VERISIGN

Updating Requirements for Caching DNS Resolution Failures

Duane Wessels, Matthew Thomas

February 2022

Presentation Overview

- Recursive DNS name servers exhibit aggressive behavior from resolution failures.
- Real world resolution failure examples with data and graphs.
- Current protocol requirements for caching resolution failures.
- Proposed updated requirements.

What is a Resolution Failure

- When a recursive name server fails to receive a useful response from any of a zone's authoritative name servers.
- Resolution failures:
  - RCODE other than NoError or NXDomain
  - Timeouts
  - Lame delegation
  - DNSSEC validation failure
- NOT resolution failures:
  - NXDomain response
  - NoError No Data

Observations of Abnormal Behavior

Facebook Outage, October 2021

- Six-hour outage of Facebook's network.
- None of Facebook's name servers responding.
- Queries on COM/NET name servers increased from 7,000 qps to 900,000 qps.
- Increase factor: **128x**

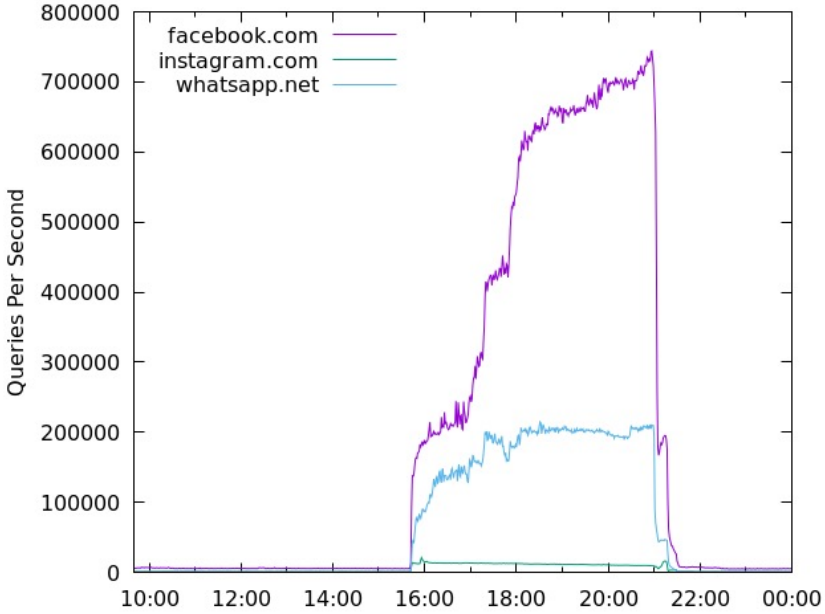

"Observations on Resolver Behavior During DNS Outages"
https://blog.verisign.com/security/facebook-dns-outage/

Botnet Research, March 2021

- Botnet domains delegated to sinkhole name servers.
- Intentionally configured one botnet domain to return SERVFAIL.
- Traffic increased from 50 qps to 60,000 qps for one botnet domain.
- Traffic increases also observed at TLD and Root name servers.
- Increase factor: **1200x**

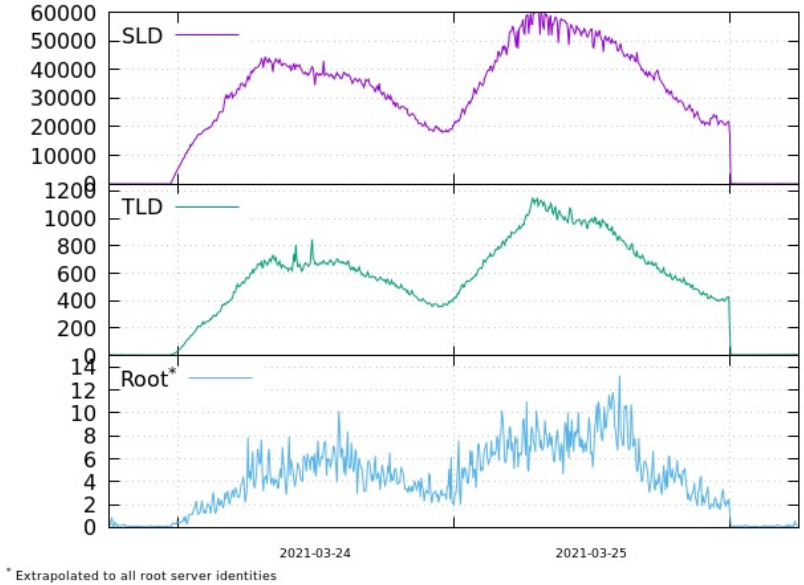

"Botnet Traffic Observed at Various Levels of the DNS Hierarchy," https://indico.dns-oarc.net/event/38/contributions/841/

KSK Post-Rollover Revocation, Jan 2019

- Pre-rollover: 15M/day
- Rollover: 75M/day
- Revocation peak: 1200M/day
- Increase factor: **80x**

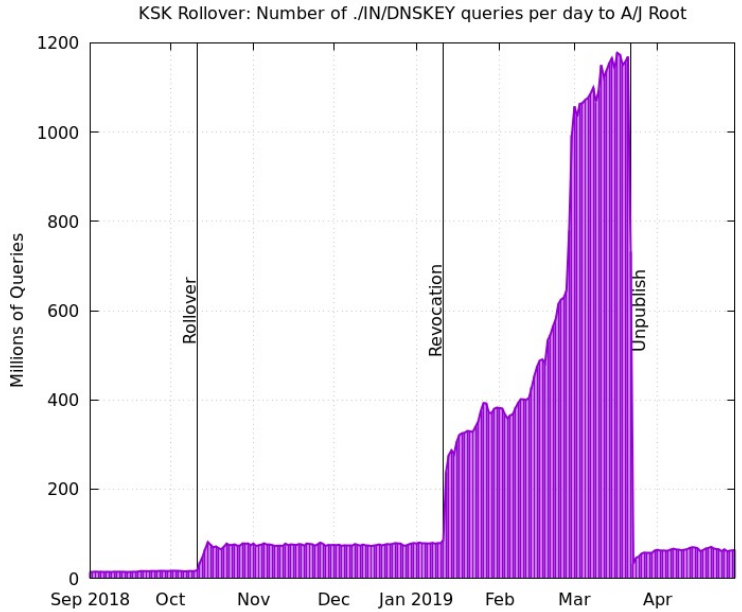

"Roll, Roll, Roll Your Root",
https://dl.acm.org/doi/10.1145/3355369.3355570

"DNSKEY Flood what does that tell us about resolvers",
https://indico.dns-oarc.net/event/31/contributions/686/

TsuNAME, March 2020

- Moura, et. al. found cyclic delegation dependencies.

```
example.com NS ns.example.nl

example.nl  NS ns.example.com
```

- Increase factor: **500x**

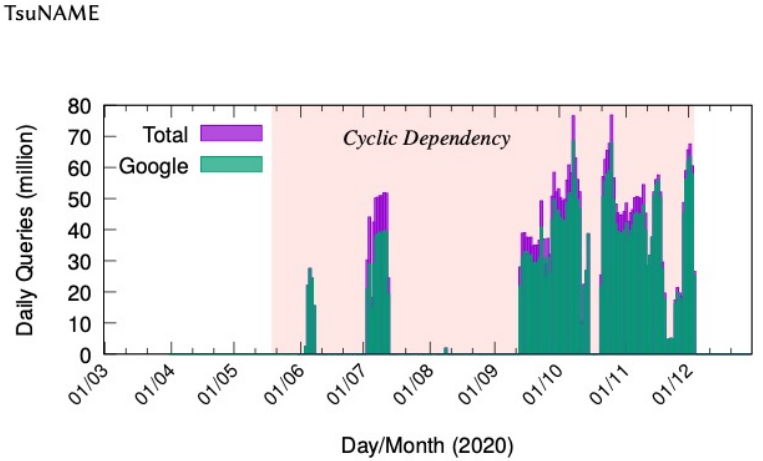

Figure 13: Query timeseries for .nl domain with cyclic dependency found with CycleHunter

"TsuNAME: exploiting misconfiguration and vulnerability to DDoS DNS"
https://dl.acm.org/doi/10.1145/3487552.3487824

More

- Survey Error Amplification, V. Dukhovni, October 2019
  - https://indico.dns-oarc.net/event/32/contributions/747/
- 2016 Dyn Attack: "At this point we are now experiencing botnet attack traffic and what is best classified as a 'retry storm'. Looking at certain large recursive platforms > 10x normal volume."
  - https://ccnso.icann.org/sites/default/files/file/field-file-attach/2017-04/presentation-oracle-dyn-ddos-dns-13mar17-en.pdf
- Roll Over and Die, December 2009
  - https://www.potaroo.net/ispcol/2010-02/rollover.html
- NXNSAttack, August 2020
  - https://www.usenix.org/conference/usenixsecurity20/presentation/afek

Existing Protocol Requirements

RFCs 2308, 4697, and 8767

RFC 2308 - Negative Caching of DNS Queries

- “negative caching should no longer be seen as an optional part of a DNS resolver.”
- Mostly focuses on NXDomain and NoData responses.
- Caching Server Failure is optional.
- Caching Dead / Unreachable servers is optional.

BCP 123 - Observed DNS Resolution Misbehavior

- “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.”
- “Iterative resolvers SHOULD cache name servers that they discover are not authoritative for zones delegated to them (i.e., lame servers).”
- “Iterative resolvers SHOULD take care to avoid sending queries at excessive rates. Implementations SHOULD support throttling logic to detect when queries are sent but no responses are received.”

RFC 8767 - Serving Stale Data to Improve DNS Resiliency

- "Attempts to refresh from non-responsive or otherwise failing authoritative nameservers are recommended to be done no more frequently than every 30 seconds."

Updating Requirements

draft-dwmtwc-dnsop-caching-resolution-failures

Types of Resolution Failures

- Server failures
- Refused response code
- Timeouts
- Delegation loops
- Alias loops and long chains
- DNSSEC validation failures
- Other?

Scope

- Resolution failures MUST be cached against the specific query tuple <query name, type, class, server IP address>.
- If all known servers for a query tuple <query name, type, class> result in resolution failure, the resolver MUST NOT send further queries for the tuple until the corresponding negative cache entries expire.

Retries and Timeouts

- A resolver MUST NOT retry more than twice (i.e., three queries in total) before considering a server unresponsive.
- No requirements placed on timeout values, which may be implementation- or configuration-dependent. It is generally expected that typical timeout values range from 3 to 30 seconds.

Negative Caching TTLs

- Resolvers MUST cache resolution failures for at least 5 seconds.
- Resolvers SHOULD employ an exponential backoff algorithm to increase the amount of time for subsequent resolution failures.
- Consistent with RFC 2308, resolution failures MUST NOT be cached for longer than 5 minutes.

Requerying Delegation Information

- Reiterate BCP 123 requirement about querying parent zones.
- Update this requirement with better language?
- For resolution failures in a zone, don't bother the parent zone more than once per X time interval?

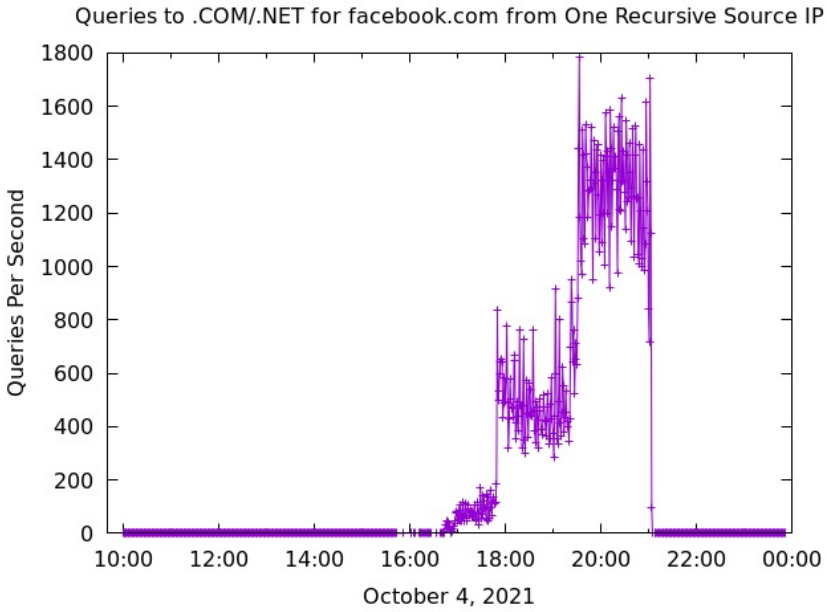

Please read and comment on the draft

draft-dwmtwc-dnsop-caching-resolution-failures

powered by

VERISIGN®