A Review of the KSK Roll

Created By: Geoff Huston

APNIC Labs

Presented By: Tim April

SSAC

Why am I presenting this?

- I ’m not part of ICANN or the PTI
- APNIC is not a root server operator
- I’m not a member of the root server cabal
- I can’t see root server query data

- Geoff is not here this week, but all of the above still apply

Why am I presenting this?

But …

I’m one of almost 4 billion consumers of the DNS, and the stability, integrity and robustness of the DNS root matters for me

So I’m an interested member of the community of DNS consumers

The Plan

- The KSK is “special”
- There is no “parent” key for the root
- Every DNSSEC-validating resolver needs to load (and trust) the new KSK
- The plan is to use “old-signs-new” approach
  - The old key signs over the new key for some minimum hold-down period
  - DNSSEC-validating resolvers are supposed to add the new key to their local trusted key collection once they have seen a stable sign-over for the hold-down period

The Plan - V1

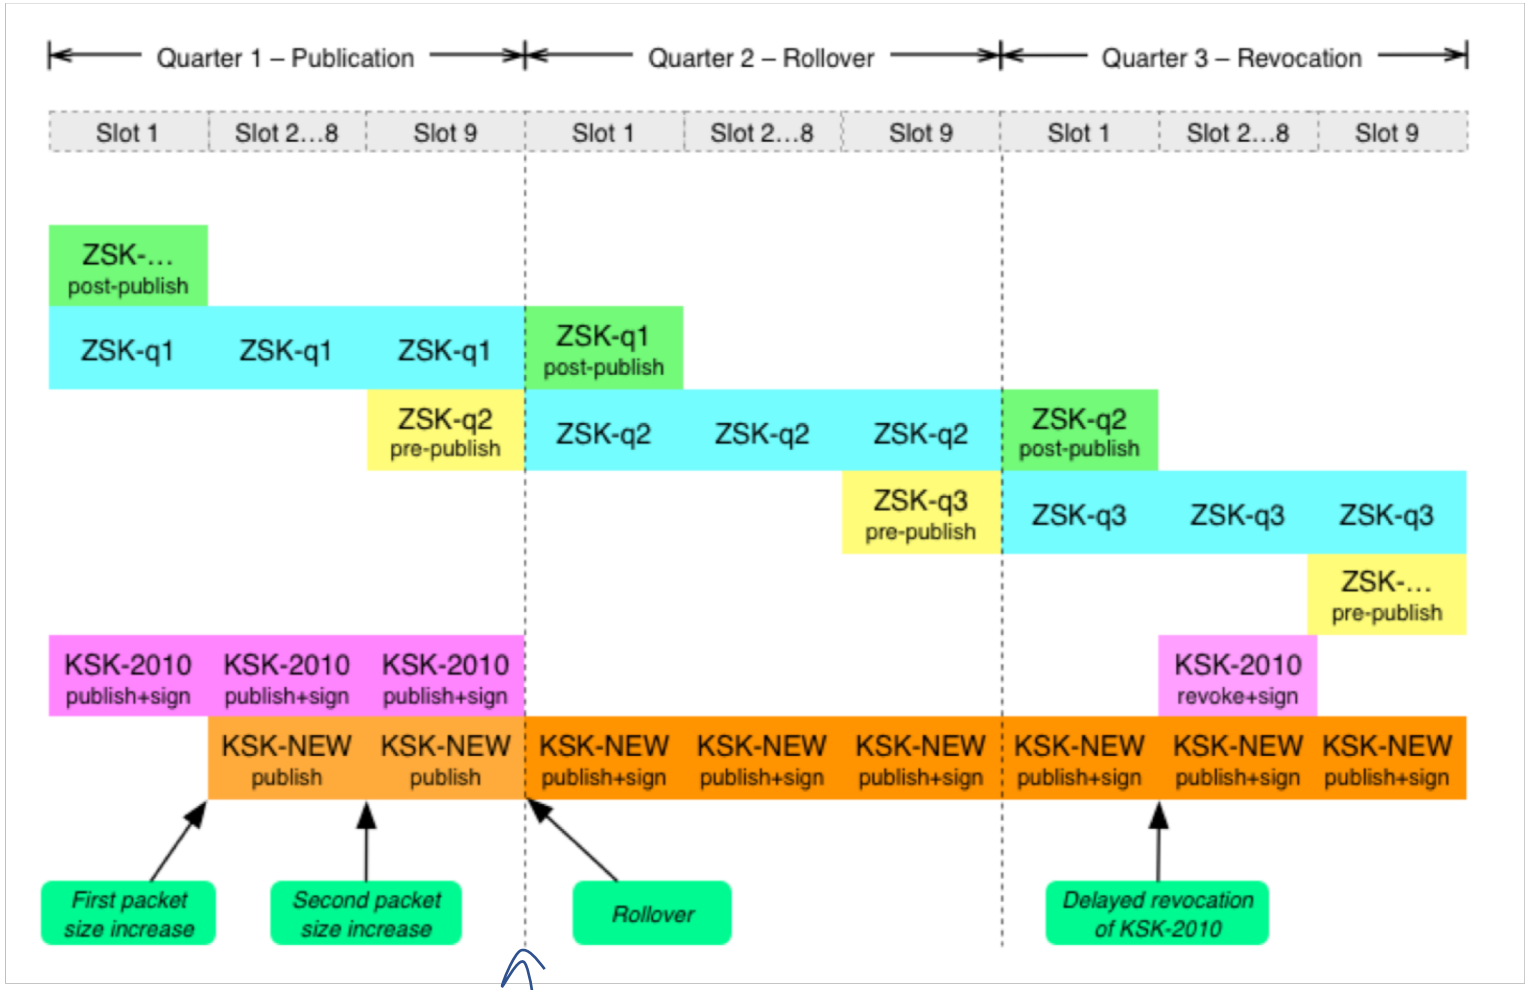

11 October 2017

The best laid plans…

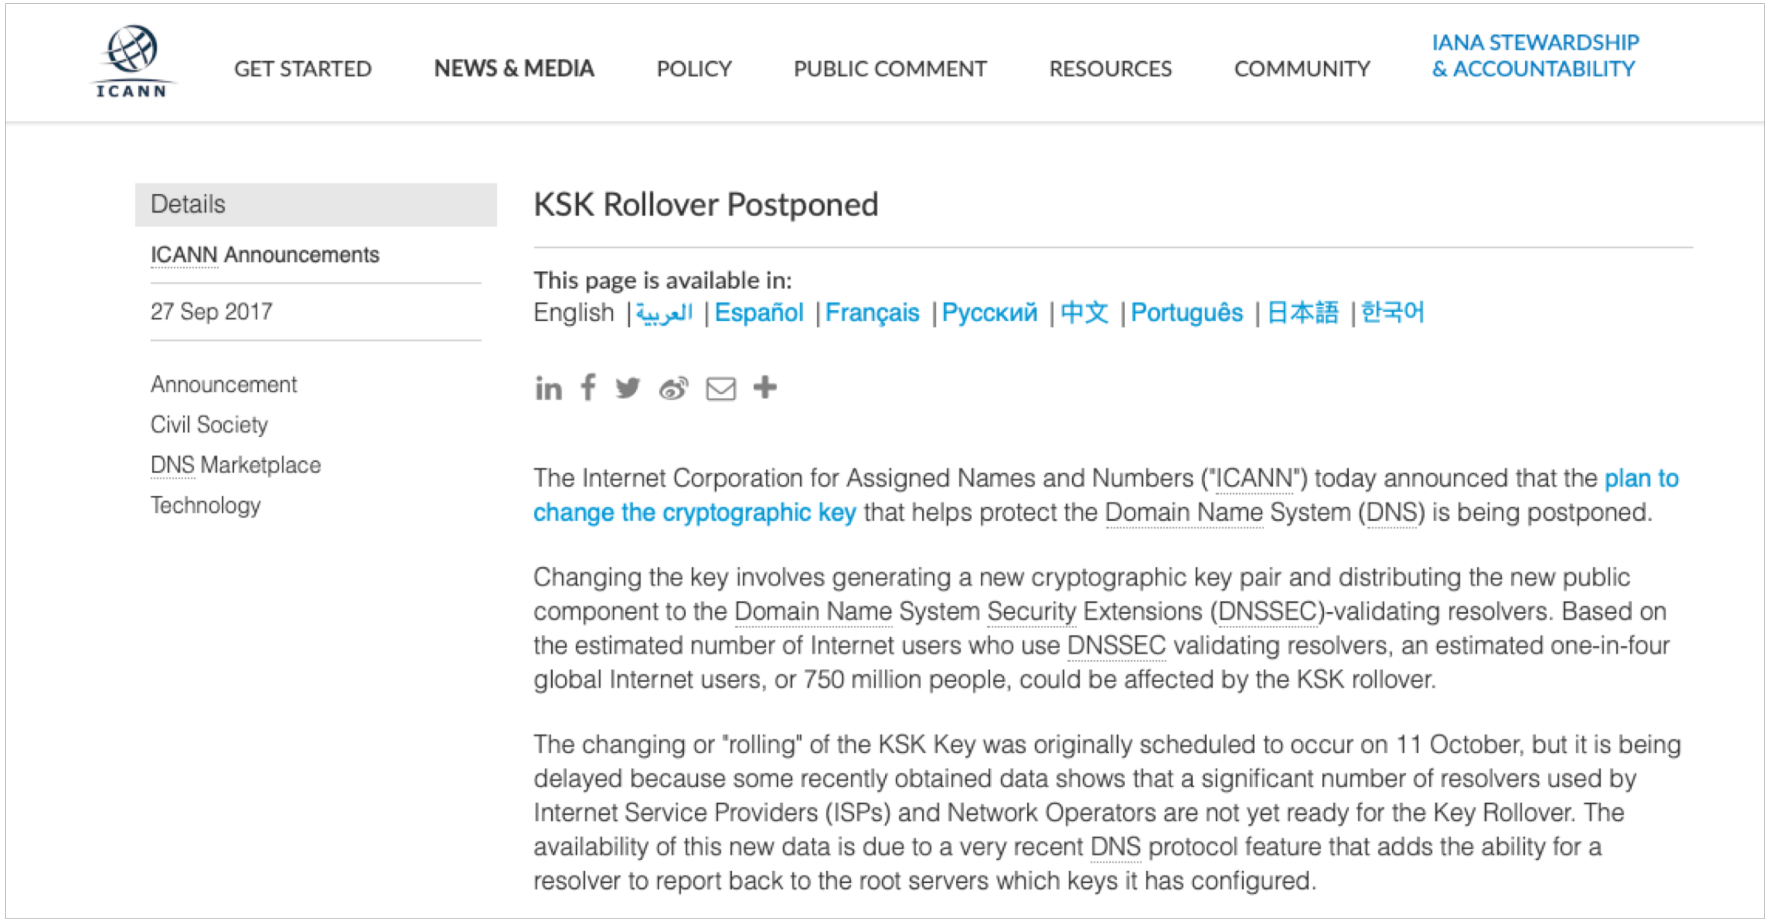

ICANN GET STARTED NEWS & MEDIA POLICY PUBLIC COMMENT RESOURCES COMMUNITY IANA STEWARDSHIP & ACCOUNTABILITY

Details

ICANN Announcements

27 Sep 2017

Announcement
Civil Society
DNS Marketplace
Technology

KSK Rollover Postponed

This page is available in:
English | العربية | Español | Français | Русский | 中文 | Português | 日本語 | 한국어

The Internet Corporation for Assigned Names and Numbers ("ICANN") today announced that the plan to change the cryptographic key that helps protect the Domain Name System (DNS) is being postponed.

Changing the key involves generating a new cryptographic key pair and distributing the new public component to the Domain Name System Security Extensions (DNSSEC)-validating resolvers. Based on the estimated number of Internet users who use DNSSEC validating resolvers, an estimated one-in-four global Internet users, or 750 million people, could be affected by the KSK rollover.

The changing or "rolling" of the KSK Key was originally scheduled to occur on 11 October, but it is being delayed because some recently obtained data shows that a significant number of resolvers used by Internet Service Providers (ISPs) and Network Operators are not yet ready for the Key Rollover. The availability of this new data is due to a very recent DNS protocol feature that adds the ability for a resolver to report back to the root servers which keys it has configured.

The Plan - V2

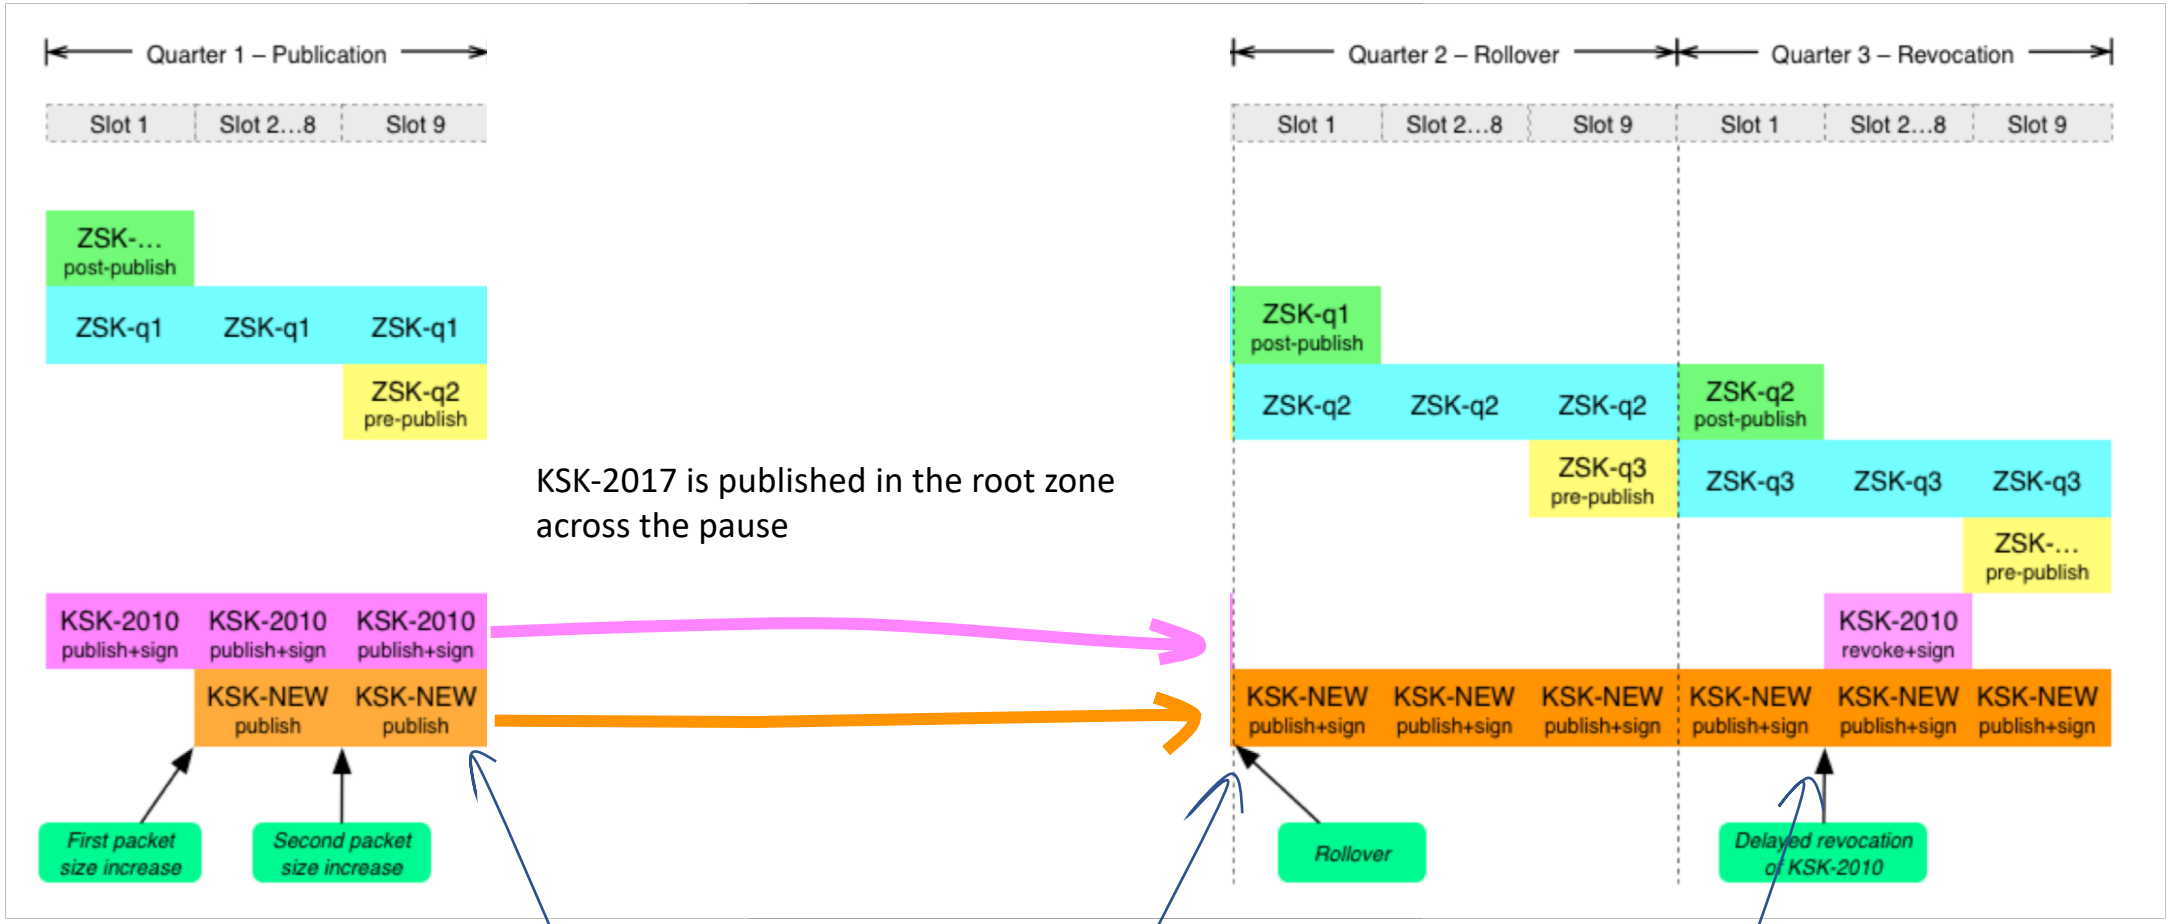

KSK-2017 is published in the root zone across the pause

27 September 2017

11 October 2018

11 January 2019

Goodbye KSK-2010

[ksk-rollover] Retention of the 2010 KSK

David Prangnell david.prangnell at iana.org
Wed Apr 24 18:57:59 UTC 2019

- Previous message: [ksk-rollover] Retention of the 2010 KSK
- Next message: [ksk-rollover] Retention of the 2010 KSK
- **Messages sorted by:** [date] [thread] [subject] [author]

To Whom It May Concern,

We have carefully reviewed the recent discussions about retaining KSK-2010 beyond its scheduled lifetime to enable a possible future as-yet-undefined technique to bootstrap a validator that has been offline for an extended period. We have decided to proceed with the deletion of the KSK-2010 as scheduled on 16 May 2019 from the Key Management Facility (KMF) East and then on 14 August 2019 from the KMF West.

What Worked

The KSK was rolled

Internet-wide DNSSEC validation levels were not significantly impacted

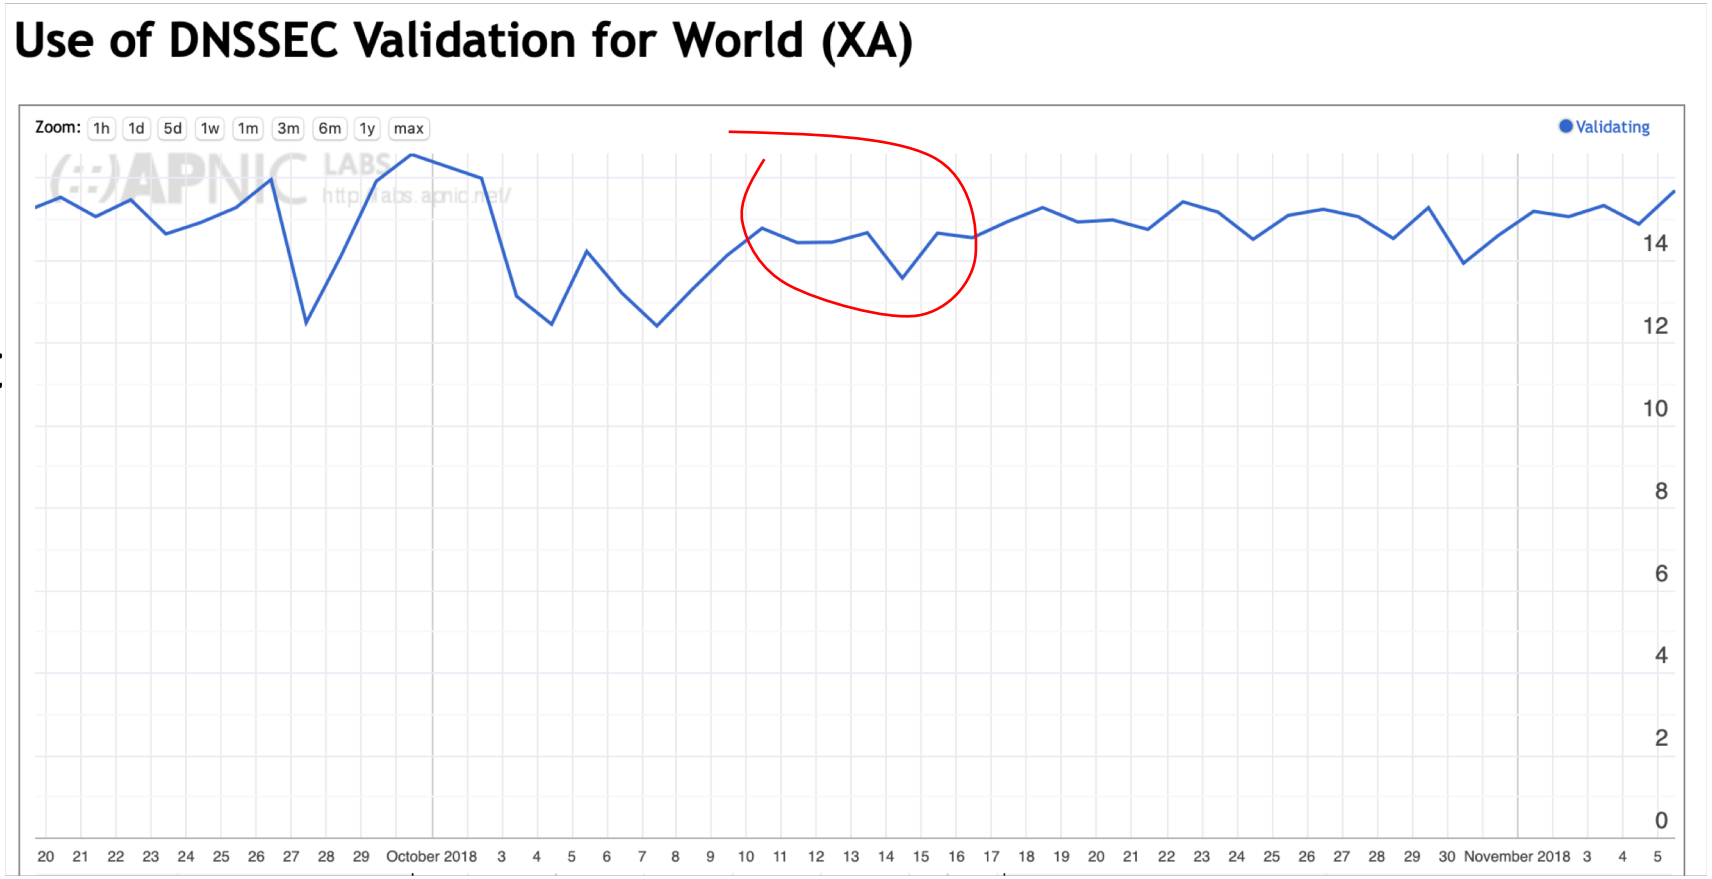

What Did Not

- The exercise was not exactly incident free
- There were issues with:
  - Predicting the impact of the KSK roll
  - Measuring the impact of the KSK roll
  - Predicting the impact of KSK revocation

KSK Measurement Objective

What number of users are at risk of being impacted by the KSK Roll?

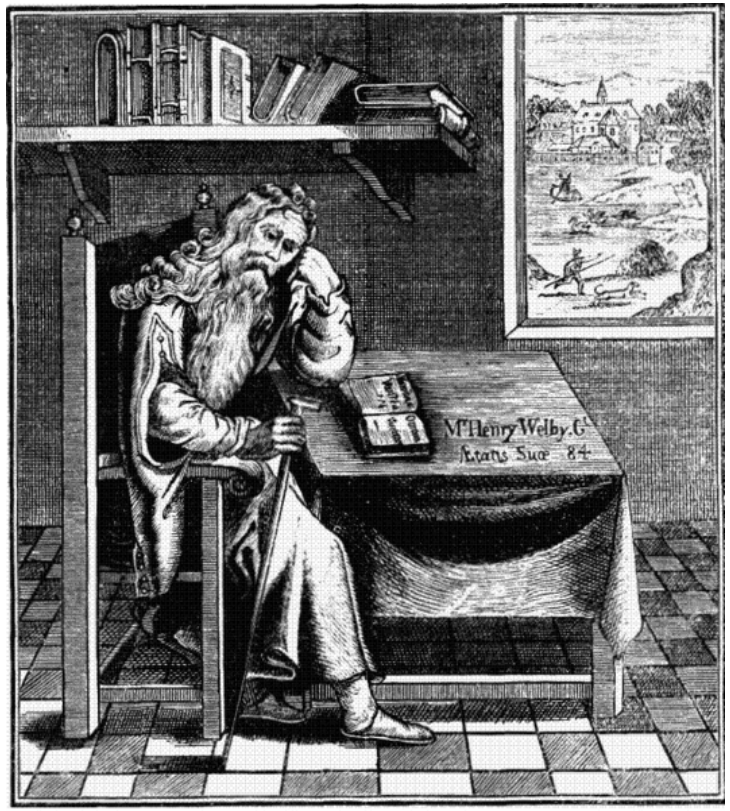

KSK Measurement Objective

What number of users are at risk of being impacted by the KSK Roll?

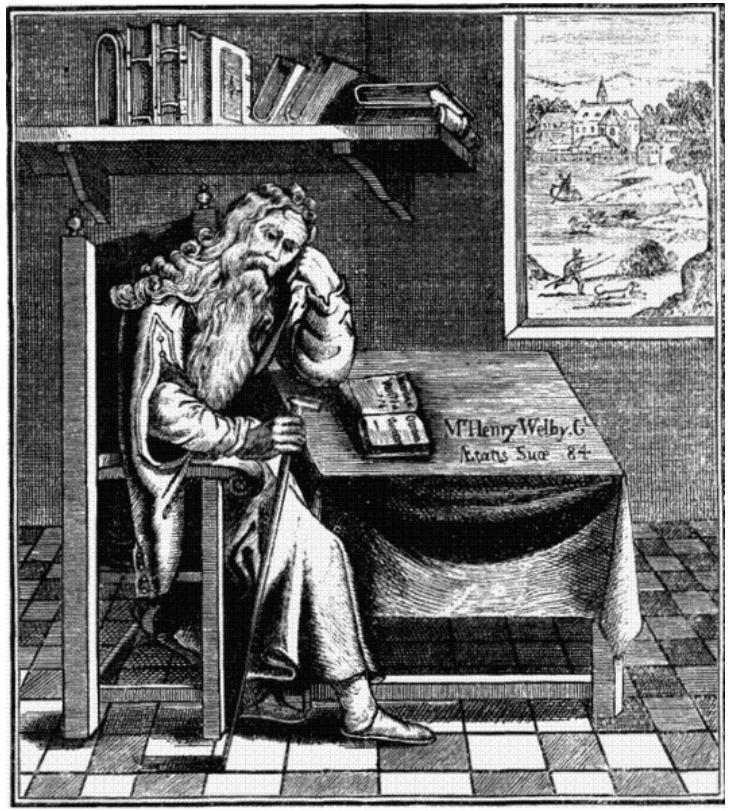

We can't do this measurement directly
We can only measure the capabilities of the DNS infrastructure and estimate the impact

Signalling via Queries

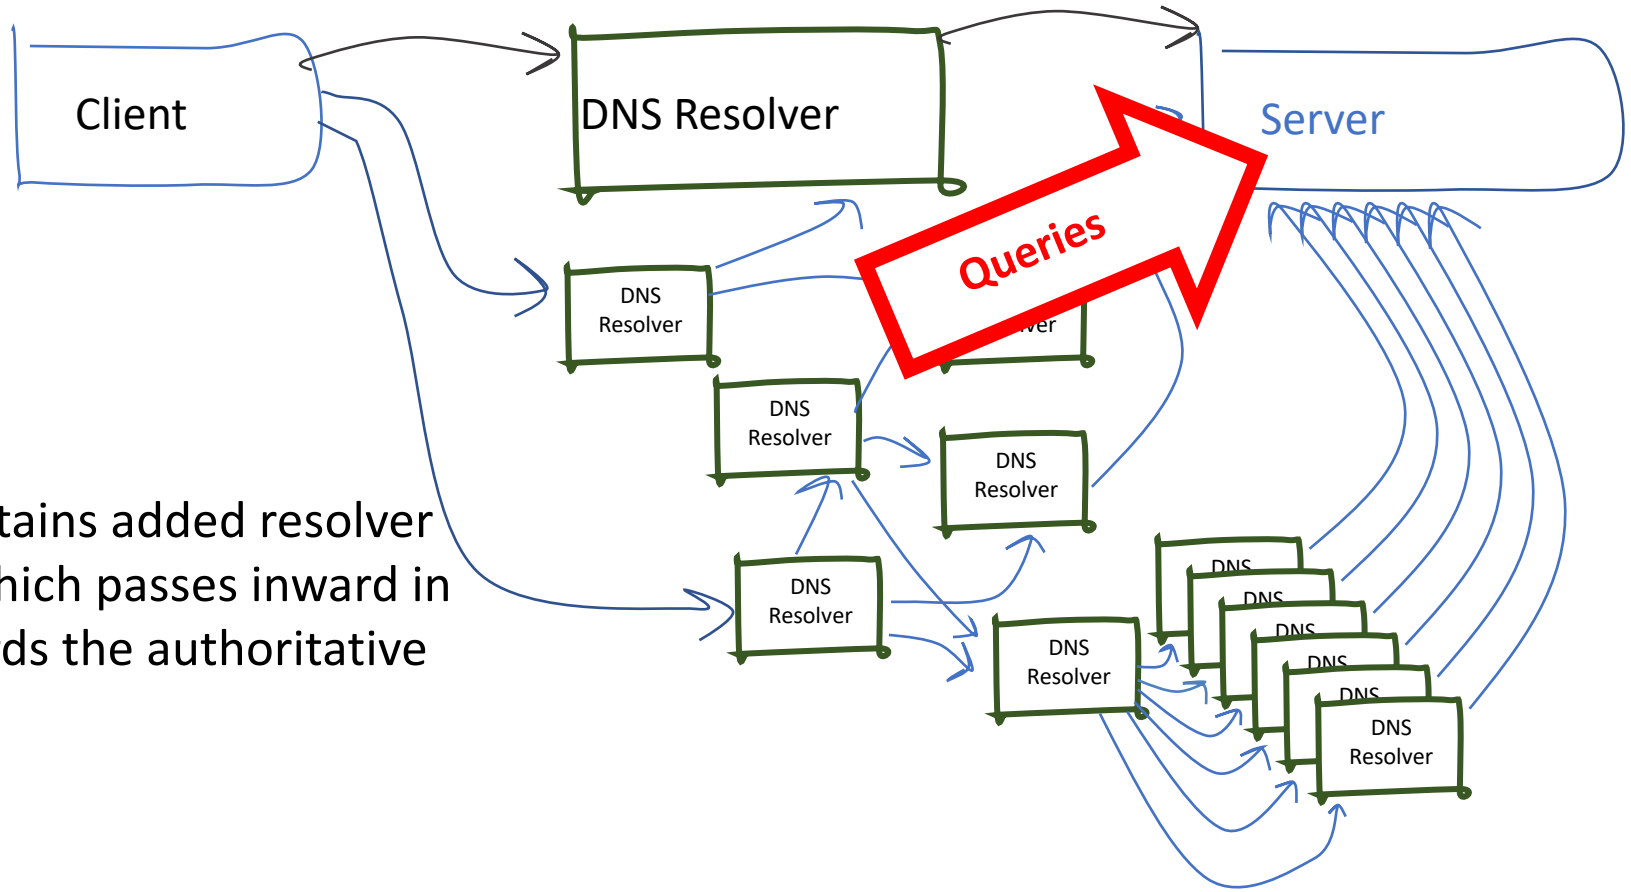

The query contains added resolver information which passes inward in the DNS towards the authoritative server(s)

Measuring Resolvers with RFC 8145

Getting resolvers to report on their local trusted key state

- A change to resolver behavior that requires deployment of new resolver code
- Resolvers that support the RFC 8145 signal mechanism periodically include the key tag of their locally trusted keys into a query directed towards the root servers

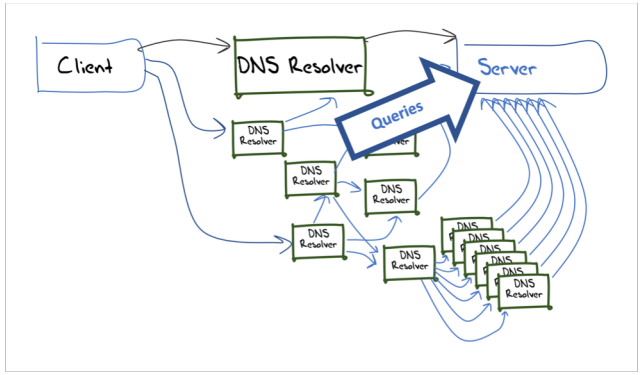

root service operators

What did ~~we~~ see at the roots?

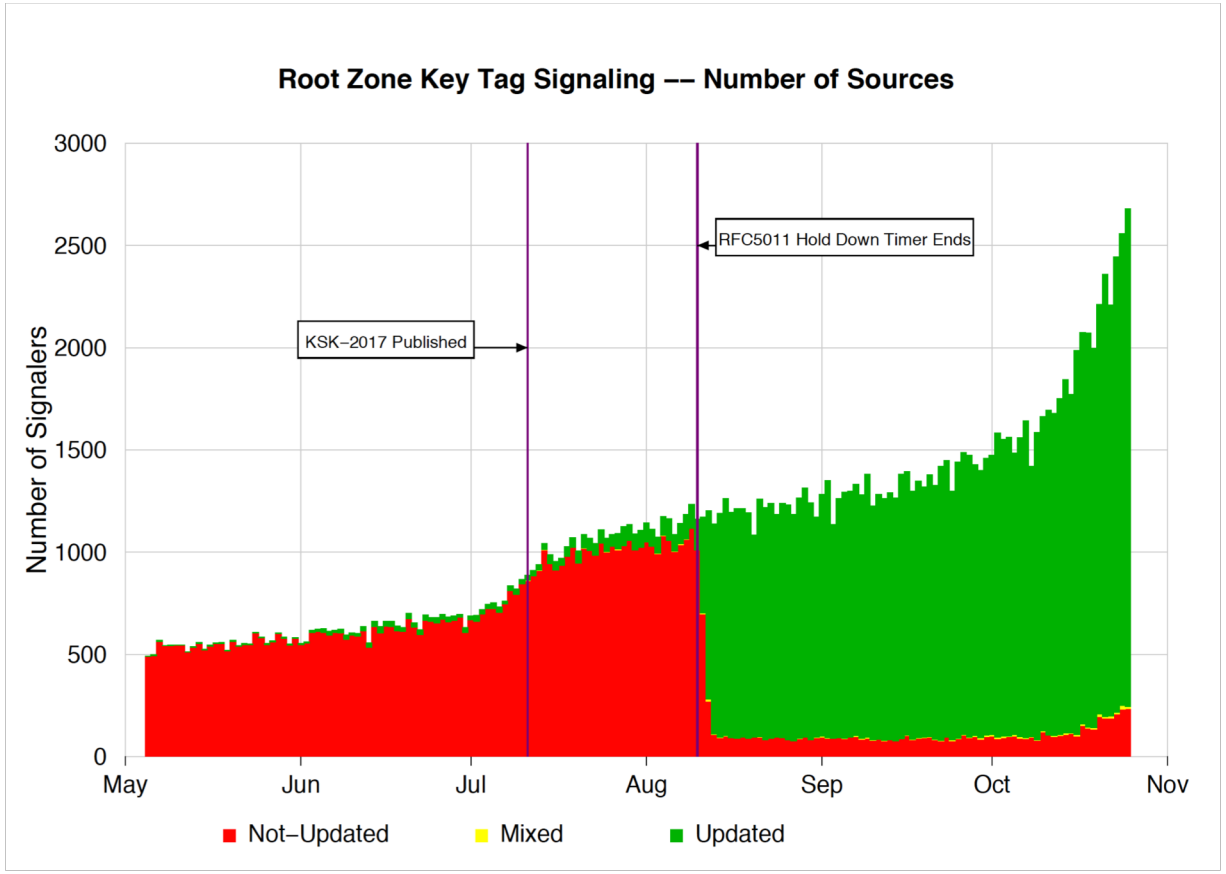

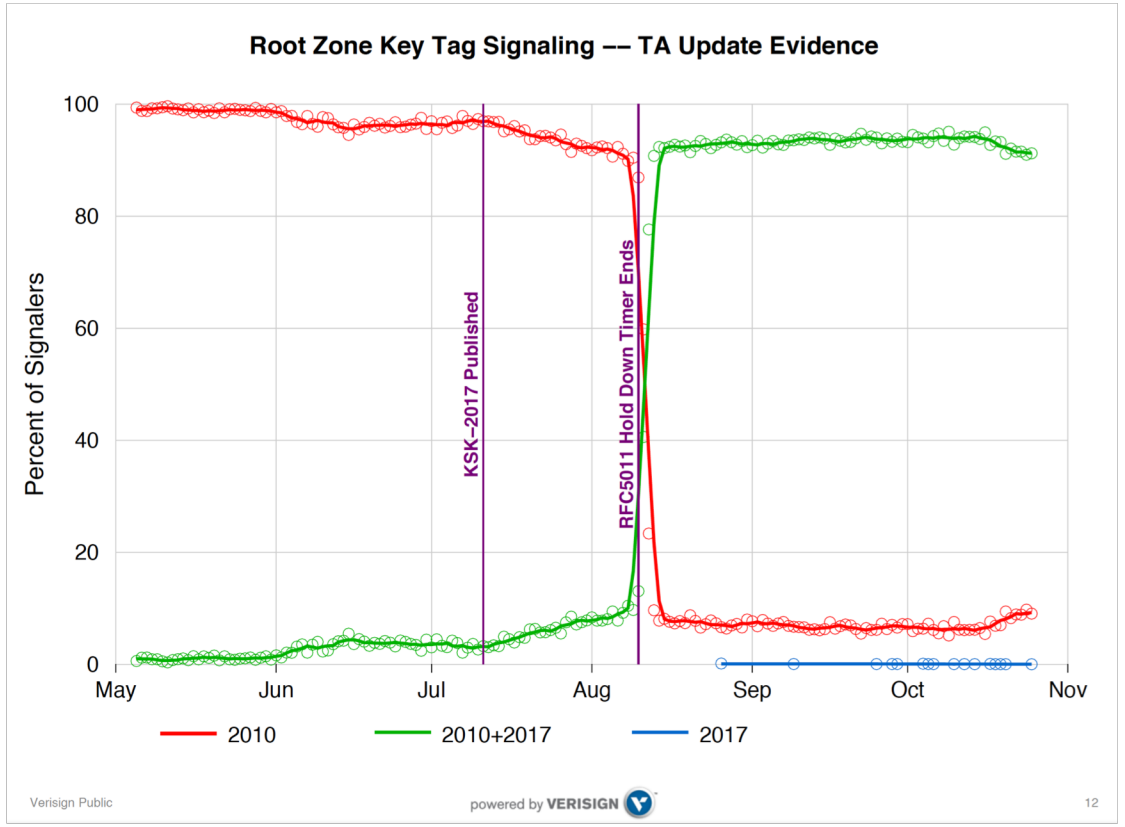

Duane Wessels VeriSign RFC 8145 Signaling Trust Anchor Knowledge In DNS Security Extensions
Presentation to DNSSEC Workshop @ ICANN 60 – 1 Nov 2017
https://schd.ws/hosted_files/icann60abudhabi2017/ea/Duane%20Wessels-VeriSign-RFC%208145-Signaling%20Trust%20Anchor%20Knowledge%20in%20DNS%20Security%20Extensions.pdf

12 months of RFC8145 signalling

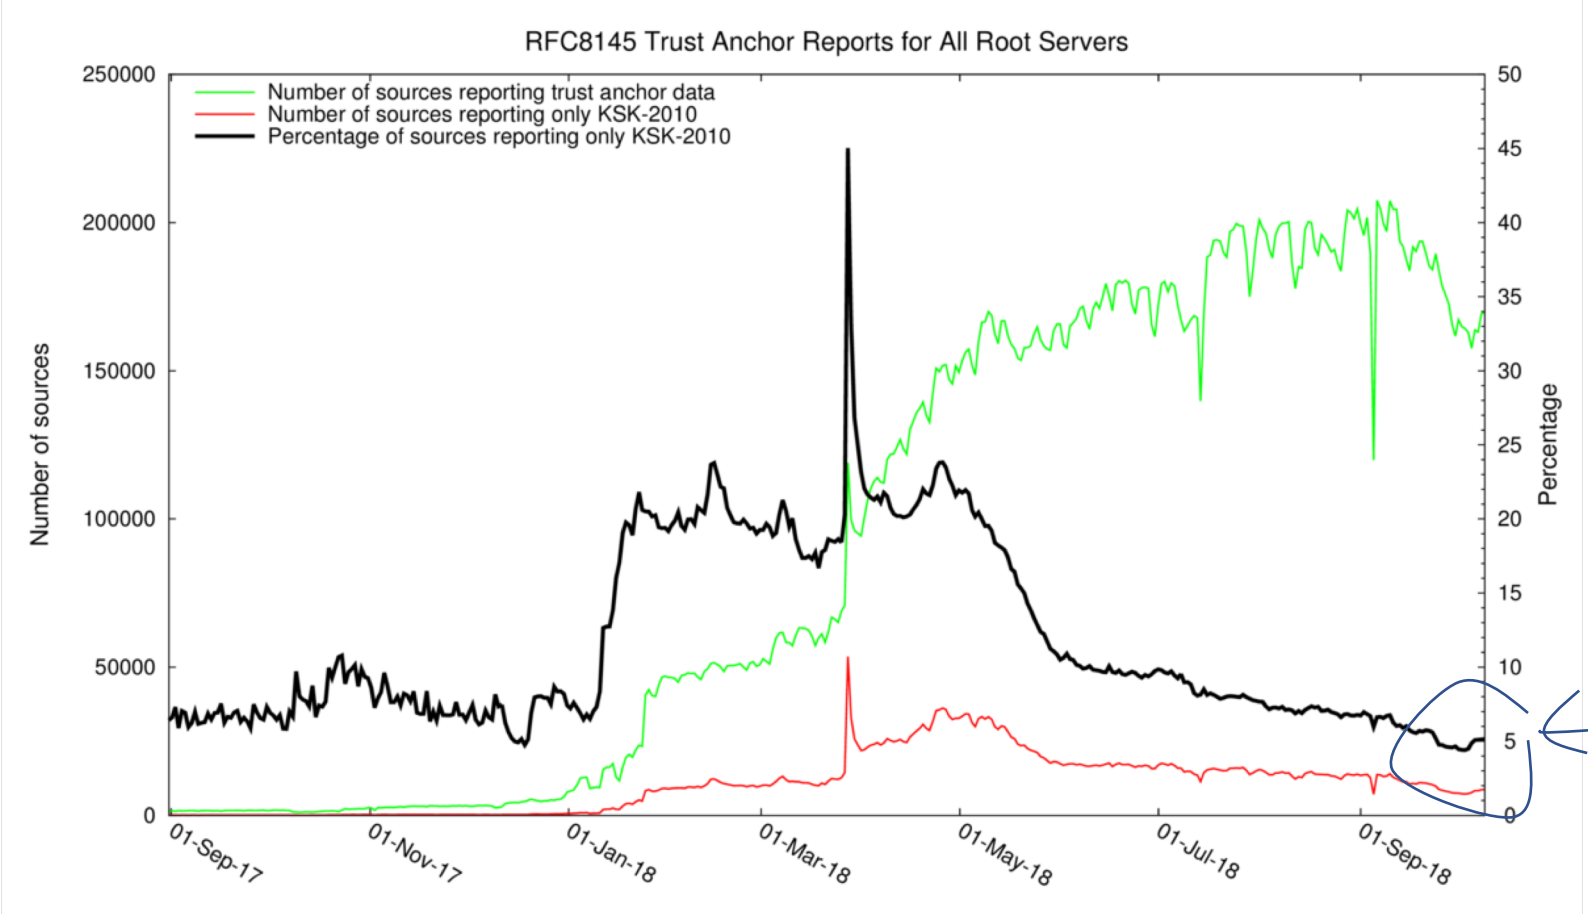

Yes, with just a few days to go this mechanism was still reporting 5% 'breakage'

http://root-trust-anchor-reports.research.icann.org

20 months of RFC8145 signalling

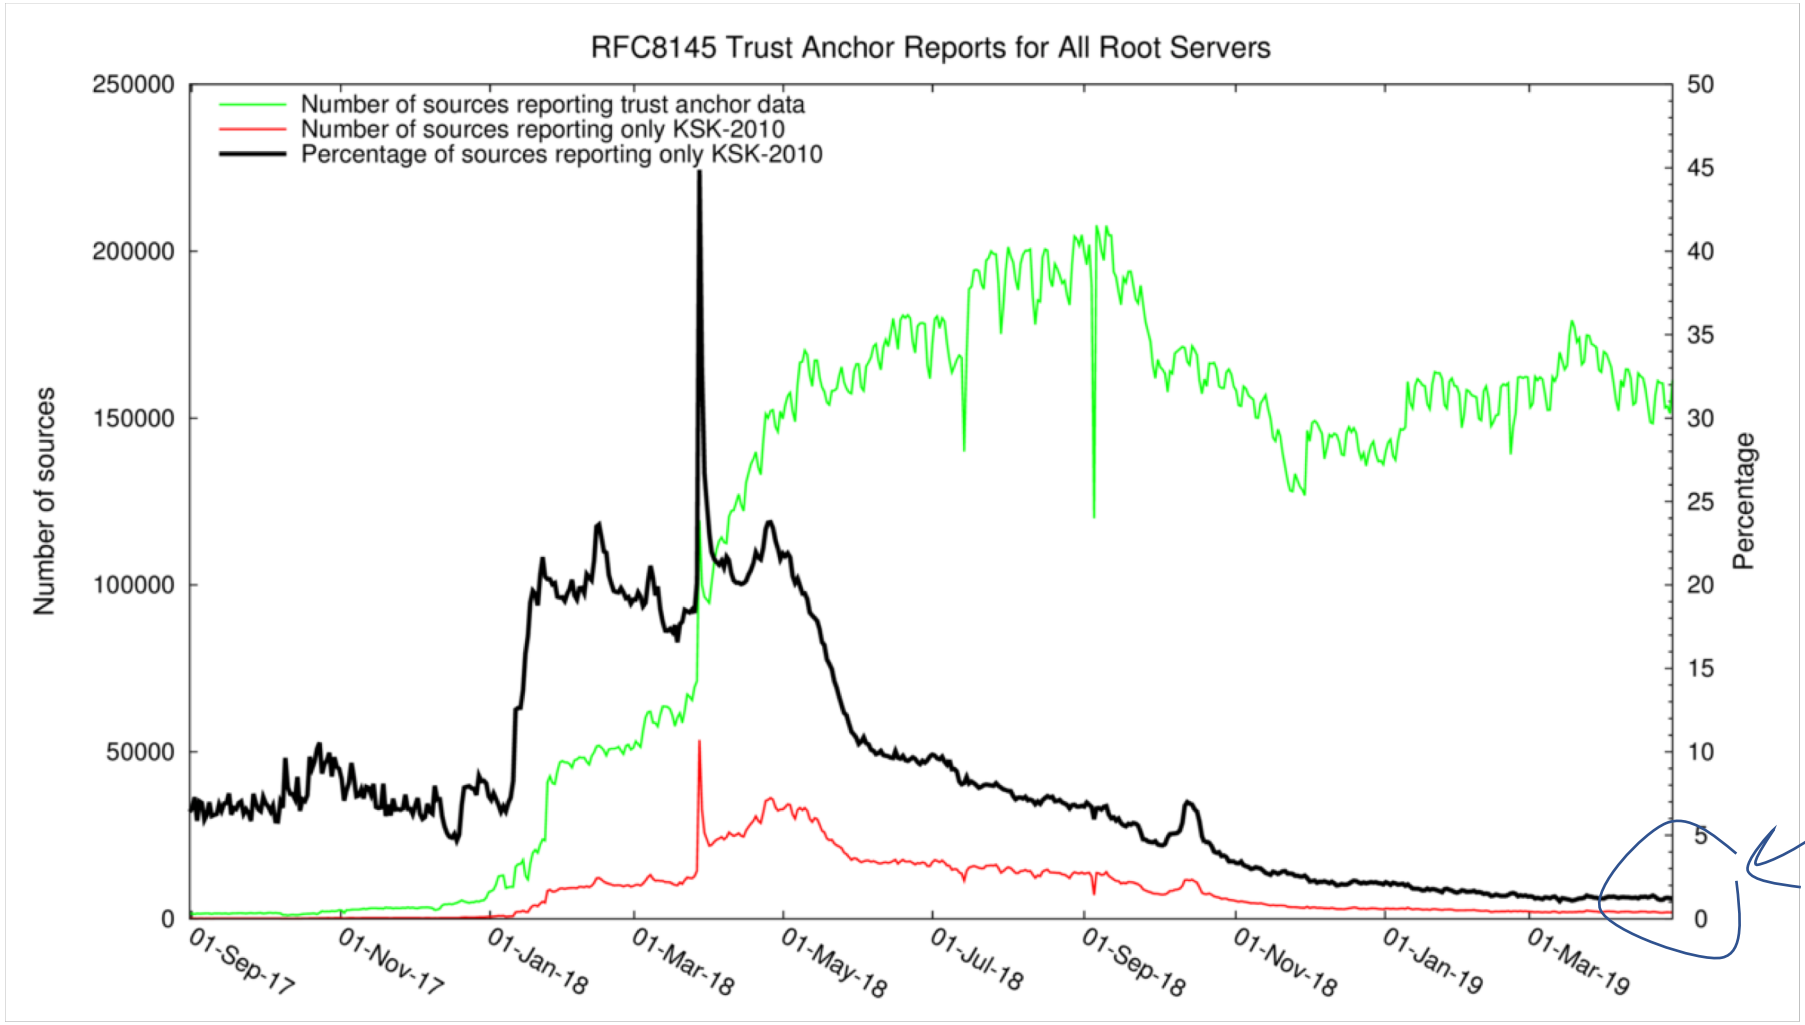

This mechanism is still reporting 1% ‘breakage’

http://root-trust-anchor-reports.research.icann.org

What is this saying?

It’s clear that there is some residual set of resolvers that are signalling that they have not yet learned to trust KSK-2017

But its **not** clear if:

- This is an accurate signal about the state of this resolver
- This is an accurate signal about the identity of this resolver
- Whether the resolver attempts DNSSEC validation
- How many users sit ‘behind’ this resolver
- Whether these uses rely solely on this resolver, or if they also have alternate resolvers that they can use

Why?

- Because the DNS does not disclose the antecedents of a query
  - If A forwards a query to B, who queries a Root Server then if the query contains an implicit signal (as in this case) then it appears that B is querying, not A
  - At no time is the user made visible in the referred query
- Because caching
  - If A and B both forward their queries via C, then it may be that one or both of these queries may be answered from C's cache
  - In this case the signal is being suppressed
- Because its actually measuring a cause, not the outcome
  - Its measuring resolvers' uptake of the new KSK, but is not able to measure the user impact of this

Signalling via Responses

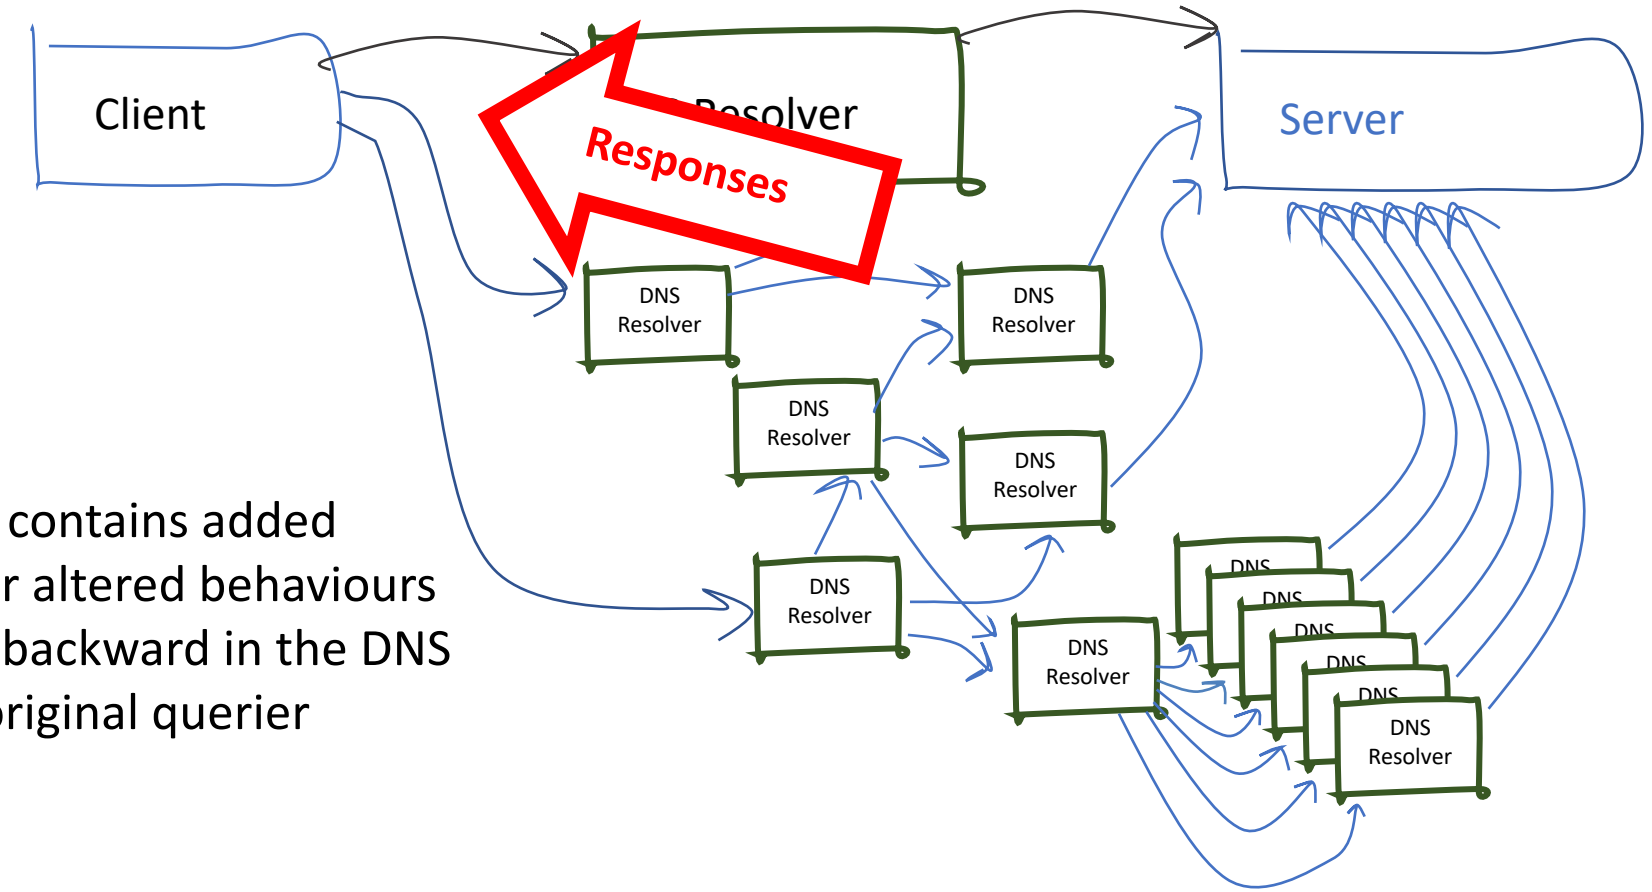

The response contains added information or altered behaviours which passes backward in the DNS towards the original querier

User-Side Measurement

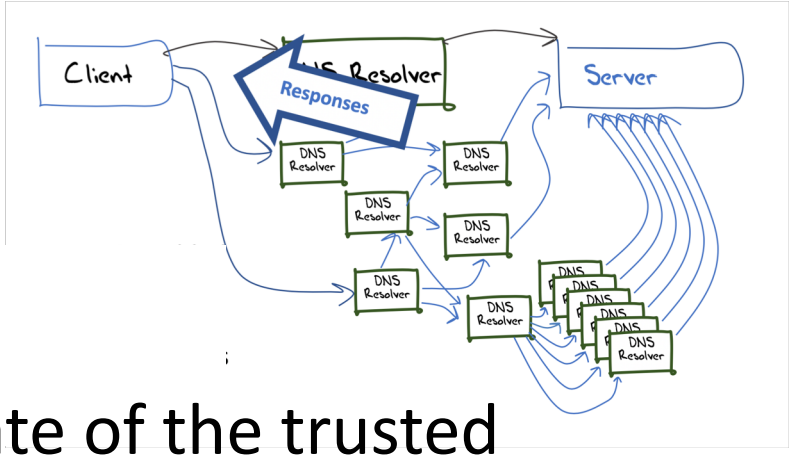

Can we devise a DNS query that could reveal the state of the trusted keys of the resolvers back to the user?

- What about a change to the resolver's reporting of validation outcome depending on the resolver's local trusted key state?
  - If a query contains the label **"root-key-sentinel-is-ta-<key-tag>**" then a validating resolver will report validation failure (SERVFAIL) if the key is NOT in the local trusted key store
  - If a query contains the label **"root-key-sentinel-not-ta-<key-tag>"** then a validating resolver will report validation failure (SERVFAIL) if the key IS in the local trusted key store

DNS + Web

- How can you tell if a user is able to resolve a DNS name?
  - Be the user (get the user to run a script of some sort)
  - Look at the DNS server AND the Web server
    - The Web object is fetched only when the DNS provides a resolution answer
    - But the opposite is not necessarily the case, so there is a noise component in such an approach

Prior to the KSK Roll

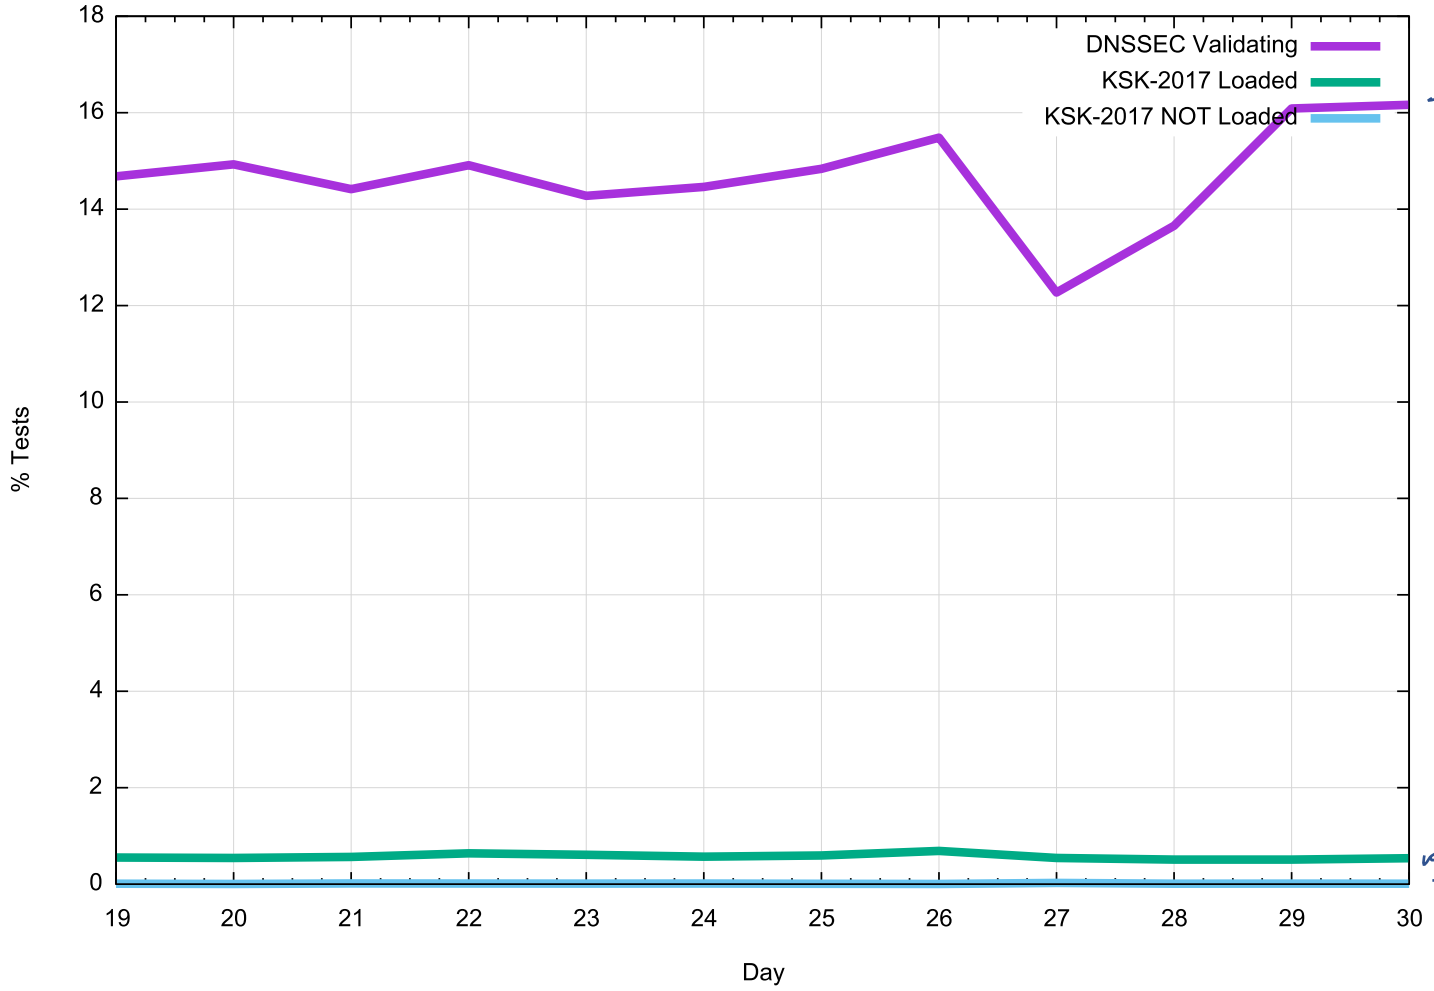

16% of users use DNSSEC-validating resolvers

15% of users do not report their KSK trust-state

0.5% of users report KSK-2017 loaded

0.005% of users report KSK-2017 NOT loaded

Possibly Affected Users

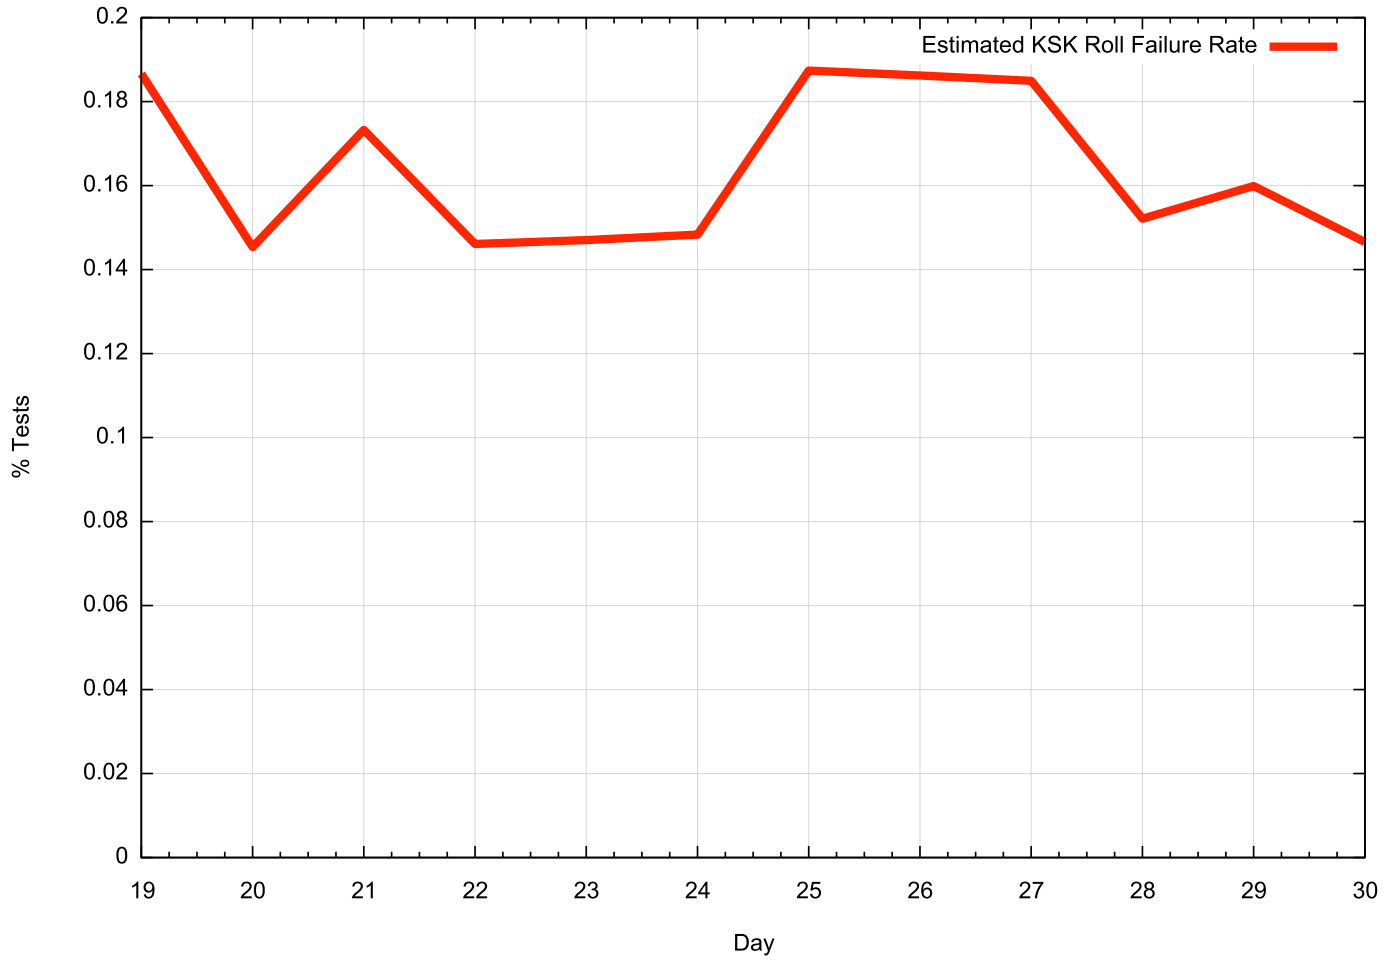

Between 0.1% to 0.2% of users are reporting that their resolvers have not loaded KSK-2017 as a trust anchor

The measurement has many uncertainties and many sources of noise so this is an upper bound of the pool of users who may encounter DNS failure due to to the KSK roll

What happened

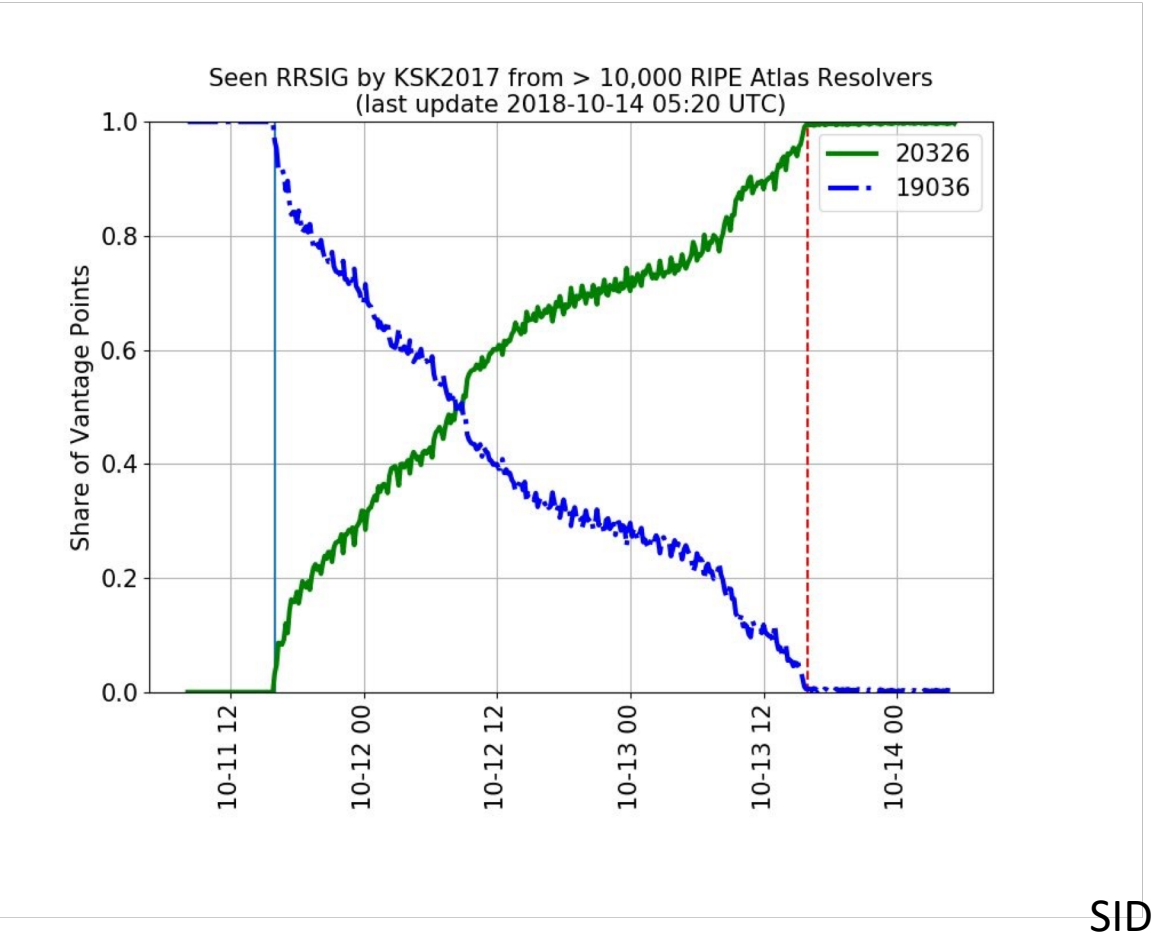

SIDN Labs Atlas Measurement

What we saw

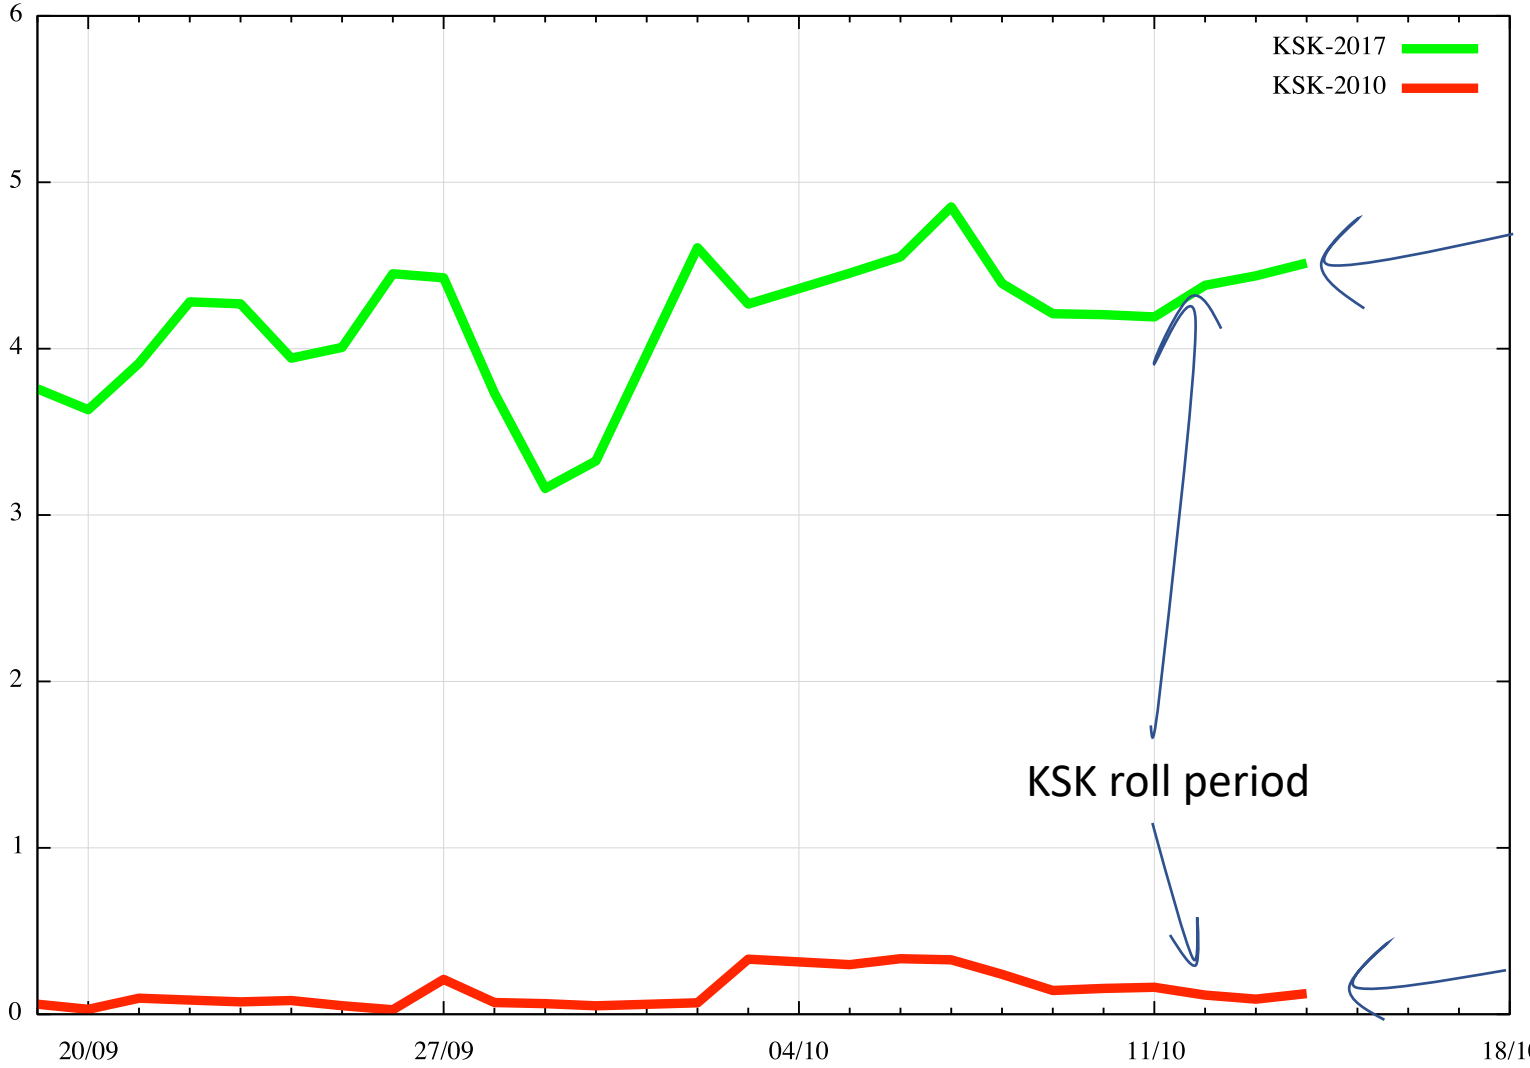

% of folk that reported “good”

KSK roll period

% of folk that reported “bad”

What we heard

Irish Examiner

IRELAND WORLD SPORT BUSINESS VIEWS LIFE PROPERTY TECH SHOWBIZ

HOT TOPICS: PITTSBURGH SYNAGOGUE ATTACK BREXIT HOMELESSNESS CLIMATE CHANGE PRESIDENTIAL EL

HOME » BREAKING NEWS » IRELAND

Eir restores broadband service saying 'we apologise again for the inconvenience'

Facebook Twitter Messenger LinkedIn WhatsApp More

Sunday, October 14, 2018 - 07:45 AM

Eir says it has resolved an internet outage that hit its service.

Customers across the country were affected by the issue late yesterday evening.

Eir has apologised to customers for the inconvenience.

In a statement released this morning, they said: "Service has been restored to those eir customers that were impacted by the internet access outage. We apologise again to our customers for the inconvenience this has caused.

"The outage was caused by a problem with an Eir DNS server that arose at approximately 14.30 on Saturday afternoon. Full service was restored around twelve hours later."

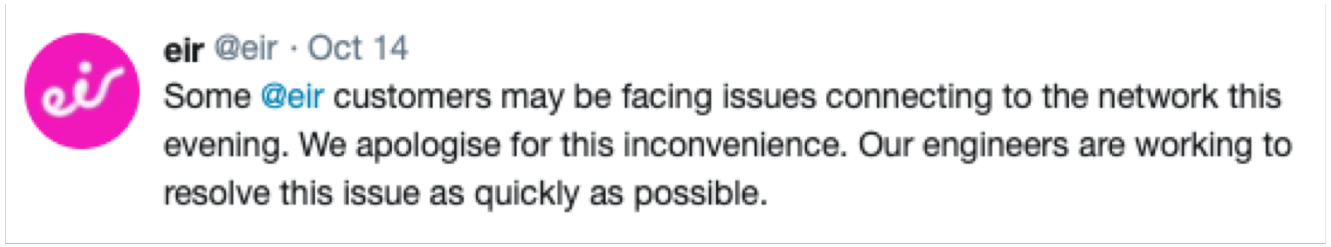

eir @eir · Oct 14

Some @eir customers may be facing issues connecting to the network this evening. We apologise for this inconvenience. Our engineers are working to resolve this issue as quickly as possible.

What happens when you lose track of the KSK?

Everything goes black

EIR - AS5466 DNSSEC Data

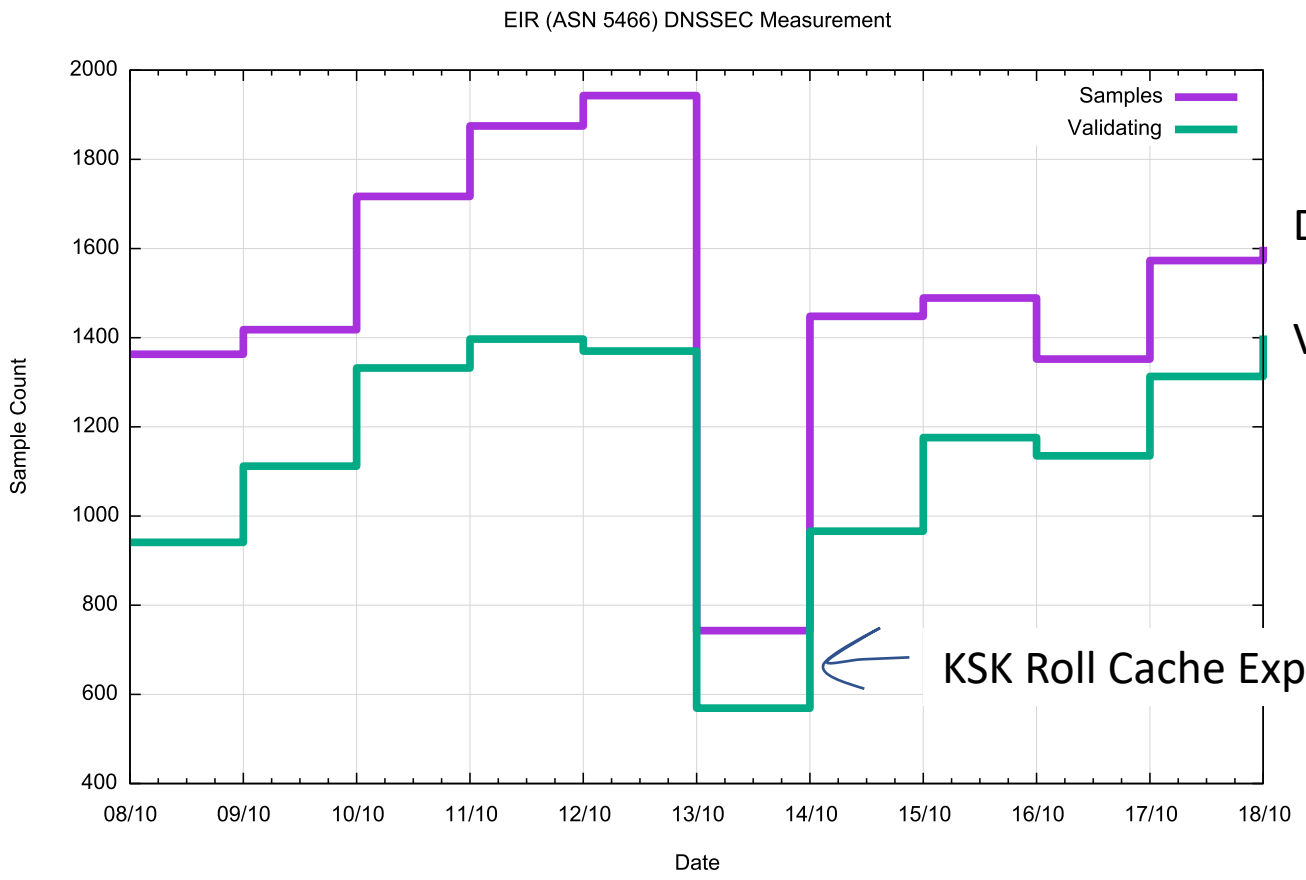

Daily Sample Count

Validating Sample Count

KSK Roll Cache Expiry

Internet DNSSEC Data

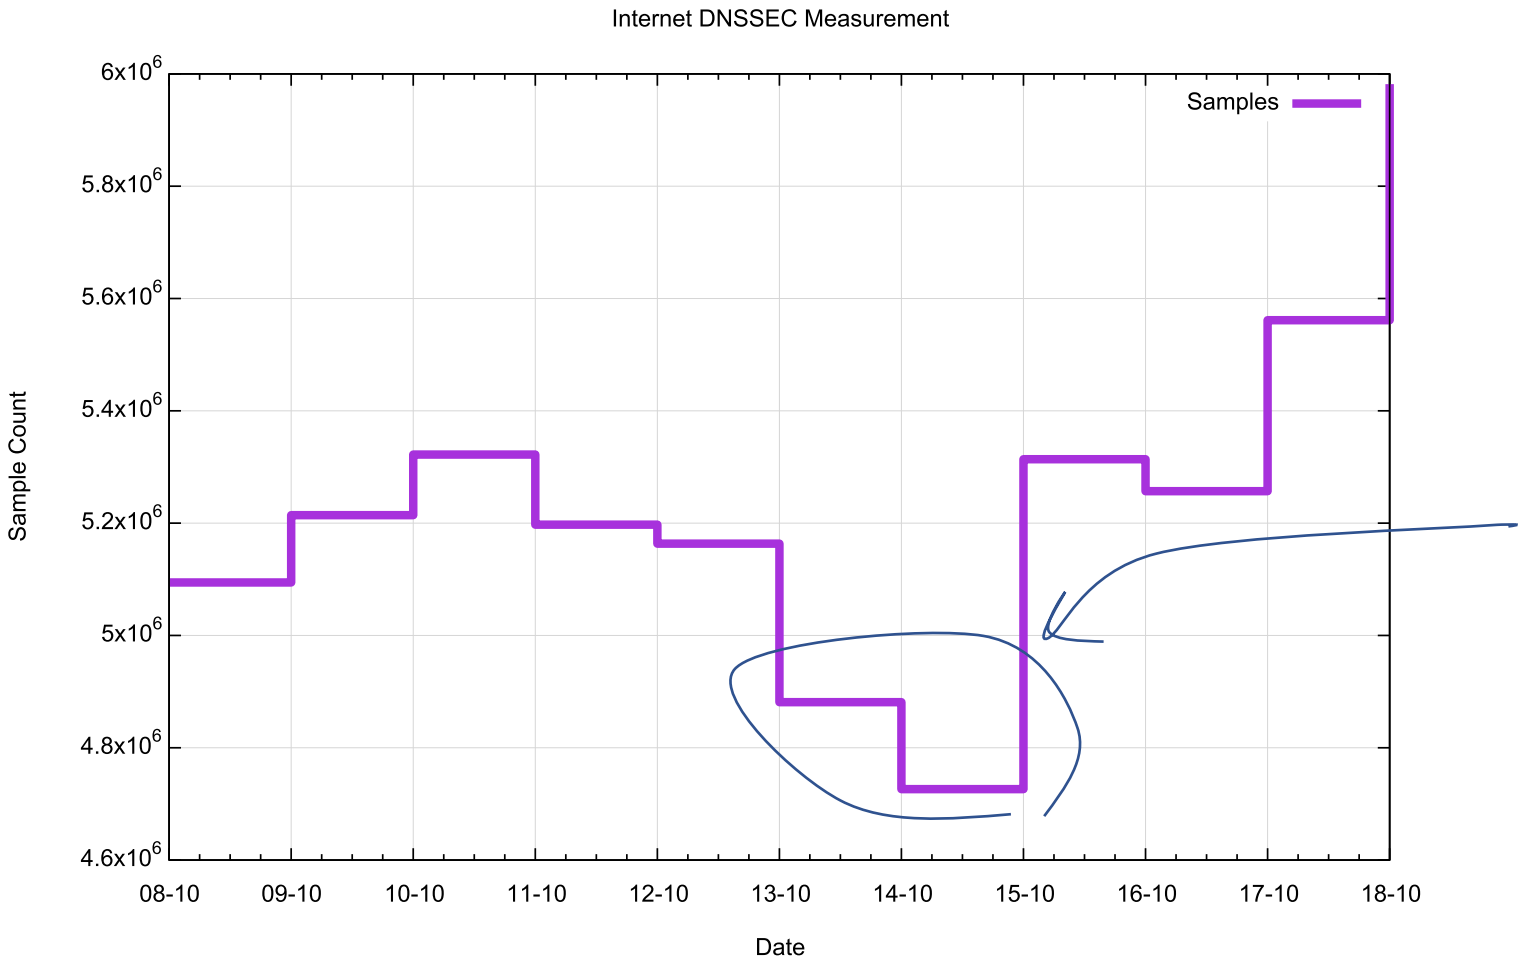

Is this part-related to the KSK Roll?

Looking for Affected Networks

- Lets use the following filter:
  - More than 400 samples / day in the lead up to the KSK roll (using weighted sample count)
  - DNSSEC validation level more than 30% prior to the KSK roll
  - Drop of more than 33% in DNSSEC validation during the KSK roll

Rank	AS	CC	Seen			Validating			As Name
			Before	During	After	Before	During	After	
1	AS2018	ZA	1,858	1,122	1,473	694	220	288	TENET, South Africa
2	AS10396	PR	1,789	1,673	1,988	1,647	276	33	COQUI-NET - DATACOM CARIBE, Puerto Rico
3	AS45773	PK	1,553	388	1,393	606	178	540	HECPERN-AS-PK PERN, Pakistan
4	AS15169	IN	1,271	438	1,286	1,209	438	1,242	GOOGLE - Google LLC, India
5	AS22616	US	1,264	503	1,526	883	377	1,014	ZSCALER- SJC, US
6	AS53813	IN	1,213	689	1,862	1,063	582	1,419	ZSCALER, India
7	AS1916	BR	1,062	94	991	326	37	277	Rede Nacional de Ensino e Pesquisa, Brazil
8	AS9658	PH	931	281	842	440	136	404	ETPI-IDS-AS-AP Eastern Telecoms, Philippines
9	AS37406	SS	888	486	972	582	365	599	RCS, South Sudan
10	AS263327	BR	882	345	438	776	289	359	ONLINE SERVICOS DE TELECOMUNICACOES, Brazil
11	AS17557	PK	835	430	777	431	277	413	Pakistan Telecommunication, Pakistan
12	AS36914	KE	834	476	937	583	354	670	KENET , Kenya
13	AS327687	UG	802	473	834	390	189	332	RENU, Uganda
14	AS680	DE	773	966	1,332	268	117	289	DFN Verein zur Foerderung, Germany
15	AS201767	UZ	761	538	729	461	200	371	UZMOBILE, Uzbekistan
16	AS37682	NG	695	401	728	593	274	568	TIZETI, Nigeria
17	AS7470	TH	674	214	507	219	94	182	True Internet, Thailand
18	AS51167	DE	670	378	479	214	78	156	CONTABO, Germany
19	AS15525	PT	600	260	593	287	125	284	MEO-EMPRESAS, Portugal
20	AS14061	GB	594	468	672	260	169	313	DigitalOcean, United Kingdom
21	AS37130	ZA	585	5	464	414	0	260	SITA, South Africa
22	AS30998	NG	583	264	484	192	54	143	NAL, Nigeria
23	AS135407	PK	569	227	457	419	207	344	TES-PL-AS-AP Trans World, Pakistan
24	AS16814	AR	565	235	456	258	120	208	NSS, Argentina
25	AS132335	IN	563	17	30	538	17	23	NETWORK-LEAPSWITCH-IN LeapSwitch Networks, India
26	AS5438	TN	559	532	579	526	171	27	ATI,Tunisia
27	AS5466	IE	547	240	401	419	184	329	EIRCOM Internet House, IE Ireland
28	AS18002	IN	538	467	614	277	176	242	WORLDPHONE-IN AS, India
29	AS37209	NG	532	109	438	269	45	194	HYPERIA, Nigeria
30	AS37100	ZA	454	161	401	168	95	131	SEACOM-AS, South Africa
31	AS5588	CZ	453	175	430	186	102	162	GTSCE GTS Central Europe, Czechia
32	AS1103	NL	446	38	363	189	7	132	SURFnet, The Netherlands
33	AS17563	PK	402	117	359	207	64	199	Nexlinx, Pakistan
34	AS327724	UG	401	120	538	208	103	266	NITA, Uganda
35	AS7590	PK	400	122	329	266	84	224	COMSATS, Pakistan

Rank	AS	CC	Seen			Validating			As Name
			Before	During	After	Before	During	After	
1	AS2018	ZA	1,858	1,122	1,473	694	220	288	TENET, South Africa
2	AS10396	PR	1,789	1,673	1,988	1,647	276	33	COQUI-NET - DATACOM CARIBE, Puerto Rico
3	AS45773	PK	1,553	388	1,393	606	178	540	HECPERN-AS-PK PERN, Pakistan
4	AS15169	IN	1,271	438	1,286	1,209	438	1,242	GOOGLE - Google LLC, India
5	AS22616	US	1,264	503	1,526	883	377	1,014	ZSCALER- SJC, US
6	AS53813	IN	1,213	689	1,862	1,063	582	1,419	ZSCALER, India
7	AS1916	BR	1,062	94	991	326	37	277	Rede Nacional de Ensino e Pesquisa, Brazil
8	AS9658	PH	931	281	842	440	136	404	ETPI-IDS-AS-AP Eastern Telecoms, Philippines
9	AS37406	SS	888	486	972	582	365	599	RCS, South Sudan
10	AS263327	BR	882	345	438	776	289	359	ONLINE SERVICOS DE TELECOMUNICACOES, Brazil
11	AS17557	PK	835	430	777	431	277	413	Pakistan Telecommunication, Pakistan
12	AS36914	KE	834	476	937	583	354	670	KENET , Kenya
13	AS327687	UG	802	473	834	390	189	332	RENU, Uganda
14	AS680	DE	773	966	1,332	268	117	289	DFN Verein zur Foerderung, Germany
15	AS201767	UZ	761	538	729	461	200	371	UZMOBILE, Uzbekistan
16	AS37682	NG	695	401	728	593	274	568	TIZETI, Nigeria
17	AS7470	TH	674	214	507	219	94	182	True Internet, Thailand
18	AS51167	DE	670	378	479	214	78	156	CONTABO, Germany
19	AS15525	PT	600	260	593	287	125	284	MEO-EMPRESAS, Portugal
20	AS14061	GB	594	468	672	260	169	313	DigitalOcean, United Kingdom
21	AS37130	ZA	585	5	464	414	0	260	SITA, South Africa
22	AS30998	NG	583	264	484	192	54	143	NAL, Nigeria
23	AS135407	PK	569	227	457	419	207	344	TES-PL-AS-AP Trans World, Pakistan
24	AS16814	AR	565	235	456	258	120	208	NSS, Argentina
25	AS132335	IN	563	17	30	538	17	23	NETWORK-LEAPSWITCH-IN LeapSwitch Networks, India
26	AS5438	TN	559	532	579	526	171	27	ATI,Tunisia
27	AS5466	IE	547	240	401	419	184	329	EIRCOM Internet House, IE Ireland
28	AS18002	IN	538	467	614	277	176	242	WORLDPHONE-IN AS, India
29	AS37209	NG	532	109	438	269	45	194	HYPERIA, Nigeria
30	AS37100	ZA	454	161	401	168	95	131	SEACOM-AS, South Africa
31	AS5588	CZ	453	175	430	186	102	162	GTSCE GTS Central Europe, Czechia
32	AS1103	NL	446	38	363	189	7	132	SURFnet, The Netherlands
33	AS17563	PK	402	117	359	207	64	199	Nexlinx, Pakistan
34	AS327724	UG	401	120	538	208	103	266	NITA, Uganda
35	AS7590	PK	400	122	329	266	84	224	COMSATS, Pakistan

These networks turned DNSSEC validation off!

Impact of the KSK Roll

- The immediate impact appears to be some 0.2% - 0.3% of users
- In 32 ISP cases service was restored with DNSSEC validation enabled
- In 3 ISP cases DNSSEC validation was turned off

But that's not the end of the story…

- The next event was the revocation of KSK-2010 at 1400 UTC, 11 January 2019
  - This was meant to be easy
  - It required no trust transition on the part of DNSSEC-validating resolvers
  - KSK-2010 was published as a signing KEY for DNSSEC with the ”revoke bit” set in the key flags
  - While the DNSKEY response was large (1,449 octets) other parts of the DNS generate larger responses for validating resolvers

And for clients the revocation was uneventful

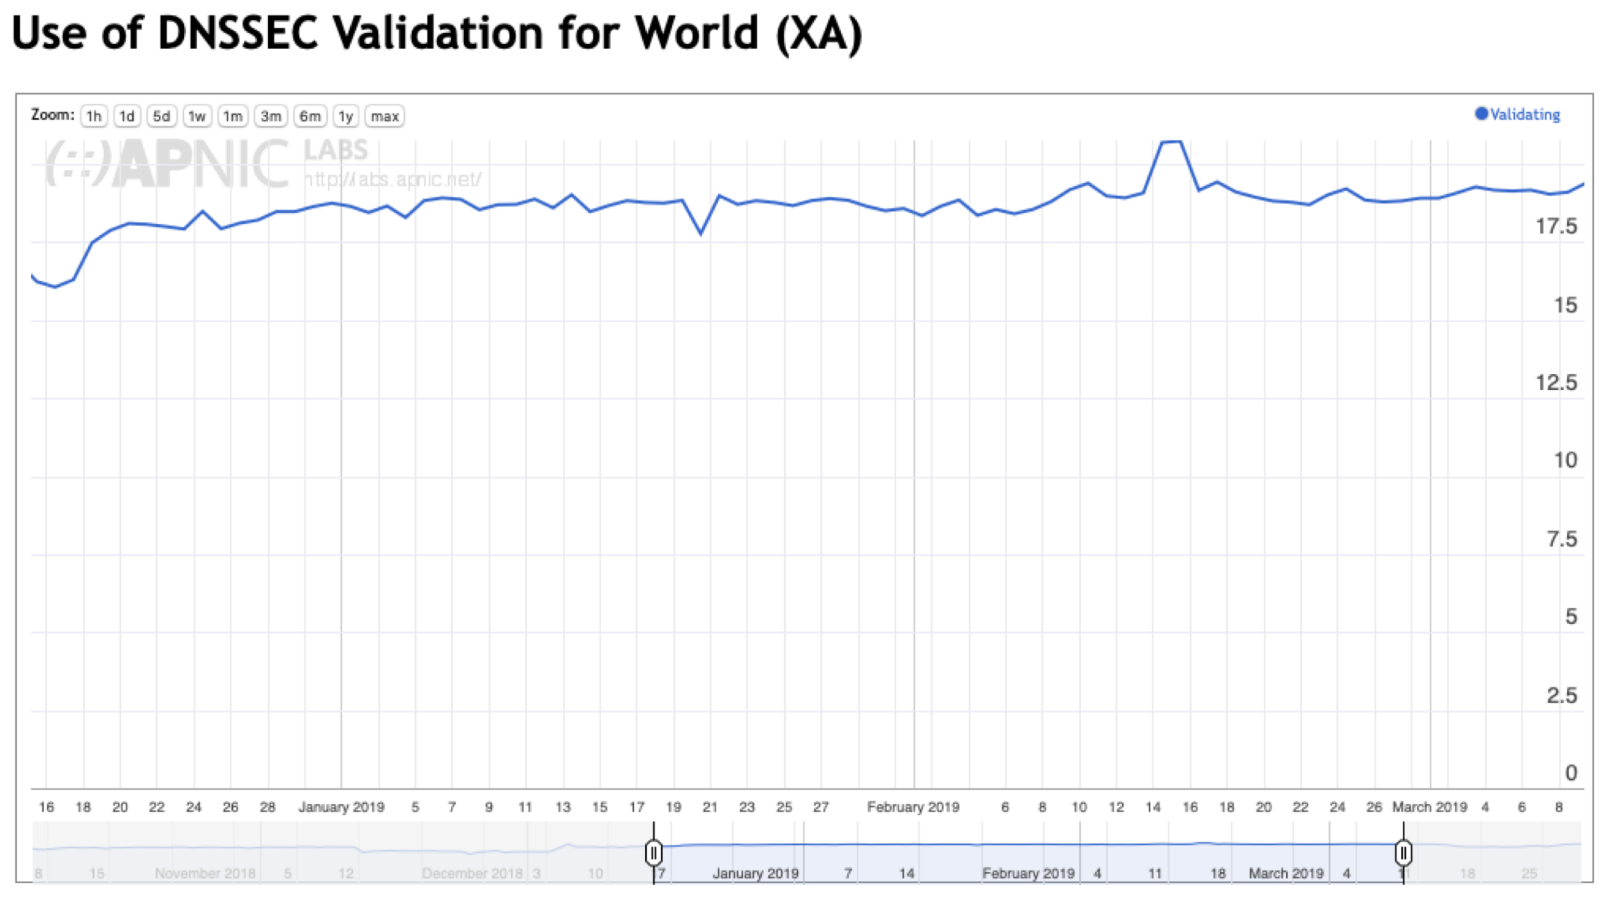

But Root Servers reported a different story

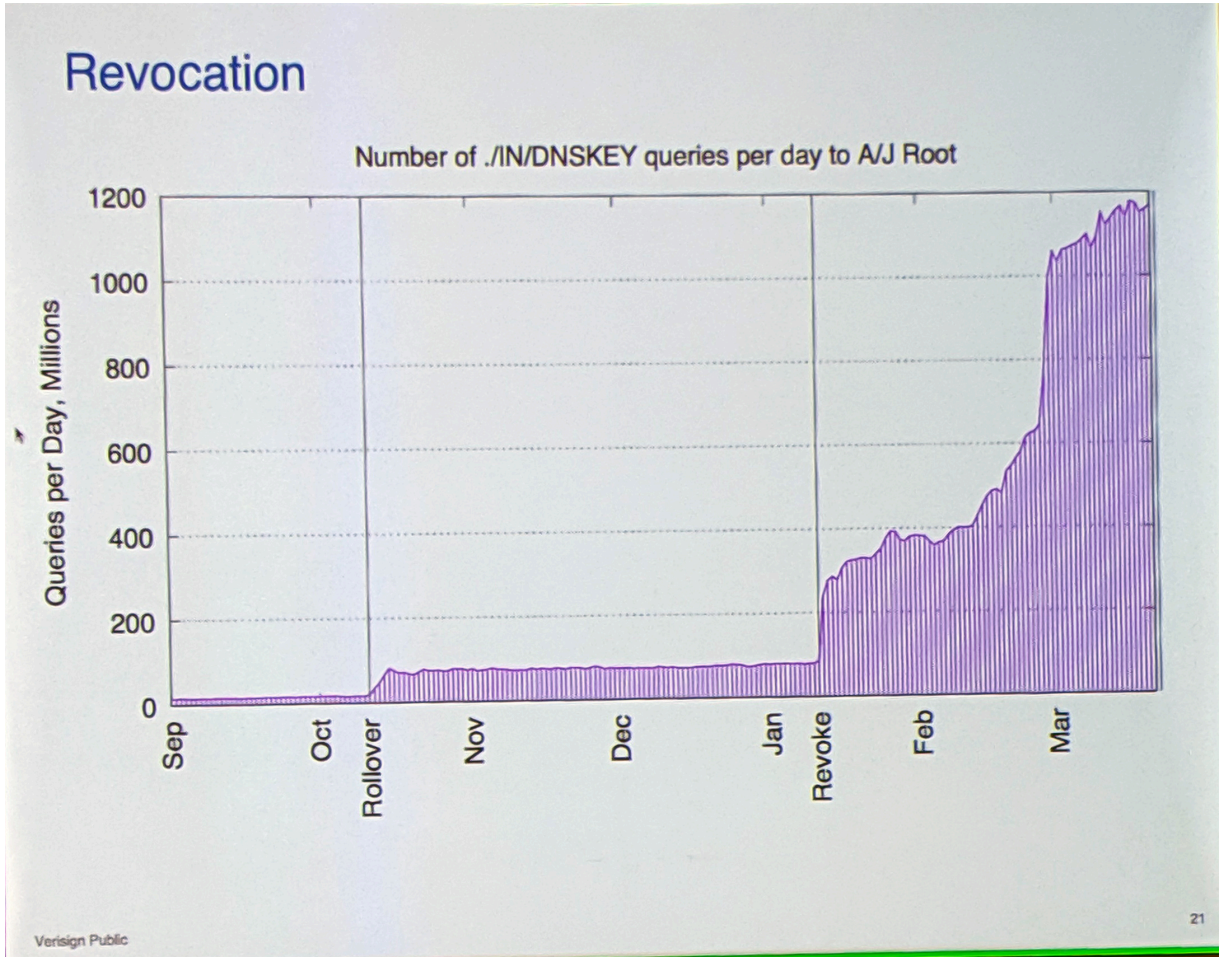

Duane Wessels, KSK Roller Post Analysis, DNS OARC, May 2019

Why?

Query spin in old versions of a popularly deployed resolver

☆ **Evan Hunt <each@isc.org>**

Re: [ksk-rollover] Description of my analysis of the too-many-KSK queries problem

To: Wes Hardaker <wjhns1@hardakers.net>, Cc: Salz, Rich via ksk-rollover <ksk-rollover@icann.org>

Thank you for this analysis.

On Wed, Apr 03, 2019 at 01:56:14PM -0700, Wes Hardaker wrote:

> Evan, at the IETF, reported in a few meetings and conversations that they had discovered a bug in bind previously that would exhibit this roll-over-and-die type behavior but that it was only present in out-of-date versions of bind (9.10 and below I believe he stated).

I think we have a case of two different bugs with superficially similar effects. I haven't yet been able to reproduce yours (maybe it's specific to Fedora somehow, or maybe I just haven't hit the right combination yet). Mine causes named to go into a tight loop sending DNSKEY queries forever, starting immediately on startup. It doesn't ever quiet down, even temporarily, and it doesn't depend on incoming queries - it just spins. Once the revoked key is removed, it stops.

Based on sheer volume, I would guess this was a bigger contributor to the observed increase in DNSKEY traffic than the bug you discovered, though yours is odd, and definitely warrants further investigation.

The looping bug was fixed in 9.10.2 and (if I recall correctly) 9.9.7, and was never in the 9.11 branch. I saw a list of "version.bind" responses from servers that were sending the most DNSKEY queries, and the worst offenders were older than that.

--
Evan Hunt -- each@isc.org
Internet Systems Consortium, Inc.

---

ksk-rollover mailing list
ksk-rollover@icann.org
https://mm.icann.org/mailman/listinfo/ksk-rollover

Lessons Learned

- Yes, we can roll the KSK!
- Yes, the extensive contact campaign helped

BUT

- The DNS is VERY opaque!
- Instrumentation was extremely challenging

What Next?

Observations

- The operation was an experiment
- DNS trust state signalling (both forms) seems to add more noise rather than clarity
- We could think about making the DNS more transparent
  - But there is a clear trade-off between greater transparency and exposing end user behaviours
  - So maybe we might not want to go there!

Observations

- Is DNSSEC validation most appropriately a resolver function or an edge function?
  - Envisaged in DANE Chain Extensions in the TLS - requires edge devices to hold the current KSK value and perform DNSSEC validation
  - Is 5011 really the best way for edge devices to maintain their KSK copy?
  - Really?

If we want to think about scaling DNSSEC validation to every host device what KSK management practices will scale?

One View

- We should perform this operation often
  - Maybe we just need to keep rolling every year
  - That way we train the manual loaders to keep up!
- We should now look at an Elliptical Curve algorithm roll
- We should now look at standing a backup KSK provision

Another View

- Why are we rolling the KSK?
- Actual key compromise might not play out in the same staged manner as a planned key roll
- If these planned key rolls are not a rehearsal for some unforeseen potential calamity then why are we deliberately adding instability into DNSSEC?
- Is doing this again going to teach us anything new?
- Is old-signs-new really the best way to do this?
- How should we scale the KSK?

Thanks!