DNSOP Working Group
Internet-Draft
Updates: 7344, 7477 (if approved)
Intended status: Standards Track
Expires: 2 February 2026

P. Thomassen
SSE - Secure Systems Engineering GmbH

1 August 2025

Clarifications on CDS/CDNSKEY and CSYNC Consistency
draft-ietf-dnsop-cds-consistency-08

Abstract

Maintenance of DNS delegations requires occasional changes of the DS and NS record sets on the parent side of the delegation. For the case of DS records, RFC 7344 provides automation by allowing the child to publish CDS and/or CDNSKEY records holding the prospective DS parameters which the parent can ingest. Similarly, RFC 7477 specifies CSYNC records to indicate a desired update of the delegation's NS (and glue) records. Parent-side entities (e.g. Registries, Registrars) can query these records from the child and, after validation, use them to update the parent-side RRsets of the delegation.

This document specifies that when performing such queries, parent-side entities MUST ensure that updates triggered via CDS/CDNSKEY and CSYNC records are consistent across the child's authoritative nameservers, before taking any action based on these records.

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 2 February 2026.

Table of Contents

1. Introduction

[RFC7344] automates DNSSEC delegation trust maintenance by having the child publish CDS and/or CDNSKEY records which hold the prospective DS parameters. Similarly, [RFC7477] specifies CSYNC records indicating a desired update of the delegation's NS and associated glue records. Parent-side entities (e.g. Registries, Registrars) can use these records to update the corresponding records of the

delegation.

For ingesting CSYNC records, [RFC7477] Section 3.1 advocates that Parental Agents limit queries to just one authoritative nameserver (as typically done in normal resolution). The corresponding Section 6.1 of [RFC7344] (CDS/CDNSKEY) contains no provision for how specifically queries for these records should be done.

Retrieving records from just one authoritative server (e.g., by directing queries towards a trusted validating resolver) works fine when all is in order. However, problems may arise if CDS/CDNSKEY/CSYNC record sets are inconsistent across authoritative nameservers, either because they are out of sync (e.g., during a KSK rollover), or because they are not controlled by the same entity (e.g., in a multi-signer setup [RFC8901]).

In such cases, if CDS/CDNSKEY/CSYNC records are retrieved from one nameserver only ("naively", without a consistency check), each nameserver can unilaterally trigger an update of the delegation’s DS or NS record set.

For example, a single provider in a multi-signer setup may (accidentally or maliciously) cause another provider’s trust anchors and/or nameservers to be removed from the delegation. This can occur both when the multi-signer configuration is temporary (during a provider change) and when it is permanent (for redundancy). In any case, a single provider should not be in the position to remove the other providers’ records from the delegation.

Similar breakage can occur when the delegation has lame nameservers, where an attacker may illegitimately initialize a DS record set and then manipulate the delegation’s NS record set at will. More detailed examples are given in Appendix A.

For a CDS/CDNSKEY/CSYNC consumer, it is generally impossible to estimate the impact of a requested delegation update unless all of the child’s authoritative nameservers are inspected. At the same time, applying an automated delegation update "MUST NOT break the current delegation" ([RFC7344], Section 4.1), i.e., it MUST NOT hamper availability or validatability of the Child’s resolution.

This document therefore specifies that parent-side entities need to ensure that the updates indicated by CDS/CDNSKEY and CSYNC record sets are plausibly consistent across the child’s nameservers, before taking any action based on these records.

Readers are expected to be familiar with DNSSEC, including [RFC4033], [RFC4034], [RFC4035], [RFC6781], [RFC7344], [RFC7477], and [RFC8901].

1.1. 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.

1.2. Terminology

Multi-provider setup: A constellation where several providers independently operate authoritative DNS service for a domain, usually for purposes of redundancy. This includes setups both with and without DNSSEC.

Multi-signer setup: A multi-provider setup for a DNSSEC-enabled domain with multiple independent signing entities [RFC8901]. Such a setup may be permanent (for redundancy) or temporary (for continuity of DNSSEC operation while changing the provider of a domain that normally uses a single one).

Otherwise, the terminology in this document is as defined in [RFC7344].

2. Updates to RFCs

[RFC7344] Section 4.1 lists acceptance rules for CDS/CDNSKEY records. This list is extended with the consistency requirements defined in this document.

[RFC7477] Sections 3.1 and 4.2 have logic for deciding from which nameserver to query CSYNC information. This logic is replaced with the CSYNC consistency requirements defined in this document.

3. Processing Requirements

Consistency requirements that apply equally to CDS/CDNSKEY and CSYNC are listed first; type-specific consistency criteria are described in separate subsections.

In order to determine plausible consistency of CDS/CDNSKEY or CSYNC RRsets across the child’s nameservers, the Parental Agent MUST fetch all IP addresses for each nameserver hostname as listed in the Child’s delegation from the Parent, using a validating resolver at one vantage point, and including glue records if available. Before acting on any CDS/CDNSKEY or CSYNC record for the child, the Parental Agent MUST have established plausible consistency by querying all of these IP addresses for the record set(s) in question, as per the guidelines spelled in the following subsections.

In all cases, consistency is REQUIRED across received responses only. (A NODATA response is a received response.)

When a response cannot be obtained from a given nameserver, the Parental Agent SHOULD attempt to obtain it at a later time, before concluding that the nameserver is permanently unreachable and removing it from consideration. A retry schedule with exponential back-off is RECOMMENDED (such as after 5, 10, 20, 40, ... minutes). To sidestep localized routing issues, the Parental Agent MAY also attempt contacting the nameserver from another vantage point.

If an inconsistent state is encountered, the Parental Agent MUST abort the operation. Specifically, it MUST NOT delete or alter any existing RRset that would have been deleted or altered, and MUST NOT create any RRsets that would have been created, had the nameservers given consistent responses.

To accommodate transient inconsistencies (e.g. replication delays), the Parental Agent MAY retry the full process, repeating all queries. A schedule with exponential back-off is RECOMMENDED.

Any pending queries can immediately be dequeued when encountering a response that confirms the status quo, either implicitly (NODATA) or explicitly. This is because any subsequent responses could only confirm that nothing needs to happen, or give an inconsistent result in which case nothing needs to happen. Queries MAY be continued across all nameservers for reporting purposes.

Existing requirements for ensuring integrity remain in effect. In particular, DNSSEC signatures MUST be requested and validated for all queries unless otherwise noted.

3.1. CDS and CDNSKEY

To retrieve a Child’s CDS/CDNSKEY RRset for DNSSEC delegation trust maintenance, the Parental Agent, knowing both the Child zone name and its NS hostnames, MUST ascertain that queries are made against all (reachable) nameservers listed in the Child’s delegation from the Parent, and ensure that each key referenced in any of the received answers is also referenced in all other received responses.

In other words, CDS/CDNSKEY records at the Child zone apex MUST be fetched directly from each (reachable) authoritative server as determined by the delegation’s NS record set. When a key is referenced in a CDS record set but not the CDNSKEY record set (or vice versa), or returned by one nameserver but is missing from at least one other nameserver’s answer, the CDS/CDNSKEY state MUST be considered inconsistent.

During initial DS provisioning (DNSSEC bootstrapping), conventional DNSSEC validation for CDS/CDNSKEY responses is not (yet) available; in this case, authenticated bootstrapping ([RFC9615]) should be used.

3.2. CSYNC

A CSYNC-based workflow generally consists of (1) querying the CSYNC (and possibly SOA) record to determine which data records shall be synchronized from child to parent, and (2) querying for these data records (e.g. NS), before placing them in the parent zone. If the below conditions are not met during these steps, the CSYNC state MUST be considered inconsistent.

When querying the CSYNC record, the Parental Agent MUST ascertain that queries are made against all (reachable) nameservers listed in the Child’s delegation from the Parent, and ensure that the record’s immediate flag and type bitmap are equal across received responses.

The CSYNC record’s SOA serial field and soaminimum flag might legitimately differ across nameservers (such as in multi-provider setups); equality is thus not required across responses. Instead, for a given response, processing of these values MUST occur with respect to the SOA record as obtained from the same nameserver. If the resulting per-nameserver assessments of whether the update is permissible do not all agree, the CSYNC state MUST be considered inconsistent.

Further, when retrieving the data record sets as indicated in the CSYNC record (such as NS or A/AAAA records), the Parental Agent MUST ascertain that all queries are made against all (reachable) nameservers listed in the delegation, and ensure that all return responses with equal rdata sets (including all empty).

Other CSYNC processing rules from [RFC7477] Section 3 remain in place without modification. For example, when the NS type flag is present, associated NS processing has to occur before potential glue updates to ensure that glue addresses match the right set of nameservers. Also, when the type bitmap contains the A/AAAA flags, corresponding address queries are only to be sent for NS hostnames "that are in bailiwick", while out-of-bailiwick NS records are ignored. For details, see [RFC7477] Sections 3.2.2 and Section 4.3.

CSYNC-based updates may cause validation or even insecure resolution to break (e.g. by changing the delegation to a set of nameservers that do not serve required DNSKEY records or do not know the zone at all). Parental Agents SHOULD check that CSYNC-based updates, if applied, do not break the delegation.

4. IANA Considerations

This document has no IANA actions.

5. Security Considerations

The level of rigor mandated by this document is needed to prevent publication of half-baked DS or delegation NS RRsets (authorized only under an insufficient subset of authoritative nameservers), ensuring that a single operator cannot unilaterally modify the delegation (add or remove trust anchors or nameservers) when other operators are present. This applies both when the setup is intentional and when it is unintentional (such as in the case of lame delegation hijacking).

As a consequence, the delegation’s records can only be modified when zones are synchronized across operators, unanimously reflecting the domain owner’s intentions. Both availability and integrity of the domain’s DNS service benefit from this policy.

In order to resolve situations in which consensus about child zone contents cannot be reached (e.g. because one of the nameserver operators is uncooperative), Parental Agents SHOULD continue to accept DS and NS/glue update requests from the domain owner via an authenticated out-of-band channel (such as EPP [RFC5730]), irrespective of the adoption of automated delegation maintenance. Availability of such an interface also enables recovery from a situation where the private key is no longer available for signing the CDS/CDNSKEY or CSYNC records in the child zone.

6. Implementation Status

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

This draft has been implemented by

* TANGO Registry Services
* CORE Registry

7. Acknowledgments

In order of first contribution or review: Viktor Dukhovni, Wes Hardaker, Libor Peltan, Oli Schacher, David Blacka, Charlie Kaufman, Michael Bauland, Patrick Mevzek, Joe Abley, Ondナ册j Caletka.

8. References

8.1. 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>.

[RFC5730] Hollenbeck, S., "Extensible Provisioning Protocol (EPP)", STD 69, RFC 5730, DOI 10.17487/RFC5730, August 2009, <https://www.rfc-editor.org/info/rfc5730>.

[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>.

[RFC7477] Hardaker, W., "Child-to-Parent Synchronization in DNS", RFC 7477, DOI 10.17487/RFC7477, March 2015, <https://www.rfc-editor.org/info/rfc7477>.

[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>.

[RFC9615] Thomassen, P. and N. Wisiol, "Automatic DNSSEC Bootstrapping Using Authenticated Signals from the Zone’s Operator", RFC 9615, DOI 10.17487/RFC9615, July 2024, <https://www.rfc-editor.org/info/rfc9615>.

8.2. Informative References

[LAME1] Akiwate, G., Jonker, M., Sommese, R., Foster, I., Voelker, G. M., Savage, S., Claffy, K., and ACM, "Unresolved Issues", Proceedings of the ACM Internet Measurement Conference, DOI 10.1145/3419394.3423623, 27 October 2020, <http://dx.doi.org/10.1145/3419394.3423623>.

[LAME2] Akiwate, G., Savage, S., Voelker, G. M., Claffy, K C, and ACM, "Risky BIZness", Proceedings of the 21st ACM Internet Measurement Conference, DOI 10.1145/3487552.3487816, 2 November 2021, <http://dx.doi.org/10.1145/3487552.3487816>.

[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>.

[RFC8901] Huque, S., Aras, P., Dickinson, J., Vcelak, J., and D. Blacka, "Multi-Signer DNSSEC Models", RFC 8901, DOI 10.17487/RFC8901, September 2020, <https://www.rfc-editor.org/info/rfc8901>.

Appendix A. Failure Scenarios

The following scenarios are examples of how things can go wrong when consistency is not enforced by the parent during CDS/CDNSKEY/CSYNC processing. Other scenarios that cause similar (or perhaps even more) harm may exist.

The common feature of these scenarios is that if one nameserver steps out of line and the parent is not careful, DNS resolution and/or validation will break down. When several DNS providers are involved, this undermines the very guarantees of operator independence that multi-provider configurations are intended to provide.

A.1. DS Breakage due to Replication Lag

If an authoritative nameserver is lagging behind during a key rollover, the parent may see different CDS/CDNSKEY RRsets depending on the nameserver contacted. This may cause old and new DS RRsets to be deployed in an alternating fashion and without the awareness of the zone maintainer, who may then inadvertently break the chain of trust by prematurely removing a DNSKEY still referenced by a (stale) CDS/CDNSKEY RRset.

While foreseen in [RFC7344] Section 6.2, the solution suggested there requires parents to keep state on CDS/CDNSKEY RRsets. This document achieves the same without this burden, and in case the parent reports consistency errors downstream, can also help detection of the child-side replication issue by the operator.

A.2. Escalation of Lame Delegation Takeover

A delegation may include a non-existent NS hostname, for example due to a typo or when the nameserver's domain registration has expired. (Re-)registering such a non-resolvable nameserver domain allows a third party to run authoritative DNS service for all domains delegated to that NS hostname, serving responses different from the legitimate ones.

This strategy for hijacking (at least part of the) DNS traffic and spoofing responses is not new, but surprisingly common [LAME1][LAME2]. It is also known that DNSSEC reduces the impact of such an attack, as validating resolvers will reject illegitimate responses due to lack of signatures consistent with the delegation's DS records.

On the other hand, if the delegation is not protected by DNSSEC, the rogue nameserver is not only able to serve unauthorized responses without detection; it is even possible for the attacker to escalate the nameserver takeover to a full domain takeover.

In particular, the rogue nameserver can publish CDS/CDNSKEY records. If those are processed by the parent without ensuring consistency with other authoritative nameservers, the delegation will, with some patience, get secured with the attacker's DNSSEC keys. Of course, as the parent窶冱 query (or sometimes queries) need to hit the attacker's nameserver, this requires some statistical luck; but eventually it will succeed. As responses served by the remaining legitimate nameservers are not signed with these keys, validating resolvers will start rejecting them.

Once DNSSEC is established, the attacker can use CSYNC to remove other nameservers from the delegation at will (and potentially add new ones under their control), or change glue records to point to the attacker's nameservers. This enables the attacker to position themself as the only party providing authoritiative DNS service for the victim domain, significantly augmenting the attack's impact.

A.3. Multi-Provider (Permanent Multi-Signer)

A.3.1. DS Breakage

While performing a key rollover and adjusting the corresponding CDS/CDNSKEY records, a provider could accidentally publish CDS/CDNSKEY records that only include its own keys.

When the parent happens to retrieve the records from a nameserver controlled by this provider, the other providers' DS records would be removed from the delegation. As a result, the zone is broken at least for some queries.

A.3.2. NS Breakage

A similar scenario affects the CSYNC record, which is used to update the delegation's NS record set at the parent. The issue occurs, for example, when a provider accidentally includes only their own set of hostnames in the local NS record set, or publishes an otherwise flawed NS record set.

If the parent then observes a CSYNC signal and fetches the flawed NS record set without ensuring consistency across nameservers, the delegation may be updated in a way that breaks resolution or silently reduces the multi-provider setup to a single-provider setup.

A.4. Bogus Provider Change (Temporary Multi-Signer)

Transferring DNS service for a domain name from one (signing) DNS provider to another, without going insecure, necessitates a brief period during which the domain is operated in multi-signer mode: First, the providers include each other's signing keys as DNSKEY and CDS/CDNSKEY records in their copy of the zone. Once the parent learns about the updated CDS/CDNSKEY record set at the old provider, the delegation's DS record set is updated. Then, after waiting for cache expiration, the new provider's NS hostnames can be added to the zone's NS record set, so that queries start balancing across both providers. (To conclude the hand-over, the old provider is removed by inverting these steps with swapped roles.)

The multi-signer phase of this process breaks when the new provider, perhaps unaware of the situation and its intricacies, fails to include the old provider's keys in the DNSKEY (and CDS/CDNSKEY) record sets. One obvious consequence of that is that whenever the resolver happens to retrieve the DNSKEY record set from the new provider, the old provider's RRSIGs do no longer validate, causing SERVFAIL to be returned.

However, an even worse consequence can occur when the parent performs their next CDS/CDNSKEY scan: It may then happen that the incorrect CDS/CDNSKEY record set is fetched from the new provider and used to update the delegation's DS record set. As a result, the old provider (who still appears in the delegation) is prematurely removed from the domain's DNSSEC chain of trust. The new DS record set authenticates the new provider's DNSKEYs only, and DNSSEC validation fails for all answers served by the old provider.

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

* draft-ietf-dnsop-cds-consistency-08

| Take into account RFC 7344 Section 6.2 for Appendix A.1 considerations

* draft-ietf-dnsop-cds-consistency-07

| Clarify that "all nameservers" means fetching all delegation NS IPs

* draft-ietf-dnsop-cds-consistency-06

| Editorial changes from Dnsdir early review
|
| Add Implementation Status

* draft-ietf-dnsop-cds-consistency-05

| Editorial overhaul

* draft-ietf-dnsop-cds-consistency-04

| Clarify that existing CSYNC NS and glue processing rules remain in place
|
| Editorial changes
|
| Clean up "multi-homing" and define "multi-provider"/"multi-signer"

* draft-ietf-dnsop-cds-consistency-03

| Clarify that CSYNC updates should not break delegations
|
| Describe consistency requirements for CSYNC soaminimum
|
| Editorial changes

* draft-ietf-dnsop-cds-consistency-02

| Retry before assuming a nameserver is permanently unreachable

* draft-ietf-dnsop-cds-consistency-01

| Make nits tool happy
|
| New failure mode: DS Breakage due to Replication Lag
|
| Point out zero overhead if nothing changed, and need for OOB interface
|
| Editorial changes
|
| Moved Failure Scenarios to appendix

* draft-ietf-dnsop-cds-consistency-00

| Point out zero overhead if nothing changed, and need for OOB interface
|
| Editorial changes.

* draft-thomassen-dnsop-cds-consistency-03

| Describe risk from lame delegations
|
| Acknowledgments
|
| Say what is being updated
|
| Editorial changes.
|
| Retry mechanism to resolve inconsistencies

* draft-thomassen-dnsop-cds-consistency-02

| Don’t ignore DoE responses from individual nameservers (instead, require consistency across all responses received)

* draft-thomassen-dnsop-cds-consistency-01

| Allow for nameservers that don’t respond or provide DoE (i.e. require consistency only among the non-empty answers received)
|
| Define similar requirements for CSYNC.
|

| Editorial changes.

* draft-thomassen-dnsop-cds-consistency-00

| Initial public draft.

Author's Address

Peter Thomassen
SSE - Secure Systems Engineering GmbH
Hauptstraテタ 3
10827 Berlin
Germany
Email: peter.thomassen@securesystems.de