Domain Name System Operations (dnsop)
Internet-Draft
Intended status: Standards Track
Expires: August 17, 2014

T. Finch
University of Cambridge

February 13, 2014

The WS resource record: dispersing trust in the DNSSEC root keys
draft-fanf-dnsop-trust-anchor-witnesses-00

Abstract

At the moment the root DNSSEC key is a single point of trust and a single point of failure for the whole system. This memo describes a mechanism for dispersing trust in the root key. Witnesses vouch for the root trust anchor by publishing WS records in the DNS. Validators only update their root trust anchors if multiple witnesses agree. The root-witnesses.arpa zone enables a validator to bootstrap trust when it has no working trust anchors other than its witnesses.

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 August 17, 2014.

Copyright Notice

Table of Contents

1. Introduction

At the moment the root DNSSEC key is a single point of trust and a single point of failure for the whole system. It has a number of problems:

- Root trust anchor rollovers using [RFC5011] require validators to be online while the rollover happens. With the current root key management plan, rollovers take a few weeks. This is uncomfortably long for emergency rollovers.

- Systems that are offline during a rollover have to use an out-of-band mechanism to update their trust anchors, relying on non-DNS sources of trust. There is no clear specification or security analysis for this process.

- The root key is a single point of failure with no standby, though its storage and management is extremely resilient and trustworthy (in stark contrast to the out-of-band trust anchor update keys).

- The concentration of trust in the root is politically uncomfortable.

This memo describes a mechanism for dispersing trust in the root key. Witnesses vouch for the root trust anchor by publishing WS records in the DNS. Validators only update their root trust anchors if multiple witnesses agree.

This mechanism has the following advantages:

- There is no single point of failure since there are many witnesses not one of which is completely trusted.

- There are no special timing constraints as in [RFC5011]. Witnessed trust anchor updates are like normal KSK rollovers.

- The mechanism is in-band, using only DNS, even for bootstrapping.
- The same procedure works for online and offline key rollovers.
- Rollovers can become routine.

There are some potential advantages:

- It can allow for a crash rollover of the root key, in the event that it is lost or compromised, with validators recovering automatically rather than having to be manually forced to fetch and authenticate the replacement trust anchor.

- It could allow a smaller root DNSKEY RRset by allowing the witnesses to vouch for the root ZSK directly instead of via a KSK. This saves the cost of high-assurance storage for the root KSK, but requires more frequent communication between the root DNSSEC key managers and the witnesses.

There are some limitations and disadvantages:

- It does not disperse trust in the root zone signing key or root zone maintenance.
- A lot more co-ordination between organizations is necessary, for the witnesses to get out-of-band authentication of new trust anchors.

This mechanism can be used to automatically update any trust anchor, though it is designed for and includes some special considerations for the root trust anchor. The root-witnesses.arpa zone is set up to enable a validator to bootstrap trust when it has no working trust anchors other than its witnesses.

1.1. Terminology

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. The WS resource record

The WS (witness signer) resource record contains a cryptographic digest of a DNSKEY record at a DNSSEC trust anchor. WS records are published separately from the trust anchor by trust anchor witnesses to show that the witnesses vouch for the trust anchor DNSKEY(s). WS records are used by validators to update their trust anchors.

The WS resource record RDATA has exactly the same wire format and presentation format as the DS RR, as described in [RFC4034] section 5. Although WS has the same syntax as DS, its semantics differ as described below.

The WS resource record type number is [TBD]

The WS RR is treated as a normal RR for signing, serving, resolving, and validating. None of the special behaviour for DS records described in [RFC4035] sections 2.4, 2.6, 3.1 applies to WS records.

Unlike DS records (where the DS record on the parent side of a zone

cut refers to a DNSKEY record on the child side of the zone cut at the same name), WS records do not explicitly indicate the name of the trust anchor DNSKEY records that they refer to. This information is part of the valiator configuration (see Section 3.1).

WS records SHOULD be placed at the apex of a witness zone. They SHOULD refer to DNSKEY records with the SEP (secure entry point) flag set [RFC3757].

3. How validators use WS records

3.1. Trust anchor configuration

A validator's configuration for a trust anchor consists of the the trust anchor owner name, and either a set of public keys or a set of DS records, as described in [RFC4035] section 4.4.

A trust anchor that is automatically updated is associated with the witnesses that vouch for it. It has a quorum value stating how many witnesses must agree before the trust anchor is updated.

Each witness is a normal statically-configured trust anchor. That is, witnesses are not updated automatically except by out-of-band configuration updates or software updates. Each witness is associated with one automatically updated trust anchor for which it vouches.

3.2. When to try a trust anchor update

The validator SHALL keep track of the DNSKEY records from the DNS at the trust anchor name. It only tracks the set of all records with the SEP flag set, or the subset of SEP keys with algorithms supported by the validator. (This is so that ZSK rollovers do not trigger trust anchor updates.)

When the validator notices that this set has changed it SHOULD attempt to update the trust anchor as described below. During the update process it SHOULD continue to serve clients and use the existing trust anchor to validate responses.

The validator MAY track the DNSKEY records persistently in order to make restarts faster. If so, it SHOULD discard any saved DNSKEY records after their RRSIG expiry time. If it does not, it SHOULD perform an update attempt at restart.

When starting, a validator can find that its existing trust anchor does not work, perhaps because a key rollover happened while it was

offline. In this situation it cannot serve clients until the update process completes successfully.

A broken trust anchor is not expected to happen during normal operations, since validation ought to work at every point in a key rollover. However, if some disaster occurs and the trust anchor private key is lost or compromised, there might be a disruptive crash key rollover.

When it sees a crash rollover, a validator will not be able to validate the new DNSKEY RRset, so will discard it and retry the query in an attempt to obtain a working version. If this problem persists the validator MAY attempt to update the trust anchor using an invalid DNSKEY response.

3.3. Trust anchor update process

Trust anchor updates are performed with respect to a DNSKEY RRset from the trust anchor owner name. This allows the validator to ensure that a successful update will lead to a working configuration.

The validator queries for the WS RRset at each of the trust anchor's witnesses. The witnesses SHOULD be queried in a random order, so that the validator avoids relying too much on a subset of the witnesses. The query process SHOULD stop when a quorum has been achieved for one or more WS RRs. The queries MAY be performed concurrently to improve performance (though it doesn't make sense to use a level of concurrency greater than the quorum size).

The witness queries follow normal DNS resolution and DNSSEC validation rules. The response from a witness MUST validate as secure using that witness's trust anchor. (The special arrangements for the root trust anchor witnesses described in Section 5.1 ensure that the requirements in this paragraph can be satisfied even when the root trust anchor is broken.)

The validator MUST ensure there are no duplicate WS RRs in the response from a witness. Duplicate RRs are not allowed (see [RFC2181] section 5), but it is particularly important to prevent duplicate WS RRs so that a witness cannot count more than once towards a quorum.

The validator SHOULD ignore a WS RR if it does not contain a valid digest of a DNSKEY record with the SEP flag set. This ensures that the validator does not count a quorum of useless WS RRs.

For each usable WS RR that the validator receives from a witness, it keeps a count of the number of responses that contained that WS RR.

A WS RR can be trusted when this count reaches the required quorum.

If the trust anchor that is being updated is configured with DS RRs, then the validator converts the trusted WS RRs into DS RRs by changing their RR TYPE fields and uses those for the new configuration. If the trust anchor is configured with public keys, then the new keys are taken from the DNSKEY RRs that are authenticated by the trusted WS RRs.

4. How witnesses publish WS records

The administrative arrangements for publishing WS records in a witness zone are analogous to publishing DS records in a parent zone.

There MUST be an out-of-band (non-DNS) communications channel between the witnesses and the owner of the zone for which they vouch. This is used to authenticate WS RRset changes.

The timing of trust anchor rollovers is the same as for KSK rollovers [RFC6781], except that instead of updating parental DS records, withess WS records must be updated.

4.1. Contents of a witness zone

- A SOA record.
- NS records and name server address records. The name server names SHOULD be in the witness zone. (Section 5.1 below explains this requirement.)
- A DNSKEY RRset. This SHALL contain exactly one record with the SEP flag set, corresponding to the witness trust anchor. The DNSKEY RRset SHALL be signed by this key. As usual, the DNSKEY RRset SHOULD contain other keys which are used as zone-signing keys.
- A WS RRset. There MAY be multiple WS records to allow for multiple digest types and/or multiple trust anchor keys.
- Other DNSSEC RRs necessary for a signed zone.
- There MAY be other records to provide information about this witness zone. There SHOULD NOT be any records unrelated to witness operations, such as delegations.

4.2. Lifecycle of witness zones

A trust anchor SHOULD have many witness zones, in order to provide resilience as well as dispersal of trust.

Each witness zone is tied to a fixed witness trust anchor. The zone lasts as long as its trust anchor. This SHOULD be at least 10 years, since old software and configurations cannot function after too many of their witnesses have been retired.

Witnesses are continually retired. It is expected that some witnesses will have to retire early, for instance, if their keys are lost or compromised, or if their host organization is no longer able to maintain them. This is OK since there are plenty of other witnesses.

New witnesses are continually introduced. Validators configured with an up-to-date set of witnesses will have a decent lifetime. Given an average witness lifetime of W years, a pool of P witnesses, and a quorum size of Q, we expect P/W witness retirements per year. A validator configuration will last until there are Q witnesses left, that is, until there have been P-Q retirements, which takes V=(P-Q)*(W/P) years. For example, if P=30, W=10, and Q=6, then V=8.

A witness organization may run multiple witness zones on a rolling replacement schedule in order to avoid a hiatus when a zone is retired. Validators SHOULD be configured to use only one witness zone from each witness organization, to avoid trusting one organization too much.

5. The root trust anchor

5.1. Locating root trust anchor witness zones

There is a bootstrapping problem when a validator has an out-of-date root trust anchor: it needs to find the name servers for the witness zones in order to be able to get the WS records that vouch for the new root trust anchor; however it is unable to validate the responses it gets while resolving the name server addresses. This section describes how to minimize this bootstrapping problem.

All root witness zones SHALL be delegated from a single parent zone, called root-witnesses.arpa. This zone is to be maintained by IANA. A delegation in this zone indicates that there are out-of-band arrangements between the root DNSSEC key managers (XXX do they have a better name?) and the witness organization allowing the witness organization to meaningfully vouch for changes to the root DNSSEC

trust anchor.

The root-witnesses.arpa zone SHOULD NOT be signed. Leaving the zone unsigned prevents the risk that validators will use some higher-level trust anchor to validate responses from a witness zone rather than the witness trust anchor itself. In particular we want to avoid a compromised root key being used to vouch for itself. The purpose of the root-witnesses.arpa zone is to contain delegation NS RRs and glue address records for the witness zones, and these records are never signed. The signed parts of delegations are the DS RRsets; omitting those prevents unsafe witness validation, but also leaves almost nothing in the root-witnesses.arpa zone to sign.

Each witness zone's name servers have names inside that witness zone so that they can be validated by the witness trust anchor without depending on any other part of the DNS.

The root-witnesses.arpa zone SHALL be served by the root name servers. This is so that a bootstrapping validating resolver can find its witnesses using just its root hints, and get a direct referral to the right witness zone name servers, again without depending on any other part of the DNS.

Since the referral from root-witnesses.arpa is to a zone for which the validating resolver has a trust anchor, it does not need to validate the delegation chain through the witness zone's parents. Depending on its validation strategy, a bootstrapping validator might need special logic to avoid validaing this delegation chain or to ignore validation failures in it, in particular when its root trust anchor is stale.

5.2. Root trust anchor witness organizations

In order to populate the root-witnesses.arpa zone, IANA has to select a number of root witness organizations who will receive delegations from this zone.

These might be selected from TLD operators, root server operators, registry service providers, accredited registrars, and others who have an interest in the security and robustness of the DNS.

6. Security Considerations

This memo aims to improve the security and robustness of the DNS. There are security-related requirements and recommendations

throughout.

7. IANA Considerations

A DNS RR type number is required for the WS RR.

This memo directs IANA to set up the root-witnesses.arpa zone and manage ongoing liaison with the root trust anchor witnesses.

8. Normative references

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2181] Elz, R. and R. Bush, "Clarifications to the DNS Specification", RFC 2181, July 1997.

[RFC3757] Kolkman, O., Schlyter, J., and E. Lewis, "Domain Name System KEY (DNSKEY) Resource Record (RR) Secure Entry Point (SEP) Flag", RFC 3757, April 2004.

[RFC4034] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", RFC 4034, March 2005.

[RFC4035] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", RFC 4035, March 2005.

[RFC5011] StJohns, M., "Automated Updates of DNS Security (DNSSEC) Trust Anchors", STD 74, RFC 5011, September 2007.

[RFC6781] Kolkman, O., Mekking, W., and R. Gieben, "DNSSEC Operational Practices, Version 2", RFC 6781, December 2012.

Appendix A. Questions

Should this this scheme be extended to be more like DLV? That is,

witness zones could vouch for any trust anchors. Validators would look up WS records by concatenating the name of the automatically updated trust anchor name and the witness zone name. There would be an implicit many-to-many relationship between witnesses and automatically updated trust anchors, instead of an explicit many-to-one relationship.

Author's Address

Tony Finch
University of Cambridge Computing Service
Roger Needham Building
7 JJ Thomson Avenue
Cambridge CB3 0RB
ENGLAND

Phone: +44 797 040 1426
Email: dot@dotat.at
URI: http://dotat.at/