

1. DNS Around Us

Before heading into the details of DNS we want to take a look at the local deployment and query some information to get familiar with it.

Tools: dig, tshark, tcpdump. *Record types:* A, AAAA, NS, DNSKEY, RRSIG.

- (a) Find the authoritative name server for HAW Hamburg.
- (b) Can a DNS query tell you whether there is IPv6 support for the HAW website?
- (c) Does the HAW domain support DNSSEC?
- (d) Request an A record for a domain using dig with 8.8.8.8 as the DNS resolver, *e.g.*, `haw-hamburg.de`, and observe the requests with a sniffer. Does this match your expectations? Compare your observation to an ANY request to `peacecorps.gov`.

2. Record Sizes of DNS Responses

DNS record sizes matter for multiple reasons. First they are often queried and second they might be abused for criminal purposes, which will be discussed in more detail in a later lecture.

Tools: dnspython, dig

Dataset: A Tranco¹ top 1M toplist from April is located in `shared-data/haw` on mobi8: `toplists/tranco/2025-04-07_tranco_top-1m.csv`

- (a) Plot a distribution of the top-level domains in the dataset.
- (b) Collect DNS records for a sample of 3k entries from the Tranco top 1M domains (or more if you have the time). Build three datasets for A, AAAA, and ANY requests.
- (c) Visualize the distribution of the record sizes and print the 10 domains with the largest responses. How do the sets differ?
- (d) Do some resolvers behave differently? The system default is 8.8.8.8, but there are many more resolvers.

3. DNSSEC Prevalence

DNSSEC extends DNS with the authentication of messages and records. Security itself is often not enough to get companies to implement new things—unless it has significant monetary benefits. Let's see how prevalent DNSSEC is among the top web sites.

Tools: python via dnspython² or dig.

- (a) How many of the Tranco top 1M deploy DNSSEC? Sample 3k entries for an estimate (consider using `DataFrame.sample`³).

¹<https://tranco-list.eu>

²<https://dnspython.readthedocs.io/en/stable/index.html>

³<https://pandas.pydata.org/pandas-docs/stable/reference/api/pandas.DataFrame.sample.html>

- (b) Visualize the DNSSEC support by TLD, *e.g.*, through a bar graph of the top 10.
- (c) Make a CDF that accumulates the share of DNSSEC supporting domains with increasing rank.

4. DNSSEC Validation

DNSSEC is useless if records cannot be validated by resolvers. In this exercise we will look into the validation of signatures using python.

Tools: dnspython, DNS Viz (<https://dnsviz.net/>).

- (a) Find one domain that supports DNSSEC and validate its signature.
- (b) Implement a python script that validates a given domain. Walking the chain of trust requires extra work that is not easily handled by the library, *skip it for now*. Here is how this could be implemented:
 - i. Instantiate a resolver,
 - ii. Find the responsible name server,
 - iii. Perform a query for the DNSKEY record,
 - iv. Validate the signature,
 - v. Handle errors accordingly. (Distinguish between validation and missing records!)
- (c) Check a sample from the Tranco list for invalid records and present your findings.