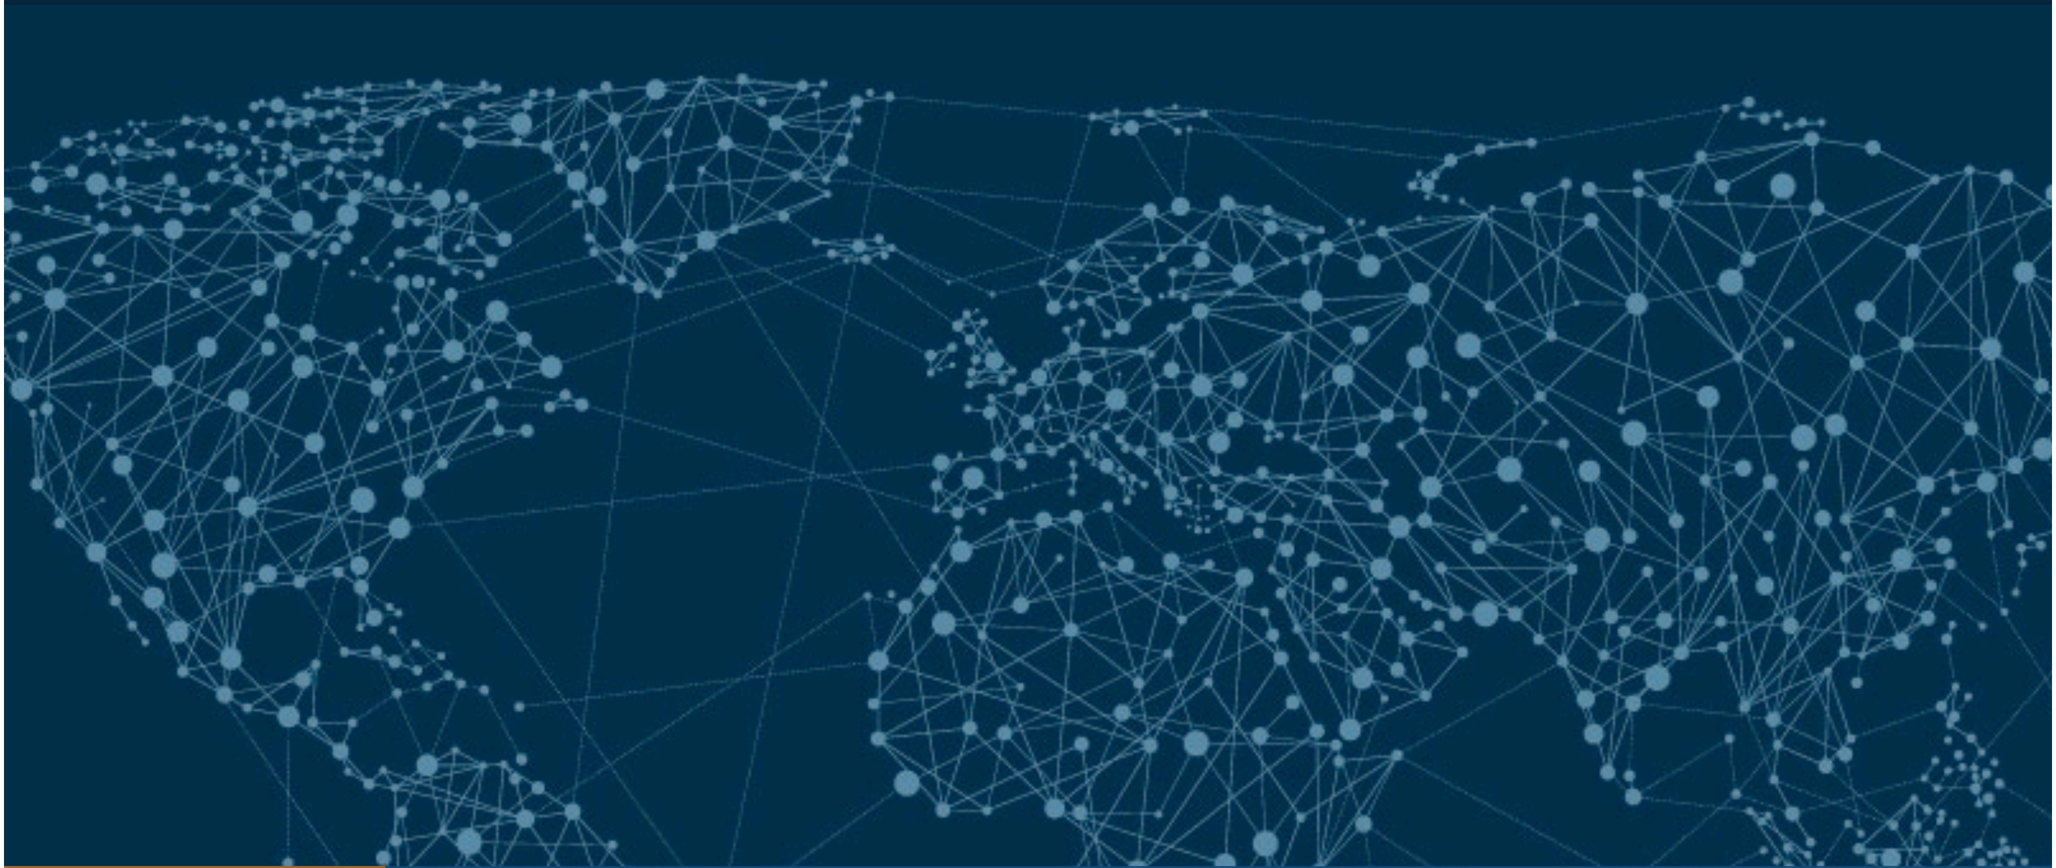

Continuous Data-driven Analysis of Root Server System Stability (CDAR)

ICANN/TNO | RSSAC Teleconference | 7 January 2016

Review background

Background

This review of the New gTLD Program for its security and stability impact was commissioned in keeping with previous commitments including advice from the GAC. Specifically, the Board committed to deferring a future round of new gTLDs unless an evaluation indicates the current round did not jeopardize the security or stability of the root zone system.

Study RFP

ICANN issued an RFP in June 2015, signing a contract with TNO in October 2015. TNO is working in a consortium with SIDN and NLnet Labs to conduct the Continuous Data-driven Analysis of Root Server System Stability (CDAR.)

Next steps

ICANN published the CDAR study plan for public comment on 2 December. The CDAR team invites comments on the data sources it seeks, as well as the study's methodology. The public comment period closes **3 February**, with an analysis of the comments expected **17 February**.

Target Timeline

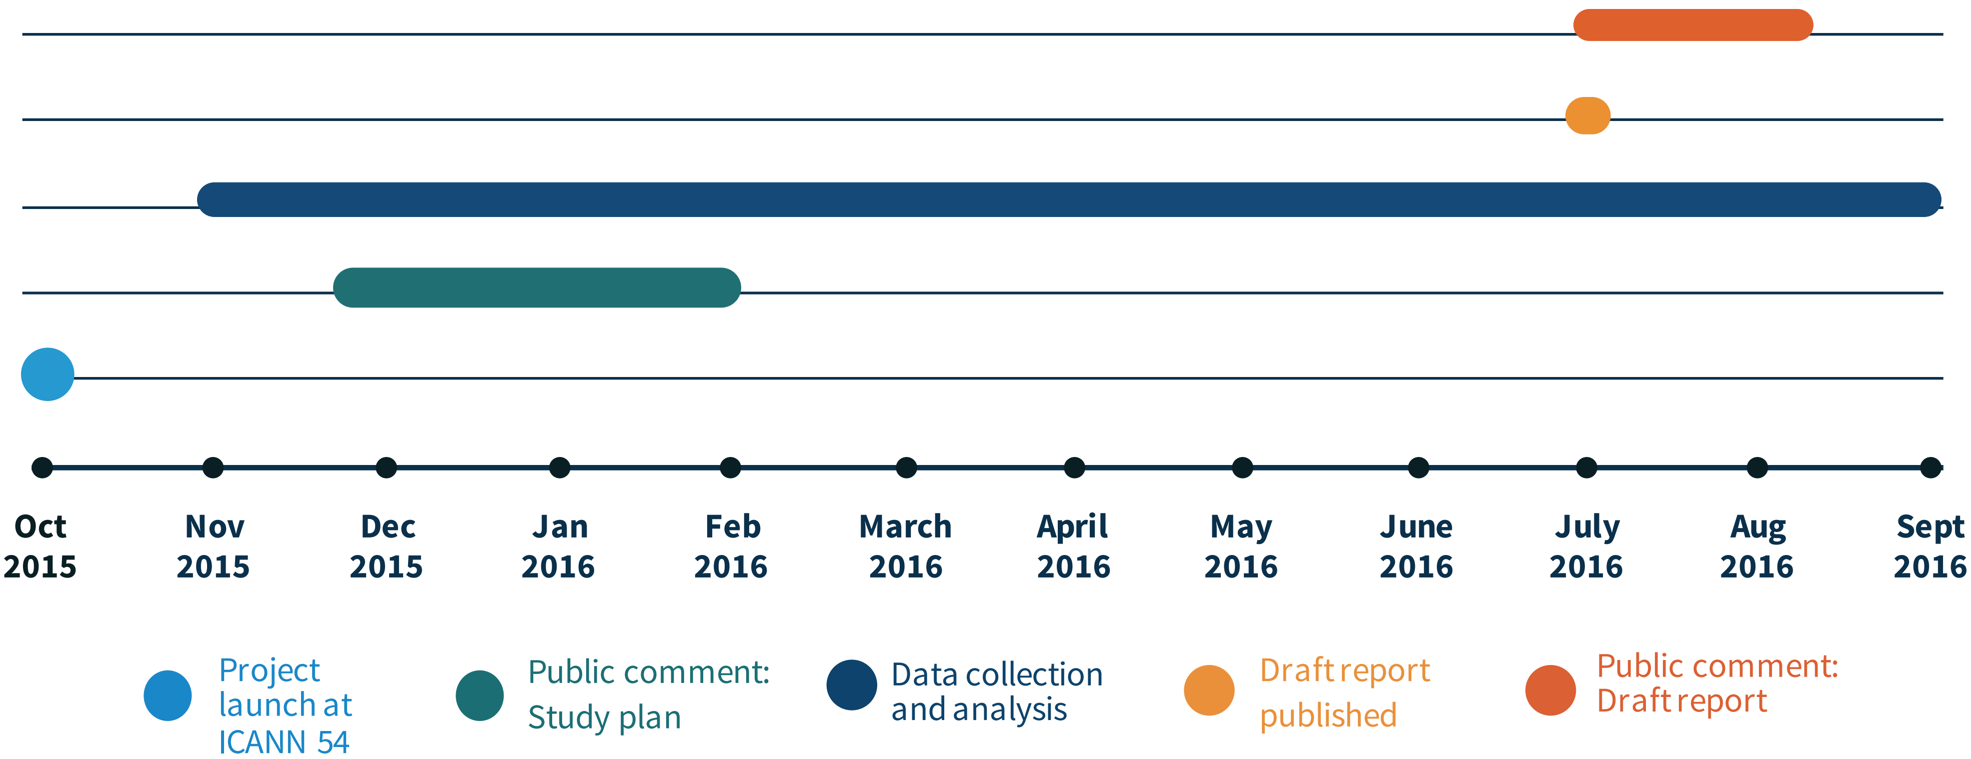

Goals and expectations

The study should provide an understanding of the technical impact to the DNS system of adding new gTLDs to the root. ICANN anticipates public comment received after publication of the first draft will inform the context and content of the final study and report. A final report is anticipated by Q2 2017.

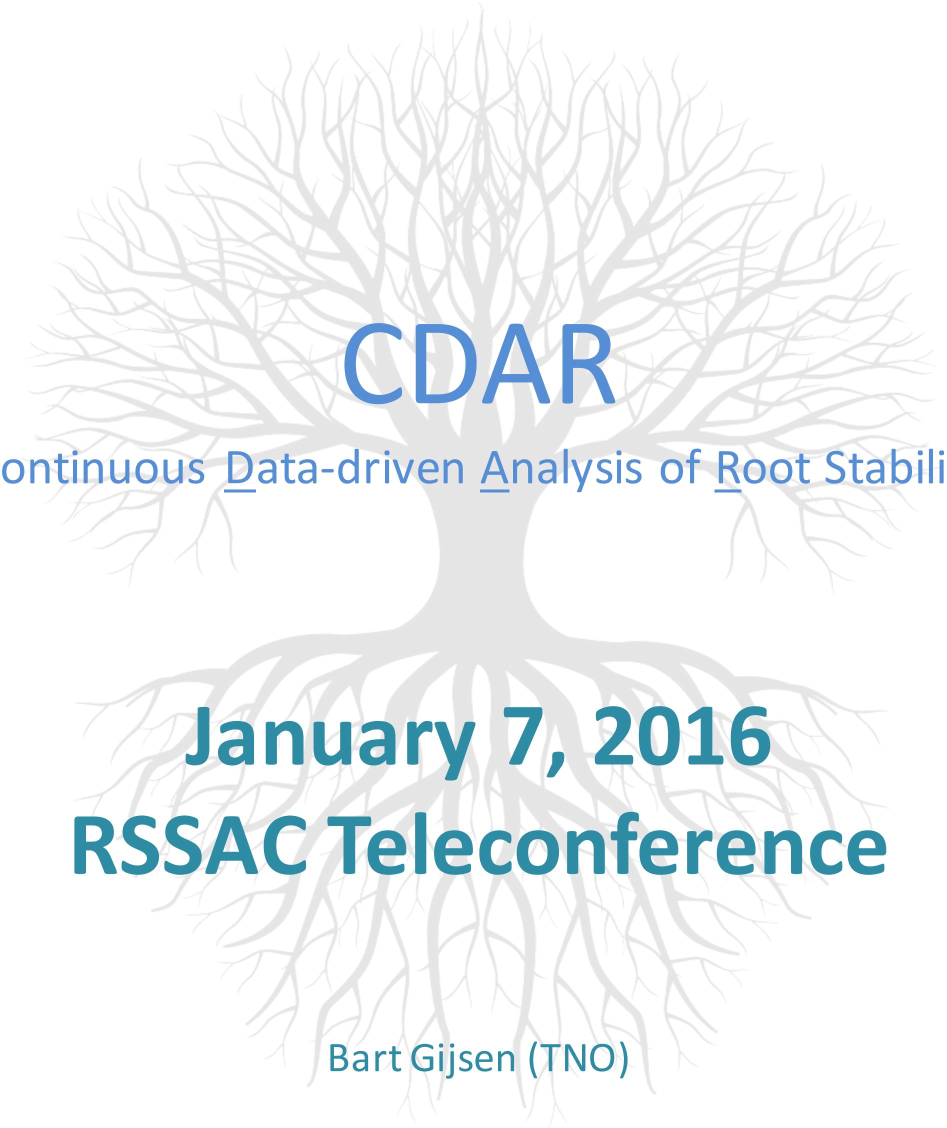

CDAR

Continuous Data-driven Analysis of Root Stability

January 7, 2016
RSSAC Teleconference

Bart Gijsen (TNO)

NLnet Labs SIDN labs TNO

CDAR Study Approach

- Initial DNS stability metrics:

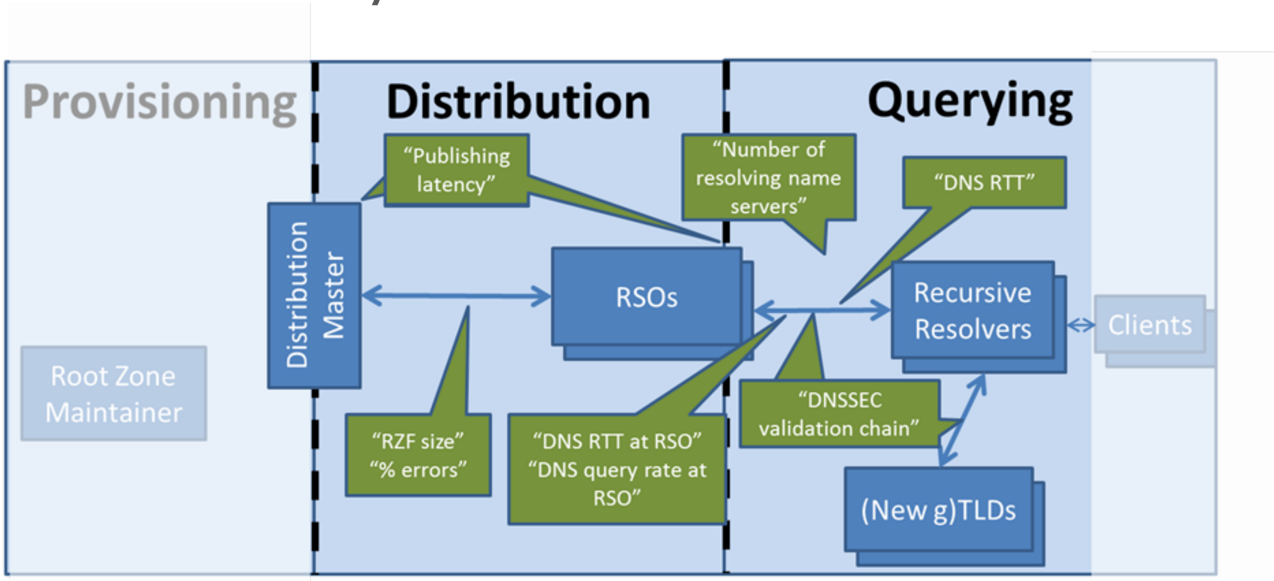

Provisioning

Distribution

Querying

Root Zone Maintainer

Distribution Master

RSOs

Recursive Resolvers

Clients

(New g)TLDs

"Publishing latency"

"Number of resolving name servers"

"DNS RTT"

"RZF size"
"% errors"

"DNS RTT at RSO"
"DNS query rate at RSO"

"DNSSEC validation chain"

- CDAR approach:
  - Analysis via passive and active measurement
    - Later: extrapolations based on a model of identified 'invariant correlations'
  - DNS threat analyses → hypotheses → measurement & analysis

Examples of hypotheses

Hyp-x.y: Increasing the number of TLDs does not significantly increase the query rate to the root

Hyp-x.y: The ratio between #domains in a TLD and query rate to the DNS root are comparable for New gTLDs and other TLDs

Hyp-x.y: Bogus traffic ending up at the Root is not increased, nor decreased by New gTLDs

Hyp-x.y: When a New gTLD is first delegated in the RZF (or when RZF data is changed) this has no significant impact on the query rate to the Root

Hyp-x.y: Ratio of TCP/UDP queries will be higher for New gTLD than for TLDs

H-x.y.z: Increased TCP/UDP ratio will have some impact on RSS server load

Etc. etc.

- Reflection from RSSAC members on hypotheses is highly appreciated

Preliminary analysis of passive measurements

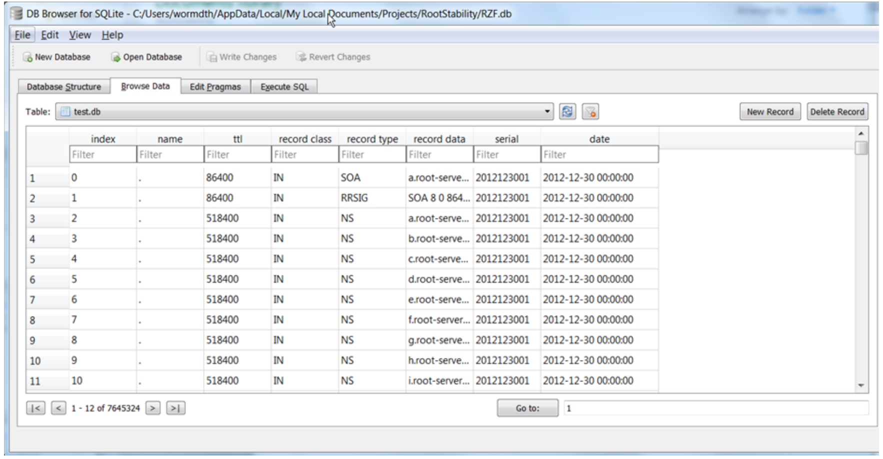

DB Browser for SQLite - C:/Users/wormdth/AppData/Local/My Local Documents/Projects/RootStability/RZF.db

	index	name	ttl	record class	record type	record data	serial	date
1	0	.	86400	IN	SOA	a.root-serve...	2012123001	2012-12-30 00:00:00
2	1	.	86400	IN	RRSIG	SOA 8 0 864...	2012123001	2012-12-30 00:00:00
3	2	.	518400	IN	NS	a.root-serve...	2012123001	2012-12-30 00:00:00
4	3	.	518400	IN	NS	b.root-serve...	2012123001	2012-12-30 00:00:00
5	4	.	518400	IN	NS	c.root-serve...	2012123001	2012-12-30 00:00:00
6	5	.	518400	IN	NS	d.root-serve...	2012123001	2012-12-30 00:00:00
7	6	.	518400	IN	NS	e.root-serve...	2012123001	2012-12-30 00:00:00
8	7	.	518400	IN	NS	f.root-server...	2012123001	2012-12-30 00:00:00
9	8	.	518400	IN	NS	g.root-serve...	2012123001	2012-12-30 00:00:00
10	9	.	518400	IN	NS	h.root-serve...	2012123001	2012-12-30 00:00:00
11	10	.	518400	IN	NS	i.root-server...	2012123001	2012-12-30 00:00:00

1 - 12 of 7645324

Zone File Repository loaded into Database

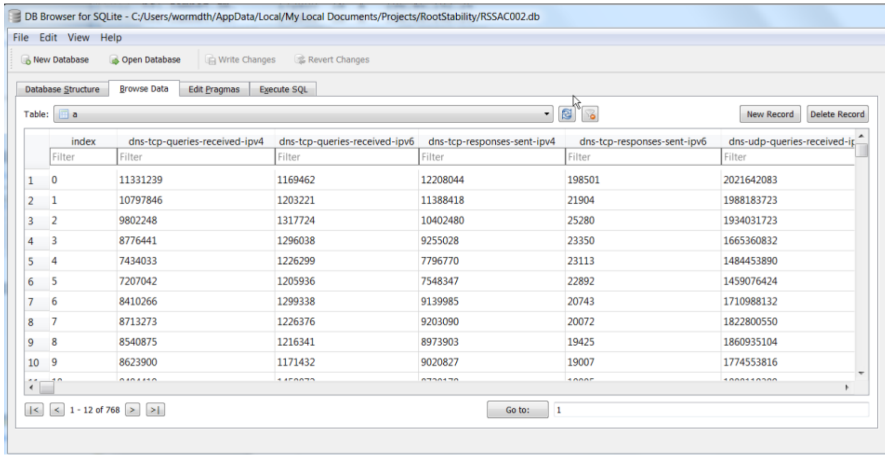

DB Browser for SQLite - C:/Users/wormdth/AppData/Local/My Local Documents/Projects/RootStability/RSSAC002.db

	index	dns-tcp-queries-received-ipv4	dns-tcp-queries-received-ipv6	dns-tcp-responses-sent-ipv4	dns-tcp-responses-sent-ipv6	dns-udp-queries-received-ip
1	0	11331239	1169462	12208044	198501	2021642083
2	1	10797846	1203221	11388418	21904	1988183723
3	2	9802248	1317724	10402480	25280	1934031723
4	3	8776441	1296038	9255028	23350	1665360832
5	4	7434033	1226299	7796770	23113	1484453890
6	5	7207042	1205936	7548347	22892	1459076424
7	6	8410266	1299338	9139985	20743	1710988132
8	7	8713273	1226376	9203090	20072	1822800550
9	8	8540875	1216341	8973903	19425	1860935104
10	9	8623900	1171432	9020827	19007	1774553816

1 - 12 of 768

RSSAC02 measurements loaded into same Database

Hypothesis: Increasing #TLDs does not increase query rate to the root

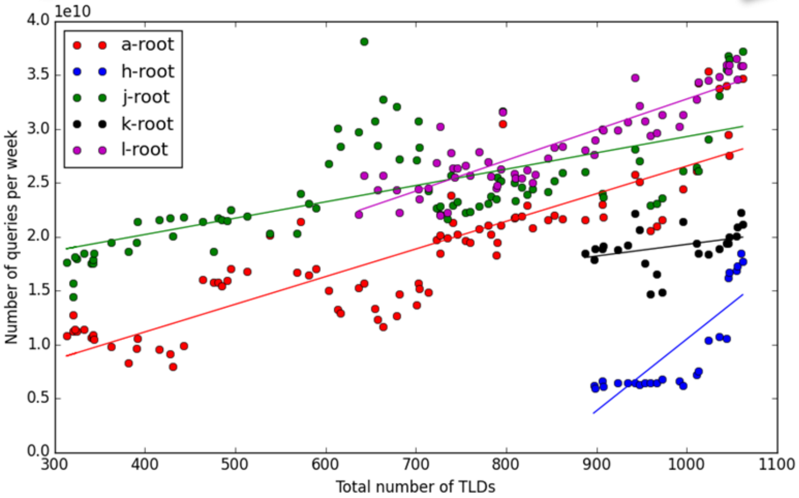

Subsequent hypotheses:

- ***Factor out 'autonomous' growth***
- ***Are #domains better predictor than #TLDs?***
- ***etc.***

Questions to RSOs/RSSAC

- Can the daily reported RSSAC002 measurements be split on a per-TLD basis?
  - Total number of queries and responses per day per TLD
  - Query and response size distribution per day per TLD
  - Etc.

- Do any of the RSSAC members have unpublished RSSAC002-like (or DITL-like) measurements tracing back for a longer period?
  - Can these be made available to the CDAR team?
  - When do other RSOs plan to publish RSSAC002 data?

Preliminary analysis of DITL data

Hyp-x.y: Bogus traffic ending up at the Root is not influenced by New gTLDs

- Analysis by CDAR team of DITL data (2015Apr13) for a, b, c, f, g, i, j, k, l, m-root
  - Queries *for valid TLDs* counted (per TLD); count = 22.339.157.183 queries
  - The counted numbers:
    - Total counted = 22.339.157.183 queries
    - Counted valid New gTLDs (at that time) = 124.746.425 queries (less than 1% of total)
    - Maximum count for a New gTLD (.club) = 5.668.551 queries (less than 0.03 % of total)
  - Next DITL in two months time: redo then?

Questions to RSOs/RSSAC

Does any of the RSSAC members have unpublished DITL alike data sets with more continuous data samples?

- Example: logs with DNS query for a few hours per week?
- Example: logs over period of several weeks during a change to RSO's infrastructure

Preliminary active measurements

Hyp-x.y: DNSSEC validation errors (broken chain) does not occur more frequent for New gTLDs, than for other TLDs

- Active measurement script by NLnet Labs:
  - Script validates all signed domains in all signed TLDs and reports any error
  - Script runs twice day
  - Since January 2012 merely 475 error reports were generated; dominated by one TLD
  - Only 27 TLDs have reported errors
  - From these measurement: DNSSEC validation error hypothesis holds
- Initial investigation has been made to use RIPE Atlas anchors for active measurements
  - Many DNSMON data already available
  - Further investigation with active measurement after measurement plan / hypotheses are worked out in more detail

Request for further interaction

- Feedback (questions / suggestions) about the study plan, hypotheses, analyses is welcome
  - Either via public comment or directly to the CDAR team
- We will contact several individual RSSAC members with more specific requests (for data)
  - Feel free to volunteer, too!
- More in depth session(s) with RSSAC members in ICANN-55 (Marrakech) meeting?

Questions and Discussion

CDAR Project Team
Bart Gijsen (TNO)
Benno Overeinder (NLnet Labs)
Cristian Hesselman (SIDN)
Daniël Worm (TNO)
Giovane Moura (SIDN)
Jaap Akkerhuis (NLnet Labs)

Coordinator
Bart Gijsen (Msc.)
+31 6 53 72 52 18
bart.gijsen@tno.nl

CDAR Home: http://www.cdar.nl