On DNSSEC Negative Responses, Lies, and Zone Size Detection

Jonathan Demke, **Casey Deccio**
Brigham Young University
OARC 31, Austin, TX
Nov 1, 2019

BYU

The Domain Name System (DNS)

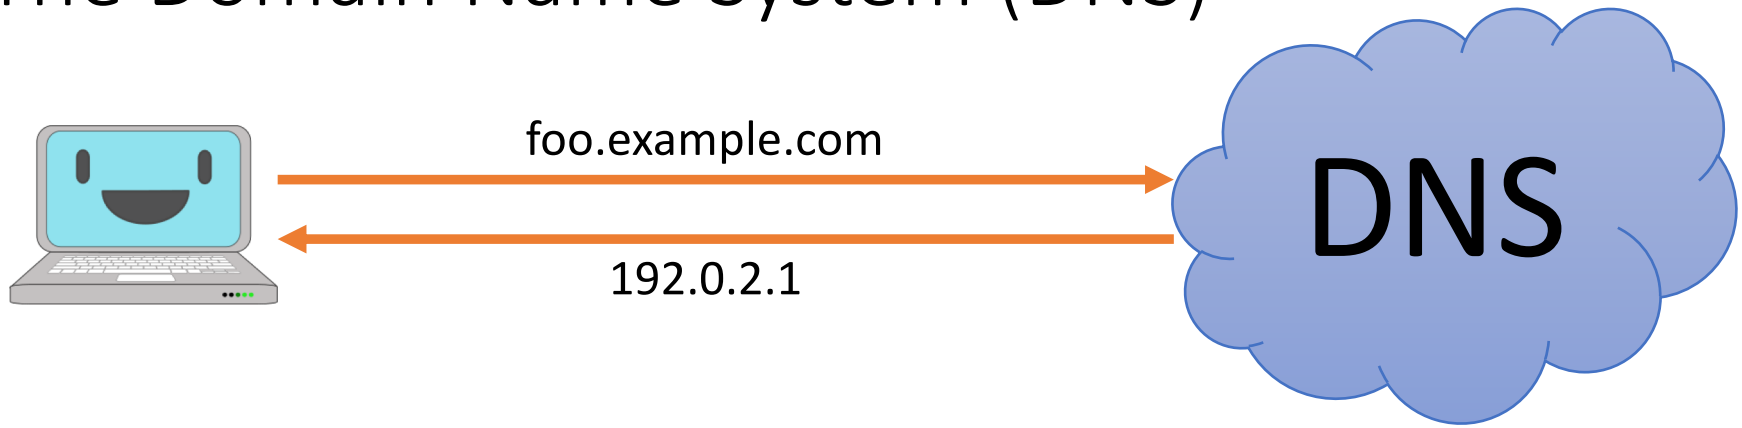

- Query-response protocol.
- Translates domain name to IP address (or other resource).

An Open Book?

- Responses will tell you whether or not a name exists.
- *But* it requires guess-and-check.
- **No general way to ask "What are all the names under example.com?"**

DNSSEC with NSEC

- NSEC records form a proof that a queried name doesn't exist.
- But they reveal the surrounding names.
- NSEC records have been used to discover and "walk" DNSSEC-signed zones [1].

[1] Osterweil, Ryan, Massey, Zhang, "Quantifying the Operational Status of the DNSSEC Deployment", IMC 2008.

DNSSEC with NSEC3

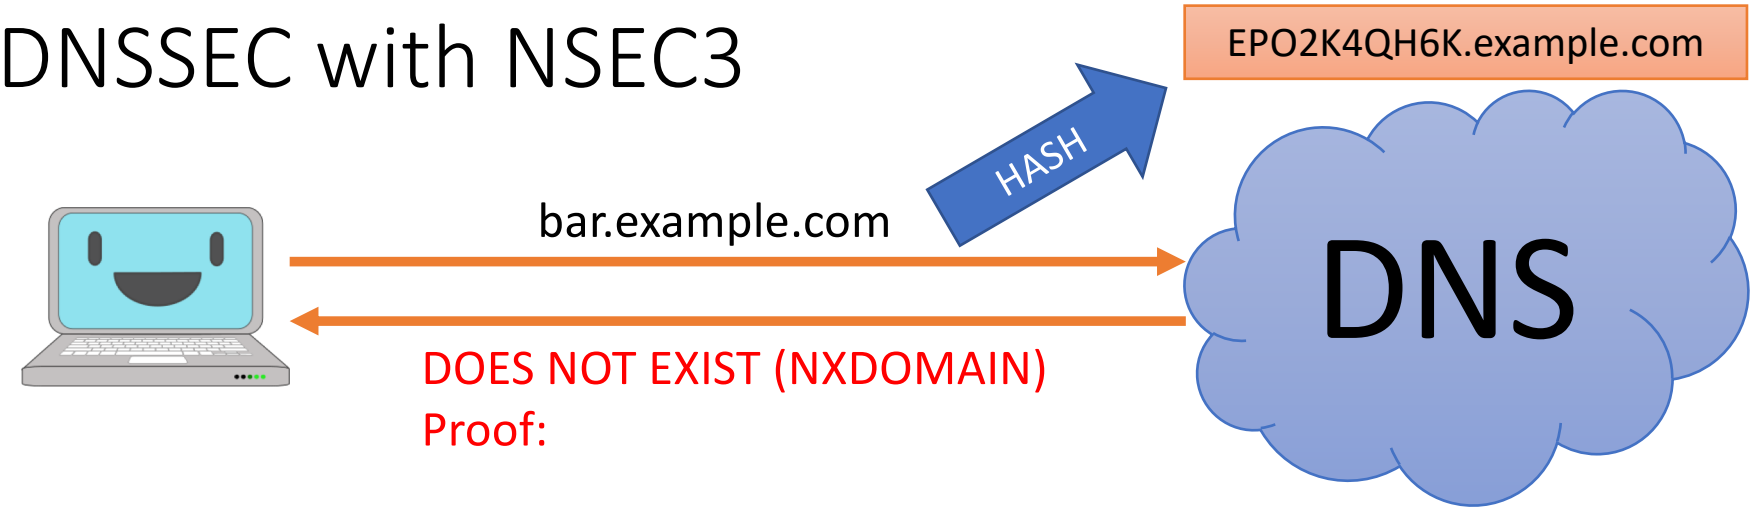

- NSEC3 records send the surrounding **hashes**.
- Hashes obfuscate the names.
- NSEC3 hashes have been broken with GPUs in a relatively short period of time [2].

[2] Wander, Schwittmann, Boelmann, Weis, “GPU-Based NSEC3 Hash Breaking”, NCA 2014.

White Lies!

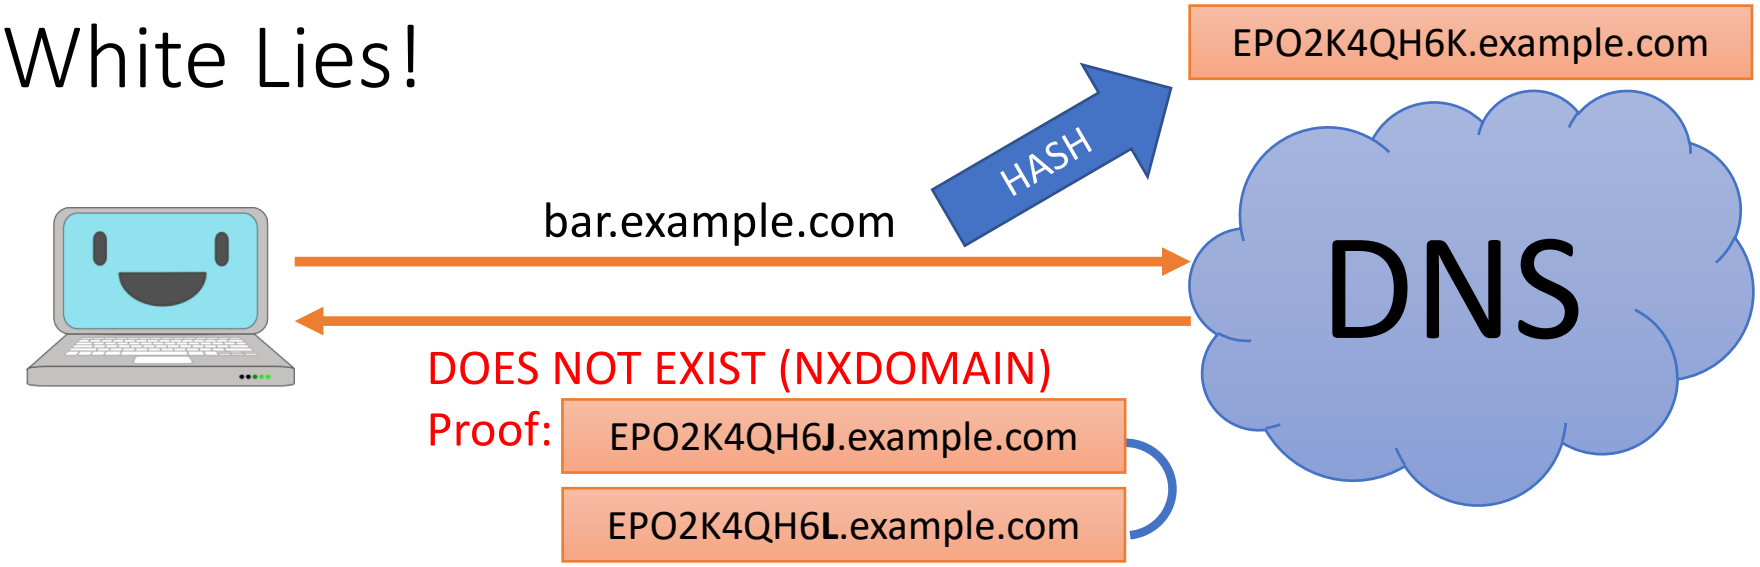

- Authoritative servers generate proof on-the-fly with minimally covering NSEC3 records [3].
- Authoritative servers must have access to **private key** – to sign the records created on-the-fly.

[3] Dan Kaminsky, “Phreebird”, https://dankaminsky.com/phreebird/ 2011.

Black Lies!

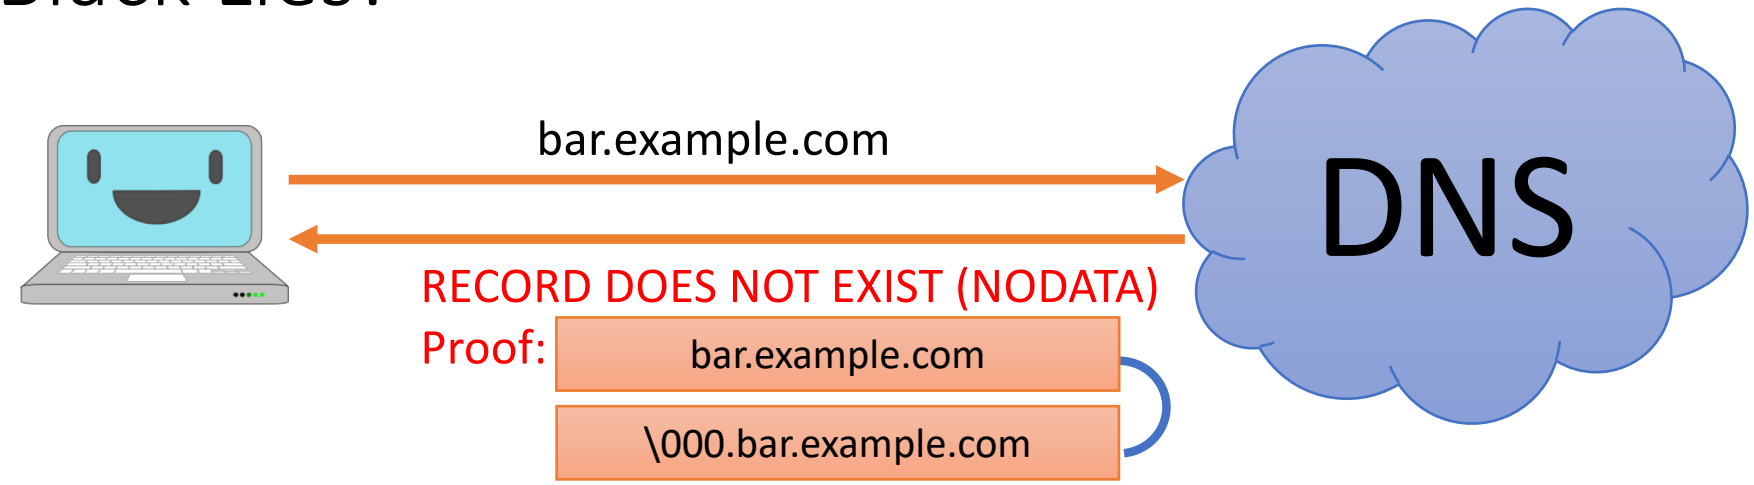

- Authoritative servers generate proof on-the-fly with minimally covering NSEC records.
- Proof is that **record** (not name) doesn't exist (NODATA).
- Authoritative servers must have access to **private key** – to sign the records created on-the-fly.

[4] Dani Grant, "Economical With The Truth: Making DNSSEC Answers Cheap", https://blog.cloudflare.com/black-lies/ 2016.

Survey of Signing Methods

- Zones extracted from zone files for 821 top-level domains (TLDs).
- 2.2M DNSSEC-signed zones discovered.
  - Presence of DS records constituted “signed” for the purpose of this study.

Traditional NSEC	Traditional NSEC3	White Lies NSEC3	Black Lies NSEC	Unclassified	Total
241,045 (11%)	1,167,219 (53%)	657,091 (30%)	48,059 (2%)	66,646 (3%)	**2,182,987**

What Else Can We Learn?

- Traditional NSEC3 comprises over half of the signed zones in our survey.
- Can we learn the **size** of an NSEC3-signed DNS zone by analyzing a few responses from a DNS server?

https://www.vectorstock.com/royalty-free-vector/man-with-question-mark-flat-icon-pictogram-vector-4920218

Analogy: Pie slices

https://www.foodnetwork.com/recipes/food-network-kitchen/apple-pie-recipe-2011423

If this slice of pie is representative:

How many guests shared the pie?

https://www.foodnetwork.com/recipes/food-network-kitchen/apple-pie-recipe-2011423

360°

90°

$$\frac{Total\ Degrees}{Slice\ Degrees} = \frac{360^o}{90^o} = 4 \text{ slices}$$

https://www.foodnetwork.com/recipes/food-network-kitchen/apple-pie-recipe-2011423

NSEC3 distance

$$d(n) = n_{owner} - n_{next}$$

NSEC3 hash space

$$H = 2^{160}$$

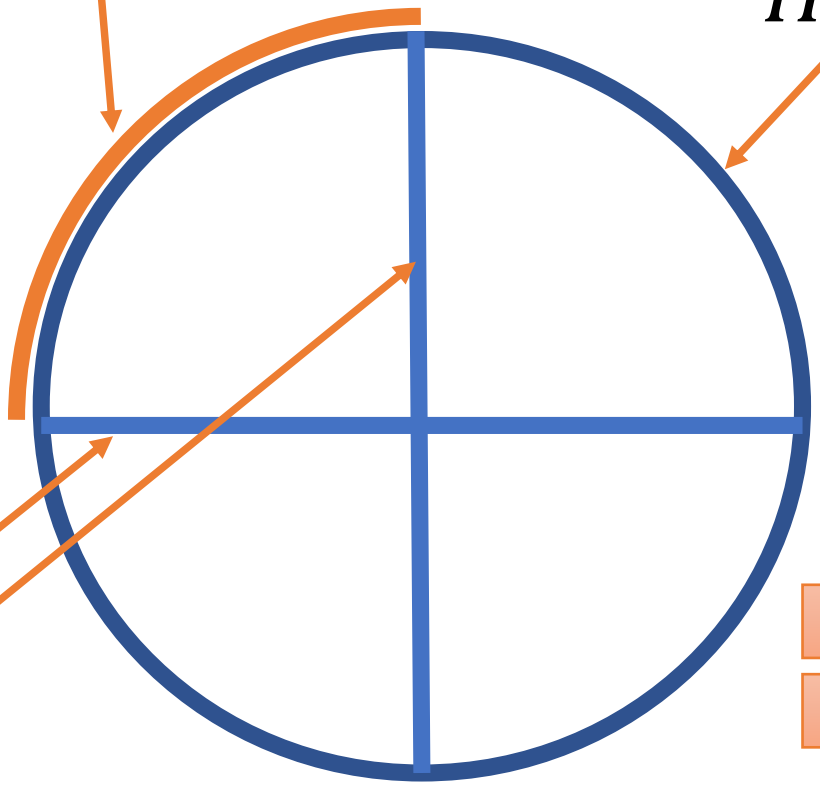

NSEC3 hash

n_{owner}

n_{next}

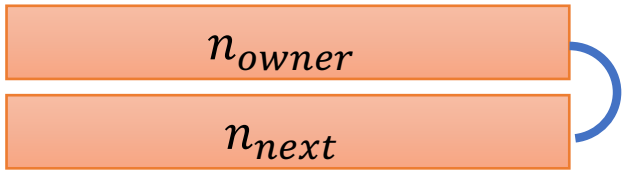

Easy as pie?  $\frac{H}{d(n)}$?

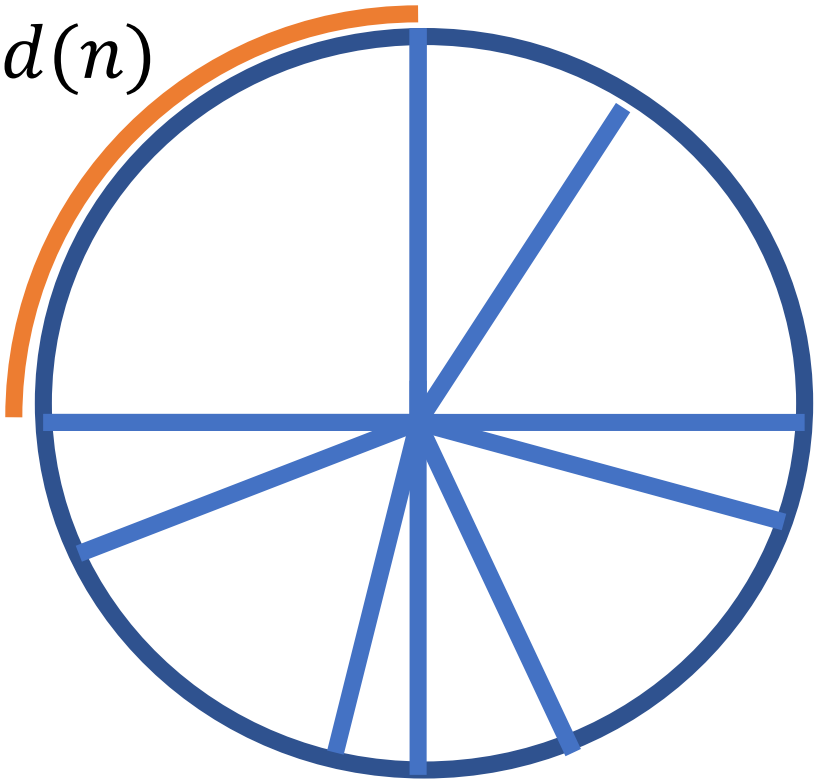

- Unlike shared pie slices, NSEC3 distances are **not equal**.
- A **sample** of distances is required.

How are NSEC3 hashes distributed?

Experiment

- Generate NSEC3 hash for 100K domain names using `nsec3hash`.
- Divide the hash space, H, into 1,024 equal-sizes bins.
- How many hashes are in each bin?

Result

- Hash values are **uniformly distributed** across hash space, H.
- The number of hashes per bin follows a normal distribution.

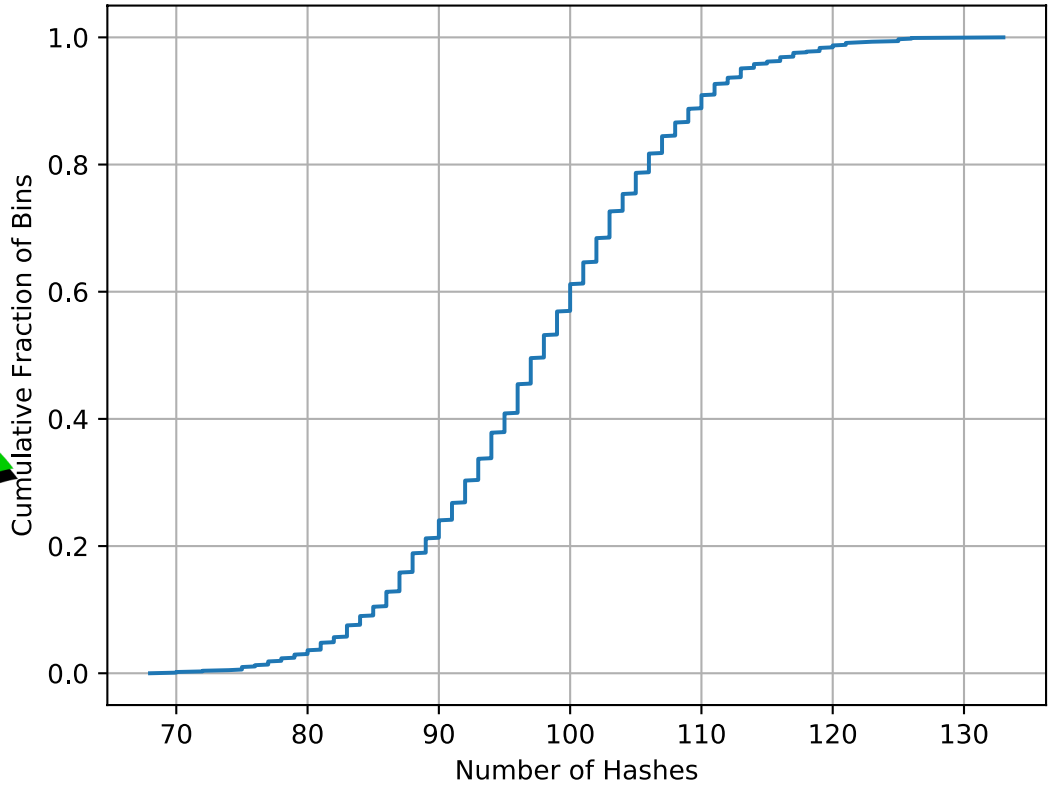

Let's test it out!

Test Zone

- 10,000 randomly-generated names.
- Signed with NSEC3.

Experiment

- Trial: 18 queries = 20 NSEC3 records
- Number of trials: 1,000
- Size: $\frac{H}{\left(\frac{\sum_{n \in N} d(n)}{|N|}\right)}$
- Error: $\frac{size_{test} - size_{actual}}{size_{actual}}$

Result

- Min, median, max all below 0 error
- Median about -0.5 (low estimate)

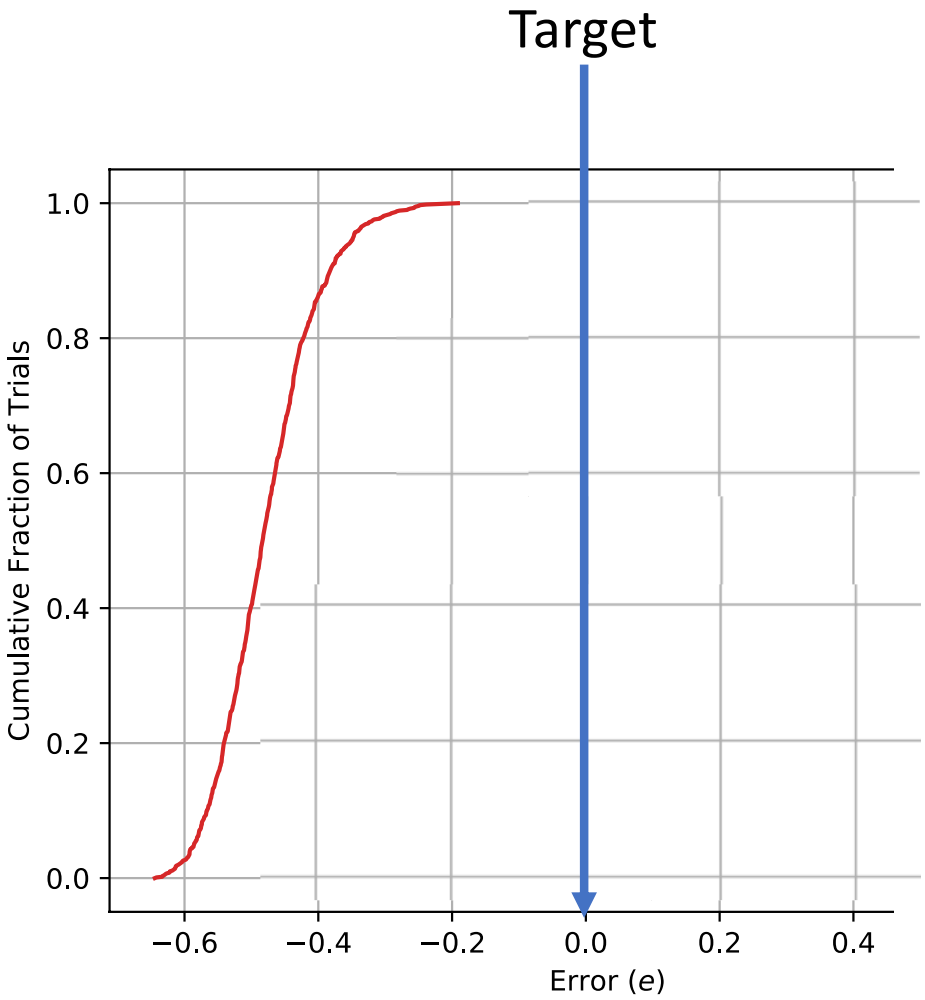

Huh?

- We measured distribution of NSEC3 **hashes**.
- We have *not* measured the distribution of NSEC3 **distances**.

Let's sign some zones!

Experiment

- 500 zones generated and signed:
  - 100 zones of size 10^2
  - 100 zones of size 10^3
  - 100 zones of size 10^4
  - 100 zones of size 10^5
  - 100 zones of size 10^6
- NSEC3 distances plotted.

Results

- Distances offset by factor of 10.
- Distribution is exponential.

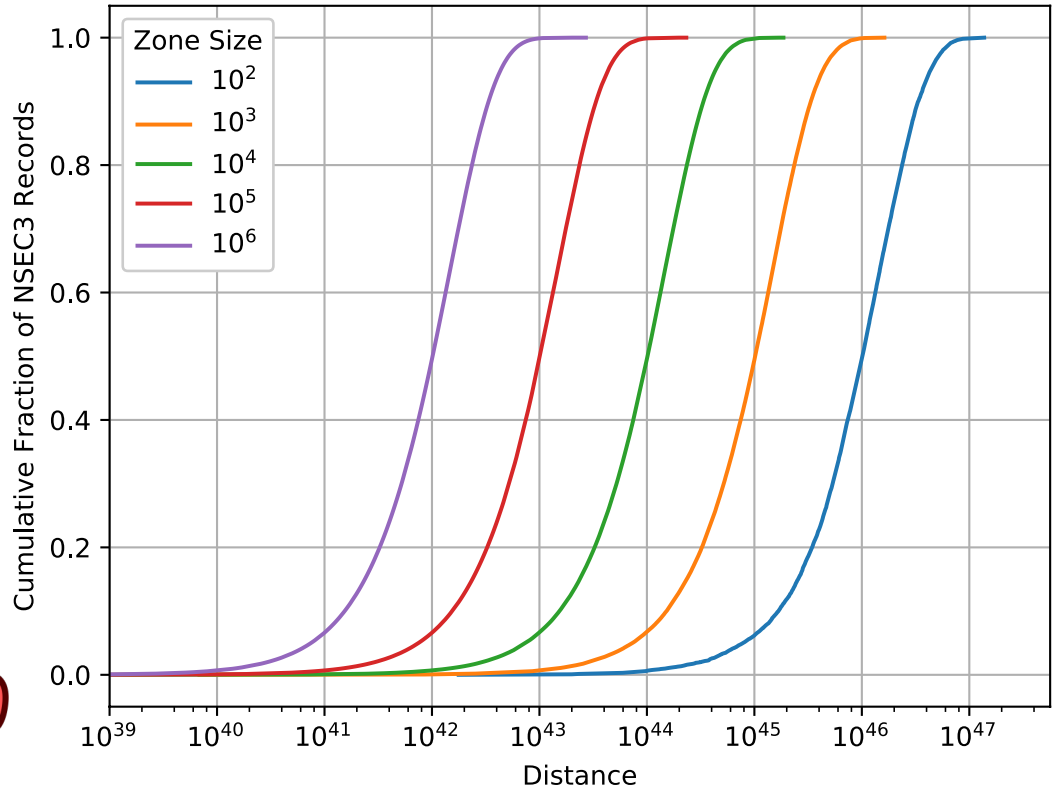

Distances and Probabilities

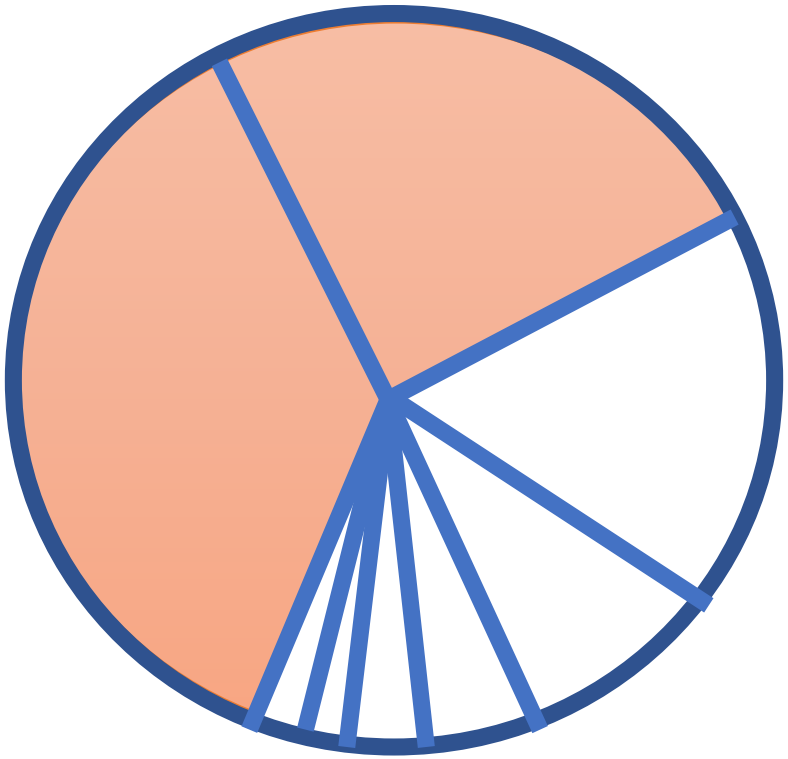

(hashes sorted roughly by distance)

- Distances are exponentially distributed.
- The probability of the hash of the queried name landing falling in a larger, non-representative distance is much higher.

NSEC3 Distance Distribution Revisited

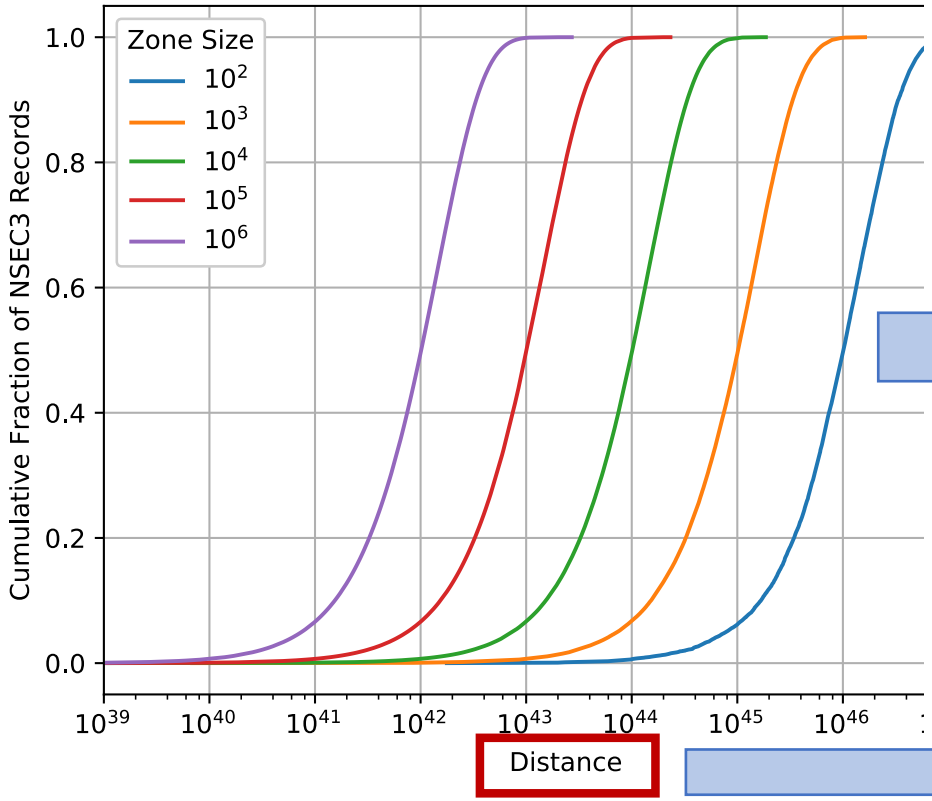

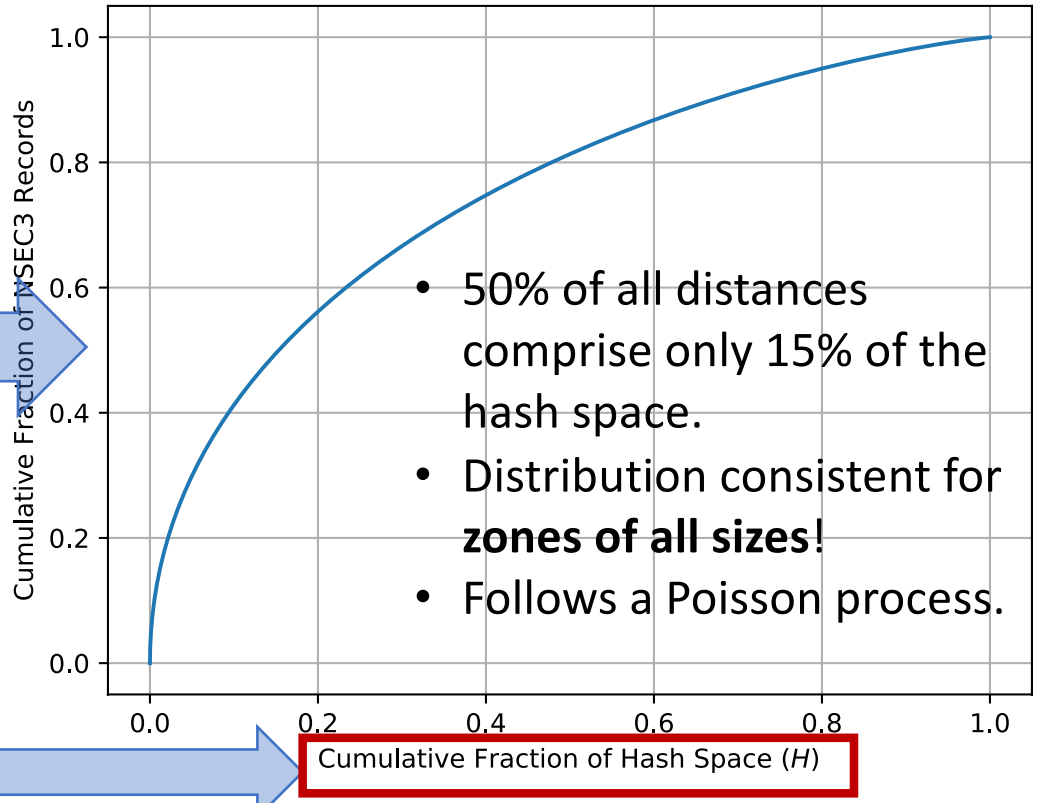

- 50% of all distances comprise only 15% of the hash space.
- Distribution consistent for **zones of all sizes**!
- Follows a Poisson process.

Weighted Average

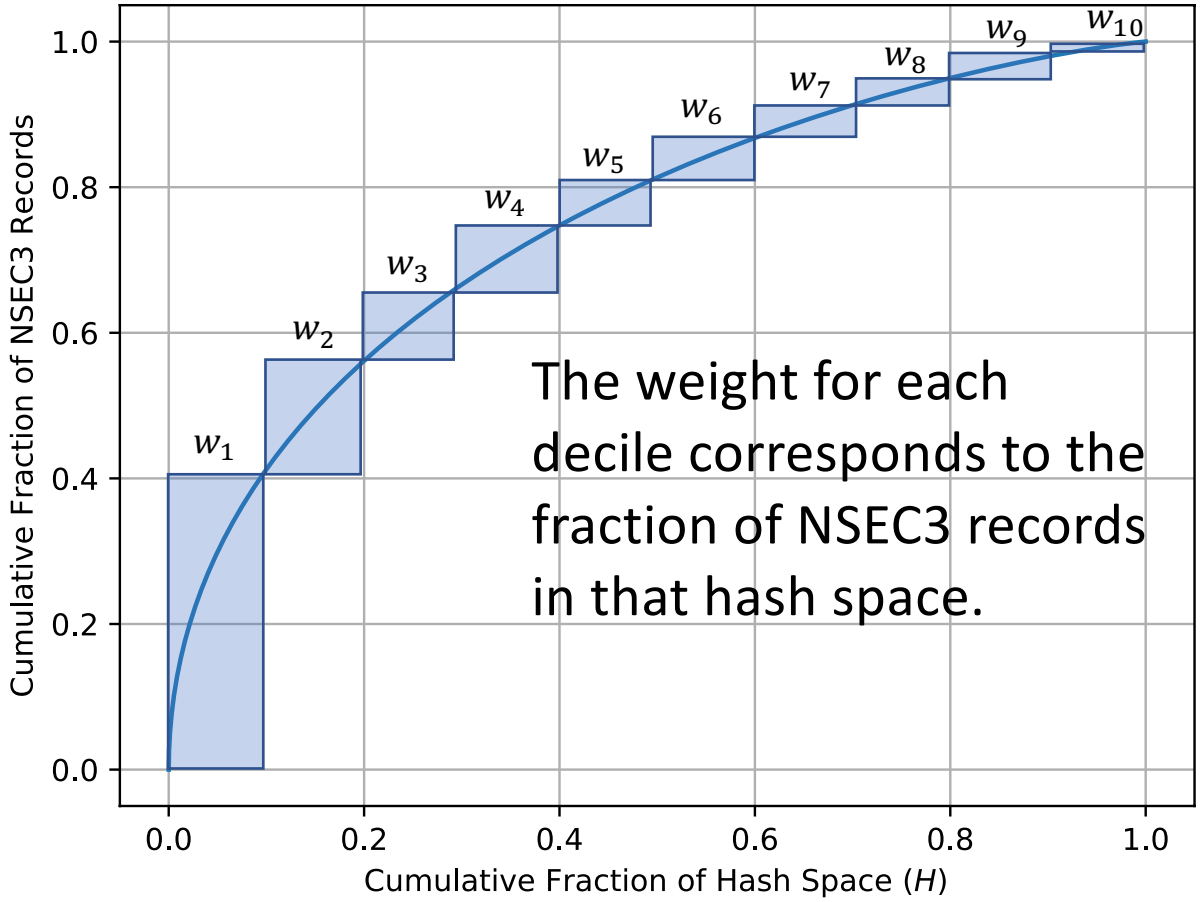

Using the weighted average

Experiment

- Zone sizes: 10^2, 10^3, 10^4, 10^5, 10^6
- Trial: 18 queries = 20 NSEC3 records
- Number of trials: 1,000
- Size: $$\frac{H}{\left(\sum_{1\leq i\leq q} \frac{w_i \sum_{n\in N_i} d(n)}{|N_i|}\right)}$$
- Error: $\frac{size_{test} - size_{actual}}{size_{actual}}$

Result

- For zones smaller than 100K, more than 75% of trials were within 20% of zone size.

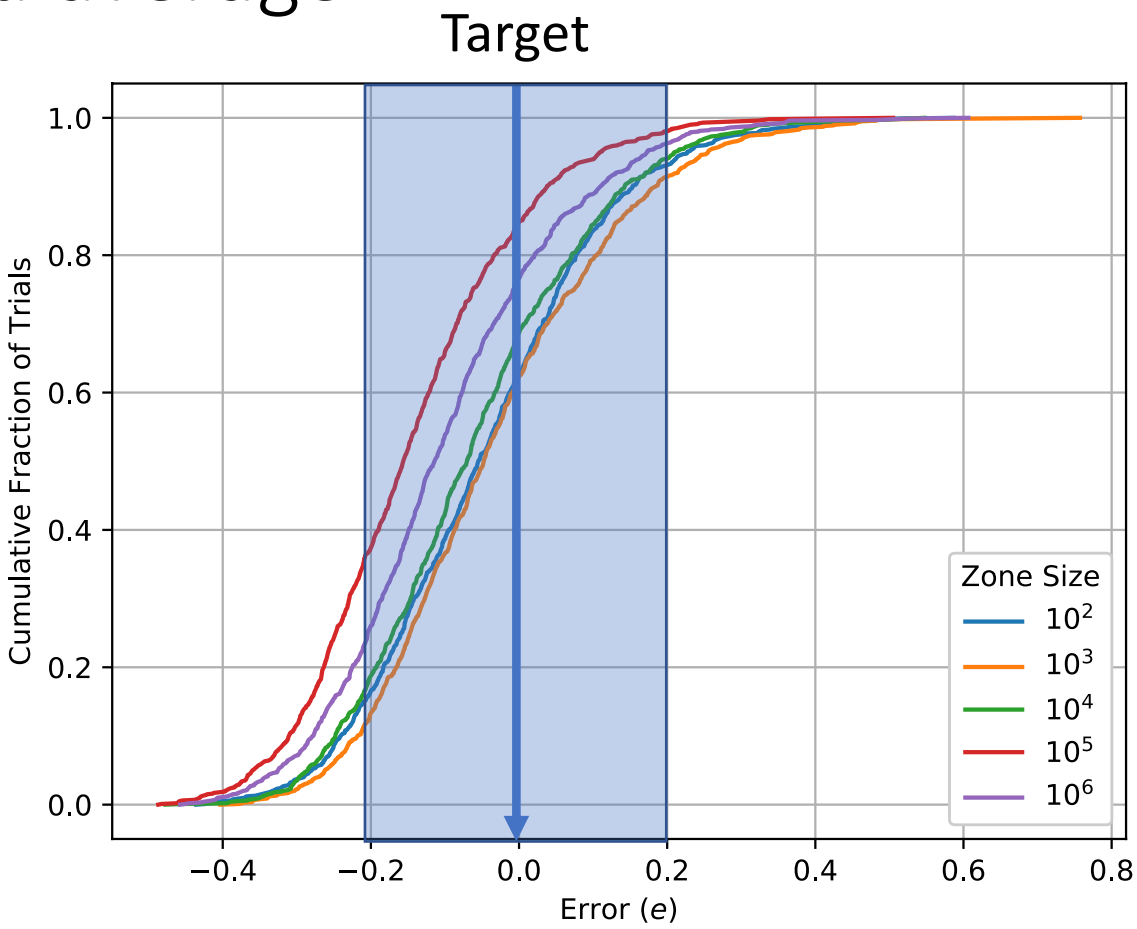

NSEC3 Zone Sizes - The Results!

- Nearly 90% of have 10 names or fewer.
- 99% of zones have 40 names or fewer.
- 1% reached up to 4M.

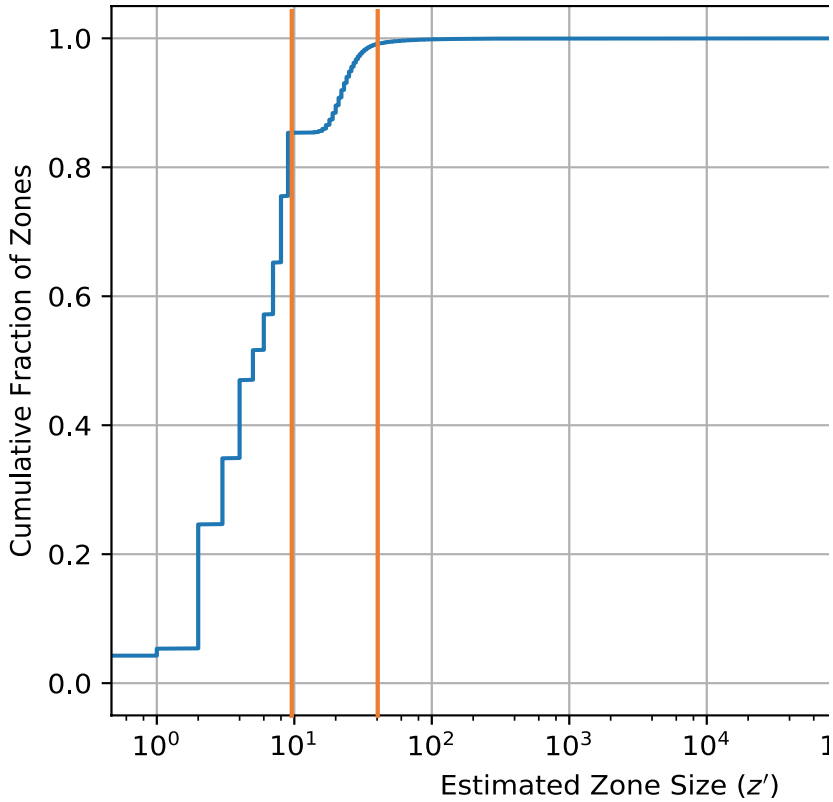

Summary

- DNSSEC provides origin authentication.
- NSEC and NSEC3 provide authenticated denial of existence – but reveal more about a DNS domain.
- Obfuscation solutions (white lies, black lies) exist, but with their own challenges.
- Sizes of traditional NSEC3-signed zones can also be estimated with few queries.

https://www.foodnetwork.com/recipes/food-network-kitchen/apple-pie-recipe-2011423

Questions?

casey@byu.edu

https://www.vectorstock.com/royalty-free-vector/man-with-question-mark-flat-icon-pictogram-vector-4920218