VERISIGN

Observations on Botnet Traffic at Various Levels of the DNS Hierarchy

Duane Wessels and Matt Thomas

Botnet Query Traffic

- This particular botnet generates DNS queries for thousands of second-level domains
  - Large subset in .COM and .NET
- Consistent and easily identifiable third-level labels
- Example
  - xxx.yyyyyyyy.$tld
- Very even distribution of queries coming from recursive clients
- Query type: 99% A, and 0.5% AAAA
- We don’t share the botnet’s name here because remediation is still ongoing at other registries

Before Sinkholing

Botnet Query Traffic

- Generated about 375,000 queries per second (peak) to Verisign's .COM and .NET name servers
- Most botnet domains hadn't been registered, so most responses are NXDOMAIN
- A small number of domains were previously registered and delegated.
- Very even distribution of queries among botnet domains
  - i.e. all domains receive essentially the same rate of queries

Pre-delegation Traffic Levels

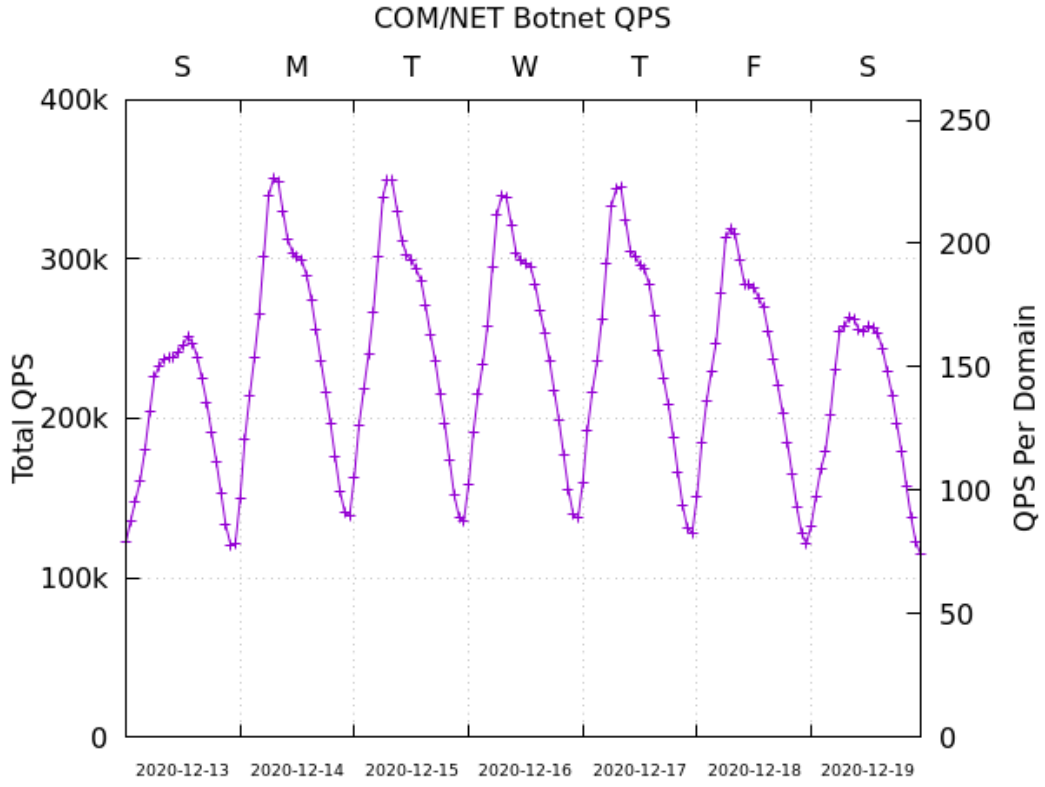

Typical Response

```
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NXDOMAIN, id: 63054
;; flags: qr aa; QUERY: 1, ANSWER: 0, AUTHORITY: 1, ADDITIONAL: 1


;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 4096
;; QUESTION SECTION:
;***.*********.com. IN A


;; AUTHORITY SECTION:
com. 900 IN SOA a.gtld-servers.net. nstld.verisign-grs.com. 1612827708 1800 900 604800
86400


;; Query time: 181 msec
;; SERVER: 192.33.14.30#53(192.33.14.30)
;; WHEN: Mon Feb 08 15:42:17 PST 2021
;; MSG SIZE  rcvd: 117
```

What Query Rate Might We Expect?

- About 71,000 unique three-label domain names
- With 900 sec negative caching TTL
- Worst case – every source querying every three-label name every 900 seconds – that would average about 79 queries per second per source
- In the graph, each dot is one query source (IP address)
- 99% of clients are below 79 QPS (vertical line)

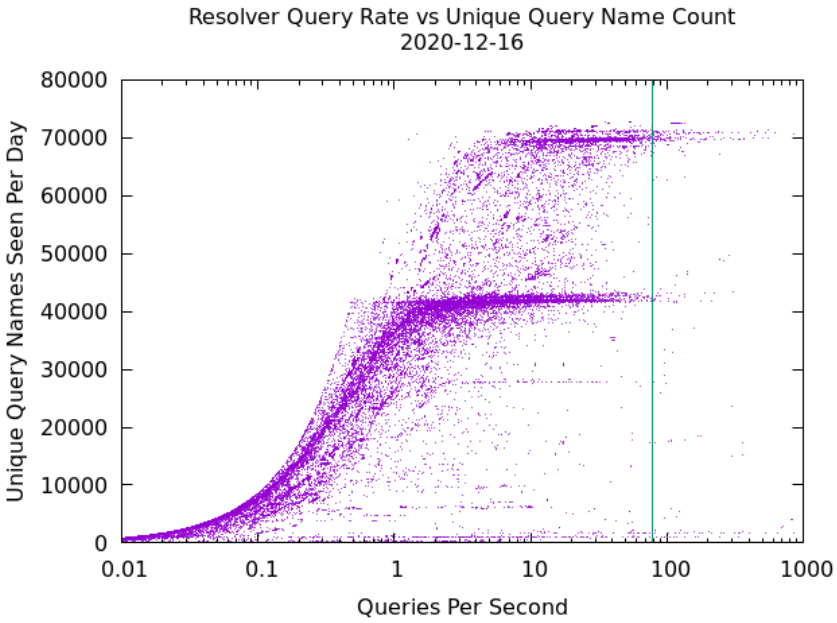

The Sinkhole

Domains Names Delegated to Sinkhole Name Servers

- Most domain names delegated to sinkhole
- A control group was left un-delegated
- Four dedicated sinkhole nameservers at three geographically distributed sites.
- 16 processor cores, 640GB Storage, 32GB RAM
- Ubuntu BIND 9.16.1
  - Configured with minimal responses
- DNS Statistics Collector[1] for data collection

[1] https://www.dns-oarc.net/tools/dsc

Before and After

- After sinkholing, TLD query rates dropped by a factor of 35-40
  - From 225 to 6 QPS/domain
- On the sinkhole name servers, query are under 50 QPS/domain

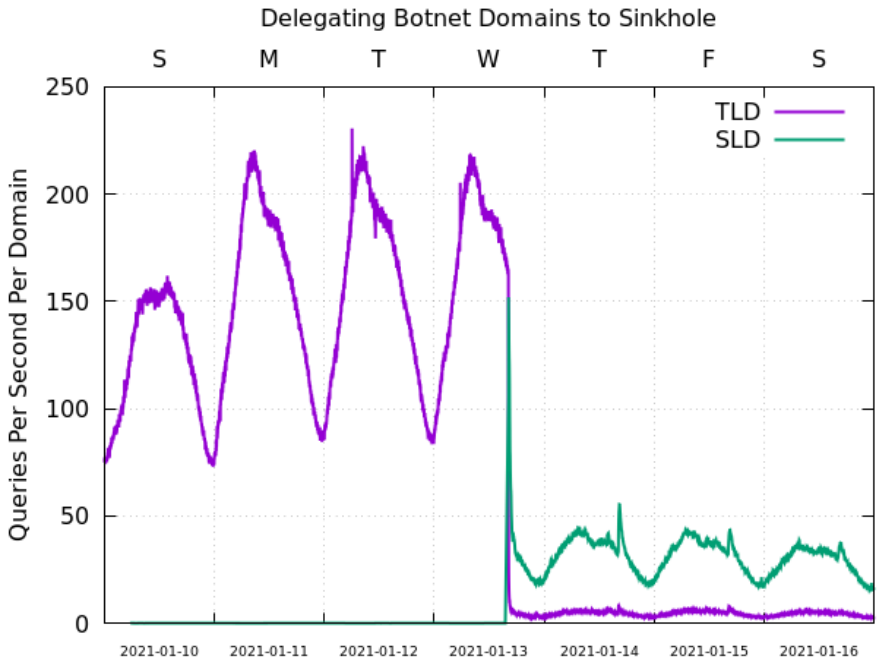

Some Experiments

Experiments

- The nature of this traffic affords us an opportunity to conduct some interesting research
- How does response code affect caching and query rates?
- How does the presence of requested data affect caching and query rates?
- How do TTLs affect caching and query rates?
- Use weekends to reset and change experiments

Types of Responses Studied

Description	RCODE	RRs in Answer?	TTL?
NOERROR / DATA	0	Yes	From Data
NOERROR / NODATA	0	No	From SOA
NXDOMAIN	3	No	From SOA
SERVFAIL	2	No	No
REFUSED	5	No	No

Different Types of Responses

- Responses with:
  - NOERROR / DATA
  - NOERROR / NODATA
  - NXDOMAIN
- All TTLs set to 24 hours
- NODATA and NXDOMAIN have identical query rates

https://imgur.com/gallery/ED6qKJv

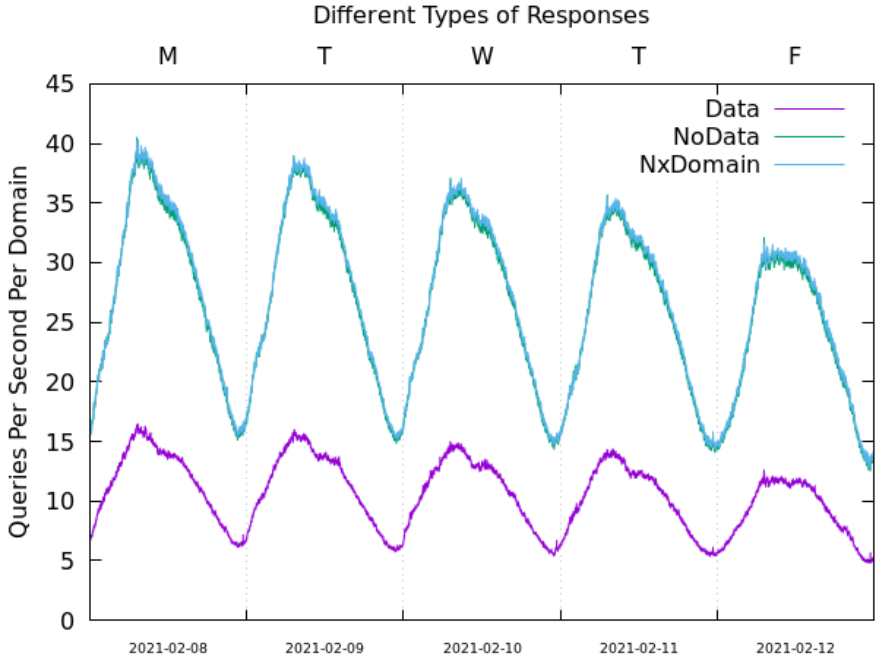

NOERROR / DATA: TTL vs Query Rate

- The second-level authoritative zones are configured to return A and AAAA records for non-apex names
- Different domain name groups are given different TTLs
- Shaded area represents daily range

TTL	SLD QPS
1 min	280
1 hour	68
1 day	34

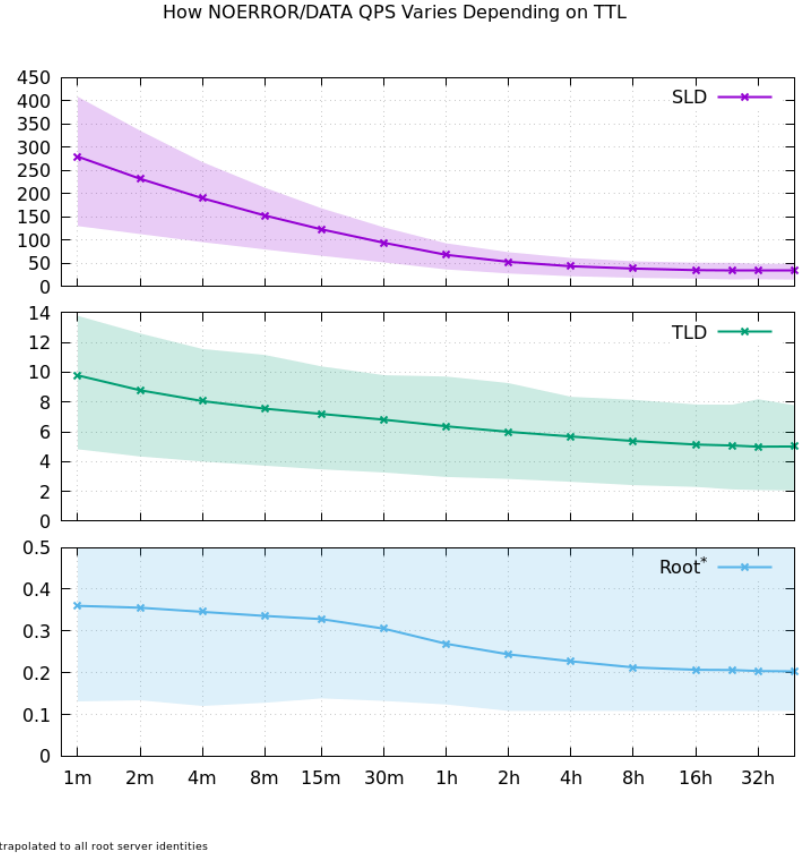

* Extrapolated to all root server identities

NOERROR / NODATA: TTL vs Query Rate

- The second-level authoritative zones are configured to return NODATA for non-apex names
- Different domain name groups are given different TTLs
- Shaded area represents daily range

TTL	SLD QPS
1 min	295
1 hour	96
1 day	86

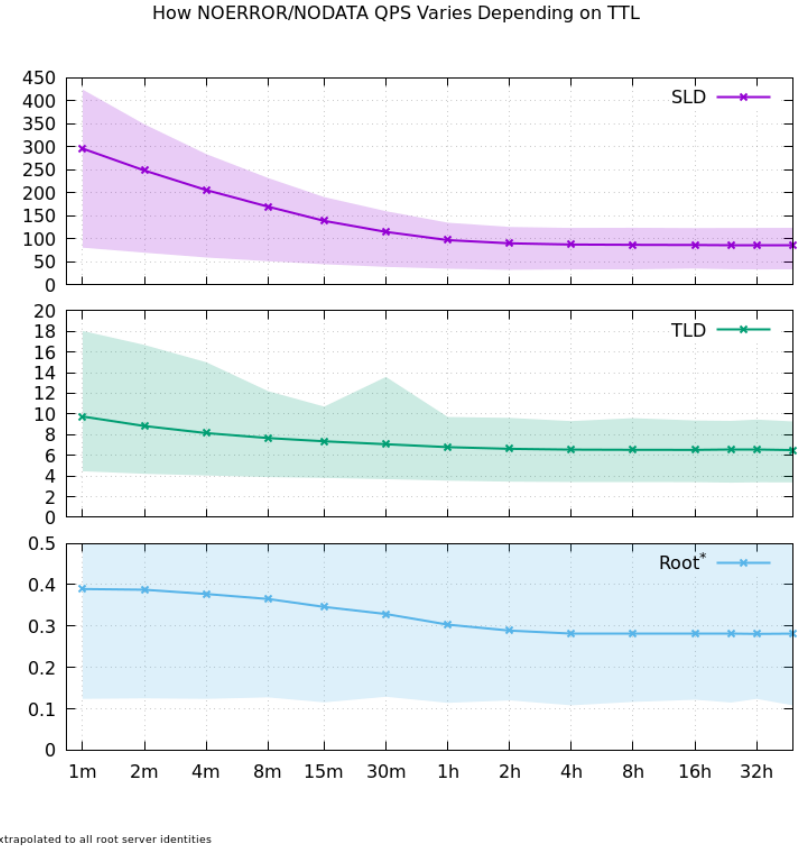

* Extrapolated to all root server identities

NXDOMAIN: TTL vs Query Rate

- The second-level authoritative zones are configured to return NXDOMAIN for non-apex names
- Different domain name groups are given different TTLs
- Shaded area represents daily range

TTL	SLD QPS
1 min	286
1 hour	92
1 day	82

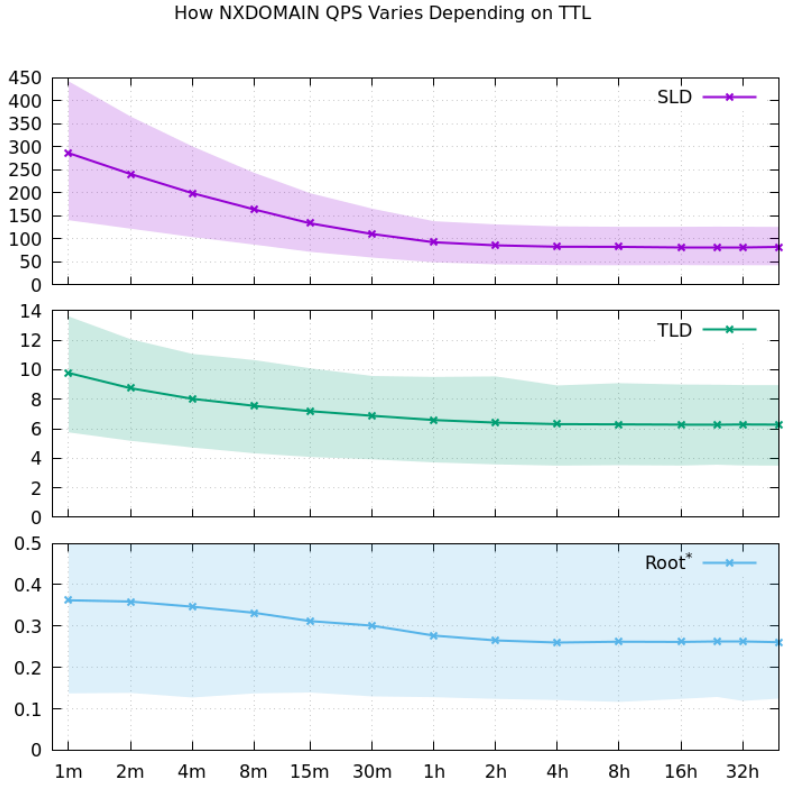

* Extrapolated to all root server identities

Comparison of Response Types With TTLs

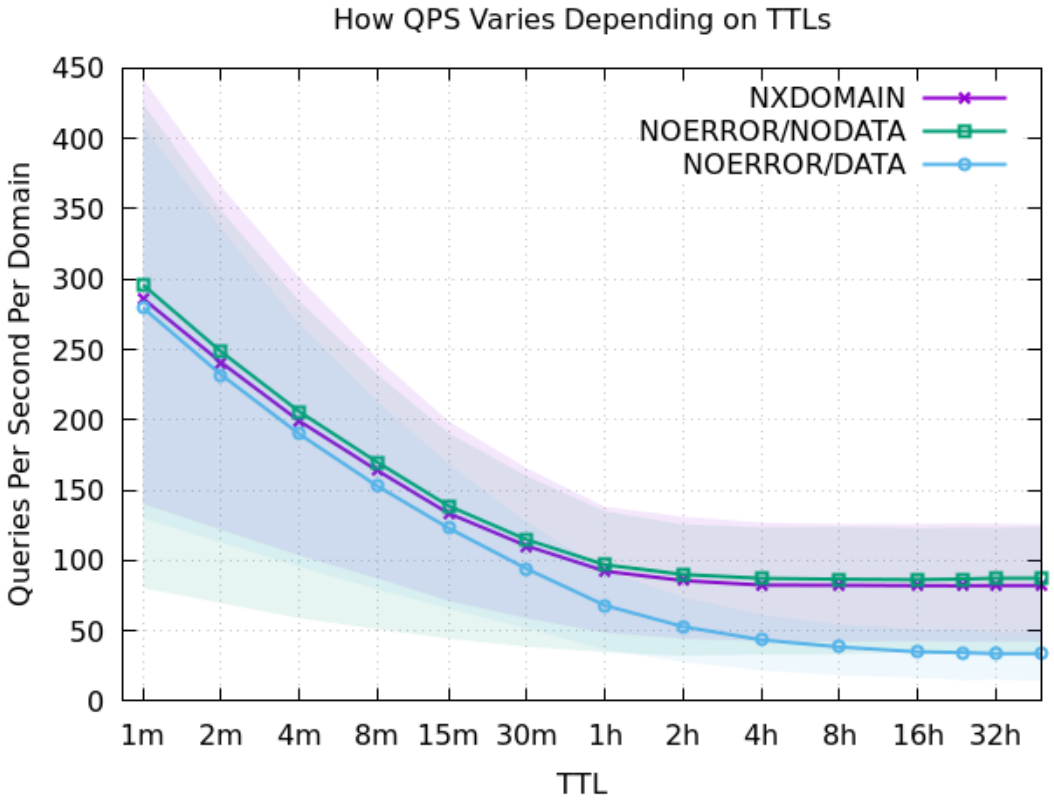

REFUSED

- Sinkhole name servers configured to return REFUSED for very small subset of domain names
- This is QPS **per domain name**
- Further investigation shows this is largely due to a small number of recursive operators.

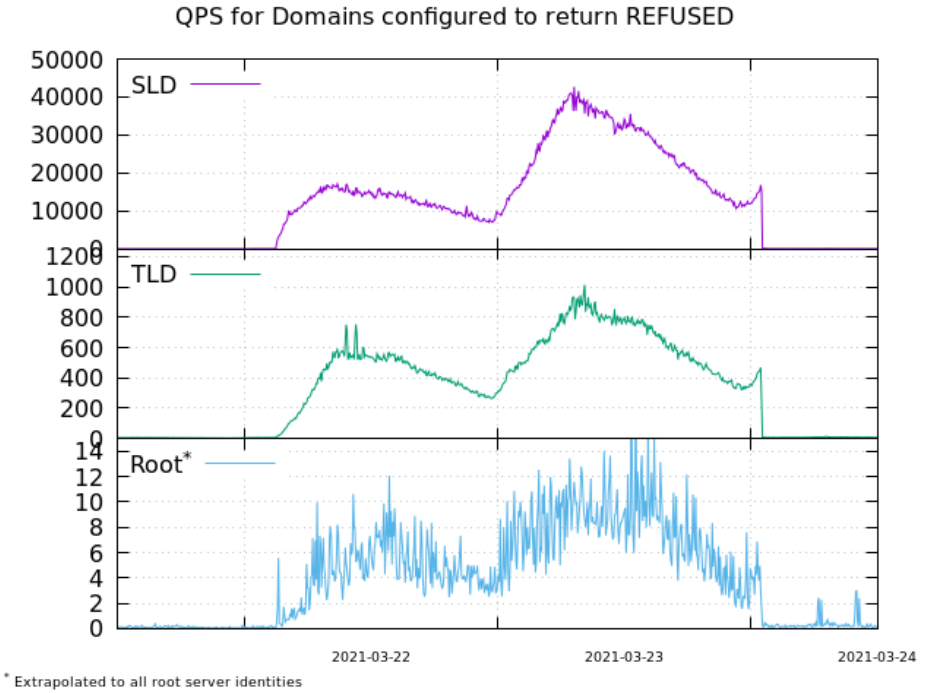

SERVFAIL

- Sinkhole name servers configured to return SERVFAIL for very small subset of domain names
- This is QPS **per domain name**
- Further investigation shows this is largely due to a small number of recursive operators.

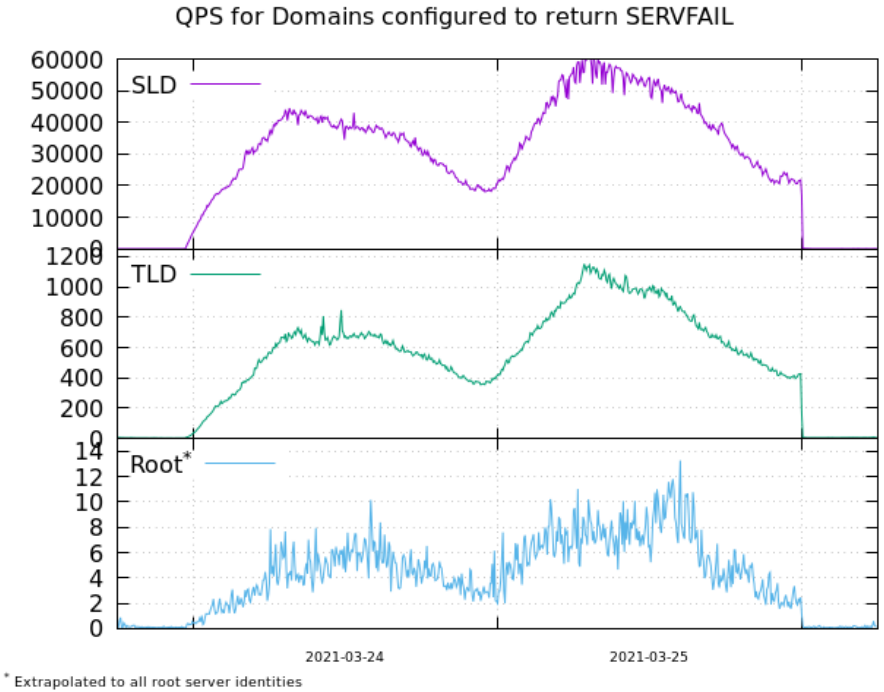

Effect of ‘minimal-responses’ on Bandwidth

- With minimal-responses on, responses are only slightly larger than queries
  - Outgoing bandwidth slightly higher than incoming
- For two days we turned off minimal responses
- Response size and outgoing bandwidth approximately doubled

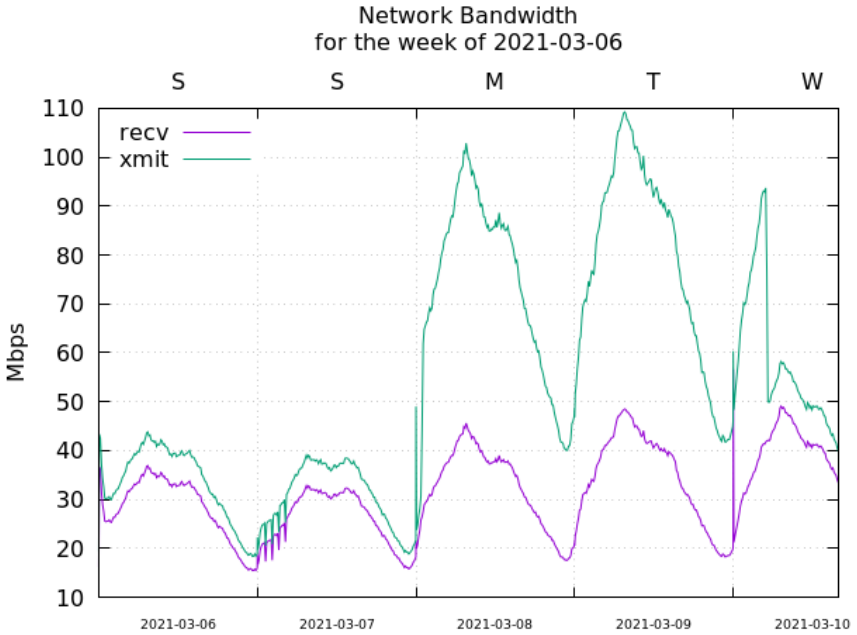

Key Observations

- No significant difference in caching of NXDOMAIN and NODATA negative responses
- For negative responses, TTLs longer than 1 hour don’t make much difference
- For positive responses, TTLs longer than 2-4 hours don’t make much difference
- REFUSED and SERVFAIL significantly amplify, rather than reduce query rates

Possible Future Experiments

- Measure effects of aggressive DNSSEC caching
- How quickly do changes to parent/TLD zone propagate?
- Differences between child / parent NS RRset

powered by

VERISIGN™