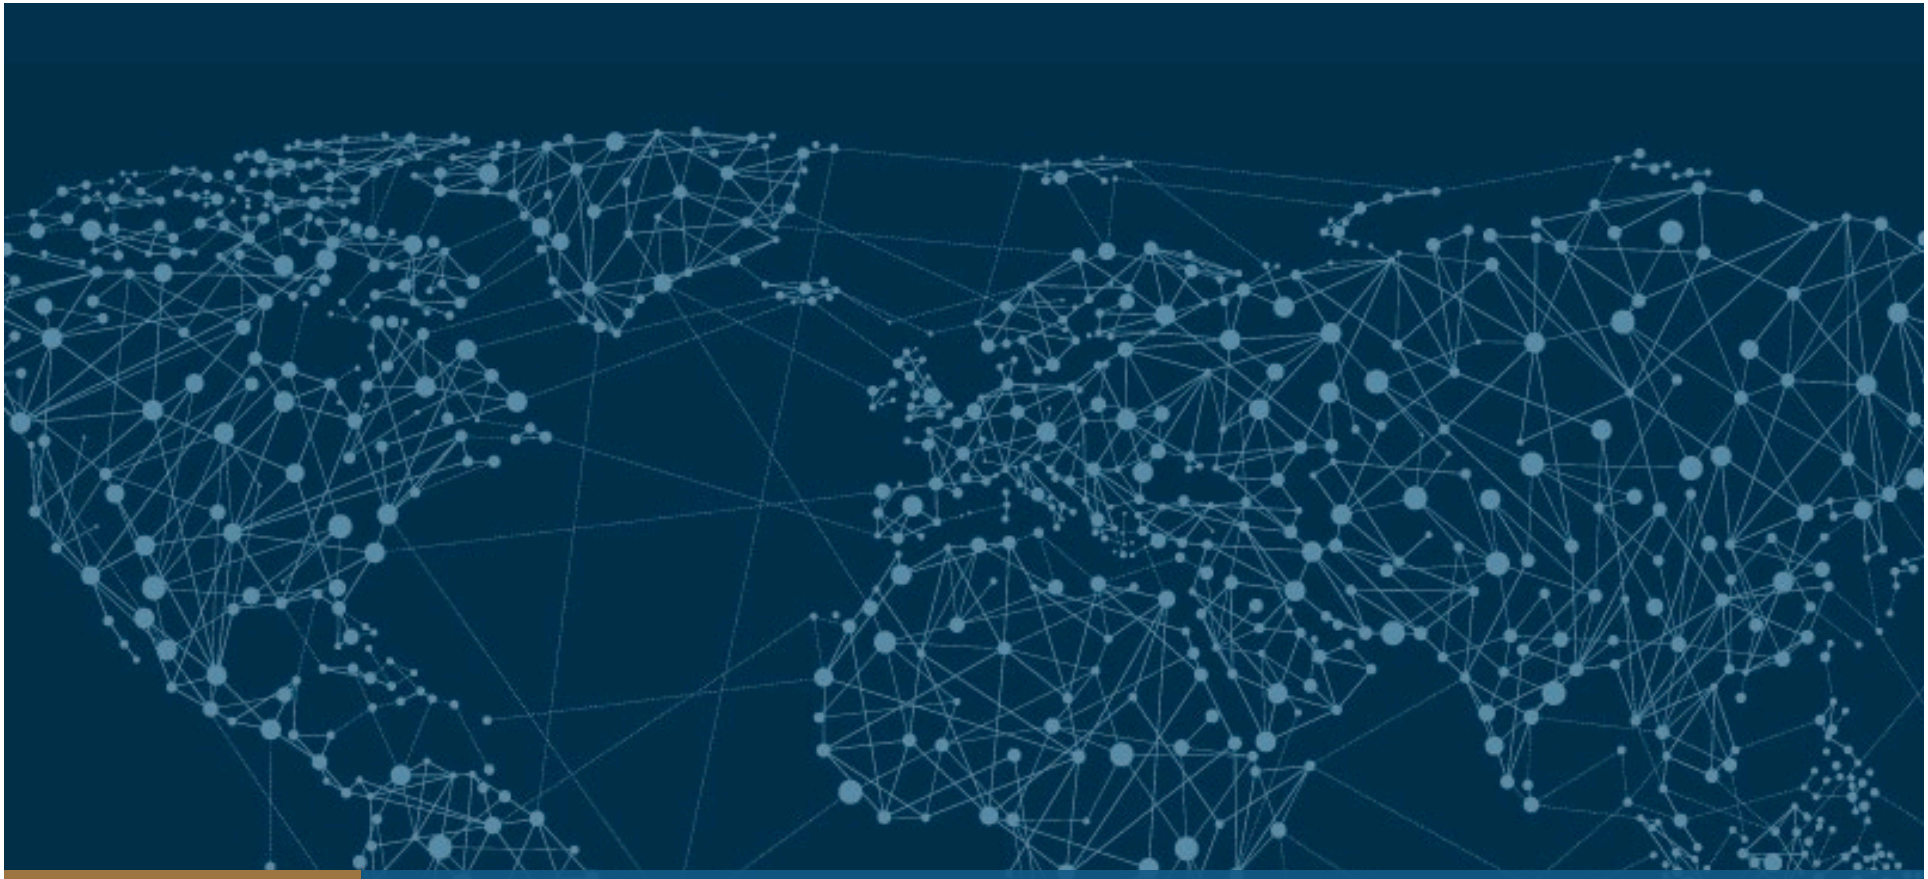

IPv6 Meets Large Responses from TLD DNS Servers

Edward Lewis | ICANN DNS Symposium | 13 May 2017

Question:

- As IPv6 fragmentation might cause problems for large DNS responses, will this be an issue when the Root Zone DNSSEC KSK rolls over?
  - During the planned rollover operation, the response to a query for the root zone's signed DNSKEY set will reach 1424 bytes in size

Why would IPv6 fragmentation be an issue?

- Unlike IPv4, in IPv6 fragmentation is only performed at the source, not at an intermediary node
  - If an intermediary node receives a packet too big, it returns an error to the sender
  - When such a notice is delivered to the original sender, the sender is supposed to retry with a smaller (source-fragmented) packets
  - But DNS servers do not retain history and cannot retry (but can mark the path's maximum transfer unit)

What could the outcome be?

- If DNS queriers can not get, or have "difficulty getting", the root zone's signed DNSKEY Resource Record set, then following the Automated Updates of DNSSEC Trust Anchors protocol (described in RFC 5011) may fail
  - There will be a rise in retry traffic, hoping the path maximum transfer unit parameter is adjusted
  - The querier will have to manually update their trust anchor set
- Or so some "think"

Can we test this?

- Yes, currently there are many TLD zones with large DNSKEY RRsets
  - Positive DNSKEY resource record set answers range from 571 bytes to 3319 bytes
  - More than 115 cases where the response is over 1280 bytes, 80 over 1500 bytes (ethernet MTU)
- We hear no complaints about large responses. Why?
  - Maybe our fear of IPv6 is misplaced

This activity is not to evaluate operators

- Criteria for success and failure is not fairly determined
- Operators sometimes need to perform actions that might make it seem like a service is down
  - Per-machine maintenance
  - DDoS mitigation
- Sometimes mid-network peering fails
  - The Internet is not completely connected after all
- What is presented does not measure an operator, it measures the receiver

Testing

- The test is first to determine whether a response to a query is received
  - A test involving non-responses is trickier than a test of "what is returned"
  - Many factors are involved in a non-response, it might not be "the fault" of the other end
- Tests are run with two elements in mind
  - Repeated once a day
  - Performed from a number of locations

The test set up

- The test tries all name servers listed in the root zone for TLDs, once per day (per test location)
  - Uses "glue addresses" provided to PTI (IANA)
  - Per address, asks for the TLD's SOA over TCP
  - Asks for the TLD's DNSKEY resource record set, signed over UDP and over TCP
- Test probe locations
  - VMs in Amsterdam, Bangalore, New York, Singapore, Toronto
  - Residential cable-TV network (i.e., "generic" home)

Why no complaining?

- The first slicing of the data is to look TLD by TLD. Check whether the signed DNSKEY resource record set is received

Amsterdam/Singapore/NYC/Toronto

- From April 25 to May 4, 2017
- Results from Amsterdam, Singapore, New York City, and Toronto are identical

At each of 4 loc	25th	26th	27th	28th	29th	30th	1st	2nd	3rd	4th	average
Total TLDs:	1532	1532	1532	1532	1532	1532	1532	1532	1532	1532	1532
Failed TCP v6:	1	1	1	1	1	1	1	1	1	1	1
Failed UDP v6:	1	1	1	1	1	1	1	1	1	1	1
Failed TCP v4:	0	0	0	0	0	0	0	0	0	0	0
Failed UDP v4:	0	0	0	0	0	0	0	0	0	0	0

- Boring! Good, but boring for a researcher.
- But, wait, there's one TLD reported as failing
  - Same one, each test, day, location

Bangalore and Residential

Bangalore	25th	26th	27th	28th	29th	30th	1st	2nd	3rd	4th	average
Total TLDs:	1532	1532	1532	1532	1532	1532	1532	1532	1532	1532	1532
Failed TCP v6:	1	2	1	1	1	2	1	1	1	1	1
Failed UDP v6:	4	4	4	2	3	6	3	3	2	2	3
Failed TCP v4:	0	0	0	0	0	0	0	0	0	0	0
Failed UDP v4:	0	0	0	0	0	0	0	0	0	0	0

Residential	25th	26th	27th	28th	29th	30th	1st	2nd	3rd	4th	average
Total TLDs:	1532	1532	1532	1532	1532	1532	1532	1532	1532	1532	1532
Failed TCP v6:	2	2	3	2	2	3	2	2	2	2	2
Failed UDP v6:	6	14	16	12	18	7	5	14	7	10	11
Failed TCP v4:	0	0	0	0	0	0	0	0	0	0	0
Failed UDP v4:	0	0	0	0	0	0	0	0	0	0	0

- Less boring (for a researcher)!
- TCP fails are low, UDP shows issues

Fails at Bangalore and Residential site

- The same lone (IDN ccTLD) TLD failing all sites
  - It appears that the TLD hasn't properly registered all their glue; they have working authoritative servers found by requesting authoritative data in DNS
- At the residential, a second (two-letter ccTLD) TLD consistently fails for TCP
- For TCP, there are spurious other fails
- UDP fails exist, for residential mostly confined to one IPv6 prefix (with a consistent set of TLDs)

One-level deeper

- What about looking server-by-server? Perhaps more can be learned about the protocol stack's behaviors

Server by server results

These results are for April 25-30, showing "average" values

Statistic	AMS	SIN	NYC	YYZ	BLR	Resid.
IPv6 Servers	6410	6410	6410	6410	6410	6410
Bad Results	36	77	34	46	116	301
No SOA	32	30	30	28	31	39
Only Fail TCP	2	0	0	0	4	149
Only Fail UDP	2	47	3	18	80	95
No TCP & UDP	1	0	0	0	1	19

Statistic	AMS	SIN	NYC	YYZ	BLR	Resid.
IPv4 Servers	7262	7262	7262	7262	7262	7262
Bad Results	50	150	55	57	131	78

Amsterdam & NYC data, April 25-May 4, 2017

Amsterdam	25th	26th	27th	28th	29th	30th	1st	2nd	3rd	4th	avg
Total TLDxIP6:	6410	6410	6412	6412	6408	6408	6408	6408	6408	6408	6409.2
Bad TLDxIPv6:	34	36	37	38	37	34	38	36	39	38	36.7
No SOA TLDxIPv6:	31	33	32	32	31	31	35	32	35	35	32.7
No TLDxIPv6 TCP:	2	2	2	3	3	1	1	1	2	0	1.7
No TLDxIPv6 UDP:	1	1	3	2	2	1	1	1	1	3	1.6
No TLDxIPv6 both:	0	0	0	1	1	1	1	2	1	0	0.7
Total TLDxIP4:	7263	7263	7264	7264	7260	7260	7260	7260	7260	7260	7261.4
Bad TLDxIPv4:	50	50	53	49	49	50	50	63	65	68	54.7

New York City	25th	26th	27th	28th	20th	30th	1st	2nd	3rd	4th	avg
Total TLDxIP6:	6410	6410	6412	6412	6408	6408	6408	6408	6408	6408	6409.2
Bad TLDxIPv6:	36	32	32	36	32	33	35	32	31	33	33.2
No SOA TLDxIPv6:	30	29	30	30	30	31	31	27	26	27	29.1
No TLDxIPv6 TCP:	2	0	0	0	0	0	2	1	0	1	0.6
No TLDxIPv6 UDP:	3	3	2	6	2	2	2	4	4	4	3.2
No TLDxIPv6 both:	1	0	0	0	0	0	0	0	1	1	0.3
Total TLDxIP4:	7263	7263	7264	7264	7260	7260	7260	7260	7260	7260	7261.4
Bad TLDxIPv4:	58	49	52	53	59	58	58	67	62	64	58

Toronto, same dates

Toronto	25th	26th	27th	28th	20th	30th	1st	2nd	3rd	4th	avg
Total TLDxIP6:	6410	6410	6412	6412	6408	6408	6408	6408	6408	6408	6409.2
Bad TLDxIPv6:	42	42	53	50	46	43	51	42	45	47	46.1
No SOA TLDxIPv6:	26	26	28	29	28	28	28	27	29	25	27.4
No TLDxIPv6 TCP:	1	0	0	0	0	0	1	1	0	0	0.3
No TLDxIPv6 UDP:	15	16	25	21	18	15	22	14	16	22	18.4
No TLDxIPv6 both:	0	0	0	0	0	0	0	0	0	0	0
Total TLDxIP4:	7263	7263	7264	7264	7260	7260	7260	7260	7260	7260	7261.4
Bad TLDxIPv4:	61	59	64	52	53	55	57	74	66	60	60.1

- Amsterdam, NYC, and Toronto VMs, fare better than Singapore, Bangalore
  - Some UDP unhappiness in Toronto

Bangalore and Singapore, April 25-May 4, 2017

Bangalore	25th	26th	27th	28th	20th	30th	1st	2nd	3rd	4th	avg
Total TLDxIP6:	6410	6410	6412	6412	6408	6408	6408	6408	6408	6408	6409.2
Bad TLDxIPv6:	119	110	99	112	107	147	162	170	146	115	128.7
No SOA TLDxIPv6:	30	30	27	31	33	34	35	33	31	28	31.2
No TLDxIPv6 TCP:	5	3	0	4	4	6	4	5	4	0	3.5
No TLDxIPv6 UDP:	84	76	72	76	69	105	120	131	110	84	92.7
No TLDxIPv6 both:	0	1	0	1	1	2	3	1	1	3	1.3
Total TLDxIP4:	7263	7263	7264	7264	7260	7260	7260	7260	7260	7260	7261.4
Bad TLDxIPv4:	138	126	135	143	121	122	136	138	141	129	132.9

Singapore	25th	26th	27th	28th	20th	30th	1st	2nd	3rd	4th	avg
Total TLDxIP6:	6410	6410	6412	6412	6408	6408	6408	6408	6408	6408	6409.2
Bad TLDxIPv6:	99	105	75	86	42	57	55	52	31	32	63.4
No SOA TLDxIPv6:	29	29	31	30	30	31	32	27	29	29	29.7
No TLDxIPv6 TCP:	0	0	0	0	0	0	0	0	0	0	0
No TLDxIPv6 UDP:	70	76	44	56	12	26	23	25	2	3	33.7
No TLDxIPv6 both:	0	0	0	0	0	0	0	0	0	0	0
Total TLDxIP4:	7263	7263	7264	7264	7260	7260	7260	7260	7260	7260	7261.4
Bad TLDxIPv4:	169	190	135	144	116	145	108	142	138	118	140.5

Residential network, same dates

Residential	25th	26th	27th	28th	29th	30th	1st	2nd	3rd	4th	average
Total TLDxIP6:	6410	6410	6412	6412	6408	6408	6408	6408	6408	6408	6409.2
Bad TLDxIPv6:	301	294	330	318	293	270	271	281	279	294	293.1
No SOA TLDxIPv6:	80	28	31	30	30	34	30	28	29	31	35.1
No TLDxIPv6 TCP:	134	149	157	170	138	144	139	131	142	147	145.1
No TLDxIPv6 UDP:	74	97	121	95	106	77	92	104	93	97	95.6
No TLDxIPv6 both:	13	20	21	23	19	15	10	18	15	19	17.3
Total TLDxIP4:	7263	7263	7264	7264	7260	7260	7260	7260	7260	7260	7261.4
Bad TLDxIPv4:	64	93	96	94	58	60	59	59	58	58	69.9

- Interesting: In IPv4, residential looks like Amsterdam, NYC, and Toronto VMs, far better than Singapore, Bangalore
- In IPv6, must "worse" in responses-received

What does this hint?

- From the still smallish data set, local conditions seem to matter in IPv6
- With the exception of one TLD, there's no other pattern of failing behavior across sites
- Deeper digging into the residential case indicated that most, if not all, errors were from one IPv6 routing prefix
  - The residential host cannot see one of the root zone servers in V6

What does this mean for the KSK rollover?

- There may be isolated reports of "breakage"
  - Potentially motivated individuals running DNSSEC validation in their residential-class environments might see trouble
  - Potentially from DNS service providers running in locations that might otherwise experience network issues (leaping to the conclusion that the Bangalore node falls into that)
    - The only other data on the Bangalore node is that access (from US east coast) and upload/download times are slower than the other VMs

What does this mean for IPv6 and DNS?

- Large responses seem not to be a primary problem
- Drop outs probably related to other, non-IPv6 (fragmentation) issues
- More work could be done (RIPE Atlas might be handy) to examine reasons for non-responses

Engage with ICANN

Thank You and Questions

Reach us at:
Email: engagement@icann.org
Website: icann.org

twitter.com/icann

soundcloud.com/icann

facebook.com/icannorg

weibo.com/ICANNorg

youtube.com/user/icannnews

flickr.com/photos/icann

linkedin.com/company/icann

slideshare.net/icannpresentations