

Workgroup: cbor
Internet-Draft:
draft-amsuess-cbor-packed-by-reference-01
Published: 11 July 2023
Intended Status: Standards Track
Expires: 12 January 2024
Authors: C. Amsüss

Packed CBOR: Table set up by reference

Abstract

Packed CBOR is a compression mechanism for Concise Binary Object Representation (CBOR) that can be used without a decompression step. This document introduces a means for setting up its tables by means of dereferencable identifiers, and introduces a pattern of using it without sending long identifiers.

Status of This Memo

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

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

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

This Internet-Draft will expire on 12 January 2024.

Copyright Notice

Copyright (c) 2023 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. Setting up the tables by reference](#)
 - [2.1. Example](#)
- [3. Nested table setups](#)
- [4. Security Considerations](#)
- [5. IANA Considerations](#)
- [6. Normative References](#)
- [Acknowledgments](#)
- [Author's Address](#)

1. Introduction

[See abstract.]

2. Setting up the tables by reference

CBOR tag TBD114 is defined with semantics similar to tags TBD113 and TBD113 from [[I-D.ietf-cbor-packed](#)] in that it sets up tables around a rump.

```
Packed-By-Reference = #6.114([count, source, rump])
rump = any
source = CRI / URI
count = (count-shared-and-argument //
  count-shared, count-argument //
  skip-shared, count-shared, skip-argument, count-argument)
count-shared-and-argument = int
count-shared = int
count-argument = int
skip-shared = int
skip-argument = int
```

The items inserted by the tables are not given explicitly, but picked out of tables known by their identifier given as source. Such a source needs to represent two lists of CBOR items, one for each kind of tables. The tag prepends some number of items out of those source lists to the tables that are used to decompress the rump.

The identifier is given as a URI string (as defined in [[RFC3986](#)] or equivalently as a CRI (as defined in [[I-D.ietf-core-href](#)]). Later iterations of this document may introduce additional options. When the source identifier is dereferencable, all considerations from [[I-D.bormann-t2trg-deref-id](#)] apply. (Simplifying: No dereferencing at runtime -- the recipient either knows it already or treats it as unknown).

If the same number of items is prepended to both tables, their count is given as a single number; otherwise, the numbers are given

separately. If a large number of items at the beginning of the source tables would not be used, the four-argument form of count defines a number of items in the source tables that are skipped before selecting items into the table. This allows keeping the indices low and therefore compact.

Source tables should be designed in such a way that commonly used items are at the start to avoid the necessity of the four-argument form.

There is no short-hand for importing all the items, as the use of such a form would rule out any extensibility in source tables.

2.1. Example

Suppose the URI "tag:example.com,2023:byref" defines the items ["price", "category", "author", "title", "fiction", 8.95, "isbn"] in both tables. Then the example in figure 3 of [[I-D.ietf-cbor-packed](#)] can be written as:

```
114([7, "tag:example.com,2023:books"
  [{"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}}}]])
```

3. Nested table setups

Assembling tables from one or more source references can easily become verbose. A pattern that reduces the verbosity while staying unambiguous are nested table setups, where the outer tables are extended to contain additional identifiers.

In this pattern, tables are set up in two stages:

The outer stage contains the CRIs or URIs that may later be used as source values. (It may also contain other items). Those tables may be set up in a way that grows over time, so that more identifiers can be added later. The inner stage is set up using tag TBD114, and the source given is a packed reference.

In this example, the initial table set up is provided by the media type, and contains these items:

*0: "This class has students with the following names"

*100: "tag:example.com,2023:english-names.txt"

*101: "tag:example.com,2023:german-names.txt"

```
114([5, 6(100) / outer item 100 /, 114([2, 6(106) / outer item 101 /, [
  6(11) / outer item 0, "This class has students with the following name
  6(0) / item 0 of german-names, "Franz" /,
  6(2) / item 0 of english-names, "George" /,
  6(1) / item 1 of german-names, "Fritz" /,
  6(7) / item 5 of english-names, "Jack" /
]]])
```

When either the outer list gets extended (to add more source identifiers) or any of the referenced source lists gets extended, decoders can treat values as undefined if their protocol allows such behavior.

Note that a constrained implementation of a decoder may not even have the fully expanded form of the URIs or CRIs available; it may only be capable of using these table entries in the source position and then find the shipped source lists.

4. Security Considerations

[TBD]

5. IANA Considerations

[TBD: Request tag.]

6. Normative References

[**I-D.bormann-t2trg-deref-id**] Bormann, C. and C. Amsüss, "The "dereferenceable identifier" pattern", Work in Progress, Internet-Draft, draft-bormann-t2trg-deref-id-01, 9 May 2023, <<https://datatracker.ietf.org/doc/html/draft-bormann-t2trg-deref-id-01>>.

[**I-D.ietf-cbor-packed**]

Bormann, C., "Packed CBOR", Work in Progress, Internet-Draft, draft-ietf-cbor-packed-09, 10 July 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-cbor-packed-09>>.

[**I-D.ietf-core-href**] Bormann, C. and H. Birkholz, "Constrained Resource Identifiers", Work in Progress, Internet-Draft,

draft-ietf-core-href-13, 10 July 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-core-href-13>>.

[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/rfc/rfc3986>>.

Acknowledgments

[TBD]

Author's Address

Christian Amsüss
Austria

Email: christian@amsuess.com