Introduction

The NCSG represents the interests of non-commercial domain name registrants and end-users in the formulation of Domain Name System policy within the Generic Names Supporting Organisation. We are proud to have individual and organizational members in over 160 countries, and as a network of academics, Internet end-users, and civil society actors, we represent a broad cross-section of the global Internet community. Since our predecessor's inception in 1999 we have facilitated global academic and civil society engagement in support of ICANN's mission, stimulating an informed citizenry and building their understanding of relevant DNS policy issues.

This comment was drafted by Eric Osterweil and James Gannon with input from Tomslin Samme-Nlar, Farrel Folly and Adisa Bolutife and approved by the NCSG Policy Committee in accordance with its charter.

Comments

As the root of trust for the global DNS, the Root zone's KSK represents both one of the Internet's most important well-known points of trust and one of its most high-value targets. Since its deployment in 2010, the maintenance and operational security practices surrounding the KSK and its use have been laudible (adherence to FIPS 140-2 level 4, public key signing ceremonies, etc.). The DNS Root KSK is the root of trust for the entire DNSSEC hierarchy, and any and all derived protections (throughout DNSSEC and into dependent protocols) are based on the security and stability of verification from this root key. However, in the recent ``Proposal for Future Root Zone KSK Rollovers,'' lessons learned from the most recent rollover seem to be poorly reflected, and the proposal seems certain to adversely affect relying party software's ability to support DNSSEC verification. The comments herein are broken into these two categories:

1. Formal process verification
2. Rollover periodicity

Process verification

The current proposal outlines the desire to create a ``predictable approach'' for rolling over the Root KSK. This is (indeed) a necessary first-step, though it is not sufficient. In addition to ensuring that the approach is predictable, care must be given to ensure that the process is **complete** and all necessary preconditions are specified, as well. Further, to ensure that DNSSEC is trustworthy, the process should be specified in such a way that it can be proven to be safe, secure, and that actions (including the handling of all foreseeable errors) must be specified and fully enumerated.

This necessity of doing this was illustrated by the rollover pause in the 2017 KSK rollover (which was completed in 2018). The possibility of the circumstances that led to that should have been foreseen and fully planned for, and the eventuality that came to pass should have been addressed by a previously planned contingency plan. A full decomposition of the rollover process (the plan, the foreseeable errors, etc.) should have evoked considerations for the error whereby validators did not appear to be learning the new KSK, and the response should have followed a precise contingency plan (such a plan remains *undescribed* today).

ICANN should not jeopardize the security and resilience of the global DNSSEC deployment by not formally detailing and verifying the Root KSK rollover process.

Rollover Periodicity

The proposal for periodic rollovers is presented without sufficient justification, and seems likely to pose a fundamental risk to the operational stability of DNSSEC's deployment. For relying party software (DNSSEC validators), it is a first-order concern to ensure that the correct Root KSK is securely learned and configured locally.

While current deployments suggest that the Root KSK is configured in millions of resolvers (globally),[1] proposals for a hyper-local root[2] and deployment of local-resolvers propose to accelerate this significantly. With such a large deployment base, and no mechanism to ensure consistency of resolvers' configured Root KSK, the proposal to accelerate the rollover process (necessarily) directly inflates the possibility that the incorrect Root KSK will be configured. in some number of resolvers.

This concern is elevated with the propensity for CPE/SOHO devices to be manufactured and sold in retail after being packaged for prolonged periods, sometimes for multiple years. What is, perhaps, more disconcerting is the decreased ability for innovation that would otherwise use DNSSEC as a verification substrate.

In particular, the DNS-based Authentication of Named Entities (DANE) suite of protocols[3],[4],[5] potentially involve relying party software (such as Mail User Agent plugins, local resolvers, etc.). Such software could easily have lifetimes and life cycles that are not well known in, or well coordinated with, the DNS community. Using RFC-5011 rollovers[6] for the Root KSK (under ideal circumstances) could only aid DANE RPs if they are deployed and able to observe every rollover (without missing any).

While such RPs could still be conformant to DNS protocols, they could fall out of synchronization with the rapid Root KSK rollovers. Increasing the churn of the Root KSK will, correspondingly, increase the uncertainty around whether these tools are able to acquire and maintain the correct Root KSK. Insufficient rigor is evident with respect to the requirements analysis needed for RP software (validating recursive resolvers) to be able to operate under such churn.

Further, proper security analyses are not published that motivate such an approach and such a schedule.

[1] https://stats.labs.apnic.net/dnssec/XA?c=XA&x=1&g=1&r=1&w=7&g=0

[2] https://tools.ietf.org/html/rfc7706

[3] https://tools.ietf.org/html/rfc6698

[4] https://tools.ietf.org/html/rfc8162

[5] https://tools.ietf.org/html/rfc7929

[6] https://tools.ietf.org/html/rfc5011

Rationale for having a standby key

We support the use of a stand-by key for the same reasons provided by the proposal. A shortened Emergency KSK rollovers as will be possible if a standby key were available, will help improve the stability and security of the Internet and its users.
Some have argued that disseminating the public key for that long subjects the key to a successful compromise but like the proposal rightly notes, 2048-bit keys are being used in web PKI without any compromise and even if quantum computing is what we were worried about, as indicated in this article (https://www.technologyreview.com/s/613596/how-a-quantum-computer-could-break-2048-bit-rsa-encryption-in-8-hours/), it'll take another 24 years before quantum computers are ready to break a 2048-bit RSA encryption.

Additional Considerations

We think the proposal should mention how long it took for the previous KSK management software to be tested thoroughly for it to be ready for use by the rollover. This will give the community an indication of whether risks related to KSK rollover delay due to change of KSK management software are possible.

Expected Changes to Processes

A more detailed process *with* error cases fully described and contingency plans illustrated (in detail) must be created and vetted before rolling the Root KSK again.

The periodic rolling of the Root KSK must not be a foregone conclusion, and proper justification (which details RP software, DNSSEC, and DANE usage) must proceed any intention to do so. Such justification must include formal security analyses of the potential harms and benefits of such periodic rolls.