Root Zone Algorithm Rollover Study

Report providing recommendations for the selection of a cryptographic algorithm for the DNS root zone and recommendations for future implementation plans.

Design Team Report
23 May 2024

TABLE OF CONTENTS

Executive Summary

The Internet Corporation for Assigned Names and Numbers (ICANN) is beginning preparations to be ready to change the Domain Name System Security Extensions (DNSSEC) signing algorithm used for the root zone. ICANN, in its role as the Internet Assigned Numbers Authority (IANA) functions operator, is working in cooperation with Verisign in its capacity as the root zone maintainer.

Consistent with common practice for the deployment and operation of DNSSEC, the root zone key signing key (KSK) is used to sign the root zone apex Domain Name System Key (DNSKEY) Resource Record Set (RRset). This set includes one or more zone signing keys (ZSKs), which are used to sign all other RRsets in the root zone. Changing the algorithm of the root zone refers to the additional step of changing the signing algorithm that has been used since 2010 when the root zone was first signed with DNSSEC. An algorithm rollover affects the operation and management of both the KSK and the ZSK.

Changing the algorithm 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 a key rollover. Adequate deployment of the new algorithm in validating resolvers is an equally critical aspect of an algorithm rollover.

In December 2014, ICANN solicited volunteers from the community to participate with the root zone management (RZM) partners in a design team to develop the root zone KSK Rollover Plans. The deliverables 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.

In January 2023, a new design team was formed to study how to change the cryptographic algorithm used to sign the root zone. This design team, the authors of this report, present the results of the study and provide a series of recommendations to inform the RZM partners in developing a detailed implementation plan for changing the algorithm.

This document contains terminology related to DNSSEC, Internet security, and networking.

1. Introduction

Managing the Domain Name System (DNS) root zone KSK is a unique responsibility and a truly global effort. IANA is tasked with creating the KSK key pair and then digitally signing information with the KSK private key. Internet users at large utilize the corresponding KSK public key to perform DNSSEC validation. Total success requires IANA to faithfully perform its duties and for all those relying on DNSSEC protections to perform their duties as well.

While IANA's work is clear – managing cryptographic assets within carefully constructed facilities, maintaining transparency and trust in operations, generating DNSSEC-signed data, and publicizing the trust anchors in use – the work performed by others may not be as obvious. For those who are running a simple validating recursive name server, all of the necessary changes may be automatically performed by built-in, updated mechanisms, but for others, configuration updates may be required. For those producing off-the-shelf software or services, there may be software or manually verified configuration work to do. Deeper into the specifics of the DNSSEC protocol implementations, there are untested and unanticipated operational considerations to take into account. Wherever there are gaps, service interruptions may occur. Due to the Internet's highly distributed nature, there is no overarching responsible authority that can guard against this.

IANA is studying what is needed to change the cryptographic signing algorithm used by the root zone. This study will help prepare all impacted parties. It needs to consider that IANA only manages certain centralized components and that this success requires actions by users all over the world. Possible motivations for changing the algorithm include:

- The algorithm of the current key is found to be weak during the key's anticipated lifetime.
- Other algorithms are more desirable for their properties.
- Ensuring the DNS community is prepared for an urgent algorithm rollover in the future by performing a non-urgent rollover first.

The design team was tasked with two goals: 1) provide guidance on how to select an algorithm; and 2) investigate how an algorithm rollover could be conducted. Drafting the detailed operational plans necessary for a rollover is out of scope for the design team.

Furthermore, this study is conducted without a specific timeframe for an algorithm rollover – providing recommendations for specific dates is out of scope of this study. Therefore, some recommendations may be intentionally vague, with the expectation that they will be refined through subsequent work.

In completing this work, the design team applied established methods used by top-level domains (TLDs) to successfully change their signing algorithm while also considering elements unique to the DNS root zone.

The report is structured to address the design team's two goals. Recommendations are made in the context of the section in which they are defined. A complete list of recommendations is reproduced in Appendix: List of Recommendations.

Unless specifically mentioned otherwise, this document uses the term "rollover" to mean an algorithm rollover.

2. Abridged History

The root zone was first signed with DNSSEC in 2010. The first change of trust anchor root keys – known as a key rollover – was completed in 2018 after several years of consultation, design, and testing. The first key rollover was initially scheduled to occur on 11 October 2017, but was delayed for one year as data showed that a significant number of resolvers used by Internet service providers and network operators were not yet ready. The first key rollover finally happened on 11 October 2018. After the rollover, in March 2019, ICANN's Office of the CTO published a Review of the 2018 DNSSEC KSK Rollover.

In November 2019, IANA published a Proposal for Future Root Zone KSK Rollovers. This sought to create a predictable approach to managing the Root Zone KSK's life cycle by establishing a standard KSK rollover schedule. Based on experience from the first KSK rollover, IANA expects to perform future rollovers by pre-publishing the new key as a standby key for about two years before using it for signing.

In January 2021, ICANN's Second Security, Stability, and Resiliency (SSR2) Review recommended that IANA work with other root zone partners and the global community to develop a consensus plan for future root DNSKEY algorithm rollovers. This recommendation was adopted by the ICANN Board on 22 July 2021 and led to the formation of the design team in January 2023.

In the first quarter of 2023, ICANN announced that the next root key rollover would use the existing algorithm. A new key was generated, however, due to issues involving obtaining secure key storage and signing equipment, this rollover was suspended. The rollover restarted in April 2024 when a new key using the existing algorithm was generated on new signing equipment. The second key rollover is expected to occur in the DNS root zone on 11 October 2026.

3. High-Level Description of Changing the Algorithm

While common operational practice has been to use one algorithm at a time, DNSSEC is designed to allow zone administrators to use keys of multiple algorithms at once. This simplifies the operator's workload and limits the size of responses. There would typically be little benefit in using several algorithms, as long as a generally supported and trusted algorithm is available.

The DNS root zone was signed in 2010 using RSA/SHA-256. The first KSK rollover did not change the DNSSEC algorithm, keeping RSA/SHA-256 and the same 2048-bit key size. Changing the DNSSEC algorithm requires a new DNSKEY that uses a different algorithm than what is currently present in the DNSKEY RRset, followed by the removal, possibly after revocation, of DNSKEY records with the incumbent algorithm.

Verisign, in its capacity as the root zone maintainer, changed the key size of the ZSK from 1024-bit to 2048-bit on 1 October 2016. While there was a noticeable change of larger signatures and DNSKEY sets, this was not a DNSSEC algorithm change.

The DNSSEC specifications define requirements for when more than one algorithm is in use. Section 2.2 of the Request For Comment (RFC) 4035 requires the existence of signatures for each RRSet using every algorithm present in DNSKEY RRset. These requirements are intended to apply to servers; Section 5.11 of RFC 6840 clarified this by specifying that validators "MUST NOT insist that all algorithms signaled in the DNSKEY RRset work." Despite this clarification, it is clear that some validating resolvers have not adhered to that recommendation, instead applying a stricter interpretation during validation.

Given the final requirement listed in section 2.2 of RFC 4035 ("Including RRSIG RRs in a Zone"), the root zone must include signatures for each RRset using both algorithms during an algorithm rollover. Double-signing the root zone would increase the size of the root zone and the size of responses as they would then include signatures for both algorithms. These changes will affect both the root name servers and the resolvers that request DNSSEC records.

Although there has been no root zone algorithm rollover to date, many TLDs have rolled to new algorithms. Before 2018, however, nearly all TLD algorithm rollovers were to new RSA variants, e.g., from algorithm 7 (RSASHA1-NSEC3-SHA1) to algorithm 8 (RSA/SHA-256). In 2018, CZ.NIC was the first TLD operator to roll to a non-RSA algorithm.

Rolling the KSK for a non-root domain, e.g., a TLD, relies on updating the Delegation Signer (DS) record at the parent level, whereas rolling the KSK for the root zone requires updating the DNS trust anchor, either through following the RFC 5011 process or through out-of-band mechanisms.

4. Potential Impacts

To explore the potential impacts of a root zone algorithm rollover on the DNS ecosystem, the impact to each of its individual components must be analyzed.

Many DNS queries originate with stub resolvers, which are generally simple library functions linked into applications. Stub resolvers, by definition, only communicate with recursive resolvers, and most do not perform DNSSEC validation. A root zone algorithm rollover is not expected to

have any direct impact on DNS stub resolvers that do not perform validation. For stub resolvers that do implement DNSSEC validation, the main concerns are whether the new algorithm is widely supported and whether the implementation follows RFC 6840, which requires ignoring unknown algorithms.

Recursive resolvers also come in validating and non-validating variants, except that it is very common for non-validating recursive resolvers to request and receive DNSSEC-related data, even when they do not validate it. Thus, all recursive resolvers may be impacted by response size changes, with the need to perform more queries over the Transmission Control Protocol (TCP) during and after an algorithm rollover, depending on the characteristics of the chosen algorithm. As with stub resolvers, validating recursive resolvers may also be negatively impacted if they do not follow RFC 6840 requirements to ignore unknown algorithms. Additionally, those that do not support a new algorithm, but correctly ignore it, will lose the protections that DNSSEC provides once keys and signatures using the old algorithm are no longer available.

Home devices and other middle boxes often participate in the DNS resolution process, either as forwarders or full resolvers. These may be impacted because (a) their software may be of lower quality due to the inexpensive nature of these products, and (b) the software is often not regularly updated. Enterprise networks also sometimes employ middle boxes to filter certain traffic in the name of improving security. These products have been known to filter DNS records with unknown parameters or values, which may be of concern for a change in algorithm.

Depending on a particular choice of a future algorithm, root name servers may be significantly impacted by increased response sizes. As discussed in Section 5.6 (Message Size Considerations), current operational practices that favor DNS message truncation over User Datagram Protocol (UDP) packet fragmentation could lead to a significant increase in DNS queries over TCP transport. The root name servers may also be impacted by increased query loads from misbehaving recursive resolvers. For example, during the old key revocation period following the 2018 KSK rollover, some resolvers were observed sending repeated DNSKEY queries at very high rates. Although this behavior did not appear to cause resolution failures, it was never fully explained.

The root zone may increase significantly in size during or after an algorithm rollover, e.g., due to doubling the number of DNSSEC signatures or longer public keys and/or signatures of the new algorithm. There is a risk that this could have an impact on root zone distribution within the Root Server System. However, as long as the expected size increase is communicated to relevant parties well in advance, the risk is considered small and manageable.

ICANN, as the root zone KSK operator through the IANA, and Verisign, as the root zone ZSK operator, will obviously be impacted by an algorithm rollover. Therefore, an algorithm rollover requires close coordination and planning by both parties. It may also be necessary for one or

both to update or add new features to the systems they use in the production and handling of root zone key material and root zone data. In addition, the quarterly root zone key signing ceremonies are impacted by the need to transmit and sign multiple key signing request documents to support backout (and extended backout) options.

5. Algorithm Selection Criteria

The primary questions around an algorithm rollover are (1) when to initiate the process and (2) selecting the successor algorithm. Both decisions are related. The timeframe when an algorithm needs to be initiated or completed may restrict the pool of available candidate algorithms or a superior candidate algorithm may arise whose benefits outweigh the operational cost associated with performing a rollover. Criteria for algorithm selection include both cryptographic attributes of possible algorithms and operational concerns of selecting and using each algorithm. Criteria for when to initiate an algorithm rollover can be reactive or proactive. The actual rollover steps will not change greatly in either situation, only the triggering event for the rollover will.

One task of the design team is to define a mechanism for selecting a suitable algorithm, the criteria that guide the selection of the algorithm, and guidance on when the criteria should be applied. Several aspects of changing the algorithm for the DNS root zone were considered:

- Cryptographic considerations – ensuring that the system as a whole has sufficient cryptographic strength.
- Protocol considerations – the extent to which existing, documented protocol elements are sufficient to accommodate a root zone algorithm rollover.
- Operational considerations – the impact on Internet users and DNS operators and the services being used.
- Impact on root zone management – the impact on the processes involved in KSK and ZSK management by the IANA and the root zone maintainer.
- Resolver considerations – the impact on both validating and non-validating resolvers
- Message size considerations – the impact of an algorithm rollover on the size of DNS messages.

Each of these areas is individually explored in the sections that follow.

5.1. Cryptographic Considerations

A set of technical criteria must be met when selecting a successor algorithm for the root zone. Some criteria may involve a trade-off compared to 2,048-bit RSA. The primary criteria are cryptographic strength and associated practical considerations. Any decision to change the signing algorithm needs to balance these two criteria, particularly if the reason for the change is weaknesses in current or future algorithms.

5.1.1. Cryptographic Strength

The successor algorithm should be of equal or greater strength (measured in equivalent symmetric strength) as the currently used algorithm. There are two types of strength considerations: short-term and long-term. For short-term strength, all of the current standardized elliptic curve algorithms for DNSSEC are stronger than the currently used 2,048-bit RSA. For long-term strength, all of the post-quantum (PQ) algorithms that are being considered by the U.S. National Institute for Standards and Technology (NIST) are expected to be stronger than the 2,048-bit RSA. For a discussion of PQ algorithm use in the DNS, see OCTO-031.

5.1.2. Practical Considerations

Different algorithms have different sizes for the public keys and signatures. Reducing the size of keys and signatures for the root zone KSK has a direct positive effect on the root service because it will reduce bandwidth usage and cause fewer requests over TCP. This is a primary reason for not switching to PQ algorithms in the near future: The keys and signatures of PQ algorithms currently under consideration by NIST are anticipated to be much larger than for current algorithms.

Another factor can be the effort necessary to sign or validate a signature; however, due to the variety of signing and resolving software, this is difficult to measure and is only one of many considerations when upgrading to cryptography that is otherwise better.

5.2. Protocol Considerations

Only DNS Security algorithms that are standardized through the Internet Engineering Task Force (IETF) and documented in the IANA registry for DNSSEC Algorithm Numbers will be considered. The registry contains algorithms for zone signing and transaction security, however only algorithms that can be used for zone signing will be considered. This registry also includes private use algorithms (253/254). These will also not be considered.

The RFC 8624, “Algorithm Implementation Requirements and Usage Guidance for DNSSEC” (and future updates), lists implementation recommendations for DNSSEC signing and validation, with varying levels of requirements for DNSSEC validation (“must not,” “not recommended,” “may,” “recommended,” and “must”). These requirements may be changed over time by subsequent RFC documents that update RFC 8624. Algorithms labeled “must not,” “not recommended,” “may,” and “recommended” should be avoided, leaving only those labeled “must.”

These requirements exist to maintain secure DNSSEC authentication over time while also ensuring interoperability. Algorithm deprecation or introduction is done gradually to allow implementations to update while remaining interoperable.

At the time of writing, two algorithms are designated “must” for both signing and validation. One is the current RSA/SHA-256 (algorithm 8). The other is the Elliptic Curve Digital Signature Algorithm (ECDSA) Curve P-256 with SHA-256 (algorithm 13).

Given the universal use of the root zone in DNS resolution, it can be assumed all validators will have to support algorithms used to sign the root zone.

Recommendation 1: The algorithm used to sign the root zone must be specified as "MUST implement" for DNSSEC validation per RFC 8624 or its successor.

The recommendation above does not restrict the root zone partners from announcing their intention to select an algorithm not currently designated as a “MUST.” While this scenario seems unlikely, root zone selection could inform an update to RFC 8624 or its successor, thereby meeting this recommendation.

5.3. Operational Considerations

For any previously unused algorithm to be considered a candidate for use in the root zone, the following operational characteristics should be reviewed.

The algorithm should have widespread support in deployed DNS software. Support of the algorithm should be measured using both public and private resolvers. It is not sufficient to assume widespread deployment solely based on the algorithm’s implementation in software, it should be demonstrated that the algorithm has sufficient global uptake.

In addition to widespread adoption, the algorithm should have been proven through use by a significant number of TLDs. When calculating TLD usage, the number of TLDs which have the algorithm deployed at the apex for a significant amount of time, as well as the number of DS records signed using the algorithm within the TLD, should be considered.

The quantitative threshold used to determine that an algorithm has sufficient deployment will evolve over time as the operational environment changes (e.g., as future TLDs are introduced into the root zone).

Execution of the algorithm rollover for the root zone has unique circumstances surrounding the deployment and uptake of the trust anchors. Once a new trust anchor has been published, there may be several factors for a given DNS operator that would delay its deployment, including:

- The implementation of DNS hardware or software may not support the new algorithm and could fail in unexpected ways.
- The signature verification algorithm for a particular DNS implementation may be non-performant.

- Hardware-based DNS solutions may require component upgrades to support the new algorithm and may be gated by long delivery times due to supply chain or distribution issues.
- Other business factors may delay the deployment of the update required to use the new algorithm.
- The implementation may not fully support RFC 5011 rollovers as expected.

Recommendation 2: The proposed algorithm must be demonstrated to have sufficient operational implementation for DNSSEC prior to a rollover being performed with the objective that significant software vendors and resolver operators support the new algorithm. The measurements necessary to be deemed "sufficient" are determined by the detailed operational plan.

As with the prior recommendation, not having sufficiently wide adoption of an algorithm does not prevent its consideration; however, by the time the rollover is performed, actions need to have been undertaken to ensure all significant validators are able to support the new algorithm.

Recommendation 3: The publication of new trust anchors should happen significantly before the introduction of the new algorithm's DNSKEY records in the root zone.

5.4. Impact on Root Zone KSK/ZSK Management

The KSK and ZSK operators must be able to sign the zone using the selected algorithm and comply with the requirements set forth by the DNSSEC Practice Statement (DPS). The 6th edition of the KSK DPS and version 2.1 of the ZSK DPS state that key material is protected using hardware security modules (HSMs) that are valid at FIPS 140-2 level 4 overall.

Since those versions of the KSK operator and ZSK operator DPS documents were published, NIST has updated the latest version of FIPS 140, "Security Requirements for Cryptographic Modules" to FIPS 140-3. The number of commercially available HSMs certified at level 4 is declining, and the HSMs currently in use by both the KSK and ZSK maintainers have been declared end-of-life by the vendor.

At the time of writing, NIST's Cryptographic Module Validation Program (CMVP) lists four active FIPS 140-2 level 4 certifications by two vendors. If historical certifications are included, the list is slightly longer, although the number of vendors remains quite low.

The lack of choice and competition at FIPS 140-2 level 4 presents risks to the overall management of the KSK and ZSK. The current design of the KSK function further restricts the list of potential vendors, such as the need for a form-factor suitable for storage in a safe. Changing the baseline requirement to FIPS 140 level 3 will accommodate additional vendors whose products already support a wider suite of algorithms.

Recommendation 4: There must be diverse support for the chosen algorithm across all aspects of the management of KSK and ZSK operations. This includes diversity in HSM products and vendors, and diversity in cryptographic implementations.

Recommendation 5: IANA should reevaluate the requirement for FIPS 140-2 level 4 certification, and any compensating control mechanisms needed to mitigate additional risk should be developed.

Note, in parallel with this study, IANA published the 7th edition of the KSK DPS that states HSMs shall meet either FIPS 140-2 or FIPS 140-3, at level 3 or higher.

5.5. Resolver Considerations

5.5.1. Impact on resolvers requesting DNSSEC resource records

Most resolvers[1] request DNSSEC records be returned in far greater numbers than DNSSEC validation is measured as being performed. This suggests that most resolvers, even those not performing DNSSEC validation, will be impacted by larger response sizes that would be sent from root servers (or forwarders relaying root zone responses) when data sets are signed by both the outgoing (old algorithm) and incoming (new algorithm) key.

A resolver can declare an upper message size limit it is able to receive over UDP, which is primarily set to minimize fragmentation. The responding server planning a response above this limit should reply with a truncated response by setting the truncated (TC) bit. The resolver is thus prompted, but not required, to retry the query in TCP and generally does. In 2020, a DNS Flag Day was held to promote the use of 1232 bytes as a default maximum UDP buffer size, suggesting that root zone responses above that size would initiate significant use of TCP. If a resolver does not opt to suggest a maximum size, the assumed limit is merely 512 bytes.

5.5.2. Impact on validating resolvers

During the active rollover period, signatures and keys with both the incoming and the outgoing algorithm are present in the root zone. Validating resolvers that support at least one of these algorithms continue to validate the signatures.

At some point, only signatures and keys of the incoming algorithm will be present in the root zone. Resolvers that do not support this algorithm, and are in accordance with RFC 4035 Section 5.2, should treat the zone as unsigned unless they are configured with a trust anchor for another algorithm.

The table below shows validators configured with a single trust anchor for an algorithm unused for signing will not accept signatures made with another algorithm.

[1] Around 70 percent of queries have the DO-flag set at A-root, and around 80% at J-root.

Validator trust anchor algorithms	Zone signed with algorithm A	Zone signed with algorithm A & B	Zone signed with algorithm B
A	OK	OK	Error
A & B	OK	OK	OK
B	Error	OK	OK

Both the cryptographic library used by the resolver and the resolver software itself must support the new algorithm.

Resolvers use multiple methods to get a new trust anchor. Resolvers developed and configured to support RFC 5011, “Automated Updates of DNSSEC Trust Anchors,” will get and trust the new key automatically if the resolver's cryptographic library supports the algorithm. These seem to be most common, but a significant amount of resolvers rely on other mechanisms for getting the keys, such as by software or operating system update, manual configuration, or querying the trust-anchor.xml file on the IANA web site.

Recommendation 6: An algorithm rollover must accommodate nearly all resolvers and thus support RFC 5011, an update to trust-anchor.xml file, as well as outreach for those who have to configure it manually.

5.6. Message Size Considerations

Recently, the DNS community has determined that fragmented DNS responses are harmful. This was the subject of the 2020 DNS Flag Day and an Internet-Draft called “Fragmentation Avoidance in DNS.” DNS software vendors and service operators now configure their software to truncate large DNS responses at sizes smaller than the standard 1500-byte Ethernet MTU. The popular open-source DNS resolver packages (BIND, Unbound, PowerDNS, Knot Resolver) all have their default UDP buffer size limit set to 1232 bytes. Additionally, many of the DNS root servers now have truncation limits between 1232 and 1472 bytes and may have different limits for IPv4 and IPv6.

To understand how an algorithm rollover may interact with these limits, four types of root zone queries and their responses were considered:

- Referrals
- NXDOMAIN
- DNSKEY
- Priming

The table below indicates the prevalence of each type of query in data from A-root and J-root for the first week of April 2023. The table also displays response sizes observed currently, where

the root zone is signed with algorithm 8 (RSA/SHA256) using 2048-bit KSKs and 2048-bit ZSKs. Referral sizes can vary depending on the query name, the number of delegated name servers, and their corresponding glue records. NXDOMAIN response sizes also vary, depending on the query name. The table below shows that only the size of the DNSKEY response changes throughout the various (non-algorithm) rollover stages because all other responses have records signed by only a single, active ZSK.

Type	Prevalence	Normal	ZSK Rollover	KSK Rollover	ZSK+KSK Rollover	KSK Revocation
Referral	38.5%	1171[2]				
NXDOMAIN	58.1%	1087[3]				
DNSKEY	0.34%	864	1139	1139	1414	1425
Priming	1.74%	1097				

Prevalence and message size for different types of root zone DNS responses. Prevalence data comes from recent A-root/J-root traffic. Response sizes are based on the current root zone RSA key parameters and as observed during the 2018 KSK rollover (no algorithm change).

During an algorithm rollover, however, all responses will increase in size if double ZSK signature methods are used. The exact amount of increases, and whether they exceed previously mentioned truncation limits, depends on the choice of the new DNSSEC algorithm.

Since DNSKEY queries comprise less than 0.5 percent of the query traffic, there is relatively little concern about DNSKEY responses exceeding truncation limits. Note, however, that for the previous KSK rollover, the root servers experienced a significant increase in DNSKEY queries, up to 10 percent of the total traffic, during the period of revocation of the old KSK.[4]

Priming queries currently comprise about 2 percent of root name server traffic. A shift to TCP for priming queries may be of concern for some root server operators.

[2] This is the size of a referral for example.com, since the .com delegation has the largest referrals for a fixed query name length.
[3] This is the size of an NXDOMAIN response for kkkkkkkkkkkkkkkkkkkkkkkkkkkkkkkkkkkkk.kkkkkkkkkkkkkkkkkkkkkkkkkkkkkkkkkkkkk since NXDOMAIN response size generally only depends on query name length and 99% of queries are for names of 70 characters or fewer.
[4] Moritz Müller, Matthew Thomas, Duane Wessels, Wes Hardaker, Taejoong Chung, Willem Toorop, and Roland van Rijswijk-Deij. 2019. Roll, Roll, Roll your Root: A Comprehensive Analysis of the First Ever DNSSEC Root KSK Rollover. In Proceedings of the Internet Measurement Conference (IMC '19). Association for Computing Machinery, New York, NY, USA, 1–14. https://doi.org/10.1145/3355369.3355570

The bigger concern, however, is referral and NXDOMAIN responses; the bulk of query traffic. If these responses exceed truncation limits, they could significantly shift the majority of query traffic to TCP. The following table shows expected response sizes during rollovers from the current algorithm (8) to other potential DNSSEC algorithms.

Response Type	Alg 8	Alg 8→13	Alg 8→14	Alg 8→15	Alg 8→16
Referral	1171	1266	1298	1266	1316
NXDOMAIN	1087	1370	1466	1370	1520
DNSKEY	864	1116	1212	1052	1152
DNSKEY with double ZSK rollover	1139	1470	1598	1374	1499
Priming	1097	1191	1223	1191	1241

Expected response sizes during DNSSEC rollovers from algorithm 8 to other potential candidate algorithms. Cells with yellow background indicate sizes above 1232 bytes, while cells with red background indicate sizes above 1400 bytes. Note that response sizes during KSK revocation are generally not a concern and are excluded from this table because keys for potential candidate algorithms are significantly smaller than RSA-2048.

Recommendation 7: The evaluation of candidate algorithms should consider response sizes and the impacts from increased TCP traffic to root name servers.

Recommendation 8: Response sizes and the amount of traffic subject to truncation, fragmentation, and TCP fallback should be accurately modeled by replaying real root server query traffic against root zones signed with candidate algorithms.

Recommendation 9: A rollover implementation team should coordinate closely with root server operators to understand any then-current operational limits with respect to message sizes, truncation limits, and TCP query capacity.

Recommendation 10: In the event that rolling from the existing RSA configuration to a new algorithm poses significant operational obstacles with respect to message size and TCP transport, the rollover implementation team may consider incorporating a temporary reduction in the size of the RSA ZSK below 2048 bits (e.g., 1536 bits[5]).

5.7. Selection Criteria Summary

The sections above provide recommendations for criteria that are expected to guide an algorithm selection process. It is important to note that an algorithm is not required to meet the

[5] Viktor Dukhovni, APNIC Blog: DNSSEC algorithms for TLDs (and everyone else)

criteria at the time it is proposed – however the criteria should be met prior to insertion into the root zone. In such a case, these criteria will inform the activities required to ensure the algorithm is suitable for the root zone.

The design team considered whether a weighting should be assigned to the selection criteria, but ultimately decided weighting was overly prescriptive and potentially difficult to define for a generic rollover. All criteria are important and must be considered when assessing an algorithm; however, some may be less important for any given choice of algorithm. One example is that message size is not a significant consideration for a rollover from ECDSA to Edwards-curve Digital Signature Algorithm (EdDSA).

The design team also chose not to define criteria that selected an absolute winner. Two or more algorithms may satisfy the criteria and both may be deemed suitable for signing the root zone. If the current algorithm is suitable, then by definition there is no need to change it. This, however, does not preclude other motivations for changing the algorithm, which would need to balance tradeoffs for other properties such as cryptographic strength, message size, or complexity of signature verification.

Implied in the definition of these criteria is that the criteria must also apply to the current algorithm. While it is unlikely that the current algorithm would suddenly become unsuitable (there would be wider impact on TLS, etc.), the suitability of the current algorithm should be assessed periodically. One potential outcome may be incorporating a review or assessment in a future version of the proposed rollover schedule, for example, performing an assessment before key generation.

Recommendation 11: The current algorithm should be assessed periodically for applicability for signing the root zone for its anticipated lifetime. A logical time is before the generation of a successor key.

6. Implementation

The design team considered three approaches when investigating how a rollover could be conducted. These approaches, the timeline as to how an algorithm rollover should be planned, and mechanisms for trust anchor distribution are described below.

6.1. Algorithm Roll Execution

The KSK rollover operational implementation plans used in 2017,[6] 2018,[7] and for future non-algorithm rollovers[8] describe eight distinct rollover phases:

A. Generation

[6] https://www.icann.org/en/system/files/files/ksk-rollover-operational-implementation-plan-22jul16-en.pdf
[7] https://www.icann.org/en/system/files/files/2018-ksk-roll-operational-implementation-plan.pdf
[8] https://www.icann.org/en/system/files/files/proposal-future-rz-ksk-rollovers-01nov19-en.pdf

B. Replication
C. First Keyset Signing
D. Publication and Standby
E. Rollover and Active
F. Revocation
G. First Deletion
H. Final Deletion

Each phase lasts one or more calendar quarters. Phases D through F involve changes to root zone content, while the surrounding phases involve activities that take place only during key signing ceremonies.

For the 2017–2018 KSK rollover, the initial expectation was that each phase would cover one calendar quarter. However, the plan did account for the possibility that it may be necessary to extend a phase, or even "back out" to a previous phase. In fact, phase D was extended for a whole year as described in the Abridged History above.

The root zone KSK operator and ZSK operator DNSSEC Practice Statements[9,10] describe a model in which each calendar quarter is divided into nine separate signature slots. In this model, slots 1 and 9 are reserved for ZSK operations, while slots 2 through 8 are reserved for KSK operations. For example, to conduct a quarterly ZSK rollover, a new ZSK is prepublished in slot 9 and an old key is post-published in the following slot 1. KSK rollovers, on the other hand, begin at the start of slot 2, and are complete by the end of slot 8.

A root zone algorithm rollover is likely to require some changes to both the phases and the slots. In a non-algorithm rollover, a new key can be introduced (published) before it is used for signing. However, during an algorithm rollover, the mandatory algorithm rules of RFC 6840 (Section 5.11) would require signatures from a new algorithm when that algorithm's key is present. In other words, a root zone algorithm rollover cannot prepublish a new key in phase D before it is used for signing in phase E. These two phases must be combined into one.

Furthermore, the same mandatory algorithm rules of RFC 6840 mean that a new algorithm's KSK and ZSK must be introduced at the same time. It is not feasible to introduce them in separate slots. It will be necessary to deviate from the traditional slot usage schedule and to update the DNSSEC Practice Statements.

Recommendation 12: A future operational implementation plan for a root zone algorithm rollover should use the previous non-algorithm rollover plans as a starting point, with phases and other details adjusted as necessary to accommodate the specifics of an algorithm rollover.

[9] https://www.iana.org/dnssec/procedures/ksk-operator/ksk-dps-20201104.html#section-6.6
[10] https://www.verisign.com/assets/ROOT_VerisignDNSSECPracticeStatement_v2.1_finalized.pdf

The operational plan should include the option to extend and/or back out phases whenever possible, to maximize flexibility.

Recommendation 13: RZM partners should collaborate closely on operational implementation plans and necessary updates to the DNSSEC Practice Statement (DPS) documents to accommodate algorithm rollover timing.

6.1.1. Classic Approach

The following picture illustrates one possible schedule for an RFC 5011-compatible algorithm rollover with double signing.

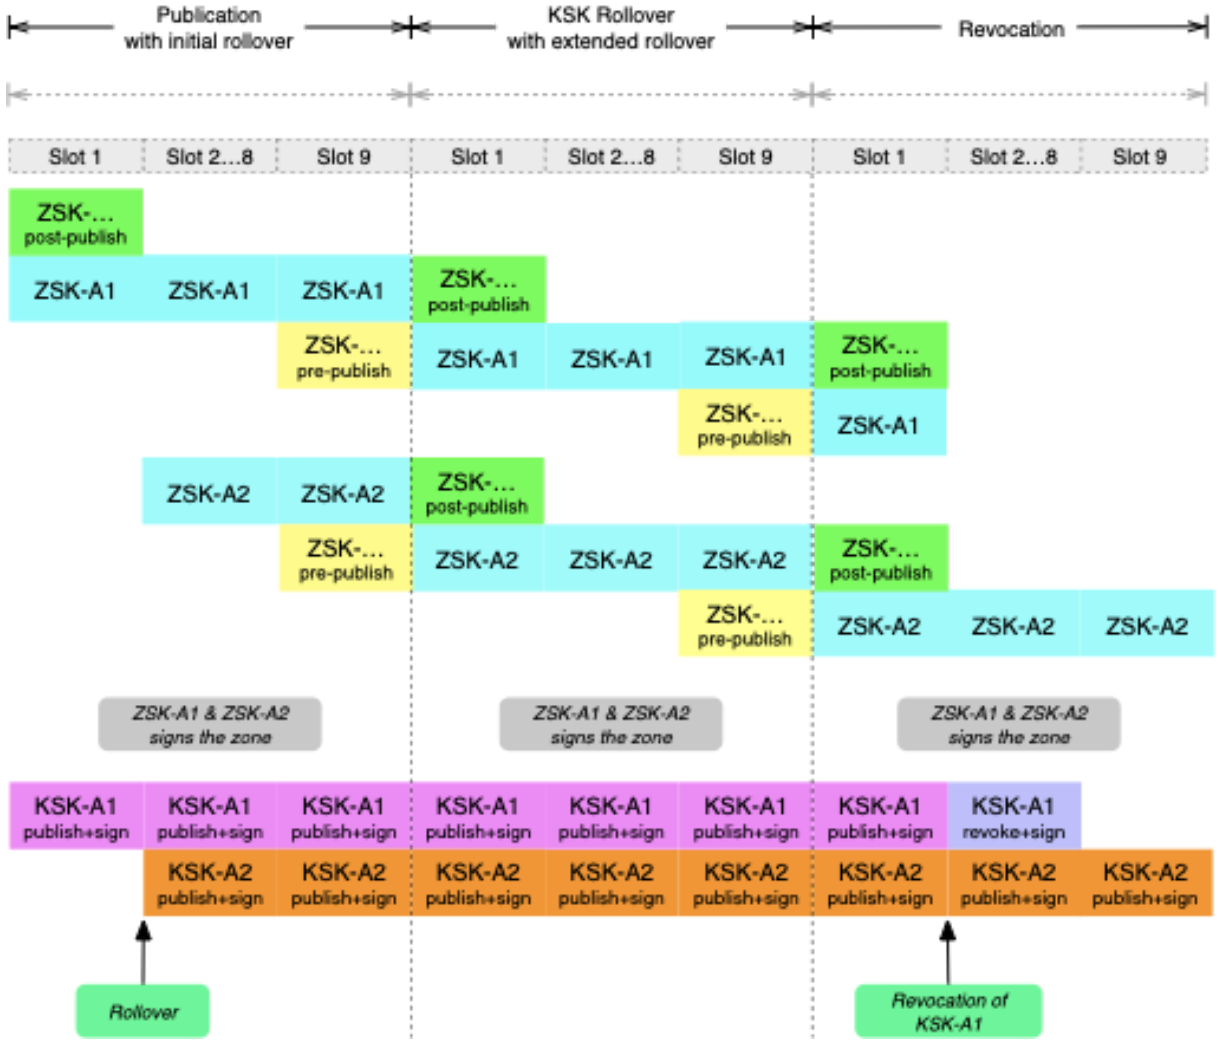

Using double signing is the simplest algorithm key rollover process and is most consistent with previous rollovers. It is also safest from a protocol point of view. A considerable downside is the larger message sizes that exceed minimums and therefore impact root server operators because of retries over TCP. The full implementation details of this approach, for example the

pre-publication of signatures, are not material to this study and are expected to be detailed in an implementation plan.

6.1.2. Alternate Approach

Should the requirement that a zone must be signed with all trust anchors' algorithms be changed in the future, then the incoming root zone KSK could be pre-published separately from the ZSK, enabling an algorithm rollover similar to a non-algorithm rollover without double RRSIGs. In addition to lowering the overall complexity of the algorithm roll, this would also lower the DNS response message sizes for priming DNSKEY and NXDOMAIN queries.

This alternate approach relies on the property that validators should evaluate and accept any available and valid path from a configured trust anchor to response data. It also requires that implementations of RFC 5011 accept a new trust anchor as long as it is signed by an existing trust anchor and does not require that the new trust anchor is also signed by itself. This is consistent with what RFC 5011 states: that new trust anchors are validated by an existing trust anchor. It does not state that they must be self signed.

The following picture illustrates one possible schedule for an RFC 5011-compatible algorithm rollover if the requirement for double-signing is relaxed. Each phase – publication, rollover, and revocation – can be prolonged with minimal adjustments if a longer phase duration is required for operational reasons.

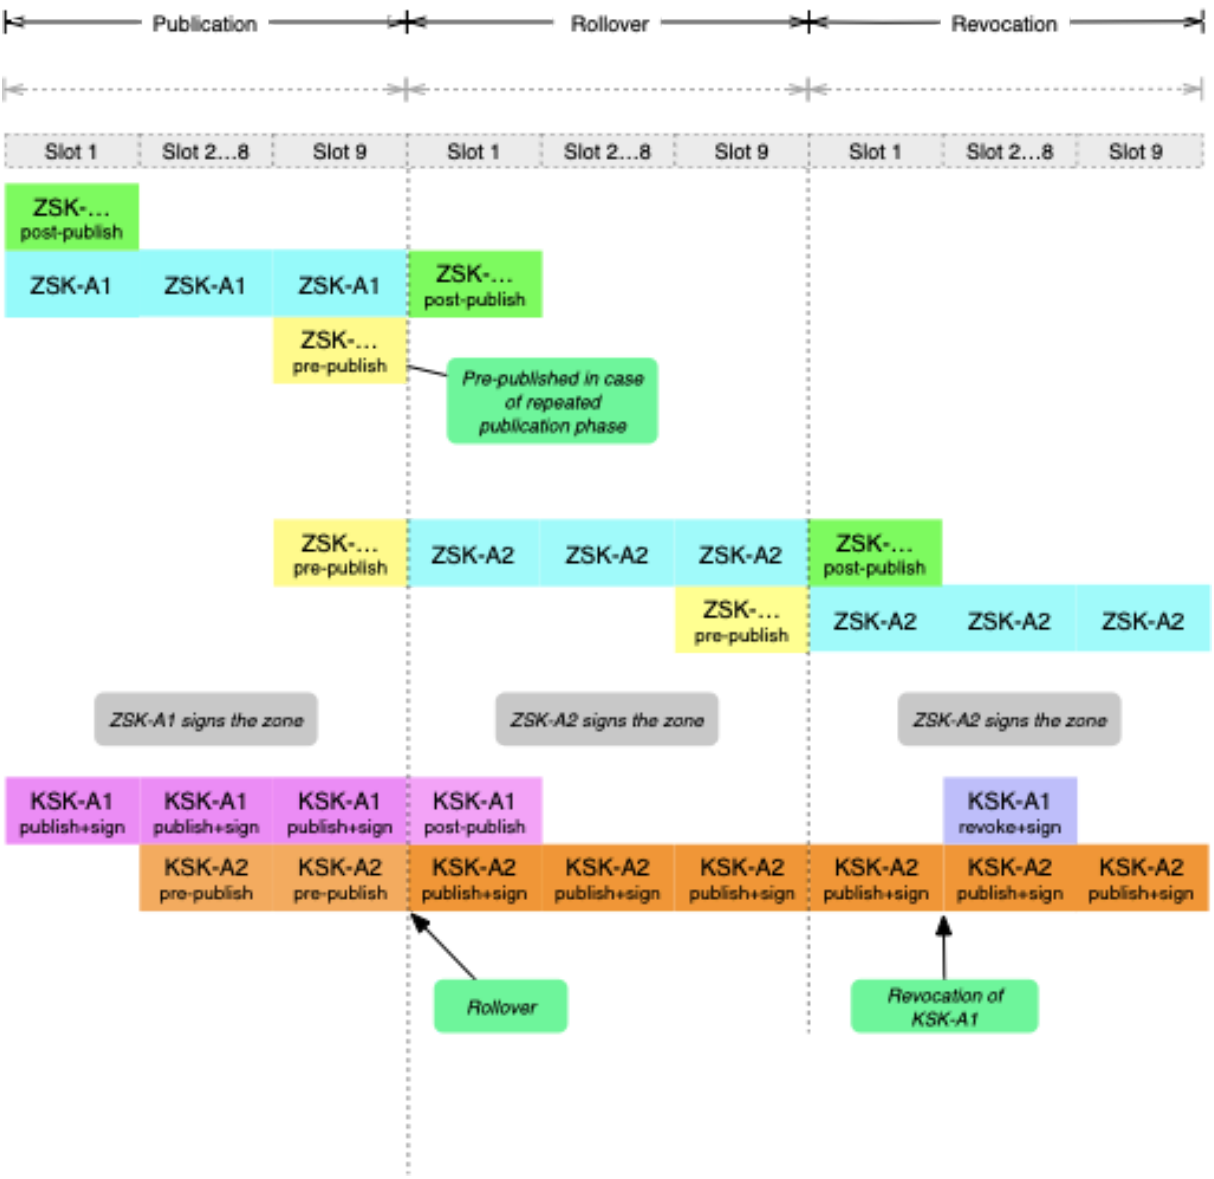

Recommendation 14: After selecting the future algorithm, it should be decided whether to begin the rollover process with a pre-publication of the trust anchor (suitable specification updates allowing) or whether to avoid pre-publication. To inform the decision, it should be assessed how double signing the root zone with the new algorithm would impact root server operators, resolver operators, potential (e.g., packet flood) attack victims, and whether pre-publication benefits outweigh the risks.

6.1.3. Hybrid Approach

It is possible to implement a combination of the classic and alternate approach where the new KSK is prepublished in the zone, signed only by the incumbent KSK and without a ZSK of the new algorithm. In other words, the hybrid approach is equivalent to the classic approach with pre-publication of the KSK.

In this scenario, resolvers that would have accepted the new trust anchor in the alternate approach will adopt the new trust anchor earlier than in the classic approach. Its adoption may be tracked per RFC 8145, “Signaling Trust Anchor Knowledge in DNS Security Extensions (DNSSEC)” and RFC 8509, “A Root Key Trust Anchor Sentinel for DNSSEC.”

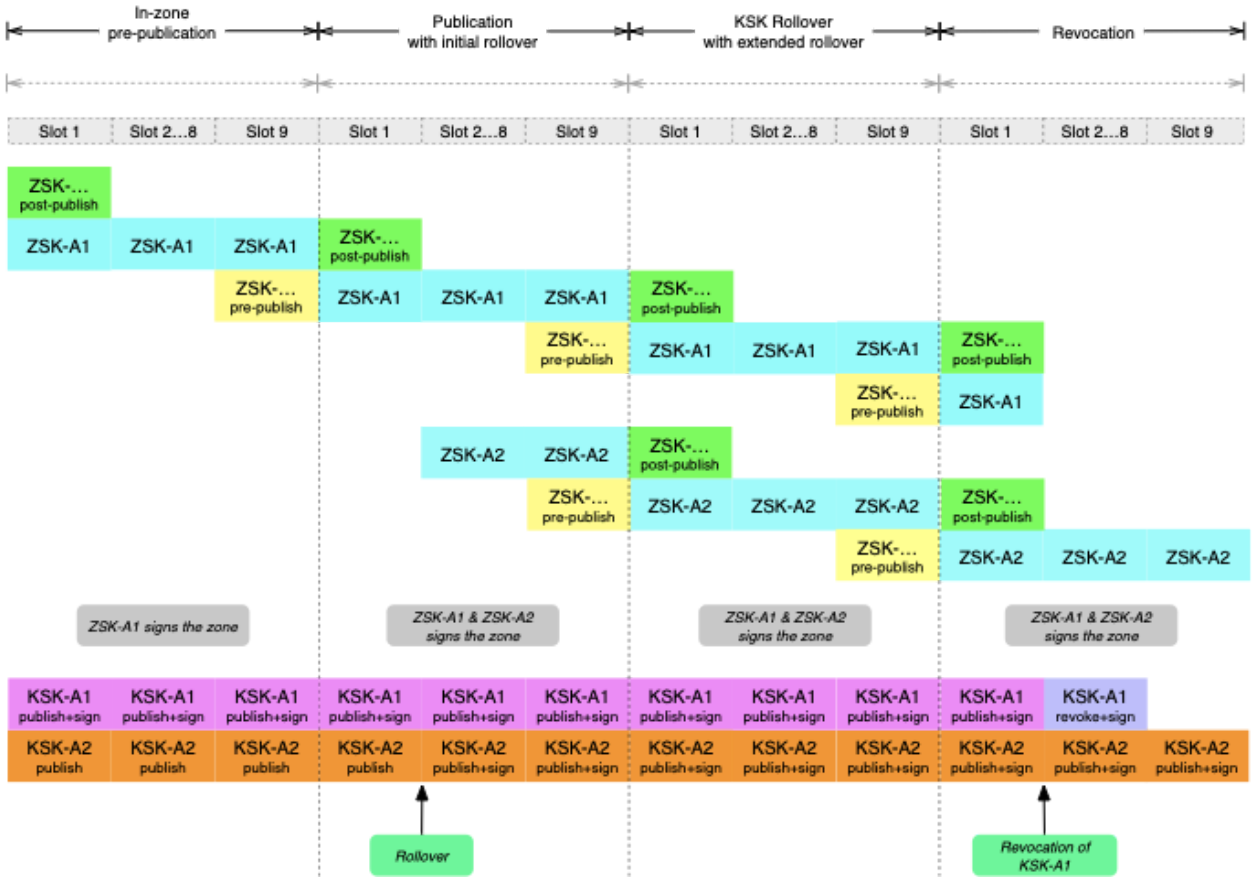

6.2. Message Size Mitigation

An algorithm rollover following the classic approach from RSA 2048-bit keys to any of the other current widely supported algorithms will lead to large responses, causing response truncation and retries over TCP. If the alternate approach is not available, the following mitigations should be considered to reduce message size:

- Suspend ZSK rollovers during some or all of the algorithm rollover phases. This will reduce the DNSKEY RRset size by removing pre- and post-published zone signing keys during the first and last slots of each quarter, approximately 20 days. This would reduce the size of DNSKEY response messages, but not impact other types of responses, which are the majority.

- Reduce the ZSK size before the start of the algorithm rollover. This would reduce the size of the DNSKEY resource record set and the size of signatures generated by that key. Initial analysis indicates that reducing the size of an RSA ZSK, described in Section 5.6 Message Size Considerations, might be sufficient to significantly reduce truncation if rolling to algorithms producing signatures with sizes comparable to ECDSA P-256/SHA-256.

If and when an algorithm roll requires a large number of bytes, for example when double signing returned answers, the root zone may become an attractive "unwitting accomplice" in distributed denial-of-service attacks. In addition to the load on the root server operators, this is another concern that requires consideration.

6.3. Timeline

According to IANA's proposed schedule for KSK rollovers, the generation and publication of the proposed successor key occurs in conjunction with the destruction of the former KSK. The effect is that the root zone is almost continuously in a phase of a rollover.

To avoid unnecessary planning and operational costs, and to reduce potential for confusion in the community, the choice to perform an algorithm rollover should occur before the generation, dissemination, or publication of a same-algorithm successor key.

In the event that an algorithm rollover becomes urgent, IANA will need to decide whether to suspend or back out of the current non-algorithm rollover or continue and complete that rollover. Performing an algorithm rollover in conjunction with a normal non-algorithm rollover introduces significant design complexity and operational impact and should be avoided.

The design team notes that a goal of the future rollover plan is to provide operators and vendors with sufficient lead time to disseminate the new trust anchor. The proposal noted two years as a sufficient lead period. The motivation behind that proposal also holds for an algorithm rollover; sufficient time must be given for operators and vendors to update and disseminate their trust anchors.

Recommendation 15: Except in an emergency, a root zone algorithm rollover should not interrupt the schedule of an in-progress non-algorithm rollover.

6.4. Trust Anchor Distribution

Root zone trust anchors are distributed by IANA, which are posted at https://www.iana.org/dnssec/files, using mechanisms described in RFC 7958, which is being updated by draft-ietf-dnsop-rfc7958bis. The defined mechanisms are algorithm agnostic and work the same for all DNSSEC algorithms.

In the first key rollover, IANA updated the trust anchor files on 3 February 2017. This followed the successful replication of the KSK to the alternate facility on 2 February 2017.
The design team expects that IANA will follow a similar process for an algorithm rollover; that is, to prepublish the new key as a trust anchor on the IANA website before publishing the DNSKEY in the DNS. This will be the first time that a second algorithm will appear in the trust anchor distribution.

IANA is unable to measure the impact of making changes to the trust-anchor file. The expectation is that the files are consumed during the distribution of resolver software and tested before its distribution. However, these files can be consumed in an operational context; incorrect assumptions or errors in implementations may lead to failures in updating the trust anchors.

Owing to the files' distribution method, IANA can observe requests to obtain the trust anchor files. Analysis of access over several days revealed a cyclical daily access pattern indicating an operational dependency on this file. Further analysis indicated that there are still requests for the trust anchor from old versions of Unbound that are known to have issues with algorithm rollovers.

Recommendation 16: Trust anchors using new algorithms should be pre-published using the same existing trust anchor distribution mechanisms as for non-algorithm rollovers.

Recommendation 17: IANA should continue to monitor access to the trust anchor files before and during a rollover and have this information available to support communication and outreach efforts.

Recommendation 18: Software with operational requirements on the root trust anchors should include software version information in their requests, such as in the User-Agent HTTP header. This information can be used to inform impact assessments.

6.5. Communications

The change of algorithm introduces scenarios not present in the first KSK rollover; for example, rolling over to unsupported algorithms, the presence of keys and signatures of different algorithms in the root zone, and resolver behavior with larger message sizes. Furthermore, software that worked before may not work with a rollover to another algorithm.

The team notes that the impact to parties depends on the implementation approach. A classic approach would likely have significant impact on the root server operators and resolvers, as a significant portion of DNS traffic to the root could switch to TCP. Mitigations to reduce the volume of TCP traffic in the classic approach require a temporary step backwards in the security profile. This could be viewed negatively by various communities. Alternatively, the alternate or hybrid approaches may also impact the operators of validating resolvers, possibly resulting in end users experiencing resolution failures, though we have not seen this in testing.

Therefore, the communication plan will depend on the choice of algorithm, implementation approach, and the timeline for the rollover. The design team recognizes the extensive communications campaign conducted during the first KSK rollover and expects an equivalent level and scope of outreach would be required for an algorithm rollover.

Recommendation 19: The implementation plan should include a comprehensive communications plan that is relevant to the rollover implementation approach.

7. Testing

For this report, the design team carried out preliminary tests to understand the behavior of resolvers during a root algorithm rollover. We focused on resolvers that rely on RFC 5011 for trust anchor management. From the previous root KSK rollover, we know that resolvers can update the trust anchor of the root zone if the algorithm of the KSK stays the same. Here, we tested if implementations of RFC 5011 can also update to a new trust anchor of a different algorithm.

We tested the primary rollover schedule as described in Section 6.1.1, as well as the alternate and hybrid approaches described in Sections 6.1.2 and 6.1.3. Additionally, we tested how resolvers behave if the root rolls to an algorithm (RSA/MD5 – Number 1 and Private Algorithm – Number 253) not supported by the recursive resolver.

The results described here are only preliminary and for a limited set of recursive resolver software brands and versions. Therefore, we recommend that before scheduling a root algorithm rollover, a larger set of software should be tested. We tested BIND 9.19.12, Unbound 1.17.1, Knot Resolver 5.6.0, Akamai Cacheserve 7.7, and Windows DNS Server 2022.

Normally, testing RFC 5011 would require waiting at least 30 days for a resolver to accept and trust a new key. Some resolver implementations provide options that allow testing on an accelerated time scale, but others do not. This presents a significant barrier to performing these types of tests. We recommend that software implementing RFC 5011 provides options and documentation to enable such testing on short time scales.

7.1. Test Results

Resolver	Primary rollover schedule	Alternate approach	Hybrid approach	Primary rollover schedule to unsupported algorithm	Primary rollover schedule to private algorithm
BIND 9.19.12	OK	OK	OK	SERVFAIL	SERVFAIL
Knot Resolver 5.6.0	OK	OK	OK	SERVFAIL	SERVFAIL
Unbound 1.17.1	OK	OK	OK	OK (insecure)	OK (insecure)
Akamai Cacheserve 7.7	OK	OK	OK	SERVFAIL	SERVFAIL
Windows DNS Server 2022	OK	OK	Not tested	Not tested	Not tested

All tested resolvers can update their trust anchor using RFC 5011 with all tested rollover approaches if they support the new algorithm. When presented with an unsupported algorithm, some implementations fail to update their trust anchors. Instead of treating the root zone as insecure, some of the tested resolvers fail validation and return the return code SERVFAIL for all queries after the rollover.

Not all resolvers support RFC 5011 and instead rely on other trust anchor distribution mechanisms. In some cases, recursive resolvers might rely on a third party (e.g., the operating system) to provide and update the trust anchors including only a key of the new algorithm. For this reason, we test how resolvers behave if the provided trust anchor file includes only a key from an unsupported algorithm. Here, we additionally test PowerDNS Recursor 4.8, which does not support RFC 5011.

Resolver	TA File with Unsupported Algorithm	TA File with Private Algorithm
BIND 9.19.12	OK (insecure)	OK (insecure)
Unbound 1.17.1	SERVFAIL	SERVFAIL
PowerDNS Recursor 4.8	SERVFAIL	SERVFAIL

Two out of the three tested resolvers treat the root zone as bogus. Only BIND treats the root zone as insecure.

We also looked into which potential validating stub resolvers could be impacted by an algorithm rollover. There are not as many available, but Apple operating systems (iOS, iPadOS, macOS) possess an API for validating domains, as well as systemd-resolve distributed with most Linux-based operating systems these days.

None of them support RFC 5011 and the Apple client has no API to give it trust anchors; they will be distributed by Apple via software distribution when the new keys appear. The Linux software suite, systemd, can be tested with multiple trust anchors but seems to fall back to the hard-coded one whenever possible. This is likely not a problem for the real rollover, but it made testing more challenging.

We also could not test private implementations of public DNS resolvers, because we could not access them to configure our testbeds.

Recommendation 20: Before scheduling an algorithm rollover, a larger set of recursive resolver software versions should be tested to assess their compatibility to the proposed rollover scheme.

Recommendation 21: Implementers of recursive resolvers, which support RFC 5011 trust anchor management, should enable and document the possibility to test RFC 5011 rollovers with a short rollover frequency. In general, it is prudent for validators to put in easy switches for test configurations.

Recommendation 22: Before scheduling the algorithm rollover there should be a live testbed with proper timing to allow public resolver operators or other validator implementers to test their implementation before the real rollover.

8. Protocol Clarifications

The ability to prepublish the new trust anchor, without introducing the public key into the root DNSKEY RRset and/or signing with it, would be helpful. It would reduce the complexity of the algorithm rollover, and in particular reduce the DNS response message sizes for priming DNSKEY and NXDOMAIN queries. This would allow for an extended time period during which operators may incorporate the new trust anchor without incurring the costs of double-signing (see “Implementation” section).

However, Section 5.11, “Mandatory Algorithm Rules” of RFC 6840, “Clarifications and Implementation Notes for DNS Security (DNSSEC)” specifies that such pre-publication of trust anchors is not permitted, as any published algorithm is required to come with a valid chain of trust (i.e., a corresponding DNSKEY and RRSIGs with that algorithm are required in the zone).

The specification currently mandates that a zone signed with DNSSEC “MUST include a DNSKEY for each algorithm present in the zone's DS RRset and expected trust anchors for the zone.” The vagueness of “expected” aside, this disregards the fact that (root) trust anchors and DS RRsets have different natures, with the former being configuration items under management of the validator operator, and the latter effectively controlled by the zone owner.

The validator operator, by choosing which trust anchors to configure, curates their own set of algorithms, each of which is individually trusted[11]. It can thus be argued that the continued presence of the old trust anchor (with a generally supported algorithm) guarantees that resolvers can validate answers from the root zone even when a new trust anchor with a different algorithm is prepublished (with no corresponding changes to the root zone yet).

This report does not establish clarifications of the DNSSEC protocol, but merely points out aspects for which clarifications may be worthwhile. However, doing so is not without risk or cost. Specifics are expected to emerge external to this report, such as via an Internet Draft, and brought to the IETF for consideration and potential publication.

Recommendation 23: Depending on the choice of algorithm and operational factors, it may be preferable if the rollover’s double-signing phase is preceded by pre-publication of the trust anchor. As such an approach is not compliant with current specifications, the design team recommends that the pre-publication method be described in an Internet Draft and brought to the IETF for consideration and potential publication.

[11] It would be a contradiction to configure trust anchors with multiple algorithms for a zone, where one algorithm (A1) is deemed preferable over another (A2), and then be afraid of downgrading from A1 to A2. Instead, only A1 should be configured as a trust anchor in this case.

9. Conclusions

The design team was tasked with two objectives: to provide guidance for the selection of an algorithm for a future rollover, and to investigate how a future rollover could be conducted.

Algorithm selection considered several areas, including cryptographical, protocol, operational, management, and resolution; however, the design team did not consider how to resolve conflict between two or more candidate algorithms. The design team expects that, when it comes time to change the algorithm, IANA will announce the selected algorithm including how it meets the criteria defined herein.

The design team notes the significant impact to the root server system caused by increased message sizes, should algorithm rollover occur following the classic approach from the current 2048-bit RSA keys.

If a rollover is required on a compressed timeline, such that the classic approach is required, reducing the size of the ZSK appears to be a viable option to limit message sizes and avoid significant volumes of message truncation. Since the size of the ZSK was changed in 2016, the design team assumes reducing the size of the ZSK is viable.

However, given the goal of developing a standard plan for an algorithm rollover, the design team recommends the community work towards support for the alternate approach. The alternate approach has the benefit of smaller message sizes, regardless of the algorithms, while maintaining the status quo for cryptographic strength. Early analysis indicates that common open-source resolvers accept any viable path; however, further analysis is required to confirm this as a viable long-term option.

Recommendation 24: The community should advance the necessary work, through protocol modifications, software updates, and outreach, to ensure the alternate implementation approach is viable.

The observation that resolvers fail following a rollover to an unsupported algorithm is a concern for the agility of DNSSEC. In the event that the current signing algorithm becomes unsuitable, IANA may be forced to decide between changing the algorithm and thus disabling resolvers, or the continued use of an unsuitable algorithm. There is nuance in this condition – it is possible that an algorithm is known to a resolver software, however the resolver depends on a system or cryptographic library where the algorithm is unavailable, by configuration or lack of implementation. This risk is low when there are two well-supported algorithms.

When rolling a key, the process is often labeled emergency or non-emergency, but such a distinction is not accurate. All rollovers need to be graceful and well-planned, limiting disruption. The distinction is whether the process is planned in advance, run according to a

pre-communicated schedule, or started immediately in response to an unexpected event, internal or external.

The current KSK and ZSK use RSA/SHA-256 with 2048-bit keys as the signing algorithm. If the community discovers a critical weakness in that algorithm that cannot be compensated for by using larger RSA keys, the signing algorithm will urgently need to be changed. The current algorithm is widely used on the Internet in other applications, so such a weakness would likely affect large parts of the Internet beyond the DNS (e.g., certification authorities, websites).

An algorithm rollover may also take place before a weakness is discovered or exploited. Because the currently used root zone signing algorithm may be at risk in the future, there is benefit to migrating to a new algorithm before migration is considered urgent. A different algorithm may offer more advantages or mitigate perceived future threats. This is a proactive rather than a reactive approach with the goal to avoid an urgent algorithm rollover.

ICANN, with the involvement of the community, will need to evaluate the tradeoffs when considering a planned algorithm rollover. The operational steps in either a planned or urgent rollover remain the same: only the timeframe, communication plan, and error tolerance may differ. An urgent algorithm rollover may need to run more quickly and be more tolerant of client failures than a planned rollover.

Recommendation 25: Periodically planned algorithm rollovers should be performed to help ensure preparedness in the event the timeline needs to be compressed to facilitate an urgent rollover.

The design team understands IANA's ambition to rollover the KSK every three years. The second root zone KSK rollover is in its early phases occurring in parallel with a change to the cryptographic hardware. This second rollover will not change the algorithm: replacing the current end-of-life hardware is time-critical and should not assume the additional risks associated with an algorithm rollover. While there exists no known weakness in the current algorithm, the preparatory time must consider potential protocol modifications and software update cycles with years of lead time.

Recommendation 26: The community should prepare to be operationally ready for a root zone algorithm rollover. The next logical time for this is following the second KSK rollover, in approximately five years.

Appendix: List of Recommendations

The recommendations in this document are repeated below in the order in which they appear in the document.

Recommendation 1: The algorithm used to sign the root zone must be specified as "MUST implement" for DNSSEC validation per RFC 8624 or its successor.

Recommendation 2: The proposed algorithm must be demonstrated to have sufficient operational implementation for DNSSEC prior to a rollover being performed with the objective that significant software vendors and resolver operators support the new algorithm. The measurements necessary to be deemed “sufficient” are determined by the detailed operational plan.

Recommendation 3: The publication of new trust anchors should happen significantly before the introduction of the new algorithm’s DNSKEY records in the root zone.

Recommendation 4: There must be diverse support for the chosen algorithm across all aspects of the management of KSK and ZSK operations. This includes diversity in HSM products and vendors, and diversity in cryptographic implementations.

Recommendation 5: IANA should reevaluate the requirement for FIPS 140-2 level 4 certification, and any compensating control mechanisms needed to mitigate additional risk should be developed.

Recommendation 6: An algorithm rollover must accommodate nearly all resolvers and thus support RFC 5011, an update to trust-anchor.xml file, as well as outreach for those who have to configure it manually.

Recommendation 7: The evaluation of candidate algorithms should consider response sizes and the impacts from increased TCP traffic to root name servers.

Recommendation 8: Response sizes and the amount of traffic subject to truncation, fragmentation, and TCP fallback should be accurately modeled by replaying real root server query traffic against root zones signed with candidate algorithms.

Recommendation 9: A rollover implementation team should coordinate closely with root server operators to understand any then-current operational limits with respect to message sizes, truncation limits, and TCP query capacity.

Recommendation 10: In the event that rolling from the existing RSA configuration to a new algorithm poses significant operational obstacles with respect to message size and TCP

transport, the rollover implementation team may consider incorporating a temporary reduction in the size of the RSA ZSK below 2048 bits (e.g., 1536 bits).

Recommendation 11: The current algorithm should be assessed periodically for applicability for signing the root zone for its anticipated lifetime. A logical time is before the generation of a successor key.

Recommendation 12: A future operational implementation plan for a root zone algorithm rollover should use the previous non-algorithm rollover plans as a starting point, with phases and other details adjusted as necessary to accommodate the specifics of an algorithm rollover. The operational plan should include the option to extend and/or back out phases whenever possible, to maximize flexibility.

Recommendation 13: RZM partners should collaborate closely on operational implementation plans and necessary updates to the DNSSEC Practice Statement (DPS) documents to accommodate algorithm rollover timing.

Recommendation 14: After selecting the future algorithm, it should be decided whether to begin the rollover process with a pre-publication of the trust anchor (suitable specification updates allowing) or whether to avoid pre-publication. To inform the decision, it should be assessed how double signing the root zone with the new algorithm would impact root server operators, resolver operators, potential (e.g., packet flood) attack victims, and whether pre-publication benefits outweigh the risks.

Recommendation 15: Except in an emergency, a root zone algorithm rollover should not interrupt the schedule of an in-progress non-algorithm rollover.

Recommendation 16: Trust anchors using new algorithms should be pre-published using the same existing trust anchor distribution mechanisms as for non-algorithm rollovers.

Recommendation 17: IANA should continue to monitor access to the trust anchor files before and during a rollover and have this information available to support communication and outreach efforts.

Recommendation 18: Software with operational requirements on the root trust anchors should include software version information in their requests, such as in the User-Agent HTTP header. This information can be used to inform impact assessments.

Recommendation 19: The implementation plan should include a comprehensive communications plan that is relevant to the rollover implementation approach.

Recommendation 20: Before scheduling an algorithm rollover, a larger set of recursive resolver software versions should be tested to assess their compatibility to the proposed rollover scheme.

Recommendation 21: Implementers of recursive resolvers, which support RFC 5011 trust anchor management, should enable and document the possibility to test RFC 5011 rollovers with a short rollover frequency. In general, it is prudent for validators to put in easy switches for test configurations.

Recommendation 22: Before scheduling the algorithm rollover there should be a live testbed with proper timing to allow public resolver operators or other validator implementers to test their implementation before the real rollover.

Recommendation 23: Depending on the choice of algorithm and operational factors, it may be preferable if the rollover's double-signing phase is preceded by pre-publication of the trust anchor. As such an approach is not compliant with current specifications, the design team recommends that the pre-publication method be described in an Internet Draft and brought to the IETF for consideration and potential publication.

Recommendation 24: The community should advance the necessary work, through protocol modifications, software updates, and outreach, to ensure the alternate implementation approach is viable.

Recommendation 25: Periodically planned algorithm rollovers should be performed to help ensure preparedness in the event the timeline needs to be compressed to facilitate an urgent rollover.

Recommendation 26: The community should prepare to be operationally ready for a root zone algorithm rollover. The next logical time for this is following the second KSK rollover, in approximately five years.

Appendix: Design Team Roster

Community Volunteers

- Howard Eland
- Moritz Muller
- Tomofumi Okubo
- Scott Rose
- Peter Thomassen
- Ralf Weber
- Yoshiro Yoneya

Root Zone Management Partners

- Roy Arends, ICANN
- Aaron Foley, IANA
- Paul Hoffman, ICANN
- Jennifer Johnson, IANA
- Matt Larson, ICANN
- Ramana Lavu, Root Zone Maintainer
- Edward Lewis, ICANN
- James Mitchell, IANA
- Andres Pavez, IANA
- Jakob Schlyter, ICANN
- Duane Wessels, Root Zone Maintainer

Acknowledgments

The design team thanks Carsten Strotmann for his help with Windows DNS.

Appendix: Data Sources & Graphs

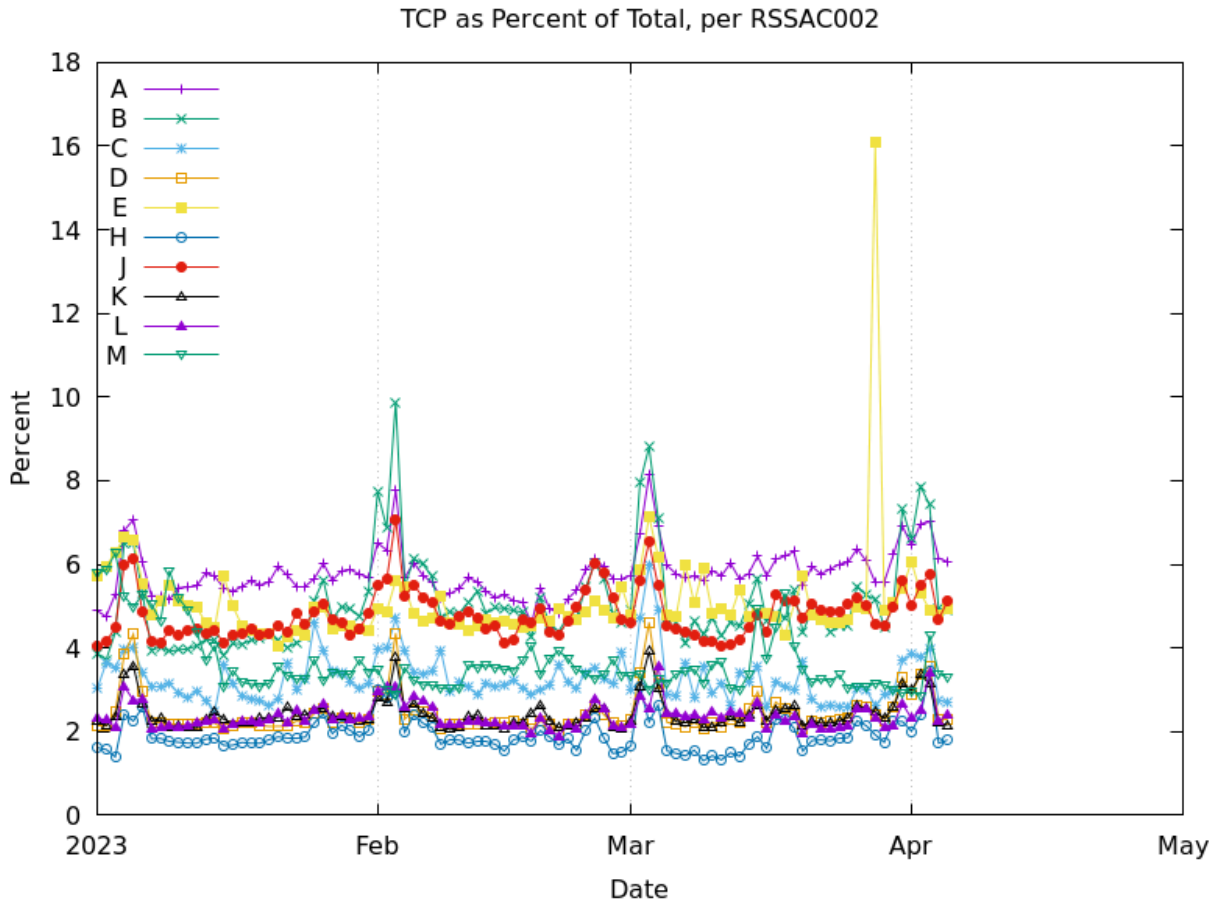

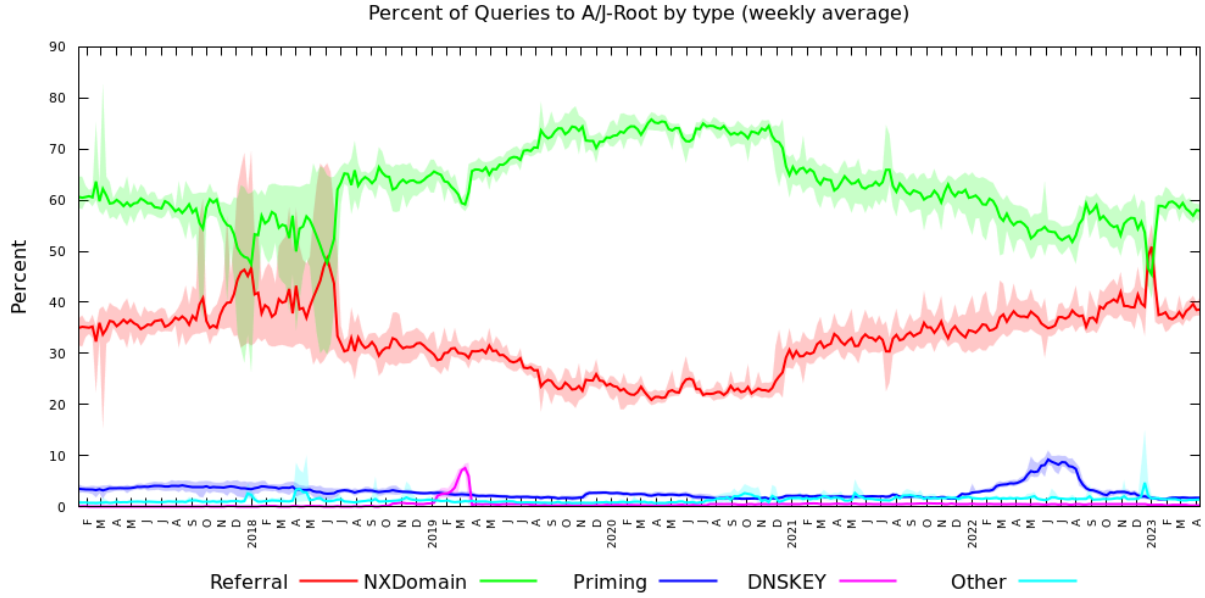
Percent of Queries to A/J-Root by type (weekly average)
Percent
Referral
NXDomain
Priming
DNSKEY
Other

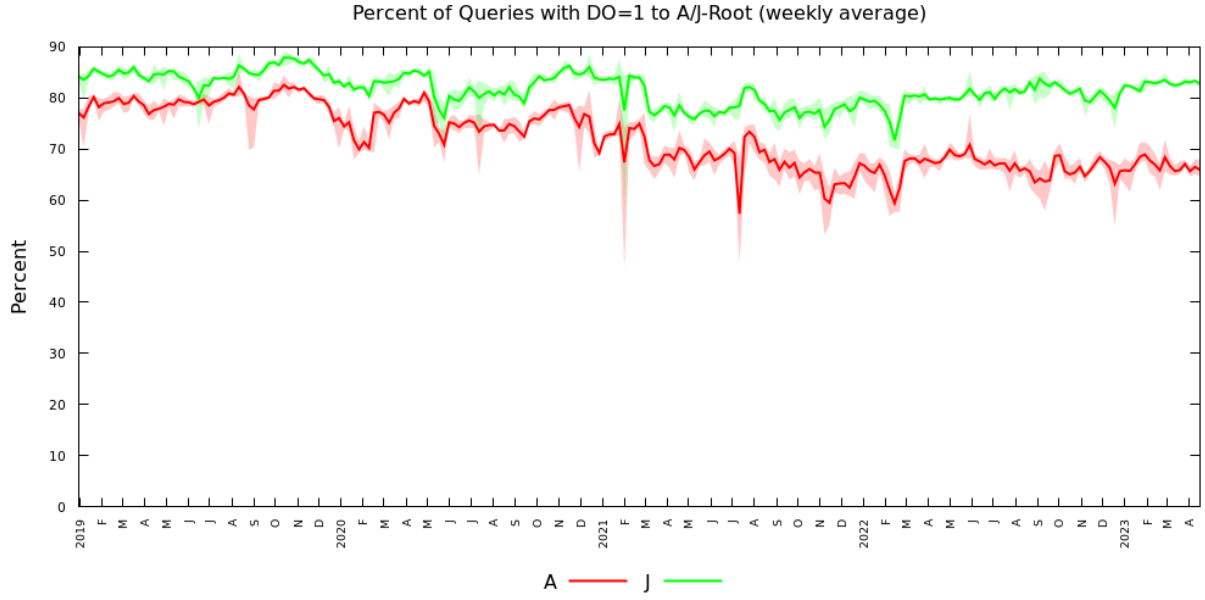
Percent of Queries with DO=1 to A/J-Root (weekly average)
Percent
A
J

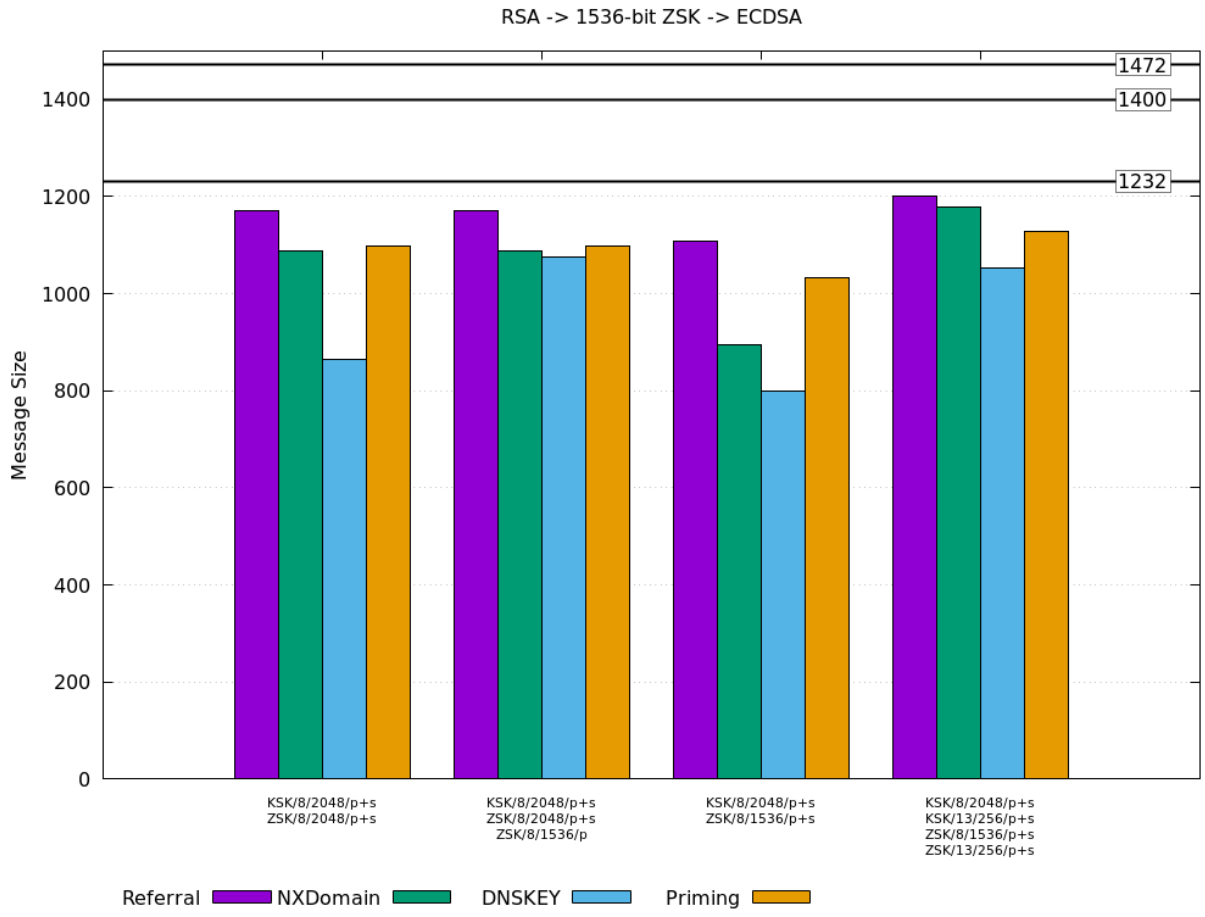
RSA -> 1536-bit ZSK -> ECDSA
Message Size
0
200
400
600
800
1000
1200
1400
1472
1400
1232
KSK/8/2048/p+s
ZSK/8/2048/p+s
KSK/8/2048/p+s
ZSK/8/2048/p+s
ZSK/8/1536/p
KSK/8/2048/p+s
ZSK/8/1536/p+s
KSK/8/2048/p+s
KSK/13/256/p+s
ZSK/8/1536/p+s
ZSK/13/256/p+s
Referral
NXDomain
DNSKEY
Priming

Appendix: Implementation Risk Analysis

The risks listed in this risk management table are primarily related to the span of time from the completion of this study until the time arrives to prepare an operational plan for changing the algorithm. Numerous assumptions may change during this time span, many involving the choice for the next DNS Security Algorithm to be used. When the operational plan is prepared, a risk management table specific to that process should be included.

Description	Impact	Likelihood	Mitigation
Algorithm roll does not start within 10 years	Analysis reported here is based on “expired” data	Possible	Recommend any operation plan rerun analysis
Protocol updates of benefit to a root algorithm roll are not accepted by community	Operational plan will have to abide by current rules	Possible	Actively work protocol updates in the IETF
Protocol updates are not sufficiently deployed in codebases prior to the need to roll	Operational plan will have to abide by current codebase restrictions	Possible	Actively support software development to stay current with protocol updates (if any)
Algorithm roll does not happen before quantum computers are able to break traditional cryptographic algorithms	Assumptions about algorithm choice will need to consider this	Unlikely	As quantum computing advances, this study will be needed again
New algorithms are standardized	Properties of these algorithms are not adequately captured by the selection criteria	Possible	Criteria for potential algorithms need to be reviewed periodically
There is an urgent need to roll the algorithm before an operational plan has been developed.	The plan may be hastily prepared, without due diligence for the sake of preserving security	Unlikely, assuming eyes are open to developments	ICANN prioritizes the work necessary to prepare themselves and the community for an algorithm

			rollover.
When the need to change the algorithm arises, there is no acceptable alternative algorithm	There would be no digital signature option	Unlikely	The problems would be beyond the DNS
When double-signing, the root zone size increase disrupts the delivery to the root server operators	File transfers would fail to impacted servers	Unlikely	Coordinate with RSOs to ensure sufficient capacity and work with protocol community to reduce load
Stringent requirements on HSMs reduce the number of available vendors to unacceptable levels	Vendor lock-in, reliance on one vendor, possibly no vendor	Possible	Review requirements to keep them "sufficient" and "necessary" while not overbearing

Appendix: Research & Data Collection

The RIPE Atlas measurement platform supports distributed, active measurements from approximately 10,000 vantage points. DNSThought is a measurement project by NLNet labs that leverages the RIPE Atlas platform. It includes several ongoing measurements relevant for root zone KSK rollovers, including resolver support for defined DNSSEC algorithms, resolver validating state, and trust anchor sentinel. DNSThought's measurement results are public. A frequent criticism of the RIPE Atlas is that the vantage points tend to be biased to the European region, which is the focus of RIPE's operation and involvement.

In 2017 the IETF published RFC 8145, "Signaling Trust Anchor Knowledge in DNS Security Extensions (DNSSEC)." It defines a method by which DNSSEC validators signal to zone operators the trust anchors that they have configured. These passive "key tag signals" are delivered as DNS queries to the trust anchor's corresponding zone. For root zone trust anchors, therefore, these signals are received by the root zone operators. Key tag signaling data played a role in postponing the scheduled 2017 KSK rollover. Criticisms of RFC 8145 include: that the data can be difficult to access and interpret, that only some DNSSEC validators provided the signals, and that it does not convey how many users are relying on the validator and whether they rely on that validator alone.

In 2018, the IETF published RFC 8509, "A Root Key Trust Anchor Sentinel for DNSSEC." A validating recursive resolver that implements this protocol will respond in specific ways to specific queries referencing trust anchor key tags. Responses from these special queries indicate whether or not the resolver is configured with the corresponding root zone trust anchors. Unlike RFC 8145 passive measurements, RFC 8509 measurements must be actively collected. The DNSThought measurement suite referenced above includes an ongoing measurement for root zone KSK key tags 19036 (2010-2018) and 20326 (2018–). The technique is also suitable for measurement via advertisement-delivery platforms.

For the 2017/2018 KSK rollover, root server operators used dnscap and its "rzkeychange" plugin. This plugin, specifically designed for collecting data during root zone DNSKEY changes, counts:

- Total number of queries
- Number of DNSKEY queries
- Number of queries over TCP
- Number of responses that set the TC flag
- Number of ICMP_UNREACH_NEEDFRAG messages
- Number of ICMP_TIMXCEED_INTRANS messages
- Number of ICMP_TIMXCEED_REASS messages

Counts are reported to a central location every 60 seconds for near-real-time monitoring.

DNS-OARC often collects packet capture data in the style of its Day In The Life (DITL) collections. Due to the nature of this data collection, i.e., that it takes a long time, it is primarily useful in post-hoc analyses.

The root server operators publish RSSAC002 metrics on a daily basis. The data is of limited use for root zone KSK rollovers due to its daily aggregation and lack of any DNSSEC-specific metrics.

One World, One Internet

Visit us at **icann.org**

@icann

facebook.com/icannorg

youtube.com/icannnews

flickr.com/icann

linkedin.com/company/icann

soundcloud.com/icann

instagram.com/icannorg