

Root KSK Roll Delay Update

CCNSO Tech Day

Roy Arends, Principal Research Scientist

30 October 2017



Background

- ⊙ When you validate DNSSEC signed DNS records, you need a Trust Anchor.
 - A Trust Anchor is a Public Key.
- ⊙ Public Keys should not live forever.
- ⊙ These Trust Anchors probably should be periodically renewed (rolled).
 - You can do this automatically or manually.
- ⊙ However, there was no way for us (ICANN) to check if you have the right key configured.
- ⊙ Therefore, a multi-year design and outreach effort ensued:
 - Design-team, blogs, outreach, presentations in various venues, plans, vendors and governments were contacted, etc., etc.

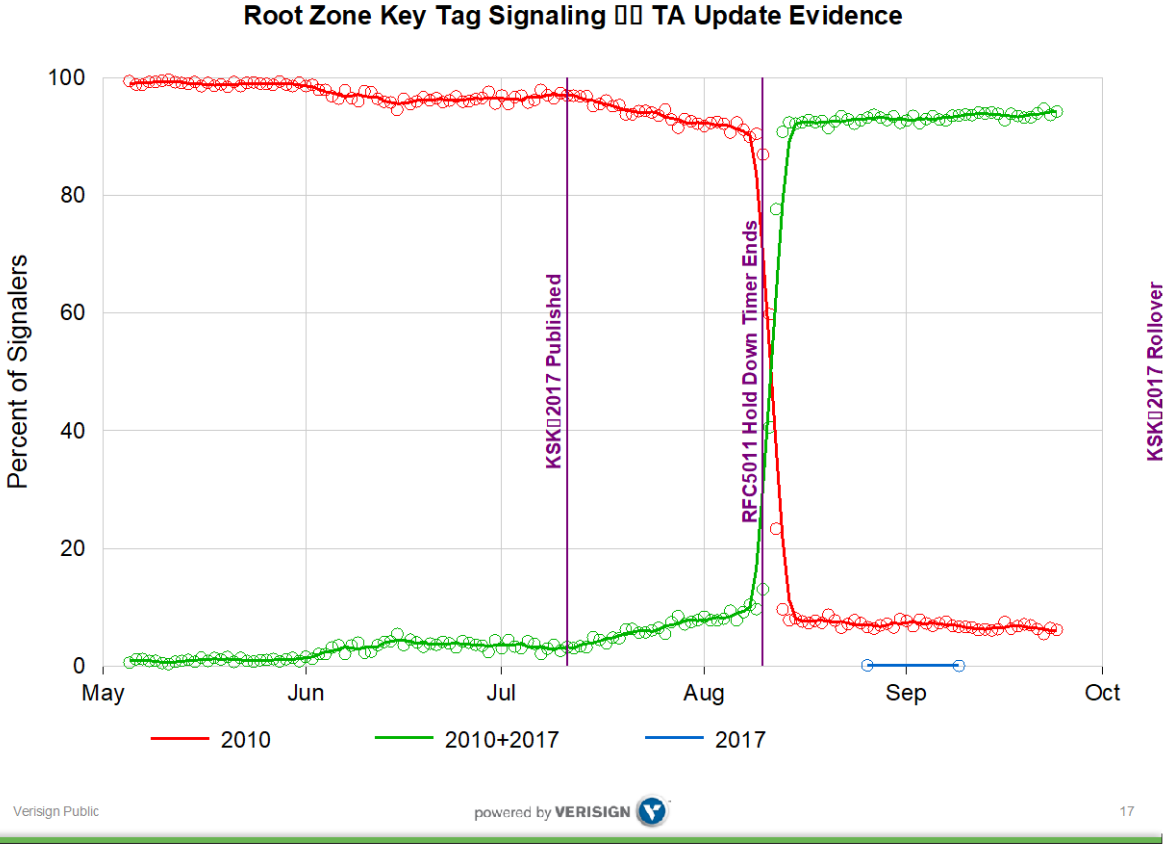
The Process

- ⦿ **11 July 2017:** Introduce the new KSK-2017.
 - Monitor if there are fundamental changes in root-server traffic
 - If not, continue, else fall back.

- ⦿ **10 August 2017:** “30 day hold-down period ends”
 - Monitor if there are fundamental changes in root-server traffic.
 - If not, continue, else fall back.

- ⦿ **19 September 2017:** DNSKEY Response size increased due to standard ZSK roll
 - Monitor if there are fundamental changes in root-server traffic.
 - If not, continue, else fall back.

The timeline in a graph




Who has KSK-2017 configured as a trust anchor?

- ⊙ Until very recently, there was no way to know which trust anchors validators have configured
- ⊙ *Signaling Trust Anchor Knowledge in DNS Security Extensions (DNSSEC)* is a recent protocol extension that can provide that information
 - Reports trust anchor key tags via EDNS option or DNS query
 - Published as RFC 8145 (April 2017)
- ⊙ Implementations
 - BIND 9.11 starting with 9.11.0b3 (28 July 2016)
 - BIND 9.10 starting with 9.10.5b1 (11 January 2017)
 - Unbound 1.6.4 (27 June 2017)
 - On by default in BIND (since 28 July 2016) and in Unbound since version 1.6.7 (10 October 2017)
 - No other known implementations

Looking for key tag signaling

- ⦿ RFC 8145 is so new and validator support so limited that the root KSK roll project team did not expect to get enough data to help with the first root KSK Roll.
 - On average, there are 4.2 Million unique addresses sending queries to root-servers.
 - Given typical deployment curves, it was assumed the dataset would be too small to statistically represent all validating resolvers.
- ⦿ However...
 - Before the introduction of KSK-2017, RFC8145-able resolvers would send KSK-2010 only.
 - After the hold down period of 30 days, RFC8145-able resolvers would send both KSK-2010 and KSK-2017.
 - Duane Wessels (Verisign, co-author of 8145) started looking at A & J root traffic for this signaling

Hey! There's data! Wait. What?




Further analysis by OCTO Research

- ⊙ ICANN OCTO Research did an analysis similar to Duane's
 - Analyzed query data from B, D, F and L root servers
 - For the entire month of September and October (until the 24th)
- ⊙ Results:
 - Total number of unique addresses reporting key tag data: **27,084** (out of 4.2 million, 0.57%)
 - Total number that only ever reports KSK-2010: **1631**
 - **6.02% of reporting validators were not ready for the KSK roll on 11 October 2017**
 - Non-zero percentage of reporting validators were announcing **only** KSK-2017 (!?)
- ⊙ Analysis is complicated
 - Dynamic resolver IPs make the situation look worse by inflating true number of sources
 - Resolvers behind forwarders make the situation look better as they obscure multiple validators behind the forwarder


Why do validators report just KSK-2010?

- ⦿ Multiple reasons suspected or confirmed:
 1. BIND reports trust anchors even if not validating
 2. Old configurations pre-dating automatic update support
 - E.g., BIND's *trusted-keys* instead of *managed-keys* or *dnssec-validation auto*
 3. Bugs in automatic update or key tag signaling support
 - E.g., announce key tags even if DNSSEC not enabled (DO=0)
 4. Operator error
 - E.g., Docker container keeps booting up with only KSK-2010 and starts 5011 all over again
- ⦿ We always knew old configurations would be an issue but never had objective data until now
- ⦿ We worried bugs and operator error were possible but didn't have evidence until now
- ⦿ Analysis is ongoing
 - Hired a contractor to try to figure out reasons for misconfiguration

Rough progress indicator



Of "Resolved" (0.8%) in previous pie



1b: The resolver had KSK-2010 configured manually without RFC 5011 updating and thought that the system was configured for RFC5011, so will fix

2a: The resolver is configured for RFC 5011 updating and the BIND journal file could not be written, so will fix

3: The resolver is forwarding for other resolvers behind it (no sub-code yet, waiting for help from operators)

Back to the plan and process

- ⦿ 19 September 2017: DNSKEY Response size increased due to standard ZSK roll
 - Monitor if there are fundamental changes in root-server traffic.
 - If not, continue, else fall back.
- ⦿ We had received Verisign's report and corroborated it with our own data.
- ⦿ From the Operational Plan:
 - “The Root Zone Management Partners might also decide to extend any phase for additional quarters. For example, if new information indicates that the next phase may lead to complications, the current phase would be prolonged. This is referred to as an extend scenario.”*
- ⦿ 27 September 2017: “Extend” scenario kicks in
 - ICANN Announces that the root KSK Rollover is delayed

- ⦿ We do not know how representative the set of validators reporting key tag data is compared to the set of all validators
- ⦿ Validators != end users (or “end systems”), and the impact on end users is what is most important
 - The design team recognized this
- ⦿ Determining number of end users/systems for a given resolver is hard
 - APNIC’s Google Ad experiment platform-based data will help
- ⦿ Mitigation is hard
 - We’ve already had a multi-year campaign to reach operators
 - Implementation-specific problems don’t make the problem easier

Next Steps

- ⦿ We postponed the root KSK roll until we can gather more information and understand the situation better
 - The delay will be at least one quarter
 - We have not yet determined how many quarters to delay
- ⦿ We will at least partially mitigate
 - Contractor hired to try to track down the 500 resolvers based on IP addresses and understand why misconfiguration is occurring
 - Data collection continues
- ⦿ We'll need to re-engage/re-tune the communications plan
 - Maybe “PLEASE DO **NOT** REMOVE KSK-2017!!”?

Engage with ICANN – Thank You and Questions



One World, One Internet

Visit us at icann.org



[@icann](https://twitter.com/icann)



facebook.com/icannorg



youtube.com/icannnews



flickr.com/icann



linkedin/company/icann



slideshare/icannpresentations



soundcloud/icann