79> (<https://github.com/anweiss/cddl/pull/79>)).

## 7. Security considerations

The security considerations of [RFC8610] apply.

While both [RFC5234] and [RFC7405] state that security is truly believed to be irrelevant to the respective document, the use of formal description techniques cannot only simplify, but sometimes also complicate a specification. This can lead to security problems in implementations and in the specification itself. As with CDDL itself, ABNF should be judiciously applied, and overly complex (or "cute") constructions should be avoided.

## 8. References

### 8.1. Normative References

- [IANA.cddl] IANA, "Concise Data Definition Language (CDDL)", <<https://www.iana.org/assignments/cddl>>.
- [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>>.
- [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>>.
- [RFC7405] Kyzivat, P., "Case-Sensitive String Support in ABNF", RFC 7405, DOI 10.17487/RFC7405, December 2014, <<https://www.rfc-editor.org/info/rfc7405>>.
- [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>>.
- [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/info/rfc8610>>.

### 8.2. Informative References

- [CDDL-RS] Weiss, A., "cddl-rs", n.d., <<https://github.com/anweiss/cddl>>.
- [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/info/rfc3339>>.
- [RFC8428] Jennings, C., Shelby, Z., Arkko, J., Keranen, A., and C. Bormann, "Sensor Measurement Lists (SenML)", RFC 8428, DOI 10.17487/RFC8428, August 2018, <<https://www.rfc-editor.org/info/rfc8428>>.

- [RFC8943] Jones, M., Nadalin, A., and J. Richter, "Concise Binary Object Representation (CBOR) Tags for Date", RFC 8943, DOI 10.17487/RFC8943, November 2020, <<https://www.rfc-editor.org/info/rfc8943>>.
- [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/info/rfc8949>>.

#### Acknowledgements

Jim Schaad suggested several improvements. The .feature feature was developed out of a discussion with Henk Birkholz. Paul Kyzivat helped isolate the need for .det.

.det is an abbreviation for "dedenting cat", but Det is also the name of a German TV Cartoon character created in the 1960s.

#### Author's Address

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

Phone: +49-421-218-63921  
Email: [cabo@tzi.org](mailto:cabo@tzi.org)

CBOR Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: 6 November 2022

M. Richardson  
Sandelman Software Works  
C. Bormann  
Universität Bremen TZI  
5 May 2022

On storing CBOR encoded items on stable storage  
draft-ietf-cbor-file-magic-12

## Abstract

This document defines a stored ("file") format for CBOR data items that is friendly to common file type recognition systems such as the Unix file(1) command.

## 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-file-magic/>.

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/>.

Source for this draft and an issue tracker can be found at  
<https://github.com/cbor-wg/cbor-magic-number>.

## 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 6 November 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 . . . . .                                | 4  |
| 1.2. Requirements for a Magic Number . . . . .            | 5  |
| 2. Protocol . . . . .                                     | 5  |
| 2.1. The CBOR Protocol Specific Tag . . . . .             | 5  |
| 2.2. Enveloping Method: CBOR Tag Wrapped . . . . .        | 6  |
| 2.2.1. Example . . . . .                                  | 7  |
| 2.3. Enveloping Method: Labeled CBOR Sequence . . . . .   | 7  |
| 2.3.1. Example . . . . .                                  | 8  |
| 3. Security Considerations . . . . .                      | 9  |
| 4. IANA Considerations . . . . .                          | 9  |
| 4.1. Labeled CBOR Sequence Tag . . . . .                  | 10 |
| 4.2. CBOR-Labeled Non-CBOR Data Tag . . . . .             | 10 |
| 4.3. CBOR Tags for CoAP Content-Format Numbers . . . . .  | 11 |
| 5. References . . . . .                                   | 11 |
| 5.1. Normative References . . . . .                       | 12 |
| 5.2. Informative References . . . . .                     | 12 |
| Appendix A. Advice to Protocol Designer . . . . .         | 14 |
| A.1. Is the on-wire format new? . . . . .                 | 15 |
| A.2. Can many items be trivially concatenated? . . . . .  | 15 |
| A.3. Are there tags at the start? . . . . .               | 16 |
| Appendix B. CBOR Tags for CoAP Content Formats . . . . .  | 16 |
| B.1. Content-Format Tag Examples . . . . .                | 18 |
| Appendix C. Example from Openswan . . . . .               | 18 |
| Appendix D. Using CBOR Labels for non-CBOR data . . . . . | 19 |
| D.1. Content-Format Tag Examples . . . . .                | 20 |
| Appendix E. Changelog . . . . .                           | 20 |
| Acknowledgements . . . . .                                | 20 |
| Contributors . . . . .                                    | 21 |
| Authors' Addresses . . . . .                              | 21 |

## 1. Introduction

Since very early in computing, operating systems have sought ways to mark which files could be processed by which programs. In Unix, everything is a stream of bytes; identifying the contents of a stream of bytes became a heuristic activity.

For instance, the Unix `file(1)` command, which has existed since 1973 [file], has for decades been able to identify many file formats based upon the contents of the file.

Many systems (Linux, macOS, Windows) will select the correct application based upon the file contents, if the system can not determine it by other means. For instance, in classical MacOS, a resource fork was maintained separately from the file data that included file type information; this way, the OS ideally never needed to know anything about the file data contents to determine the media type.

Many other systems do this by file extensions. Many common web servers derive the media-type information from file extensions.

Having a media type associated with the file contents can avoid some of the brittleness of this approach. When files become disconnected from their type information, such as when attempting to do forensics on a damaged system, then being able to identify the type of information that is stored in file can become very important.

A common way to identify the type of a file from its contents is to place a "magic number" at the start of the file contents [MAGIC]. It is noted that in the media type registration template [RFC6838], a magic number is asked for, if available, as is a file extension.

A challenge for the `file(1)` command is often that it can be confused by the encoding vs. the content. For instance, an Android "apk" (as used to transfer and store an application) may be identified as a ZIP file. Additionally, both OpenOffice and MSOffice files are ZIP files of XML files, and may also be identified as a ZIP file.

As CBOR becomes a more and more common encoding for a wide variety of artifacts, identifying them as just "CBOR" is probably not sufficient. This document provides a way to encode a magic number into the beginning of a CBOR format file. As a CBOR format may use a single CBOR data item or a CBOR sequence of data items [RFC8742], two possible methods of enveloping data are presented; a CBOR Protocol designer will specify one. (A CBOR Protocol is a specification which uses CBOR as its encoding.)

This document also gives advice to designers of CBOR Protocols on choosing one of these mechanisms for identifying their contents. This advice is informative.

A third method is also proposed by which this CBOR format prepended tag is used to identify non-CBOR files. This third method has been placed in Appendix D because it is not about identifying media types containing CBOR-encoded data items. This includes a simple way to derive a magic number to content-formats as defined by [RFC7252], even if not in CBOR form.

Examples of CBOR Protocols currently under development include Concise Software Identification Tags (CoSWID, [I-D.ietf-sacm-coswid]) and Entity Attestation Tokens (EAT, [I-D.ietf-rats-eat]). COSE itself [RFC8152] is considered infrastructure. The encoding of public keys in CBOR as described in [I-D.ietf-cose-cbor-encoded-cert] as `_C509_` would benefit from being an identified CBOR Protocol.

A major inspiration for this document is observing the disarray in certain ASN.1 based systems where most files are PEM encoded; these are then all identified by the extension "pem", confusing public keys, private keys, certificate requests, and S/MIME content.

While the envelopes defined in this specification add information to how data conforming to CBOR Protocols are stored in files, there is no requirement that either type of envelope be transferred on the wire. However, there are some protocols which may benefit from having such a magic number on the wire if they are presently using a different (legacy) encoding scheme. The presence of the identifiable magic sequence can be used to signal that a CBOR Protocol is being used as opposed to a legacy scheme.

## 1.1. Terminology

Byte is a synonym for octet. The term "byte string" refers to the data item defined in [STD94].

The term "file" is understood to stand in a general way for a stored representation that is somewhat detached from the original context of usage of that representation; its usage in this document encompasses similar units of storage that may have different identification schemes such as partitions or media blocks.

The term "diagnostic notation" refers to the human-readable notation for CBOR data items defined in Section 8 of [STD94] and Appendix G of [RFC8610].



The term CDDL (Concise Data Definition Language) refers to the language defined in [RFC8610].

The function TN(ct) is defined in Appendix B.

## 1.2. Requirements for a Magic Number

A magic number is ideally a fingerprint that is unique to a specific CBOR protocol, present in the first few (small multiple of 4) bytes of the file, which does not change when the contents change, and does not depend upon the length of the file.

Less ideal solutions have a pattern that needs to be matched, but in which some bytes need to be ignored. While the Unix file(1) command can be told to ignore certain bytes, this can lead to ambiguities.

## 2. Protocol

This Section presents two enveloping methods. Both use CBOR Tags in a way that results in a deterministic first 8 to 12 bytes. Which one is to be used is up to the CBOR Protocol designer to determine; see Appendix A for some guidance.

### 2.1. The CBOR Protocol Specific Tag

In both enveloping methods, CBOR Protocol designers need to obtain a CBOR tag for each kind of object that they might store in files. As there are more than 4 billion available 4-byte tags, there should be little issue in allocating a few to each available CBOR Protocol.

The IANA policy for 4-byte CBOR Tags is First Come First Served, so all that is required is a simple interaction (e.g., via web or email) with IANA, having filled in the small template provided in Section 9.2 of [STD94]. In the template, it is suggested to include a reference to this specification (RFC XXXX) alongside the Description of semantics.

```
// (Note to RFC Editor: Please replace all occurrences of "RFC XXXX"  
// with the RFC number of the present specification and remove this  
// note.)
```

Allocation of the CBOR tag needs to be initiated by the designer of the CBOR Protocol, who can provide a proposed tag number. In order to be in the four-byte range, and so that there are no leading zero bytes in the four-byte encoding of the tag number, the value needs to be in the range 0x01000000 (decimal 16777216) to 0xFFFFFFFF (decimal 4294967295) inclusive. It is further suggested to avoid values that have an embedded zero byte in the four bytes of their binary representation (such as 0x12003456), as these may confuse implementations that treat the magic number as a C string.

The use of a sequence of four US-ASCII [RFC20] codes which are mnemonic to the protocol is encouraged, but not required (there may be reasons to encode other information into the tag; see Appendix B for an example). For instance, Appendix C uses "OPSN" which translates to the tag number 1330664270 registered for it.

For CBOR data items that form a representation that is described by a CoAP Content-Format Number (Section 12.3 of [RFC7252], Registry CoAP Content-Formats of [IANA.core-parameters]), a tag number has proactively been allocated in Section 4.3 (see Appendix B for details and examples).

## 2.2. Enveloping Method: CBOR Tag Wrapped

The CBOR Tag Wrapped method is appropriate for use with CBOR protocols that encode a single CBOR data item. This data item is enveloped into two nested tags:

The outer tag is a Self-described CBOR tag, 55799, as described in Section 3.4.6 of [STD94].

The tag content of the outer tag is a second CBOR tag whose tag number has been allocated to describe the specific Protocol involved, as discussed in Section 2.1. The tag content of this inner tag is the single CBOR data item.

This method wraps the CBOR data item as CBOR tags usually do. Applications that need to send the stored CBOR data item across a constrained network may wish to remove the two tags if the type is understood from the protocol context, e.g., from a CoAP Content-Format Option (Section 5.10.3 of [RFC7252]). A CBOR Protocol specification may therefore pick the specific cases where the CBOR Tag Wrapped enveloping method is to be used. For instance, it might specify its use for storing the representation in a local file or for Web access, but not within protocol messages that already provide the necessary context.

### 2.2.1. Example

To construct an example without registering a new tag, we use the Content-Format number registered in [RFC8428] for application/senml+cbor (as per Registry Content-Formats of [IANA.core-parameters]), the number 112.

Using the technique described in Appendix B, this translates into the tag TN(112) = 1668546929.

With this tag, the SenML-CBOR pack [{0: "current", 6: 3, 2: 1.5}] would be enveloped as (in diagnostic notation):

```
55799(1668546929([0: "current", 6: 3, 2: 1.5]))
```

Or in hex:

|                |                    |
|----------------|--------------------|
| d9 d9f7        | # tag(55799)       |
| da 63740171    | # tag(1668546929)  |
| 81             | # array(1)         |
| a3             | # map(3)           |
| 00             | # unsigned(0)      |
| 67             | # text(7)          |
| 63757272656e74 | # "current"        |
| 06             | # unsigned(6)      |
| 03             | # unsigned(3)      |
| 02             | # unsigned(2)      |
| f9 3e00        | # primitive(15872) |

At the representation level, the unique fingerprint for application/senml+cbor is composed of the 8 bytes d9d9f7da63740171 hex, after which the unadorned CBOR data (81... for the SenML data) is appended.

### 2.3. Enveloping Method: Labeled CBOR Sequence

The Labeled CBOR Sequence method is appropriate for use with CBOR Sequences as described in [RFC8742].

This method prepends a newly constructed, separate data item to the CBOR Sequence, the `_label_`.

The label is a nesting of two tags, similar to but distinct from the CBOR Tag Wrapped methods, with an inner tag content of a constant byte string. The total length of the label is 12 bytes.

1. The outer tag is the self-described CBOR Sequence tag, 55800.

2. The inner tag is a CBOR tag, from the First Come First Served space, that uniquely identifies the CBOR Protocol. As with CBOR Tag Wrapped, the use of a four-byte tag is encouraged that encodes without zero bytes.
3. The tag content is a three byte CBOR byte string containing 0x42\_4f\_52 ('BOR' in diagnostic notation).

The outer tag in the label identifies the file as being a CBOR Sequence, and does so with all the desirable properties explained in Section 3.4.6 of [STD94]. Specifically, it does not appear to conflict with any known file types, and it is not valid Unicode in any Unicode encoding.

The inner tag in the label identifies which CBOR Protocol is used, as described above.

The inner tag content is a constant byte string which is represented as 0x43\_42\_4f\_52, the ASCII characters "CBOR", which is the CBOR encoded data item for the three-byte string 0x42\_4f\_52 ('BOR' in diagnostic notation).

The actual CBOR Protocol data then follow as the next data item(s) in the CBOR Sequence, without a need for any further specific tag. The use of a CBOR Sequence allows the application to trivially remove the first item with the two tags.

Should this file be reviewed by a human (directly in an editor, or in a hexdump display), it will include the ASCII characters "CBOR" prominently. This value is also included simply because the inner nested tag needs to tag something.

#### 2.3.1. Example

To construct an example without registering a new tag, we use the Content-Format number registered in [RFC9177] for application/missing-blocks+cbor-seq (as per Registry Content-Formats of [IANA.core-parameters]), the number 272.

Using the technique described in Appendix B, this translates into the tag TN(272) = 1668547090.

This is a somewhat contrived example, as this is not a media type that is likely to be committed to storage. Nonetheless, with this tag, missing blocks list 0, 8, 15 would be enveloped as (in diagnostic notation):

```
55800(1668547090('BOR')),  
0,  
8,  
15
```

Or in hex:

```
# CBOR sequence with 4 elements  
d9 d9f8      # tag(55800)  
  da 63740212 # tag(1668547090)  
    43        # bytes(3)  
      424f52  # "BOR"  
00 # unsigned(0)  
08 # unsigned(8)  
0f # unsigned(15)
```

At the representation level, the unique fingerprint for application/missing-blocks+cbor-seq is composed of the 8 bytes d9d9f8da63740212 hex, after which the unadorned CBOR sequence (00... for the missing block list given) is appended.

### 3. Security Considerations

This document provides a way to identify CBOR Protocol objects. Clearly identifying CBOR contents in files may have a variety of impacts.

The most obvious is that it may allow malware to identify interesting stored objects, and then exfiltrate or corrupt them.

Protective applications (that check data) cannot rely on the applications they try to protect (that use the data) to make exactly the same decisions in recognizing file formats. (This is an instance of a check vs. use issue.) For example, end-point assessment technologies should not solely rely on the labeling approaches described in this document to decide whether to inspect a given file. Similarly, depending on operating systems configurations and related properties of the execution environment the labeling might influence the default application used to process a file in a way that may not be predicted by a protective application.

### 4. IANA Considerations

These IANA considerations are entirely about CBOR Tags, in the registry CBOR Tags of [IANA.cbor-tags].

Section 4.1 documents the allocation that was done for a CBOR tag to be used in a CBOR sequence to identify the sequence (an example for using this tag is found in Appendix C). Section 4.3 allocates a CBOR tag for each actual or potential CoAP Content-Format number (examples are in Appendix B).

#### 4.1. Labeled CBOR Sequence Tag

IANA has allocated tag 55800 as the tag for the Labeled CBOR Sequence Enveloping Method from the CBOR Tags Registry. IANA is asked to update this tag registration to point to this document.

This tag is from the First Come/First Served area.

The value has been picked to have properties similar to the 55799 tag (Section 3.4.6 of [STD94]).

The hexadecimal representation of the encoded tag head is:  
0xd9\_d9\_f8.

This is not valid UTF-8: the first 0xd9 introduces a three-byte sequence in UTF-8, but the 0xd9 as the second value is not a valid second byte for UTF-8.

This is not valid UTF-16: the byte sequence 0xd9d9 (in either endian order) puts this value into the UTF-16 high-half zone, which would signal that this a 32-bit Unicode value. However, the following 16-bit big-endian value 0xf8.. is not a valid second sequence according to [RFC2781]. On a little-endian system, it would be necessary to examine the fourth byte to determine if it is valid. That next byte is determined by the subsequent encoding, and Section 3.4.6 of [STD94] has already determined that no valid CBOR encodings result in valid UTF-16.

Data Item:  
tagged byte string

Semantics:  
indicates that the file contains CBOR Sequences

#### 4.2. CBOR-Labeled Non-CBOR Data Tag

IANA is requested to allocate tag 55801 as the tag for the CBOR-Labeled Non-CBOR Data Enveloping Method (Appendix D) from the CBOR Tags Registry. IANA is asked to update this tag registration to point to this document.

This tag is from the First Come/First Served area.

The value has been picked to have properties similar to the 55799 tag (Section 3.4.6 of [STD94]).

The hexadecimal representation of the encoded tag head is:  
0xd9\_d9\_f9.

This is not valid UTF-8: the first 0xd9 introduces a three-byte sequence in UTF-8, but the 0xd9 as the second value is not a valid second byte for UTF-8.

This is not valid UTF-16: the byte sequence 0xd9d9 (in either endian order) puts this value into the UTF-16 high-half zone, which would signal that this a 32-bit Unicode value. However, the following 16-bit big-endian value 0xf9.. is not a valid second sequence according to [RFC2781]. On a little-endian system, it would be necessary to examine the fourth byte to determine if it is valid. That next byte is determined by the subsequent encoding, and Section 3.4.6 of [STD94] has already determined that no valid CBOR encodings result in valid UTF-16.

Data Item:  
tagged byte string

Semantics:  
indicates that the file starts with a CBOR-Labeled Non-CBOR Data label.

#### 4.3. CBOR Tags for CoAP Content-Format Numbers

IANA is requested to allocate the tag numbers 1668546817 (0x63740101) to 1668612095 (0x6374ffff) as follows:

Data Item:  
byte string or any CBOR data item (see Appendix B of RFC XXXX)

Semantics:  
the representation of content-format  $ct < 65025$  is indicated by tag number  
 $TN(ct) = 0x63740101 + (ct / 255) * 256 + ct \% 255$

Reference:  
RFC XXXX

The Registry for Content-Formats of [IANA.core-parameters] has been defined in Section 12.3 of [RFC7252].

#### 5. References

## 5.1. Normative References

- [C] International Organization for Standardization, "Information technology Programming languages C", ISO/IEC 9899:2018, Fourth Edition, June 2018, <<https://www.iso.org/standard/74528.html>>.
- [RFC8742] Bormann, C., "Concise Binary Object Representation (CBOR) Sequences", RFC 8742, DOI 10.17487/RFC8742, February 2020, <<https://www.rfc-editor.org/info/rfc8742>>.
- [STD94] 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/info/rfc8949>>.

## 5.2. Informative References

- [file] Wikipedia, "file (command)", 20 January 2021, <[https://en.wikipedia.org/wiki/File\\_%28command%29](https://en.wikipedia.org/wiki/File_%28command%29)>.
- [I-D.ietf-cose-cbor-encoded-cert] Mattsson, J. P., Selander, G., Raza, S., Höglund, J., and M. Furuheid, "CBOR Encoded X.509 Certificates (C509 Certificates)", Work in Progress, Internet-Draft, draft-ietf-cose-cbor-encoded-cert-03, 10 January 2022, <<https://www.ietf.org/archive/id/draft-ietf-cose-cbor-encoded-cert-03.txt>>.
- [I-D.ietf-rats-eat] Lundblade, L., Mandyam, G., and J. O'Donoghue, "The Entity Attestation Token (EAT)", Work in Progress, Internet-Draft, draft-ietf-rats-eat-12, 24 February 2022, <<https://www.ietf.org/archive/id/draft-ietf-rats-eat-12.txt>>.
- [I-D.ietf-sacm-coswid] Birkholz, H., Fitzgerald-McKay, J., Schmidt, C., and D. Waltermire, "Concise Software Identification Tags", Work in Progress, Internet-Draft, draft-ietf-sacm-coswid-21, 7 March 2022, <<https://www.ietf.org/archive/id/draft-ietf-sacm-coswid-21.txt>>.
- [IANA.cbor-tags] IANA, "Concise Binary Object Representation (CBOR) Tags", <<https://www.iana.org/assignments/cbor-tags>>.



- [IANA.core-parameters] IANA, "Constrained RESTful Environments (CoRE) Parameters",  
<<https://www.iana.org/assignments/core-parameters>>.
- [MAGIC] Ritchie, D., "archive (library) file format", in Bell Labs, Unix Programmer's Manual, First Edition: File Formats, 3 November 1971,  
<<https://www.bell-labs.com/usr/dmr/www/man51.pdf#page=4>>.
- [RFC20] Cerf, V., "ASCII format for network interchange", STD 80, RFC 20, DOI 10.17487/RFC0020, October 1969,  
<<https://www.rfc-editor.org/info/rfc20>>.
- [RFC2781] Hoffman, P. and F. Yergeau, "UTF-16, an encoding of ISO 10646", RFC 2781, DOI 10.17487/RFC2781, February 2000,  
<<https://www.rfc-editor.org/info/rfc2781>>.
- [RFC6838] Freed, N., Klensin, J., and T. Hansen, "Media Type Specifications and Registration Procedures", BCP 13, RFC 6838, DOI 10.17487/RFC6838, January 2013,  
<<https://www.rfc-editor.org/info/rfc6838>>.
- [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>>.
- [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>>.
- [RFC8152] Schaad, J., "CBOR Object Signing and Encryption (COSE)", RFC 8152, DOI 10.17487/RFC8152, July 2017,  
<<https://www.rfc-editor.org/info/rfc8152>>.
- [RFC8428] Jennings, C., Shelby, Z., Arkko, J., Keranen, A., and C. Bormann, "Sensor Measurement Lists (SenML)", RFC 8428, DOI 10.17487/RFC8428, August 2018,  
<<https://www.rfc-editor.org/info/rfc8428>>.
- [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/info/rfc8610>>.

- [RFC9177] Boucadair, M. and J. Shallow, "Constrained Application Protocol (CoAP) Block-Wise Transfer Options Supporting Robust Transmission", RFC 9177, DOI 10.17487/RFC9177, March 2022, <<https://www.rfc-editor.org/info/rfc9177>>.
- [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)", ITU-T Recommendation X.690, ISO/IEC 8825-1, February 2021.

## Appendix A. Advice to Protocol Designer

This document introduces a choice between wrapping a single CBOR data item into a (pair of) identifying CBOR tags, or prepending an identifying encoded CBOR data item (which in turn contains a pair of identifying CBOR tags) to a CBOR Sequence (which might be single data item).

Which should a protocol designer use?

In this discussion, one assumes that there is an object stored in a file, perhaps specified by a system operator in a configuration file.

For example: a private key used in COSE operations, a public key/certificate in C509 ([I-D.ietf-cose-cbor-encoded-cert]) or CBOR format, a recorded sensor reading stored for later transmission, or a COVID-19 vaccination certificate that needs to be displayed in QR code form.

Both the Labeled CBOR Sequence and the wrapped tag can be trivially removed by an application before sending the CBOR content out on the wire.

The Labeled CBOR Sequence can be slightly easier to remove as in most cases, CBOR parsers will return it as a unit, and then return the actual CBOR item, which could be anything at all, and could include CBOR tags that `_do_` need to be sent on wire.

On the other hand, having the Labeled CBOR Sequence in the file requires that all programs that expect to examine that file are able to skip what appears to be a CBOR item with two tags nested around a three-byte byte string. The three byte entry is not of the format the program would normally have processed, so it may be a surprise. On the other hand, CBOR parsers are generally tolerant of tags that appear: many of them will process extra tags, making unknown tags available as meta information. A program that is not expecting those tags may just ignore those extra tags.

As an example of where there was a problem with previous security systems, "PEM" format certificate files grew to be able to contain multiple certificates by simple concatenation. The PKCS1 format [RFC8017] could also contain a private key object followed by a one or more certificate objects: but only when in PEM format. Annoyingly, when in binary DER format ([X.690], which like CBOR is self-delimiting), concatenation of certificates was not compatible with most programs as they did not expect to read more than one item in the file.

The use of CBOR Tag Wrapped format is easier to retrofit to an existing format with existing and unchangeable stored format for a single CBOR data item. This new sequence of tags is expected to be trivially ignored by many existing programs when reading CBOR from files or similar units of storage, even if the program only supports decoding a single data item (and not a CBOR sequence). But, a naive program might also then transmit the additional tags across the network. Removing the CBOR Tag Wrapped format requires knowledge of the two tags involved. Other tags present might be needed.

For a representation matching a specific media-type that is carried in a CBOR byte string, the byte string head will already have to be removed for use as such a representation, so it should be easy to remove the enclosing tag heads as well. This is of particular interest with the pre-defined tags provided by Appendix B for media-types with CoAP Content-Format numbers.

Here are some considerations in the form of survey questions:

A.1. Is the on-wire format new?

If the on-wire format is new, then it could be specified with the CBOR Tag Wrapped format if the extra eight bytes are not a problem. The stored format is then identical to the on-wire format.

If the eight bytes are a problem on the wire (and they often are if CBOR is being considered), then the Labeled CBOR Sequence format should be adopted for the stored format.

A.2. Can many items be trivially concatenated?

If the programs that read the contents of the file already expect to process all of the CBOR data items in the file (not just the first), then the Labeled CBOR Sequence format may be easily retrofitted.

The program involved may throw errors or warnings on the Labeled CBOR Sequence if they have not yet been updated, but this may not be a problem.

There are situations where multiple objects may be concatenated into a single file. If each object is preceded by a Labeled CBOR Sequence label then there may be multiple such labels in the file.

A protocol based on CBOR Sequences may specify that Labeled CBOR Sequence labels can occur within a CBOR Sequence, possibly even to switch to data items following in the sequence that are of a different type.

If the CBOR Sequence based protocol does not define the semantics for or at least tolerate embedded labels, care must be taken when concatenating Labeled CBOR Sequences to remove the label from all but the first part.

As an example from legacy PEM encoded PKIX certificates, many programs accept a series of PKIX certificates in a single file in order to set up a certificate chain. The file would contain not just the End-Entity (EE) certificate, but also any subordinate certification authorities (CA) needed to validate the EE. This mechanism actually only works for PEM encoded certificates, and not DER encoded certificates. One of the reasons for this specification is to make sure that CBOR encoded certificates do not suffer from this problem.

As an example of mixing of types, some TLS server programs also can accept both their PEM encoded private key, and their PEM encoded certificate in the same file.

If only one item is ever expected in the file, the use of Labeled CBOR Sequence may present an implementation hurdle to programs that previously just read a single data item and used it.

#### A.3. Are there tags at the start?

If the Protocol expects to use other tags at its top-level, then the use of the CBOR Tag Wrapped format may be easy to explain at the same place in the protocol description.

#### Appendix B. CBOR Tags for CoAP Content Formats

Section 5.10.3 of [RFC7252] defines the concept of a Content-Format, which is a short 16-bit unsigned integer that identifies a specific content type (media type plus optionally parameters), optionally together with a content encoding.

Outside of a transfer protocol that indicates the Content-Format for a representation, it may be necessary to identify the Content-Format of the representation when it is stored in a file, in firmware, or when debugging.

This specification allocates CBOR tag numbers 1668546817 (0x63740101) to 1668612095 (0x6374FFFF) for the tagging of representations of specific content formats.

Using tags from this range, a byte string that is to be interpreted as a representation of Content-Format number *ct*, with  $ct < 65025$  ( $255 \times 255$ ), can be identified by enclosing it in a tag with tag number  $TN(ct)$  where:

$$TN(ct) = 0x63740101 + (ct / 255) * 256 + ct \% 255.$$

(where +, \*, / and % stand for integer addition, multiplication, division and remainder as in the programming language C [C].)

This formula avoids the use of zero bytes in the representation of the tag number.

Note that no tag numbers are assigned for Content-Format numbers in the range 65025  $ct$  65535. (This range is in the range reserved by Section 12.3 of [RFC7252] for experimental use. The overlap of 25 code points between this experimental range with the range this appendix defines tag numbers for can be used for experiments that want to employ a tag number.)

Exceptionally, when used immediately as tag content of one of the tags 55799, 55800, or 55801, the tag content is as follows:

Tag 55799 (Section 2.2): One of:

1. The CBOR data item within the representation (without byte string wrapping). This only works for Content Formats that are represented by a single CBOR data item in identity content-coding.
2. The data items in the CBOR sequence within the representation, without byte string wrapping, but wrapped in a CBOR array. This works for Content Formats that are represented by a CBOR sequence in identity content-coding.

Tags 55800 (Section 2.3) or 55801 (Appendix D): the byte string 'BOR', signifying that the representation of the given content-format follows in the file, in the way defined for these tags.

### B.1. Content-Format Tag Examples

Registry Content-Formats of [IANA.core-parameters] defines content formats that can be used as examples:

- \* As discussed in Section 2.2.1, Content-Format 112 stands for media type application/senml+cbor (no parameters). The corresponding tag number is TN(112) = 1668546929.

So the following CDDL snippet can be used to identify application/senml+cbor representations:

```
senml-cbor = #6.1668546929(bstr)
```

Note that a byte string is used as the type of the tag content, because a media type representation in general can be any byte string.

- \* Content-Format 272 stands for media type application/missing-blocks+cbor-seq, a CBOR sequence [RFC9177].

The corresponding tag number is TN(272) = 1668547090.

So the following CDDL snippet can be used to identify application/missing-blocks+cbor-seq representations as embedded in a CBOR byte string:

```
missing-blocks = #6.1668547090(bstr)
```

### Appendix C. Example from Openswan

The Openswan IPsec project has a daemon ("pluto"), and two control programs ("addconn", and "whack"). They communicate via a Unix-domain socket, over which a C-structure containing pointers to strings is serialized using a bespoke mechanism. This is normally not a problem as the structure is compiled by the same compiler; but when there are upgrades it is possible for the daemon and the control programs to get out of sync by the bespoke serialization. As a result, there are extra compensations to deal with shutting the daemon down. During testing, it is sometimes the case that upgrades are backed out.

In addition, when doing unit testing, the easiest way to load policy is to use the normal policy reading process, but that is not normally loaded in the daemon. Instead, the IPC that is normally sent across the wire is compiled/serialized and placed in a file. The above magic number is included in the file, and also on the IPC in order to distinguish the "shutdown" command CBOR operation.

In order to reduce the problems due to serialization, the serialization is being changed to CBOR. Additionally, this change allows the IPC to be described by CDDL, and for any language that encode to CBOR can be used.

IANA has allocated the tag 1330664270, or 0x4f\_50\_53\_4e for this purpose. As a result, each file and each IPC is prefixed with a CBOR Tag Sequence.

In diagnostic notation:

```
55800(1330664270(h'424F52'))
```

Or in hex:

```
d9 d9f8      # tag(55800)
da 4f50534e  # tag(1330664270)
  43         # bytes(3)
    424f52   # "BOR"
```

#### Appendix D. Using CBOR Labels for non-CBOR data

The CBOR-Labeled non-CBOR data method is appropriate for adding a magic number to a non-CBOR data format, particularly one that can be described by a Content-Format tag (Appendix B).

This method prepends a CBOR data item to the non-CBOR data; this data item is called the "header" and, similarly to the Labeled CBOR-Sequence label, consists of two nested tags around a constant byte string for a total of 12 bytes.

1. The outer tag is the CBOR-Labeled Non-CBOR Data tag, 55801.
2. The inner tag is a CBOR tag, from the First Come First Served space, that uniquely identifies the CBOR Protocol. As with CBOR Tag Wrapped, the use of a four-byte tag is encouraged that encodes without zero bytes.
3. The tag content is a three byte CBOR byte string containing 0x42\_4F\_52 ('BOR' in diagnostic notation).

The outer tag in the label identifies the file as being file as being prefixed by a non-CBOR data label, and does so with all the desirable properties explained in Section 3.4.6 of [STD94]. Specifically, it does not appear to conflict with any known file types, and it is not valid Unicode in any Unicode encoding.

The inner tag in the label identifies which non-CBOR Protocol is used.

The inner tag content is a constant byte string which is represented as 0x43\_42\_4f\_52, the ASCII characters "CBOR", which is the CBOR encoded data item for the three-byte string 0x42\_4f\_52 ('BOR' in diagnostic notation).

The actual non-CBOR Protocol data then follow directly appended to the CBOR representation of the header. This allows the application to trivially remove the header item with the two nested tags and the byte string.

As with the Labeled CBOR Sequence {#sequences}, this choice of the tag content places the ASCII characters "CBOR" prominently into the header.

#### D.1. Content-Format Tag Examples

Registry Content-Formats of [IANA.core-parameters] defines content formats that can be used as examples:

- \* Content-Format 432 stands for media type application/td+json (no parameters). The corresponding tag number is TN(432) = 1668547250.

So the following CDDL snippet can be used to identify a CBOR-Labeled non-CBOR data for application/td+json representations:

```
td-json-header = #6.55801(#6.1668547250('BOR'))
```

- \* Content-Format 11050 stands for media type application/json in deflate content-coding.

The corresponding tag number is TN(11050) = 1668557910.

So the following CDDL snippet can be used to identify a CBOR-Labeled non-CBOR data for application/json representations compressed in deflate content-coding:

```
json-deflate-header = #6.55801(#6.1668557910('BOR'))
```

#### Appendix E. Changelog

#### Acknowledgements

The CBOR WG brainstormed this protocol on January 20, 2021 via a number of productive email exchanges on the mailing list.



Contributors

Josef 'Jeff' Sipek  
Email: jeffpc@josefsipek.net

Authors' Addresses

Michael Richardson  
Sandelman Software Works  
Email: mcr+ietf@sandelman.ca

Carsten Bormann  
Universität Bremen TZI  
Postfach 330440  
D-28359 Bremen  
Germany  
Phone: +49-421-218-63921  
Email: cabo@tzi.org

CBOR Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: 25 April 2022

M. Richardson  
Sandelman Software Works  
C. Bormann  
Universität Bremen TZI  
22 October 2021

CBOR tags for IPv4 and IPv6 addresses and prefixes  
draft-ietf-cbor-network-addresses-13

Abstract

This specification defines two CBOR Tags for use with IPv6 and IPv4 addresses and prefixes.

// RFC-EDITOR-please-remove: This work is tracked at  
// <https://github.com/cbor-wg/cbor-network-address>

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 25 April 2022.

Copyright Notice

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

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<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. Terminology . . . . .                           | 3  |
| 3. Protocol . . . . .                              | 3  |
| 3.1. Three Forms . . . . .                         | 3  |
| 3.1.1. Addresses . . . . .                         | 3  |
| 3.1.2. Prefixes . . . . .                          | 3  |
| 3.1.3. Interface Definition . . . . .              | 4  |
| 3.2. IPv6 . . . . .                                | 4  |
| 3.3. IPv4 . . . . .                                | 5  |
| 4. Tag validity . . . . .                          | 6  |
| 4.1. Deterministic Encoding . . . . .              | 6  |
| 4.2. Encoder Considerations for Prefixes . . . . . | 6  |
| 4.3. Decoder Considerations for Prefixes . . . . . | 7  |
| 4.3.1. Example implementation . . . . .            | 7  |
| 5. CDDL . . . . .                                  | 8  |
| 6. Security Considerations . . . . .               | 9  |
| 7. IANA Considerations . . . . .                   | 10 |
| 7.1. Tag 54 - IPv6 . . . . .                       | 10 |
| 7.2. Tag 52 - IPv4 . . . . .                       | 10 |
| 7.3. Tags 260 and 261 . . . . .                    | 10 |
| 8. References . . . . .                            | 10 |
| 8.1. Normative References . . . . .                | 10 |
| 8.2. Informative References . . . . .              | 11 |
| Appendix A. Changelog . . . . .                    | 11 |
| Acknowledgements . . . . .                         | 11 |
| Authors' Addresses . . . . .                       | 11 |

## 1. Introduction

[RFC8949] defines a number of CBOR Tags for common items. Tags 260 and 261 were later defined in drafts listed with IANA [IANA.cbor-tags]. These tags were intended to cover addresses (260) and prefixes (261). Tag 260 distinguishes between IPv6, IPv4, and MAC [RFC7042] addresses only through the length of the byte string, making it impossible, for example, to drop trailing zeros in the encoding of IP addresses. Tag 261 was not documented well enough for use.

This specification defines tags 54 and 52 achieving an explicit indication of IPv6 or IPv4 by the tag number. These new tags are intended to be used in preference to tags 260 and 261. They provide formats for IPv6 and IPv4 addresses, prefixes, and addresses with prefixes, achieving an explicit indication of IPv6 or IPv4. The prefix format omits trailing zeroes in the address part. (Due to the complexity of testing, the value of omitting trailing zeros for the pure address format was considered non-essential and support for that is not provided in this specification.) This specification does not deal with MAC addresses (Section 2 of [RFC7042]) such as they are used for Ethernet.

## 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. Protocol

### 3.1. Three Forms

#### 3.1.1. Addresses

These tags can be applied to byte strings to represent a single address.

This form is called the Address Format.

#### 3.1.2. Prefixes

When applied to an array that starts with an unsigned integer, they represent a CIDR-style prefix of that length.

When the Address Format (i.e., without prefix) appears in a context where a prefix is expected, then it is to be assumed that all bits are relevant. That is, for IPv4, a /32 is implied, and for IPv6, a /128 is implied.

This form is called the Prefix Format.

### 3.1.3. Interface Definition

When applied to an array that starts with a byte string, which stands for an IP address, followed by an unsigned integer giving the bit length of a prefix built out of the first length bits of the address, they represent information that is commonly used to specify both the network prefix and the IP address of an interface.

The length of the byte string is always 16 bytes (for IPv6) and 4 bytes (for IPv4).

This form is called the Interface Format.

Interface Format definitions support an optional third element to the array, which is to be used as the IPv6 Link-Local zone identifier from Section 6 of [RFC4007]; for symmetry this is also provided for IPv4 as in [RFC4001] and [RFC6991]. The zone identifier may be an integer, in which case it is to be interpreted as the interface index. It may be a text string, in which case it is to be interpreted as an interface name.

As explained in [RFC4007] the zone identifiers are strictly local to the node. They are useful for communications within a node about connected addresses (for instance, where a link-local peer is discovered by one daemon, and another daemon needs to be informed). They may also have utility in some management protocols.

In the cases where the Interface Format is being used to represent only an address with a zone identifier, and no interface prefix information, then the prefix length may be replaced with the CBOR "null" (0xF6).

### 3.2. IPv6

IANA has allocated tag 54 for IPv6 uses. (This is the ASCII code for '6'.)

An IPv6 address is to be encoded as a sixteen-byte byte string (Section 3.1 of [RFC8949], major type 2), enclosed in Tag number 54.

For example:

54(h'20010db81234deedbeefcafeacefeed')

An IPv6 prefix, such as 2001:db8:1234::/48 is to be encoded as a two element array, with the length of the prefix first. See Section 4 for the detailed construction of the second element.

For example:

```
54([48, h'20010db81234'])
```

An IPv6 address combined with a prefix length, such as being used for configuring an interface, is to be encoded as a two element array, with the (full-length) IPv6 address first and the length of the associated network the prefix next; a third element can be added for the zone identifier.

For example:

```
54([h'20010db81234deedbeefcafeacefeed', 56])
```

The address-with-prefix form can be reliably distinguished from the prefix form only in the sequence of the array elements.

Some example of a link-local IPv6 address with a 64-bit prefix:

```
54([h'fe80000000000020202fffffe030303', 64, 'eth0'])
```

with a numeric zone identifier:

```
54([h'fe80000000000020202fffffe030303', 64, 42])
```

An IPv6 link-local address without a prefix length:

```
54([h'fe80000000000020202fffffe030303', null, 42])
```

Zone identifiers may be used with any kind of IP address, not just Link-Local addresses. In particular, they are valid for multicast addresses, and there may still be some significance for Globally Unique Addresses (GUA).

### 3.3. IPv4

IANA has allocated tag 52 for IPv4 uses. (This is the ASCII code for '4'.)

An IPv4 address is to be encoded as a four-byte byte string (Section 3.1 of [RFC8949], major type 2), enclosed in Tag number 52.

For example:

```
52(h'c0000201')
```

An IPv4 prefix, such as 192.0.2.0/24 is to be encoded as a two element array, with the length of the prefix first. See Section 4 for the detailed construction of the second element.

For example:

```
52([24, h'c00002'])
```

An IPv4 address combined with a prefix length, such as being used for configuring an interface, is to be encoded as a two element array, with the (full-length) IPv4 address first and the length of the associated network the prefix next; a third element can be added for the zone identifier.

For example, 192.0.2.1/24 is to be encoded as a two element array, with the length of the prefix (implied 192.0.2.0/24) last.

```
52([h'c0000201', 24])
```

The address-with-prefix form can be reliably distinguished from the prefix form only in the sequence of the array elements.

#### 4. Tag validity

This section discusses when a tag 54 or tag 52 is valid (Section 5.3.2 of [RFC8949]). As with all CBOR tags, validity checking can be handled in a generic CBOR library or in the application. A generic CBOR library needs to document whether and how it handles validity checking.

The rule ip-address-or-prefix in Figure 1 shows how to check the overall structure of these tags and their content, the ranges of integer values, and the lengths of byte strings. An instance of tag 52 or 54 is valid if it matches that rule and, for ipv6-prefix and ipv4-prefix, the considerations of Sections 4.2 and 4.3.

##### 4.1. Deterministic Encoding

The tag validity rules, combined with the rules in Section 4.2.1 of [RFC8949], lead to deterministic encoding for tags 54 and 52 and require no further Additional Deterministic Encoding Considerations as per Section 4.2.2 of [RFC8949].

##### 4.2. Encoder Considerations for Prefixes

For the byte strings used as the second element in the array representing a prefix:

(1) An encoder MUST set any unused bytes, and any unused bits in the final byte, if any, to zero. Unused bytes/bits are bytes/bits that are not covered by the prefix length given. So for example, 2001:db8:1230::/44 MUST be encoded as:

```
54([44, h'20010db81230'])
```

even though variations like:

```
54([44, h'20010db81233'])
```

```
54([44, h'20010db8123f'])
```

```
54([44, h'20010db8123012'])
```

start with the same 44 bits, but are not valid.

(Analogous examples can be constructed for IPv4 prefixes.)

(2) An encoder MUST then omit any right-aligned (trailing) sequence of bytes that are all zero.

There is no relationship between the number of bytes omitted and the prefix length. For instance, the prefix 2001:db8::/64 is encoded as:

```
54([64, h'20010db8'])
```

#### 4.3. Decoder Considerations for Prefixes

A decoder MUST check that all unused bits encoded in the byte string `ipv6-prefix-bytes/ipv4-prefix-bytes`, i.e., the bits to the right of the prefix length, are zero.

A decoder MUST also check that the byte string does not end in a zero byte.

Since encoders are required to remove zero-valued trailing bytes, a decoder MUST handle the case where a prefix length specifies that more bits are relevant than are actually present in the byte-string.

As an example, `::/128` is encoded as

```
54([128, h''])
```

##### 4.3.1. Example implementation

A recommendation for prefix decoder implementations is to first create an array of 16 (or 4) zero bytes.



Then taking whichever is smaller between (a) the length of the included byte-string, and (b) the number of bytes covered by the prefix-length rounded up to the next multiple of 8: fail if that number is greater than 16 (or 4), and then copy that many bytes from the byte-string into the byte array.

Finally, looking at the number of unused bits in the last byte (if any) of the range covered by the prefix length, check that any unused bits in the byte string are zero:

```
unused_bits = (8 - (prefix_length_in_bits % 8)) % 8;
if (length_in_bytes > 0 &&
    (address_bytes[length_in_bytes - 1] & ~(0xFF << unused_bits))
    != 0)
    fail();
```

## 5. CDDL

For use with CDDL [RFC8610], the typenames defined in Figure 1 are recommended:

```

ip-address-or-prefix = ipv6-address-or-prefix /
                        ipv4-address-or-prefix

ipv6-address-or-prefix = #6.54(ipv6-address /
                                ipv6-address-with-prefix /
                                ipv6-prefix)
ipv4-address-or-prefix = #6.52(ipv4-address /
                                ipv4-address-with-prefix /
                                ipv4-prefix)

ipv6-address = bytes .size 16
ipv4-address = bytes .size 4

ipv6-address-with-prefix = [ipv6-address,
                            ipv6-prefix-length / null,
                            ?ip-zone-identifier]
ipv4-address-with-prefix = [ipv4-address,
                            ipv4-prefix-length / null,
                            ?ip-zone-identifier]

ipv6-prefix-length = 0..128
ipv4-prefix-length = 0..32

ipv6-prefix = [ipv6-prefix-length, ipv6-prefix-bytes]
ipv4-prefix = [ipv4-prefix-length, ipv4-prefix-bytes]

ipv6-prefix-bytes = bytes .size (uint .le 16)
ipv4-prefix-bytes = bytes .size (uint .le 4)

ip-zone-identifier = uint / text

```

Figure 1: CDDL types for tags 54 and 52

## 6. Security Considerations

This document provides an CBOR encoding for IPv4 and IPv6 address information. Any applications using these encodings will need to consider the security implications of these data in their specific context. For example, identifying which byte sequences in a protocol are addresses may allow an attacker or eavesdropper to better understand what parts of a packet to attack.

Applications need to check the validity (Section 4) of a tag before acting on any of its contents. If the validity checking is not done in the generic CBOR decoder, it needs to be done in the application; in any case it needs to be done before the tag is transformed into a platform-specific representation that could conceal validity errors.

The right-hand bits of the prefix, after the prefix-length, are set to zero by this protocol. (Otherwise, a malicious party could use them to transmit covert data in a way that would not affect the primary use of this encoding. Such abuse is detected by tag validity checking, and can also be detected by examination of the raw protocol bytes.)

## 7. IANA Considerations

IANA has allocated two tags from the Specification Required area of the Concise Binary Object Representation (CBOR) Tags [IANA.cbor-tags]:

### 7.1. Tag 54 - IPv6

Data Item: byte string or array  
Semantics: IPv6, [prefixlen,IPv6], [IPv6,prefixpart]

### 7.2. Tag 52 - IPv4

Data Item: byte string or array  
Semantics: IPv4, [prefixlen,IPv4], [IPv4,prefixpart]

### 7.3. Tags 260 and 261

IANA is requested to add the note "DEPRECATED in favor of 52 and 54 for IP addresses" to registrations 260 and 261

## 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>>.
- [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>>.
- [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/info/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/info/rfc8949>>.

## 8.2. Informative References

- [IANA.cbor-tags] IANA, "Concise Binary Object Representation (CBOR) Tags", <<https://www.iana.org/assignments/cbor-tags>>.
- [RFC4001] Daniele, M., Haberman, B., Routhier, S., and J. Schoenwaelder, "Textual Conventions for Internet Network Addresses", RFC 4001, DOI 10.17487/RFC4001, February 2005, <<https://www.rfc-editor.org/info/rfc4001>>.
- [RFC4007] Deering, S., Haberman, B., Jinmei, T., Nordmark, E., and B. Zill, "IPv6 Scoped Address Architecture", RFC 4007, DOI 10.17487/RFC4007, March 2005, <<https://www.rfc-editor.org/info/rfc4007>>.
- [RFC6991] Schoenwaelder, J., Ed., "Common YANG Data Types", RFC 6991, DOI 10.17487/RFC6991, July 2013, <<https://www.rfc-editor.org/info/rfc6991>>.
- [RFC7042] Eastlake 3rd, D. and J. Abley, "IANA Considerations and IETF Protocol and Documentation Usage for IEEE 802 Parameters", BCP 141, RFC 7042, DOI 10.17487/RFC7042, October 2013, <<https://www.rfc-editor.org/info/rfc7042>>.

## Appendix A. Changelog

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

\* 03

\* 02

\* 01 added security considerations about covert channel

## Acknowledgements

Roman Danyliw, Donald Eastlake, Ben Kaduk, Barry Leiba, and Éric Vyncke reviewed the document and provided suggested text. Jürgen Schönwälder helped finding the history of IPv4 zone identifiers.

## Authors' Addresses

Michael Richardson  
Sandelman Software Works

Email: [mcr+ietf@sandelman.ca](mailto:mcr+ietf@sandelman.ca)

Carsten Bormann  
Universität Bremen TZI  
Germany

Email: [cabo@tzi.org](mailto:cabo@tzi.org)

Network Working Group  
Internet-Draft  
Intended status: Informational  
Expires: 22 October 2022

C. Bormann  
Universität Bremen TZI  
20 April 2022

Packed CBOR  
draft-ietf-cbor-packed-05

## Abstract

The Concise Binary Object Representation (CBOR, RFC 8949 == STD 94) 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.

CBOR does not provide any forms of data compression. CBOR data items, in particular when generated from legacy data models, often allow considerable gains in compactness when applying data compression. While traditional data compression techniques such as DEFLATE (RFC 1951) can work well for CBOR encoded data items, their disadvantage is that the receiver needs to uncompress the compressed form to make use of the data.

This specification describes Packed CBOR, a simple transformation of a CBOR data item into another CBOR data item that is almost as easy to consume as the original CBOR data item. A separate decompression step is therefore often not required at the receiver.

## Note to Readers

This is a working-group draft of the CBOR working group of the IETF, <https://datatracker.ietf.org/wg/cbor/about/> (<https://datatracker.ietf.org/wg/cbor/about/>). Discussion takes places on the GitHub repository <https://github.com/cbor-wg/cbor-packed> (<https://github.com/cbor-wg/cbor-packed>) and on the CBOR WG mailing list, <https://www.ietf.org/mailman/listinfo/cbor> (<https://www.ietf.org/mailman/listinfo/cbor>).

## 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 22 October 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 . . . . .               | 2  |
| 1.1. Terminology . . . . .              | 3  |
| 2. Packed CBOR . . . . .                | 4  |
| 2.1. Packing Tables . . . . .           | 4  |
| 2.2. Referencing Shared Items . . . . . | 5  |
| 2.3. Referencing Affix Items . . . . .  | 6  |
| 2.4. Discussion . . . . .               | 7  |
| 3. Table Setup . . . . .                | 8  |
| 3.1. Basic Packed CBOR . . . . .        | 9  |
| 4. IANA Considerations . . . . .        | 10 |
| 5. Security Considerations . . . . .    | 11 |
| 6. References . . . . .                 | 12 |
| 6.1. Normative References . . . . .     | 12 |
| 6.2. Informative References . . . . .   | 12 |
| Appendix A. Examples . . . . .          | 13 |
| Acknowledgements . . . . .              | 17 |
| Author's Address . . . . .              | 18 |

## 1. Introduction

The Concise Binary Object Representation (CBOR, [STD94]) 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.

CBOR does not provide any forms of data compression. CBOR data items, in particular when generated from legacy data models, often allow considerable gains in compactness when applying data compression. While traditional data compression techniques such as DEFLATE [RFC1951] can work well for CBOR encoded data items, their disadvantage is that the receiver needs to uncompress the compressed form to make use of the data.

This specification describes Packed CBOR, a simple transformation of a CBOR data item into another CBOR data item that is almost as easy to consume as the original CBOR data item. A separate decompression step is therefore often not required at the receiver.

This document defines the Packed CBOR format by specifying the transformation from a Packed CBOR data item to the original CBOR data item; it does not define an algorithm for a packer. Different packers can differ in the amount of effort they invest in arriving at a minimal packed form; often, they simply employ the sharing that is natural for a specific application.

Packed CBOR can make use of two kinds of optimization:

- \* item sharing: substructures (data items) that occur repeatedly in the original CBOR data item can be collapsed to a simple reference to a common representation of that data item. The processing required during consumption is limited to following that reference.
- \* affix sharing: data items (strings, containers) that share a prefix or suffix (affix) can be replaced by a reference to a common affix plus the rest of the data item. For strings, the processing required during consumption is similar to following the affix reference plus that for an indefinite-length string.

A specific application protocol that employs Packed CBOR might allow both kinds of optimization or limit the representation to item sharing only.

Packed CBOR is defined in two parts: Referencing packing tables (Section 2) and setting up packing tables (Section 3).

### 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.



Packed reference: A shared item reference or an affix reference.

Shared item reference: A reference to a shared item as defined in Section 2.2.

Affix reference: A reference that combines an affix item as defined in Section 2.3.

Affix: Prefix or suffix.

Packing tables: The triple of a shared item table, a prefix table, and a suffix table.

Current set: The packing tables in effect at the data item under consideration.

Expansion: The result of applying a packed reference in the context of given Packing tables.

The definitions of [STD94] apply. Specifically: The term "byte" is used in its now customary sense as a synonym for "octet"; "byte strings" are CBOR data items carrying a sequence of zero or more (binary) bytes, while "text strings" are CBOR data items carrying a sequence of zero or more Unicode code points, encoded in UTF-8 [STD63].

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, in the plain text form of this document, the operator "^" stands for exponentiation, and, in the HTML and PDF versions, subtraction and negation are rendered as a hyphen ("-", as are various dashes).

## 2. Packed CBOR

This section describes the packing tables, their structure, and how they are referenced.

### 2.1. Packing Tables

At any point within a data item making use of Packed CBOR, there is a Current Set of packing tables that applies.

There are three packing tables in a Current Set:

- \* Shared item table
- \* Prefix table

## \* Suffix table

Without any table setup, all these tables are empty arrays. Table setup can cause these arrays to be non-empty, where the elements are (potentially themselves packed) data items. Each of the tables is indexed by an unsigned integer (starting from 0). Such an index may be derived from information in tags and their content as well as from CBOR simple values.

## 2.2. Referencing Shared Items

Shared items are stored in the shared item table of the Current Set.

The shared data items are referenced by using the reference data items in Table 1. When reconstructing the original data item, such a reference is replaced by the referenced data item, which is then recursively unpacked.

| reference                 | table index    |
|---------------------------|----------------|
| Simple value 0-15         | 0-15           |
| Tag 6(unsigned integer N) | $16 + 2*N$     |
| Tag 6(negative integer N) | $16 - 2*N - 1$ |

Table 1: Referencing Shared Values

As examples in CBOR diagnostic notation (Section 8 of [STD94]), the first 22 elements of the shared item table are referenced by `simple(0)`, `simple(1)`, ... `simple(15)`, `6(0)`, `6(-1)`, `6(1)`, `6(-2)`, `6(2)`, `6(-3)`. (The alternation between unsigned and negative integers for even/odd table index values -- "zigzag encoding" -- makes systematic use of shorter integer encodings first.)

Taking into account the encoding of these referring data items, there are 16 one-byte references, 48 two-byte references, 512 three-byte references, 131072 four-byte references, etc. As CBOR integers can grow to very large (or very negative) values, there is no practical limit to how many shared items might be used in a Packed CBOR item.

Note that the semantics of Tag 6 depend on its tag content: An integer turns the tag into a shared item reference, whereas a string or container (map or array) turns it into a prefix reference (see Table 2). Note also that the tag content of Tag 6 may itself be packed, so it may need to be unpacked to make this determination.

### 2.3. Referencing Affix Items

Prefix items are stored in the prefix table of the Current Set; suffix items are stored in the suffix table of the Current Set. We collectively call these items affix items; when referencing, which of the tables is actually used depends on whether a prefix or a suffix reference was used.

| prefix reference                  | table index    |
|-----------------------------------|----------------|
| Tag 6(suffix)                     | 0              |
| Tag 225-255(suffix)               | 1-31           |
| Tag 28704-32767(suffix)           | 32-4095        |
| Tag 1879052288-2147483647(suffix) | 4096-268435455 |

Table 2: Referencing Prefix Values

| suffix reference                  | table index   |
|-----------------------------------|---------------|
| Tag 216-223(prefix)               | 0-7           |
| Tag 27647-28671(prefix)           | 8-1023        |
| Tag 1811940352-1879048191(prefix) | 1024-67108863 |

Table 3: Referencing Suffix Values

Affix data items are referenced by using the data items in Table 2 and Table 3. The tag number indicates the table used (prefix or suffix) and a table index (an unsigned integer); the tag content contains a "rump item". When reconstructing the original data item, such a reference is replaced by a data item constructed from the referenced affix data item (affix, which might need to be recursively unpacked first) "concatenated" with the tag content (rump, again possibly recursively unpacked).

- \* For a rump of type array and map, the affix also needs to be an array or a map. For an array, the elements from the prefix are prepended to the rump array, while the elements from a suffix are appended. For a map, the entries in the affix are added to those of the rump; prefix and suffix references differ in how entries

with identical keys are combined: for prefix references, an entry in the rump with the same key as an entry in the affix overrides the one in the affix, while for suffix references, an entry in the affix overrides an entry in the rump that has the same key.

NOTE: One application of the rule for prefix references is to supply default values out of a dictionary, which can then be overridden by the entries in the map supplied as the rump value. Note that this pattern provides no way to remove a map entry from the prefix table entry.

- \* For a rump of one of the string types, the affix also needs to be one of the string types; the bytes of the strings are concatenated as specified (prefix + rump, rump + suffix). The result of the concatenation gets the type of the rump; this way a single affix can be used to build both byte and text strings, depending on what type of rump is being used.

As a contrived (but short) example, if the prefix table is ["foobar", "foob", "fo"], the following prefix references will all unpack to "foobart": 6("t"), 224("art"), 225("obart") (the last example is not an optimization).

Taking into account the encoding, there is one single-byte prefix reference, 31 ( $2^5-2^0$ ) two-byte references, 4064 ( $2^{12}-2^5$ ) three-byte references, and 26843160 ( $2^{28}-2^{12}$ ) five-byte references for prefixes. 268435455 ( $2^{28}$ ) is an artificial limit, but should be high enough that there, again, is no practical limit to how many prefix items might be used in a Packed CBOR item. The numbers for suffix references are one quarter of those, except that there is no single-byte reference and 8 two-byte references.

Rationale: Experience suggests that prefix packing might be more likely than suffix packing. Also for this reason, there is no intent to spend a 1+0 tag value for suffix packing.

## 2.4. Discussion

This specification uses up a large number of Simple Values and Tags, in particular one of the rare one-byte tags and two thirds of the one-byte simple values. Since the objective is compression, this is warranted only based on a consensus that this specific format could be useful for a wide area of applications, while maintaining reasonable simplicity in particular at the side of the consumer.

A maliciously crafted Packed CBOR data item might contain a reference loop. A consumer/decompressor MUST protect against that.

Different strategies for decoding/consuming Packed CBOR are available.

For example:

- \* the decoder can decode and unpack the packed item, presenting an unpacked data item to the application. In this case, the onus of dealing with loops is on the decoder. (This strategy generally has the highest memory consumption, but also the simplest interface to the application.) Besides avoiding getting stuck in a reference loop, the decoder will need to control its resource allocation, as data items can "blow up" during unpacking.
- \* the decoder can be oblivious of Packed CBOR. In this case, the onus of dealing with loops is on the application, as is the entire onus of dealing with Packed CBOR.
- \* hybrid models are possible, for instance: The decoder builds a data item tree directly from the Packed CBOR as if it were oblivious, but also provides accessors that hide (resolve) the packing. In this specific case, the onus of dealing with loops is on the accessors.

In general, loop detection can be handled in a similar way in which loops of symbolic links are handled in a file system: A system-wide limit (often 31 or 40 indirections for symbolic links) is applied to any reference chase.

NOTE: The present specification does nothing to help with the packing of CBOR sequences [RFC8742]; maybe such a specification should be added.

### 3. Table Setup

The packing references described in Section 2 assume that packing tables have been set up.

By default, all three tables are empty (zero-length arrays).

Table setup can happen in one of two ways:

- \* By the application environment, e.g., a media type. These can define tables that amount to a static dictionary that can be used in a CBOR data item for this application environment. Note that, without this information, a data item that uses such a static dictionary can be decoded at the CBOR level, but not fully

unpacked. The table setup mechanisms provided by this document are defined in such a way that an unpacker can at least recognize if this is the case.

- \* By one or more tags enclosing the packed content. These can be defined to add to the packing tables that already apply to the tag. Usually, the semantics of the tag will be to prepend items to one of the tables. Note that it may be useful to leave a particular efficiency tier alone and only prepend to a higher tier; e.g., a tag could insert shared items at table index 16 and shift anything that was already there further down in the array while leaving index 0 to 15 alone. Explicit additions by tag can combine with application-environment supplied tables that apply to the entire CBOR data item.

For table setup, the present specification only defines a single tag, which operates by prepending to the (by default empty) tables.

We could also define a tag for dictionary referencing (or include that in the basic packed CBOR), but the desirable details are likely to vary considerably between applications. A URI-based reference would be easy to add, but might be too inefficient when used in the likely combination with an ni: URI [RFC6920].

### 3.1. Basic Packed CBOR

A predefined tag for packing table setup is defined in CDDL [RFC8610] as in Figure 1:

```
Basic-Packed-CBOR = #6.51([[*shared-item], [*prefix-item],  
                           [*suffix-item], rump])  
rump = any  
prefix-item = any  
suffix-item = any  
shared-item = any
```

Figure 1: Packed CBOR in CDDL

(This assumes the allocation of tag number 51 for this tag.)

The arrays given as the first, second, and third element of the content of the tag 51 are prepended to the tables for shared items, prefixes, and suffixes that apply to the entire tag (by default empty tables).

The original CBOR data item can be reconstructed by recursively replacing shared, prefix, and suffix references encountered in the rump by their expansions.

Packed item references in the newly constructed (low-numbered) parts of the table need to be interpreted in the number space of that table (which includes the, now higher-numbered, inherited parts), while references in any existing, inherited (higher-numbered) part continue to use the (more limited) number space of the inherited table.

#### 4. IANA Considerations

In the registry "CBOR Tags" [IANA.cbor-tags], IANA is requested to allocate the tags defined in Table 4.

| Tag                   | Data Item   | Semantics                  | Reference              |
|-----------------------|---|----------------------------|------------------------|
| 6                     | integer (for shared); text string, byte string, array, map (for prefix) | Packed CBOR: shared/prefix | draft-ietf-cbor-packed |
| 225-255               | text string, byte string, array, map                                    | Packed CBOR: prefix        | draft-ietf-cbor-packed |
| 28704-32767           | text string, byte string, array, map                                    | Packed CBOR: prefix        | draft-ietf-cbor-packed |
| 1879052288-2147483647 | text string, byte string,   | Packed CBOR: prefix        | draft-ietf-cbor-packed |

|                       |   |                           |                        |
|-----------------------|---|---------------------------|------------------------|
|                       | array,<br>map                                       |                           |                        |
| 216-223               | text<br>string,<br>byte<br>string,<br>array,<br>map | Packed<br>CBOR:<br>suffix | draft-ietf-cbor-packed |
| 27647-28671           | text<br>string,<br>byte<br>string,<br>array,<br>map | Packed<br>CBOR:<br>suffix | draft-ietf-cbor-packed |
| 1811940352-1879048191 | text<br>string,<br>byte<br>string,<br>array,<br>map | Packed<br>CBOR:<br>suffix | draft-ietf-cbor-packed |

Table 4: Values for Tag Numbers

In the registry "CBOR Simple Values" [IANA.cbor-simple-values], IANA is requested to allocate the simple values defined in Table 5.

| Value | Semantics           | Reference              |
|-------|---------------------|------------------------|
| 0-15  | Packed CBOR: shared | draft-ietf-cbor-packed |

Table 5: Simple Values

## 5. Security Considerations

The security considerations of [STD94] apply.

Loops in the Packed CBOR can be used as a denial of service attack, see Section 2.4.

As the unpacking is deterministic, packed forms can be used as signing inputs. (Note that if external dictionaries are added to cbor-packed, this requires additional consideration.)



## 6. References

### 6.1. Normative References

- [IANA.cbor-simple-values]  
IANA, "Concise Binary Object Representation (CBOR) Simple Values",  
<<https://www.iana.org/assignments/cbor-simple-values>>.
- [IANA.cbor-tags]  
IANA, "Concise Binary Object Representation (CBOR) Tags",  
<<https://www.iana.org/assignments/cbor-tags>>.
- [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>>.
- [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>>.
- [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/info/rfc8610>>.
- [STD94] 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/info/rfc8949>>.

### 6.2. Informative References

- [RFC1951] Deutsch, P., "DEFLATE Compressed Data Format Specification version 1.3", RFC 1951, DOI 10.17487/RFC1951, May 1996, <<https://www.rfc-editor.org/info/rfc1951>>.
- [RFC6920] Farrell, S., Kutscher, D., Dannewitz, C., Ohlman, B., Keranen, A., and P. Hallam-Baker, "Naming Things with Hashes", RFC 6920, DOI 10.17487/RFC6920, April 2013, <<https://www.rfc-editor.org/info/rfc6920>>.
- [RFC7049] Bormann, C. and P. Hoffman, "Concise Binary Object Representation (CBOR)", RFC 7049, DOI 10.17487/RFC7049, October 2013, <<https://www.rfc-editor.org/info/rfc7049>>.

- [RFC8742] Bormann, C., "Concise Binary Object Representation (CBOR) Sequences", RFC 8742, DOI 10.17487/RFC8742, February 2020, <<https://www.rfc-editor.org/info/rfc8742>>.
- [STD63] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, DOI 10.17487/RFC3629, November 2003, <<https://www.rfc-editor.org/info/rfc3629>>.

## Appendix A. Examples

The (JSON-compatible) CBOR data structure depicted in Figure 2, 400 bytes of binary CBOR, could lead to a packed CBOR data item depicted in Figure 3, ~309 bytes. Note that this particular example does not lend itself to prefix compression.

```
{ "store": {  
  "book": [  
    { "category": "reference",  
      "author": "Nigel Rees",  
      "title": "Sayings of the Century",  
      "price": 8.95  
    },  
    { "category": "fiction",  
      "author": "Evelyn Waugh",  
      "title": "Sword of Honour",  
      "price": 12.99  
    },  
    { "category": "fiction",  
      "author": "Herman Melville",  
      "title": "Moby Dick",  
      "isbn": "0-553-21311-3",  
      "price": 8.99  
    },  
    { "category": "fiction",  
      "author": "J. R. R. Tolkien",  
      "title": "The Lord of the Rings",  
      "isbn": "0-395-19395-8",  
      "price": 22.99  
    }  
  ],  
  "bicycle": {  
    "color": "red",  
    "price": 19.95  
  }  
}
```

Figure 2: Example original CBOR data item

```

51(["price", "category", "author", "title", "fiction", 8.95, "isbn"],
  / 0      1      2      3      4      5      6 /
  [], [],
  [{"store": {
    "book": [
      {simple(1): "reference", simple(2): "Nigel Rees",
        simple(3): "Sayings of the Century", simple(0): simple(5)},
      {simple(1): simple(4), simple(2): "Evelyn Waugh",
        simple(3): "Sword of Honour", simple(0): 12.99},
      {simple(1): simple(4), simple(2): "Herman Melville",
        simple(3): "Moby Dick", simple(6): "0-553-21311-3",
        simple(0): simple(5)},
      {simple(1): simple(4), simple(2): "J. R. R. Tolkien",
        simple(3): "The Lord of the Rings",
        simple(6): "0-395-19395-8", simple(0): 22.99}],
    "bicycle": {"color": "red", simple(0): 19.95}}}]

```

Figure 3: Example packed CBOR data item

The (JSON-compatible) CBOR data structure below has been packed with shared item and (partial) prefix compression only.

```

{
  "name": "MyLED",
  "interactions": [
    {
      "links": [
        {
          "href":
            "http://192.168.1.103:8445/wot/thing/MyLED/rgbValueRed",
          "mediaType": "application/json"
        }
      ],
      "outputData": {
        "valueType": {
          "type": "number"
        }
      },
      "name": "rgbValueRed",
      "writable": true,
      "@type": [
        "Property"
      ]
    },
    {
      "links": [
        {
          "href":

```

```
        "http://192.168.1.103:8445/wot/thing/MyLED/rgbValueGreen",
        "mediaType": "application/json"
    },
    ],
    "outputData": {
        "valueType": {
            "type": "number"
        }
    },
    "name": "rgbValueGreen",
    "writable": true,
    "@type": [
        "Property"
    ]
},
{
    "links": [
        {
            "href":
                "http://192.168.1.103:8445/wot/thing/MyLED/rgbValueBlue",
            "mediaType": "application/json"
        }
    ],
    "outputData": {
        "valueType": {
            "type": "number"
        }
    },
    "name": "rgbValueBlue",
    "writable": true,
    "@type": [
        "Property"
    ]
},
{
    "links": [
        {
            "href":
                "http://192.168.1.103:8445/wot/thing/MyLED/rgbValueWhite",
            "mediaType": "application/json"
        }
    ],
    "outputData": {
        "valueType": {
            "type": "number"
        }
    },
    "name": "rgbValueWhite",
```

```
    "writable": true,
    "@type": [
      "Property"
    ]
  },
  {
    "links": [
      {
        "href":
          "http://192.168.1.103:8445/wot/thing/MyLED/ledOnOff",
        "mediaType": "application/json"
      }
    ],
    "outputData": {
      "valueType": {
        "type": "boolean"
      }
    },
    "name": "ledOnOff",
    "writable": true,
    "@type": [
      "Property"
    ]
  },
  {
    "links": [
      {
        "href":
          "http://192.168.1.103:8445/wot/thing/MyLED/colorTemperatureChanged",
        "mediaType": "application/json"
      }
    ],
    "outputData": {
      "valueType": {
        "type": "number"
      }
    },
    "name": "colorTemperatureChanged",
    "@type": [
      "Event"
    ]
  }
],
"@type": "Lamp",
"id": "0",
"base": "http://192.168.1.103:8445/wot/thing",
"@context":
  "http://192.168.1.102:8444/wot/w3c-wot-td-context.jsonld"
```

```
}

```

Figure 4: Example original CBOR data item

```
51([/shared/["name", "@type", "links", "href", "mediaType",
/ 0 1 2 3 4 /
"application/json", "outputData", {"valueType": {"type":
/ 5 6 7 /
"number"}}, ["Property"], "writable", "valueType", "type"],
/ 8 9 10 11 /
/prefix/ ["http://192.168.1.10", 6("3:8445/wot/thing"),
/ 6 225 /
225("/MyLED/"), 226("rgbValue"), "rgbValue",
/ 226 227 228 /
{simple(6): simple(7), simple(9): true, simple(1): simple(8)}],
/ 229 /
/suffix/ [],
/rump/ {simple(0): "MyLED",
"interactions": [
229({simple(2): [{simple(3): 227("Red"), simple(4): simple(5)}],
simple(0): 228("Red")}),
229({simple(2): [{simple(3): 227("Green"), simple(4): simple(5)}],
simple(0): 228("Green")}),
229({simple(2): [{simple(3): 227("Blue"), simple(4): simple(5)}],
simple(0): 228("Blue")}),
229({simple(2): [{simple(3): 227("White"), simple(4): simple(5)}],
simple(0): "rgbValueWhite"}),
{simple(2): [{simple(3): 226("ledOnOff"), simple(4): simple(5)}],
simple(6): {simple(10): {simple(11): "boolean"}}, simple(0):
"ledOnOff", simple(9): true, simple(1): simple(8)},
{simple(2): [{simple(3): 226("colorTemperatureChanged"),
simple(4): simple(5)}], simple(6): simple(7), simple(0):
"colorTemperatureChanged", simple(1): ["Event"]},
simple(1): "Lamp", "id": "0", "base": 225(""),
"@context": 6("2:8444/wot/w3c-wot-td-context.jsonld")})])

```

Figure 5: Example packed CBOR data item

## Acknowledgements

CBOR packing was originally invented with the rest of CBOR, but did not make it into [RFC7049], the predecessor of [STD94]. Various attempts to come up with a specification over the years didn't proceed. In 2017, Sebastian Käbisch proposed investigating compact representations of W3C Thing Descriptions, which prompted the author to come up with essentially the present design.

Author's Address

Carsten Bormann  
Universität Bremen TZI  
Postfach 330440  
D-28359 Bremen  
Germany  
Phone: +49-421-218-63921  
Email: cabo@tzi.org

Network Working Group  
Internet-Draft  
Intended status: Informational  
Expires: 20 November 2021

C. Bormann  
Universität Bremen TZI  
B. Gamari  
Well-Typed  
H. Birkholz  
Fraunhofer SIT  
19 May 2021

Concise Binary Object Representation (CBOR) Tags for Time, Duration, and  
Period  
draft-ietf-cbor-time-tag-00

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.

Note to Readers

Version -00 of the individual submission that led to the present draft opened up the possibilities provided by extended representations of time in CBOR. Version -01 consolidated this draft to non-speculative content, the normative parts of which were believed will stay unchanged during further development of the draft. This version was provided to aid the registration of the CBOR tag immediately needed. Further versions of the individual submission made use of the IANA allocations registered and made other editorial updates. Now a WG document, future versions could re-introduce some of the material from the initial submission, but in a more concrete form.

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 20 November 2021.

#### Copyright Notice

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

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<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. Terminology . . . . .                        | 3  |
| 2. Objectives . . . . .                           | 3  |
| 3. Time Format . . . . .                          | 4  |
| 3.1. Key 1 . . . . .                              | 5  |
| 3.2. Keys 4 and 5 . . . . .                       | 5  |
| 3.3. Keys -3, -6, -9, -12, -15, -18 . . . . .     | 5  |
| 3.4. Key -1: Time Scale . . . . .                 | 5  |
| 3.5. Clock Quality . . . . .                      | 6  |
| 3.5.1. ClockClass (Key -2) . . . . .              | 6  |
| 3.5.2. ClockAccuracy (Key -4) . . . . .           | 7  |
| 3.5.3. OffsetScaledLogVariance (Key -5) . . . . . | 7  |
| 3.5.4. Uncertainty (Key -7) . . . . .             | 7  |
| 3.5.5. Guarantee (Key -8) . . . . .               | 7  |
| 4. Duration Format . . . . .                      | 8  |
| 5. Period Format . . . . .                        | 8  |
| 6. CDDL typenames . . . . .                       | 9  |
| 7. IANA Considerations . . . . .                  | 9  |
| 8. Security Considerations . . . . .              | 9  |
| 9. References . . . . .                           | 10 |

|                                       |    |
|---------------------------------------|----|
| 9.1. Normative References . . . . .   | 10 |
| 9.2. Informative References . . . . . | 11 |
| Acknowledgements . . . . .            | 11 |
| Authors' Addresses . . . . .          | 11 |

## 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.

(TBD: Expand on text from abstract here.)

### 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.

## 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 time scale. Tags 0 and 1 are for UTC; however, some interchanges are better performed on TAI. Other time scales 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).

Not currently addressed, but possibly covered by the definition of additional map keys for the map inside the tag:

- \* 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. (TBD: 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 with the extended time format.)
- \* Additional indication of intents about the interpretation of the time given, in particular for future times. Intents might include information about time zones, daylight savings times, etc.

Additional tags are defined for durations and periods.

### 3. Time Format

An extended time is indicated by CBOR tag 1001, which tags 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. Implementations MUST ignore key/value types they do not understand for negative integer and text string values of the key. Not understanding key/value for unsigned keys is an error.

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

- \* a higher precision time offset to be added to the base time,
- \* a reference time scale and epoch different from the default UTC and 1970-01-01
- \* information about clock quality parameters, such as source, accuracy, and uncertainty

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.

While this document does not define supplementary text keys, a number of unsigned and negative-integer keys are defined below.

### 3.1. Key 1

Key 1 indicates a 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.

### 3.2. Keys 4 and 5

Keys 4 and 5 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.

### 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       |
| -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

### 3.4. Key -1: Time Scale

Key -1 is used to indicate a time scale. 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].

If key -1 is not present, time scale value 0 is implied. Additional values can be registered in the (TBD define name for time scale registry); values MUST be integers or text strings.

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

```
t      = t      - 2208988800
utc    ntp

t      = t      + 315964819
tai    gps
```

Figure 1: Converting Common Offset Time Scales

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

#### 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 approx48 + |_2cdotlog {accovers} - epsilon_|
      acc              10
```

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.

#### 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 time scale given (e.g., leap seconds).

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

```
| (TBD: Clearly, ISO8601 durations are rather different; we do
| not want to use these.)
```

#### 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:

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

| (Issue: should start/end be given the two-element treatment, or start/duration?)

## 6. CDDL typenames

For the use with the CBOR Data Definition Language, CDDL [RFC8610], 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

In the registry [IANA.cbor-tags], IANA has allocated the tags in Table 2 from 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".

| (TBD: Add registry for time scales. Add registry for map keys and allocation policies for additional keys.)

## 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

- [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", <<http://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>>.
- [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>>.
- [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>>.
- [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/info/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/info/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](http://pubs.opengroup.org/onlinepubs/9699919799/basedefs/V1_chap04.html#tag_04_16)>.

## 9.2. Informative References

- [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/info/rfc3161>>.
- [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/info/rfc8575>>.

## Acknowledgements

### Authors' Addresses

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

Phone: +49-421-218-63921  
Email: [cabo@tzi.org](mailto:cabo@tzi.org)

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

Email: [ben@well-typed.com](mailto:ben@well-typed.com)

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

Email: [henk.birkholz@sit.fraunhofer.de](mailto:henk.birkholz@sit.fraunhofer.de)