Root Zone KSK Rollover Plan

Design Team Report - March 7, 2016

This document is the product of a Design Team coordinated by ICANN that includes volunteers not affiliated with any of the Root Zone Management (RZM) Partners (U.S. DOC NTIA, Verisign or ICANN) as well as representatives from the RZM Partners themselves. A complete list of the members of the Design Team can be found in Section 12.

Table of Contents

1 Executive Summary 4
1.1 DNSSEC-Related Terminology 5
1.2 Other Security Terms 6
1.3 Other Networking Terms 6
1.4 Summary of Recommendations 7
1.5 Audience 9
1.6 Document Scope 9
2 Abridged History 10
2.1 Deployment of DNSSEC in the Root Zone 10
2.2 Root Zone KSK Rollover Public Comment 11
2.3 Root Zone KSK Rollover Preliminary Discussion in 2013 11
2.4 SSAC Advisory on DNSSEC Key Rollover in the Root Zone 12
2.5 ICANN Convenes Root Zone KSK Rollover Design Team 12
2.6 Root Zone KSK Rollover Public Comment 12
3 High-level Description of Rolling a KSK 16
4 Design Team Approach 18
4.1 Operational Considerations 18
4.2 Protocol Considerations 19
4.2.1 Root Zone Trust Anchor Configuration 19
4.3 Impact on Root Zone KSK Management 24
4.4 Cryptographic Considerations 25
4.4.1 Finite Field Cryptography 25
4.4.2 Elliptic Curve Cryptography 26
4.4.3 Conclusion 26
4.5 Coordination and Communication 27
4.5.1 Coordination with the Technical Community and Channel Partners 27
4.5.2 Coordination with Root Server Operators 28
4.5.3 Coordination between KSK Operator and ZSK Operator 29
5 Impact on Validating Resolvers 31
5.1 Packet Size Considerations 31
5.1.1 Measurement Experiment 32
5.1.2 Test Results 33
5.1.3 Conclusion 34
5.2 DNSSEC Validation Behavior 35
5.2.1 Test Results 35
5.2.2 Conclusion 36
6 Testing 38
6.1 Testing for Impact 38
6.2 Self-Test Facilities 38

6.3 KSK and ZSK Maintainer Software and Process Modification Interoperability Testing....39
7 Implementation....40
7.1 Publication of the Incoming KSK....41
7.2 Rollover to the Incoming KSK....42
7.3 Revocation of the Incumbent KSK....42
7.4 Response Packet Size Impact....42
7.5 Deploying Root Server by Root Server....44
8 Rollback....46
8.1 Thresholds....47
9 Schedule for the Root Zone KSK Rollover....49
10 Risk Analysis....51
11 Design Team Roster....53
11.1 Community Volunteers....53
11.2 Root Zone Management Partners....53
12 References....54
13 Channel Partners....56
13.1 Software Producers....56
13.1.1 Pending....56
13.2 System Integrators....57
13.2.1 Linux....57
13.2.2 BSD....57
13.2.3 Others....57
13.3 Public Resolver Operators....58

1 Executive Summary

The Internet Corporation for Assigned Names and Numbers (ICANN) is planning to perform a Root Zone Domain Name System Security Extensions (DNSSEC[1]) key signing key (KSK) rollover. ICANN, in its role as the Internet Assigned Numbers Authority (IANA) Functions Operator, is working in cooperation with the other Root Zone Management (RZM) Partners. The RZM Partners are Verisign, as the Root Zone Maintainer, and the U.S. Department of Commerce's National Telecommunications and Information Administration (NTIA), as the Root Zone Administrator.[2]

Consistent with common practice for the deployment and operation of DNSSEC, the Root Zone KSK is used to sign the root zone apex Domain Name System Key (DNSKEY) resource record set (RRset). That set includes one or more Zone Signing Keys (ZSKs), which are used to sign all other RRsets in the root zone. Rolling the Root Zone KSK refers to changing the key that has been in use since 2010 (when the root zone was first signed with DNSSEC). Changing the key means generating a new cryptographic key pair and distributing the new public component. Adequate distribution of the new public component is the most critical aspect of the key rollover.

In December 2014, ICANN solicited volunteers from the community to participate with the RZM Partners in a Design Team to develop the Root Zone KSK Rollover Plan. The deliverables for this work were a comprehensive set of technical and operational recommendations intended to guide the RZM Partners in producing a detailed implementation plan for executing the first Root Zone KSK rollover. This document should be reviewed as a recommended plan for providing those deliverables.

This document contains terminology related to DNSSEC, Internet security and networking. The following sections give definitions of relevant terminology.

[1] See RFC 4033, RFC 4034 and RFC 4035.

[2] This plan has been developed in accordance with and in recognition of the current RZM structure as currently dictated by the IANA functions contract and the Cooperative Agreement between NTIA and Verisign. The Design Team and RZM Partners recognize that the IANA Stewardship Transition efforts underway may have implications for the KSK Rollover Plan and the involvement of NTIA in any future process. However, the technical details and considerations are largely independent of the transition effort and its end result.

1.1 DNSSEC-Related Terminology

Term	Shorthand	Explanation
Automated updates of DNSSEC trust anchors	RFC 5011	A method for automatically updating the trust anchors used by a validator, published by the Internet Engineering Task Force as an Internet Standard (STD 74).
Delegation signer resource record	DS	A DNSSEC-related RRset that indicates the KSKs currently used by a delegation (or for the root zone, the KSKs of a top-level domain (TLD)).
DNSKEY RRset		The set of keys used in a zone, including the roles of KSK and ZSK, represented as a set of DNSKEY resource records published in the zone.
DNSSEC Policy and Practice Statement	DPS	A document describing the procedural framework for how DNSSEC is implemented and operated for a particular set of zones.
(DNSSEC) validator		Software that performs security checks on DNSSEC responses, including verifying the signatures on data.
Double-signing		The inclusion of two signatures for a single RRset each generated using a different key, usually the old and new key involved in a key rollover. Ordinarily one signature is sufficient for an RRset.
Extension mechanisms for DNS	EDNS or EDNS(0)	Currently defined in RFC 6891 (Internet Standard STD 75), this specification provides a means to extend or expand the original DNS protocol format. EDNS(0) refers to the first set of extensions.
Key rollover		The act of changing from one cryptographic key to another in an orderly way.
Key signing key	KSK	A public-private key pair whose role is to produce a verifiable signature of the set of keys in use in a DNS zone (the DNSKEY RRset). This key is special in the context of the root zone of the DNS, as the key cannot be signed by the parent zone's ZSK, as is the case with other DNS KSKs.
Proof of non-existence	NSEC or NSEC3	DNSSEC-defined resource records used to indicate securely that no data exists for the question asked.
Resource record set	RRset	A unit of data stored in the DNS, the smallest unit that is signed by a DNSSEC key.
Root Server System Advisory Committee	RSSAC	Chartered in the ICANN bylaws, this advisory committee gives advice about the root server System to the ICANN community.

Term	Shorthand	Explanation
Root zone key ceremonies		In-person events held regularly where the private component of a Root Zone KSK key pair is generated, used or destroyed. A formal process is used when witnesses are desired to observe the practices.
Trust anchors		The public components of one or more KSK key pairs that are trusted by a validator.
Zone signing key	ZSK	A public-private key pair whose role is to produce signatures for all RRsets of data in a DNS zone other than the DNSKEY RRset. This key is verified by having the zone's KSK sign the ZSK.

1.2 Other Security Terms

Term	Shorthand	Explanation
Cryptographic Message Syntax Standard	PKCS#7	RFC 2315: PKCS #7: Cryptographic Message Syntax, Version 1.5
The Directory: Public Key and Attribute Certificate Frameworks	X.509	ITU-T standard for management of public-private keys (recommendation ITU-T X.509 | ISO/IEC 9594-8).
Key Signing Request	KSR	A data structure containing requests for signatures over keys, specifically DNSKEY RRsets to be signed by the KSK with specific signature validity periods.
OpenPGP	OpenPGP	An encryption and decryption standard that provides cryptographic privacy and authentication for data communication (RFC 4880: OpenPGP Message Format).
Signed Key Response	SKR	A data structure containing DNSSEC signatures satisfying a corresponding KSR.

1.3 Other Networking Terms

Term	Shorthand	Explanation
Maximum transmission unit	MTU	The maximum number of bytes that can be in data sent over a portion of the Internet. Path MTU refers to the lowest MTU of all portions used in an end-to-end trip across the Internet.
Transmission Control Protocol	TCP	Connection-oriented, byte-order guaranteed-transport protocol for sending data across the Internet.
User Datagram Protocol	UDP	A context-free, best-effort transport protocol for sending data across the Internet.

1.4 Summary of Recommendations

Recommendation 1: The Root Zone KSK rollover should follow the procedures described in RFC 5011 to update the trust anchors during KSK rollover.

Recommendation 2: ICANN should identify key DNS software vendors and work closely with them to formalize processes to ensure that trust anchor distribution using vendor-specific channels is robust and secure.

Recommendation 3: ICANN should identify key DNS systems integrators and work closely with them to formalize processes to ensure that trust anchor distribution using integrator-specific channels is robust and secure.

Recommendation 4: ICANN should take an active role in promoting proper root zone trust anchor authentication, including highlighting the information posted on ICANN's IANA website.

Recommendation 5: Root Zone KSK rollover should require no substantive changes to existing KSK management and usage processes to retain the high standards of transparency associated with them.

Recommendation 6: All changes to the root zone DNSKEY RRsets must be aligned with the 10-day slots described in the KSK Operator's DPS.

Recommendation 7: The existing algorithm and key size for the incoming KSK for the first Root Zone KSK rollover should be maintained.

Recommendation 8: The choice of algorithm and key size should be reviewed in the future for subsequent Root Zone KSK rollovers.

Recommendation 9: ICANN, in cooperation with the RZM Partners, should design and execute a communications plan to raise awareness of the Root Zone KSK rollover, including outreach to the global technical community through appropriate technical meetings and to "Channel Partners" such as those identified in this document.

Recommendation 10: ICANN should request that RSSAC coordinate a review of the detailed timetable for the KSK rollover period before it is published, and should accommodate reasonable requests to modify that timetable in the event that any root server operator identifies operational reasons to do so.

Recommendation 11: ICANN should coordinate with RSSAC and the RZM Partners to ensure that real-time communications channels are used to ensure good operational

awareness of the root server system for each change in the root zone that involves the addition or removal of a KSK.

Recommendation 12: ICANN should coordinate with RSSAC to request that the root server operators carry out data collection that will inform subsequent analysis and help characterize the operational impact of the KSK rollover, and that the plans and products of that data collection be made available for third-party analysis.

Recommendation 13: The RZM Partners should ensure that any future increase in ZSK size is carefully coordinated with KSK rollovers, such that the two exercises are not carried out concurrently.

Recommendation 14: To support a number of potential operational contingencies that may require rollback of changes to the root zone during each phase of the KSK key roll, SKRs using the incumbent KSK, SKRs using both the incumbent and the incoming KSK, and SKRs using the incoming KSK should be generated. The Design Team also recommends that the double-signing approach is the preferred mechanism to respond to a requirement to perform a rollback in Quarter 2 of the key roll procedure.

Recommendation 15: The RZM Partners should undertake or commission a measurement program that is capable of measuring the impact of changes to resolvers' DNSSEC validation behavior, and also capable of estimating the population of endpoints that are negatively impacted by changes to resolvers' validation behavior.

Recommendation 16: Rollback of any step in the key roll process should be initiated if the measurement program indicated that a minimum of 0.5% of the estimated Internet end-user population has been negatively impacted by the change 72 hours after each change has been deployed into the root zone.

Recommendation 17: It is recommended that the KSK rollover process should begin on 1 April 2016, beginning with a nine-month period to generate the new KSK and use the existing scheduled KSK access ceremonies in the period from March to December 2016 to generate the new KSK, copy it to the secondary facility, and prepare the key material to be used in the key roll. The actions associated with changes to the root zone, using the steps and associated timetable as described in "Schedule for the Root Zone KSK Rollover" of this report will begin on 1 January 2017. The publication of the new KSK should be incorporated into the root zone on 11 January 2017, and the old KSK withdrawn and the new KSK to be used in its place on 1 April 2017. If the outcome of the process to evaluate acceptance of the new KSK meets the acceptance criteria described in "Rollback" of this report, then the

old KSK should be revoked starting on 11 July 2017 and the revocation should be removed from the root zone 70 days thereafter, on 19 September 2017.

1.5 Audience

This document is intended for a technical audience, and in particular an audience familiar with the DNS and DNSSEC protocols, operational aspects of the DNS, and the processes associated with the use of DNSSEC in the root zone.

1.6 Document Scope

This document aims to frame and provide a set of recommendations to guide the RZM Partners in their development of a detailed implementation plan for rolling the Root Zone KSK.

2 Abridged History

2.1 Deployment of DNSSEC in the Root Zone

In 2009, the RZM Partners collaborated[3] to deploy DNSSEC in the root zone, which culminated in the first publication of a validatable, signed root zone in July 2010. The Root Zone KSK currently in use was generated in the first KSK ceremony held in a Key Management Facility (KMF) managed by ICANN in Culpeper, Virginia, USA. The key materials were subsequently transported to a second ICANN KMF in El Segundo, California, USA and, once it was verified that they had been securely transported, the public component of the KSK key pair was published in the root zone apex DNSKEY RRset and, separately, as trust anchors for retrieval using non-DNS protocols.[4]

The requirements for generating and maintaining the Root Zone KSK, as well as the respective responsibilities of each of the RZM Partners, were specified by NTIA.[5] The procedures by which those requirements were met by the Root Zone Maintainer and the IANA Functions Operator were published in separate DNSSEC Policy and Practice Statements (DPS).[6]

The IANA Functions Contract between NTIA and ICANN was modified in July 2010 to include responsibilities associated with Root Zone KSK management, and those requirements have been carried forward in subsequent revisions of that contract.[7] The Cooperative Agreement between NTIA and Verisign was also amended in July 2010 to reflect Verisign's Root Zone ZSK Operator responsibilities.[8]

The IANA Functions Contract requires ICANN to perform a Root Zone KSK rollover, but does not provide requirements or a detailed timeline or implementation plan. The Root Zone KSK Operator DPS contains this statement, laying a requirement for a rollover in Section 6.5:

"Each RZ KSK will be scheduled to be rolled over through a key ceremony as required, or after 5 years of operation."

---

[3] Details of DNSSEC deployment in the root zone are published at http://www.root-dnssec.org/.
[4] http://www.root-dnssec.org/wp-content/uploads/2010/07/draft-icann-dnssec-trust-anchor-01.txt
[5] "Testing and Implementation Requirements for the Initial Deployment of DNSSEC in the Authoritative Root Zone," 29 October 2009, http://www.ntia.doc.gov/files/ntia/publications/dnssec_requirements_102909.pdf
[6] https://www.iana.org/dnssec, https://www.verisigninc.com/en_US/repository/index.xhtml
[7] http://www.ntia.doc.gov/page/iana-functions-purchase-order
[8] http://www.ntia.doc.gov/files/ntia/publications/amendment31_07062010.pdf

2.2 Root Zone KSK Rollover Public Comment

On 8 March 2013, ICANN opened a Public Comment period seeking feedback with respect to the execution of a Root Zone KSK rollover[9]. Six organizations and 15 individuals responded. In its summary of the responses,[10] ICANN identified seven recommendations for the RZM Partners to consider:

- A set of tests and measurements, with a test-bed, should be established before embarking on a RFC 5011 KSK rollover. Lines of communication need to be established during testing phases and methods for success evaluation constructed.
- The KSK rollover should be performed as soon as practical with an emphasis on preparedness.
- Measurements and monitoring are the key modes highlighted to gauge the technical and end-user impact of a KSK rollover should one be implemented.
- KSK rollover should take place regularly.
- Public notifications to multiple, diverse stakeholder groups should be made in advance of a KSK rollover event, providing significant advance notice.
- Further investigation is needed on operational stability, repeated KSK rollovers and the likelihood of and impact of non-compliance with RFC 5011.

2.3 Root Zone KSK Rollover Preliminary Discussion in 2013

The RZM Partners convened a meeting in late July 2013 to discuss options for rolling the Root Zone KSK. The team identified the need for a key rollover procedure to be carried out in distinct steps over a conservative time period, the benefits of extensive community outreach, and the notion of a modified RFC 5011 rollover schedule with delayed revocation. These high-level principles were presented at the IETF DNS Operations (DNSOP) working group meeting at IETF 87.[11]

[9] https://www.icann.org/public-comments/root-zone-consultation-2013-03-08-en
[10] https://www.icann.org/en/system/files/files/report-comments-root-zone-consultation-08apr14-en.pdf
[11] http://www.ietf.org/proceedings/87/slides/slides-87-dnsop-6.pdf

2.4 SSAC Advisory on DNSSEC Key Rollover in the Root Zone

In November 2013, the ICANN Stability and Security Advisory Committee (SSAC) published SAC063,[12] concerning the KSK rollover. The report covered the risks involved and the state of the code base at that time (in particular, open source DNS implementations). The report recommended:

- Communication action to publicize the Root Zone KSK key rollover
- Testing to collect and analyze resolver behaviors
- Creation of metrics for what would be acceptable levels of "breakage" in a Root Zone KSK key rollover
- Definition of rollback measures in the event of excess "breakage"
- Collection of information to inform future key roll exercises of this nature.

The SSAC report highlighted three themes that will be covered later in this document. First, a rough estimate of 1.1% of those relying on DNSSEC-enabled DNS could be negatively impacted by even a well-managed Root Zone KSK rollover. Second, the state of support for automated DNSSEC trust anchor updates, described in RFC 5011, was present but unpredictable. And thirdly, the size of DNS responses was a concern when it related to the occurrence of underlying UDP packet fragmentation and reversion to TCP queries.

2.5 ICANN Convenes Root Zone KSK Rollover Design Team

In December 2014, ICANN solicited volunteers from the community to participate with the RZM Partners in a Design Team to develop the Root Zone KSK Rollover Plan, which is presented in this document.

2.6 Root Zone KSK Rollover Public Comment

On 6 August 2015, ICANN opened a Public Comment period seeking feedback with respect to the Design Team's draft report on recommendations concerning the proposed Root Zone

[12] https://www.icann.org/en/system/files/files/sac-063-en.pdf

KSK rollover.[13] Five organizations and nine individuals responded. The comments highlighted a number of aspects of the draft report:

- The report noted the need for a communications plan, but the report did not provide adequate details of what this communications plan would be. Respondents stressed the need for a broad-reaching communications program to be undertaken in conjunction with the KSK roll process.
- With respect to the timing and justification of the KSK roll, some responses advocated the timely roll of the KSK, while others advocated delay.
- The absence of any active testing and signaling of a resolver's capability to follow RFC 5011 was a source of stated concern to respondents.
- The inability to measure the extent to which resolvers were able to pick up an announced new KSK using RFC 5011 procedures was also a source of concern.
- The draft report did not clearly explain its stated preference to roll the key using the existing RSA 2048-bit key, as compared to using a different cryptographic protocol.
- Respondents voiced concern that the draft report did not clearly describe fallback options and the criteria for acceptance of the new KSK.

The devising of a communications plan is beyond the scope of the matters to be specifically considered by the Design Team. The report recommends (Recommendations 2, 3, 4, 9, 11) that the Root Zone Partners develop and implement such a plan, and that an appropriate level of resources be expended to undertake the plan. Such a plan should be executed in conjunction with a detailed timetable of a proposed KSK roll itself, so that the communications can focus on particular critical changeover dates.

Regarding the timing of the KSK roll, comments have been received that advocate the timely roll of the KSK, and other comments that recommend delay. It is noted that Section 3.2.2 of RFC 6781 sets forth the arguments for retaining a trust anchor KSK and only rolling it in the event of a suspected compromise, and also argues that a trust KSK that is rolled regularly creates its own operational habit and operational robustness. In assessing these arguments, RFC 6781 argues for a position of regular rollover of trust anchor KSKs. The Design Team is unaware of what specific objectives would be achieved by delaying a KSK roll. The Design Team is also in broad agreement with the arguments presented in RFC

[13] https://www.icann.org/public-comments/root-ksk-2015-08-06-en

6781 that a regular process of rolling the KSK in a way that minimizes known risks results in a more robust operational environment where both planned and the potential for unplanned KSK rolls are an intrinsic part of the operational environment for root zone management.

The Design Team has been unable to come up with a scenario that allows production resolvers to test their capability to follow a RFC 5011-style key rollover. There are test harnesses in operation, but they all involve the resolver being tested to use a different root hints configuration and a different initial trusted key value. As a result of this bench testing, there is a strong degree of confidence that if a resolver is configured to use RFC 5011-managed trusted keys, then it will correctly follow the process and trust new keys as long as they are introduced in the manner described by RFC 5011. The Design Team is advocating the adoption of an approach to roll the KSK that uses a procedure that is informed by RFC 5011 and RFC 6781, and also informed by the initial efforts to design this process in 2012 and 2013.

Measurement actions can be divided into pre- and post-roll activities. Measuring before the actual key roll has proved to be challenging. The potential to signal whether a validating resolver that relies on a configured trust anchor for the root zone follows the implicit key roll signals defined in RFC 5011 has been the subject of further investigation the Design Team. The conclusion is that it is not possible to devise such a signal or test in the current environment. In other words, when a new KSK is published in the root zone, it is not possible to use a third-party measurement technique to determine which resolvers have picked up the new KSK, nor is it possible at this juncture to determine which resolvers have not picked up the new KSK. Two IETF Internet-draft documents[14] [15] propose to add explicit trust anchor signaling into the DNS specification. Either approach, if adopted, would add some further visibility to the situation, but would also complicate the analysis. The measurements to be performed immediately following the key roll are not specified in the report, other than noting that such measurements should be performed (Recommendation 12) and used to inform the process.

[14] https://tools.ietf.org/html/draft-wkumari-dnsop-trust-management-01
[15] https://tools.ietf.org/html/draft-wessels-edns-key-tag-00

A change of the cryptographic protocol used by the KSK would add one more element of variability to a process that already has a number of unknowns. It is the case that ECSDA allows the response to DNSKEY queries to the root zone to have a smaller response, but this must be offset against the compounding of risk factors by introducing a new protocol that has already been seen to have some level of acceptance issues by the existing set of DNS resolvers.

Measurement and potential fallback procedures were not directly addressed in the draft report. "Rollback" of this report addresses this topic.

3 High-level Description of Rolling a KSK

The plan to roll the KSK is not far removed from plans for rolling any other KSK (as described in RFC 6781), and it follows these steps:

1. An incoming KSK key pair (public and private) is generated. It is noted that this may involve a number of iterations of the process used to access the stored key material, namely one to generate and store the new key at the primary storage facility, and a second to record a copy of the key material at the second storage facility. These access ceremonies are scheduled every three months, so this key generation process will take no less than three months, and prudently estimated to take six months to complete.

2. The public component of the incoming KSK is placed in the DNSKEY RRset of the root zone (as described in RFC 6781) and made available to relying parties. The Root Zone DNSKEY RRset is signed by the incumbent KSK.

3. In a deviation from other zones, the public component of the incoming Root Zone KSK sits in a state where it all concerned accept that it is indeed the next KSK by virtue of being published in the root zone and signed by the incumbent KSK, in accordance with the key introduction procedures specified in RFC 5011. In addition to being accepted via RFC 5011 processes, the new Root Zone KSK public key is to be made available on various electronic and non-electronic media well in advance to allow developers and operators opting out of RFC 5011 to include the new trust anchor in their configurations.

4. The signing process switches from using the incumbent KSK to sign the DNSKEY RRset to using the incoming KSK. At the same time as the signing key is changed to the incoming KSK, the incumbent KSK is removed from the published DNS Root Zone (without revocation).

5. After a period to assess the impact of introducing the incoming KSK, and after confirming that acceptance criteria have been met, the public component of the incumbent Root Zone KSK is reintroduced into the Root Zone DNSKEY RRset for the purpose of marking it as revoked as per RFC 5011 procedures.

As noted in Section 9, these steps are intended to be timed to avoid the periods where the ZSK is rolled, in order to avoid periods where the response to DNSKEY requests for the root zone becomes significantly larger than in ordinary operating conditions.

4 Design Team Approach

The Design Team has considered several aspects of a Root Zone KSK rollover, and produced recommendations from each area of study to guide the Root Zones Partners in implementing this rollover plan.

- Operational considerations—the impact on users of the Internet and the operators of the DNS systems, and services used by those users
- Protocol considerations—the extent to which existing, documented protocol elements are sufficient to accommodate a Root Zone KSK rollover
- Impact on Root Zone KSK Management—the impact on the processes involved in KSK Management by the IANA Functions Operator
- Cryptographic considerations—ensuring that the system as a whole has sufficient cryptographic strength
- Communication and coordination with all involved parties

Each of these areas is individually explored in the sections that follow. A detailed technical rollover solution is also provided as an illustration of how the recommendations might be followed, and intended as a starting point for the RZM Partners as they finalize their implementation plan.

4.1 Operational Considerations

The Root Zone KSK rollover is anticipated to affect Internet users and DNS operators. When the public component of the incoming KSK is added to the Root Zone DNSKEY RRset, the size of the response to queries for the Root Zone DNSKEY RRset will grow. When the incumbent KSK private key is no longer used to sign the Root Zone DNSKEY RRset, validation using the corresponding public key will no longer be possible.

With an increased response to DNSKEY queries, it is possible that fragmentation of UDP packets may occur with slightly different results over IPv4 and over IPv6. It is known that deployed middleware exists on the Internet that considers IP datagram fragments to be anomalous and filters them. For DNS, which maintains no state regarding sent responses, this means a client might not get a validly formed, reassembled response. There is also the potential for a larger UDP response to exceed the query's specified DNS payload buffer

size, therefore increasing the prevalence of truncated responses and subsequent re-query using TCP. The main point of concern here is not necessarily the fallback to using a truncated UDP response and the re-query using TCP per se, but that resolvers may be located behind network filters that block connection attempts to TCP port 53.

Once the incumbent KSK no longer signs the DNSKEY RRset, with the implication that the incoming KSK is generating this signature, a DNSSEC validating resolver with only the incumbent KSK configured as a trust anchor will fail to validate signed DNSSEC responses. The validating resolver will "fail shut," meaning that it will regard all signed DNS responses as invalid, and return a "SERVFAIL" response code to its own clients when it attempts to validate a response.

A DNS client that exclusively uses validating resolvers that fail to pick up the incoming KSK, or fail to receive the larger responses during the key roll process, will be unable to validate any signed DNS responses. This will appear to the end client as a form of Internet outage where domain names are unresolvable. When similar situations have happened before, the side effect is increased calls to customer support centers, which imposes an additional load on ISPs' customer support and operational management roles.

ICANN should plan communications to be coordinated with the introduction of the incoming KSK, as well as the switch from the incumbent to incoming KSK for signature generation (Recommendation 8).

4.2 Protocol Considerations

4.2.1 Root Zone Trust Anchor Configuration

There are two kinds of trust anchor configurations to take into consideration:

- Trust anchors in online validating resolvers
- Trust anchors in devices/systems that are offline during the rollover and brought online later

Online validating resolvers might automatically update DNSSEC trust anchors, described in RFC 5011, if the DNS software used supports this mechanism and is configured to use it to update the Root Zone KSK.

Online validating resolvers that are unable or unwilling to use automated updates of DNSSEC trust anchors will need to be updated manually during the KSK rollover. The manual update should follow the timing of RFC 5011 mechanism—the new trust anchor

must be added to the configuration of such validating resolver in the PUBLISH period of the Rollover (for details, see "Implementation"), and the incumbent trust anchor must not be removed before the incumbent Root Zone KSK is revoked. The mechanisms for retrieving the new trust anchor are the same as for the offline devices.

Recommendation 1: The Root Zone KSK rollover should follow the procedures described in RFC 5011 to update the trust anchors during KSK rollover.

Devices that are offline during the Root Zone KSK rollover will have to be updated manually if they are brought online after the rollover is finished. Such devices, in essence, have to be bootstrapped as if they were newly installed.

Most generally, the process by which any device prepares to be able to perform DNSSEC validation should follow an approach that reduces the opportunity for an inappropriate trust anchor to be used. General advice for such devices is currently being circulated within the IETF in an Internet Draft entitled "Establishing an Appropriate Root Zone DNSSEC Trust Anchor at Startup," [16] but more review is needed to arrive at a stable consensus document that provides advice to implementers.

The Design Team supports and recommends community discussion and review of this Internet Draft within the IETF, with the goal of publishing a stable, peer-reviewed specification in the RFC series.

There are several use-cases of retrieving up-to-date trust anchors, which are explored briefly below.

4.2.1.1 Current and Future Availability of RFC 5011

In the preceding text there is mention of resolvers that are "unable or unwilling" to rely upon RFC 5011's approach. This section is meant to provide some background on that phrase.

The spirit of RFC 5011's add-hold timer is important. The timer is included to prevent a falsely presented key from gaining acceptance. In other words, if an entity wants to present a false KSK, it might succeed in publishing the key. In that event, the true authority will be able to disclaim the false key before any reliance is built on it.

Resistance to RFC 5011 in resolvers is not based on questions related to the design of the update mechanism. Rather, resistance is rooted in a few operational realities. Configuration management is a major concern when operating a fleet of servers and relies on the

[16] https://tools.ietf.org/id/draft-jabley-dnsop-validator-bootstrap-00.html

"pushing outwards" of managed configuration files. RFC 5011's update mechanism runs counter to that, with the configured fleet machines learning new trust anchor data from responses to queries, thereby diverging from the original centrally managed configuration.

With that in mind, large operators may have a manual process in place, a process that will make use of various automated mechanisms. An example of such an automated system might be a tool that follows RFC 5011's update mechanism, but pushes configuration changes to its fleet of servers directly, so that the servers follow the tool's directives rather than the key roll signals that are implicit in the published root zone. In a brief, informal survey, large operators will count on vetting the new Root Zone KSK a few different ways, including human-to-human communication to establish trust. This is the reason alternatives to RFC 5011 are proposed.

Digging deeper into the operationalization of RFC 5011, a few gaps have been identified. The first gap involves remote verification of a successful RFC 5011 process. The second gap involves the ability to test deployments in light of the add-hold timer.

What is needed is a means for the trust anchors in use at a resolver to be made known to the source of the trust. Given the backdrop of pervasive monitoring, the intent is not to have knowledge of specific resolver configuration and capabilities, but first to confirm that the RFC 5011 process was sufficiently followed and to have an idea of when it is acceptable to commit to the incoming Root Zone KSK.

Also identified is a need to speed up the ability to perform a functional test, one that shows the RFC 5011 steps happening, although not adhering to the needed security model. Specifically, tools need to be able to override the specified add-hold timer to allow for a shorter setting during testing. Providing a "test-safe" mechanism to ensure that the test add-hold timer is not used in production is desirable. This is a suggestion for tool developers and DNS software vendors.

4.2.1.2 Other Trust Anchor Formats

Ever since the initial signing of the Root Zone, ICANN has made the trust anchor in non-DNS formats available online[17]. These alternative formats provide an out-of-band method to distribute and obtain the root zone trust anchor, i.e., a method that does not rely upon DNSSEC validation.

[17] https://data.iana.org/root-anchors/

4.2.1.3 Standby Keys

Given that scheduled KSK rollovers are likely to require at least some relying parties to follow manual processes to update trust anchors, there might be an advantage in generating standby KSKs intended for future use so that the corresponding trust anchors can be distributed alongside the active trust anchor, well in advance of the time when they are expected to be used.

Such standby keys would also be extremely useful in the event that an emergency KSK roll was necessary, e.g., due to compromise of the incumbent KSK. The short timeframe in which an emergency KSK roll might need to be executed might eliminate the possibility of using automatic mechanisms like those described in RFC 5011, further increasing the amount of manual work required to distribute new trust anchors and the corresponding risk of undesirable validation failures.

To be meaningful in an emergency key roll, standby keys, once generated, must be stored and managed in a way that is significantly different from the storage and management of the incumbent KSK. The risk factors of a compromise of the incumbent KSK are the determining factors of what "significantly different" means. Regardless of the differences in storage and management, the security of standby keys must be (at least) equivalent to that of the operating incumbent KSK.

The Design Team did not carry out a full assessment of the potential for standby KSKs to be useful, including a comprehensive risk analysis of what differences in key management would be required for a standby KSK set to be useful in the event of compromise of the incumbent KSK. The Design Team suggests that such an analysis be carried out, however, and that the use of a standby KSK set be incorporated into future planning for KSK rollover if that analysis confirms that a standby KSK set would be useful.

4.2.1.4 DNS Software Vendors

Trust anchors may be packaged with DNS software by its vendor (either open source or proprietary/commercial). The software vendor will have to issue a new version of the trust anchor set to keep the software current.

It is important that trust anchors distributed in this way are authentic, and take advantage of the verification mechanisms that already exist to ensure the integrity of software on an end-system. Software vendors require a robust and efficient method to ensure that the trust anchors they distribute with their software are authentic, since the impact of distributing non-

authentic keys is potentially significant, especially if they are signed with code-signing keys as part of a vendor's software update strategy.

Recommendation 2: ICANN should identify key DNS software vendors and work closely with them to formalize processes to ensure that trust anchor distribution using vendor-specific channels is robust and secure.

4.2.1.5 System Integrators

One distribution method of DNSSEC trust anchors is via system integrators, for example, a package maintainer or an operating system vendor. In this case, the system integrators will provide updated packages for all copies of trust anchors in the system. There are efforts in several Linux distributions to provide a package with one authoritative copy of the trust anchor.

Recommendation 3: ICANN should identify key DNS system integrators and work closely with them to formalize processes to ensure that trust anchor distribution using integrator-specific channels is robust and secure.

4.2.1.6 System Administrators

System administrators can manually download DNSSEC trust anchors from ICANN's IANA website while installing or updating software. Current root zone trust anchors are provided by the IANA Functions Operator on a dedicated website[18] for information pertaining to DNSSEC in the root zone. Determining the authenticity of downloaded trust anchors is critical to establishing trust in DNSSEC. To support verifying authenticity of various types of digital signatures, in the form of OpenPGP, PKCS#7 and a X.509 certificate containing the root key, are also published at the same dedicated website.

Although determining authenticity is extremely important, it is often overlooked and further underspecified. When processes for supporting the authenticity proofs were made available for public review, there was a low volume of substantive comment, which undermines the effort to adequately support authenticity. It seems possible that additional review (with backward-compatible changes, where appropriate) is merited. As mentioned before, the Design Team supports community discussion and review of the Internet Draft entitled "DNSSEC Trust Anchor Publication for the Root Zone" (cited earlier) within the IETF, with the goal of publishing a stable, peer-reviewed specification in the RFC series.

[18] https://www.iana.org/dnssec/files

Recommendation 4: ICANN should take an active role in promoting proper root zone trust anchor authentication, including highlighting the information posted on ICANN's IANA website.

4.3 Impact on Root Zone KSK Management

As described in the "DNSSEC Practice Statement for the Root Zone KSK Operator," the Root Zone KSK Operator signs each of the root zone's apex DNSKEY RRsets by way of a KSR supplied by the Root Zone ZSK Operator. The result is a SKR containing a set of signed DNSKEY RRsets provided to the Root Zone Maintainer that encompasses the forthcoming key rollovers of the ZSK.

These processes are well documented and, in the case of actions that take place during KSK ceremonies, subject to external audit and widespread observation. The Design Team considers it highly advantageous to avoid any substantive changes to processes as a result of the rolling of the KSK in order to avoid disruption to a process that is, in its current form, already well understood.

Recommendation 5: Root Zone KSK rollover should require no substantive changes to existing KSK management and usage processes to retain the high standards of transparency associated with them.

Each KSR covers a time cycle of one calendar quarter (three months or roughly 90 days) and is divided into 9 slots of 10 days each. If the time cycle is more than 90 days, the last slot in the cycle is expanded to fill the period. Therefore, all changes to the root zone DNSKEY RRset, e.g., adding and/or removing keys as required by a key rollover, should be aligned with these 10-day periods to minimize any substantive changes in the processes used to publish a signed root zone.

Recommendation 6: All changes to the root zone DNSKEY RRsets must be aligned with the 10-day slots described in the KSK Operator's DPS.

With the standard periods, the root DNSKEY RRset packet response size increases with the first and last slot in each time cycle. The first slot contains the post-published ZSK from the previous time cycle, whereas the last slot contains the pre-published ZSK for the next time cycle.

To minimize potential issues related to larger DNS responses sizes, it is desirable to schedule a rollover that can keep the DNSKEY RRset response size as small as possible. For a detailed examination of response size issues, with accompanying recommendations,

see "Implementation". For a Root Zone KSK rollover schedule designed with the aforementioned considerations in mind, see " Schedule for the Root Zone KSK Rollover".

4.4 Cryptographic Considerations

The Design Team considered the question of whether there were sufficiently compelling grounds to consider a change in key size or algorithm for the KSK. A compelling reason might stem from questions regarding the cryptographic strength of the chosen key size or algorithm.

With the initial publication of SP 800-57, “Recommendation for Key Management, Part 1,” in 2005, the U.S. National Institute of Standards and Technology (NIST) announced the intent to raise minimum cryptographic strengths. However, in the five years between the publication and the proposed end date, factoring techniques have not progressed as quickly as anticipated. There is nothing to suggest that there is an urgency to use longer key lengths for the Root Zone KSK.

4.4.1 Finite Field Cryptography

The 2048-bit asymmetric RSA key is considered to be equivalent to a 103-bit symmetric key, according to the “ECRYPT II 2012 Yearly Report on Algorithms and Keysizes.”[19] The same report recommends using at least 96 bits of security for about 10 years of protection. The NIST “Recommendation for Key Management, Part 1: General (Revision 4)”[20] considers the 2048-bit RSA key to be equivalent of 112 bits of security and considers this strength to be acceptable for use in the period from 2014 to 2030. The French Agence nationale de la sécurité des systèmes d'information (ANSSI) document “Référentiel Général de Sécurité”[21] also considers the 204-bit RSA key to be safe to use until 2030.

The signed content in the root zone is typically short lived as the DNSKEY signature periods are measured in days (about 15 days), and the Design Team believes that the 2048-bit RSA key should be safe for a five more years unless there is a significant technological breakthrough in the large integer factorization area.

[19] http://www.ecrypt.eu.org/ecrypt2/documents/D.SPA.20.pdf
[20] http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-57pt1r4.pdf
[21] http://www.ssi.gouv.fr/uploads/2015/01/RGS_v-2-0_B1.pdf

4.4.2 Elliptic Curve Cryptography

Another algorithm option available for DNSSEC is the Elliptic Curve Digital Signature Algorithm (ECDSA) that is defined in RFC 6605.[22] ECDSA has some properties that would make it desirable to use as an algorithm for Root Zone KSK. The keys are much smaller, while keeping equivalent strength to RSA keys. The current estimates are that ECDSA with curve P-256 has an approximate equivalent strength to RSA with 3072-bit keys (NIST) or 3248-bit keys (ECRYPT II). However, the algorithm was standardized for use in DNSSEC only relatively recently—RFC 6605 was published in 2012—and measurements described later in this document have observed that support for ECDSA in validators is not as widespread as the support for RSA (see "Impact on Validating Resolvers").

The IETF Crypto Forum Research Group (CFRG) is also working on a new "Elliptic Curves for Security" RFC that adds new Elliptic Curves security, and it also voices some concerns from the crypto community about the generation and potential weaknesses of the curves used by ECDSA. It is desirable to let the CFRG finish the work on the document before switching to a new Elliptic Curve algorithm for signing the root zone.

4.4.3 Conclusion

Based on the guidance described above, the Design Team found that there is no pressing need to change either the algorithm or the size of the KSK from 2048-bit RSA. The Design Team also learned of a DNS validating resolver implementation that requires all digital signatures in the root zone to be signed by all algorithms matching the configured trust anchors, therefore the rollover to a different algorithm would require coordinated changes in both the KSK and ZSK for the root zone.[23] This implementation provides another practical reason to avoid a change in algorithm at this time. The Design Team has contacted the vendor of this DNS resolver implementation regarding the issue and the vendor's requirement, and there is the expectation that it will be relaxed in the future.

For these reasons, the incoming KSK for the first KSK rollover should be a 2048-bit RSA key, but changes in algorithm and/or key length may be worth considering for subsequent KSK rollovers.

Recommendation 7: The existing algorithm and key size for the incoming KSK for the first Root Zone KSK rollover should be maintained.

[22] https://tools.ietf.org/html/rfc6605
[23] The Design Team is now aware this requirement was addressed in a subsequent software release.

Recommendation 8: The choice of algorithm and key size should be reviewed in the future for subsequent Root Zone KSK rollovers.

4.5 Coordination and Communication

The devising of a communications plan is beyond the scope of the matters to be specifically considered by the Design Team. The report is recommending that such a plan be developed and implemented by the Root Zone Partners and that an appropriate level of resources be expended to undertake the plan. Such a plan should be executed in conjunction with a detailed timetable of a proposed KSK roll itself, so that the communications can focus on particular critical changeover dates.

4.5.1 Coordination with the Technical Community and Channel Partners

ICANN should design and execute a communications plan to raise awareness of the Root Zone KSK roll. Awareness ought to be raised within technical forums such as those at which the original deployment of DNSSEC in the root zone was presented.

The term “Channel Partners” refers to external organizations that facilitate the use of DNSSEC independent of the management of the root zone. These partners “channel” the value of signing the root zone out from the RZM Partners into the global public Internet.

The Channel Partners are segmented into three general areas. First are the enablers, those implementing DNSSEC validation software, concerned with, among other items, implementing RFC 5011. Second are distributors of software and systems that include DNSSEC validation software, primarily concerned with distributing copies of the Root Zone KSK. Third are operators of DNSSEC validating systems that make use of the Root Zone KSK.

To facilitate communication, the Design Team recommends that for each Channel Partner, if willing, a contact should be kept on file, and updates on the KSK roll will be given to these contacts. This contact list is not intended to be exclusive or to exchange material that is not otherwise publicly available. The contact list is intended to allow for a sampling of the awareness of steps in the Root Zone KSK roll. The list should, however, remain closed to allow Channel Partners to manage the awareness of their selected contact information.

Recommendation 9: ICANN, in cooperation with the RZM Partners should design and execute a communications plan to raise awareness of the Root Zone KSK rollover, including outreach to the global technical community through appropriate technical meetings and to “Channel Partners” such as those identified in this document.

4.5.2 Coordination with Root Server Operators

Any structural change in the contents of the root zone has the potential to affect operational behavior of individual root servers. The initial provisioning of IPv6 address (AAAA) glue in the root zone and the subsequent deployment of DNSSEC are examples of changes that were made with consultation and close coordination with the root server operators, since those changes triggered changes in query patterns. Therefore, prudence with critical infrastructure dictates a conservative approach to any change in the event that there are unexpected consequences that might degrade the performance of the root server system as a whole.

The experiments conducted as part of the preparation of this document suggest that a KSK rollover event will cause no harmful effects; however, as with the earlier examples of structural change mentioned above, a conservative approach is recommended.

The Design Team suggests that individual root server operators might treat particular events within the KSK rollover period as they would treat a significant, planned, operational event, issuing public status notices and coordinating with other root server operators using the normal real-time channels used for such events. Such events should include the period surrounding the addition of a new, incoming KSK to the root zone apex DNSKEY RRset, and the removal of the outgoing KSK from the same RRset.

The Design Team suggests that real-time communication channels between individual root server operators and ICANN, and between ICANN and the other RZM Partners be similarly exercised around the same events to ensure that any expected effect can be identified and shared promptly.

A detailed timetable for the KSK rollover period should be reviewed by the root server operators before it is finalized and published, in order to ensure that it does not conflict with any other plans that might reduce the ability of an individual root server operator to provide the desired level of operational coverage. Effort should be made to adjust the timing of the rollover to avoid operational conflicts, as far as is practical.

Recommendation 10: ICANN should request that RSSAC coordinate a review of the detailed timetable for the KSK rollover period before it is published, and should accommodate reasonable requests to modify that timetable in the event that any root server operator identifies operational reasons to do so.

Recommendation 11: ICANN should coordinate with RSSAC and the RZM Partners to ensure that real-time communications channels are used to ensure good operational

awareness of the root server system for each change in the root zone that involves the addition or removal of a KSK.

Data collection by root server operators over the course of the KSK rollover facilitate understanding of the operational impact of a KSK rollover on validators and on the root servers themselves. Since the root server system is diverse both in architecture and distribution around the Internet, it is understood that opportunities for long time-based data collection by individual root server operators will involve various constraints that are difficult to characterize succinctly for the system as a whole. It is also understood that baseline data collection capabilities already exist to satisfy the tactical requirements of monitoring service conditions in real time, as the KSK rollover proceeds.

When DNSSEC was initially deployed in the root zone, a substantial data collection exercise was carried out, and the resulting data proved useful in off-line analysis of the reaction of the DNS as a whole to the structural changes taking place in the root zone, including analysis by third parties, facilitated by the Domain Name System Operations Analysis and Research Center (DNS-OARC).[24] A similar exercise is warranted for the first KSK rollover.

Recommendation 12: ICANN should coordinate with RSSAC to request that the root server operators carry out data collection that will inform subsequent analysis and help characterize the operational impact of the KSK rollover, and that the plans and products of that data collection be made available for third-party analysis.

4.5.3 Coordination between KSK Operator and ZSK Operator

Responsibility for the management of the Root Zone KSK and ZSK are separately assigned to the IANA Functions Operator and the Root Zone Maintainer, respectively. The two roles are managed separately.

The Root Zone ZSK is currently a 1024-bit RSA key, as specified in the ZSK Maintainer's DPS.[25] It is possible that the Root Zone Maintainer will increase the ZSK key size in the future.

The ZSK is regularly rolled on a 90-day schedule, and it is expected that this will continue as usual during the KSK rollover period. Since the KSK rollover period is expected to be longer

[24] https://www.dns-oarc.net/
[25] http://www.verisigninc.com/assets/dps-zsk-operator-1527.pdf

than 90 days, there will be periods during which the root zone apex DNSKEY RRset may contain four keys, depending on the final plan.

Increasing the ZSK size during a key rollover event might trigger different behavior in validators for part of the KSK rollover period, since response sizes will increase with ZSK size. This might complicate efforts to identify, understand and mitigate any operational problems that arise.

Any decision relating to ZSK size is outside the scope of this document. However, the Design Team recommends that ICANN coordinate with the Root Zone Maintainer to ensure that any future increase in ZSK size is carefully coordinated with KSK rollovers, such that the two exercises are not carried out concurrently.

Recommendation 13: The RZM Partners should ensure that any future increase in ZSK size is carefully coordinated with KSK rollovers, such that the two exercises are not carried out concurrently.

5 Impact on Validating Resolvers

5.1 Packet Size Considerations

The DNS is defined to operate over the UDP and TCP transport protocols. UDP was preferred in the design of the DNS protocol due to the lower overhead of UDP when compared to TCP, particularly in terms of maintaining connect states on a server. However, there is a limitation imposed by this protocol choice. In the original definition of DNS, RFC 1035, UDP responses were limited to 512 octets. The 512-octet limit is observed in software still in use today, either honoring or enforcing that limit.

Through the extension mechanism for DNS, EDNS(0), originally defined in an RFC published in August 1999 (RFC 2671, updated by RFC 6891) a DNS requestor is able to inform the DNS server that it can handle UDP response sizes larger than 512 octets. The requestor places its maximum UDP payload size (not the IP packet size, but the DNS message size) in the query, and the server is required to respond with a UDP response, where the DNS payload is no larger than the specified buffer size. If this is not possible, then the server sets the truncate bit in the response to indicate that truncation has occurred. If the truncated response includes a valid DNS message, the requestor may elect to use the truncated response. Otherwise, the requestor opens a TCP session to the server and repeats the query over TCP.

DNS systems that make use of DNSSEC must signal their ability to do so using the DO (DNSSEC OK) flag in the EDNS pseudo-header. Since the operational impact considered in this document is entirely concerned with systems that are DNSSEC-capable, the systems involved are EDNS(0)-capable (because DNSSEC requires EDNS(0) support) and hence not restricted to the 512 octet limit.

A client may initiate a transaction in TCP, but common requestor behavior is to initiate the transaction in UDP and use the truncate bit in the UDP response to indicate that the requestor should use TCP for a re-query.

UDP packet fragmentation is treated differently in IPv4 and IPv6. When a packet is too large for the underlying IP packet transmission medium, the IP packet may be fragmented. In this case, the trailing fragments use the same IP level leader (including the UDP protocol number field), but specifically exclude the UDP pseudo-header in the trailing fragments. In IPv4, the original sender or any intermediate router, may fragment an IP packet, unless the "Don't Fragment" IP flag is set. In IPv6, only the original sender may fragment an IP packet.

If an intermediate router cannot forward a packet onto the next hop interface, then, in IPv6, the router will generate an ICMPv6 diagnostic packet with the MTU size of the next hop interface and the leading part of the packet, and will pass this information back to the packet sender.

When using UDP, a sender does not maintain a buffer of unacknowledged data, so the IPv6 sender, when receiving this message, cannot retransmit the original data. Empirical data appears to suggest that a common response by many IPv6 implementations is to generate a host entry in the local IPv6 forwarding table, and record the received MTU in this table for some locally determined cache time. This implies that any subsequent attempts to send an IPv6 UDP packet to this destination will use this MTU value to determine how to fragment the outgoing packet.

5.1.1 Measurement Experiment

An experiment has been designed and set up to reproduce the environment of the root server situation with a goal of evaluating the impact of large packet sizes on resolvers and users.

This was achieved by using an online advertisement platform to trigger DNS resolvers to pose unique queries to an authoritative name server configured to respond to queries for two zones with different response sizes. It is believed that the resolvers that pose the query to the authoritative name server in this test are largely the same set of resolvers that would be expected to query the root zone.

To test whether a resolver could receive a large response, the advertisement queried for a target domain name. The target domain name itself would return a normal-sized response. But to get to the target response, the resolver had to receive a large intermediate response first. If the resolver succeeded in even asking for the target domain name's information, then the test showed that the resolver could handle the large intermediate response.

The test also involved the retrieval of a web object from the experiment's web server, allowing the experiment to match the addresses used in the web retrieval (the end user's IP address) to the addresses used by the name resolvers in posing the DNS query.

In this test, a 1,444-octet DNS response was used.

5.1.2 Test Results

In a five-day period during May 2015, some 7.26 million end systems successfully fetched a small control record, and of these, some 7.17 million systems successfully fetched the test record, a difference of approximately 90,000 users, or 1% of the sample set, who failed to fetch the 1,444 octet DNS test record.

These end systems used some 83,000 different DNS resolver IP addresses. Of these, 94% of the resolvers successfully obtained both the control record and the test record. Of the 4,251 resolvers who retrieved the control record but failed to retrieve the test record, 3,396 resolvers used the EDNS(0) extension with the DNSSEC OK bit set, which triggered the 1,444 octet response. Of these failing resolvers, 3,110 resolvers were observed only a single time during the experiment, while 826 resolvers exhibited the failure condition more than once. This implies that 1% of resolvers seen in this experiment failed to retrieve a large response two or more times, while a further 3% of resolvers who failed to retrieve the large response were only seen a single time, which is insufficient to conclude with any assurance that they would fail consistently with large responses. This 1% of resolvers that failed consistently two or more times were used by slightly less than 3,000 end systems, or 0.04% of the sampled end system population.

Some 5,237 resolvers used IPv6 addresses in this test (6% of the total), and 830 of those resolvers failed to retrieve the test record (21% of the failing resolvers). These data suggest a potential issue with some IPv6 resolvers and their handling of MTU sizes.

In terms of measuring the change in query load with larger responses, the control name (with a 93-octet response size) was queried 16.4 million times, and 475 queries were observed using TCP. The test name (with a 1,444 octet response size) was queried 18.6 million times, and 1.2 million of these queries were made over TCP, or some 6.5% of the total query count for the test name. There is a difference in the total number of queries made to the control record versus the total number of queries to the test record. The difference can be explained by resolvers responding to receiving truncated responses for the test record by sending another query over TCP. This result correlates reasonably well with the distribution of UDP buffer sizes offered in the EDNS(0) extensions of the UDP queries. When serving larger responses, an authoritative server can anticipate a higher query load, and a higher proportion of queries over TCP.

5.1.3 Conclusion

Approximately 1% of DNS resolvers that set the DNSSEC OK flag in their queries appear to be unable to receive a DNS response of 1,444 octets (experimental uncertainty factors mean that the upper bound on this number is 6% of all resolvers). Within this set of resolvers, resolvers using IPv6 as a transport protocol are disproportionately represented. It is possible that this failure rate is due to the presence of various forms of DNS-intercepting middleware, or in the case of IPv6, due to potential mishandling of ICMP6 “Packet Too Big” messages. However, the precise nature of the failures cannot be established from within this experimental methodology.

Resolvers failing to receive responses serve a very small proportion of users. The number of users who use DNS resolvers that are consistently unable to resolve a DNS name when DNS responses of this size are involved appears to be 0.04% of all users (experimental uncertainty factors mean that the upper bound on this number is 1% of all users).

These experiments tested a DNS response of 1,444 octets. It is noted that other parts of the DNS already provide significantly larger responses than the size being contemplated here, and these response sizes do not appear to have generated public attention or visible comment. For example, a comparable DNSKEY query for the .org name on the 6th June 2015 generated a 1,625-octet response containing two 2048-bit RSA KSKs, two 1024-bit RSA ZSKs and three signatures—one by each KSK and one by one of the ZSKs. Any validating resolvers that are incapable of receiving such large DNS responses would be unable to validate the signature of either the DS record or the NSEC3 record (which are used to signal the non-existence of a DS record) for each delegation in the .org zone, effectively causing DNS resolution failures for delegations in .org.

The Design Team is not aware of any operational problems that domain name holders in .org might be experiencing related to the size of DNSKEY DNS response packet of the .org name. Even after taking into account the very small number of signed zones within .org, this lack of any operational reports about resolution failure in .org domain names would indicate that response size is unlikely to present a significant operational issue for the Root Zone KSK rollover.

One difference to note between the test case and the .org situation is that only resolvers that actually perform validation will query for the large DNSKEY RRset. In the test case, all resolvers signaling DNSSEC OK would try to fetch the large response. As described in "Impact on Validating Resolvers", it appears that less than 30% of resolvers setting DNSSEC OK in the original query subsequently perform validation of the response. It is

possible that those resolver operators that have turned on validation have been more diligent in identifying and correcting any network-related issues which may prevent them from retrieving large response packets, as these resolvers would be more prone to experience such problems. Other resolvers, not doing validation, would only under relatively rare circumstances encounter large response packets, and may not be aware of such limitations imposed upon them by their network environment.

It is reasonable to infer that the vast majority of those failing to receive the large response in the tests are non-validating resolvers, which would not be affected by the increase in size of the DNSKEY resource record of the root zone.

In summary, these tests indicate that less than 0.04% of users may be impacted by a larger response size during a Root Zone KSK rollover, but this is an estimate with a high uncertainty factor, and related observations drawn from TLDs with large key sets would tend to indicate that this is an upper bound on the extent of impact from the larger response size.[26]

5.2 DNSSEC Validation Behavior

There are three aspects of DNSSEC validation behavior to measure. The first is retrieval of the DNSSEC digital signatures (setting the DNSSEC OK flag of the EDNS(0) options in the query), the second is the validation function, where a chain of trust is created from the root key to the name being validated, and the third is whether the user's name resolution configuration will accept a DNSSEC validation failure as a definitive failure or whether the query will be referred to another resolver.

5.2.1 Test Results

Using the experiment described above (Section 6.1.1), in May 2015, some 85% to 90% of users were observed to pass their queries to resolvers where the resultant queries observed at an authoritative name server for an uncached name have the EDNS(0) option included in the query and also have the DNSSEC OK flag set.

Some 24% of the same sampled user population performed subsequent queries that illustrate that the resolver was validating the response using DNSSEC by following the chain of interlocking signatures back up the name delegation hierarchy to the Root Zone KSK.

[26] Further details and results of the experiment are described at http://www.potaroo.net/ispcol/2015-05/ksk.html.

Some 11% of the same sampled user population corresponds to end-user behavior that will respond to a DNSSEC validation failure from the previous pass by passing the query to a different resolver that does not perform DNSSEC validation.

This suggests that any change in DNSSEC validation procedures has the potential to impact approximately one quarter of the Internet's user population.

Of these, a little less than one half of this pool of users already interpret DNSSEC validation failure (signaled by SERVFAIL) as a signal to present the same query to a different resolver that does not perform DNSSEC validation. For this pool of 11% of the Internet's users, the change of the Root Zone KSK may potentially involve an unrecognized Root Zone KSK and validation failure, but these users have demonstrated that they already interpret SERVFAIL by using an alternate resolver. The outcome could potentially involve a longer time to resolve DNSSEC-signed names, but would not result in the inability to resolve the name at all.

The remaining 13% of users who do not revert to a non-validating resolver when receiving a SERVFAIL response are potentially at risk of being unable to resolve a DNSSEC-signed name, if the resolvers used by the user are incapable of following the signals provided through the RFC 5011 key rollover process.

5.2.2 Conclusion

It is not possible to use this measurement process to test whether resolvers are capable of following an RFC 5011 process to automatically pick up a new Root Zone KSK value. The best that can be done here is to quantify the user population who use resolvers that perform DNSSEC validation, and hence use resolvers that will either support RFC 5011 or need manual intervention to load the new Root Zone KSK at the appropriate point in time.

Some 24% of users use resolvers that perform DNSSEC validation, and will therefore be potentially impacted by a Root Zone KSK roll. Failure to validate will return a SERVFAIL response, and 11% of all users use a collection of resolvers where a SERVFAIL response from one resolver will cause the query to be resolved by a non-validating resolver. This implies that 13% of all users may be impacted by a Root Zone KSK roll if their resolver is not RFC 5011 aware and the resolver administrator does not load the new Root Zone KSK at the appropriate time.

However, many of these users are using one of the larger DNSSEC validating resolver services that are understood to be RFC 5011 aware (such as Comcast's DNS resolvers), so this 13% figure is an upper bound on the population of users who may be impacted in this way, and the true figure is considered to be far lower than this number.

6 Testing

There are two elements related to testing. One is the activity of measuring the impact of the KSK roll on the general operations of the Internet for the purposes of assessing the level of negative impact that might halt the operation. The other is the activity related to preparing relying parties for the operation, including test-bed resources for self-testing. Self-testing may be conducted by Channel Partners developing software and/or operators deploying fleets of servers, or anyone else interested.

6.1 Testing for Impact

Tests run for other portions of this report measuring validation success have uncovered some reaction to DNSSEC validation failures. Using evidence that some queries start with DNSSEC and then "failover" to DNS, whether this practice increases (or falls) as the KSK is rolled can be one means to assessing damage. This so-called damage might otherwise go unnoticed, but could be a valuable metric when observing the impact of the Root Zone KSK key roll operation. End users (at a screen) most likely do not detect this, and thus never open a ticket to a service provider help desk.

Tests that detect this ought to be run on a periodic basis (monthly) from now until the end (successful or not) of the Root Zone KSK key roll operation. Pre-roll, the tests will give us a baseline from which to compare.

In addition to automated testing, contact with Channel Partners during the Root Zone KSK key roll will be needed to provide explicit, real-time or near-real time information. They will want to choose times when staffing is adequate and when they can give sufficient notice to affected parties.

6.2 Self-Test Facilities

In enabling relying parties to self-test, there should be a test platform mimicking the operational platform at an accelerated rate of roll. Besides having servers running RFC 5011 at an accelerated rate with signed false root zones, the trust anchors in other data structures should be present at the same path names. This will encourage better tools to be produced, such as tools to assist in vetting a key or discover what is in a validator (for local or remote consumption).

This can help with education about new algorithms by allowing the insertion and removal of keys of different parameters.

Timing is an important issue. Running a test in which time is compressed (i.e., instead of waiting 30 days, wait only 5 minutes) is needed to allow for reasonable observation of the functioning of the process. Testing in normal (i.e., wall-clock) time is also needed, this is a better estimation of processing load as well as being a more faithful estimation of the true operational environment.

And finally, fidelity to the root system has to be addressed. Whether or not the whole root zone is used as data or a representative, false zone is a topic to be considered.

There are existing examples of such test beds[27, 28] that may be used as a model for future testing.

6.3 KSK and ZSK Maintainer Software and Process Modification Interoperability Testing

Since the KSK rollover process requires modifications to existing schedules, processes, and possibly software supporting KSK operations, thorough testing of these changes must be performed prior to commencement of rollover, including but not limited to: key generation, signed DNSKEY RRset generation, DNSSEC validation, KSR/SKR exchange, any fallback mechanisms, and Key Ceremony rehearsals.

[27] http://keyroll.systems/
[28] http://icksk.dnssek.info/fauxroot.html

7 Implementation

The proposed key rollover process was first conceived in July 2013 and has since been refined. The process described here should be considered a draft. Before implementation, RZM Partners might want to make further improvements.

The process is divided into three phases:

1. Publication of the incoming Root Zone KSK
2. Change to signing with incoming Root Zone KSK ("the rollover")
3. Revocation of the incumbent Root Zone KSK

Revocation of the incumbent Root Zone KSK is deliberately delayed to allow for a rollback, should any problems with the incoming Root Zone KSK arise after the incumbent Root Zone KSK has been removed from the key set. The process aims to be compliant with RFC 5011, with extended windows for adding the incoming KSK and revoking the incumbent KSK. This process explicitly allows for the option to defer the revocation of the incumbent Root Zone KSK for an indefinite period, allowing for the case where there are unforeseen issues observed with the rollover process that require a change to the planned key rollover process.

Figure 1 below shows an overview of the three quarters during which the process takes place. Note that the numbering of the quarters is relative to the start of the process, not tied to a calendar. For example, Quarter 1 and Q1 do not necessarily mean the period January to March. The incoming KSK is noted as "KSK-NEW", the incumbent KSK is "KSK-2010."

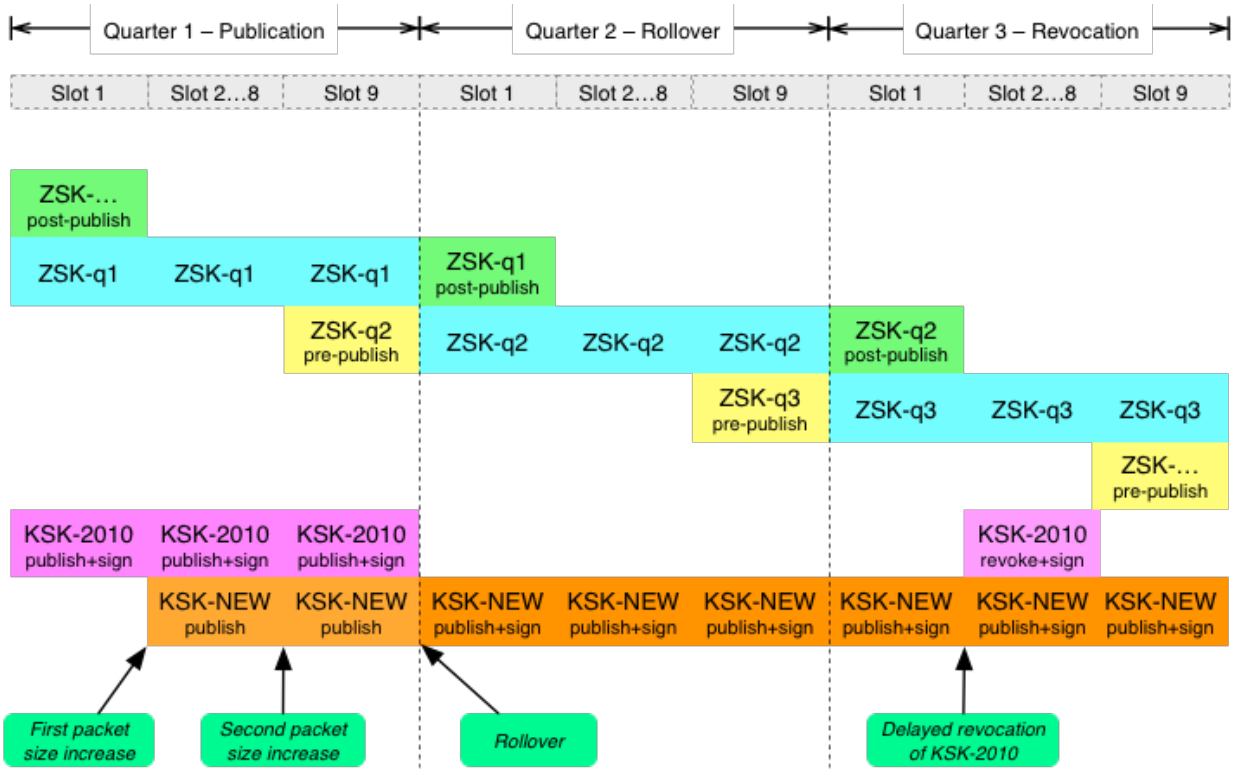

Figure 1. Rollover Scheduling

7.1 Publication of the Incoming KSK

The incoming KSK is added to the DNSKEY RRset at Q1 slot 2, but is not yet used for signing. This is a provisional publication phase in order for the incoming KSK to be picked up by RFC 5011-compliant validators. The incoming KSK is published (and signed by the incumbent KSK) in the root zone for a total of 80 days before used for signing. Manually configured trust anchors are expected to be updated to include the incoming KSK before or during this time period.

An RFC 5011-compliant rollover requires that a new key be published during a period of no less than 30 days ("add hold-down time"). If the proposed 80-day publication period is deemed insufficiently long, it is possible to insert one or more additional publication quarters before rolling the key.

During the publication quarter of the incoming KSK, DNSSEC validating resolvers will see the packet size of a response to a query for the root zone DNSKEY RRset (response packet size) increase from 736 octets to 1,011 octets. This increase compares the sizes of the DNSKEY set between (and not during) the roll of the ZSK. During the last slot of Q1, the

slot when the regularly scheduled ZSK roll begins, the response packet size is increases from 833 octets to 1,158 octets.

7.2 Rollover to the Incoming KSK

After the incoming KSK has been introduced, it is used to sign the root DNSKEY RRset, starting at Q2 slot 1. This quarter is just like any other quarter, except that all DNSKEY RRsets are signed with only the incoming KSK. The only time that the DNSKEY RRset would be signed by both the incumbent and incoming KSKs is during the optional revocation period, described below.

7.3 Revocation of the Incumbent KSK

If the incumbent KSK is to be revoked as described in RFC 5011, the incumbent KSK is published with the revoke bit and signed by both the incumbent and the incoming KSK.

Revocation of the incumbent KSK is optional. If revocation is desired, publication of the revoked incumbent KSK is performed starting at Q3 slot 2 through Q3 slot 8.

During a revocation, the response packet size increases from 736 octets to 1,297 octets.

7.4 Response Packet Size Impact

A desired objective is to avoid UDP fragmentation as far as possible. Table 1 shows some relevant response size constraints.

Table 1. Packet Size Thresholds

Size	Threshold
512 octets	Minimum DNS payload size that must be supported by DNS
1,232 octets	Largest DNS payload size of an unfragmentable IPv6 DNS UDP packet
1,452 octets	Largest DNS payload size of an unfragmented Ethernet IPv6 DNS UDP packet
1,472 octets	Largest DNS payload size of an unfragmented Ethernet IPv4 DNS UDP packet

Results of testing presented earlier indicate potential problems with some IPv6 resolvers and their handling of large responses. The first and most important size constraint is therefore the threshold of an unfragmentable IPv6 DNS UDP packet, which implies a DNSKEY response packet size of at most 1,232 octets.

This first threshold is only reached during the optional revocation phase, where the incumbent Root Zone KSK has to be reintroduced and flagged with the revoke bit. For full compliance with RFC 5011, during the revocation phase, it is a requirement to double-sign

the DNSKEY RRset with both the incoming Root Zone KSK and the incumbent Root Zone KSK. Double-signing the RRset will result in the response size exceeding 1,232 octets.

The largest single response packet for the root zone is the signed DNSKEY RRset. Table 2 contains an overview of the DNSKEY response packet size during the proposed roll, as well as a comparison with the non-roll response packet sizes. (The color coding in Table 2 corresponds to Figure 2.)

Table 2. Packet Sizes During Rollover

Time	DNSKEY During Roll	RRSIG During Roll	DNSKEY Response Size During Roll	DNSKEY Response Size During Non-Roll
Q1 slot 1	1· KSK + 2· ZSK	1· KSK	883 octets	883 octets
Q1 slots 2 to 8	2· KSK + 1· ZSK	1· KSK	1,011 octets	736 octets
Q1 slot 9	2· KSK + 2· ZSK	1· KSK	1,158 octets	883 octets
Q2 slot 1	1· KSK + 2· ZSK	1· KSK	883 octets	883 octets
Q2 slots 2 to 8	1· KSK + 1· ZSK	1· KSK	736 octets	736 octets
Q2 slot 9	1· KSK + 2· ZSK	1· KSK	883 octets	883 octets
Q3 slot 1	1· KSK + 2· ZSK	1· KSK	883 octets	883 octets
Q3 slots to 8	2· KSK + 1· ZSK	2· KSK	1,297 octets	736 octets
Q3 slot 9	1· KSK + 2· ZSK	1· KSK	883 octets	883 octets

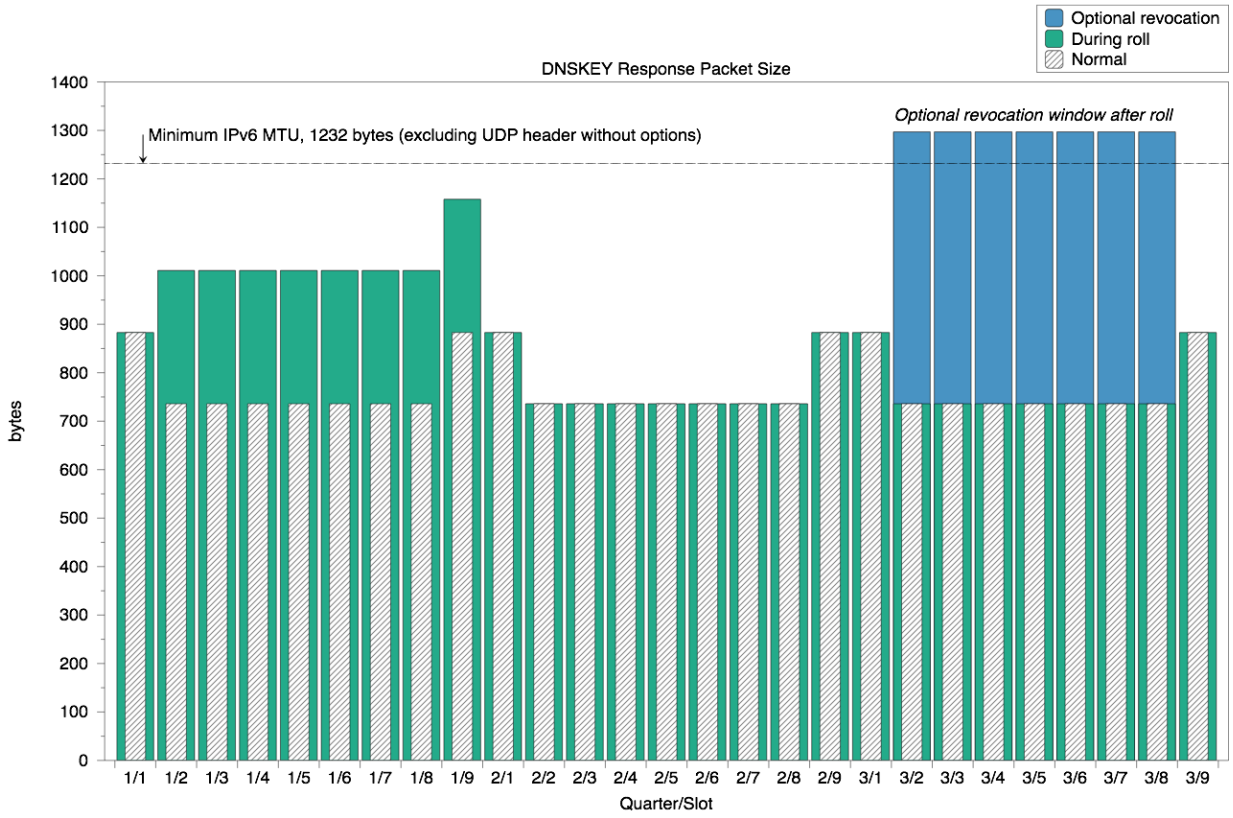

Figure 2. DNSKEY Response Packet Sizes

Risks associated with avoiding revoking the outgoing key have not been thoroughly discussed, but the revocation phase can be viewed as optional at this stage. One option could be to update the RFC 5011 in this respect, and to not require double-signing for revoking an outgoing key. This revision would have the added benefits that a lost or destroyed key can be revoked. Not having to double-sign with the outgoing key could also facilitate future key rollovers, algorithm changes and changes in key lengths. However, due to the time to redefine, publish, develop and distribute code, as well as move the code into production, this option is not deemed feasible for this KSK key rollover.

7.5 Deploying Root Server by Root Server

The 2010 introduction of DNSSEC happened root server by root server. A preliminary version of the DNSSEC signed zone appeared on one server in January 2010, another root server in February, two more root servers in March, and so on. The goal was to allow recursive servers (or anything sending queries to the root servers) the ability to try DNSSEC first and fallback if the answers weren't acceptable.

This strategy was proposed for the Root Zone KSK roll, but was dismissed for a number of reasons. With the goal of mitigating problems related to the new Root Zone KSK and an ability to measure the adoption of the new trust anchor over time, the following realities stood in the way.

In face of DNSSEC validation failure, the reaction by the validating recursive server varies from tool to tool. Some tools are known to be very aggressive when retrying, some not so, and some don't bother at all.

Detecting whether a recursive server (or any query source) has made an explicit decision to prefer one root server over another is known to be impractical. In ordinary circumstances there is insufficient tracking of query sources at the root servers to detect recursive servers preferring one root server over another root server. The Day in the Life of the Internet (DITL) collection[29] performed annually by DNS-OARC runs for a short period of time, is an enormous undertaking and still has never managed to cover all of the root servers in any time period.

A final consideration is the time span available to incrementally introduce the new trust anchor. There are only 70 days in any quarter outside of a Root Zone ZSK rollover. Adding the incoming KSK (to the first server) requires 40 days, leaving just 30 more days to complete the task within one ZSK rollover period. The original incremental deployment stretched for more than four months.

[29] https://www.dns-oarc.net/ditl/2011

8 Rollback

In case there are serious problems detected immediately after the introduction of the incoming KSK, alternative DNSKEY RRsets that include the incumbent KSK should be available for deployment. These RRsets are in SKR format and can be produced using the same Root Zone KSK key ceremonies as the non-rollback RRsets.

Rollback SKRs containing DNSKEY RRsets need to be prepared for all three quarters of the key roll process.

During Q1, where the incoming KSK is being introduced to the root zone, rollback consists of removal of that key introduction. In this case, the rollback SKR consists of DNSKEY RRsets with the incumbent KSK and the current ZSK(s), signed by the incumbent KSK. The incoming KSK is omitted in this rollback SKR.

The rollback measures for Q2 could take one of two forms.

The first is to restore the root zone DNSKEY set to contain only the previous incumbent KSK and be signed by this KSK. DNS resolvers that are using automated trust anchor management according to RFC 5011 would have kept the previous incumbent KSK as a "missing" trusted key, and its reappearance would allow these resolvers to immediately resume trusting this key for validation. It is presumed that the level of impact that has triggered this rollback is due to clients behind resolvers using a manually managed key regime who have not updated their trusted key set, and still have not added the incoming KSK to their trusted key set. These resolvers would be able to validate responses in this scenario once the previous incumbent KSK was restored to the root zone. However, there is also the potential for resolvers with manually managed keys to have performed a swap of the previous incumbent key to the incoming key at the time of the rollover. While this could be considered an operationally imprudent move, nevertheless, there is a consideration that clients behind resolvers configured in such a way would be negatively impacted by this rollback measure.

The second form of rollback is to add the previous KSK to the DNSKEY RRset and to use this key to also sign the DNSKEY RRset. DNS resolvers who are using automated trust anchor management per RFC 5011 would be able to validate DNS responses using either the previous incumbent key or the incoming key as a trust anchor. Those resolvers using manually managed keys with either the previous incumbent or the incoming KSK configured as trust anchors would be able to validate DNS responses. While this appears to offer no

negative impacts in terms of resolvers with manually or automatically managed trust anchor keys, the downside of this approach is the increased size of the DNSKEY response, which would then contain two DNSKEY entries and signatures, similar to the scenario described in the key revocation phase in "Implementation". While the use of the two KSKs as a rollback response would have no understood drawback in terms of resolvers using manually managed keys, the larger response size of 1,297 octets may have some negative impacts. However, as noted in "Impact on Validating Resolvers", it is estimated that this level of impact is of the order of less than 0.04% of all users.

The Design Team recommends that root zone key material be prepared for both rollback scenarios. The Design Team also recommends that the double-signing approach is the preferred mechanism to respond to a requirement to perform a rollback.

During Q3, where the incumbent (outgoing KSK) is placed back into the root zone with the revocation flag set, then rollback would logically consist of removal of the incumbent (outgoing) KSK from the root zone. The rollback SKR consists of the DNSKEY RRset with the incoming KSK and the current ZSKs, signed by the incoming KSK.

8.1 Thresholds

In setting a threshold of impact of a change in the DNSSEC properties of the root zone, any change that causes an immediate impact on more than a predetermined proportion of the Internet's endpoint population would be a clear signal for implementation of some form of rollback as described above.

What is appropriate in this context is a clear understanding of the DNSSEC validation behavior of resolvers, and the population of the endpoint client set of each of the larger resolvers, before the start of changes to the root zone for the key roll. A measurement approach should be able to detect the set of resolvers whose DNSSEC validation behavior has changed during the various phases of the key roll, and be able to estimate population of clients who have been affected by this change of resolver behavior.

Determining what is a minimum threshold level of impact that can trigger rollback is acknowledged to be a relatively imprecise exercise, but as a starting point, the Design Team offers an initial threshold of an estimated 0.5% of all users being impacted by the change in the root zone at a time of 72 hours after each change being deployed into the root zone for the KSK roll.

Recommendation 14: To support a number of potential operational contingencies that may require rollback of changes to the root zone during each phase of the KSK

key roll, SKRs using the incumbent KSK, SKRs using both the incumbent and the incoming KSK, and SKRs using the incoming KSK should be generated. The Design Team also recommends that the double-signing approach is the preferred mechanism to respond to a requirement to perform a rollback in Quarter 2 of the key roll procedure.

Recommendation 15: The RZM Partners should undertake or commission a measurement program that is capable of measuring the impact of changes to resolvers' DNSSEC validation behavior, and also capable of estimating the population of endpoints that are negatively impacted by changes to resolvers' validation behavior.

Recommendation 16: Rollback of a step in the key roll process should be initiated if the measurement program indicated that a minimum of 0.5% of the estimated Internet end-user population has been negatively impacted by the change 72 hours after each change has been deployed into the root zone.

9 Schedule for the Root Zone KSK Rollover

Given the existing operational environment, there are four days in the calendar year when a new Root Zone KSK can take over from the incumbent KSK. Those four days are the first days of quarters: 1 January, 1 April, 1 July and 1 October. Picking a specific date for the change has two components—what is operationally reasonable and what is compatible with the current discussions regarding the IANA transition.[30]

Operationally reasonable means that the dates involved should avoid weekends, holidays that affect work schedules, and times when operations staff is operating on a thin margin. Given the need to align three dates with a global audience, not every consideration may be readily accommodated in every case.

It is noted that Section 3.2.2 of RFC 6781 sets forth the arguments for retaining a trust anchor KSK and only rolling it in the event of a suspected compromise, and also argues that a trust anchor KSK that is rolled regularly creates its own operational habit and operational robustness. In assessing these arguments, RFC 6781 argues for a position of regular rollover of trust anchor KSKs. The Design Team is unaware of specific objectives to be achieved by delaying a KSK roll, and without a specific objective, little is gained in proposing an indefinite delay in the KSK roll. The Design Team is in broad agreement with the arguments presented in RFC 6781 that a regular process of rolling the KSK, in a manner that minimizes known risks, results in a more robust operational environment where both planned and the potential for unplanned KSK rolls are an intrinsic part of the operational environment for RZM.

The schedule proposed here uses information available to the Design Team at the time of preparation of this report. Root Zone Managers may have to include other considerations, and may need to alter the proposed schedule to ensure that all considerations relating to the stability of this key roll are adequately addressed.

The schedule allows nine months for the preparation of the new KSK, its storage in the KSK storage facilities and the generation of signature material that will be used in the key rollover. This period allows for the rollover to make use of the existing key access

[30] http://www.ntia.doc.gov/press-release/2014/ntia-announces-intent-transition-key-internet-domain-name-functions

ceremonies and should not require the calling of ad hoc assembling of the Trusted Community Representatives outside of the regularly scheduled key access ceremonies. This report does not specify the steps to be taken to prepare the new KSK and the preparation of signature material to be used at these ceremonies—this planning activity is part of the responsibilities of the KSK manager.

Recommendation 17: It is recommended that the KSK rollover process should being on 1 April 2016, beginning with a nine-month period to generate the new KSK and use the existing scheduled KSK access ceremonies in the period from March to December 2016 to generate the new KSK, copy it to the secondary facility, and prepare the key material to be used in the key roll. The actions associated with changes to the root zone, using the steps and associated timetable as described in "Implementation" of this report will begin on 1 January 2017. The publication of the new KSK should be incorporated into the root zone on 11 January 2017, and the old KSK withdrawn and the new KSK to be used in its place on 1 April 2017. If the outcome of the process to evaluate acceptance of the new KSK meets the acceptance criteria described in "Rollback" of this report, then the old KSK should be revoked starting on 11 July 2017 and the revocation should be removed from the root zone 70 days thereafter, on 19 September 2017.

10 Risk Analysis

Risks Associated with Insufficient Preparation

Description	Impact	Likelihood	Mitigation
Roll of KSK with same algorithm, hash and size will not be sufficient in the eyes of stakeholders.	Low	Unlikely	Plan another roll once the first one is complete; if different parameters are needed, change them.
Network operators will not be aware of the change (i.e., NOC gets trouble tickets, needs to know how to react).	Moderate	Likely	In communications plan; operator focus
Network operators and software developers (or "all Channel Partners") will not have (access to) adequate testing environments.	Moderate	Likely	Set up an ICANN RFC 5011 testbed with accelerated and in-time rolls; perform other testing.
Ability to centrally test during progress not feasible	Low	Likely	Develop distributed test approaches; develop contact list.
Lack of deterministic criteria to make go/no-go decision	Low	Likely	Prepare communications and testing, feasibility studies of mechanisms used in the field, long-term efforts to develop measurements updated trust anchor acceptance.

Automated Trust Anchor Mechanism Doesn't Work or Is Inadequate

Description	Impact	Likelihood	Mitigation
RFC 5011 not enabled everywhere.	Moderate	Likely	Alternative trust anchor management approaches
RFC 5011 incompletely implemented.	Moderate	Unlikely	Contact software developers; verify understanding of RFC 5011.
Validator bootstrap process incompletely implemented.	Moderate	Unlikely	Contact system integrators and trust anchor handlers.
Trust anchor sets not available from ICANN's IANA website.	Low	Unlikely	Monitoring of availability
Equipment with out-of-sync trust anchor sets due to lack of maintenance.	Low	Likely	Communications plan

Removal of Incumbent KSK Causes Validation Failures

Description	Impact	Likelihood	Mitigation
Automated trust anchor protocol insufficiently followed (by any participant in the process)	Low	Likely	Testing, communication; provide resources for operators to speed remediation.
Elevated traffic due to retry-in-face-of-failure	Low	Unlikely	Examine "roll-over-and-die" lingering effects, negative caching recommendations

Addition of Incoming KSK Causes DNS Message Size to Exceed Limits

Description	Impact	Likelihood	Mitigation
Transition of keysets causes oversized datagrams.	Moderate	Unlikely	Thorough planning of transition by examining size of messages
Confusion over IPv6 fragmentation handling in DNS software	Low	Unlikely	Examination and testing of DNS software

Operational Errors Occur

Description	Impact	Likelihood	Mitigation
Botched KSK roll will end momentum for DNSSEC adoption.	High	Unlikely	Design and review carefully.
Indefinitely postponing a key rollover increases the impact if it becomes urgent.	High	Unlikely	Commit to a Root Zone KSK roll.
Once begun, can never return to the current acceptable state.	High	Unlikely	Define a fallback plan.
Incumbent KSK (private component) is not sufficiently destroyed.	Low	Unlikely	Commit to completing the plan.

11 Design Team Roster

11.1 Community Volunteers

- Joe Abley, Dyn, Inc., CA
- Jaap Akkerhuis, NLnet Labs, NL
- John Dickinson, Sinodun Internet Technologies, UK
- Geoff Huston, APNIC, AU
- Ondrej Sury, CZ.NIC, CZ
- Paul Wouters, Red Hat / No Hats, CA
- Yoshiro Yoneya, JPRS, JP

11.2 Root Zone Management Partners

- David Conrad, ICANN
- Edward Lewis, ICANN
- Richard Lamb, ICANN
- Alain Durand, ICANN
- Hayley Laframboise, ICANN
- Elise Gerich, ICANN
- Kim Davies, ICANN
- Roy Arends, ICANN
- Jakob Schlyter, ICANN
- Fredrik Ljunggren, ICANN
- Brad Verd, Verisign
- Duane Wessels, Verisign
- David Blacka, Verisign
- Al Bolivar, Verisign
- Tim Polk, U.S. DOC NIST
- Scott Rose, U.S. DOC NIST
- Doug Montgomery, U.S. DOC NIST
- Ashley Heineman, U.S. DOC NTIA
- Vernita Harris, U.S. DOC NTIA

12 References

- "DNSSEC Practice Statement for the Root Zone KSK Operator," (5 November 2012) https://www.iana.org/dnssec/icann-dps.txt
- "DNSSEC Practice Statement for the Root Zone ZSK Operator," (21 October 2010) https://www.verisigninc.com/assets/dps-zsk-operator-1527.pdf
- "DNSSEC Trust Anchor Publication for the Root Zone" https://tools.ietf.org/html/draft-jabley-dnssec-trust-anchor
- "ECRYPT II 2012 Yearly Report on Algorithms and Keysizes" http://www.ecrypt.eu.org/ecrypt2/documents/D.SPA.20.pdf
- "The EDNS Key Tag Option" https://tools.ietf.org/html/draft-wessels-edns-key-tag-00
- "Establishing an Appropriate Root Zone DNSSEC Trust Anchor at Startup" https://tools.ietf.org/html/draft-jabley-dnsop-validator-bootstrap
- "Recommendation for Key Management, Part 1: General (Revision 4)" NIST Special Publication http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-57pt1r4.pdf
- "Référentiel Général de Sécurité" French Agence nationale de la sécurité des systèmes d'information (ANSSI) http://www.ssi.gouv.fr/uploads/2015/01/RGS_v-2-0_B1.pdf
- RFC 5011: "Automated Updates of DNS Security (DNSSEC) Trust Anchors" https://tools.ietf.org/html/rfc5011https://tools.ietf.org/html/rfc5011
- RFC 6605: "Elliptical Curve Digital Signature Algorithm (DSA) for DNSSEC" https://tools.ietf.org/html/rfc6605
- RFC 6781: "DNSSEC Operational Practices" https://tools.ietf.org/html/rfc6781

- SAC063: “SSAC Advisory on DNSSEC Key Rollover in the Root Zone” https://www.icann.org/en/system/files/files/sac-063-en.pdf
- “Signalling of DNS Security (DNSSEC) Trust Anchors” (draft) https://tools.ietf.org/html/draft-wkumari-dnsop-trust-management-01

13 Channel Partners

The term “Channel Partners” refers to external organizations that independently either enable or convey the value of managing the Root Zone KSK. These organizations have no formal relationship with the RZM Partners, yet coordination is essential to some extent. For each organization, appropriate contacts are to be maintained to exchange status and other information related to the change of the Root Zone KSK.

The Channel Partners are listed in no particular order.

13.1 Software Producers

The substantive communication with these partners pertains to the implementation (or not) of RFC 5011 trust anchor management in software. The set of partners are those with validating recursive cache servers. Contact information with these organizations is not listed in this document.

- ISC BIND (http://www.isc.org)
- NLnet Lab Unbound (https://nlnetlabs.nl)
- Microsoft Windows Server (https://www.microsoft.com/)
- Nominum’s Vantio (http://nominum.com/caching-dns/)
- DNSMASQ (http://www.thekelleys.org.uk/dnsmasq/doc.html)
- IRONSIDES (http://ironsides.martincarlisle.com)
- Infoblox (http://www.infoblox.com/)
- Secure64 DNS Cache (http://www.secure64.com/)

13.1.1 Pending

The following set of partners have discussed, but not released, DNSSEC validating recursive cache servers. They are on a list to be included if code is distributed. (Other DNS recursive cache servers without DNSSEC support do not depend on the Root Zone KSK.)

- CZ.NIC TBD recursive server (aside from Knot)
- PowerDNS TBD

13.2 System Integrators

These Channel Partners convey the Root Zone KSK as part of configuration data involving, in some cases, the DNS software previously mentioned. The expectation is that these organizations will review the incoming Root Zone KSK and include it in their software updates.

13.2.1 Linux

- Red Hat Enterprise Linux (RHEL) RPMs
- Micro Focus International's SUSE (RPMs)
- Fedora
- CentOS
- Debian and Canonical (Ubuntu) APT
- MontaVista Linux

13.2.2 BSD

- FreeBSD ports
- NetBSD pkgsrc
- OpenBSD ports

13.2.3 Others

- Apple iOS, OS X
- Google Android, Chrome OS
- Microsoft
- Cisco
- Juniper
- Belkin
- Cisco Linksys
- Wind River Real-Time Operating System (RTOS)
- QNX Neutrino (RTOS)
- OpenVMS
- OpenWrt

13.3 Public Resolver Operators

These partners are reported to run recursive DNS servers, in some cases validating DNSSEC. The expectation is that these would include the Root Zone KSK as configuration data, therefore there may be internal reviews that need to know of the incoming Root Zone KSK.

- Google Public DNS
- OpenDNS
- Neustar DNS Advantage
- Norton ConnectSafe
- Level 3
- CensurfriDNS
- Comodo
- Dyn Internet Guide
- Liquid Telecom

In addition to the preceding list of operators with public resolvers, selected based on accepting traffic from anywhere in the Internet (so far as can be seen), there are partners that operate public resolvers with restrictions on their relying party base. As these partners are identified, they will also be offered notifications of Root Zone KSK events.

One World, One Internet

ICANN.ORG