2018 KSK Rollover Back Out Plan

Revised from 2017 KSK Rollover Back Out Plan

ICANN Office of the Chief Technical Officer
27 April 2018

Introduction

This document describes anticipated deviations from the Operational Implementation Plan based on anomalies occurring while executing that plan. It describes the process to be followed, including data collection, applicable criteria, and the steps involved in changing the contents of the DNS. It updates the 2017 KSK Roll Back Out Plan.

At the beginning of 2017, there were three type of events that might require a back out: updating the trust anchor XML file, changing the KSKs in the root DNSKEY RRset, and an increase in the size of DNS responses to priming queries. All three of those events already happened without noticeable issues.

During the KSK rollover process, the ZSK is rolled regularly. The KSK rollover process is designed to work independent of the regular ZSK roll. A back out scenario will not hamper a regular ZSK roll. When there is a back out, it will appear as if the latest step of the KSK roll did not happen while the ZSK has actually rolled forward. For example, a back out from D9 to C9 changes the KSKs from what they were in D9 to what they were in C9, but the ZSKs after back out remain the same.

Summary of Events

The following critical events have been identified and will be referred to in this document:

Event	Date	Slot	Response Size	Description
1	2017-03-08	B8		Add KSK-2017 to TA XML
2	2017-07-11	D2	1139 bytes	Publish KSK-2017 in the DNSKEY RRset
3	2017-09-19	D9	1414 bytes	Response size increase
4	**2018-10-11**	E2	1139 bytes	KSK-2017 signed the DNSKEY RRset
5	**2019-01-11**	F2		Set KSK-2010 as not valid in TA XML
6	**2019-01-11**	F2	1424 bytes	Publish KSK-2010 as revoked
7	**2019-03-22**	F9	1139 bytes	Remove KSK-2010 from the DNSKEY RRset

Events 1, 2, and 3 already happened. Events 4, 5, 6, and 7 (marked in bold type in the "Date" column) are part of the plan for resuming the root KSK roll, with each date moved one year into the future. The specific details for these events are covered in the updated back out plan.

Response Size Changes

Response size is considered to be a sensitive parameter in designing the KSK roll because above some threshold responses in UDP may be fragmented and/or repeated over TCP. The threshold is subject to many factors in the query/response path rendering a fixed limit impossible to establish. Despite many reasons why large response sizes ought to have negative impacts, in practice there have been many cases of TLDs operating with a large DNSKEY RR set without reports of problems. However, to be operationally conservative, there is a general desire to minimize the response size and concerns over growing the size beyond what has been proven to work.

The next increase in response size for priming queries occurs at the start of slot F2, where it increases from 1414 to 1425 bytes. This is where KSK-2010 is published as revoked and signs the DNSKEY RRset. If a back out is ordered during or after F2, the result is that the rollover is brought to the next stage (G). The F2 back out is a keyset that contains the current ZSK and KSK-2017. The revoked KSK-2010 and its signature are absent from the DNSKEY RRset (and signatures), and the response size becomes 864 bytes. The impact of this back out is that KSK-2010, while not present in any DNSKEY set, is not actually revoked from the set of trust anchors for any resolver

relying on RFC 5011 that did not see the revoked key; however it is revoked for those resolvers that saw the revoked key before the back out. Chronologically, this is event 6.

Back out Criteria: Possibly related to fragmentation issues, significant retries for DNSKEY records (exempting retries related to fallback from UDP to TCP). Operators (relying on automated updates) reporting that their validators are not seeing the new trust anchors.

DNSKEY RRset Changes

The next change in the KSK part is in slot E2 when the RRSIG generated by KSK-2010 is replaced by the RRSIG by KSK-2017 (that is, the DNSKEY RRset is signed with KSK-2017 instead of KSK-2010). The back out for slot E2 is to go back to phase D: use the RRSIG for KSK-2010 to sign the DNSKEY RRset. The impact of using this back out is that phase E needs to be postponed until the cause of the damage is found and can be mitigated. Chronologically, this is event 4.

The third change in the KSK part is in slot F2, when KSK-2010 is revoked. The back out for slot F2 is to go forward to phase G: remove the revoked KSK-2010 and its signature. The impact of using this back out is that KSK-2010 is not revoked in validators that have not observed the revoked KSK-2010; for these validators, the KSK roll is not complete because they still have KSK-2010 as a trust anchor. Chronologically, this is event 6.

Lastly, there is a fourth change in the KSK part in slot F9, where the revoked KSK-2010 is permanently removed from the keyset. The back out scenario here is to go to slot G9, which happens to be exactly the same as slot F9: a keyset with KSK-2017, signed by KSK-2017. Chronologically, this is event 7.

Back out Criteria: Higher frequency retries for DNSKEY records by a significant rise in queries from a large distribution of ASNs. Operators (relying on automated updates) reporting that their validators are not seeing the new trust anchors. Reports of distress.

Changes Runbook

1. Status of automated monitoring system verified and phone bridge opened.
2. Root Zone Maintainer (RZM) inserts a root zone with a DNSKEY RR set "of concern" (tied to the events), or the IANA web site is changed.
3. Automated monitoring continues and watched carefully.
4. If there is any sign of a problem, whether reported by an operator looking at a monitor or reported otherwise, the bridge is alerted.
    a. Any reports are evaluated to see if a report meets the criteria for executing a backout.

    b. The backout plan leading to the quickest stable state is enacted.
    c. Adjustments needed to achieve a longer term stable state are determined.
    d. Post-mortem examination of the change begins.
5. After 72 or more (perhaps related to TTL of the affected records) the phone bridge closes.
6. If there is a reason to reopen the phone bridge, it is opened again and step (4) above continues.
7. If there was a problem, before long-term decisions are made a root cause analysis meeting and report are completed and reviewed.

Coordination of Events

At each event, the Root Zone Management Partners will deploy a telephone bridge. During this call the changes are deployed. Additionally, the behavior of resolvers is tracked by continuously analyzing traffic to a subset of root servers whose traffic is available to the IANA Functions Operator and the Root Zone Maintainer.

Backout Recovery

This document does not define the process that deals with recovery after a backout. Invoking a backout is a significant event that is used as a last resort.