

dnsop  
Internet-Draft  
Updates: [7583](#) (if approved)  
Intended status: Standards Track  
Expires: April 19, 2018

W. Hardaker  
USC/ISI  
W. Kumari  
Google  
October 16, 2017

**Security Considerations for [RFC5011](#) Publishers**  
**draft-ietf-dnsop-rfc5011-security-considerations-06**

Abstract

This document extends the [RFC5011](#) rollover strategy with timing advice that must be followed in order to maintain security. Specifically, this document describes the math behind the minimum time-length that a DNS zone publisher must wait before signing exclusively with recently added DNSKEYs. It contains much math and complicated equations, but the summary is that the key rollover / revocation time is much longer than intuition would suggest. If you are not both publishing a DNSSEC trust anchor, and using [RFC5011](#) to update that trust anchor, you probably don't need to read this document.

This document also describes the minimum time-length that a DNS zone publisher must wait after publishing a revoked DNSKEY before assuming that all active [RFC5011](#) resolvers should have seen the revocation-marked key and removed it from their list of trust anchors.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 19, 2018.

## Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

<a href="#">1.</a>	Introduction . . . . .	<a href="#">3</a>
<a href="#">1.1.</a>	Document History and Motivation . . . . .	<a href="#">3</a>
<a href="#">1.2.</a>	Safely Rolling the Root Zone's KSK in 2017/2018 . . . . .	<a href="#">3</a>
<a href="#">1.3.</a>	Requirements notation . . . . .	<a href="#">3</a>
<a href="#">2.</a>	Background . . . . .	<a href="#">3</a>
<a href="#">3.</a>	Terminology . . . . .	<a href="#">4</a>
<a href="#">4.</a>	Timing Associated with <a href="#">RFC5011</a> Processing . . . . .	<a href="#">5</a>
<a href="#">4.1.</a>	Timing Associated with Publication . . . . .	<a href="#">5</a>
<a href="#">4.2.</a>	Timing Associated with Revocation . . . . .	<a href="#">5</a>
<a href="#">5.</a>	Denial of Service Attack Considerations . . . . .	<a href="#">6</a>
<a href="#">5.1.</a>	Enumerated Attack Example . . . . .	<a href="#">6</a>
<a href="#">5.1.1.</a>	Attack Timing Breakdown . . . . .	<a href="#">7</a>
<a href="#">6.</a>	Minimum <a href="#">RFC5011</a> Timing Requirements . . . . .	<a href="#">8</a>
<a href="#">6.1.</a>	Timing Requirements For Adding a New KSK . . . . .	<a href="#">8</a>
<a href="#">6.1.1.</a>	addHoldDownTime . . . . .	<a href="#">8</a>
<a href="#">6.1.2.</a>	sigExpirationTime . . . . .	<a href="#">9</a>
<a href="#">6.1.3.</a>	activeRefresh . . . . .	<a href="#">9</a>
<a href="#">6.1.4.</a>	activeRefreshOffset . . . . .	<a href="#">9</a>
<a href="#">6.1.5.</a>	safetyMargin . . . . .	<a href="#">9</a>
<a href="#">6.1.6.</a>	Fully expanded equation . . . . .	<a href="#">10</a>
<a href="#">6.1.7.</a>	Timing Constraint Summary . . . . .	<a href="#">10</a>
<a href="#">6.1.8.</a>	Additional Considerations . . . . .	<a href="#">11</a>
<a href="#">6.2.</a>	Timing Requirements For Revoking an Old KSK . . . . .	<a href="#">11</a>
<a href="#">6.2.1.</a>	Example Results . . . . .	<a href="#">12</a>
<a href="#">7.</a>	IANA Considerations . . . . .	<a href="#">12</a>
<a href="#">8.</a>	Operational Considerations . . . . .	<a href="#">12</a>
<a href="#">9.</a>	Security Considerations . . . . .	<a href="#">13</a>
<a href="#">10.</a>	Acknowledgements . . . . .	<a href="#">13</a>
<a href="#">11.</a>	Normative References . . . . .	<a href="#">13</a>
<a href="#">Appendix A.</a>	Real World Example: The 2017 Root KSK Key Roll . . . . .	<a href="#">14</a>
	Authors' Addresses . . . . .	<a href="#">14</a>



## **1. Introduction**

[RFC5011] defines a mechanism by which DNSSEC validators can update their list of trust anchors when they've seen a new key published in a zone. However, [RFC5011](#) [intentionally] provides no guidance to the publishers of DNSKEYs about how long they must wait before switching to exclusively using recently published keys for signing records, or how long they must wait before ceasing publication of a revoked key. Because of this lack of guidance, zone publishers may derive incorrect assumptions about safe usage of the [RFC5011](#) DNSKEY advertising, rolling and revocation process. This document describes the minimum security requirements from a publisher's point of view and is intended to complement the guidance offered in [RFC5011](#) (which is written to provide timing guidance solely to a Validating Resolver's point of view).

### **1.1. Document History and Motivation**

To verify this lack of understanding is wide-spread, the authors reached out to 5 DNSSEC experts to ask them how long they thought they must wait before signing a zone exclusively with a new KSK [[RFC4033](#)] that was being introduced according to the 5011 process. All 5 experts answered with an insecure value, and we determined that this lack of operational guidance is causing security concerns today and wrote this companion document to [RFC5011](#). We hope that this document will rectify this understanding and provide better guidance to zone publishers that wish to make use of the [RFC5011](#) rollover process.

### **1.2. Safely Rolling the Root Zone's KSK in 2017/2018**

One important note about ICANN's [currently upcoming] 2017/2018 KSK rollover plan for the root zone: the timing values chosen for rolling the KSK in the root zone appear completely safe, and are not affected by the timing concerns introduced by this draft

### **1.3. Requirements notation**

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

## **2. Background**

The [RFC5011](#) process describes a process by which a [RFC5011](#) Validating Resolver may accept a newly published KSK as a trust anchor for validating future DNSSEC signed records. It also describes the process for publicly revoking a published KSK. This document



augments that information with additional constraints, from the DNSKEY publication and revocation's points of view. Note that this document does not define any other operational guidance or recommendations about the [RFC5011](#) process and restricts itself to solely the security and operational ramifications of switching to exclusively using recently added keys or removing a revoked keys too soon.

Failure of a DNSKEY publisher to follow the minimum recommendations associated with this draft will result in potential denial-of-service attack opportunities against validating resolvers. Failure of a DNSKEY publisher to publish a revoked key for a long enough period of time may result in [RFC5011](#) Validating Resolvers leaving that key in their trust anchor storage beyond the key's expected lifetime.

### **3. Terminology**

**Trust Anchor Publisher** The entity responsible for publishing a DNSKEY that can be used as a trust anchor.

**Zone Signer** The owner of a zone intending to publish a new Key-Signing-Key (KSK) that will become a trust anchor by validators following the [RFC5011](#) process.

**[RFC5011](#) Validating Resolver** A DNSSEC Validating Resolver that is using the [RFC5011](#) processes to track and update trust anchors. Sometimes referred to as a "[RFC5011](#) Resolver"

**Attacker** An entity intent on foiling the [RFC5011](#) Validator's ability to successfully adopt the Zone Signer's new DNSKEY as a new trust anchor or to prevent the [RFC5011](#) Validator from removing an old DNSKEY from its list of trust anchors.

**sigExpirationTime** The amount of time remaining before any existing RRSIG's Signature Expiration time is reached. Note that for organizations pre-creating signatures this time may be fairly lengthy unless they can be significantly assured their signatures can not be replayed at a later date. sigExpirationTime will fundamentally be the RRSIG's Signature Expiration time minus the RRSIG's Signature Inception time when the signature is created.

Also see [Section 2 of \[RFC4033\]](#) and [\[RFC7719\]](#) for additional terminology.



#### **4. Timing Associated with [RFC5011](#) Processing**

These sections define a high-level overview of [[RFC5011](#)] processing. These steps are not sufficient for proper [RFC5011](#) implementation, but provide enough background for the reader to follow the discussion in this document. Readers need to fully understand [[RFC5011](#)] as well to fully comprehend the importance of this document.

##### **4.1. Timing Associated with Publication**

[RFC5011](#)'s process of safely publishing a new key and then making use of that key falls into a number of high-level steps to be performed by the Trust Anchor Publisher. This document discusses the following scenario, which the principle way [RFC5011](#) is currently being used (even though [Section 6 of RFC5011](#) suggests having a stand-by key available):

1. Publish a new DNSKEY in the zone, but continue to sign the zone with the old one.
2. Wait a period of time.
3. Begin to exclusively use recently published DNSKEYs to sign the appropriate resource records.

This document discusses step 2 of the above process. Some interpretations of [RFC5011](#) have erroneously determined that the wait time is equal to [RFC5011](#)'s "hold down time". [Section 5](#) describes an attack based on this (common) erroneous belief, which can result in a denial of service attack against the zone.

##### **4.2. Timing Associated with Revocation**

[RFC5011](#)'s process of advertising that an old key is to be revoked from [RFC5011](#) validating resolvers falls into a number of high-level steps:

1. Set the revoke bit on the DNSKEY to be revoked.
2. Sign the revoked DNSKEY with itself.
3. Wait a period of time.
4. Remove the revoked key from the zone.

This document discusses step 3 of the above process. Some interpretations of [RFC5011](#) have erroneously determined that the wait time is equal to [RFC5011](#)'s "hold down time". This document describes





an attack based on this (common) erroneous belief, which results in a revoked DNSKEY potentially remaining as a trust anchor in a [RFC5011](#) validating resolver long past its expected usage.

## 5. Denial of Service Attack Considerations

If an attacker is able to provide a [RFC5011](#) Validating Resolver with past responses, such as when it is in-path or able to perform any number of cache poisoning attacks, the attacker may be able to leave compliant [RFC5011](#)-Validating Resolvers without an appropriate DNSKEY trust anchor. This scenario will remain until an administrator manually fixes the situation.

The time-line below illustrates this situation.

### 5.1. Enumerated Attack Example

The following example settings are used in the example scenario within this section:

TTL (all records) 1 day

sigExpirationTime 10 days

Zone resigned every 1 day

Given these settings, the sequence of events in [Section 5.1.1](#) depicts how a Trust Anchor Publisher that waits for only the [RFC5011](#) hold time timer length of 30 days subjects its users to a potential Denial of Service attack. The timing schedule listed below is based on a Trust Anchor Publisher publishing a new Key Signing Key (KSK), with the intent that it will later become a trust anchor. We label this publication time as "T+0". All numbers in this sequence refer to days before and after this initial publication event. Thus, T-1 is the day before the introduction of the new key, and T+15 is the 15th day after the key was introduced into the fictitious zone being discussed.

In this dialog, we consider two keys within the example zone:

K\_old An older KSK and Trust Anchor being replaced.

K\_new A new KSK being transitioned into active use and expected to become a Trust Anchor via the [RFC5011](#) process.



### **5.1.1. Attack Timing Breakdown**

The steps shows an attack that foils the adoption of a new DNSKEY by a 5011 Validating Resolver when the Trust Anchor Publisher that starts signing and publishing with the new DNSKEY too quickly.

T-1 The K\_old based RRSIGs are being published by the Zone Signer. [It may also be signing ZSKs as well, but they are not relevant to this event so we will not talk further about them; we are only considering the RRSIGs that cover the DNSKEYs in this document.] The Attacker queries for, retrieves and caches this DNSKEY set and corresponding RRSIG signatures. Note that for simplicity we assume the signer is not pre-signing and creating "valid in the future" signature sets that may be stolen and replayed even later.

T+0 The Zone Signer adds K\_new to their zone and signs the zone's key set with K\_old. The [RFC5011](#) Validator (later to be under attack) retrieves this new key set and corresponding RRSIGs and notices the publication of K\_new. The [RFC5011](#) Validator starts the (30-day) hold-down timer for K\_new. [Note that in a more real-world scenario there will likely be a further delay between the point where the Zone Signer publishes a new RRSIG and the [RFC5011](#) Validator notices its publication; though not shown in this example, this delay is accounted for in the final solution below]

T+5 The [RFC5011](#) Validator queries for the zone's keyset per the [RFC5011](#) Active Refresh schedule, discussed in [Section 2.3 of RFC5011](#). Instead of receiving the intended published keyset, the Attacker successfully replays the keyset and associated signatures recorded at T-1. Because the signature lifetime is 10 days (in this example), the replayed signature and keyset is accepted as valid (being only 6 days old, which is less than sigExpirationTime) and the [RFC5011](#) Validator cancels the hold-down timer for K\_new, per the [RFC5011](#) algorithm.

T+10 The [RFC5011](#) Validator queries for the zone's keyset and discovers a signed keyset that includes K\_new (again), and is signed by K\_old. Note: the attacker is unable to replay the records cached at T-1, because they have now expired. Thus at T+10, the [RFC5011](#) Validator starts (anew) the hold-timer for K\_new.

T+11 through T+29 The [RFC5011](#) Validator continues checking the zone's key set at the prescribed regular intervals. During this period, the attacker can no longer replay traffic to their benefit.



T+30 The Zone Signer knows that this is the first time at which some validators might accept K\_new as a new trust anchor, since the hold-down timer of a [RFC5011](#) Validator not under attack that had queried and retrieved K\_new at T+0 would now have reached 30 days. However, the hold-down timer of our attacked [RFC5011](#) Validator is only at 20 days.

T+35 The Zone Signer (mistakenly) believes that all validators following the Active Refresh schedule ([Section 2.3 of RFC5011](#)) should have accepted K\_new as a the new trust anchor (since the hold down time (30 days) + the query interval [which is just 1/2 the signature validity period in this example] would have passed). However, the hold-down timer of our attacked [RFC5011](#) Validator is only at 25 days (T+35 minus T+10); thus the [RFC5011](#) won't consider it a valid trust anchor addition yet, as the required 30 days have not yet elapsed.

T+36 The Zone Signer, believing K\_new is safe to use, switches their active signing KSK to K\_new and publishes a new RRSIG, signed with K\_new, covering the DNSKEY set. Non-attacked [RFC5011](#) validators, with a hold-down timer of at least 30 days, would have accepted K\_new into their set of trusted keys. But, because our attacked [RFC5011](#) Validator now has a hold-down timer for K\_new of only 26 days, it failed to accept K\_new as a trust anchor. Since K\_old is no longer being used to sign the zone's DNSKEYs, all the DNSKEY records from the zone will be treated as invalid. Subsequently, all of the records in the DNS tree below the zone's apex will be deemed invalid by DNSSEC.

## **6. Minimum [RFC5011](#) Timing Requirements**

### **6.1. Timing Requirements For Adding a New KSK**

Given the attack description in [Section 5](#), the correct minimum length of time required for the Zone Signer to wait after publishing K\_new but before exclusively using it and newer keys is:

```
addWaitTime = addHoldDownTime
              + sigExpirationTime
              + activeRefresh
              + activeRefreshOffset
              + safetyMargin
```

#### **6.1.1. addHoldDownTime**

The addHoldDownTime is defined in [Section 2.4.1 of \[RFC5011\]](#) as:



The add hold-down time is 30 days or the expiration time of the original TTL of the first trust point DNSKEY RRSset that contained the new key, whichever is greater. This ensures that at least two validated DNSKEY RRSets that contain the new key MUST be seen by the resolver prior to the key's acceptance.

#### **[6.1.2.](#) sigExpirationTime**

sigExpirationTime is defined in [Section 3](#).

#### **[6.1.3.](#) activeRefresh**

activeRefresh time is defined by [RFC5011](#) by

A resolver that has been configured for an automatic update of keys from a particular trust point MUST query that trust point (e.g., do a lookup for the DNSKEY RRSset and related RRSIG records) no less often than the lesser of 15 days, half the original TTL for the DNSKEY RRSset, or half the RRSIG expiration interval and no more often than once per hour.

This translates to:

```
activeRefresh = MAX(1 hour,  
                    MIN(sigExpirationTime / 2,  
                        MAX(TTL of K_old DNSKEY RRSset) / 2,  
                        15 days)  
                    )
```

#### **[6.1.4.](#) activeRefreshOffset**

The activeRefreshOffset term must be added for situations where the activeRefresh value is not a factor of the addHoldDownTime. Specifically, activeRefreshOffset will be "addHoldDownTime % activeRefresh", where % is the mathematical mod operator (calculating the remainder in a division problem). This will frequently be zero, but could be nearly as large as activeRefresh itself. For simplicity, setting the activeRefreshOffset to the activeRefresh value itself is always safe.

#### **[6.1.5.](#) safetyMargin**

The safetyMargin is an extra period of time to account for caching, network delays, etc. A suggested operational value for this is 2 \* MAX(TTL of all records) unless the TTLs are shorter than an hour, at which point they start affecting the calculations in the MIN() clause





in the activeRefresh timer in [Section 6.1.3](#). Thus, we suggest a safetyMargin of at least:

$$\text{safetyMargin} = \text{MAX} (1.5 \text{ hours}, 2 * \text{MAX}(\text{TTL of all records}))$$

[RFC5011](#) also discusses a retryTime value for failed queries. Our equation cannot take into account undeterministic failure situations, so it might be wise to extend the safetyMargin by some factor of retryTime, which is defined in [RFC5011](#) as:

$$\begin{aligned} \text{retryTime} = \text{MAX} (1 \text{ hour}, \\ \text{MIN} (1 \text{ day}, \\ .1 * \text{TTL of K\_old DNSKEY RRset}, \\ .1 * \text{sigExpirationTime})) \end{aligned}$$

#### [6.1.6](#). Fully expanded equation

The full expanded equation, with activeRefreshOffset set to activeRefresh for simplicity, is:

$$\begin{aligned} \text{addWaitTime} = & \text{addHoldDownTime} \\ & + \text{sigExpirationTime} \\ & + 2 * \text{MAX}(1 \text{ hour}, \\ & \quad \text{MIN}(\text{sigExpirationTime} / 2, \\ & \quad \quad \text{MAX}(\text{TTL of K\_old DNSKEY RRSet}) / 2, \\ & \quad \quad 15 \text{ days}) \\ & ) \\ & + (\text{addHoldDownTime} \% \text{activeRefresh}) \\ & + \text{MAX}(1.5 \text{ hours}, 2 * \text{MAX}(\text{TTL of all records})) \end{aligned}$$

#### [6.1.7](#). Timing Constraint Summary

The important timing constraint introduced by this memo relates to the last point at which a validating resolver may have received a replayed original DNSKEY set, containing K\_old and not K\_new. The next query of the [RFC5011](#) validator at which K\_new will be seen without the potential for a replay attack will occur after the publication time plus sigExpirationTime. Thus, the latest time that a [RFC5011](#) Validator may begin their hold down timer is an "Active Refresh" period after the last point that an attacker can replay the K\_old DNSKEY set. The worst case scenario of this attack is if the attacker can replay K\_old seconds before the (DNSKEY RRSIG Signature Validity) field of the last K\_old only RRSIG.



### **[6.1.8.](#) Additional Considerations**

Note: our notion of addWaitTime is called "Itrp" in [Section 3.3.4.1 of \[RFC7583\]](#). The equation for Itrp in [RFC7583](#) is insecure as it does not include the sigExpirationTime listed above. The Itrp equation in [RFC7583](#) also does not include the 2\*TTL safety margin, though that is an operational consideration and not necessarily as critical.

#### **[6.1.8.1.](#) Example Results**

For the parameters listed in [Section 5.1](#), the activeRefreshOffset is 0, since 30 days is evenly divisible by activeRefresh (1/2 day), and our resulting addWaitTime is:

```
addWaitTime = 30
              + 10
              + 1 / 2
              + 2 * (1)          (days)

addWaitTime = 42.5              (days)
```

This addWaitTime of 42.5 days is 12.5 days longer than just the hold down timer.

### **[6.2.](#) Timing Requirements For Revoking an Old KSK**

It is important to note that this issue affects not just the publication of new DNSKEYs intended to be used as trust anchors, but also the length of time required to continuously publish a DNSKEY with the revoke bit set. Both of these publication timing requirements are affected by the attacks described in this document, but with revocation the key is revoked immediately and the addHoldDown timer does not apply. Thus the minimum amount of time that a Trust Anchor Publisher must wait before removing a revoked key from publication is:

```
remWaitTime = sigExpirationTime
              + MAX(1 hour,
                    MIN((sigExpirationTime) / 2,
                        MAX(TTL of K_old DNSKEY RRSets) / 2,
                          15 days),
                    1 hour)
              + 2 * MAX(TTL of all records)
```

Note that the activeRefreshOffset time does not apply to this equation.



Note that our notion of `remWaitTime` is called "Irev" in [Section 3.3.4.2 of \[RFC7583\]](#). The equation for Irev in [RFC7583](#) is insecure as it does not include the `sigExpirationTime` listed above. The Irev equation in [RFC7583](#) also does not include the `2*TTL` safety margin, though that is an operational consideration and not necessarily as critical.

Note also that adding `retryTime` intervals to the `remWaitTime` may be wise, just as it was for `addWaitTime` in [Section 6](#).

### [6.2.1](#). Example Results

For the parameters listed in [Section 5.1](#), our example:

```
remwaitTime = 10
               + 1 / 2
               + 2 * (1)          (days)

remwaitTime = 12.5              (days)
```

Note that for the values in this example produce a length shorter than the recommended 30 days in [RFC5011](#)'s [section 6.6](#), step 3. Other values of `sigExpirationTime` and the original TTL of the `K_old` DNSKEY RRSet, however, can produce values longer than 30 days.

Note that because revocation happens immediately, an attacker has a much harder job tricking a [RFC5011](#) Validator into leaving a trust anchor in place, as the attacker must successfully replay the old data for every query a [RFC5011](#) Validator sends, not just one.

## [7](#). IANA Considerations

This document contains no IANA considerations.

## [8](#). Operational Considerations

A companion document to [RFC5011](#) was expected to be published that describes the best operational practice considerations from the perspective of a zone publisher and Trust Anchor Publisher. However, this companion document has yet to be published. The authors of this document hope that it will at some point in the future, as [RFC5011](#) timing can be tricky as we have shown, and a BCP is clearly warranted. This document is intended only to fill a single operational void which, when left misunderstood, can result in serious security ramifications. This document does not attempt to document any other missing operational guidance for zone publishers.



## **9. Security Considerations**

This document, is solely about the security considerations with respect to the Trust Anchor Publisher of [RFC5011](#) trust anchors / DNSKEYs. Thus the entire document is a discussion of Security Considerations when adding or removing DNSKEYs from trust anchor storage using the [RFC5011](#) process.

For simplicity, this document assumes that the Trust Anchor Publisher will use a consistent RRSIG validity period. Trust Anchor Publishers that vary the length of RRSIG validity periods will need to adjust the sigExpirationTime value accordingly so that the equations in [Section 6](#) and [Section 6.2](#) use a value that coincides with the last time a replay of older RRSIGs will no longer succeed.

## **10. Acknowledgements**

The authors would like to especially thank to Michael StJohns for his help and advice and the care and thought he put into [RFC5011](#) itself. We would also like to thank Bob Harold, Shane Kerr, Matthijs Mekking, Duane Wessels, Petr Petr Spacek, Ed Lewis, and the dnsop working group who have assisted with this document.

## **11. Normative References**

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", [RFC 4033](#), DOI 10.17487/RFC4033, March 2005, <<http://www.rfc-editor.org/info/rfc4033>>.
- [RFC5011] StJohns, M., "Automated Updates of DNS Security (DNSSEC) Trust Anchors", STD 74, [RFC 5011](#), DOI 10.17487/RFC5011, September 2007, <<http://www.rfc-editor.org/info/rfc5011>>.
- [RFC7583] Morris, S., Ihren, J., Dickinson, J., and W. Mekking, "DNSSEC Key Rollover Timing Considerations", [RFC 7583](#), DOI 10.17487/RFC7583, October 2015, <<https://www.rfc-editor.org/info/rfc7583>>.
- [RFC7719] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", [RFC 7719](#), DOI 10.17487/RFC7719, December 2015, <<http://www.rfc-editor.org/info/rfc7719>>.





## [Appendix A](#). Real World Example: The 2017 Root KSK Key Roll

In 2017, ICANN expects to (or has, depending on when you're reading this) roll the key signing key (KSK) for the root zone. The relevant parameters associated with the root zone at the time of this writing is as follows:

addHoldDownTime:	30 days
Old DNSKEY sigExpirationTime:	21 days
Old DNSKEY TTL:	2 days

Thus, sticking this information into the equation in [Section 6](#) yields (in days):

```
addWaitTime = 30
              + (21)
              + MAX(MIN((21) / 2,
                        MAX(2 / 2,
                          15 days)),
                    1 hour)
              + 2 * MAX(2)

addWaitTime = 30 + 21 + MAX(MIN(11.5, 1, 15)), 1 hour) + 4

addWaitTime = 30 + 21 + 1 + 4

addWaitTime = 56 days
```

Note that we use a activeRefreshOffset of 0, since 30 days is evenly divisible by activeRefresh (1 day).

Thus, ICANN should wait a minimum of 56 days before switching to the newly published KSK (and 26 days before removing the old revoked key once it is published as revoked). ICANN's current plans are to wait 70 days before using the new KEY and 69 days before removing the old, revoked key. Thus, their current rollover plans are sufficiently secure from the attack discussed in this memo.

### Authors' Addresses

Wes Hardaker  
USC/ISI  
P.O. Box 382  
Davis, CA 95617  
US

Email: [ietf@hardakers.net](mailto:ietf@hardakers.net)



Warren Kumari  
Google  
1600 Amphitheatre Parkway  
Mountain View, CA 94043  
US

Email: [warren@kumari.net](mailto:warren@kumari.net)