

# A Concise Binary Object Representation (CBOR) of DNS Messages

draft-lenders-dns-cbor

(<https://datatracker.ietf.org/doc/draft-lenders-dns-cbor/>)

---

**Martine S. Lenders** (martine-lenders@tu-dresden.de),  
Carsten Bormann, Thomas C. Schmidt, Matthias Wählisch

IETF CBOR WG Interim Meeting, 2023-10-04

# Outline

Motivation

Objectives and Definition

Progress

- EDNS OPT Pseudo-RRs

Preliminary Evaluation


Name Compression Ideas

Next Steps

# Motivation: DNS in Constrained Networks

Packet size in DoC exceeds 802.15.4 PDU depending on queried name length

⇒ Fragmentation



DNS over CoAP (`draft-ietf-core-dns-over-coap`) messages for different name lengths

# Motivation: DNS in Constrained Networks

Packet size in DoC exceeds 802.15.4 PDU depending on queried name length

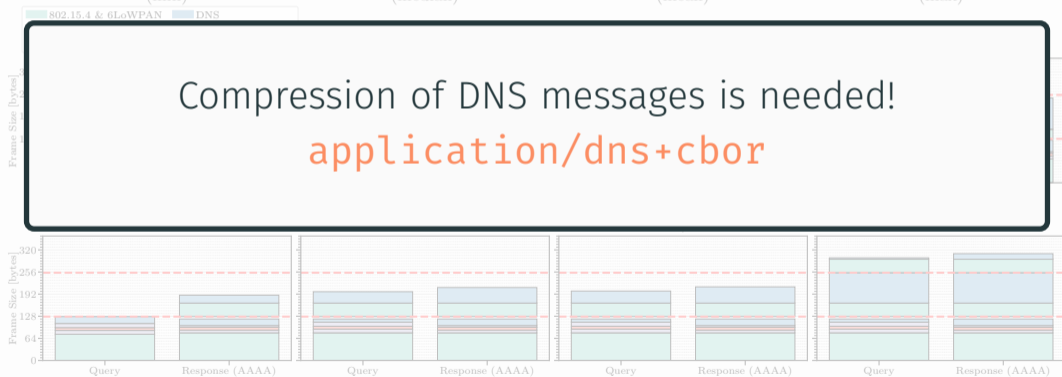
⇒ Fragmentation

Name length = 2 chars  
(min)

Name length = 24 chars  
(median)

Name length = 25.9 chars  
(mean)

Name length = 83 chars  
(max)



DNS over CoAP (draft-ietf-core-dns-over-coap) messages for different name lengths

# Objectives of draft-lenders-dns-cbor (application/dns+cbor)

Reduce packet sizes of DNS queries and replies with conciseness and compression:

1. Encoding of DNS messages in CBOR (conciseness)
2. Omit (redundant) DNS fields in DNS queries and responses (conciseness)
3. Address and name compression using packed CBOR (compression, optional)

## Changes to DNS+CBOR Draft Since IETF 117 and -03

- + Provide format description for EDNS OPT Pseudo-RRs
  - Simplify CDDL to more idiomatic style
  - Other clean-up and housekeeping
- Remove DNS transaction IDs

# EDNS OPT Pseudo-RRs

```
opt-rr = [  
  ? udp-payload-size: uint .default 512,  
  options: [* opt],  
  ? opt-rcode-v-flags,  
]  
opt = (  
  ocode: uint,  
  odata: bstr,  
)  
opt-rcode-v-flags = (  
  flags: uint .default 0,  
  ? opt-rcode-v,  
)  
opt-rcode-v = (  
  rcode: uint .default 0,  
  ? version: uint .default 0,  
)
```

For example wire-format of 23 bytes

```
00 00 29 02 00 00 00 00  
00 00 0c 00 0a 00 08 96  
0c db 3b 79 af 95 26
```

becomes 14 bytes

```
141(                                     |d8 8d  
  [                                       | 81  
    [                                       | 82  
      10,                                 | 0a  
                                           | 48  
      h'960cdb3b79af9526'|          96 0c db 3b  
    ]                                       |          79 af 95 26  
  )                                       |
```

# Evaluation: Data Corpus

## YourThings<sup>1</sup>

## IoTFinder<sup>2</sup>

## MonIoTr<sup>3</sup>

- Collected throughout 2019
- 90 consumer IoT devices from 50 vendors
  - + Phones, Tablets, PCs, ... (in IoTFinder & YourThings)
- 1.20 million queries
- 2.74 million responses
  - 2.07 million w/o mapped query (IoTFinder contains only responses)
  - 0.66 million w/ mapped query

---

<sup>1</sup>O. Alrawi, C. Lever, M. Antonakakis, and F. Monrose. 2019. **SoK: Security Evaluation of Home-Based IoT Deployments**. In *IEEE S&P 2019*. 1362–1380.

<sup>2</sup>R. Perdisci, T. Papastergiou, O. Alrawi, and M. Antonakakis. 2020. **IoTFinder: Efficient Large-Scale Identification of IoT Devices via Passive DNS Traffic Analysis**. In *IEEE EuroS&P 2020*. 474–489.

<sup>3</sup>J. Ren, D.J. Dubois, D. Choffnes, A.M. Mandalari, R. Kolcun, and H. Haddadi. 2019. **Information Exposure for Consumer IoT Devices: A Multidimensional, Network-Informed Measurement Approach**. In *Proc. of the Internet Measurement Conference (IMC)*. ACM.



- Implementation: <https://github.com/netd-tud/cbor4dns>
- Applied compression to data corpus
  - Elide question section in response if query is present in data corpus
  - For all: unpacked and packed
    - `unpacked` “classic” CBOR+DNS (`application/cbor+dns`)
    - `packed` packed CBOR+DNS (`application/cbor+dns;packed=1`)

# Compression Ratios & Byte Savings

Compression ratio

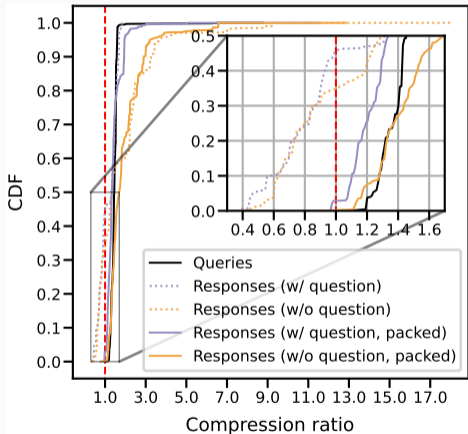
$$\frac{\text{len(wire format)}}{\text{len(CBOR format)}}$$

Byte savings


$$\text{len(wire format)} - \text{len(CBOR format)}$$

# Compression Ratios & Byte Savings

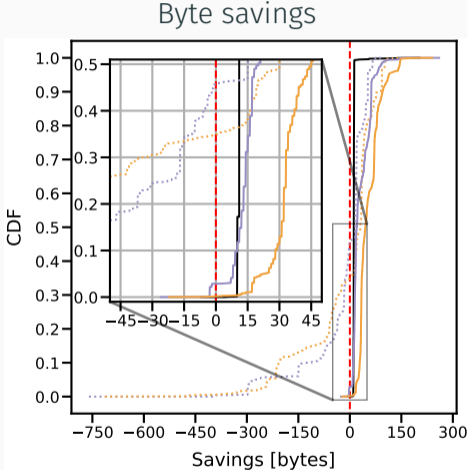

## Compression ratio



## Byte savings


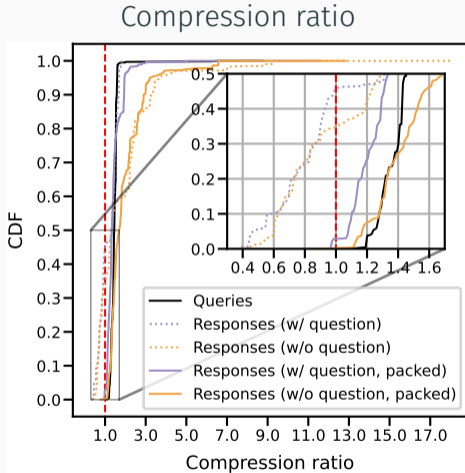


# Compression Ratios & Byte Savings

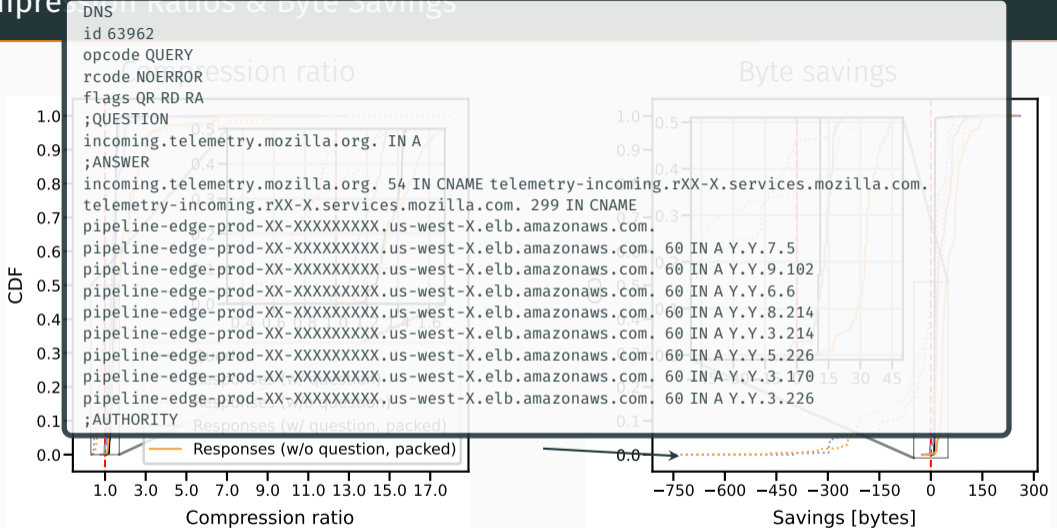




# Compression Ratios & Byte Savings



# Compression Ratios & Byte Savings

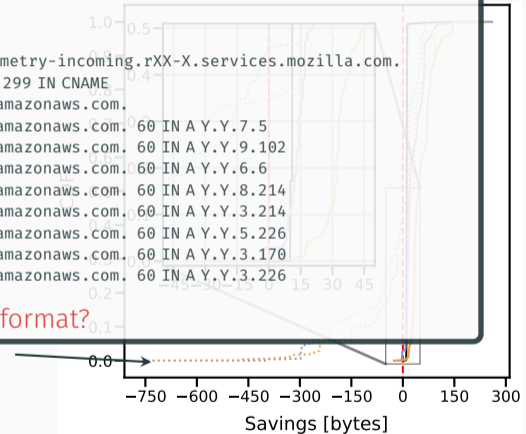
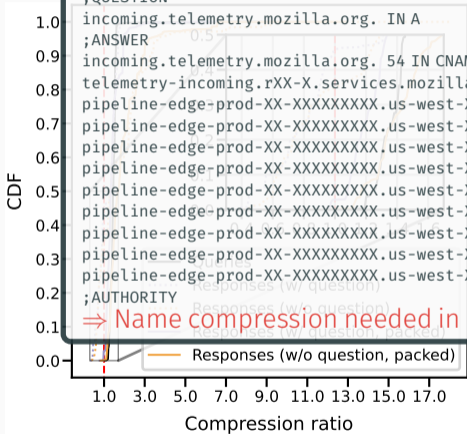


# Compression Ratios & Byte Savings

DNS  
id 63962  
opcode QUERY  
rcode NOERROR  
flags QR RD RA  
;QUESTION

```
incoming.telemetry.mozilla.org. IN A  
;ANSWER  
incoming.telemetry.mozilla.org. 54 IN CNAME telemetry-incoming.rXX-X.services.mozilla.com.  
telemetry-incoming.rXX-X.services.mozilla.com. 299 IN CNAME  
pipeline-edge-prod-XX-XXXXXXXXX.us-west-X.elb.amazonaws.com.  
pipeline-edge-prod-XX-XXXXXXXXX.us-west-X.elb.amazonaws.com. 60 IN A Y.Y.7.5  
pipeline-edge-prod-XX-XXXXXXXXX.us-west-X.elb.amazonaws.com. 60 IN A Y.Y.9.102  
pipeline-edge-prod-XX-XXXXXXXXX.us-west-X.elb.amazonaws.com. 60 IN A Y.Y.6.6  
pipeline-edge-prod-XX-XXXXXXXXX.us-west-X.elb.amazonaws.com. 60 IN A Y.Y.8.214  
pipeline-edge-prod-XX-XXXXXXXXX.us-west-X.elb.amazonaws.com. 60 IN A Y.Y.3.214  
pipeline-edge-prod-XX-XXXXXXXXX.us-west-X.elb.amazonaws.com. 60 IN A Y.Y.5.226  
pipeline-edge-prod-XX-XXXXXXXXX.us-west-X.elb.amazonaws.com. 60 IN A Y.Y.3.170  
pipeline-edge-prod-XX-XXXXXXXXX.us-west-X.elb.amazonaws.com. 60 IN A Y.Y.3.226
```

⇒ Name compression needed in base format?





- Two ideas based on name components
  1. DNS-wire-format-style: Reference components within CBOR object, *less idiomatic*
  2. CBOR-packed „lite“: Reference components in pre-defined table, *more idiomatic*

## Name Compression Idea 1: Reference within CBOR object

- Loosely based on Christian's idea<sup>1</sup>:
  - Current draft: domain-name = tstr .regex "([^.]+[.])\*[^.]+"
  - Proposal: domain-name = (\*text) ; lowercase, NFC labels (cf. core-href)
- Use componified name to also allow for uint in name, *i.e.*, reference

```
[ ["example", "org", 12, 1],  
  [[3600, 12, "_coap", "_udp", "local"]],  
  [[3600, 2, "ns1", i], [3600, 2, "ns2", i]],  
  [ [  
    j, 3600,  
    h'20010db8000000000000000000000001'  
  ], [  
    k, 3600,  
    h'20010db8000000000000000000000002'  
  ], [  
    l, 3600,  
    h'20010db800000000000000000000003535'  
  ] ] ]
```

---

<sup>1</sup><https://mailarchive.ietf.org/arch/msg/cbor/JOHCCB0zC46PrSq-61MMeV--mpU/>

## Name Compression Idea 1: Reference within CBOR object

- Loosely based on Christian's idea<sup>1</sup>:
  - Current draft: domain-name = tstr .regex "<sup>^</sup>.<sup>+</sup>[<sup>^</sup>.]<sup>+</sup>[<sup>^</sup>.]<sup>+</sup>"
  - Proposal: domain-name = (\*text) ; lowercase, NFC labels (cf. core-href)
- Use componified name to also allow for uint in name, i.e., reference

```
[ ["example", "org", 12, 1],  
  [[3600, 12, "_coap", "_udp", "local"]],  
  [[3600, 2, "ns1", i], [3600, 2, "ns2", i]],  
  [ [  
    j, 3600,  
    h'20010db80000000000000000000001'  
  ], [  
    k, 3600,  
    h'20010db80000000000000000000002'  
  ], [  
    l, 3600,  
    h'20010db8000000000000000000003535'  
  ] ] ]
```

The diagram illustrates the name compression mechanism. It shows a nested CBOR array structure. The first level is an array with four elements. The second element is an array of labels: `[[3600, 12, "_coap", "_udp", "local"]]`. The third element is an array of pairs: `[[3600, 2, "ns1", i], [3600, 2, "ns2", i]]`. The fourth element is an array of three sub-arrays. Each sub-array contains a label and a hash: `[j, 3600, h'20010db80000000000000000000001']`, `[k, 3600, h'20010db80000000000000000000002']`, and `[l, 3600, h'20010db8000000000000000000003535']`. Colored arrows indicate that the labels 'j', 'k', and 'l' are associated with their respective hash values, and the 'i' labels in the second level point to the 'ns1' and 'ns2' labels in the third level.

<sup>1</sup><https://mailarchive.ietf.org/arch/msg/cbor/JOHCCB0zC46PrSq-61MMeV--mpU/>

# Name Compression Idea 1: Reference within CBOR object

- Loosely based on Christian's idea<sup>1</sup>:
  - Current draft: domain-name = tstr .regex "<sup>i</sup>([^.]+[.])<sup>j</sup>\*[^.]+"
  - Proposal: domain-name = (\*text) ; lowercase, NFC labels (cf. core-href)
- Use componified name to also allow for uint in name, i.e., reference

```
[ ["example", "org", 12, 1],  
  [[3600, 12, "_coap", "_udp", "local"]],  
  [[3600, 2, "ns1", i], [3600, 2, "ns2", i]],  
  [ [  
    j, 3600,  
    h'20010db8000000000000000000000001'  
  ], [  
    k, 3600,  
    h'20010db8000000000000000000000002'  
  ], [  
    l, 3600,  
    h'20010db80000000000000000000003535'  
  ] ] ]
```

- What do *i*, *j*, *k*, *l* represent?
  - Offset in CBOR binary?
  - E.g., *i*-th string in CBOR object?
- How to detect if name ends?

<sup>1</sup><https://mailarchive.ietf.org/arch/msg/cbor/JOHCCB0zC46PrSq-61MMeV--mpU/>

## Name Compression Idea 2: Reference components in pre-defined table

- Combining `name_ref` proposal in upcoming paper by Lemogue, Martinez, Toutain, and Bouabdallah as well as Christian's proposal

```
[ [{"example", "org"},
  [{"_coap", "_udp", "local"},
  ["ns1", 0],
  ["ns2", 0]],
 [0, 12, 1],
 [[3600, 12, 1]],
 [[3600, 2, 2], [3600, 2, 2, 3]],
 [ [
   1, 3600,
   h'20010db8000000000000000000000001'
 ], [
   2, 3600,
   h'20010db8000000000000000000000002'
 ], [
   3, 3600,
   h'20010db800000000000000000000003535'
 ] ] ]
```

## Name Compression Idea 2: Reference components in pre-defined table

- Combining `name_ref` proposal in upcoming paper by Lemogue, Martinez, Toutain, and Bouabdallah as well as Christian's proposal

```
[ ["example", "org"],  
  [ "_coap", "_udp", "local"],  
  [ "ns1", 0],  
  [ "ns2", 0]],  
 [0, 12, 1],  
 [[3600, 12, 1]],  
 [[3600, 2, 2], [3600, 2, 2, 3]],  
 [ [ [ 1, 3600,  
      h'20010db8000000000000000000000001'  
    ], [ 2, 3600,  
      h'20010db8000000000000000000000002'  
    ], [ 3, 3600,  
      h'20010db800000000000000000000003535'  
    ] ] ]
```

## Name Compression Idea 2: Reference components in pre-defined table

- Combining `name_ref` proposal in upcoming paper by Lemogue, Martinez, Toutain, and Bouabdallah as well as Christian's proposal

```
[ [ ["example", "org"],  
  [ "_coap", "_udp", "local"],  
  [ "ns1", 0 ],  
  [ "ns2", 0 ] ],  
 [ 0, 12, 1 ],  
 [ [ 3600, 12, 1 ] ],  
 [ [ 3600, 2, 2 ], [ 3600, 2, 2, 3 ] ],  
 [ [ [ 1, 3600,  
  h'20010db8000000000000000000000001'  
 ], [ [ 2, 3600,  
  h'20010db8000000000000000000000002'  
 ], [ [ 3, 3600,  
  h'20010db8000000000000000000000003535'  
 ] ] ]
```

- Introduces more overhead
- Very much like packed CBOR ...
- Should base format just be packed CBOR as well?

## Next Steps

- Rotate flags? (QR flag in default flags, DO flag in extended flags are MSB)



In CBOR: **19 80 00**

Rotated CBOR: **01**

- Specify and evaluate name compression
- Runtime analysis