DNSOP Working Group
Internet-Draft
Updates: 7344, 8078 (if approved)
Intended status: Standards Track
Expires: 19 December 2022

P. Thomassen
deSEC, Secure Systems Engineering
N. Wisiol
deSEC, Technische Universität Berlin

17 June 2022

Automatic DNSSEC Bootstrapping using Authenticated Signals from the Zone's Operator
draft-ietf-dnsop-dnssec-bootstrapping-01

Abstract

This document introduces an in-band method for DNS operators to publish arbitrary information about the zones they are authoritative for, in an authenticated fashion and on a per-zone basis. The mechanism allows managed DNS operators to securely announce DNSSEC key parameters for zones under their management, including for zones that are not currently securely delegated.

Whenever DS records are absent for a zone's delegation, this signal enables the parent's registry or registrar to cryptographically validate the CDS/CDNSKEY records found at the child's apex. The parent can then provision DS records for the delegation without resorting to out-of-band validation or weaker types of cross-checks such as "Accept after Delay" ([RFC8078]).

This document updates [RFC8078] and replaces its Section 3 with Section 3.2 of this document.

[Ed note: This document is being collaborated on at https://github.com/desec-io/draft-ietf-dnsop-dnssec-bootstrapping/ (https://github.com/desec-io/draft-ietf-dnsop-dnssec-bootstrapping/). The authors gratefully accept pull requests.]

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/.

Copyright Notice

Table of Contents

1. Introduction

Securing a DNS delegation for the first time requires that the Child's DNSSEC parameters be conveyed to the Parent through some trusted channel. While the communication conceptually has to occur between the Parent registry and the DNSSEC key holder, what exactly that means and how the communication is coordinated traditionally depends on the relationship the Child has with the Parent.

A typical situation is that the key is held by the Child DNS Operator; the communication thus often involes this entity. In addition, depending on the circumstances, it may also involve the Registrar, possibly via the Registrant (for details, see [RFC7344], Appendix A).

As observed in [RFC7344], these dependencies often result in a manual process that is susceptible to mistakes and/or errors. In addition, due to the annoyance factor of the process, involved parties may avoid the process of getting a DS record set published in the first place.

To alleviate these problems, automated provisioning of DS records has been specified in ([RFC8078]). It is based on the Parental Agent (registry or registrar) fetching DNSSEC key parameters in the form of CDS and CDNSKEY records ([RFC7344]) from the Child zone's apex, and validating them somehow. This validation can be done using DNSSEC itself if the objective is to update an existing DS record set (such as during key rollover). However, when bootstrapping a DNSSEC delegation, the Child zone has no existing DNSSEC validation path, and other means to ensure the CDS/CDNSKEY records' legitimacy must be found.

For lack of a comprehensive DNS-innate solution, either out-of-band methods have been used so far to complete the chain of trust, or

cryptographic validation has been entirely dispensed with, in exchange for weaker types of cross-checks such as "Accept after Delay" ([RFC8078] Section 3.3). [RFC8078] does not define an in-band validation method for enabling DNSSEC.

This document aims to close this gap by introducing an in-band method for DNS Operators to publish arbitrary information about the zones they are authoritative for, in an authenticated manner and on a per-zone basis. The mechanism allows managed DNS Operators to securely announce DNSSEC key parameters for zones under their management. The Parent can then use this signal to cryptographically validate the CDS/CDNSKEY records found at an insecure Child zone's apex, and upon success secure the delegation.

While applicable to the vast majority of domains, the protocol does not support certain edge cases, such as excessively long Child zone names, or DNSSEC bootstrapping for domains with in-bailick nameservers only (see Section 3.4).

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

1.1. Terminology

This section defines the terminology used in this document.

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

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

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

Parent The domain in which the Child is registered.

Parental Agent The entity that has the authority to insert DS records into the Parent zone on behalf of the Child. (It could the the registry, registrar, a reseller, or some other authorized entity.)

Signaling Domain A hostname from the Child's NS record set, prefixed with the label _signal. There are as many Signaling Domains as there are distinct NS targets.

Signaling Name The labels that are prefixed to a Signaling Domain in order to identify a Signaling Type and a Child zone's name (see Section 2.2).

Signaling Record A DNS record located at a Signaling Name under a Signaling Domain. Signaling Records are used by the Child DNS Operator to publish information about the Child.

Signaling Type A signal type identifier, such as _dsboot for DNSSEC bootstrapping.

Signaling Zone The zone which is authoritative for a given Signaling Record.

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. Signaling

This section describes the general mechanism by which a Child DNS Operator can publish an authenticated signal about a Child zone. Parental Agents (or any other party) can then discover and process the signal. Authenticity is ensured through standard DNSSEC validation.

2.1. Chain of Trust

If a Child DNS Operator implements the protocol, each Signaling Zone MUST be signed and securely delegated, i.e. have a valid DNSSEC chain of trust.

For example, when publishing a signal that relates to a Child zone with NS records ns1.example.net and ns2.example.org, the Child DNS Operator needs to ensure that a valid DNSSEC chain of trust exists for the zone(s) that are authoritative for the Signaling Domains _signal.ns1.example.net and _signal.ns2.example.org.

2.2. Signaling Names

To publish a piece of information about the Child zone in an authenticated fashion, the Child DNS Operator MUST publish one or more Signaling Records at a Signaling Name under each Signaling Domain.

Signaling Records MUST be accompanied by RRSIG records created with the corresponding Signaling Zone's key(s). The type and contents of these Signaling Records depend on the type of signal.

The Signaling Name identifies the Child and the Signaling Type. It is identical to the Child name (but with the final root label removed), prefixed with a label containing the Signaling Type.

3. Bootstrapping a DNSSEC Delegation

When the Child zone's CDS/CDNSKEY RRsets are used for setting up initial trust, they need to be authenticated. This is achieved by co-publishing the Child's CDS/CDNSKEY records as an authenticated signal from the Child DNS Operator. The Parent can discover and validate this signal, thus transferring trust from the Child DNS Operator to the Child zone.

Child DNS Operators and Parental Agents who wish to use CDS/CDNSKEY records for DNSSEC bootstrapping SHOULD support the authentication protocol described in this section.

3.1. Signaling Consent to Act as the Child's Signer

To confirm its willingness to act as the Child's delegated signer and authenticate the Child's CDS/CDNSKEY RRsets, the Child DNS Operator MUST co-publish them at the corresponding Signaling Name under each out-of-bailiwick Signaling Domain (Section 2.2). For simplicity, the Child DNS Operator MAY also co-publish the Child's CDS/CDNSKEY RRsets under Signaling Domains that are in bailiwick, although those Signaling Domains are not used for validation (Section 3.2).

Unlike the CDS/CDNSKEY records at the Child's apex, Signaling Records MUST be signed with the corresponding Signaling Zone's key(s). Their contents MUST be identical to the corresponding records published at the Child's apex.

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

3.1.1. Example

For the purposes of bootstrapping the Child zone example.co.uk with NS records ns1.example.net, ns2.example.org, and ns3.example.co.uk, the required Signaling Domains are _signal.ns1.example.net and _signal.ns2.example.org.

In the zones containing these domains, the Child DNS Operator authenticates the CDS/CDNSKEY records found at the Child's apex by co-publishing them at the names:

_dsboot.example.co.uk._signal.ns1.example.net
_dsboot.example.co.uk._signal.ns2.example.org

The records are accompanied by RRSIG records created using the key(s) of the respective Signaling Zone.

Publication of Signaling Records under the in-bailiwick domain _signal.ns3.example.co.uk is not required.

3.2. Validating CDS/CDNSKEY Records for DNSSEC Bootstrapping

This section replaces Section 3 of [RFC8078].

To validate a Child's CDS/CDNSKEY RRset for DNSSEC bootstrapping, the Parental Agent, knowing both the Child zone name and its NS hostnames, MUST execute the following steps:

1. verify that the Child is not currently securely delegated and that at least one of its nameservers is out of bailiwick;

2. query the CDS/CDNSKEY records at the Child zone apex directly from each of the authoritative servers as determined by the delegation's NS record set (without caching);

3. query the CDS/CDNSKEY records located at the Signaling Name under each out-of-bailiwick Signaling Domain using a trusted DNS resolver and enforce DNSSEC validation;

4. check (separately by record type) that all record sets retrieved in Steps 2 and 3 have equal contents;

If the above steps succeed without error, the CDS/CDNSKEY records are successfully validated, and the Parental Agent can proceed with the publication of the DS record set under the precautions described in [RFC8078], Section 5.

If, however, an error condition occurs, in particular:

* in Step 1: the Child is already securely delegated or has in-bailiwick nameservers only;

* in Step 2: any failure during the retrieval of the CDS/CDNSKEY records located at the Child apex from any of the authoritative nameservers;

* in Step 3: any failure to retrieve the CDS/CDNSKEY RRsets located at the Signaling Name under any Signaling Domain, including failure of DNSSEC validation, or unauthenticated data (AD bit not set);

* in Step 4: inconsistent responses (for at least one of the types),

including a record set that is empty in one of Steps 2 or 3, but non-empty in the other;

the Parental Agent MUST abort the procedure.

3.2.1. Example

To verify the CDS/CDNSKEY records for the Child example.co.uk, the Parental Agent (assuming that the Child delegation's NS records are ns1.example.net, ns2.example.org, and ns3.example.co.uk)

1. checks that the Child domain is not yet securely delegated;

2. queries CDS/CDNSKEY records for example.co.uk directly from ns1.example.net, ns2.example.org, and ns3.example.co.uk (without caching);

3. queries and validates the CDS/CDNSKEY records located at (see Section 2.2; ns3.example.co.uk is ignored because it is in bailiwick)

_dsboot.example.co.uk._signal.ns1.example.net
_dsboot.example.co.uk._signal.ns2.example.org

4. checks that the CDS/CDNSKEY record sets retrieved in Steps 2 and 3 agree across responses.

If all these steps succeed, the Parental Agent can proceed to publish a DS record set as indicated by the validated CDS/CDNSKEY records.

The Parental Agent does not use in-bailiwick Signaling Names during validation because they cannot have a pre-established chain of trust at bootstrapping time, so are not useful for bootstrapping. Consequently, if all NS hostnames are in bailiwick, validation cannot be completed, and DS records are not published.

3.3. Triggers

Parental Agents SHOULD trigger the procedure described in Section 3.2 once one of the following conditions is fulfilled:

* The Parental Agent receives a new or updated NS record set for a Child;

* The Parental Agent encounters Signaling Records during a proactive, opportunistic scan (e.g. daily queries for the Signaling Records of some or all of its delegations);

* Any other condition as deemed appropriate by local policy.

Most types of discovery (such as daily scans of delegations) are based directly on the delegation's NS record set. In this case, these NS names can be used as is by the bootstrapping algorithm (Section 3.2) for querying Signaling Records.

Some discovery methods, however, do not imply reliable knowledge of the Child's NS record set. For example, when discovering Signaling Names by performing an NSEC walk or zone transfer for a Signaling Domain, the Parental Agent MUST NOT assume that the nameserver(s) under whose Signaling Domain(s) a Signaling Name appears is in fact authoritative for the corresponding Child.

In this case (and in other cases alike where some list of "bootstrappable domains" is retrieved elsewhere), the Parental Agent MUST ascertain that the Child's delegation actually contains the nameserver hostname seen during discovery, and ensure that Signaling Record queries are only made against the proper set of nameservers as listed in the Child's delegation from the Parent.

3.4. Limitations

As a consequence of Step 3 in Section 3.2, DS bootstrapping does not work for fully in-bailiwick delegations, as no pre-existing chain of trust to the Child domain is available during bootstrapping.

The protocol is further restricted by the fact that the fully qualified Signaling Names fit within the general limits that apply to DNS names (such as their length and label count).

4. Operational Recommendations

4.1. Child DNS Operator

Signaling Domains SHOULD be delegated as zones of their own, so that the Signaling Zone's apex coincides with the Signaling Domain (such as _signal.ns1.example.net). While it is permissible for the Signaling Domain to be contained in a Signaling Zone of fewer labels (such as example.net), a zone cut ensures that bootstrapping activities do not require modifications of the zone containing the nameserver hostname.

To keep the size of the Signaling Zones minimal and bulk processing efficient (such as via zone transfers), Child DNS Operators SHOULD remove Signaling Records which are found to have been acted upon.

4.2. Parental Agent

It is RECOMMENDED to perform queries within Signaling Domains (Section 3.2) with an (initially) cold resolver cache or to limit the TTL of cached records to the interval between scans, as to retrieve the most current information regardless of TTL. (When a batch job is used to attempt bootstrapping for a large number of delegations, the cache does not need to get cleared in between.)

5. Implementation Status

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

5.1. Child DNS Operator-side

* Knot DNS supports manual creation of non-apex CDS/CDNSKEY records.
* PowerDNS supports manual creation of non-apex CDS/CDNSKEY records.
* Proof-of-concept Signaling Domains with several thousand Signaling Names exist at _signal.ns1.desec.io and _signal.ns2.desec.org.
* Another DNS operator has implemented the protocol (synthesizing Signaling Records for a significant number of domains).
* The authors are planning to develop a tool for automatic generation of signaling records.

5.2. Parental Agent-side

* A tool to retrieve and process Signaling Records for bootstrapping purposes, either directly or via zone walking, is available at https://github.com/desec-io/dsbootstrap (https://github.com/desec-io/dsbootstrap). The tool outputs the validated DS records which then can be added to the Parent zone.
* Some registries/registrars (e.g. .cl, GoDaddy) are working on implementations of the protocol.

6. Security Considerations

The protocol adds authentication to the CDS/CDNSKEY-based bootstrapping concept of [RFC8078], while removing nothing. Its security level is therefore strictly higher than that of existing approaches described in that document (e.g. "Accept after Delay"). Apart from this general improvement, the same Security Considerations apply as in [RFC8078].

The level of rigor in Section 3.2 is needed to prevent publication of a half-baked DS RRset (authorized only under a subset of NS hostnames). This ensures, for example, that an operator in a multi-homed setup cannot enable DNSSEC unless all other operators agree.

Because the parents of a Signaling Domain (such as the corresponding TLD registry) are in control of its chain of trust, they are also able to undermine the signal's authenticity. To mitigate this risk, it is RECOMMENDED to increase the effort required to collude for taking control of all Signaling Domains, by diversifying the path from the root to each nameserver. This is best achieved by using different and independently operated TLDs for each one. (TLD-independent NS hostnames are advisable anyways in DNS operations, in order to prevent the TLD from becoming a single point of failure.) Furthermore, as the Child DNS Operator has authoritative knowledge of the Child's CDS/CDNSKEY records, it can readily detect fraudulent provisioning of DS records.

7. IANA Considerations

Per [RFC8552], IANA is requested to add the following entries to the "Underscored and Globally Scoped DNS Node Names" registry:

RR Type	_NODE NAME	Reference
CDS	_signal	[draft-ietf-dnsop-dnssec-bootstrapping]
CDNSKEY	_signal	[draft-ietf-dnsop-dnssec-bootstrapping]

8. Acknowledgements

Thanks to Brian Dickson, Ondřej Caletka, John R. Levine, Christian Elmerot, Oli Schacher, Donald Eastlake, and Libor Peltan for reviewing draft proposals and offering comments and suggestions.

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

9. 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>.

[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>.

[RFC8552] Crocker, D., "Scoped Interpretation of DNS Resource Records through "Underscored" Naming of Attribute Leaves", BCP 222, RFC 8552, DOI 10.17487/RFC8552, March 2019, <https://www.rfc-editor.org/info/rfc8552>.

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

* draft-ietf-dnsop-dnssec-bootstrapping-01

| Allow bootstrapping when some (not all) NS hostnames are in bailiwick.
|
| Clarified Operational Recommendations according to operator feedback.

|
| Turn loose Security Considerations points into coherent text.
|
| Do no longer suggest NSEC-walking Signaling Domains. (It does not work well due to the Signaling Type prefix. What's more, it's unclear who would do this: Parents know there delegations and can do a targeted scan; others are not interested.)
|
| Editorial changes.
|
| Added IANA request.
|
| Introduced Signaling Type prefix (_dsboot), renamed Signaling Name infix from _dsauth to _signal.

* draft-ietf-dnsop-dnssec-bootstrapping-00

| Editorial changes.

* draft-thomassen-dnsop-dnssec-bootstrapping-03

| Clarified importance of record cleanup by moving paragraph up.
|
| Pointed out limitations.
|
| Replace [RFC8078] Section 3 with our Section 3.2.
|
| Changed _boot label to _dsauth.
|
| Removed hashing of Child name components in Signaling Names.
|
| Editorial changes.

* draft-thomassen-dnsop-dnssec-bootstrapping-02

| Reframed as an authentication mechanism for RFC 8078.
|
| Removed multi-signer use case (focus on RFC 8078 authentication).
|
| Triggers need to fetch NS records (if not implicit from context).
|
| Improved title.
|
| Recognized that hash collisions are dealt with by Child apex
| check.

* draft-thomassen-dnsop-dnssec-bootstrapping-01

| Add section on Triggers.
|
| Clarified title.
|
| Improved abstract.
|
| Require CDS/CDNSKEY records at the Child.
|
| Reworked Signaling Name scheme.
|
| Recommend using cold cache for consumption.

|
| Updated terminology (replace "Bootstrapping" by "Signaling").
|
| Added NSEC recommendation for Bootstrapping Zones.
|
| Added multi-signer use case.
|
| Editorial changes.

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

| Initial public draft.

Authors' Addresses

Peter Thomassen
deSEC, Secure Systems Engineering
Berlin
Germany
Email: peter@desec.io

Nils Wisiol
deSEC, Technische Universität Berlin
Berlin
Germany
Email: nils@desec.io