Workgroup: CBOR Working Group
Internet-Draft:
draft-ietf-cbor-network-addresses-10
Published: 6 October 2021
Intended Status: Standards Track
Expires: 9 April 2022
Authors: M. Richardson
Sandelman Software Works
C. Bormann
Universität Bremen TZI

CBOR tags for IPv4 and IPv6 addresses and prefixes

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 9 April 2022.

Table of Contents

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 6- or 8-byte Ethernet addresses.

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 interface identifier Section 4 of [RFC3542]. This may be an integer, in which case it is to be interpreted as the interface index. This may be a string, in which case it is to be interpreted as an interface name.

In the cases where the Interface Format is being used to represent only an address with an interface identifier, and no interface prefix information, then the prefix length may be replaced with the CBOR "false" (0xF4).

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'20010db81234deedbeefcafefacefeed')
```

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. Trailing zero bytes MUST be omitted.

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.

For example:

```
54([h'20010db81234deedbeefcafefacefeed', 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'fe8000000000020202fffffffe030303', 64, 'eth0'])
```

with a numeric interface identifier:

```
54([h'fe8000000000020202fffffffe030303', 64, 42])
```

An IPv6 link-local address without a prefix length:

```
54([h'fe8000000000020202fffffffe030303', false, 42])
```

Interface identifiers may be used with any kind of IPv6 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. Trailing zero bytes MUST be omitted.

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.

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. Encoder Considerations for Prefixes

For the byte strings used in representing prefixes, an encoder MUST 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'])
```

An encoder MUST take care to set all trailing bits in the final byte to zero, if any. While decoders are expected to ignore them, such garbage entities could be used as a covert channel, or may reveal the state of what would otherwise be private memory contents. So for example, 2001:db8:1230::/44 MUST be encoded as:

```
52([44, h'20010db81230'])
```

even though variations like:

```
54([44, h'20010db81233'])
54([45, h'20010db8123f'])
```

would be parsed in the exact same way; they MUST be considered invalid.

The same considerations apply to IPv4 prefixes.

5. Decoder Considerations for Prefixes

A decoder MUST consider all bits to the right of the prefix length to be zero.

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 a pathological case, ::/128 can be encoded as

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

A recommendation for 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 array.

Finally, looking at the last three bits of the prefix-length in bits (that is, the prefix-length modulo 8), use a static array of 8 values to force the lower, non-relevant bits to zero, or simply:

```
unused_bits = (8 - (prefix_length_in_bits & 7)) % 8;
if (length_in_bytes > 0)
  address_bytes[length_in_bytes - 1] &= (0xFF << unused_bits);
```

A particularly paranoid decoder could examine the lower non-relevant bits to determine if they are non-zero, and reject the prefix. This would detect non-compliant encoders, or a possible covert channel.

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

6. 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-value,
                            ?ipv6-interface-identifier]
ipv4-address-with-prefix = [ipv4-address, ipv4-prefix-length]

ipv6-prefix-value  = ipv6-prefix-length
                   / false
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)

ipv6-interface-identifier = uint / tstr
```

Figure 1

7. 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 this 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.

The right-hand bits of the prefix, after the prefix-length, are ignored by this protocol. 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 would be detected by examination of the raw protocol bytes. Users of this encoding should be aware of this possibility.

There are many ways in which the encodings may be invalid: wrong byte lengths (too long, too short), or invalid prefix lengths (greater than 32 for IPv4, greater than 128 for IPv6, negative values, etc.) These are all invalid and this error needs to be signaled to the application, and the entire content thrown away.

8. IANA Considerations

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

8.1. Tag 54 - IPv6

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

8.2. Tag 52 - IPv4

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

8.3. 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

9. References

9.1. 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>.

9.2. Informative References

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

[RFC3542] Stevens, W., Thomas, M., Nordmark, E., and T. Jinmei, "Advanced Sockets Application Program Interface (API) for IPv6", RFC 3542, DOI 10.17487/RFC3542, May 2003, <https://www.rfc-editor.org/info/rfc3542>.

[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 Eric Vyncke reviewed the document and provided suggested text.

Authors' Addresses

Michael Richardson
Sandelman Software Works

Email: mcr+ietf@sandelman.ca

Carsten Bormann
Universität Bremen TZI
Germany

Email: cabo@tzi.org