What part of "NO" is so hard to understand?

Geoff Huston
APNIC Labs

We were looking elsewhere…

- We were setting up a measurement experiment that was looking at the extent of support for aggressive NSEC caching (RFC8198) in the DNS
- The experiment setup involved presenting to the user a DNS name that did not exist from a signed zone, so that we would pass an NSEC record to a DNSSEC-aware resolver
- But was was intriguing was that we were seeing many more queries for the non-existent name than we had expected

What we saw:

- We used an online ad to get users to query for a **unique** non-existent DNS name
- And then counted the number of queries we saw for these names

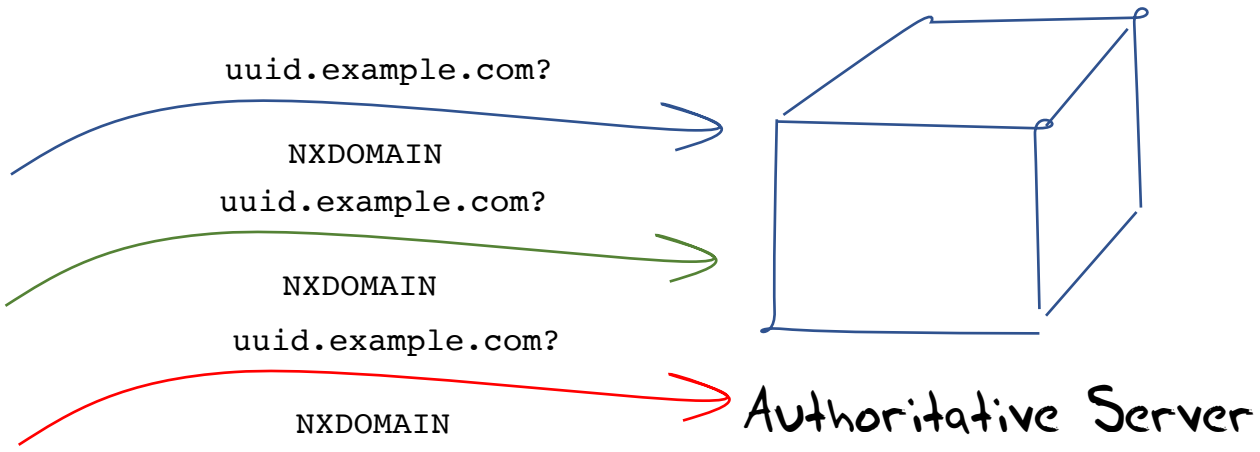

What we saw:

- We used an online ad to get users to query for a **unique** non-existent DNS name
- And then counted the number of queries we saw for these names

Queried Names: 60,210,983

DNS Queries: 142,631,272

- That's an average of 2.37 queries per non-existent name!

Why so many queries?

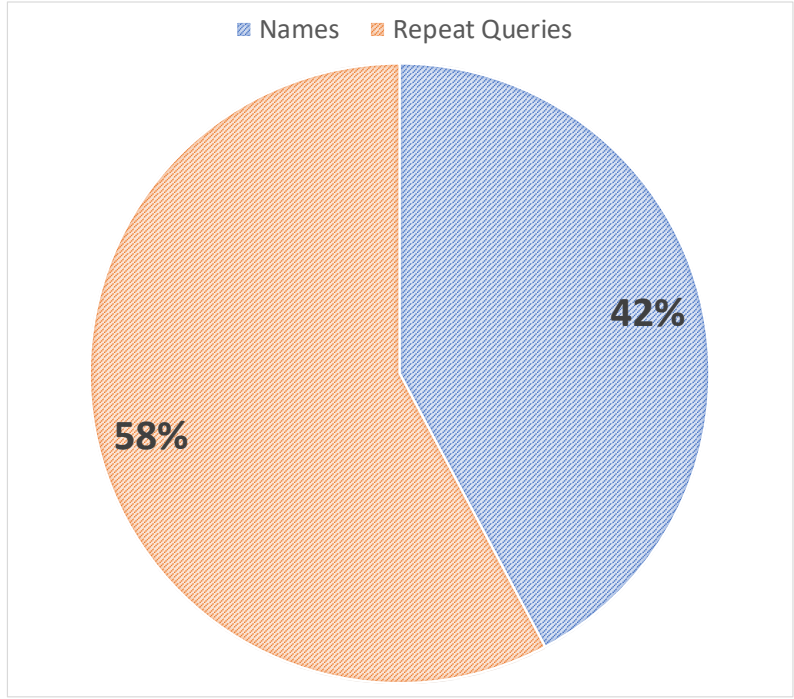

Expectations

- If “NO means NO” then naïvely we would expect to see 1 query per name, not 2.37 queries per name
- But maybe that’s just too naïve these days…
- After all - the DNS is now SO clever that this just couldn’t be due to random DNS insanity – could it?

Happy Eyeballs and the DNS

A ‘happy eyeballs” dual stack client will launch 2 DNS queries back-to-back (roughly), for A and AAAA records of the name

- 23% of clients asked both A and AAAA records
- 3% asked only AAAA records (*)
- 74% asked only for A records

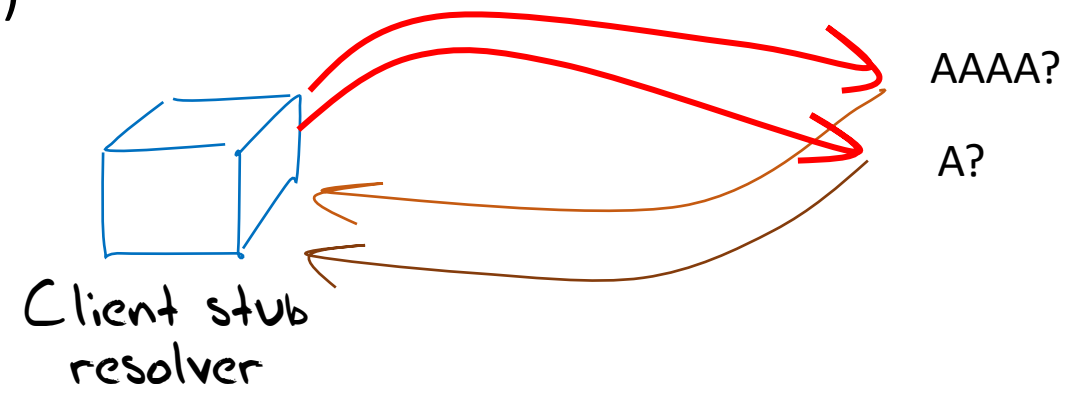

* If the client asks for an AAAA record and waits for a response before asking for the A record then the NXDOMAIN response will stop the connection process and any subsequent A query will not be performed

Factoring Happy Eyeballs

- If we split out the A and AAAA queries the experiment launched 73,537,852 DNS resolution ‘events’
- We saw 142,631,272 DNS queries, or an average of **1.93** queries per name

That's better, but still unexpectedly high

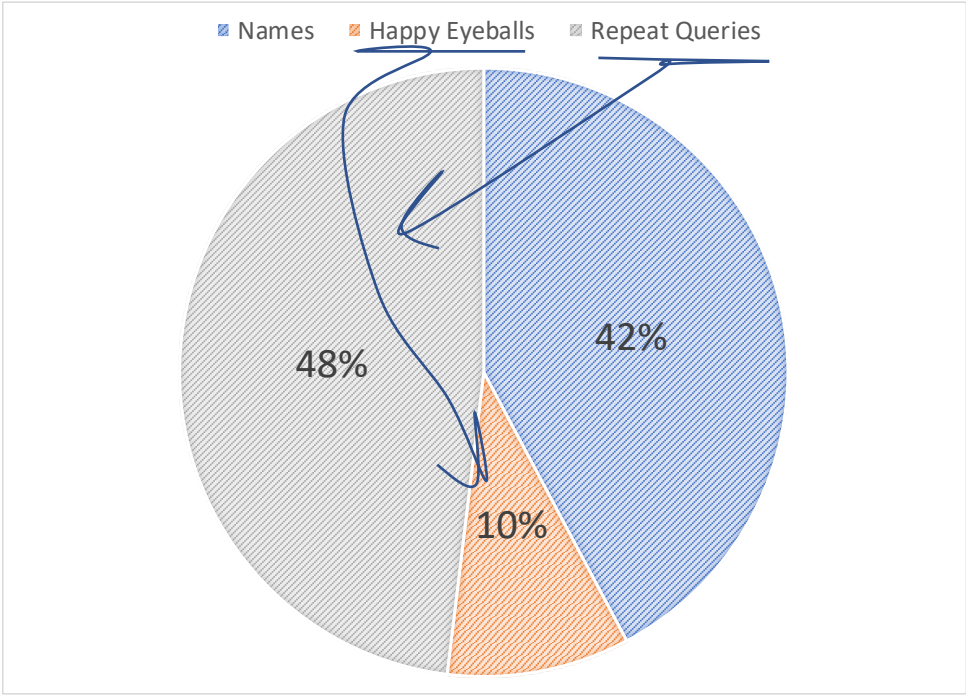

Single vs Multiple Queries

- 36,059,484 resolution ‘events’ were completed with just 1 query to the authoritative server (49% of all resolution events)
- If there were multiple queries for a name (>= 2 queries), then the average of the multiple queries was 2.84 queries

Distribution of Queries

- Are averages misleading here?
- Is this a generic issue of re-queries across a large set of queries?
- Or a small number of queries that are the subject of a total frenzy of re-queries?

Re-query Distribution

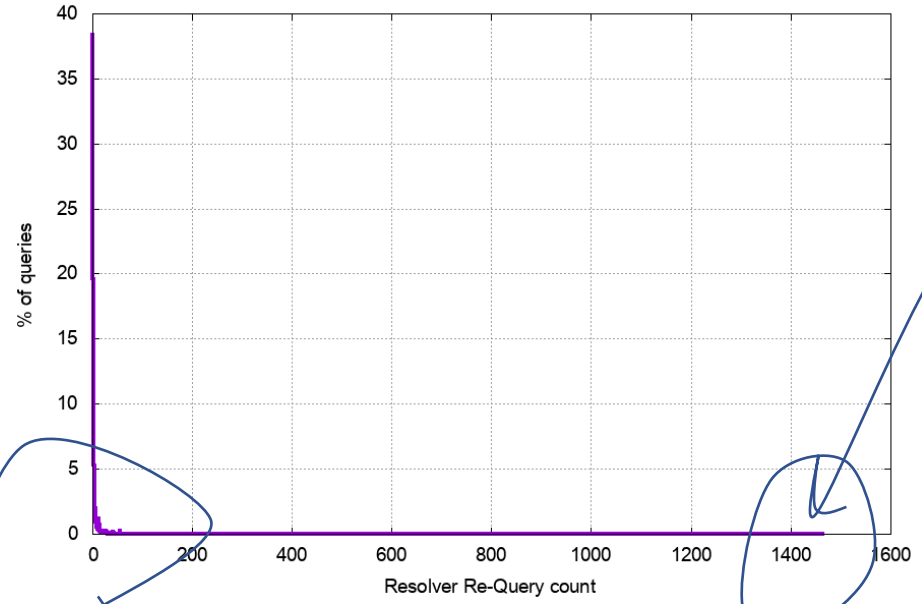

!! Seriously????

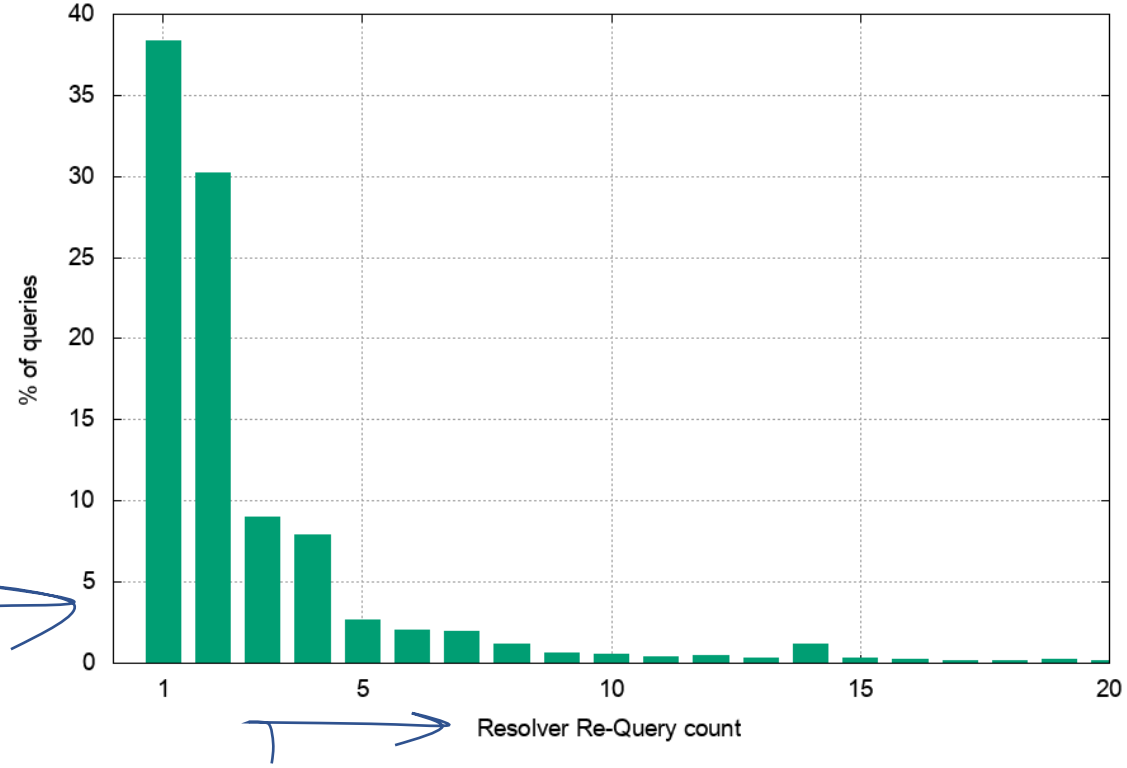

32% of queries were parts of query sequences of 3 or more

Does UDP suck THAT much?

- Why is the total re-query rate at 51% of tests?
- Surely DNS over UDP is not THAT bad
  - The servers are responding to every query
  - The signed response is 603 bytes in size
  - We are using a distributed setup of servers to localize DNS transactions
  - So why are the servers seeing 51% of tests generate 2 or more queries?

Re-Query Time Intervals

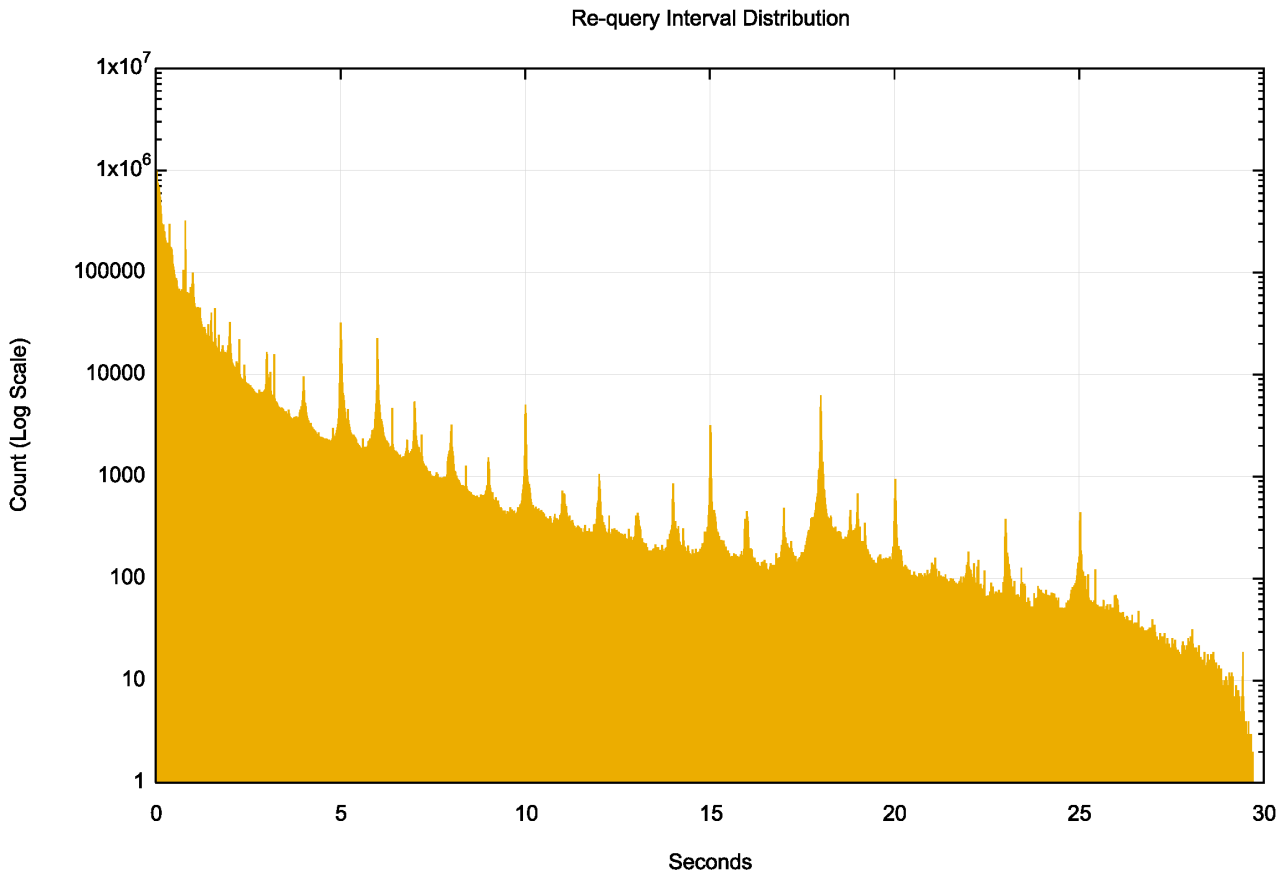

There are strong local peaks at regular 1 second intervals – this would appear to be an end host / stub resolver re-query behaviour

Re-Query Time Intervals

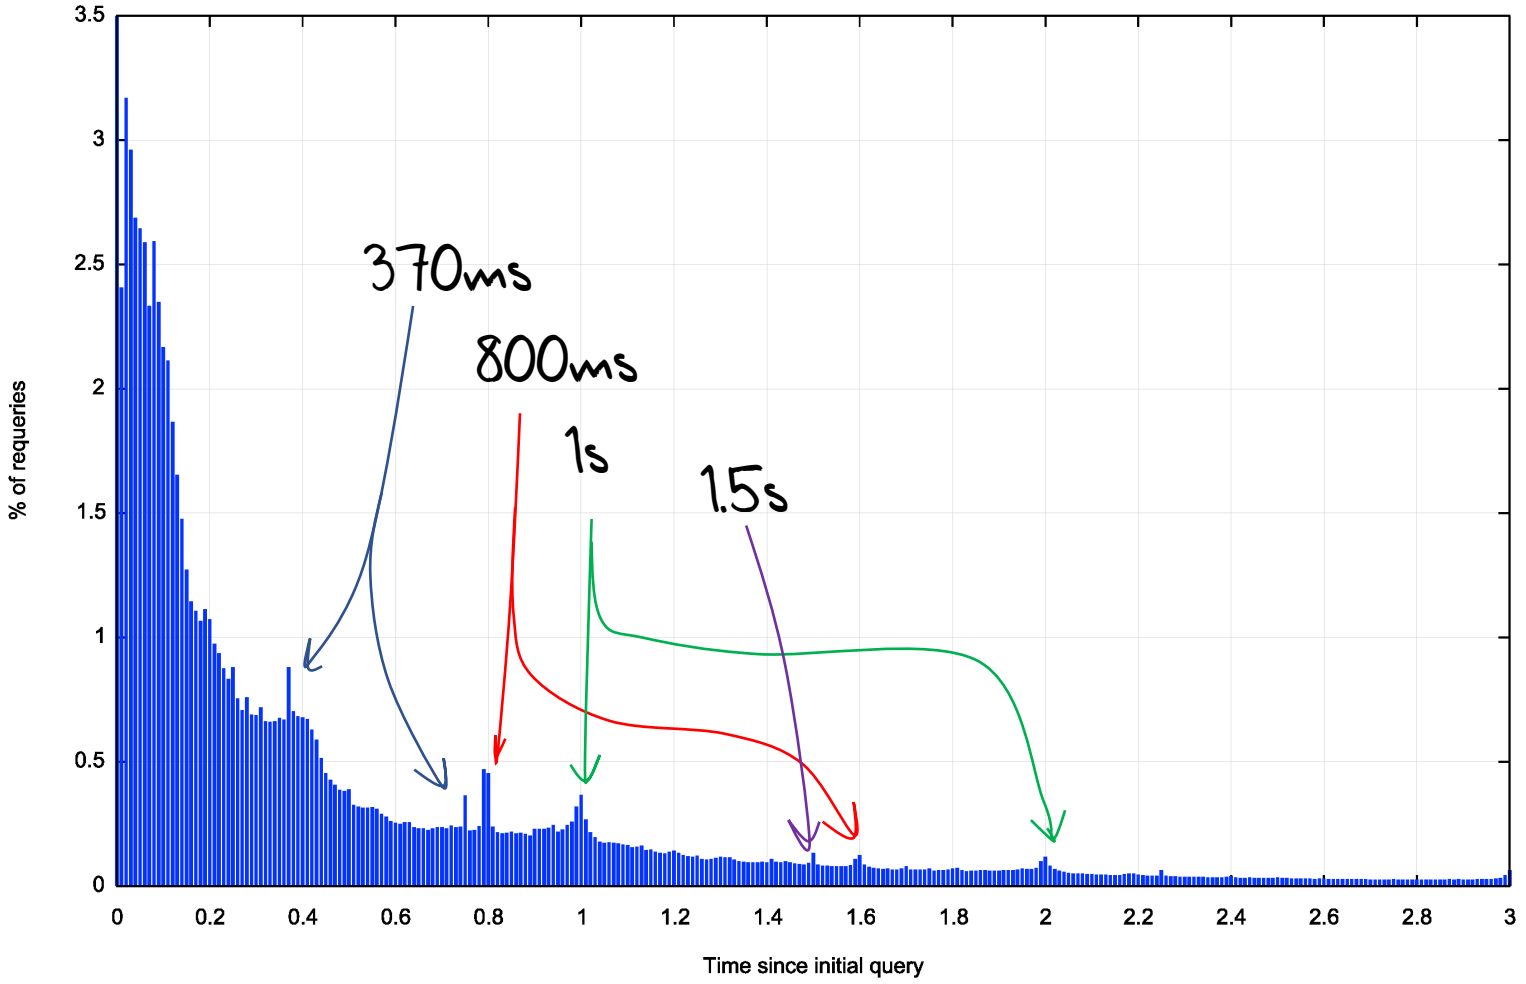

There are local re-query peaks here at 370 ms, 800ms, 1 sec and 1.5sec

It is likely that these time intervals represent stub and/or recursive resolver re-query timers

DNSSEC?

- This is a DNSSEC-signed non-existent name
  - i.e. the NXDOMAIN also has an NSEC record if DO is set in the query
- Is DNSSEC a factor in the excessive re-query volume?
  - i.e. is the additional time to validate causing requery timers to trigger?
- We added an unsigned non-existent name to the test set

Re-Queries per resolver IP addresss

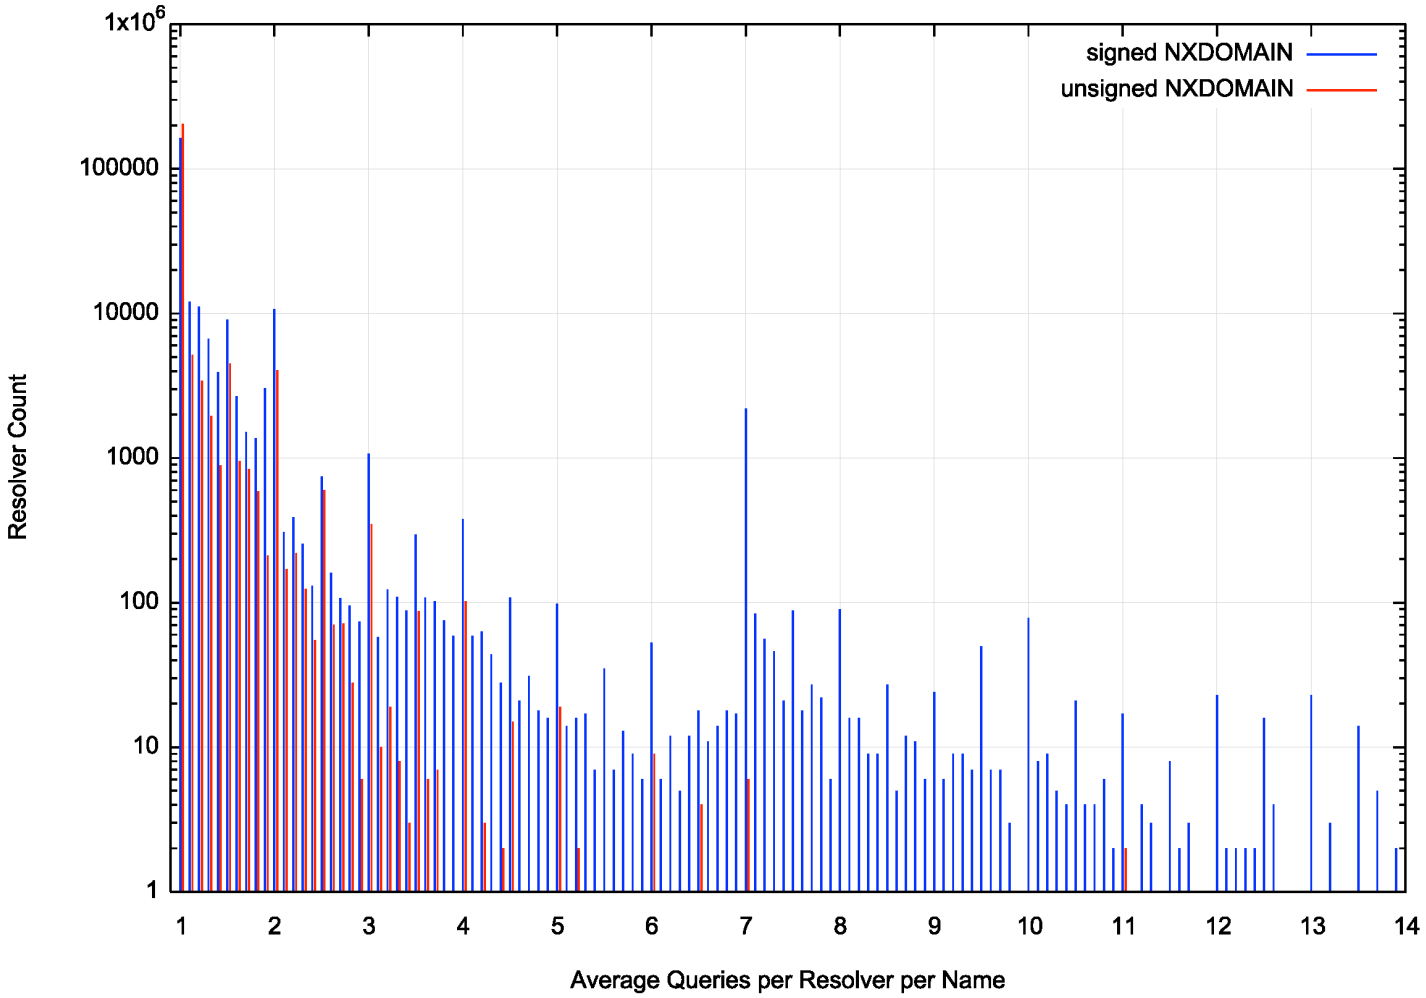

The log scale exaggerates the effect, but we observe that a DNSSEC-signed NXDOMAIN response generates a higher repeat query profile from the same resolver IP address

Signed vs Unsigned

	Signed	Unsigned
Experiments	65,686,452	69,251,349
A/AAAA	81,057,694	84,979,990
Queries	153,697,947	122,665,888
Single Query Exps	47,694,930	60,061,746
Ratio	59%	71%
Multi-Query Exps	33,092,764	24,918,244
Re-Query Rate	3.19	2.51

Split out the 'happy eyeballs' factor

DNSSEC validation adds delay, and in around **12%** of cases this additional delay causes the resolver system to re-query the name

DNSSEC!

- About 12% of cases of re-query for non-existent names appear to be related to recursive resolver’s DNSSEC validation of NSEC records

stupid Resolver "farms"

We also see query patterns of the form:

Resolver	Query Time
7x.xxx.0.178	0.752
6z.zzz.161.146	0.865
7x.xxx.0.230	0.980
6z.zzz.161.220	1.094
7x.xxx.0.188	1.201
6z.zzz.161.182	1.319
7x.xxx.0.180	1.430
6z.zzz.161.144	1.542
7x.xxx.0.226	1.650
7x.xxx.0.138	1.654
6z.zzz.161.134	1.762
6z.zzz.161.222	1.775

It appears that some resolver farms operate by farming the query across all members of the farm. This pattern seen here shows two such cases where different IP addresses in the same subnet repeat the initial query at approximately 100ms intervals

How common is this form of subnet-based query repetition?

Re-query Profile

Re-queries	59,782,873
Same IP Address	17,848,729
Same Subnet	41,111,416

Wow!

Some 70÷ of the re-queries are from resolvers that share the same subnet prefix!!!

This points to some outstanding issues with resolver farm management

It may sound odd but…

- **Is NO worse than YES?**
- Is NXDOMAIN part of the issue here?
- Is the re-query rate lower if the name exists in the DNS?

NXDOMAIN vs A/AAAA re-queries

NXDOMAIN Signed:

41% of experiments generate multiple queries

39% of queries are re-queries (avg of 3.19 queries per experiment)

A/AAAA Signed:

18% of experiments generate multiple queries

13% of queries are re-queries (avg of 5.81 queries per experiment)

Fewer re-query events but each event has more queries!

NXDOMAIN Unsigned:

39% of experiments generate multiple queries

29% of queries are re-queries (avg of 2.51 queries per experiment)

A/AAAA Unsigned

38% of experiments generate multiple queries

36% of queries are re-queries (avg of 3.03 queries per experiment)

WTF?

- A DNSSEC-signed NXDOMAIN response generates many more re-queries than a DNSSEC-signed A / AAAA response

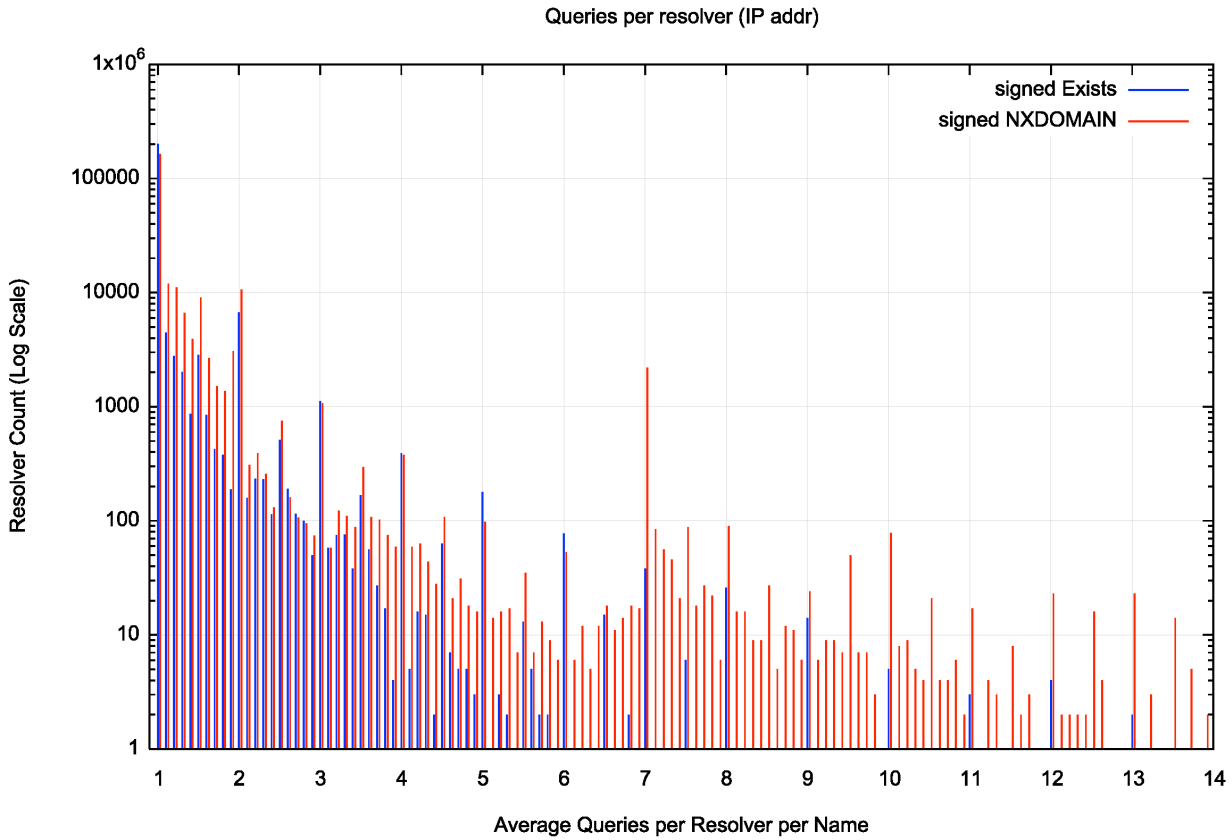

Pulling it back together

- Why are there so many DNS repeat queries?

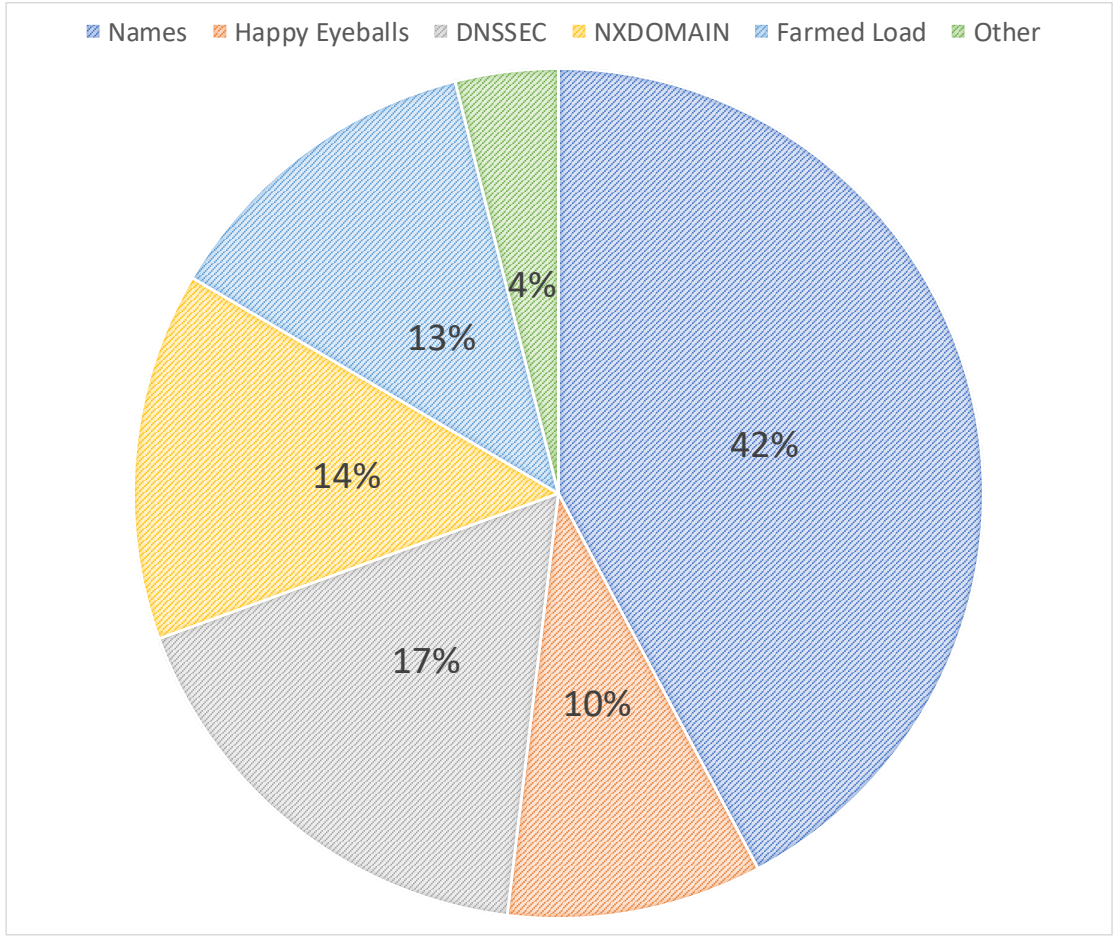

Pulling it back together

- Why are there so many DNS repeat queries?
  - Happy Eyeballs

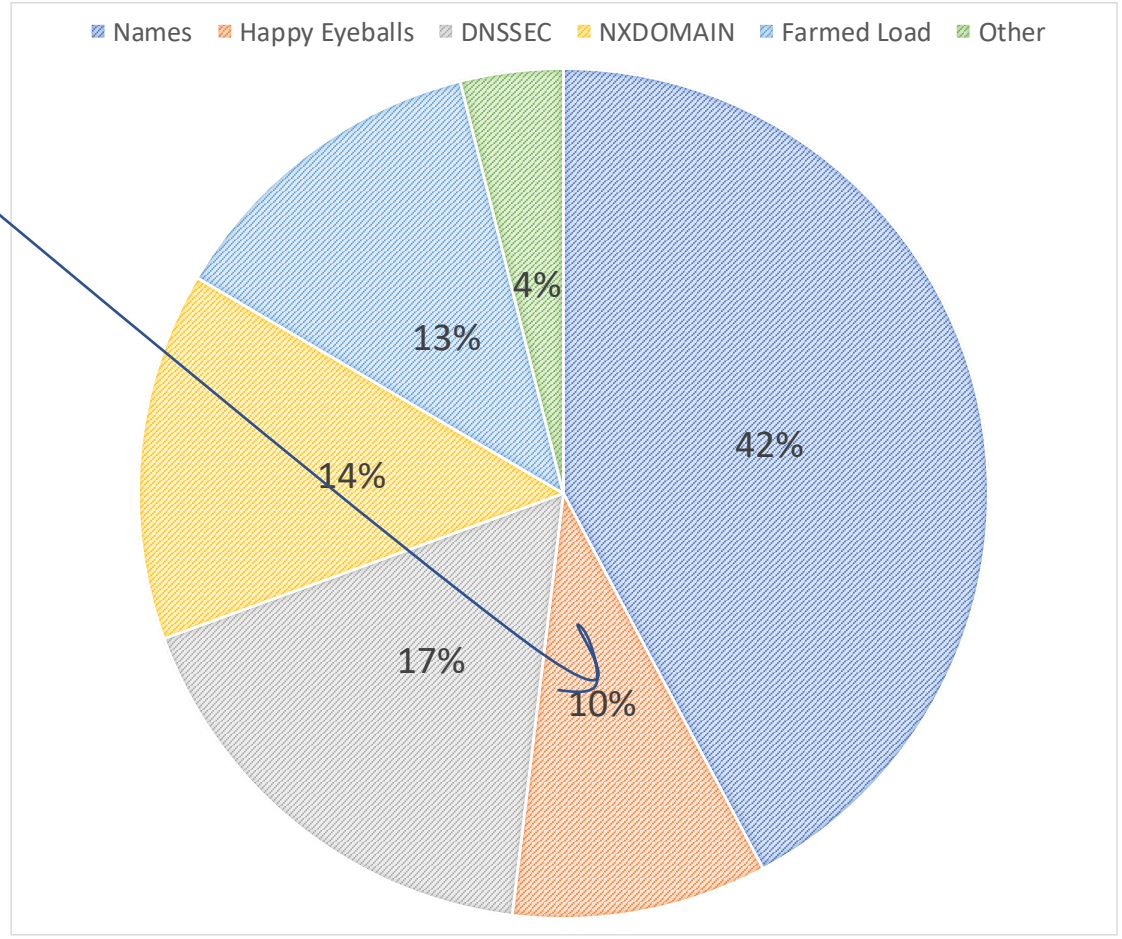

Pulling it back together

- Why are there so many DNS repeat queries?
  - Happy Eyeballs
  - DNSSEC

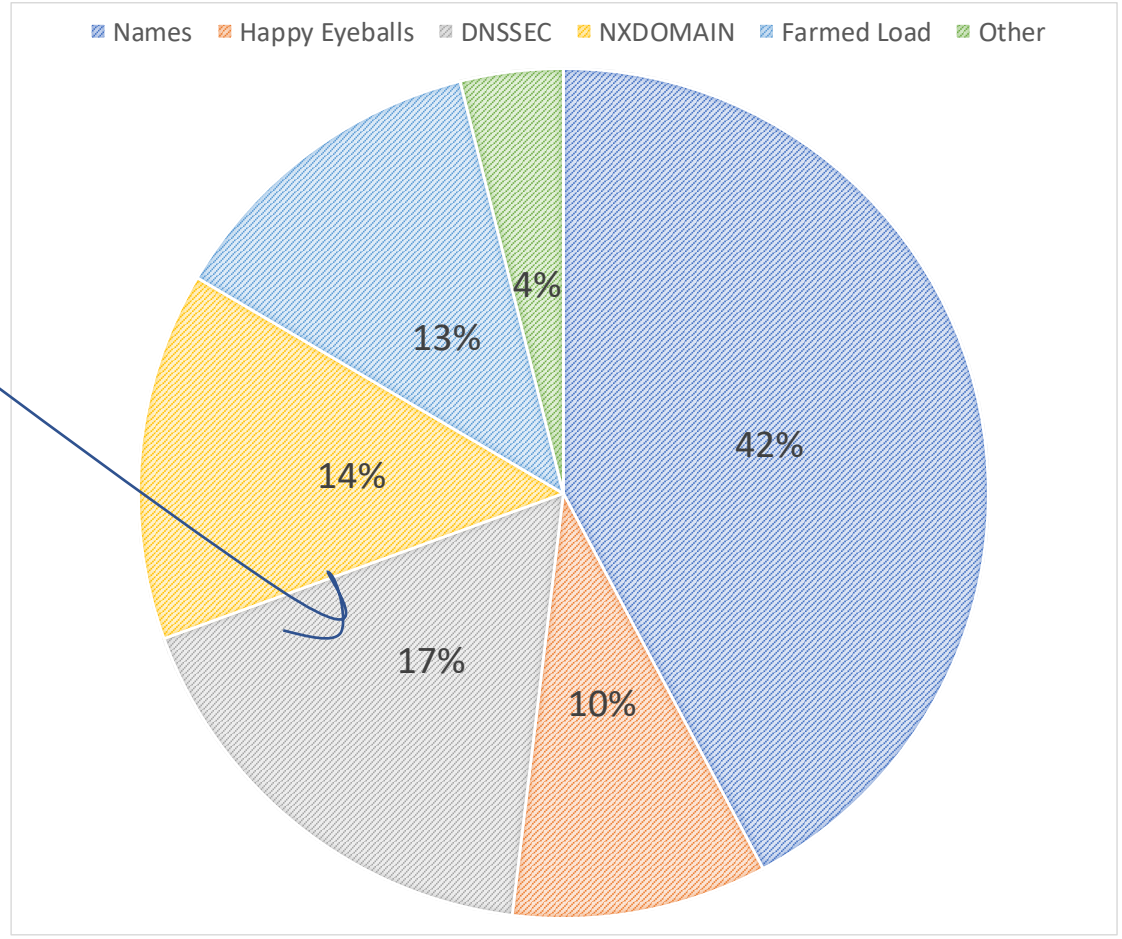

Pulling it back together

- Why are there so many DNS repeat queries?
  - Happy Eyeballs
  - DNSSEC
  - NXDOMAIN

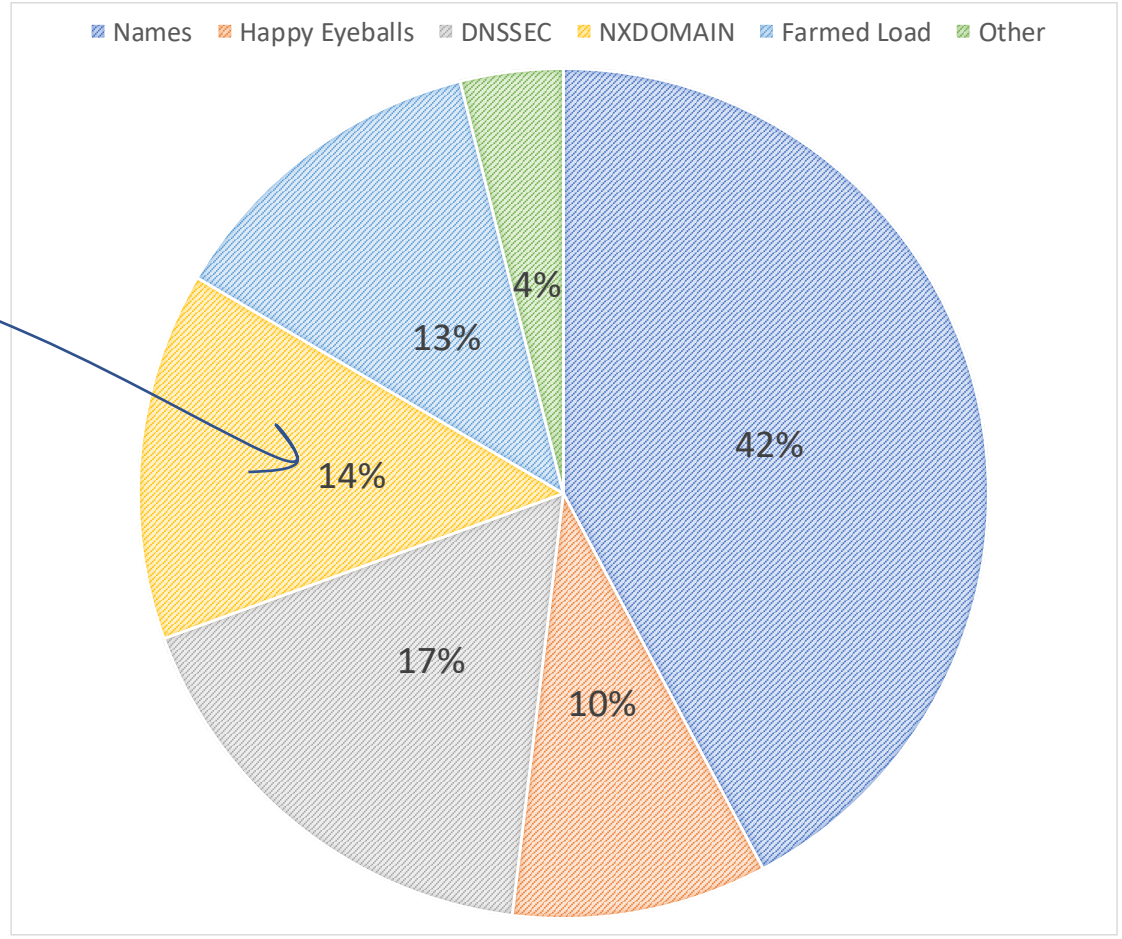

Pulling it back together

- Why are there so many DNS repeat queries?
  - Happy Eyeballs
  - DNSSEC
  - NXDOMAIN
  - stupidity down on the farm

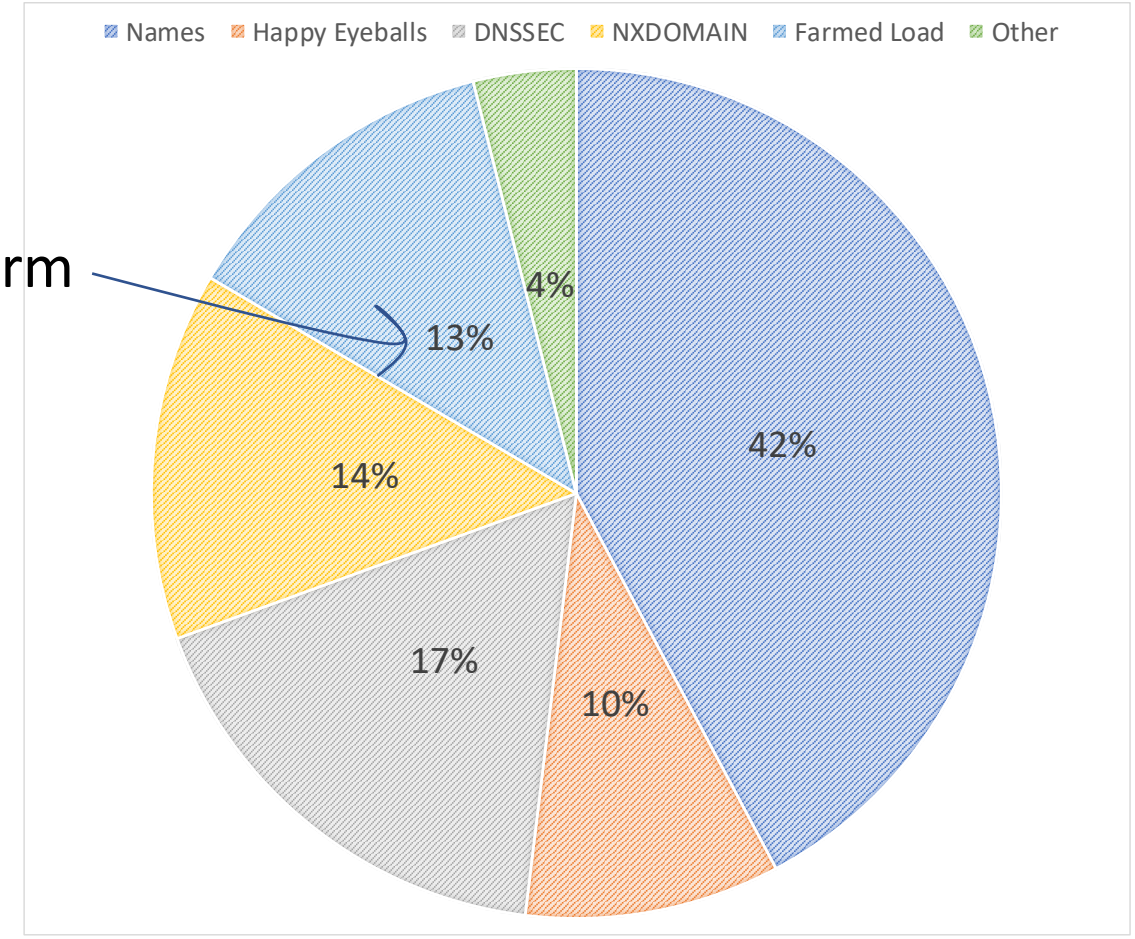

Pulling it back together

- Why are there so many DNS repeat queries?
  - Happy Eyeballs
  - DNSSEC
  - NXDOMAIN
  - Stupidity down on the farm
  - <reasons>

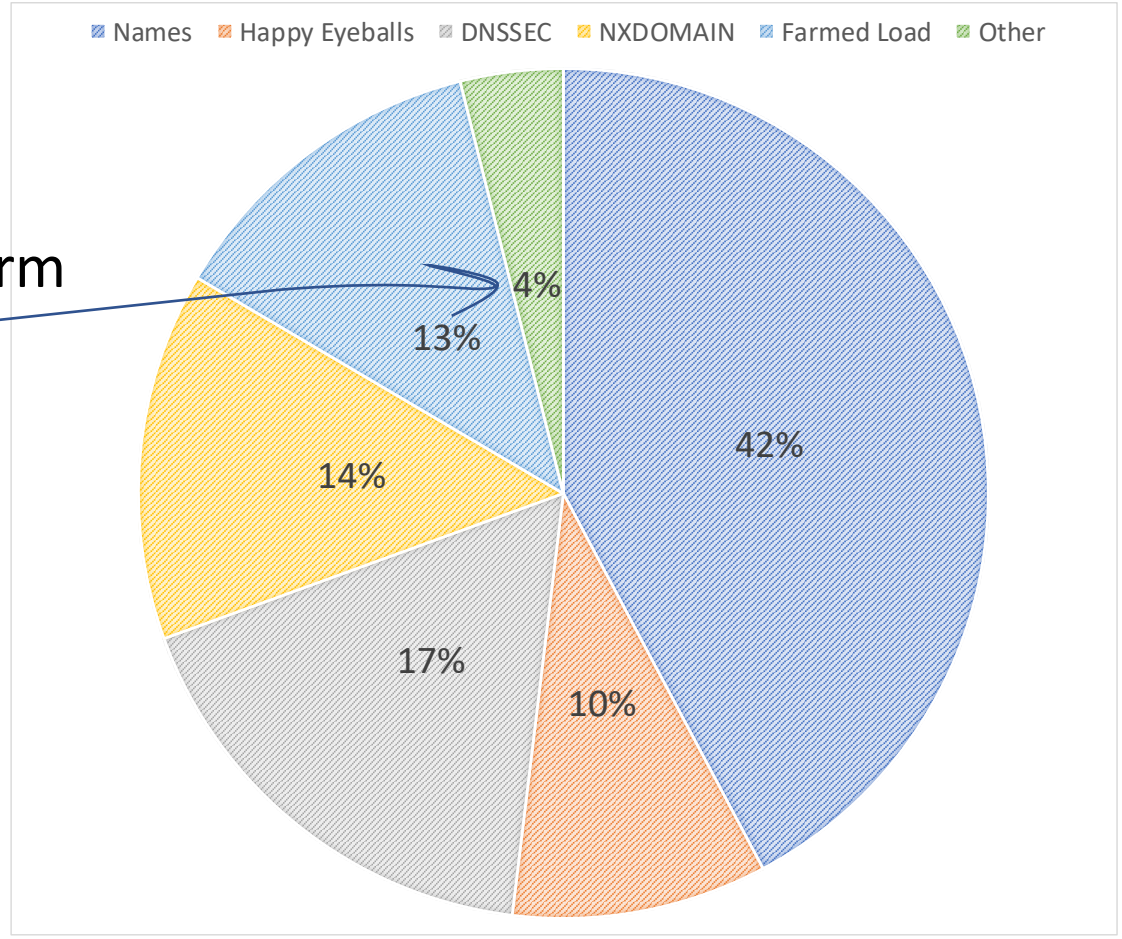

Thanks!