Testing Over 1000 gTLDs for EDNS0

Or *A Funny Thing Happened on the Way to the Testing Room*

Edward Lewis
ICANN

FOSDEM 2019/DNS Developer Room
3 February 2019

My Experience with the ISC EDNS Compliance Test Code

- 10,973,106,002 calls to DiG
  - That's over 10 Billion, almost 11 Billion

- Relax – that didn't really happen, that is a virtual count

Code Version Caveat

- The test code I used was sent to me by the author
  - No "version number"
  - A shell script to generate a report (genreport) and a host of supplemental files to produce graphs

- Towards the finish of my analysis I learned of code on github
  - Written in C, "more efficient", different results
  - https://gitlab.isc.org/isc-projects/DNS-Compliance-Testing
  - Didn't have time to push that into my test environment

Why?

- ICANN has contractual relationships with about 80% of the top-level domains
  - Part of this gives staff researchers access to data, like zone files
  - We have an interest in a secure, stable and reliable DNS and Internet
  - We don't have a interest is measuring one TLD against another

- The EDNS0 compliance tester, in the context of the DNS Flag Day presented an opportunity to run a tool over all the data we have
  - This was seen as a major undertaking
  - Our reporting on the data is not the usual though, no alerts of "breakage"

What comes out of this effort?

- Small insight into EDNS0 protocol compliance
  - But not very detailed at the scale of the testing

- The platform built to "eat zone files" and launch the test program can be used for other testing
  - We would like to improve the manageability of the DNS ecosystem via better measurement and analysis

- Related to a lesson from the KSK rollover project
  - The state of tooling for management is "in need of development"

The Workload

- From gTLD zone files plus root and arpa:
- Data-take date: 18 January 2019

Number of	Count
Root + TLD zone files	1,228
Delegations (NS sets)	193,825,454
NS resource records (not sets)	457,887,042
Glue resource records (all address)	3,255,827
IPv4 Glue records	3,198,649
IPv6 Glue records	57,178

Graphs

- You won't see any
- Maybe I'm just too lazy to graph
- They are (almost) all "long tail distributions" anyway (yawn!)
- There would be "thousands" of graphs, audience's attention would drown in them

Ground rule: On naming names

- I don't, except "root" and "arpa"
- This is about the protocol, not the industry
- One of the gTLDs represents half of the processing time and two-thirds of all the "virtual test cases"
- None of the delegations studied are in the ccTLD tree
- Only one delegation is related to the reverse map tree ("arpa")

- I've also changed IP addresses (which makes that data really dull)

Looking at names of nameservers

Number of	Count	Percent
Nameserver (names)	3,168,952	
IDN named nameservers	3,357	< 1 %
ASCII named nameservers	3,165,595	> 99 %
Nameserver in gTLDs	2,706,669	85 %
IDN-named, in-gTLD nameservers	3,298	< 1 %
Nameserver in ccTLDs	462,283	15 %
IDN-named, in-ccTLD nameservers	59	~0 %

Looking at glue records

Number of	Count
Glue resource records (all address)	3,255,827
IPv4 Glue records	3,198,649
IPv6 Glue records	57,178
Glue Addresses (unique)	2,726,352
IPv4	2,677,195
IPv6	49,157
Nameservers with IPv4 and IPv6	43,369
Nameservers with IPv4 Only	2,394,361
Nameservers with IPv6 Only	5,183

Glue addresses per nameserver

Glue/Nameserver	Count	Glue/Nameserver	Count
0	726,039	7	789
1	2,314,201	8	861
2	53,681	9	64
3	3,334	10	44
4	69,167	11	49
5	511	12	10
6	164	13	38

The "sticky thing" (I.e., "Glue" – bad pun)

- The previous slides "stuck" to the glue records
  - The EDNS test tool would use glue records for the "dig @ parameter"

- If there are no glue records for a nameserver, the tool will "dig" at the default recursive server for authoritative addresses
  - Analysis of the addresses in the results mixes glue addresses and authoritative addresses

- In future work, studying glue and studying authoritative address sets separately is a goal

Looking at zones per nameserver (long tail)

Zones delegated to Server	Servers Counted
1 (server serves only the one zone)	1,952,398
2	248,665
3	139,192
4	96,307
5	72,713
...	
10	28,789
...	
Max – next slide	

"Big fish" – nameservers with lots of zones

- Names are faked (except for the "1" and "2"'s), the numbers are real

Nameserver2.vendor1.example. has 4,014,724 zones
Nameserver1.vendor1.example. has 4,014,702 zones

Nameserver1.vendor2.example. has 3,889,501 zones
Nameserver2.vendor2.example. has 3,885,716 zones

Nameserver1.vendor3.example. has 3,223,197 zones
Nameserver2.vendor3.example. has 3,222,739 zones

Without real names, this seems less dramatic...

"Multi-tenant"

- Some operators use one name server name to server many zones
- Some operators use "vanity" name server names relying on the same IP address, this will be seen later
- Some operators use "vanity" names and addresses all hosted on the same process (same DNS server), this is not obvious in this study

Looking at TLDs represented per nameserver

TLDs represented on a Server	Servers Counted
1 (all server's zones are in one TLD)	2,352,370
2	355,250
3	190,748
4	94,888
5	56,610
...	
10	6,044
...	
537	1
539 (maximum seen)	2

"Compressing" the tests

- The test software in use takes these parameters
  - <zone> <address> <nameserver>
  - It would run once for each zone on each nameserver's addresses
  - Common-held assumption, server would behave the same for all zones as far as EDNS0 is concerned (but not if testing for, say, lame delegations)
- The mean number of zones for a nameserver is about 144(.49)
  - I "simplistically" expected a "gain" of 144:1

Estimating the expected load on testing

- The mean number of addresses for each nameserver is “complicated”
  - 726,039 servers have no glue (addresses), test software will look up addresses
  - 2,442,913 servers have addresses, mean is 1.116 (excluding glueless)
- Expected tests: 2,726,352 tests plus unknown number more tests for 726,039 glueless servers
- Of the glueless servers
  - Some will have no addresses (including NXDOMAIN)
  - Some will likely have more than one
  - But there's no way to tell ahead of time
- Why estimate? To know how many VM's are needed to launch the test

Launching the tests

- How many test results were pulled back (from probe machines)?
  - Tests: 3,533,474
  - Tests for gluefull servers: 2,726,352
  - Tests for glueless servers with addresses: 596,647
  - Tests for glueless servers with no addresses: 210,475

- How many "virtual" tests?
  - Expanded results: 999,793,566
  - Gain of about 283:1 (not 144:1)
  - I hadn’t accounted for the address multiplier effect

"Found" addresses for the glueless

- For glueless: 726,039 servers: 596,647 total addresses

Addresses found per Glueless Server	Servers Counted
0 (NXDOMAIN or NoError/NoAnswer)	210,475
1	465,564
2	35,242
3	989
4	13,101
5	152
...	
33	1
58	1

There's interesting things to study in addresses

- But this data is "muddy" – mixture of glue addresses and authoritative answers
- Just a surface look – look for where the most name servers claim the same address
- IPv4 (of course) but in IPv6 too
- Hilbert Curves would be useful here

Looking at the V4 addresses - names in nets

- Addresses 1,152,553
- Servers 3,239,077
- Singletons 692,847 (addresses having just one name)

IPv4 Prefix	Named servers
A1.B1/16	332,993
A1.B1.C1/16	331,502
A1.B1.C1.D1/24	331,502
A2.B2/16	46,480
A3.B3/16	32,593
A3.B3.C2/16	23,241

IPv4 Prefix	Named servers
A6.B6/16	21,128
A6.B6.C6/16	5,871
A6.B6.C7/16	5,536
A6.B6.C8/16	5,038
A6.B6.C9/16	4,513

Remarkable address reuse (these are V4/32's)

Names	Address	Names	Address
21,058	A/24.124	20,581	A/24.252
21,009			A/24.249
20,975			A/24.250
20,933	A/24.120	20,463	A/24.125
20,884	A/24.251	20,410	A/24.123
20,860	A/24.246	20,409	A/24.243
20,796	A/24.122	11,651	B/24.114
20,765	A/24.128	11,642	C/24.226
20,688	A/24.121	8,092	D/16.X.20
20,646	A/24.119	7,695	D/16.Y.20

331,476 named servers in :
"A/24".118-125 and 246-253

Looking at the V6 addresses - names in nets

- Addresses 48,449
- Servers 83,922
- Singletons 40,167

IPv6 Prefix	Named servers
add1::/48	2386
add2::/48	1752
add2::3/128	1751
add3::/48	1751
add3::3/128	1751

Other Counts for a Single IPv6/128
867
864
626
551
551
551
548
416

Finally, the EDNS0 results

- The purpose of the test was to decide if a zone will "suffer"
  - But there are 193,825,454 zones
  - And this presentation is past the DNS Flag Day
  - And it's hard to juggle the different nameservers for a zone, v4 vs v6
  - And it's hard to determine whether a nameserver is "good enough"

- So, some gross results, covering all zone/nameserver/address combinations
  - And for just some of the 11 experiments

dns= experiment

```
dns=
 910,824,012 ( 91%) ok
  36,787,293 (  3%) timeout
  27,423,704 (  2%) refused,nosoa
  10,583,952 (  1%) nosoa
  10,169,837 (  1%) servfail,nosoa
 997,423,084 (100%) Total dns
```

edns= experiment

```
edns=
 764,496,023 ( 76%) ok
 147,075,831 ( 14%) noopt
  35,857,545 (  3%) timeout
  26,589,459 (  2%) refused,nosoa
 997,423,084 (100%) Total edns
```

do= experiment

```
do=
 772,963,243 ( 77%) ok
 139,029,488 ( 13%) noopt
  35,222,260 (  3%) timeout
  27,039,326 (  2%) refused,nosoa
 997,423,084 (100%) Total do
```

ednstcp= experiment

```
ednstcp=
 583,738,539 ( 58%) ok
 274,638,598 ( 27%) timeout
  62,613,595 (  6%) noopt
  24,082,648 (  2%) connection-refused
  18,862,925 (  1%) refused
 997,423,084 (100%) Total ednstcp
```

address not found failures

```
excuse=
    1,744,564 ( 73%) no address records
      625,918 ( 26%) no address records found
    2,370,482 (100%) Total excuse
```

Lessons learned

- Having a tool to scan for protocol compliance is interesting
  - Having a clear "yes/no" output is helpful
- More importantly, this spurred a framework to test other protocol compliance questions
  - Glue: match reality (vs. Authoritative answer)? reverse map agree?
  - Lame delegations: do servers really host the zones, and vice versa
- Accuracy in measurements: avoiding timeouts, "freshness" in collecting
- Goal: to improve ability to manage the DNS system, optimize tool development and perhaps protocol development

So, how did I get the "almost 11 Billion DiGs"?

- 997,423,084 (100%) Total dns tests
  - These did 11 digs
- The glueless: 726,039
  - These added 1 dig (regardless of whether they resulted in test above)

```
$ dc

997423084 11 * 726039 2 * + p

10973106002
```

- In other words: 10,973,106,002 or "about 11 billion"

Engage with ICANN

Thank You and Questions

Visit us at **icann.org**
Email: edward.lewis@icann.org

@icann

linkedin/company/icann

facebook.com/icannorg

slideshare/icannpresentations

youtube.com/icannnews

soundcloud/icann

flickr.com/icann

instagram.com/icannorg