Negative Caching of DNS Resolution Failures

draft-dwmtwc-dnsop-caching-resolution-failures

IETF 113 March 2022

The Problem

- Recursive name servers are bad at caching resolution failures
  - timeouts
  - SERVFAIL / REFUSED
  - validation failures
  - loops

Facebook outage, October 4, 2021

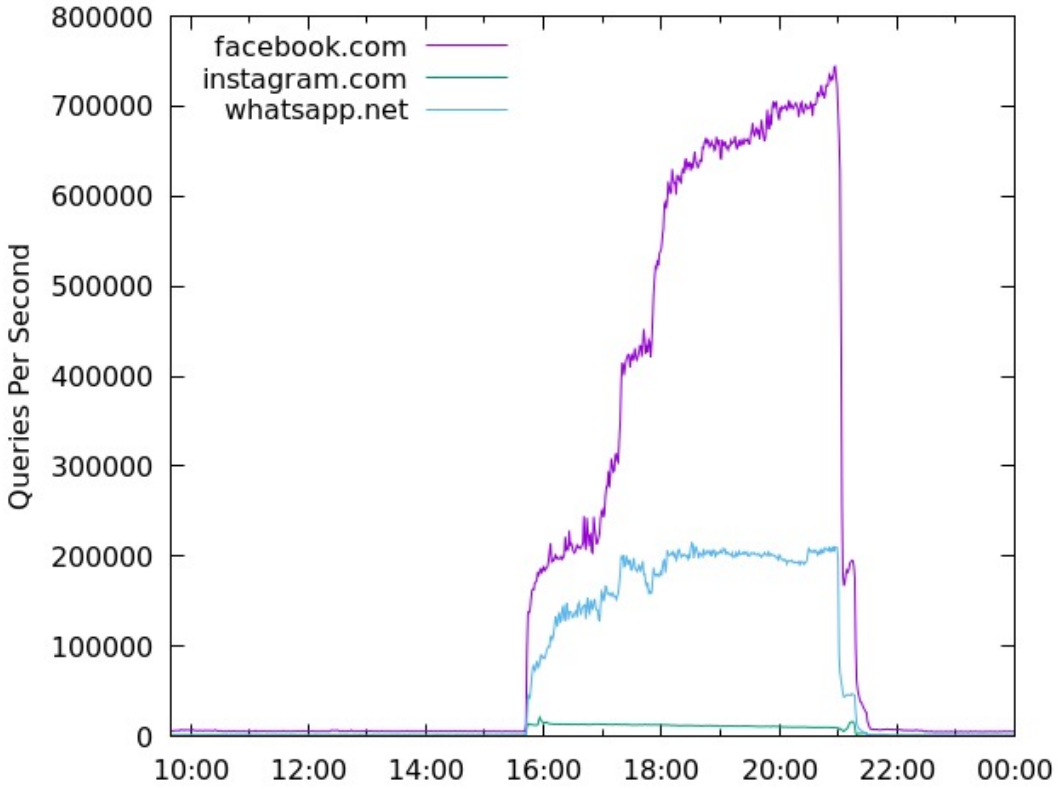

Evidence and Incidents

- Facebook outage (2021): 128x increase in queries to com/net
- Botnet domain research (2021): 1200x increase from SERVFAIL responses
- TsuNAME (2021): 500x increase from cyclic delegation
- NXNSAttack (2020): 1620x increase
- KSK Rollover post-revocation (2019): 80x increase in DNSKEY queries at root servers
- Dhukovni Survey Error Amplification (2019)
- Dyn Attack (2016)
- Roll Over and Die (2009)

Existing requirements

- 2308: “negative caching should no longer be seen as an optional part of a DNS resolver”
  - but caching timeouts & server failures is OPTIONAL
- 4697: “An iterative resolver MUST NOT send a query for the NS RRSet of a non-responsive zone to any of the name servers for that zone's parent zone.”
  - plus some other SHOULDs
- 8767: "Attempts to refresh from non-responsive or otherwise failing authoritative nameservers are recommended to be done no more frequently than every 30 seconds."

Proposed New Requirements

- Resolvers MUST cache (all) resolution failures, for at least 5 seconds, and not longer than 5 minutes
- Resolvers SHOULD employ exponential backoff
- A resolver MUST NOT retry more than twice (three queries total) before considering a server unresponsive
- Reiterate and/or strengthen requirements about re-querying parent name servers of a zone experiencing failures

Discussion and Request for Adoption