

Workgroup: CBOR Working Group
Internet-Draft: draft-ietf-cbor-file-magic-10
Published: 7 March 2022
Intended Status: Standards Track
Expires: 8 September 2022
Authors: M. Richardson C. Bormann
 Sandelman Software Works Universität Bremen TZI
 On storing CBOR encoded items on stable storage

Abstract

This document defines an on-disk format for CBOR objects that is friendly to common on-disk 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 8 September 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](#)
 - [1.1. Terminology](#)
 - [1.2. Requirements for a Magic Number](#)
- [2. Protocol](#)
 - [2.1. The CBOR Protocol Specific Tag](#)
 - [2.2. Enveloping Method: CBOR Tag Wrapped](#)
 - [2.2.1. Example](#)
 - [2.3. Enveloping Method: Labeled CBOR Sequence](#)
- [3. Advice to Protocol Developers](#)
 - [3.1. Is the on-wire format new?](#)
 - [3.2. Can many items be trivially concatenated?](#)
 - [3.3. Are there tags at the start?](#)
- [4. Security Considerations](#)
- [5. IANA Considerations](#)
 - [5.1. Labeled CBOR Sequence Tag](#)
 - [5.2. CBOR-Labeled Non-CBOR Data Tag](#)
 - [5.3. CBOR Tags for CoAP Content-Format Numbers](#)
- [6. References](#)
 - [6.1. Normative References](#)
 - [6.2. Informative References](#)
- [Appendix A. CBOR Tags for CoAP Content Formats](#)
 - [A.1. Content-Format Tag Examples](#)
- [Appendix B. Example from Openswan](#)
- [Appendix C. Using CBOR Labels for non-CBOR data](#)
 - [C.1. Content-Format Tag Examples](#)
- [Appendix D. Changelog](#)
- [Acknowledgements](#)
- [Contributors](#)
- [Authors' Addresses](#)

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 (related to the TCP/IP socket model, where all network connectivity also 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 that includes media type ("MIME type") information and therefore ideally never needs to know anything about the file.

But, many other systems do this by file extensions. Many common web servers derive the MIME-type information from file extensions.

Having a media type associated with the file is a better solution in general. When files become disconnected from their type information, such as when attempting to do forensics on a damaged system, then being able to identify a file type can become very important.

It is noted that in the media type registration, that a magic number is asked for, if available, as is a file extension.

A challenge for the `file(1)` program is often that it can be confused by the encoding vs the content. For instance, an Android "apk" 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. Two possible methods of enveloping data are presented: a CBOR Protocol author will specify one. (A CBOR Protocol is a specification which uses CBOR as its encoding.)

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 an appendix because it is not about identifying media types containing CBOR-encoded data items.

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

Examples of CBOR Protocols currently under development include CoSWID [[I-D.ietf-sacm-coswid](#)] and EAT [[I-D.ietf-rats-eat](#)]. COSE itself [[RFC8152](#)] is considered infrastructure, however the encoding of public keys in CBOR as described in [[I-D.ietf-cose-cbor-encoded-cert](#)] would be an identified CBOR Protocol as well.

A major inspiration for this document is observing the mess 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.

In addition to the on-disk identification aspects, 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 signals that CBOR is being used as opposed to a legacy scheme.

In addition, for convenience, [Appendix C](#) defines a simple way to retroactively add a magic number to content-formats as defined by [[RFC7252](#)], even if not in CBOR form.

1.1. Terminology

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

The term "diagnostic notation" refers to the human-readable notation for CBOR data items defined in [Section 8](#) of [[RFC8949](#)] and [Appendix G](#) of [[RFC8610](#)].

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

1.2. Requirements for a Magic Number

A magic number is ideally a fingerprint that is unique to a 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. Which one is to be used is up to the CBOR Protocol author to determine. Both use CBOR Tags in a way that results in a deterministic first 8 to 12 bytes.

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 on disk. 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 an email to IANA, having filled in the small template provided in [Section 9.2](#) of [\[RFC8949\]](#).

This tag needs to be allocated by the author of the CBOR Protocol. In order to be in the four-byte range, and so that there are no leading zeros, the value needs to be in the range 0x01000000 (decimal 16777216) to 0xFFFFFFFF (decimal 4294967295). It is further suggested to avoid values that have an embedded zero byte in the four bytes of their binary representation (such as 0x12003456).

The use of a sequence of four US-ASCII codes which are mnemonic to the protocol is encouraged, but not required.

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 already been allocated in [Section 5.3](#) (see [Appendix A](#) 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 [\[RFC8949\]](#).

The tag content of that tag is a second CBOR Tag that has been allocated to describe the specific Protocol involved, as described above.

This method wraps the CBOR value as tags usually do. Applications that need to send the CBOR value across a constrained link may wish to remove the two tags if the use is implicitly understood.

Whether these two tags should be removed for specific further processing is a decision made by the CBOR Protocol specification.

2.2.1. Example

To construct an example without registering a new tag, we use the technique described in [Appendix A](#) to translate the Content-Format number registered for application/senml+cbor, the number 112, into the tag $1668546560+112 = 1668546672$.

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

```
55799(1668546672([0: "current", 6: 3, 2: 1.5]))
```

Or in hex:

d9 d9f7	# tag(55799)
da 63740070	# tag(1668546672)
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 d9d9f7da63740070 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 [\[RFC8949\]](#). 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.

3. Advice to Protocol Developers

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 or CBOR format, a recorded sensor reading stored for later transmission, or a COVID 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. Programs which might not expect the Labeled CBOR Sequence, but which would operate without a problem would include any program that expects to process CBOR Sequences from the file.

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 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 (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 on-disk 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 disk, 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 A](#) for media-types with CoAP Content-Format numbers.

Here are some considerations in the form of survey questions:

3.1. 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 disk 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 on-disk storage.

3.2. 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.

3.3. Are there tags at the start?

If the Protocol expects to use other tags values at the top-level, then it may be easier to explain if the Labeled CBOR Sequence format is used.

4. Security Considerations

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

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

5. IANA Considerations

[Section 5.1](#) 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 B](#)). [Section 5.3](#) allocates a CBOR tag for each actual or potential CoAP Content-Format number (examples are in [Appendix A](#)).

5.1. Labeled CBOR Sequence Tag

IANA has allocated tag 55800 as the tag for the Labeled CBOR Sequence Enveloping Method.

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 [[RFC8949](#)]).

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 [\[RFC8949\]](#) 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

5.2. 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 C](#)).

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 [\[RFC8949\]](#)).

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 [\[RFC8949\]](#) 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.

5.3. CBOR Tags for CoAP Content-Format Numbers

IANA is requested to allocate the tag numbers 1668546560 (0x63740000) to 1668612095 (0x6374FFFF) as follows:

Data Item:

byte string

Semantics:

for each tag number NNNNNNNN, the representation of content-format (RFC7252) NNNNNNNN-1668546560

Reference:

RFCthis

6. References

6.1. Normative References

[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, December 2020. <<https://www.rfc-editor.org/info/std94>>

6.2. Informative References

[file] Wikipedia, "file (command)", 20 January 2021, <https://en.wikipedia.org/wiki/File_%28command%29>.

[I-D.ietf-core-new-block] Boucadair, M. and J. Shallow, "Constrained Application Protocol (CoAP) Block-Wise Transfer Options Supporting Robust Transmission", Work in Progress, Internet-Draft, draft-ietf-core-new-block-14, 26 May 2021, <<https://www.ietf.org/archive/id/draft-ietf-core-new-block-14.txt>>.

[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.core-parameters] IANA, "Constrained RESTful Environments (CoRE) Parameters", <<https://www.iana.org/assignments/core-parameters>>.

[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>>.

[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>>.

[RFC8152] Schaad, J., "CBOR Object Signing and Encryption (COSE)", RFC 8152, DOI 10.17487/RFC8152, July 2017, <<https://www.rfc-editor.org/info/rfc8152>>.

[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>>.

Appendix A. 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 on disk, in firmware, or when debugging.

This specification allocates CBOR tag numbers 1668546560 (0x63740000) to 1668612095 (0x6374FFFF) for the tagging of representations of specific content formats. The tag content tagged with tag number NNNNNNNN (in above range) is a byte string that is to be interpreted as a representation of the content format NNNNNNNN-1668546560.

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 C](#)): the byte string 'BOR', signifying that the representation of the given content-format follows in the file, in the way defined for these tags.

A.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 1668546672 (i.e., 1668546560+112).

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

```
senml-cbor = #6.1668546672(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 [[I-D.ietf-core-new-block](#)].

The corresponding tag number is 1668546832 (i.e., 1668546560+272).

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 = #6.1668546832(bstr)
```

Appendix B. 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 C. 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 A](#)).

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 [\[RFC8949\]](#). 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.

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 to indicate the nature of the header. This value is also included simply because the two tags need to tag something.

C.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 1668546992 (i.e., 1668546560+432).

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(#6.1668546992('BOR'))
```

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

The corresponding tag number is 1668557610 (i.e., 1668546560+11050).

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(#6.1668557610('BOR'))
```

Appendix D. 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](tel:+49-421-218-63921)

Email: cabo@tzi.org