A Concise Binary Object Representation (CBOR) of DNS Messages

draft-lenders-dns-cbor

Packed CBOR Requirements and Evaluation

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

IETF 123, CBOR WG Session, 2025-07-25

Outline

Recap

Evaluation

Conclusion

Recap: application/dns+cbor uses Packed CBOR

- Name compression using table setup tag TBD28259:
  - “virtual” packing table containing all name suffixes in the message
  - Repeated name suffixes are shared value references to that table
  - Splicing integration[1] allows for insertion of sequences into existing lists
- Further compression (`;packed=1`) using table setup tag 113:
  - Straight references for (IPv6) prefix compression
  - Shared value references for repeated values (*e.g.*, TTLs, RR types, RR classes)

[1] https://github.com/cbor-wg/cbor-packed/pull/25

Recap: Changes in Packed CBOR Since We Started

- Reduction of number of allocated tags / simple values
- Current Packed CBOR is parametrized (as of -15):
  - *A*: Number of simple values used for shared reference
  - *B*: Number of tags 1+1 used for low-index straight references (beyond tag 6)
  - *C*: Number of tags 1+1 used for low-index inverted references (beyond tag 6)
- Higher index straight and inverted references are now 6([*N*, rump]) instead of 1+*b* tags ($b > 1$, $N \geq 0$ straight, $N < 0$ inverted)
- Proposals, for parameters:
  - **Original**: $A = 16$, $B = 32$, $C = 8$
  - **Radical** (IETF 122): $A = x$, $B = 0$, $C = 0$
  - **Maybe reduced** (IETF 122): $A = x$, $B = 16$, $C = 8$
  - **Economic** (2025-05-14 interim): $A = x$, $B = 8$, $C = 8$
  - **Latest** (2025-07-09 interim): $A = 12$, $B = 8$, $C = 8$

Evaluation: Data Corpus

IoT data sets

YourThings[1] | IoTFinder[2] | MonIoTr[3]

- Collected throughout 2019
- DNS & mDNS (DNS-SD) traffic
- 90 consumer devices from 50 vendors
- 0.2 million queries, 1.3 million responses
- 2336 unique queried names

[1] 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.

[2] 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.

[3] 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.

Evaluation: Analysis of Data Corpus

Common name suffixes per message:

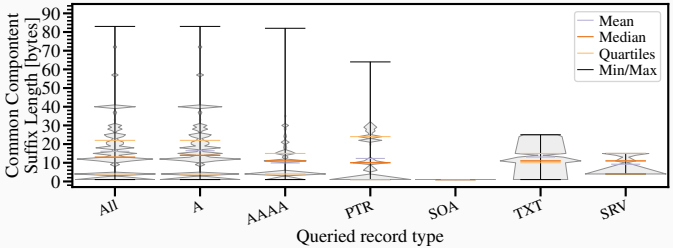

Common IP prefixes addresses per message:

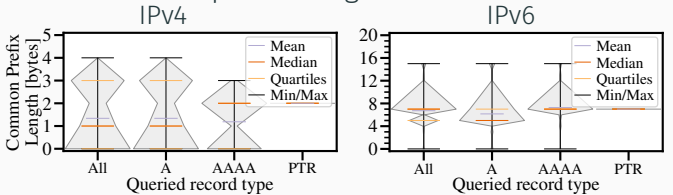

Reminder: Preliminary Evaluations for *A*

- Carsten showed results from this dataset already in the May interim
- Given *A*: What share of messages has all suffixes covered for name compression?

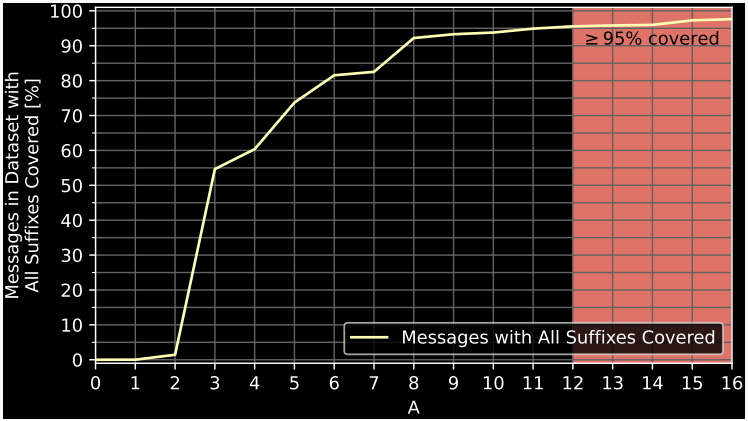

Reminder: Preliminary Evaluations for A

- Carsten showed results from this dataset already in the May interim
- Given A: What share of messages has all suffixes covered for name compression?

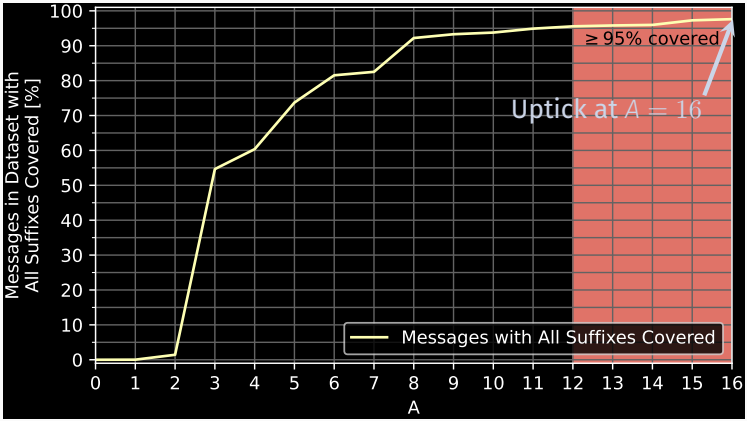

Evaluation: Setup

What are the actual message sizes?

- Name compression only uses shared value references
  $\Rightarrow$ Only *A* relevant
- `;packed=1` does not use inverted references
  $\Rightarrow$ Only *A* and *B* relevant
- Parametrized `cbor4dns` accordingly[2]

[2] https://github.com/netd-tud/cbor4dns/tree/parametrized-packed

Evaluation: Name Compression Only

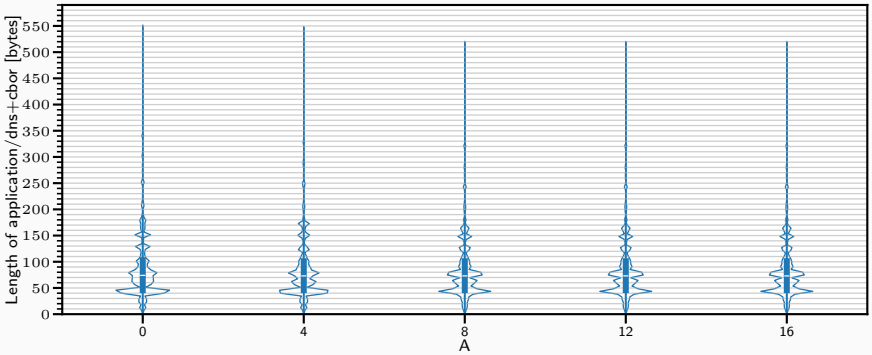

Message sizes of name compression for various values of *A*.

Evaluation: Name Compression Only

- Let’s look at statistics for details

Message sizes of name compression for various values of *A*.

Evaluation: Name Compression Only, Statistics

A	μ	σ	min	$q_{0.1}$	$q_{0.2}$	$q_{0.3}$	$q_{0.4}$	$q_{0.5}$	$q_{0.6}$	$q_{0.7}$	$q_{0.8}$	$q_{0.9}$	$q_{0.95}$	max
0	85.5	56.4	5	41	43	47	63	74	80	94	125	164	187	548
4	83.9	54.7	5	40	42	47	61	73	80	90	124	159	181	545
8	83.2	53.4	5	40	42	47	61	73	79	90	122	157	180	517
12	83.1	53.3	5	40	42	47	61	73	79	90	121	157	180	517
16	83.1	53.3	5	40	42	47	61	73	79	90	121	157	180	517

Length of application/dns+cbor for selected A in bytes.

Group 1 (worst)
Group 2
Group 3
Group 4 (best)

μ Mean
σ Standard diviation
q_x x-quantile

$q_{0.5}$ $1/2$-quantile (Median)
min Minimum
max Maximum

Evaluation: Name Compression Only, Statistics

A	μ	σ	min	$q_{0.1}$	$q_{0.2}$	$q_{0.3}$	$q_{0.4}$	$q_{0.5}$	$q_{0.6}$	$q_{0.7}$	$q_{0.8}$	$q_{0.9}$	$q_{0.95}$	max
0	85.5	56.4	5	41	43	47	63	74	80	94	125	164	187	548
4	83.9	54.7	5	40	42	47	61	73	80	90	124	159	181	545
8	83.2	53.4	5	40	42	47	61	73	79	90	122	157	180	517
12	83.1	53.3	5	40	42	47	61	73	79	90	121	157	180	517
16	83.1	53.3	5	40	42	47	61	73	79	90	121	157	180	517

Length of application/dns+cbor for selected A in bytes.

Group 1 (worst)
Group 2
Group 3
Group 4 (best)

μ Mean
σ Standard diviation
q_x x-quantile

$q_{0.5}$ $1/2$-quantile (Median)
min Minimum
max Maximum

Evaluation: application/dns+cbor;packed=1

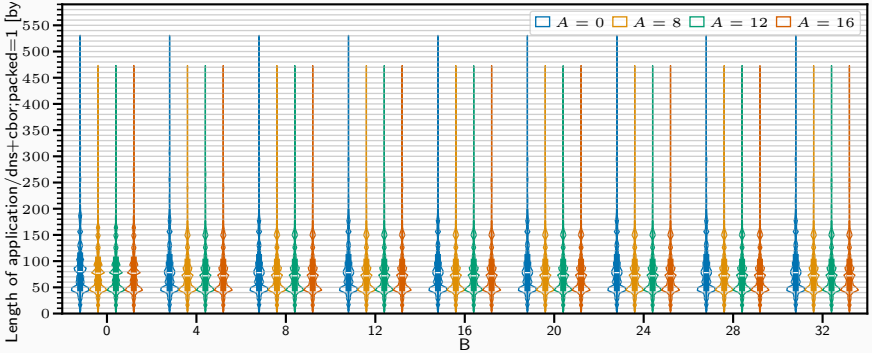

Message sizes of name compression for various values of *A* and *B*.

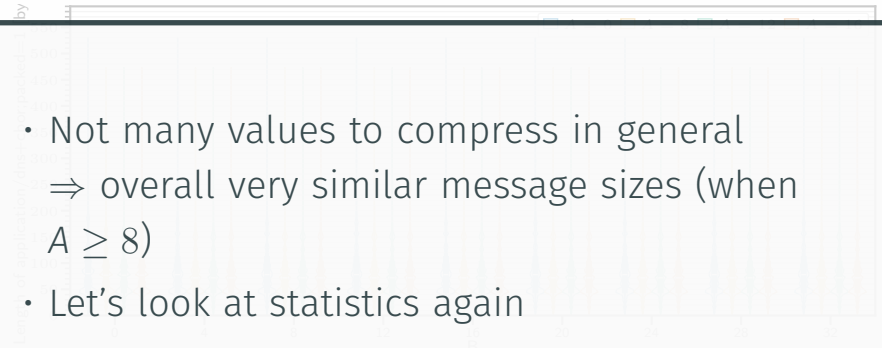

Message sizes of name compression for various values of A and B.

- Not many values to compress in general
  $\Rightarrow$ overall very similar message sizes (when $A \geq 8$)
- Let's look at statistics again

Evaluation: application/dns+cbor;packed=1, Statistics for $A = 8$

B	μ	σ	min	$q_{0.1}$	$q_{0.2}$	$q_{0.3}$	$q_{0.4}$	$q_{0.5}$	$q_{0.6}$	$q_{0.7}$	$q_{0.8}$	$q_{0.9}$	$q_{0.95}$	max
0	84.9	51.6	7	42	44	49	61	77	81	95	121	157	182	468
4	82.8	50.3	7	42	44	49	58	72	81	90	116	151	182	468
8	82.6	49.7	7	42	44	49	58	72	81	89	115	151	182	468
12	82.6	49.7	7	42	44	49	58	72	81	89	115	151	182	468
16	82.6	49.7	7	42	44	49	58	72	81	89	115	151	182	468
32	82.6	49.7	7	42	44	49	58	72	81	89	115	151	182	468

Length of application/dns+cbor;packed=1 for selected B in bytes for $A = 8$.

Group 1 (worst)	μ	Mean	$q_{0.5}$	$1/2$-quantile (Median)
Group 2	σ	Standard diviation	min	Minimum
Group 3 (best)	q_x	x-quantile	max	Maximum

Evaluation: application/dns+cbor;packed=1, Statistics for $A = 12$

B	μ	σ	min	$q_{0.1}$	$q_{0.2}$	$q_{0.3}$	$q_{0.4}$	$q_{0.5}$	$q_{0.6}$	$q_{0.7}$	$q_{0.8}$	$q_{0.9}$	$q_{0.95}$	max
0	84.9	51.5	7	42	44	49	61	77	81	95	121	157	182	468
4	82.7	50.2	7	42	44	49	58	72	81	90	116	151	182	468
8	82.5	49.6	7	42	44	49	58	72	81	89	115	151	182	468
12	82.5	49.6	7	42	44	49	58	72	81	89	115	151	182	468
16	82.5	49.6	7	42	44	49	58	72	81	89	115	151	182	468
32	82.5	49.6	7	42	44	49	58	72	81	89	115	151	182	468

Length of application/dns+cbor;packed=1 for selected B in bytes for $A = 12$.

Group 1 (worst)
Group 2
Group 3 (best)

μ Mean
σ Standard diviation
q_x x-quantile

$q_{0.5}$ $1/2$-quantile (Median)
min Minimum
max Maximum

Evaluation: application/dns+cbor;packed=1, Statistics for $A = 12$

B	μ	σ	min	$q_{0.1}$	$q_{0.2}$	$q_{0.3}$	$q_{0.4}$	$q_{0.5}$	$q_{0.6}$	$q_{0.7}$	$q_{0.8}$	$q_{0.9}$	$q_{0.95}$	max
0	84.9	51.5	7	42	44	49	61	77	81	95	121	157	182	468
4	82.7	50.2	7	42	44	49	58	72	81	90	116	151	182	468
8	82.5	49.6	7	42	44	49	58	72	81	89	115	151	182	468
12	82.5	49.6	7	42	44	49	58	72	81	89	115	151	182	468
16	82.5	49.6	7	42	44	49	58	72	81	89	115	151	182	468
32	82.5	49.6	7	42	44	49	58	72	81	89	115	151	182	468

Length of application/dns+cbor;packed=1 for selected B in bytes for $A = 12$.

Group 1 (worst)
Group 2
Group 3 (best)

μ Mean
σ Standard diviation
q_x x-quantile

$q_{0.5}$ $1/2$-quantile (Median)
min Minimum
max Maximum

Evaluation: application/dns+cbor;packed=1, Statistics for $A = 16$

B	μ	σ	min	$q_{0.1}$	$q_{0.2}$	$q_{0.3}$	$q_{0.4}$	$q_{0.5}$	$q_{0.6}$	$q_{0.7}$	$q_{0.8}$	$q_{0.9}$	$q_{0.95}$	max
0	84.8	51.5	7	42	44	49	61	77	81	95	121	157	182	468
4	82.7	50.2	7	42	44	49	58	72	81	90	115	151	182	468
8	82.5	49.5	7	42	44	49	58	72	81	89	115	151	182	468
12	82.5	49.5	7	42	44	49	58	72	81	89	115	151	182	468
16	82.5	49.5	7	42	44	49	58	72	81	89	115	151	182	468
32	82.5	49.5	7	42	44	49	58	72	81	89	115	151	182	468

Length of application/dns+cbor;packed=1 for selected B in bytes for $A = 16$.

Group 1 (worst)
Group 2
Group 3 (best)

μ Mean
σ Standard diviation
q_x x-quantile

$q_{0.5}$ $1/2$-quantile (Median)
min Minimum
max Maximum

Evaluation: application/dns+cbor;packed=1, Statistics for $A = 16$

B	μ	σ	min	$q_{0.1}$	$q_{0.2}$	$q_{0.3}$	$q_{0.4}$	$q_{0.5}$	$q_{0.6}$	$q_{0.7}$	$q_{0.8}$	$q_{0.9}$	$q_{0.95}$	max
0	84.8	51.5	7	42	44	49	61	77	81	95	121	157	182	468
4	82.7	50.2	7	42	44	49	58	72	81	90	115	151	182	468
8	82.5	49.5	7	42	44	49	58	72	81	89	115	151	182	468
12	82.5	49.5	7	42	44	49	58	72	81	89	115	151	182	468
16	82.5	49.5	7	42	44	49	58	72	81	89	115	151	182	468
32	82.5	49.5	7	42	44	49	58	72	81	89	115	151	182	468

Length of application/dns+cbor;packed=1 for selected B in bytes for $A = 16$.

Group 1 (worst)	μ	Mean	$q_{0.5}$	$1/2$-quantile (Median)
Group 2	σ	Standard diviation	min	Minimum
Group 3 (best)	q_x	x-quantile	max	Maximum

Conclusion: Which *A*, *B*, *C* should be picked?

- In general compression achieved even with radical approach (not shown)
  - between 10 and 226 bytes saved with only name compression
  - between 8 and 314 bytes saved with packed
- **A**: Should be ≥ 8; 12 a good compromise, 16 also has benefits
- **B**: Micro-optimizations, but $B \geq 8$ seems most sensible
- **C**: not evaluated