

# Resolver Data Prefetch

Wouter C.A. Wijngaards

[wouter@NLnetLabs.nl](mailto:wouter@NLnetLabs.nl)

# NLnet Labs

*To develop Open Source Software and Open Standards for the benefit of the Internet.*

Paraphrased Article I of the Foundations Charter

# Introduction

- Draft-wkumari-dnsop-hammer-00
- Premise: fetch data that is about to expire so that client does not get spike in latency when lookup needs to be redone when TTL expires on data.
  - Expected DNS cache response latency:  
local ISP roundtrip
  - Expected DNS re-lookup latency:  
local ISP roundtrip + authority server roundtrip
- Implementation: prefetch data that is about to expire. Give clients cached previous result and start prefetch.

# Unbound Implementation





- Not the same as draft-hammer
- Differences:
  - Lookup in last 10% of TTL (not in last x seconds).
  - Only needs one query in last of 10% TTL to start lookup.
  - Strange cache policy during prefetch.
  - No prefetch for `OriginalMessageTTL`  $\leq 9$  sec
  - '10%' not configurable (compiled constant).

# Real Data

- Kindly provided by SURFnet.
- Data from prefetch in a production resolver.
  - Statistics have been graphed with cacti
- With Unbound, “prefetch: yes” in unbound.conf
- The graphs are one-week runs.
- Num.prefetch is 1.5% of total qps, 39 of 2689.6 qps.  
(cache hit rate is 84.9%).
- Very little difference

# Prefetch:no

# Prefetch:yes



# Design Decisions

- Keep load off authority servers
  - 10% shorter TTL means +10% qps
- Keep load off local resolver
  - 10% shorter TTLs thus max +10% work
  - Not configurable 10% because clients shouldn't burden authorities too much
- Do not prefetch unless there is real interest by clients
  - Wait for client query
  - Hammer has a similar design
- No 'sticky resolver': domain transfer do not stick to old
- Prefetch 'popular domains' but calculation what 'popular' means is strange.

# 'Popular'


- With  $N$  queries in the RR-TTL interval, and  $N/TTL > \text{threshold}$  0.1, it is popular
- Also 'less popular but high TTL' answers can be prefetched
- The authority owner can influence prefetch by changing the TTLs
- Unbound does not measure  $N/TTL$ 
  - query in last 10% is seen as indication that the expected value for  $N/TTL$  (assuming semi-flat qps graph) exceeds threshold.
  - Authority owners can double the TTL from  $10/qps$  to also get prefetched in 'standard deviation' cases.



# Cache Policy

- Relative prime numbers in the cache update policy
  - If you only update RRsets that are about to expire (as unbound does)  
then the TTL that remains creates a relative-prime-number seeking algorithm with repeated prefetches
  - Because the relative prime number may be very low, also update other RRsets to increase result TTL; 60s leeway in unbound.
- If you prefetch data 'pretend' that cached nameserver information that is about to expire is not available when the prefetch lookup is performed

# Questions



If you use our software then consider to support NLnet Labs' mission Financially:

<http://www.nlnetlabs.nl/labs/contributors/>