DNSOP Working Group
Internet-Draft
Intended status: Standards Track
Expires: 1 January 2022

P. Thomassen
deSEC, Secure Systems Engineering

30 June 2021

DNSSEC Bootstrapping
draft-thomassen-dnsop-dnssec-bootstrapping-00

Abstract

This document describes an authenticated in-band method for automatic signaling of a DNS zone's delegation signer information from the zone's DNS operator. The zone's registrar or registry may subsequently use this signal for automatic DS record provisioning in the parent zone.

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 https://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 1 January 2022.

Table of Contents

1. Introduction 2
1.1. Terminology 3
1.2. Requirements Notation 4
2. Description 4
2.1. Preconditions 4
2.1.1. Example 4
2.1.2. Zone Cut Clarification 5
2.2. Bootstrapping Method 5
2.2.1. Steps Taken by the Child DNS Operator 5
2.2.2. Steps Taken by the Parental Agent 6
3. Implementation Status 7
4. Security Considerations 8
5. IANA Considerations 8
6. Acknowledgements 9
7. Normative References 9
Appendix A. Change History (to be removed before final publication) 10
Author's Address 10

1. Introduction

TODO remove: this section is inspired by [RFC7344], Section 1.

The first time a DNS Operator signs a zone, they need to communicate the keying material to the Parent. Depending on the desires of the Parent, the Child might send their DNSKEY record, a DS record, or both.

So far, out-of-band methods are typically used to complete the chain of trust. In-band methods exist, in particular based on the CDS and CDNSKEY record types as specified in [RFC7344] and [RFC8078]. However, such communication is only authenticated when performing a rollover of the Child's keys represented in the parent. An authenticated in-band channel for enabling DNSSEC so far has been missing.

How the keying material is conveyed to the Parent during initial DNSSEC bootstrapping depends on the relationship the Child has with the Parent. In many cases this is a manual process -- and not an easy one. The communication has to occur between the DNS Operator and, depending on the circumstances, the Registry or the Registrar,

possibly via the Registrant (for details, see [RFC7344], Appendix A). Any manual process is susceptible to mistakes and/or errors. In addition, due to the annoyance factor of the process, Operators may avoid the process of getting a DS record set published at the Parent.

DNSSEC provides data integrity to information published in DNS; thus, DNS publication can be used to automate maintenance of delegation information. This document describes a method to automate publication of inital DS records for a hitherto insecure delegation.

Readers are expected to be familiar with DNSSEC, including [RFC4033], [RFC4034], [RFC4035], [RFC6781], [RFC7344], and [RFC8078].

This document describes a method for automated provisioning of the delegation trust information and proposes a polled/periodic trigger for simplicity. Some users may prefer a different trigger. These alternate additional triggers are not discussed in this document.

1.1. Terminology

The terminology we use is defined in this section. The highlighted roles are as follows:

Child The entity on record that has the delegation of the domain from the Parent.

Parent The domain in which the Child is registered.

Child DNS Operator The entity that maintains and publishes the zone information for the Child DNS.

Parental Agent The entity that the Child has a relationship with to change its delegation information.

Bootstrapping Domain Given an authoritative nameserver hostname from the Child's NS record set, that hostname prefixed the label "_boot".

Signaling Name A Bootstrapping Domain prefixed with a label encoding the Child's name.

CDS/CDNSKEY This notation refers to CDS and/or CDNSKEY, i.e., one or both.

Base32hex Encoding "Base 32 Encoding with Extended Hex Alphabet" as per [RFC4648].

1.2. Requirements Notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Description

When setting up initial trust, the child generally wants to enable global validation. As long as the child is insecure, DNS answers can be forged. The goal is to promote the child from insecure to secure as soon as reasonably possible by the parent. This means that the period from the child's publication of CDS/CDNSKEY RRset to the parent publishing the synchronized DS RRset should be as short as possible.

This goal is achieved by transferring trust from the Child DNS Operator.

2.1. Preconditions

In order to use this technique, the following conditions have to be met:

1. The Child DNS Operator SHOULD publish CDS/CDNSKEY records at the Child's apex, as described in [RFC7344].

2. Each Bootstrapping Domain MUST be part of a securely delegated zone, i.e. has a valid DNSSEC chain of trust from the root.

3. The Child DNS Operator MUST be able to maintain and publish DNS information in these zone (i.e. under the Bootstrapping Domains).

For operational or other reasons, a Bootstrapping Domain MAY coincide with a zone cut.

2.1.1. Example

When performing DNSSEC bootstrapping for the Child zone "example.com" using NS records "ns1.example.net" and "ns2.example.net", the Child DNS Operator

1. should publish CDS/CDNSKEY records at "example.com";

2. needs to ensure that a valid DNSSEC chain of trust exists for the zone(s) that are authoritative for the Bootstrapping Domains "_boot.ns1.example.net" and "_boot.ns2.example.net";

3. must be able to maintain and publish DNS information in these zones.

2.1.2. Zone Cut Clarification

A Bootstrapping Domain such as "_boot.ns1.example.net" may be a zone of its own, in which case it needs to be secure and under the control of the Child DNS Operator. If the Bootstrapping Domain does not coincide with a zone cut, these conditions are instead imposed on the containing zone (such as "example.net").

The "Bootstrapping Domain" terminology is necessary to describe the bootstrapping mechanism without regard to whether there is a zone cut at these names or not.

2.2. Bootstrapping Method

2.2.1. Steps Taken by the Child DNS Operator

To perform DNSSEC bootstrapping for the Child zone, the Child DNS Operator MUST (re-)publish the Child's CDS/CDNSKEY records at the corresponding Signaling Name under each Bootstrapping Domain (see example below). These records belong to the autoritative zone of the Bootstrapping Domain, and as such they MUST be signed with that zone's keys, and MUST NOT be signed with the Child zone's keys.

The Signaling Name contains a label identifying the Child's name. This label MUST be equal to the SHA-256 hash digest of the Child's name in "Base 32 Encoding with Extended Hex Alphabet", as specified in [RFC4648]. Trailing padding characters ("=") MUST be dropped.

Previous uses of CDS/CDNSKEY records are specified at the apex only ([RFC7344], Section 4.1). This protocol extends the use of these record types on non-apex owner names for the purpose of DNSSEC bootstrapping. To avoid the possibility of semantic collision, there MUST NOT be a zone cut at a Signaling Name.

TODO Remove Note 1: The purpose of the hash function is to avoid the possibility of exceeding the maximum length of a DNS name. This could occur if the Child name was used as is.

TODO Remove Note 2: The encoding choice is like in NSEC3, except that SHA-256 is used instead of SHA-1. This is to prevent other tenants in shared hosting environments from creating collisions.

2.2.1.1. Example

To bootstrap the Child zone "example.com" using NS records "ns1.example.net" and "ns2.example.net", the Bootstrapping Domains are "_boot.ns1.example.net" and "_boot.ns2.example.net". The Child DNS Operator thus (re-)publishes the Child's CDS/CDNSKEY records under the names

```
kdsqdtnelusqanhnhg8o0d72ekf6gbtbjsmj1aojq895b1me353g._boot.ns1.example.net
kdsqdtnelusqanhnhg8o0d72ekf6gbtbjsmj1aojq895b1me353g._boot.ns2.example.net
```

where "kdsqdtnelusqanhnhg8o0d72ekf6gbtbjsmj1aojq895b1me353g" is the unpadded Base32hex Encoding of "example.com". The records are accompanied by RRSIG records created using the key(s) of the zone which is authoritative for the respective Bootstrapping Domain.

TODO remove: Should hash input include trailing dot? (Command was: "echo -n example.com | openssl dgst -binary -sha256 | base32hex | tr -d =")

2.2.2. Steps Taken by the Parental Agent

When the Parental Agent receives a new NS record set (or additionally at any other time considered appropriate), the Parental Agent, knowing both the Child zone name and its NS hostnames,

1. MUST query the CDS/CDNSKEY records located at each of the Signaling Names (using standard DNS resolution);

2. MUST perform DNSSEC validation of all responses retrieved in Step 1;

3. SHOULD query the CDS/CDNSKEY records located at the Child zone apex, directly from each of the authoritative nameservers as given in the Child NS record set;

4. MUST checks that all CDS/CDNSKEY record sets retrieved in Steps 1 and 3 have equal record contents;

5. SHOULD derive a DS record set from the retrieved CDS/CDNSKEY record sets and publish it in the Parent zone, as to secure the Child's delegation.

If an error condition occurs during Steps 1--4, in particular:

* DNS resolution failure during retrieval of CDS/CDNSKEY records from any Signaling Name (Step 1), or failure of DNSSEC validation (Step 2),

* Failure to retrieve CDS/CDNSKEY records located at the Child apex from all of the Child's authoritative nameservers (Step 3),

* Inconsistent responses (Step 4),

the Parental Agent MUST NOT proceed to Step 5.

2.2.2.1. Example

To bootstrap the Child zone "example.com" using NS records "ns1.example.net" and "ns2.example.net", the Parental Agent

1. queries CDS/CDNSKEY records, using standard DNS resolution, for the names

kdsqdtnelusqanhnhg8o0d72ekf6gbtbjsmj1aojq895b1me353g._boot.ns1.example.net
kdsqdtnelusqanhnhg8o0d72ekf6gbtbjsmj1aojq895b1me353g._boot.ns2.example.net

2. performs DNSSEC validation of the responses retrieved in Step 1;

3. queries CDS/CDNSKEY records for "example.com" directly from "ns1.example.net" and "ns2.example.net";

4. checks that CDS record sets retrieved in Step 1 agree across responses and also with the CDS record sets retrieved in Step 3; ditto for CDNSKEY;

5. publishes a DS record set according to the information retrieved in the previous steps.

2.2.2.2. Opt-out

As a special case of Step 4 failure, the Child MAY opt out from DNSSEC bootstrapping by publishing a CDS/CDNSKEY record with algorithm 0 and other fields as specified in [RFC8078], Section 4, at its apex. (This opt-out mechanism is without regard to whether the Child DNS Operator signs the zones and publishes records at the Signaling Names.)

3. Implementation Status

Note to the RFC Editor: please remove this entire section before publication.

* PowerDNS supports manual creation of CDS/CDNSKEY records on non-apex names.

* TODO Proof of concept

4. Security Considerations

Thoughts (to be expanded):

* We use at least one established chain of trust (via the secure delegations of the zones containing the NS hostnames). As a result,
  - communication is authenticated;
  - process is immediate (no need for observing CDS/CDNSKEY records via TCP for several days);
  - an active on-wire attacker cannot tamper with the delegation.
* When validating against CDS/CDNSKEY records at the Child's apex, the security level of the method is strictly higher than the "accept CDS/CDNSKEY after a while"-approch that is already in use at several ccTLD registries ("Accept after Delay", [RFC8078], Section 3.3). This is because the method described here adds stronger guarantees, but removes nothing. Perhaps this means that co-publication of CDS/CDNSKEY at the Child apex should be mandatory. (This in turn may interact somehow with the Child's opt-out option.)
* Actors in the chain(s) of trust of the zone(s) used for bootstrapping (the DNS Operator themselves, plus entities further up in the chain) can undermine the protocol. However,
  - that's also possible in the case of CDS/CDNSKEY (see previous point);
  - if the Child DNS Operator doesn't control the zones in which its NS hostnames live (including their nameservers' A records) because the path from the root is untrusted, you probably don't want to trust that operator as a whole;
  - when bootstrapping is done upon receipt of a new NS record set, the window of opportunity is very small (and easily monitored by the Child DNS operator);
  - mitigation exists by diversifying e.g. the nameserver hostname's TLDs, which is advisable anyways.

5. IANA Considerations

TODO: reserve "_boot"?

This document has no IANA actions.

6. Acknowledgements

Thanks to TODO for reviewing draft proposals and offering comments and suggestions.

Thanks also to Steve Crocker, Hugo Salgado, and Ulrich Wisser for early-stage brainstorming.

7. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[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, <https://www.rfc-editor.org/info/rfc4033>.

[RFC4034] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", RFC 4034, DOI 10.17487/RFC4034, March 2005, <https://www.rfc-editor.org/info/rfc4034>.

[RFC4035] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", RFC 4035, DOI 10.17487/RFC4035, March 2005, <https://www.rfc-editor.org/info/rfc4035>.

[RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", RFC 4648, DOI 10.17487/RFC4648, October 2006, <https://www.rfc-editor.org/info/rfc4648>.

[RFC6781] Kolkman, O., Mekking, W., and R. Gieben, "DNSSEC Operational Practices, Version 2", RFC 6781, DOI 10.17487/RFC6781, December 2012, <https://www.rfc-editor.org/info/rfc6781>.

[RFC7344] Kumari, W., Gudmundsson, O., and G. Barwood, "Automating DNSSEC Delegation Trust Maintenance", RFC 7344, DOI 10.17487/RFC7344, September 2014, <https://www.rfc-editor.org/info/rfc7344>.

[RFC8078] Gudmundsson, O. and P. Wouters, "Managing DS Records from the Parent via CDS/CDNSKEY", RFC 8078, DOI 10.17487/RFC8078, March 2017, <https://www.rfc-editor.org/info/rfc8078>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

Appendix A. Change History (to be removed before final publication)

* draft-thomassen-dnsop-dnssec-bootstrapping-00

| Initial public draft.

Author's Address

Peter Thomassen
deSEC, Secure Systems Engineering
Berlin
Germany

Email: peter@desec.io