Root Zone KSK Rollover Plan

Design Team Draft Report - Updated August 4, 2015

1 Overview

ICANN is preparing a plan to perform a Root Zone DNSSEC Key-Signing Key (KSK) rollover. The rollover operation is being planned by ICANN, in its role as the IANA Functions Operator, in cooperation with the other Root Zone Management (RZM) partners. The partners are Verisign, as the Root Zone Maintainer, and the US Department of Commerce's National Telecommunications and Information Administration (NTIA), as the Root Zone Administrator.[1]

Rolling the Root Zone KSK refers to changing the key that has been in use since 2010 when the Root Zone was first signed according to the definition of DNS Security Extensions (DNSSEC)[2]. Changing the key means generating a new cryptographic secret component and distributing a new public component. Adequate distribution of the new public component is the most critical aspect of the key rollover operation.

This document is being made available for Public Comment and is a draft report of the deliberations of a Design Team that consists of a panel of recruited volunteer experts in DNS and DNSSEC, along with the Root Zone Management partners. The state of this document is a draft, to be amended by input from the Internet community during ICANN's open Public Comment and further deliberations. Following the ensuing conversations to be had, a final report will be issued.

2 Table of Contents

[1] This draft plan is developed in accordance with and/or in recognition of the current root zone management 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.

[2] See RFC 4033, RFC 4034 and RFC 4035

3 Executive Summary

ICANN, as the IANA Functions Operator, in cooperation with Verisign as the Root Zone Maintainer and the US Department of Commerce’s National Telecommunications and Information Administration (NTIA) as the Root Zone Administrator, together known as the Root Zone Management (RZM) partners, has sought to develop a plan for rolling the Root Zone Key-Signing Key (KSK).

In accordance with DNSSEC, the Root Zone KSK is used to sign the Root Zone DNSKEY resource record set. That set includes the Zone-Signing Key (ZSK), which is used to sign all other resource record sets (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 according to DNSSEC). Changing the key means generating a new cryptographic secret component and distributing a 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, as presented in this document. 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 draft plan intended to provide those deliverables.

3.1 DNS Terminology

This document relates to technical details of the DNS and DNSSEC. So that definitions of DNSSEC-related terms (jargon) are readily available, definitions of some relevant items are included in Table 1 below.

Term	Shorthand	Explanation

Term	Shorthand	Explanation
Resource Record Set	RRSet	A unit of data stored in the DNS, the smallest unit that is signed by a DNSSEC key
Key Signing Key	KSK	A public-private key pair[3] whose role is to produce a verifiable signature of the set of keys in use in a DNS zone. This role is special because DNSSEC requires this kind of public key to be distributed external to the DNS protocol
Zone Signing Key	ZSK	A public-private key pair whose role is to produce signatures for all other sets of data in a DNS zone. This key is not distributed outside of the DNS protocol
DNSKEY RRset		The set of keys used in a zone, including the roles of KSK and ZSK, a set of DNSKEY resource records
Key Rollover		The act of changing from one cryptographic key to another in an orderly fashion
(DNSSEC) Validator		Software that performs security checks on DNSSEC responses, including verifying the signatures on data as one step
Trust Anchors		A stored public KSK trusted absolutely by a validator
Automated Updates of DNSSEC Trust Anchors	RFC 5011	One method for automatically updating the trust anchors in a validator
Double-signing		The inclusion of two signatures for an RRset, usually the old and new key involved in a rollover. Ordinarily one signature is sufficient for an RRset

[3] Ferguson, Niels; Schneier, Bruce (2003). *Practical Cryptography*. Wiley. ISBN 0-471-22357-3.

Term	Shorthand	Explanation
Root Server System Advisory Committee	RSSAC	Chartered in the ICANN by-laws, provides advice regarding the Root Server System to the ICANN community
Extension Mechanisms for DNS	EDNS or EDNS(0)	Currently defined in RFC 6891, provides a means to extend or expand the original DNS protocol format. EDNS(0) refers to the first set of extensions
Delegation Signer Resource Record	DS	DNSSEC record indicating the KSK in use by a sub-delegation (or for the Root Zone, the KSK of a top-level domain)
Negative Answer	NSEC or NSEC3	DNSSEC-defined resource records used to indicate data does not exist for the question asked
DNSSEC Practices Statement	DPS	A document describing specifics of DNSSEC processing for a zone.
Key Ceremonies		Events in which the private key is used, inside an HSM, to generate signatures. A formal process is used when witnesses are desired to observe the practices.

Table 1. DNS and DNSSEC Terminology

3.2 Other Security Terms

Term	Shorthand	Explanation
OpenPGP	OpenPGP	A means for management of public-private keys. RFC 4880: *OpenPGP Message Format*
Cryptographic Message Syntax Standard	PKCS#7	RFC 2315: *PKCS #7: Cryptographic Message Syntax - Version 1.5*

Term	Shorthand	Explanation
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 sets to be signed by the KSK
Signed Key Response	SKR	A data structure containing private key-generated signatures, specifically KSK signatures for DNSKEY sets

Table 2. Other Security Terms

3.3 Other Networking Terms

A few other terms are used than might need definition for a general audience

Term	Shorthand	Explanation
User Datagram Protocol	UDP	A context-free, best-effort transport protocol for sending data across the Internet
Transmission Control Protocol	TCP	Connection-oriented, octet-order guaranteed transport protocol for sending data across the Internet
Maximum Transfer Unit	MTU	The maximum number of octets 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

Table 3. Other Networking Terms

3.4 Summary of Recommendations

Recommendation 1: The Root Zone KSK Rollover should follow the procedures described in RFC 5011 to update the Trust Anchors during Key Signing Key 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 in order 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 minimize the time to recover due to difficulties involving the incoming KSK, an SKR generated only by the incumbent KSK should be generated in parallel with the SKR generated by the incoming KSK.

Recommendation 15: The RZM partners should develop and document the process of having to use the incumbent KSK generated SKR.

3.5 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.

3.6 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.

4 Abridged History

4.1 Deployment of DNSSEC in the Root Zone

In 2009, the RZM partners collaborated[4] 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 portion of the KSK was published in the Root Zone and as trust anchors.

[4] Details of DNSSEC deployment in the Root Zone are published at http://www.root-dnssec.org/

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 specify 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."

4.2 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:

1. 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.
2. The KSK rollover should be performed as soon as practical with an emphasis on preparedness.

[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
[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

3. Measurements and monitoring are the key modes highlighted to gauge the [technical and end-user] impact of a KSK rollover should one be implemented.
4. KSK rollover should take place regularly.
5. Public notifications to multiple, diverse stakeholder groups should be made in advance of a KSK rollover event, providing significant advance notice.
6. Further investigation is needed on operational stability, repeated KSK rollovers and [the likelihood of and impact of] non-compliance with RFC 5011.

4.3 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].

4.4 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 as well as the state of the code base at that time (open source DNS implementations in particular.) The report recommended communication action to publicize the Root Zone KSK key rollover, encouraged testing to collect and analyze resolver behaviors, the 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”, and the 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, aka RFC 5011, is present but unpredictable. And thirdly, that the size of DNS responses has been thought to be a concern when it relates to the occurrence of underlying UDP packet fragmentation and reversion to TCP queries.

4.5 ICANN Convenes Root Zone KSK Rollover Design Team

[11] http://www.ietf.org/proceedings/87/slides/slides-87-dnsop-6.pdf
[12] https://www.icann.org/en/system/files/files/sac-063-en.pdf

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, as presented in this document.

5 High-level Description of Rolling a KSK

The plan derived in July 2013, which is not far removed from plans for rolling any other KSK, follows these steps:

1) An incoming KSK key pair (public and private) is generated.

2) The incoming KSK public key is placed in the Root Zone and/or made available to relying parties.

3) In a deviation from other zones, the new Root Zone KSK public key sits in a state where it becomes accepted by all concerned that it is indeed the next KSK. In addition to passively being accepted, the new Root Zone KSK public key is made available on various electronic and non-electronic media to allow resolver operators and developers that have servers that do not support RFC 5011 time to include the new trust anchor in their systems and products. (For "other zones", this step is replaced by informing the holder of the DS record that there is an incoming KSK.)

4) The signing process switches from using the incumbent KSK private key to the incoming KSK private key.

5) The incoming KSK is now in a state of transition as the signatures generated by the incumbent KSK expire or otherwise disappear from the operational view.

6) The incumbent KSK public key is removed from the Root Zone (without revocation).

7) In another deviation from normal operations, the incumbent Root Zone KSK is reintroduced for the purpose of marking it revoked as per RFC 5011 guidelines. This separate step is designed to accommodate ZSK operations, which include rolls of that key without over-sizing DNS responses for the Root Zone complete key set.

6 Design Team Approach

The Design Team considered several aspects of a Root Zone KSK Rollover, and produced recommendations from each area of study to guide the development of an implementation plan by the Root Zone Partners.

- Operational Considerations: the impact on end-users of the Internet and the operators of the DNS systems, and services used by those end-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.

6.1 Operational Considerations

Impact on end-users of the Internet and operators of DNS systems are anticipated to occur during two of the steps above. When the incoming KSK public key is added to the Root Zone, the size of the response for the root DNSKEY set will grow. When the incumbent KSK private key is no longer generating signatures, validation using that public key will cease to work as expected.

With an enlarged response to DNSKEY, it is possible that fragmentation of UDP packets may occur with slightly different results over IPv4 and over IPv6. Already there are Internet components that consider fragments to be anomalous and filter them. For DNS, which maintains no state regarding sent responses, this means a client might not get an expected response. There is also the potential for a larger UDP response to exceed the query's specified DNS payload buffer size, therefore increasing the level of truncated responses and the subsequent re-query using TCP.

Once the incumbent KSK no longer signs the Zone Signing Key, with the implication that the incoming KSK is generating signatures, a DNSSEC validator with only the incumbent KSK configured as a trust anchor will fail to validate signed DNSSEC responses. The validator will "fail shut" meaning that it will regard all signed DNS responses as invalid.

An end 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 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 (see Recommendation 8).

6.2 Protocol Considerations

6.2.1 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 use *Automated Updates of DNS Security (DNSSEC) Trust Anchors* as described in RFC 5011, if the DNS software used supports this mechanism and is configured to use this mechanism to update the Root Zone Key Signing Key.

The online Validating Resolvers that are unable or unwilling to use Automated Updates of DNS Security Trust Anchors will need to be updated manually during the Key Signing Key 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 (see Section 11 for details), and the incumbent Trust Anchor must not be removed before the Root Zone is signed with the incoming Root Zone KSK. Furthermore, in terms of following prudent operational practice, the incumbent Trust Anchor should 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 and they are described below.

Recommendation 1: The Root Zone KSK Rollover should follow the procedures described in RFC 5011 to update the Trust Anchors during Key Signing Key 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 in an Internet Draft, entitled "*DNSSEC Trust Anchor Publication for the Root Zone*" within the IETF[13], but more review is needed in order to arrive at a stable consensus document that provides advice to implementers.

The Design Team supports community discussion and review of the 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.

6.2.1.1 Further Discussion of RFC 5011

In preceding text there is mention of resolvers "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, they 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 "pushing outwards" of managed configuration files. RFC 5011's update mechanism runs counter to that, with the configured fleet machines learning new data, diverging from the centrally managed configuration.

With that in mind, large operators will have a manual process in place, a process that will make use of various automated mechanisms. One automated system might be a tool that follows RFC 5011's update mechanism. 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

[13] http://tools.ietf.org/html/draft-jabley-dnsop-validator-bootstrap-00

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 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 to be aimed towards tool developers and DNS software vendors.

6.2.1.2 Other Trust Anchor Formats

Ever since the initial signing of the Root Zone, ICANN has made available the trust anchor in non-DNS formats via a website[14]. These trust anchors provide a non-critical-path means to distribute and receive the root zone trust anchor, i.e., a means outside of DNS operations. (The website does require accessing the DNS to reach the files.) Given the non-critical-path consideration, new trust anchors can be distributed. At some point in the future, it is possible to add trust anchors of different DNSSEC cryptographic algorithms[15] to emphasize new capabilities needed. This can also be a means for pre-populating resolvers in advance of an emergency-triggered rollover.

6.2.1.3 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 fashion are authentic, and take advantage of whatever verification mechanisms 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

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

[15] https://www.iana.org/assignments/dns-sec-alg-numbers/dns-sec-alg-numbers.xhtml#dns-sec-alg-numbers-1

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.

6.2.1.4 Systems Integrators

One distribution method of DNSSEC trust anchors is via systems integrator, for example, a package maintainer or an operating system vendor. In this case, the systems integrator 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 systems integrators and work closely with them to formalize processes to ensure that trust anchor distribution using integrator-specific channels is robust and secure.

6.2.1.5 System Administrators

Systems 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[16] 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. This undermines the effort to adequately support authenticity. It seems possible that additional review (with backwards-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.

[16] Listed at https://www.iana.org/dnssec/files

Further, observed retrievals of authentication-supporting digital signatures suggests that few, if any, relying parties have been making use of the digital signatures. Gaining trust is not simply done by providing digital signatures, but comes from active promotion.

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.

6.3 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.

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 processes in order 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. Because of this, all changes to the Root Zone DNSKEY RRset, e.g., adding and/or removing keys as required by a key rollover, must 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. A detailed examination of response size issues, with accompanying recommendations, appears later in this document. A Root Zone KSK rollover schedule designed with the aforementioned considerations in mind is also included later in this document.

6.4 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 ground might stem from questions regarding the cryptographic strength of the chosen key size or algorithm.

With the initial publication of SP 800-57, part 1 (*Recommendation for Key Management*) in 2005, the US 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.

6.4.1 Finite Field Cryptography

The 2048 bit asymmetric RSA key is considered to be equivalent to 103 bits symmetric key in ECRYPT II's 2012 Yearly Report on Algorithms and Key Sizes[17]. The same report recommends using at least 96 bits of security for ~10 year protection. The NIST *Recommendation for Key Management-Part 1: General (Revision 3)*[18] 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) *Référentiel Général de Sécurité*[19] also considers the 2048 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 (~15 days), and the Design Team believes that the 2048 bit RSA key should be safe for a further five years unless there is a significant technological breakthrough in the large integer factorization area.

[17] http://www.ecrypt.eu.org/ecrypt2/documents/D.SPA.20.pdf
[18] http://csrc.nist.gov/publications/nistpubs/800-57/sp800-57_part1_rev3_general.pdf
[19] http://www.ssi.gouv.fr/uploads/2015/01/RGS_v-2-0_B1.pdf

6.4.2 Elliptic Curve Cryptography

Another algorithm option available for DNSSEC is the Elliptic Curve Digital Signature Algorithm (ECDSA) that is defined in RFC 6605[20]. ECDSA has some properties that would make it desirable to use as an algorithm for Root Zone Key Signing Key. 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 (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 Section 7 – Operational Considerations).

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.

6.4.3 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 the Root Zone to be signed by all algorithms matching the configured Trust Anchors and thus the rollover to a different algorithm would require a different approach than for rolling the KSK. This provides further practical motivation to avoid a change in algorithm at this time. The Design Team has contacted the vendor regarding the issue and the vendor's requirement, and there is the expectation that it will be relaxed for future, unscheduled, KSK rollovers.

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 Design Team recommends maintaining the existing algorithm and key size for the incoming KSK for the first Root Zone KSK rollover.

[20] https://tools.ietf.org/html/rfc6605

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

6.5 Coordination and Communication

6.5.1 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 upcoming 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.

In order 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 roll of the KSK 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.

6.5.2 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.

Understanding the operational impact of a KSK rollover on validators and on the root servers themselves is facilitated by data collection by root server operators over the course of the KSK rollover. 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 DNS-OARC[21]. 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.

6.5.3 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[22]. 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 normal during the KSK rollover period; since the KSK rollover period is expected to be longer 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, we recommend that ICANN coordinate with the Root Zone Maintainer to ensure that

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

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.

7 Impact on Validating Resolvers

7.1 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 a response to indicate that the requestor should use TCP for the 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 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.

7.1.1 The Measurement Experience

An experiment has been designed and set up to reproduce the environment of the root server situation in order to evaluate what impact large packet sizes might have 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 who 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 in order 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.

7.1.2 Test Results

In a 5 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 who 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) while 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.

7.1.3 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, but 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 Key Signing Keys, two 1024 bit RSA Zone Signing Keys and three signatures – one by each Key Signing Key and one by one of the Zone Signing Keys. 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 as 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 section 8.2 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 who 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 who failed 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.[23]

7.2 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.

7.2.1 Test Results

Using the experiment described above (Section 7.1.1), 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.

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.

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

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.

7.2.2 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.

8 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.

8.1 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 that might otherwise go unnoticed but could be a valuable metric when observing the impact of the Root Zone KSK key roll operation. Users (at a screen) likely do not detect this and thus never open 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 keyroll 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 or near-real time information. This is a motivating factor to provide advance notice to impacted parties, avoid time spans when staffs are thin, and prefer times when contacts can easily be made.

8.2 Self-Test Facilities

As far as 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" ought to be present at the same path names. This will encourage better tools to be produced, such as tools to help assist in vetting a key, tools to 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. Faster than real-time is needed to allow for reasonable observation of the process. But in real-time is beneficial too to reduce the effects of testing.

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 consideration.

There are existing examples of such test beds[24, 25] that may be used as a model for future testing.

8.3 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.

9 Implementation

The proposed key rollover process was first conceived shortly in July 2013 and has since then been vetted and refined. The process described here should be considered as a draft and may be further improved by the RZM partners before implementation.

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. E.g., Quarter 1 and Q1 do not necessarily mean

[24] http://keyroll.systems/
[25] http://icksk.dnssek.info/fauxroot.html

January to March. The incoming KSK is noted as "KSK-NEW", the incumbent KSK is "KSK-2010".

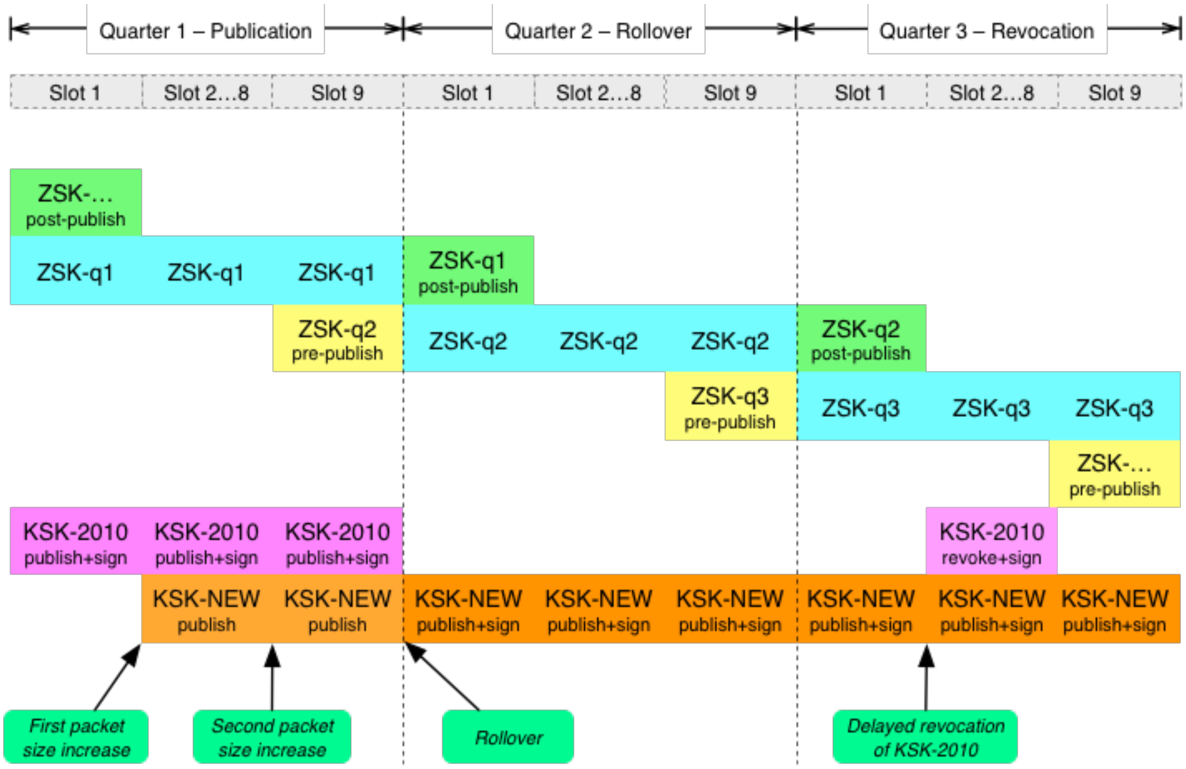

Figure 1. Rollover Scheduling

9.1 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 prior to 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 notional increase is based on a comparison of the size of a DNS response at this stage if no key rollover was underway to the size during the key rollover process.) During the last slot of Q1, at the ZSK rollover, the response packet size is increased from 833 octets to 1,158 octets.

9.2 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.

9.3 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.

9.4 Response Packet Size Impact

A desired objective is to avoid UDP fragmentation as far as possible, and the following are some relevant response size constraints:

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

Table 4. Packet Size Thresholds

Results of testing presented earlier indicate potential problems with some IPv6 resolvers and their handling of large responses. The first and most present 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 re-introduced and flagged with the revoke bit. For full compliance with RFC 5011, it is a requirement to double-sign the DNSKEY RRset with both the incoming Root Zone KSK and the incumbent Root Zone KSK during the revocation phase. 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. The table below 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.

Time	DNSKEY during roll	RRSIG during roll	DNSKEY response size during roll	DNSKEY response size during non-roll
Q1 slot 1	1x KSK + 2xZSK	1x KSK	883 octets	883 octets
Q1 slot 2 … 8	2x KSK + 1xZSK	1x KSK	1,011 octets	736 octets
Q1 slot 9	2x KSK + 2xZSK	1x KSK	1,158 octets	883 octets
Q2 slot 1	1x KSK + 2xZSK	1x KSK	883 octets	883 octets
Q2 slot 2 … 8	1x KSK + 1xZSK	1x KSK	736 octets	736 octets
Q2 slot 9	1x KSK + 2xZSK	1x KSK	883 octets	883 octets
Q3 slot 1	1x KSK + 2xZSK	1x KSK	883 octets	883 octets
Q3 slot 2 … 8	2x KSK + 2xZSK	2x KSK	1,297 octets	736 octets
Q3 slot 9	1x KSK + 2xZSK	1x KSK	883 octets	883 octets

Table 5. Packet Sizes During Rollover

(The color coding in the above table corresponds to the graphic below.)

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 re-define, publish, develop and distribute code, as well as press the code into operations, this option is not deemed feasible for this KSK key rollover.

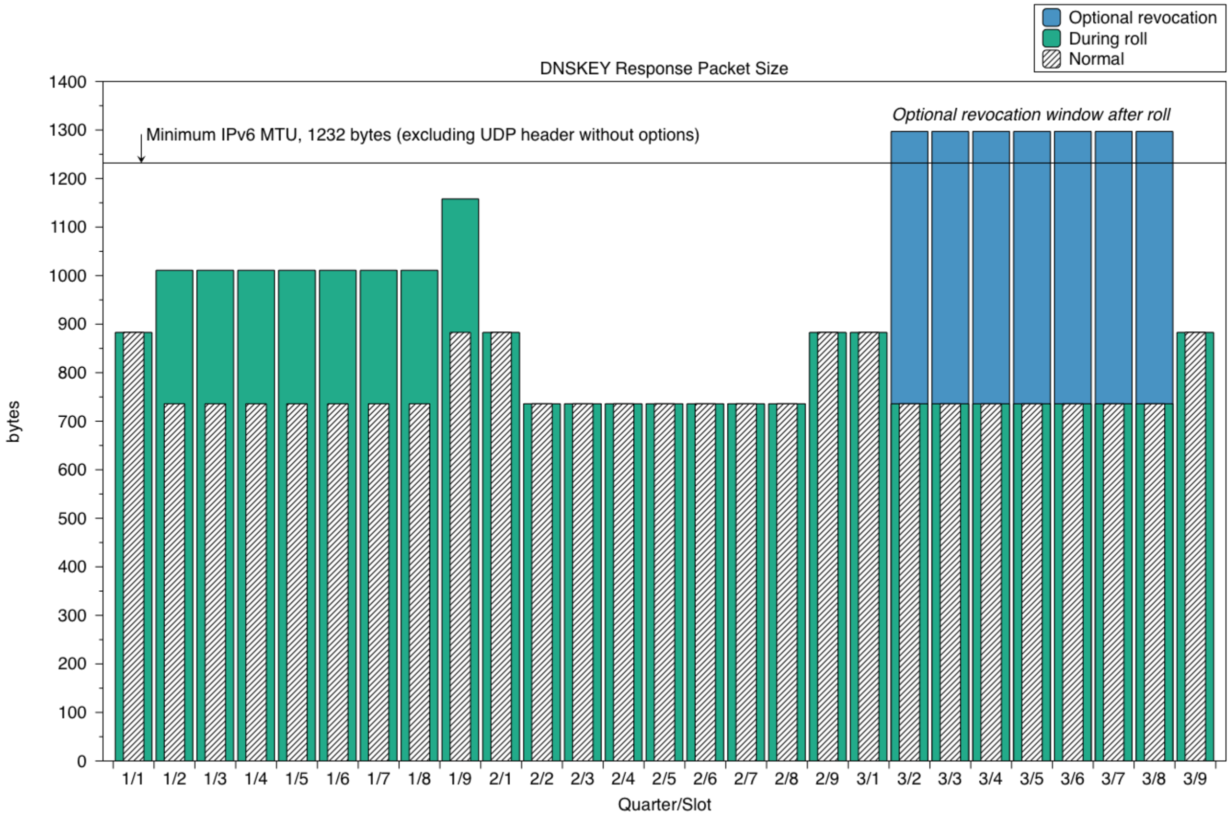

Figure 2. DNSKEY Response Packet Sizes

9.5 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 quickly 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 DITL collection[26] 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. Adding the incoming KSK (to the first server) requires 40 days, leaving just 30 more days to complete the task within one ZSK roll period. The original incremental deployment stretched for more than 4 months.

10 Rollback

In case there are serious problems detected after the introduction of the incoming KSK, DNSKEY RRsets signed by only the incumbent KSK should be prepared and ready for deployment. These RRsets are in *Signed Key Response* (SKR) format and can be produced using the same Root Zone KSK key ceremonies as the non-rollback RRsets. Criteria for such a rollback needs to be developed further by the RZM partners.

Recommendation 14: To minimize the time to recover due to difficulties involving the incoming KSK, an SKR generated only by the incumbent KSK should be generated in parallel with the SKR generated by the incoming KSK.

Recommendation 15: The RZM partners should develop and document the process of having to use the incumbent KSK generated SKR.

Rollback SKRs containing DNSKEY RRsets need to be prepared for all quarters of the process. During Q1 and Q2, 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. During Q3 the rollback SKR consists of DNSKEY RRsets with the incoming KSK and the current ZSK(s), signed by the incoming KSK. The revoked incumbent KSK is omitted.

Thresholds

Tests to date of DNSSEC deployment indicate that such tests have a margin of error of about 5%. This is taken to mean that any statement relating to an amount of

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

damage occurring will have to recognize that 5% of the population (people or recursive servers - depending on how the measurements are conducted) may suffer some degraded performance without detection. From this, a defining a specific metric is not viewed as the one way to go about defining a trigger for rollback.

Further, it is not clear what form damage will take. It could be an errant deployment, an errant strain of code, an errant procedure or a random act of the Internet. For this reason, maintaining contact with channel partners and opening up means for reporting problems is the first step, using judgment then to react to the reports.

Besides severity and spread of damage, it is not clear, as in there are many use cases, whether rollback would cause more damage than going forward, mitigating problems as the are detected.

11 When?

Given the existing operational environment, there are four days in the calendar year when a new Root Zone KSK can take over for the incumbent. Those four days are the first days of quarters, or the firsts of January, April, July and 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.[27]

Operationally reasonable means that the dates involved should avoid weekends, holidays which impact work schedules, and times when operations staffs are operating on a thin margin. Given the need to align three dates with a global audience, not all of this may be accommodated. Adding to the challenge, in 2016 and 2017, each quarter begins on a Friday, Saturday or Sunday. No quarter begins on any other day of the week until 2018. (The fourth quarter of 2015, October 1, begins on a Thursday but there will not be a plan in place, much less needed testing accomplished to have a key roll on that date.)

A non-technical impact is the planned IANA stewardship transition. This makes recommending a specific date impractical at this moment.

12 Risk Analysis

12.1 Risks associated with Insufficient Preparation

Description	Impact	Likelihood	Mitigation

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

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; 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	Need to prepare communications and testing; feasibility studies of mechanisms used in field; long-term effort to develop measurement of updated trust anchor acceptance

12.2 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 web site	Low	Unlikely	Monitoring of availability
Equipment with out-of-sync trust anchor sets via lack of maintenance	Low	Likely	Communications plan

12.3 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[28]" lingering effects; negative caching recommendations

12.4 Addition of Incoming KSK causes DNS Message Size to Exceed Limits

Description	Impact	Likelihood	Mitigation
Transition of keysets causes over-sized 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

12.5 Operational Errors Occur

Description	Impact	Likelihood	Mitigation
Botched KSK roll will end momentum for DNSSEC adoption	High	Unlikely	Careful design/review

[28] http://iepg.org/2010-03-ietf77/dnssec-goes-wrong.pdf, http://www.potaroo.net/ispcol/2010-02/rollover.html

Indefinitely postponing a key rollover increases the impact if it becomes urgent	High	Unlikely	Commitment 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

13 Design Team Roster

13.1 Community Volunteers

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

13.2 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, US DoC NIST
- Scott Rose, US DoC NIST
- Doug Montgomery, US NIST
- Ashley Heineman, US DoC NTIA
- Vernita Harris, US DoC NTIA

14 References

- RFC 5011: Automated Updates of DNS Security (DNSSEC) Trust Anchors
  https://tools.ietf.org/html/rfc5011
- SAC063: SSAC Advisory on DNSSEC Key Rollover in the Root Zone
  https://www.icann.org/en/system/files/files/sac-063-en.pdf
- DNSSEC Practice Statement for the Root Zone KSK Operator
  https://www.iana.org/dnssec/icann-dps.txt
- DNSSEC Practice Statement for the Root Zone ZSK Operator
- 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
- Establishing an Appropriate Root Zone DNSSEC Trust Anchor at Startup
  https://tools.ietf.org/html/draft-jabley-dnsop-validator-bootstrap

15 Appendix: 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.

15.1 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's BIND (http://www.isc.org)
- NLNetLab's 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/)

15.1.1 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's TBD recursive server (aside from Knot)
- PowerDNS TBD

15.2 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.

15.2.1 Linux

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

15.2.2 BSD

- FreeBSD ports
- NetBSD pkgsrc
- OpenBSD ports

15.2.3 Others

- Apple iOS, OS X
- Google Android, ChromeOS
- Microsoft
- Cisco
- Juniper
- Belkin
- Cisco / Linksys
- Wind River (RTOS)
- QNX (RTOS)
- OpenVMS
- OpenWRT

15.3 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, hence there may be internal reviews that need to know of the incoming Root Zone KSK.

- Google Public DNS
- OpenDNS
- Neustar DNSAdvantage
- Symantec ConnectSafe
- Level 3
- Censurfridns
- Comodo

- Dyn Internet Guide
- Liquid Telecom

In addition to the above 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.